



The University of Hong Kong



# How Machines Learn (Without Being Taught)

Michael I. Shamos, Ph.D., J.D.  
School of Computer Science  
Carnegie Mellon University

# **As the artificial intelligence race heats up, China is taking action to move toward its goal of rapidly surpassing America**

WASHINGTON — When a Google computer program beat the world's best player of an ancient Chinese board game last May, it might have seemed like an incremental milestone.

But for some, the success of the program known as AlphaGo marked more than a man-vs.-machine clash. It set up a broader race between China and the United States over artificial intelligence, a competition that could mold the future of humankind just as the widespread arrival of electricity did in the last century.

The stakes are high. Advances in artificial intelligence could add trillions of dollars to a major economy and give an edge on the battlefield, shifting empires and global power.

# Machine Learning

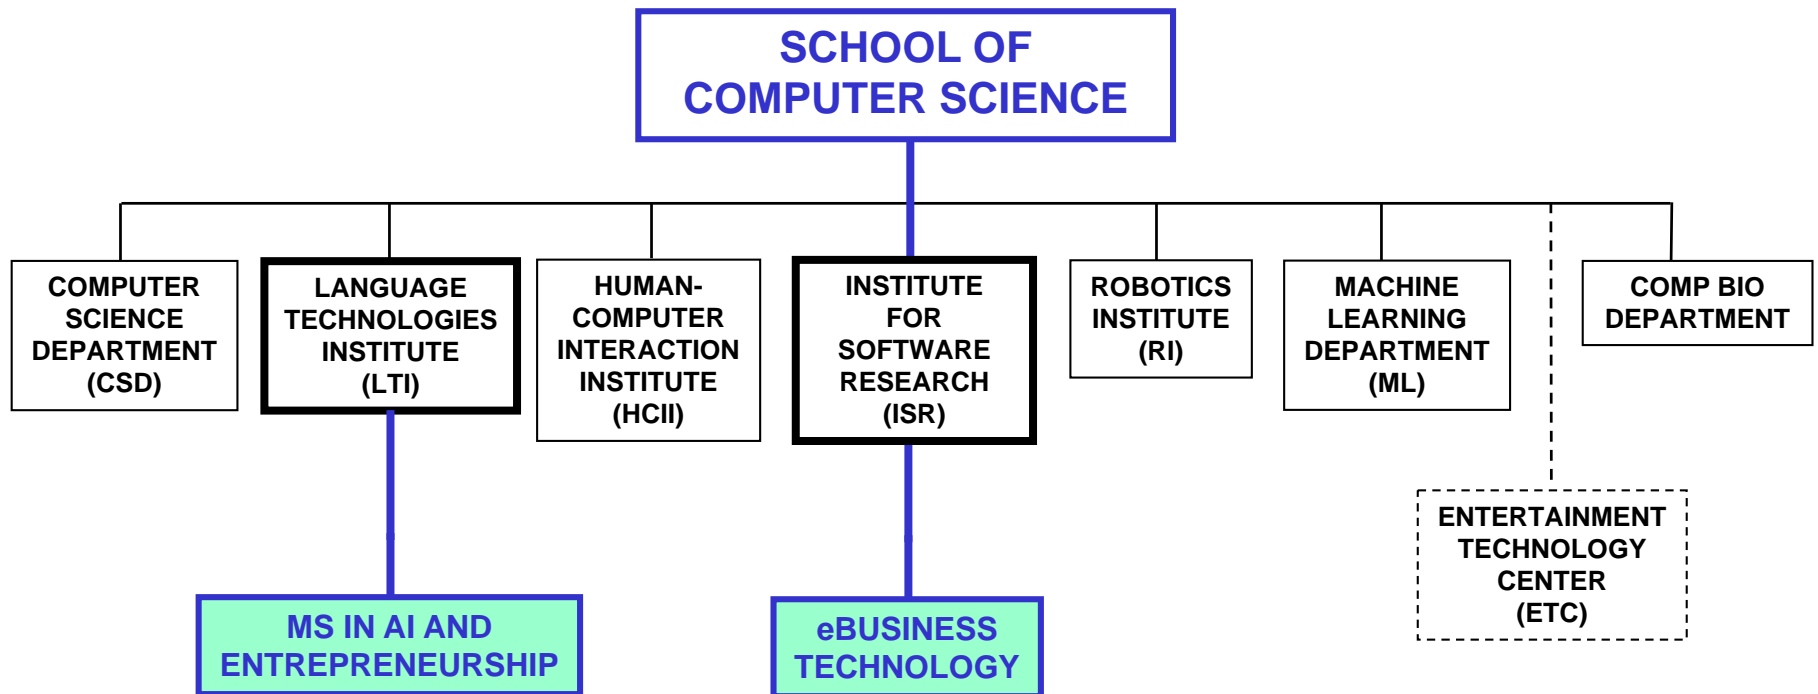
- The computer is incredibly fast, accurate and stupid.  
Man is unbelievably slow, inaccurate and brilliant.  
The marriage of the two is a challenge and opportunity beyond imagination.

-- Stuart G. Walesh, author  
and consultant

# Background

- Ph.D., Yale University (computer science, 1978)
- J.D., Duquesne University (law, 1981)
- Carnegie Mellon computer science faculty since 1975
- Visiting Professor, University of Hong Kong (2001- ),  
Electronic Payment Systems
  
- Director, Master's Program in eBusiness Technology,  
roughly equivalent to HKU Ecom/Icomp
- Incoming Director, MS in Artificial Intelligence and  
Entrepreneurship

# Carnegie Mellon School of Computer Science



# Machine Learning

- A computer program “learns” from experience if its performance on a task improves based on that experience.

-- paraphrased from Carnegie  
Mellon Professor Tom Mitchell

# Machine Learning Examples

Self-driving cars



By Steve Jurvetson (CC BY 2.0)

Recommendation systems



[https://commons.wikimedia.org/wiki/File:Netflix\\_Logo.svg](https://commons.wikimedia.org/wiki/File:Netflix_Logo.svg) (public domain)



<https://i10.kr.jp/EBLW6G> (CC BY 2.0)


Spam Filtering

Text Recognition



[http://commons.wikimedia.org/wiki/File:American\\_book\\_company\\_1916\\_letter\\_envelope-2.JPG#filelinks](http://commons.wikimedia.org/wiki/File:American_book_company_1916_letter_envelope-2.JPG#filelinks) (public domain)

Biology

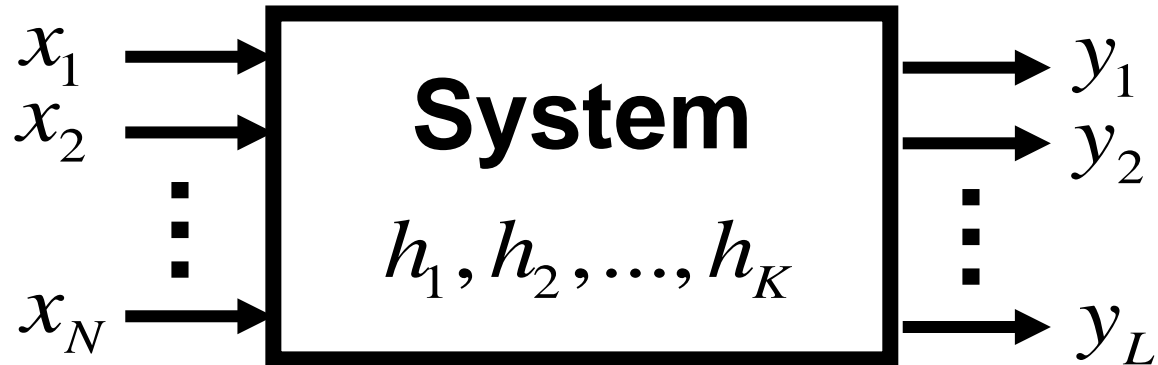


# Types of Machine Learning

- No learning
  - Static computer program. Always performs the same way. Changes made by humans.
- Supervised learning
  - The program is given examples of inputs and desired outputs. “Trains” itself to perform well.
- Unsupervised learning
  - Program given only inputs and must discover patterns in the data.
- Reinforcement learning
  - Program is given only inputs, but gets rewards for good outputs. Objective: maximize reward.



