

# VBM683

# Machine Learning

Pinar Duygulu

Slides are adapted from  
Dhruv Batra, David Sontag, Aykut Erdem

# Quotes

- “If you were a current computer science student what area would you start studying heavily?”
  - Answer: Machine Learning.
  - “The ultimate is computers that learn”
  - Bill Gates, Reddit AMA
- “Machine learning is the next Internet”
  - Tony Tether, Director, DARPA
- “Machine learning is today’s discontinuity”
  - Jerry Yang, CEO, Yahoo

# Two definitions of learning

(1) Learning is the acquisition of knowledge about the world.

*Kupfermann (1985)*

(2) Learning is an adaptive change in behavior caused by experience.

*Shepherd (1988)*


# What is Machine Learning?

- [Arthur Samuel, 1959]
  - Field of study that gives computers
  - the ability to learn without being explicitly programmed
- [Kevin Murphy] algorithms that
  - automatically detect patterns in data
  - use the uncovered patterns to predict future data or other outcomes of interest
- [Tom Mitchell] algorithms that
  - improve their performance (P)
  - at some task (T)
  - with experience (E)

# What is Machine Learning?

- If you are a Scientist

**Data**



**Understanding**

- If you are an Engineer / Entrepreneur
  - Get lots of data
  - Machine Learning
  - ???
  - Profit!

# Acquisitions

## Google snaps up object recognition startup

DNNr

Google has ac  
Toronto, who

by Josh Lowensohn

2 / 0

Google has acqui  
research compan

Topic: Cloud

Follow via:

## Microsoft acquires legal-focused machine-learning vendor Equivio

**Summary:** Microsoft has purchased Equivio, maker of a machine-learning platform for the legal industry, for an undisclosed amount.



By Mary Jo Foley for All About Microsoft | January 20, 2015 -- 16:24 GMT (08:24 PST)

Follow @maryjofoley

95.6K followers


Get the ZDNet Cloud newsletter now

Microsoft has purchased Equivio, an eDiscovery/compliance vendor with a specialization in text analysis, for an undisclosed amount.

Microsoft officials announced the acquisition of the Israeli company -- its first acquisition of 2015 using more of its offshore cash -- on January 20.

**Update:** The Wall Street Journal reported back in October last year that Microsoft planned to buy Equivio for \$200 million.

**Update No. 2:** A Microsoft spokesperson said the \$200 million estimate was inflated and incorrect, but declined to provide a different figure.



« Search needs a shake-up

Songbirds use grammar rules »

## Machine Learning Startup Acquired by ai-one

### Press Release

For Immediate Release: August 4, 2011

San Diego artificial intelligence startup acquired by leading  
ring SDKs as market for advanced

oday that it acquired Auto-Semantics, a local start-up  
es to corporate IT departments. The acquisition is the  
nd acquisitions by ai-one that consolidates its  
ing market for machine learning technologies.



FOUNDED  
2011


#### OVERVIEW

DeepMind is a cutting edge artificial intelligence company. We combine the best techniques from machine learning and systems neuroscience to build powerful general-purpose learning algorithms. Founded by Demis Hassabis, Shane Legg and Mustafa Suleyman, the company is based in London and supported by some of the most iconic technology entrepreneurs and investors of the past decade. Our first commercial ...

# What is Machine Learning?

- Let's say you want to solve Character Recognition
- Hard way: Understand handwriting/characters

Aa Bb Cc Dd  
Ee Ff Gg Hh  
Ii Jj Kk Ll  
Mm Nn Oo Pp  
Qq Rr Ss Tt  
Uu Vv Ww Xx  
Yy Zz




(C) Dhruv Batra

Image Credit: <http://www.linotype.com/6896/devanagari.html>

# What is Machine Learning?

- Let's say you want to solve Character Recognition
- Hard way: Understand handwriting/characters
  - Latin
  - Devanagari
  - Symbols: <http://detexify.kirelabs.org/classify.html>





# What is Machine Learning?

- Let's say you want to solve Character Recognition
- Hard way: Understand handwriting/characters
- Lazy way: Throw data!