# A Computational System

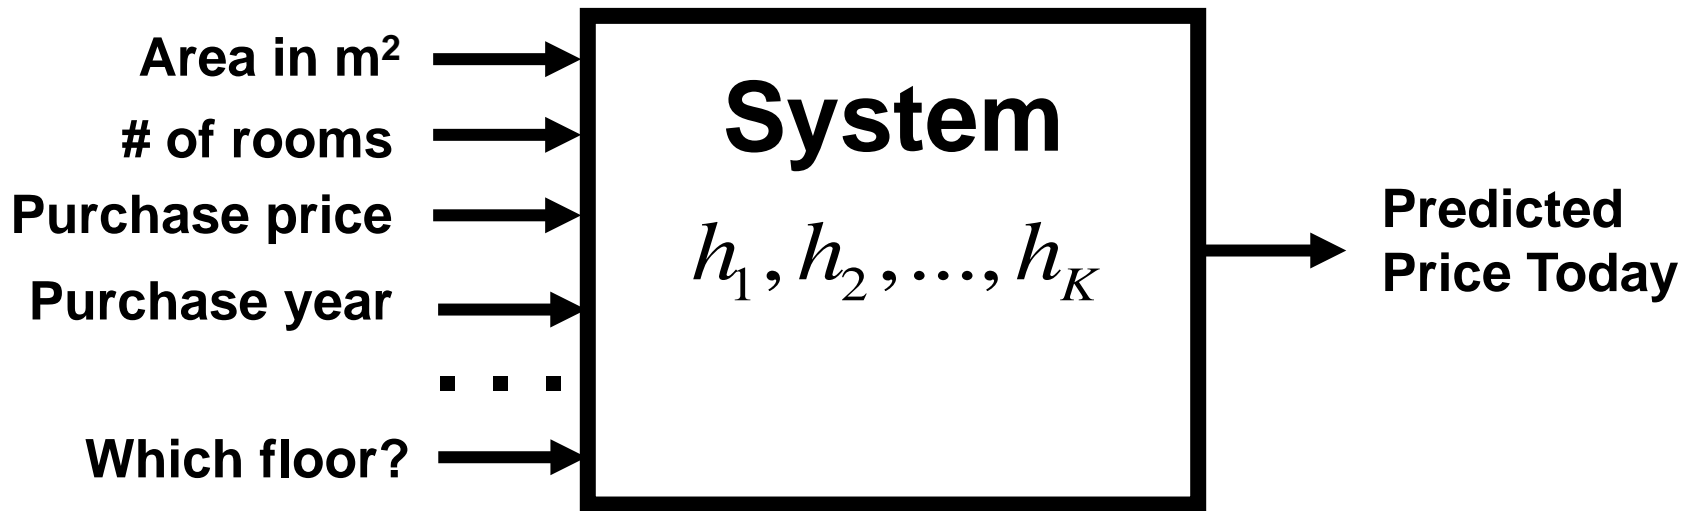


**Inputs:**  $\mathbf{x} = (x_1, x_2, \dots, x_N)$

**Internal Variables:**  $\mathbf{h} = (h_1, h_2, \dots, h_K)$

**Outputs:**  $\mathbf{y} = (y_1, y_2, \dots, y_L)$

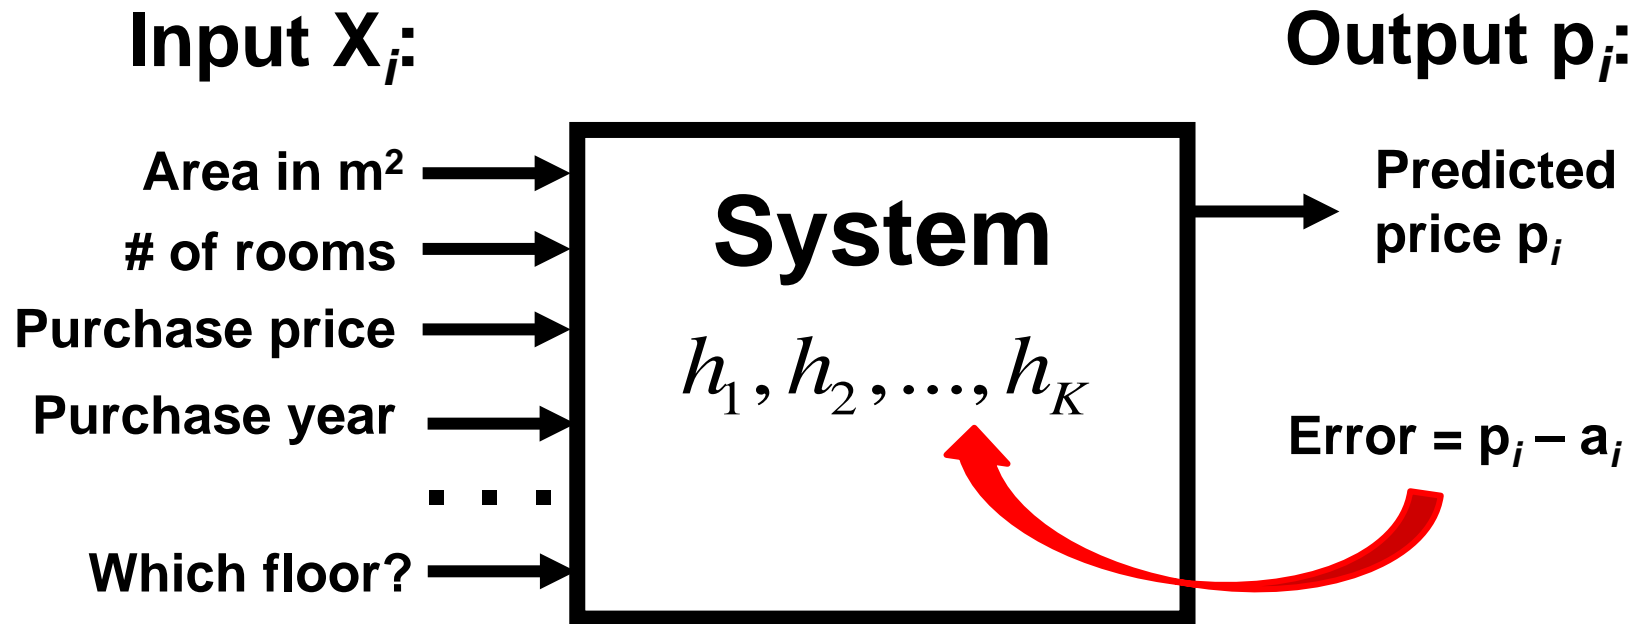
# Predicting Apartment Prices



**Classic non-learning approach: construct a model of apartment prices and write a computer program**

**No learning. If the model is inaccurate, we need a new model and a new program**

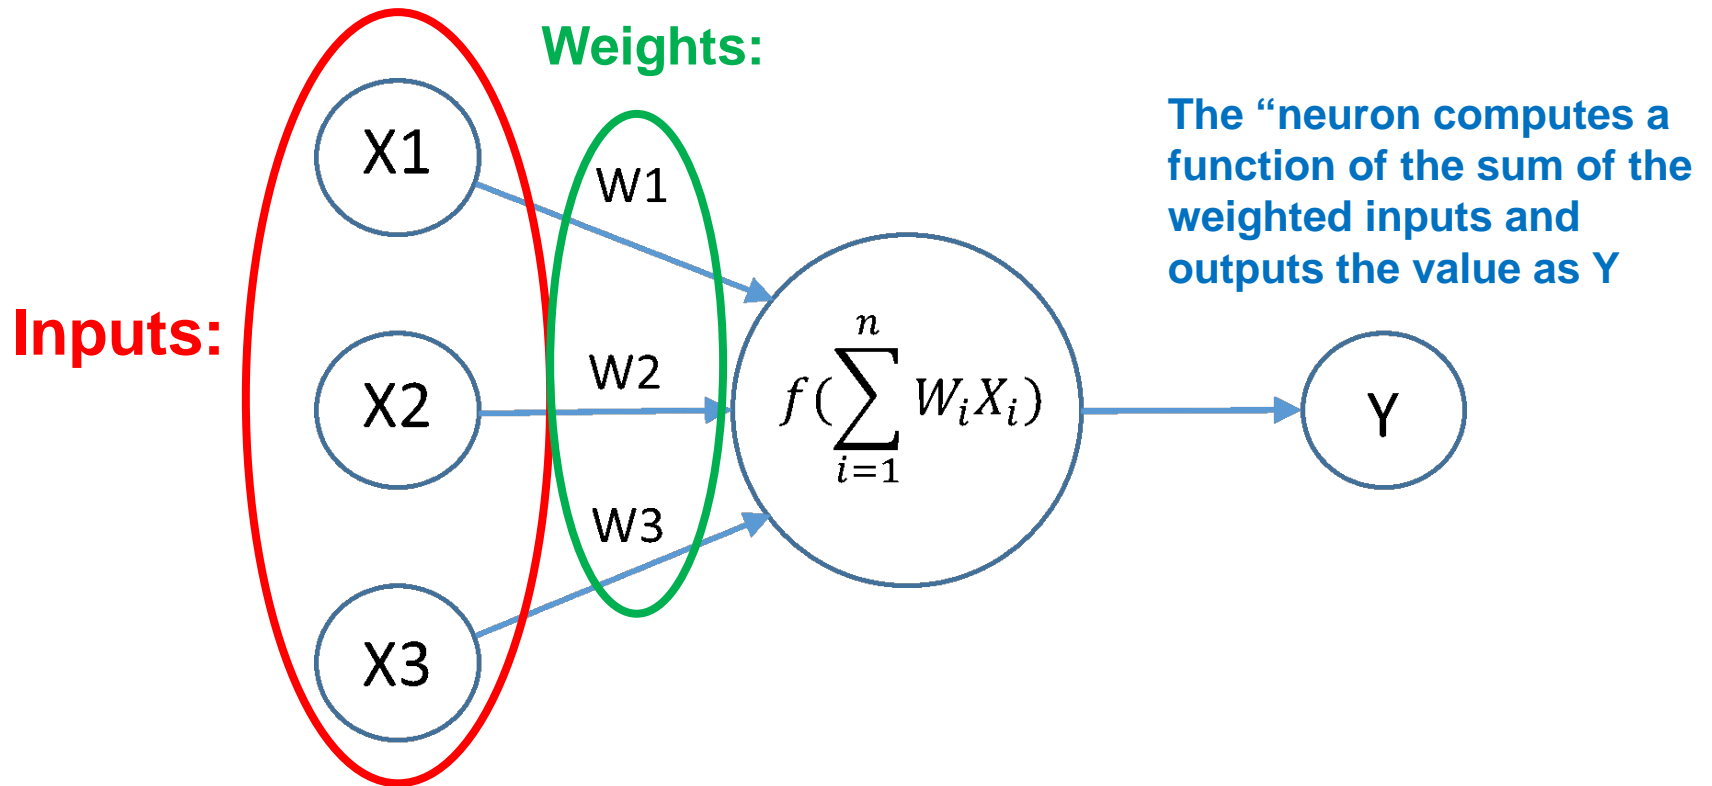
# Learning to Predict Apartment Prices



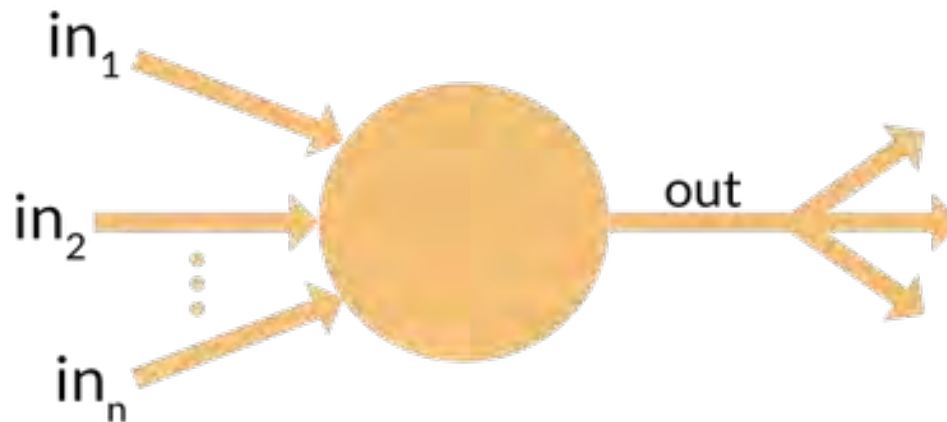

**Supervised learning approach: use a large number  $M$  of actual price examples ( $X_i$ , actual price  $a_i$ )**

**Compare the predicted price  $p$  to the actual price  $a$ , and modify the program to reduce the error  $e = p - a$**

# A “Neuron”




# Biological Basis of Neurons





SOURCE: [QUORA.COM](https://www.quora.com)

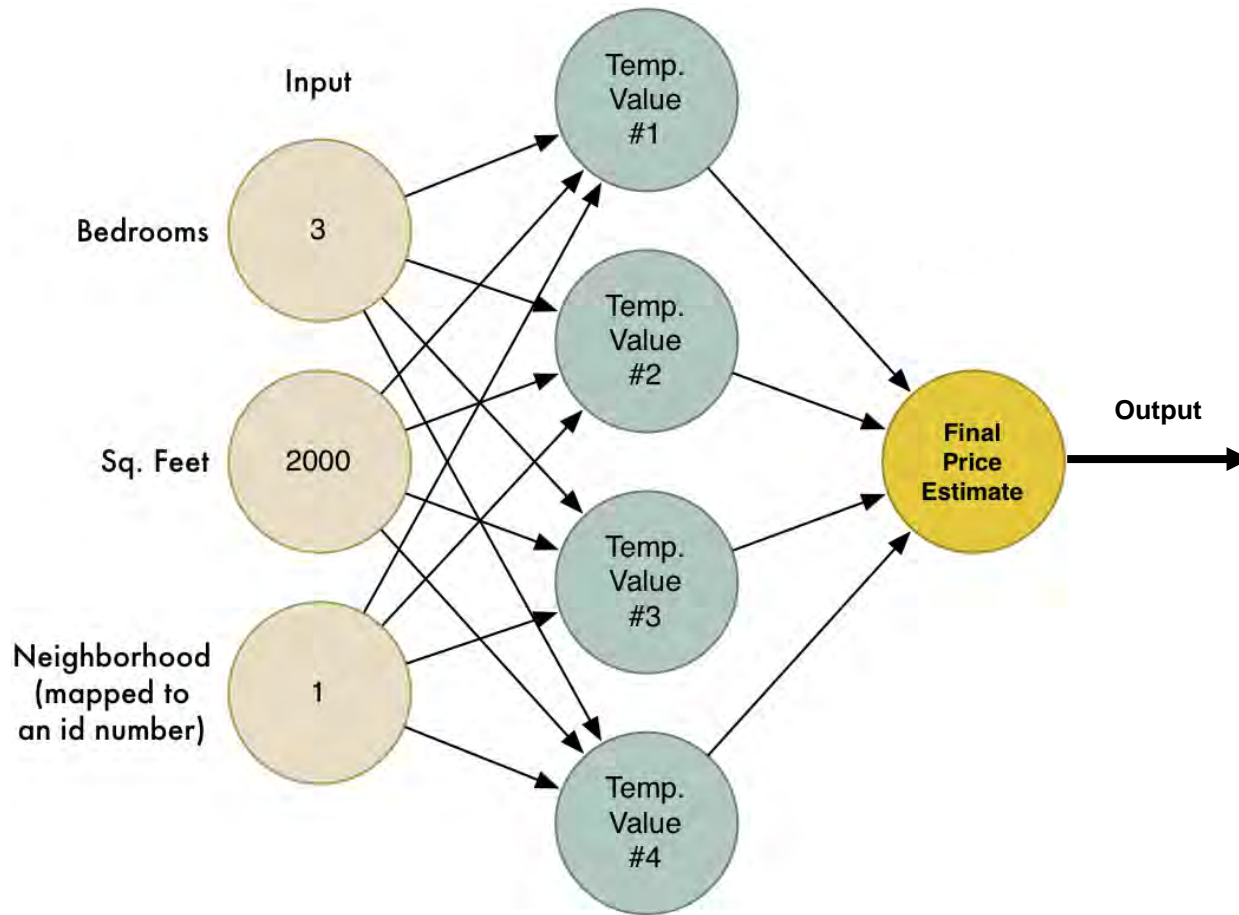
# A “Neural Network”



# A Neural Network is a Computational System



# Neural Network for Price Estimation



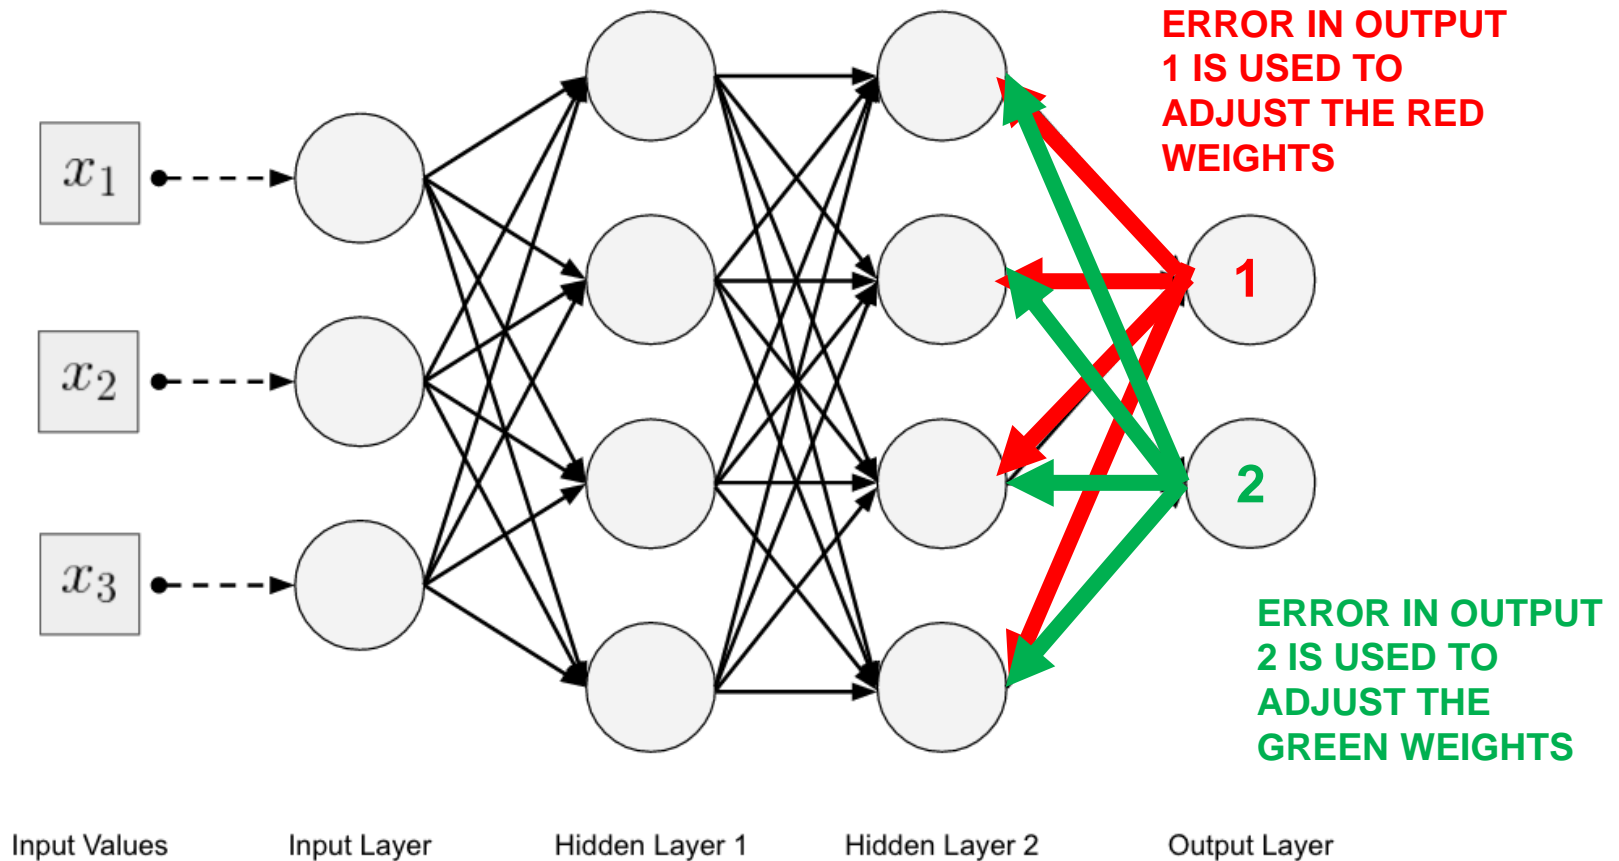


# Neural Networks Can Learn

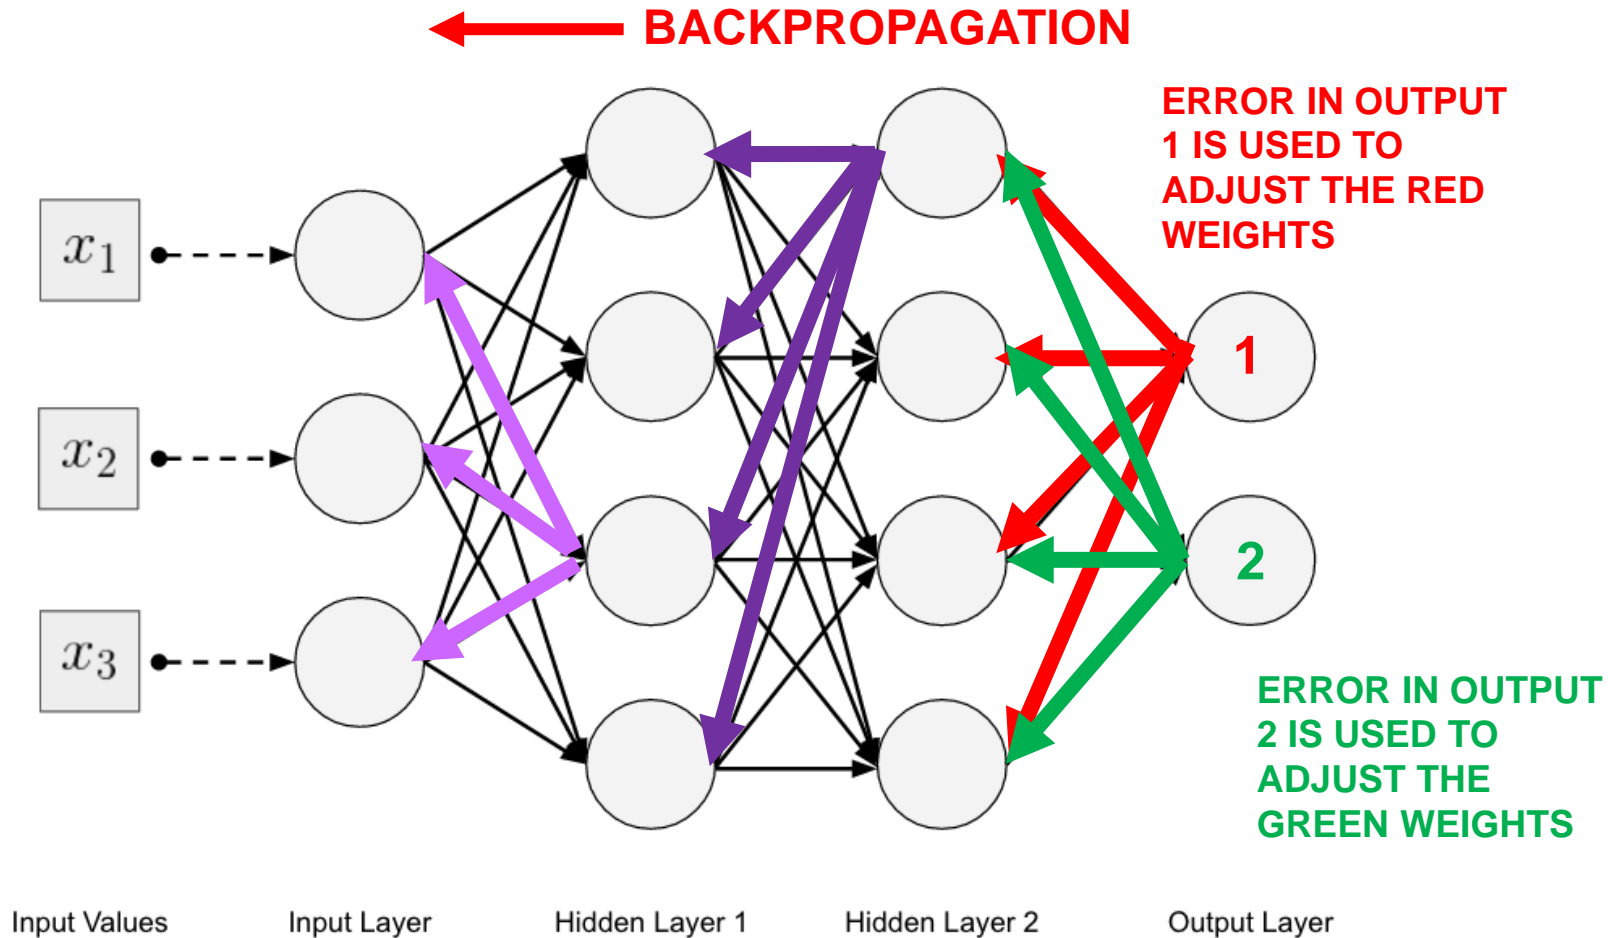
## A learning algorithm:

- Given an input and a known desired output, run the neural network to see the actual output
- $\text{Error} = \text{desired output} - \text{actual output}$
- Use the error to modify the weights in the network
- This is called “training” the network


# Neural Networks Can Learn



# Neural Networks Can Learn



# What an Image Looks Like to a Machine



$x = (255, 255, 255, 255, \dots)$

A sequence of red-green-blue (RGB) color intensity values

$(0, 0, 0) = \text{black}$

$(255, 255, 255) = \text{white}$

$(255, 98, 89) = \text{a shade of pink}$

$2^{24} = 16 \text{ million possible values for each pixel}$

For a 1000 x 1000 pixel image, 16 trillion possible inputs

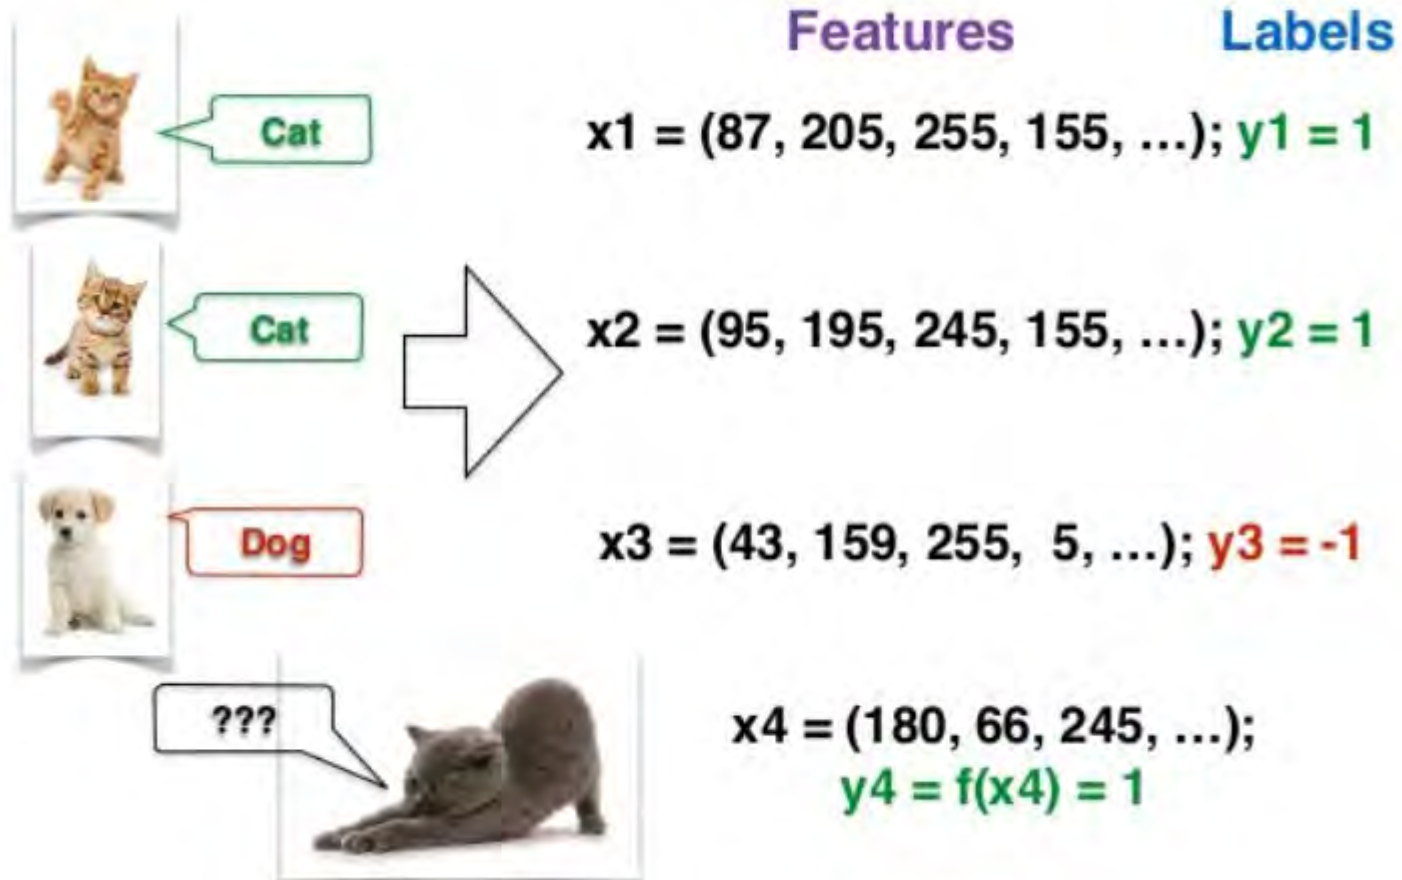
SOURCE: DMYTRO FISHMAN

# Variations of Cat



SOURCE:POO KUAN HOONG

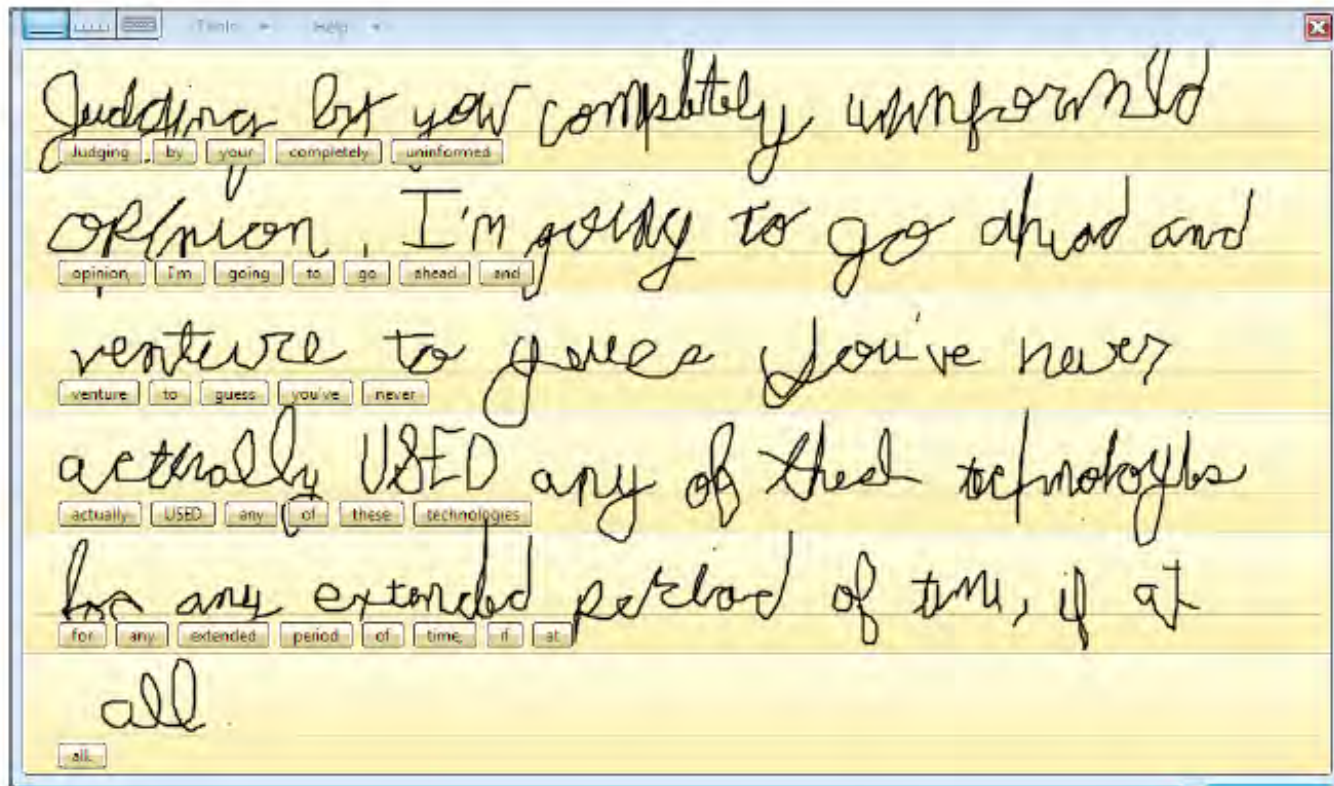
# Supervised Learning



SOURCE: DMYTRO FISHMAN

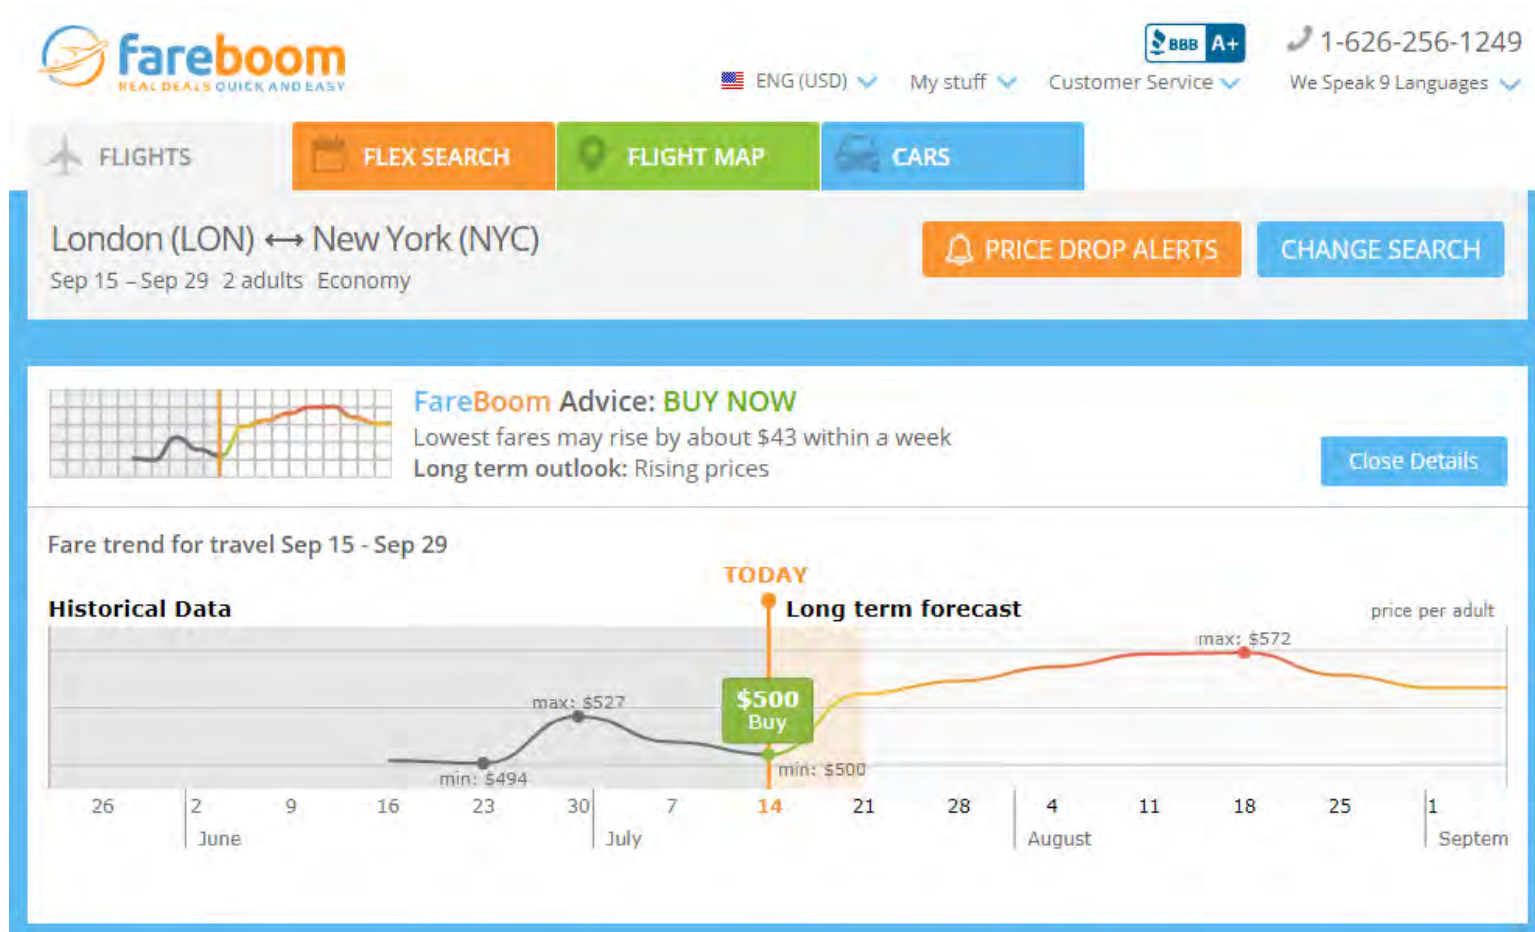
# Supervised Learning Application

## Handwriting Recognition on a PDA



SOURCE: E. ALPAYDIN

# Supervised Learning Application Predicting Airline Ticket Prices





# Unsupervised Learning

- No training data
- Network must detect similarities or patterns in the inputs

# Example: Clustering News Stories

**Categories / Titles**

**India >**

- Tamil Nadu Assembly passes dhoti bill**  
Zee News - 9 minutes ago  
Chennai: A bill seeking to remove the ban on wearing dhoti and other Indian traditional attire imposed by recreational clubs and others was on Wednesday passed by the Tamil Nadu Assembly.
- HSGMC members clash with police in Kurukshetra**  
Times of India - 10 minutes ago  
CHANDIGARH: Tension escalated in Kurukshetra district on Wednesday as members of the Haryana Sahitya Akademi clashed with police after they "fired" a guard, where Amritar-based members were ...
- Manjula Chellur sworn in as first woman Chief Justice of Calcutta High Court**  
The Hindu - 7 hours ago  
Justice Manjula Chellur was sworn in as the first woman Chief Justice of Calcutta High Court here on Tuesday. Governor Keshari Nath Tripathi administered the oath of office at Raj Bhavan during the day and Chief Minister Mamata Banerjee was present on ...

**World >**

- Gaza truce holds for the second day, diplomacy on to secure peace**  
Daily News & Analysis - 20 minutes ago  
Palestinians were still coming to terms with the scale of devastation in Gaza on the second day of a ceasefire between Israel and Hamas on Wednesday, even as the two sides engaged in diplomacy in Egypt to hammer out a long-term truce.