# Example: Netflix Challenge

- Goal: Predict how a viewer will rate a movie
- 10% improvement = 1 million dollars



(C) Dhruv Batra


Slide Credit: Yaser Abu-Mostapha

# Essence of Machine Learning:


- A pattern exists
- We cannot pin it down mathematically
- We have data on it

# Comparison

- **Traditional Programming**



- **Machine Learning**



# Why Study Machine Learning?

## Engineering Better Computing Systems

- Develop systems
  - too difficult/expensive to construct manually
  - because they require specific detailed skills/knowledge
  - *knowledge engineering bottleneck*
- Develop systems
  - that adapt and customize themselves to individual users.
  - Personalized news or mail filter
  - Personalized tutoring
- Discover new knowledge from large databases
  - Medical text mining (e.g. migraines to calcium channel blockers to magnesium)
  - *data mining*

# Why Study Machine Learning?


## Cognitive Science

- Computational studies of learning may help us understand learning in humans
  - and other biological organisms.
- Hebbian neural learning
  - “Neurons that fire together, wire together.”

# Why Study Machine Learning?

The Time is Ripe

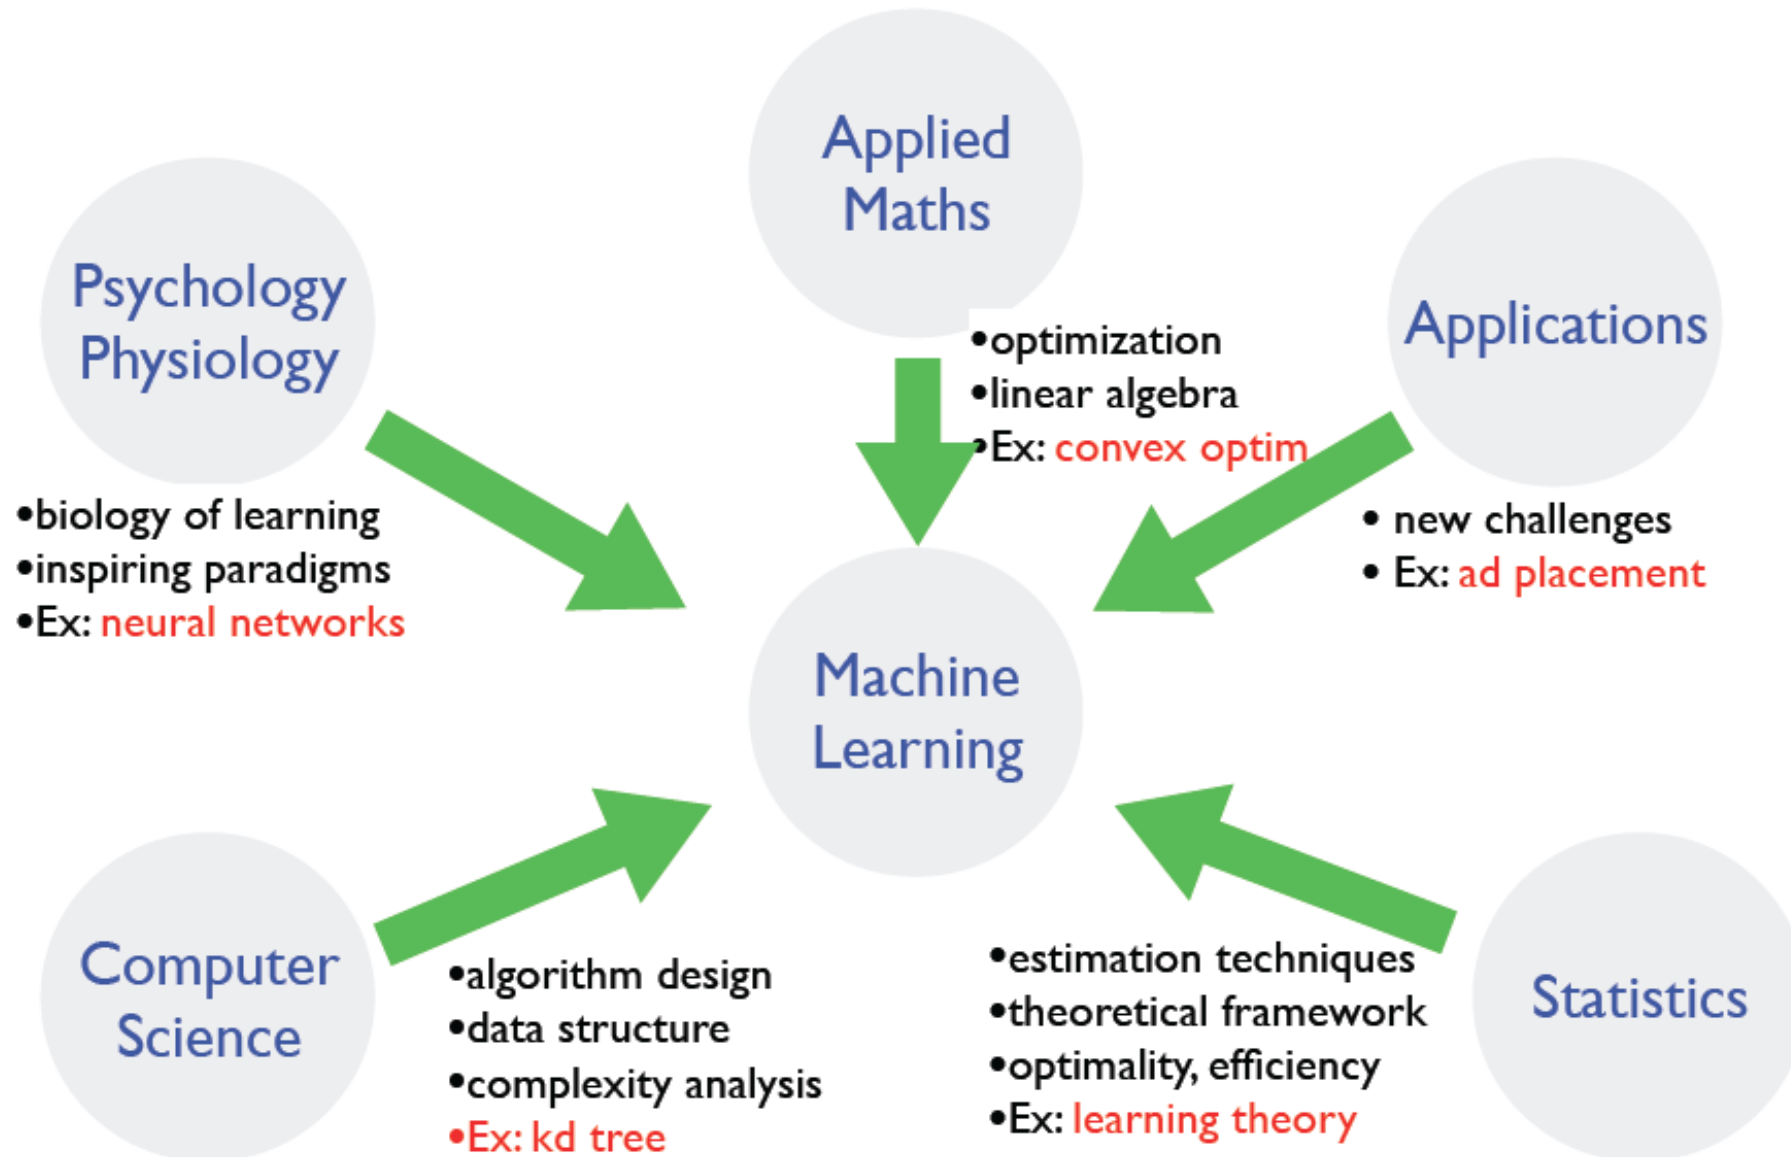
- More compute power
- More data
- Better algorithms /models



(C) Dhruv Batra

Figure Credit: Banko & Brill, 2011

# Where does ML fit in?





# A Brief History of AI



A Proposal for the Dartmouth Summer Research  
Project on Artificial Intelligence.