- Death toll in China quake nears 600; 4200 people evacuated**  
Firstpost - 3 hours ago  
Beijing: The death toll in a powerful 6.5 magnitude earthquake that hit southwest China neared 600 on Wednesday as 4,200 people were evacuated after artificial lakes formed by the landslides posed a risk of flooding.
- Australian commuters tip train cars to rescue man trapped in station gap**  
National Post - 18 minutes ago  
SYDNEY, Australia - Dozens of people helped rescue a fellow commuter in Australia by pushing against train carriages to free the man whose leg had slipped between the platform and the train.

**Business >**


- Indian rupee posts biggest single-day fall in 6-1/2 months**  
Reuters India - 35 minutes ago  
... \* Rupee ends at 61.4950/5050 per dlr vs 60.8450/8550 on Tues. \* Heavy debt and equity selling by foreign funds hurts rupee. \* INR likely to stabilise going ahead, says trader.
- Sensex ends nearly 1 percent lower; banks fall**  
Reuters India - 6 minutes ago  
MUMBAI (Reuters) - The BSE Sensex and the Nifty fell nearly 1 percent on Wednesday, snapping two days of gains, as lenders such as State Bank of India and ICICI Bank slumped on worries that a sell-off in government bonds would hit the value of their debt ...
- Adani Port net up 36% at Rs 568 cr in Q1**  
Hindu Business Line - 47 minutes ago  
Adani Port and SEZ Ltd (ADPEZL) on Wednesday reported ...

**Technology >**

- Apple to Hold iPhone Event On September 9**  
Indiatimes.com - 28 minutes ago  
Apple Inc has scheduled a "big" media event related to the iPhone for September 9, technology news website Re/code said, without citing sources.
- Regulating WhatsApp and the likes: Indian telcos are coming across as cry babies**  
Firstpost - 38 minutes ago  
Indian telecom companies have been quite vocal about their dislike for over-the-top (OTT) services such as messaging and calling apps eating into one of their biggest revenue streams - voice calling and SMSes.
- Hive UI on Xolo Play 8X-1000 first impressions: Not compelling enough to make ...**  
Firstpost - 38 minutes ago  
The new Xolo Play 8X-1000 smartphone is a mid-range device with a 5.5-inch screen and a 13MP camera. It is powered by the Xolo 8800 processor and runs on Android 4.4.2 KitKat. The device is priced at Rs 12,999. The first impressions are mixed, with some praising the design and performance, while others feel the UI is not compelling enough to make it stand out in the market.


# Unsupervised News Clustering

Model P(document)

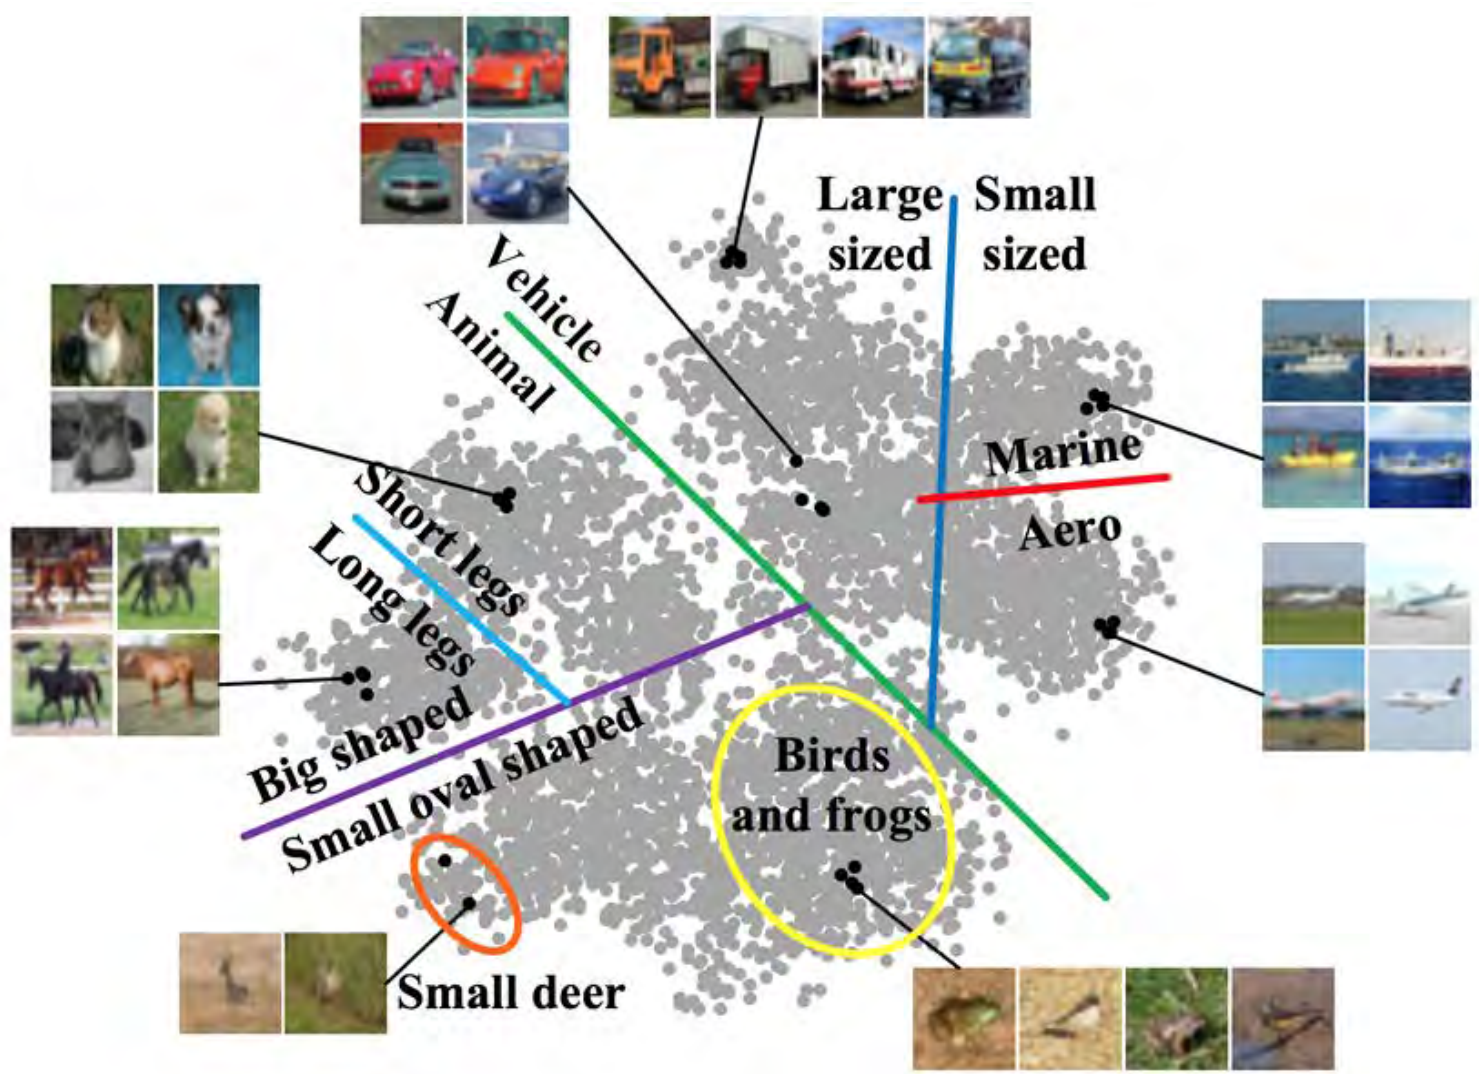


Bag of words

Reuters dataset: 804,414  
newswire stories: **unsupervised**

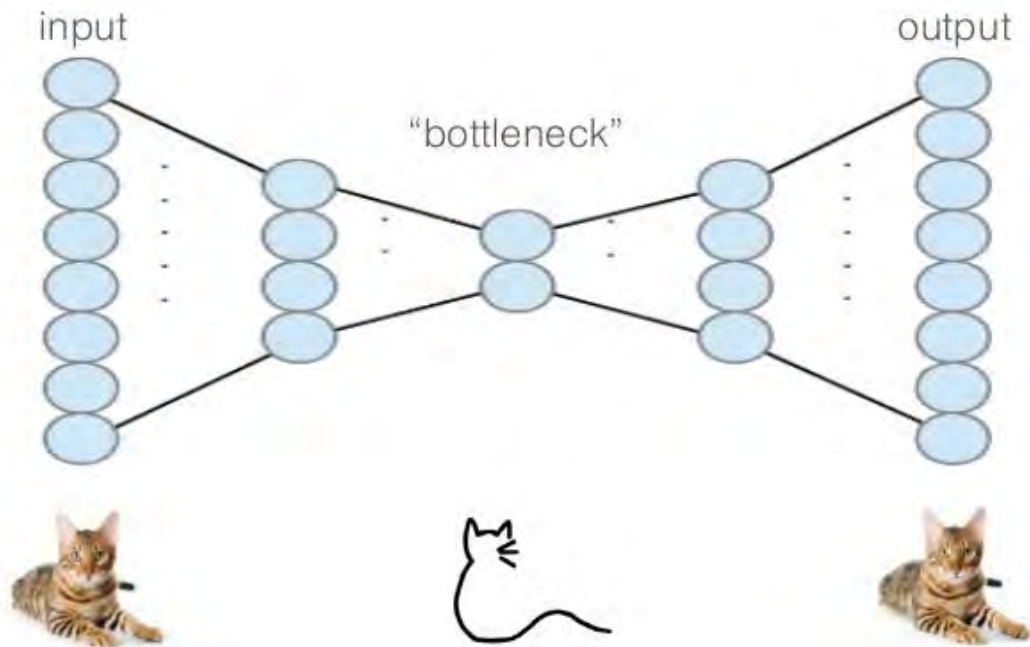


# Unsupervised Image Recognition




# Autoencoders


- Idea: “compress” patterns to represent them with fewer features in a “code.” Train the net to reproduce the original patterns just from the code.
- Gives a much more robust recognizer.




# Autoencoders



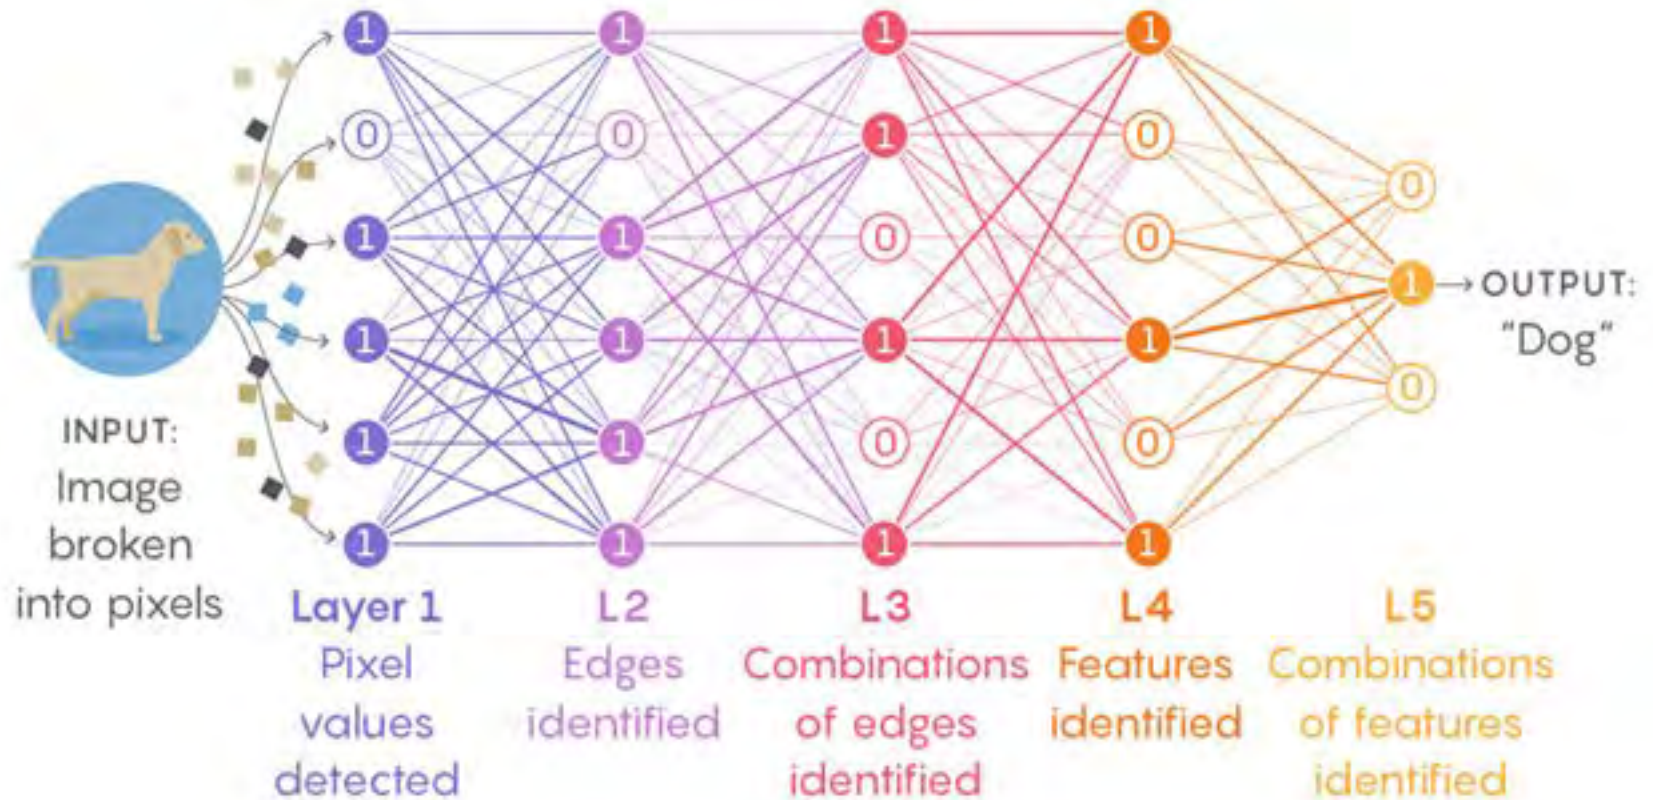
**Better: Add noise!**



SOURCE: [ARDEN DERTAT](#)