(John McCarthy)



(C) Dhruv Batra

# A Brief History of AI

- “We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire.”
- The study is to proceed on the basis of the conjecture that 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.
- An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.
- We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”


# Why is AI hard?



(C) Dhruv Batra

Slide Credit: <http://karpathy.github.io/2012/10/22/state-of-computer-vision/>


# What humans see



(C) Dhruv Batra

Slide Credit: Larry Zitnick


# What computers see



(C) Dhruv Batra

Slide Credit: Larry Zitnick


“I saw her duck”



“I saw her duck”



“I saw her duck”









(C) Dhruv Batra

Image Credit: Liang Huang



“I saw her duck with a telescope...”



# We've come a long way...



What is Jeopardy?

- <http://youtu.be/Xqb66bdsQlw?t=53s>

Challenge:

- [http://youtu.be/\\_429UIzN1JM](http://youtu.be/_429UIzN1JM)

• Watson Demo:

- [http://youtu.be/WFR3lOm\\_xhE?t=22s](http://youtu.be/WFR3lOm_xhE?t=22s)

• Explanation

- [http://youtu.be/d\\_yXV22O6n4?t=4s](http://youtu.be/d_yXV22O6n4?t=4s)

• Future: Automated operator, doctor assistant, finance

# AlphaGo vs Lee Sedol



Slide Credit: Aykut Erdem



# NVIDIA BB8 AI Car



Meet NVIDIA BB8

End to End Learning for Self-Driving Cars

Matthew Bejorral  
NVIDIA Corporation  
Hoboken, NJ 07735

David Dai Tran  
NVIDIA Corporation  
Hoboken, NJ 07735

Ben Flagg  
NVIDIA Corporation  
Hoboken, NJ 07735

Tim Miller  
NVIDIA Corporation  
Hoboken, NJ 07735

David Dworkowski  
NVIDIA Corporation  
Hoboken, NJ 07735

Pranav Goyal  
NVIDIA Corporation  
Hoboken, NJ 07735

Abhishek Prasad  
NVIDIA Corporation  
Hoboken, NJ 07735

Karel Hala  
NVIDIA Corporation  
Hoboken, NJ 07735

Bernhard Peters  
NVIDIA Corporation  
Hoboken, NJ 07735

Lawrence D. Jacob  
NVIDIA Corporation  
Hoboken, NJ 07735

Kim Zhang  
NVIDIA Corporation  
Hoboken, NJ 07735

John Xiao  
NVIDIA Corporation  
Hoboken, NJ 07735

Abstract

We trained a convolutional neural network (CNN) to map raw pixels from a single frame-lens camera directly to steering commands. This end-to-end approach proved surprisingly powerful. With extensive testing data from humans, the system learns to drive in traffic on local roads with or without lane markings and on highways. It also operates in areas with unclear visual guidance such as in parking lots and on unimproved roads.

The system automatically learns internal representations of the necessary processing steps such as detecting useful road features with only the human steering angle as the training signal. We never explicitly trained it to detect, for example, the width of roads.

Compared to explicit decompositions of the problem, such as lane marking detection, path planning, and control, our end-to-end system optimizes all processing steps simultaneously. We argue that this will eventually lead to better performance and smaller systems. Better performance will result because the internal components self-optimize to maximize overall system performance, instead of optimizing human-selected intermediate criteria, e.g., lane detection. Such criteria undeniably are selected for ease of human interpretation which doesn't necessarily guarantee maximum system performance. Smaller networks are possible because the system learns to solve the problem with the minimal number of processing steps.

We used an NVIDIA DevBox and Tesla T4 for training and an NVIDIA DRIVE™ PX2 self-driving car computer also running Tesla T4 for determining where to drive. The system operates at 50 frames per second (FPS).

[https://www.youtube.com/watch?v=mCmO\\_5ZxdvE](https://www.youtube.com/watch?v=mCmO_5ZxdvE)  
<https://www.youtube.com/watch?v=LVBBKppAaV4>

Slide Credit: Aykut Erdem

# ML in a Nutshell

- Tens of thousands of machine learning algorithms
  - Hundreds new every year
- Decades of ML research oversimplified:
  - All of Machine Learning:
  - Learn a mapping from input to output  $f: X \rightarrow Y$
  - $X$ : emails,  $Y$ : {spam, notspam}

# ML in a Nutshell

- Input:  $x$  (images, text, emails...)
- Output:  $y$  (spam or non-spam...)
- (Unknown) Target Function
  - $f: X \rightarrow Y$  (the “true” mapping / reality)
- Data
  - $(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$
- Model / Hypothesis Class
  - $g: X \rightarrow Y$
  - $y = g(x) = \text{sign}(w^T x)$

# ML in a Nutshell

- Every machine learning algorithm has three components:
  - Representation / Model Class
  - Evaluation / Objective Function
  - Optimization

# Representation / Model Class

- Decision trees
- Sets of rules / Logic programs
- Instances
- Graphical models (Bayes/Markov nets)
- Neural networks
- Support vector machines
- Model ensembles
- Etc.



# Evaluation / Objective Function

- Accuracy
- Precision and recall
- Squared error
- Likelihood
- Posterior probability
- Cost / Utility
- Margin
- Entropy
- K-L divergence
- Etc.


# Optimization

- Discrete/Combinatorial optimization
  - greedy search
  - Graph algorithms (cuts, flows, etc)
- Continuous optimization
  - Convex/Non-convex optimization
  - Linear programming

# Types of Learning

- Supervised learning
  - Training data includes desired outputs
- Unsupervised learning
  - Training data does not include desired outputs
- Weakly or Semi-supervised learning
  - Training data includes a few desired outputs
- Reinforcement learning
  - Rewards from sequence of actions

# Spam vs Regular Email

**Sebring, Tracy**   
To: Batra, Dhruv  
ECE 4424 proposal

January 21, 2015 2:53 PM  
[Hide Details](#)

CUSP has approved ECE 4424 with the  
copy of the proposal with these items ad  
Thanks!!!  
Tracy

**nadia bamba**  
To: undisclosed recipients: ;  
Reply-To: nadia bamba  
From Miss Nadia BamBa,

January 19, 2015 5:57 AM  
[Hide Details](#)

From Miss Nadia BamBa,

Greeting, Permit me to inform you of my desire of going into business relationship with you. I am Nadia BamBa the only Daughter of late Mr and Mrs James BamBa, My father was a director of cocoa merchant in Abidjan, the economic capital of Ivory Coast before he was poisoned to death by his business associates on one of their outing to discus on a business deal. When my mother died on the 21st October 2002, my father took me very special because i am motherless.

Before the death of my father in a private hospital here in Abidjan, He secretly called me on his bedside and told me that he had a sum of \$6, 8000.000(SIX Million EIGHT HUNDRED THOUSAND), Dollars) left in a suspense account in a Bank here in Abidjan, that he used my name as his first Daughter for the next of kin in deposit of the fund.

He also explained to me that it was because of this wealth and some huge amount of money That his business associates supposed to balance him from the deal they had that he was poisoned by his business associates, that I should seek for a God fearing foreign partner in a country of my choice where I will transfer this money and use it for investment purposes, (such as real estate Or Hotel management).please i am honourably seeking your assistance in the following ways.

- 1) To provide a Bank account where this money would be transferred to.
- 2) To serve as the guardian of this Money since I am a girl of 19 years old.
- 3)Your private phone number's and your family background's that we can know each order more.

Moreover i am willing to offer you 15% of the total sum as compensation for effort input after the successful transfer of this fund to your designated account overseas,

Anticipating to hear from you soon.  
Thanks and God Bless.  
Best regards.

VS

# Intuition

- Spam Emails
  - a lot of words like
    - “money”
    - “free”
    - “bank account”
    - “viagara” ... in a single email
- Regular Emails
  - word usage pattern is more spread out

# Simple Strategy: Let us count!

## This is X

I know you will be very much interested, kindly provide me with the details below.