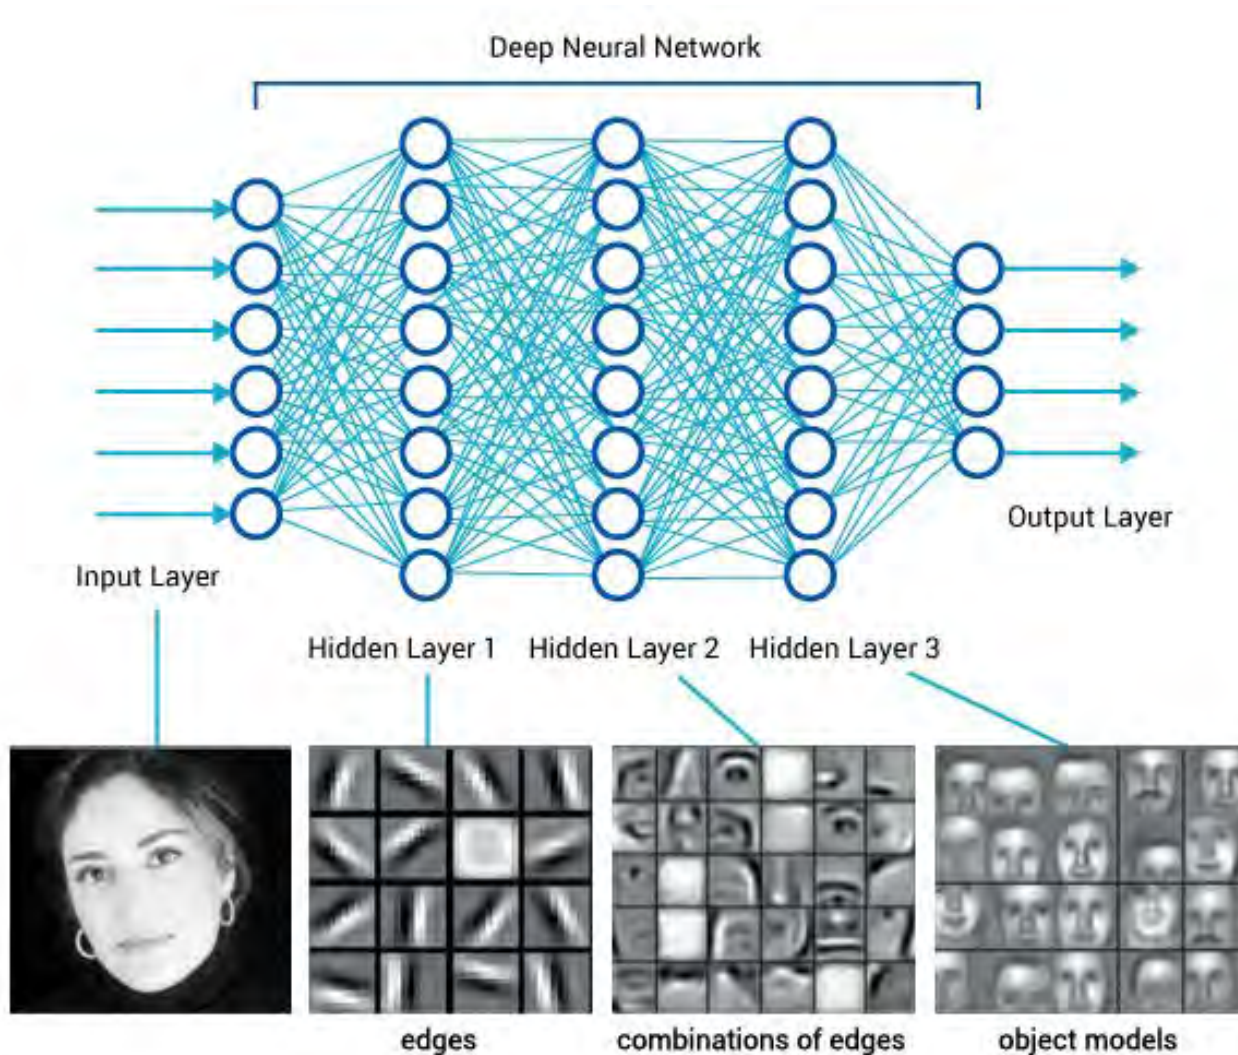


# A Deep Neural Network

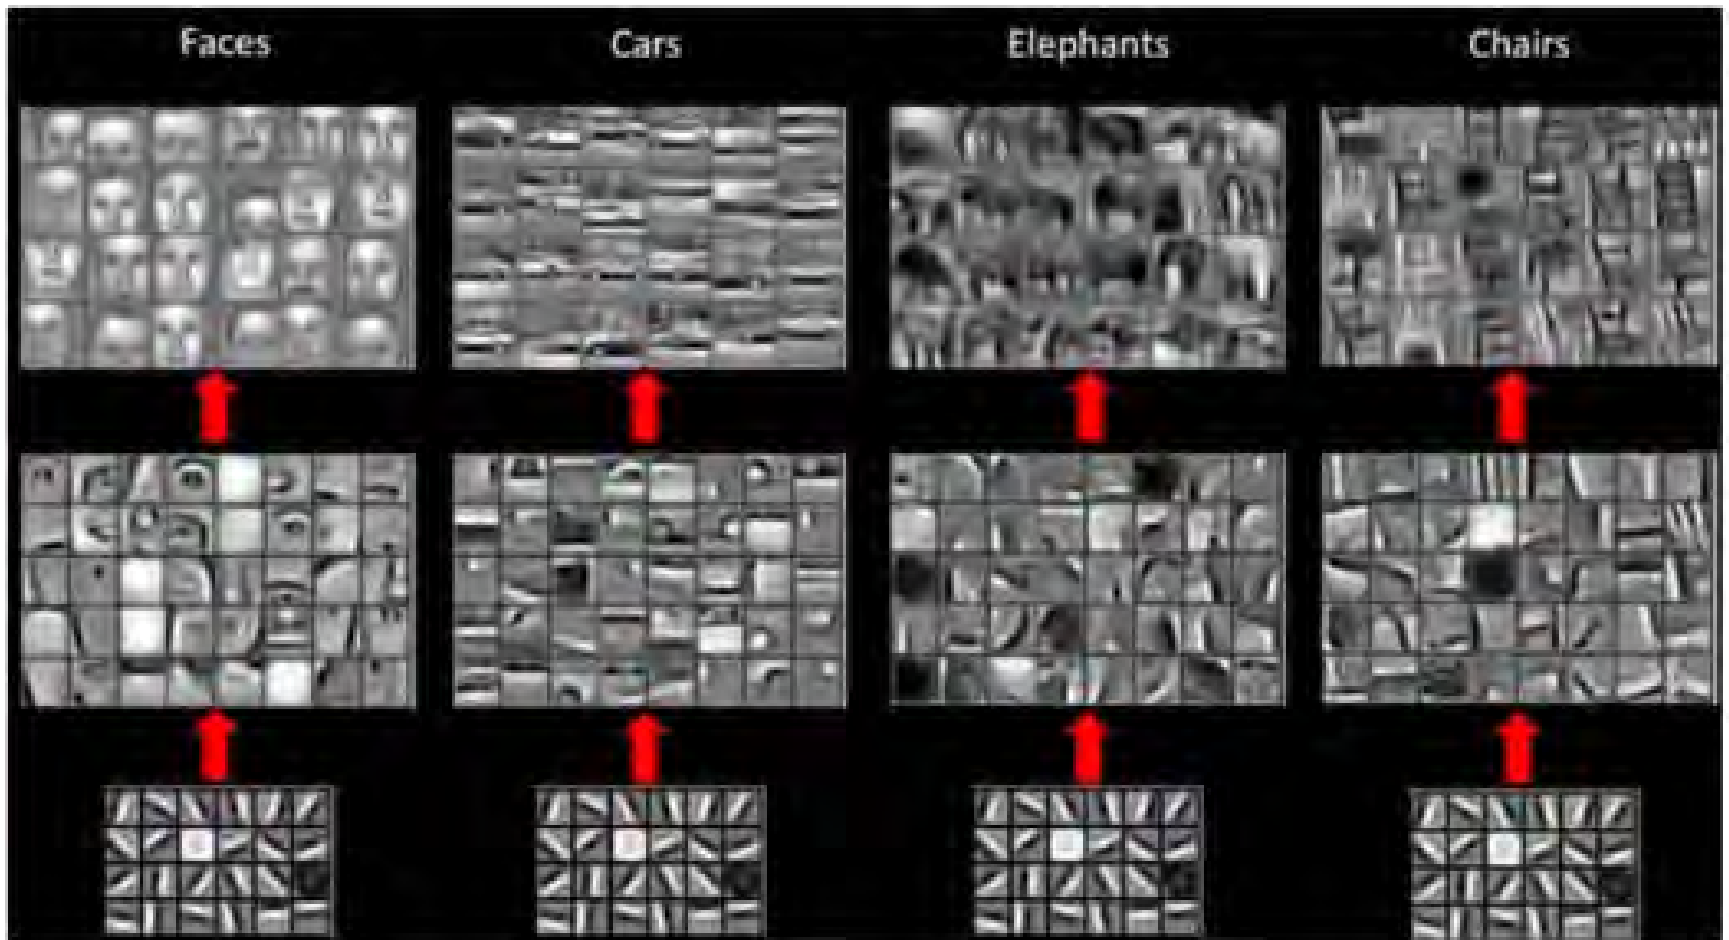




# Deep Neural Network



# Deep Image Recognition

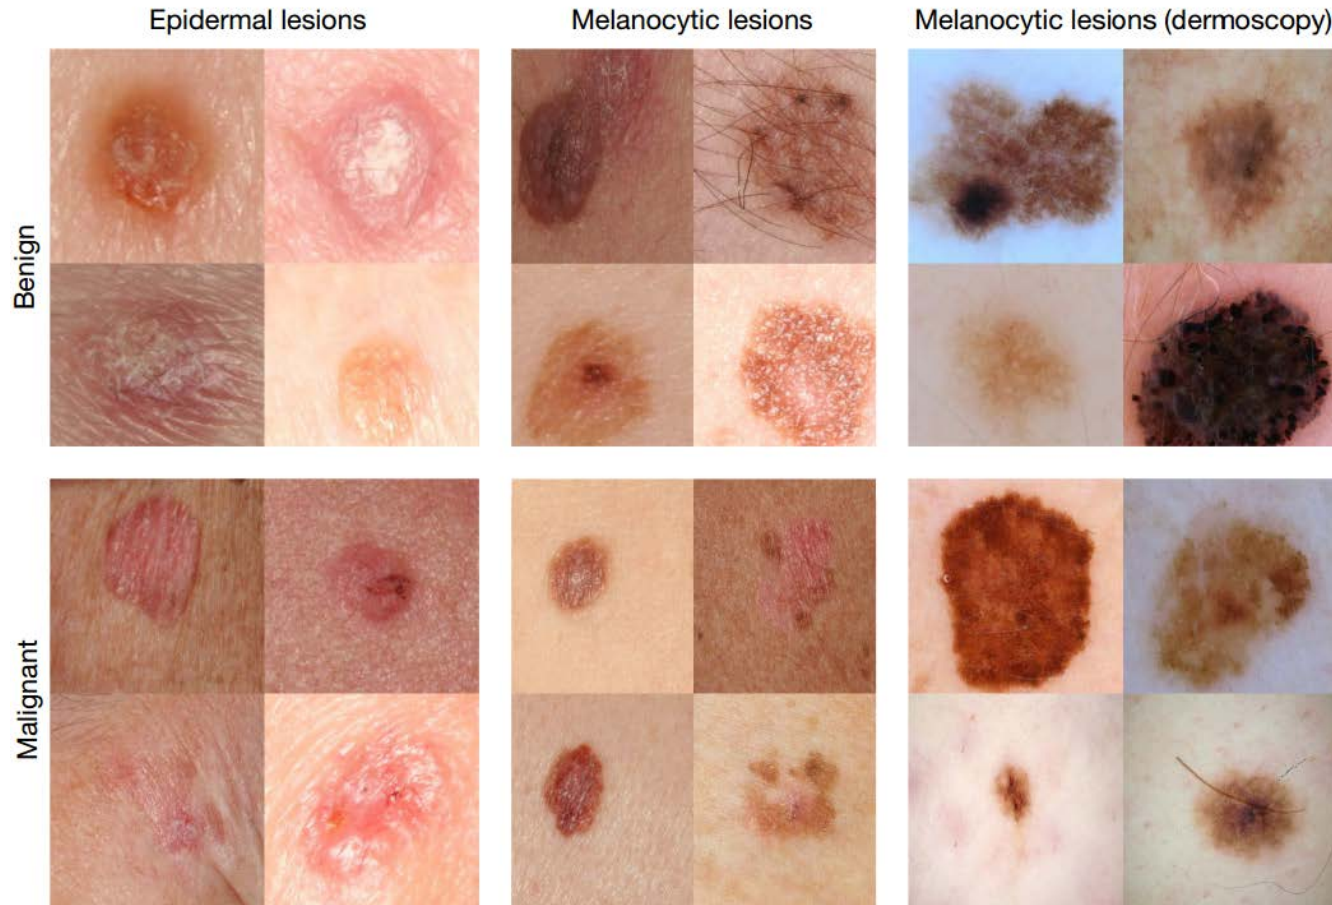


SOURCE: [DATASKEPTIC](#)

# Application: Diagnosing Skin Cancer

- Stanford researchers collected 130,000 images of skin lesions representing over 2,000 different diseases
- Used the data as a training set on a deep neural network using only pixels and disease labels as inputs
- Performs as well as expert dermatologists, better than non-experts
- Projection: 6.3 billion smartphones by the year 2021
- Can provide low-cost universal access to diagnostic procedures

# Application: Diagnosing Skin Cancer



SOURCE: [EXTREMETECH.COM](http://EXTREMETECH.COM)

Basal cell carcinomas



- Epidermal benign
- Epidermal malignant
- Melanocytic benign
- Melanocytic malignant

Squamous cell carcinomas




Melanomas



Nevi



Seborrheic keratoses



SOURCE: [EXTREMETECH.COM](http://EXTREMETECH.COM)

# Deep Learning Applications

- Uber estimates arrival time by training a neural network on millions of previous trips
- UberEATS estimates food preparation time to allow prediction of final delivery time
- Recommendation engines: Amazon, Netflix (estimated value: 20B HKD)
- Google Maps: analysis of 80 billion street view images to recognize house numbers and street signs
- Facebook DeepFace facial recognition



## Meet the winners of the biggest ever face-recognition challenge

Everyone from government agencies to police forces are looking for software to track us in airports or spot us in CCTV images. But much of this technology is developed behind closed doors – how can we know if any of it really works?