First Name.....  
Surname.....  
Address.....  
City.....  
State/Province.....  
Country.....  
Telephone No.....  
Occupation .....  
Date of Birth (date/m/yr) .....  
Copy of International Passport Or ID card.....

$$\begin{pmatrix} \text{free} & 100 \\ \text{money} & 2 \\ \vdots & \vdots \\ \text{account} & 2 \\ \vdots & \vdots \end{pmatrix}$$



**From:** Ross Girshick  
**Subject:** Re: hey  
**Date:** January 17, 2013 7:48:18 PM EST  
**To:** Dhruv Batra

Hi Dhruv,

sorry for the high latency. I just got back from Singapore last night ar

$$\begin{pmatrix} \text{free} & 1 \\ \text{money} & 1 \\ \vdots & \vdots \\ \text{account} & 2 \\ \vdots & \vdots \end{pmatrix}$$

# Final Procedure



$$\begin{pmatrix} \text{free} & 100 \\ \text{money} & 2 \\ \vdots & \vdots \\ \text{account} & 2 \\ \vdots & \vdots \end{pmatrix}$$

Why these words?

$$\begin{pmatrix} 100 \times 0.2 \\ 2 \times 0.3 \\ \vdots \\ 2 \times 0.3 \\ \vdots \end{pmatrix}$$

= 3.2



Confidence /  
performance guarantee?

= 1.03

$$\begin{pmatrix} 100 \times 0.01 \\ 2 \times 0.02 \\ \vdots \\ 2 \times 0.01 \\ \vdots \end{pmatrix}$$

Why linear combination?

Where do the weights  
come from?

# Types of Learning

- Supervised learning
  - Training data includes desired outputs
- Unsupervised learning
  - Training data does not include desired outputs
- Weakly or Semi-supervised learning
  - Training data includes a few desired outputs
- Reinforcement learning
  - Rewards from sequence of actions



# Tasks

## Supervised Learning



## Unsupervised Learning



# Classification

- From data to discrete classes



# Spam filtering

data

prediction

★ **Osman Khan** to Carlos [show details](#) Jan 7 (6 days ago) [Reply](#) ▼

sounds good  
+ok

Carlos Guestrin wrote:  
Let's try to chat on Friday a little to coordinate and more on Sunday in person?

Carlos

## Welcome to New Media Installation: Art that Learns

★ **Carlos Guestrin** to 10615-announce, Osman, Miche [show details](#) 3:15 PM (8 hours ago) [Reply](#) ▼

Hi everyone,

Welcome to New Media Installation:Art that Learns

The class will start tomorrow.  
\*\*\*Make sure you attend the first class, even if you are on the Wait List.\*\*\*  
The classes are held in Doherty Hall C316, and will be Tue, Thu 01:30-4:20 PM.

By now, you should be subscribed to our course mailing list: [10615-announce@cs.cmu.edu](mailto:10615-announce@cs.cmu.edu).  
You can contact the instructors by emailing: [10615-instructors@cs.cmu.edu](mailto:10615-instructors@cs.cmu.edu)

**Natural \_LoseWeight SuperFood Endorsed by Oprah Winfrey, Free Trial 1 bottle, pay only \$5.95 for shipping mfw rik** [Spam](#) | [X](#)

★ **Jaquelyn Halley** to nherlein, bcc: thehorney, bcc: anç [show details](#) 9:52 PM (1 hour ago) [Reply](#) ▼

=== Natural WeightLOSS Solution ===

Vital Acai is a natural WeightLOSS product that Enables people to lose wieght and cleansing their bodies faster than most other products on the market.

Here are some of the benefits of Vital Acai that You might not be aware of. These benefits have helped people who have been using Vital Acai daily to Achieve goals and reach new heights in there dieting that they never thought they could.

- \* Rapid WeightLOSS
- \* Increased metabolism - BurnFat & calories easily!
- \* Better Mood and Attitude
- \* More Self Confidence
- \* Cleanse and Detoxify Your Body
- \* Much More Energy
- \* BetterSexLife
- \* A Natural Colon Cleanse