To answer this question, the Intelligence Advanced Research Projects Activity (IARPA) and the US National Institute of Standards and Technology (NIST) have been running the biggest face-recognition competition to date.

The Face Recognition Prize Challenge tested two tasks: face verification and face search. Face verification is what phone manufacturers such as Apple – whose iPhone X, out last week, can be unlocked with your face – are trying to master. The software must say whether a face matches that of a known person. Face search is the harder problem. It requires finding every image of a person in a database of maybe millions of images.


The winner of the face-verification task was a company called Ntech whose FindFace product can match a person's face correctly 99.9 per cent of the time.



# Facebook DeepFace



# Reinforcement Learning



Positive Reinforcement



Decision  
Buy / Sell

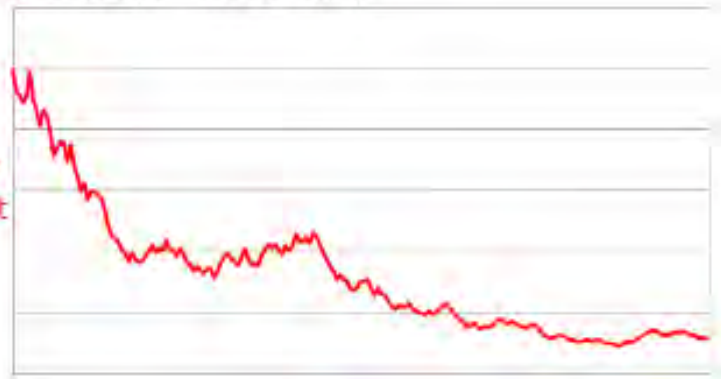


Negative Reinforcement


Result one - make money



Result two - lose money



# Reinforcement Learning Example: Atari Breakout




- Reward function is “Score”: number of targets removed

# Google DeepMind Video

(play to 2:03)

# Reinforcement Learning in Go



- In 2017, AlphaGo Master defeated the world Go champion, Ke Jie. He called it “God.”
- A later version, AlphaGo Zero, can now beat AlphaGo Master

# AlphaGo Zero


- No knowledge of Go except rules for legal moves
- Reward function: number of stones remaining at end of game
- AlphaGo Zero played a huge number of games against itself to maximize its reward
- 1.6 million games per day
- Outputs were used to train a neural network
- Hardware cost: 25 million USD

# AlphaGo Zero Progress

---

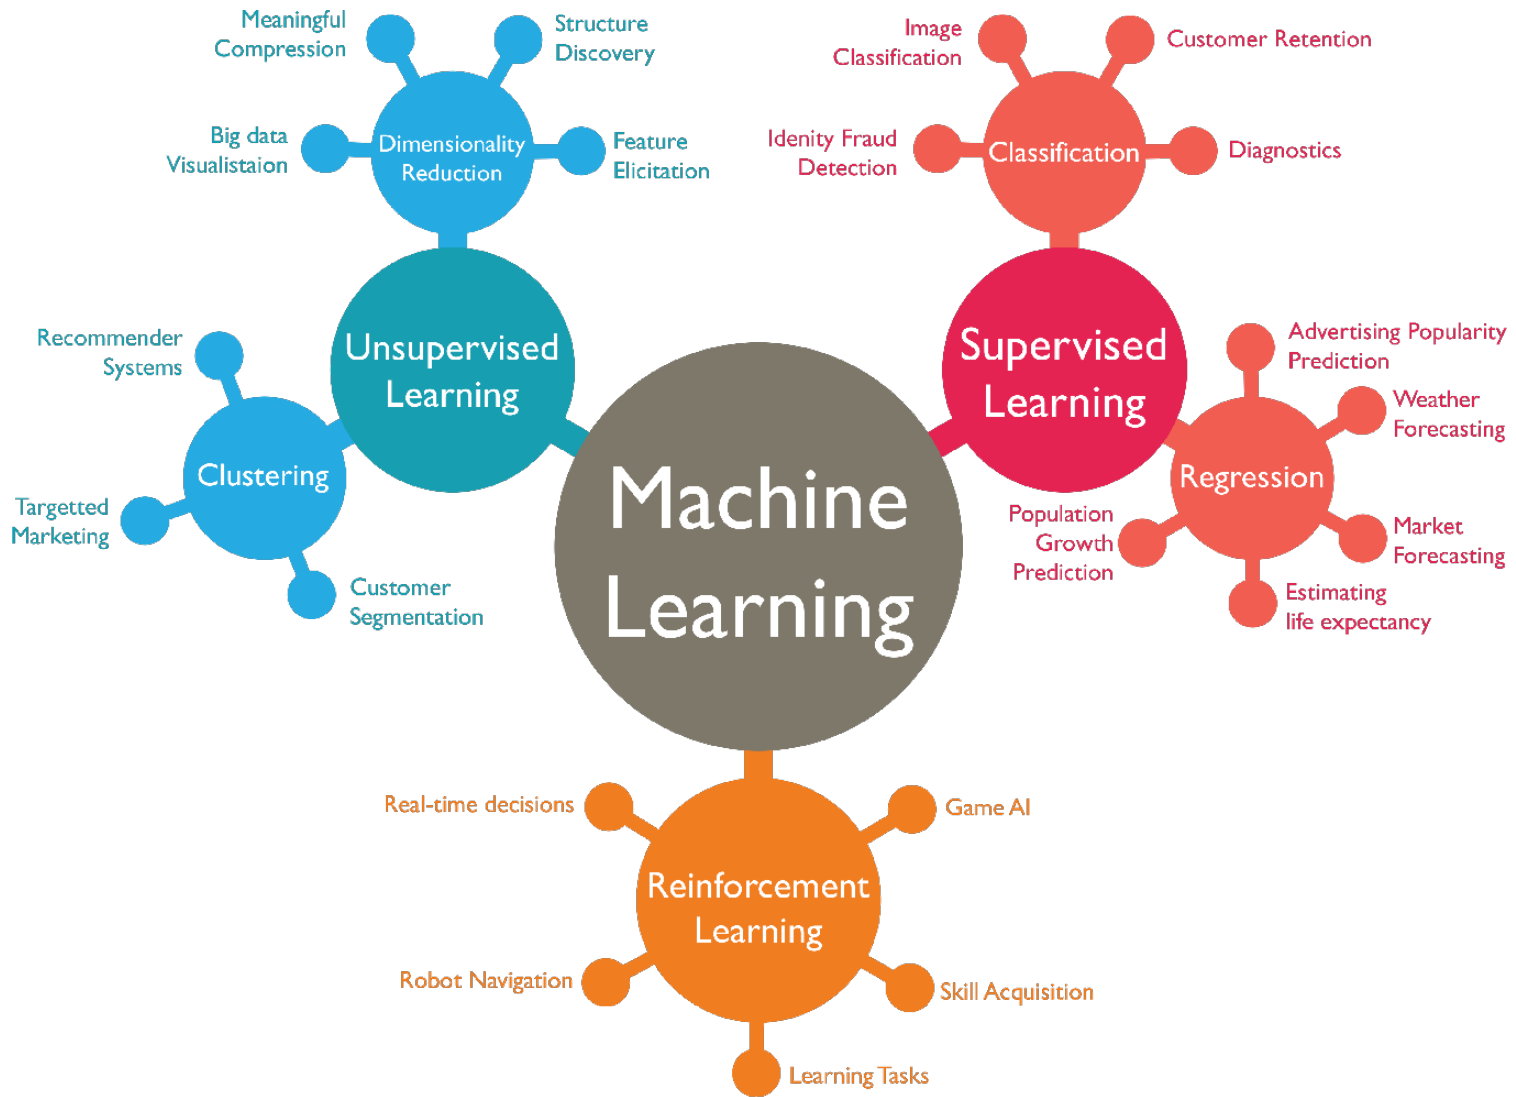
SOURCE: [DEEPMIND.COM](https://www.deepmind.com)

# AlphaGo Progress



SOURCE: [DEEPMIND.COM](https://www.deepmind.com)





SOURCE: [NORMSHIELD.COM](https://www.normshield.com)

# Machine Intelligence LANDSCAPE

## CORE TECHNOLOGIES

### ARTIFICIAL INTELLIGENCE



### DEEP LEARNING



### MACHINE LEARNING



### NLP PLATFORMS



### PREDICTIVE APIS



### IMAGE RECOGNITION



### SPEECH RECOGNITION



## RETHINKING ENTERPRISE

### SALES



### SECURITY / AUTHENTICATION



### FRAUD DETECTION



### HR / RECRUITING



### MARKETING



### PERSONAL ASSISTANT



### INTELLIGENCE TOOLS



## RETHINKING INDUSTRIES

### ADTECH



### AGRICULTURE



### EDUCATION



### FINANCE



### LEGAL



### MANUFACTURING



### MEDICAL



### OIL AND GAS



### MEDIA / CONTENT



### CONSUMER FINANCE



### PHILANTHROPIES



### AUTOMOTIVE



### DIAGNOSTICS



### RETAIL



## RETHINKING HUMANS / HCI

### AUGMENTED REALITY



### GESTURAL COMPUTING



### ROBOTICS



### EMOTIONAL RECOGNITION




## SUPPORTING TECHNOLOGIES

### HARDWARE



### DATA PREP



### DATA COLLECTION





# 60+ STARTUPS USING DEEP LEARNING

## CORE AI: COMPUTER VISION



## CORE AI: OTHER



## BI, SALES & CRM



## CORE AI: VOICE INTERFACE



## ROBOTICS & AUTO




## HEALTHCARE



## SECURITY



## OTHER



istock.com/a-image

## E-COMMERCE



## ACQUIRED (2014-2016YTD)



# Q&A