Spam  
vs.  
Not Spam

# Face detection



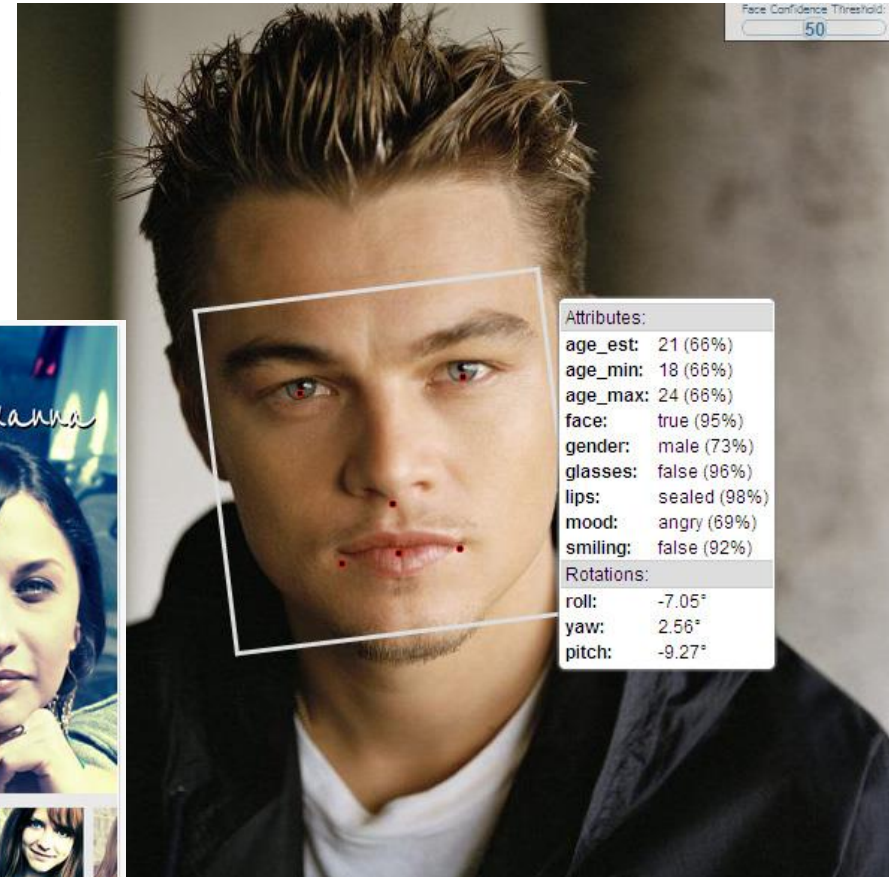
Example training images  
for each orientation



# Face Recognition



<http://developers.face.com/tools/>



Slide Credit: Noah Snaveley



# Speech Recognition



# Weather prediction



# Image Classification

- Im2tags; Im2text
- <http://deeplearning.cs.toronto.edu/>



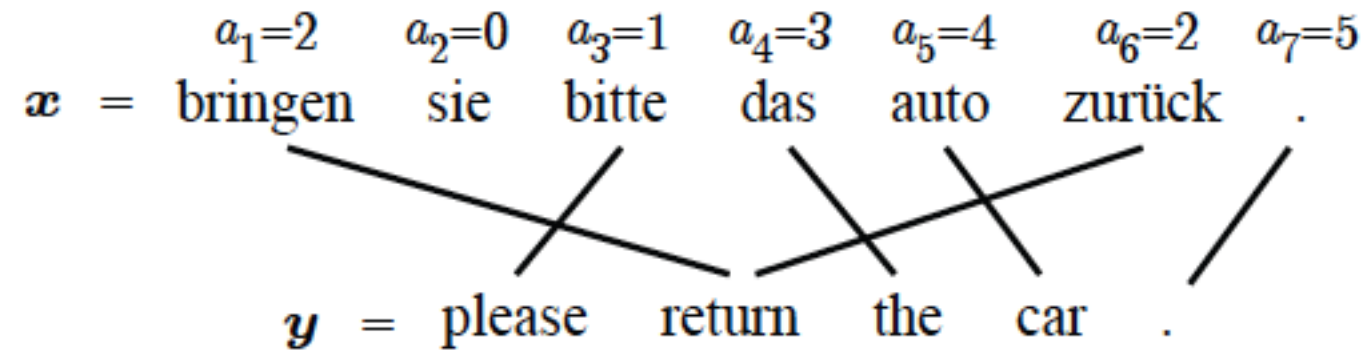
(C) Dhruv Batra



Pizza  
Wine  
Stove

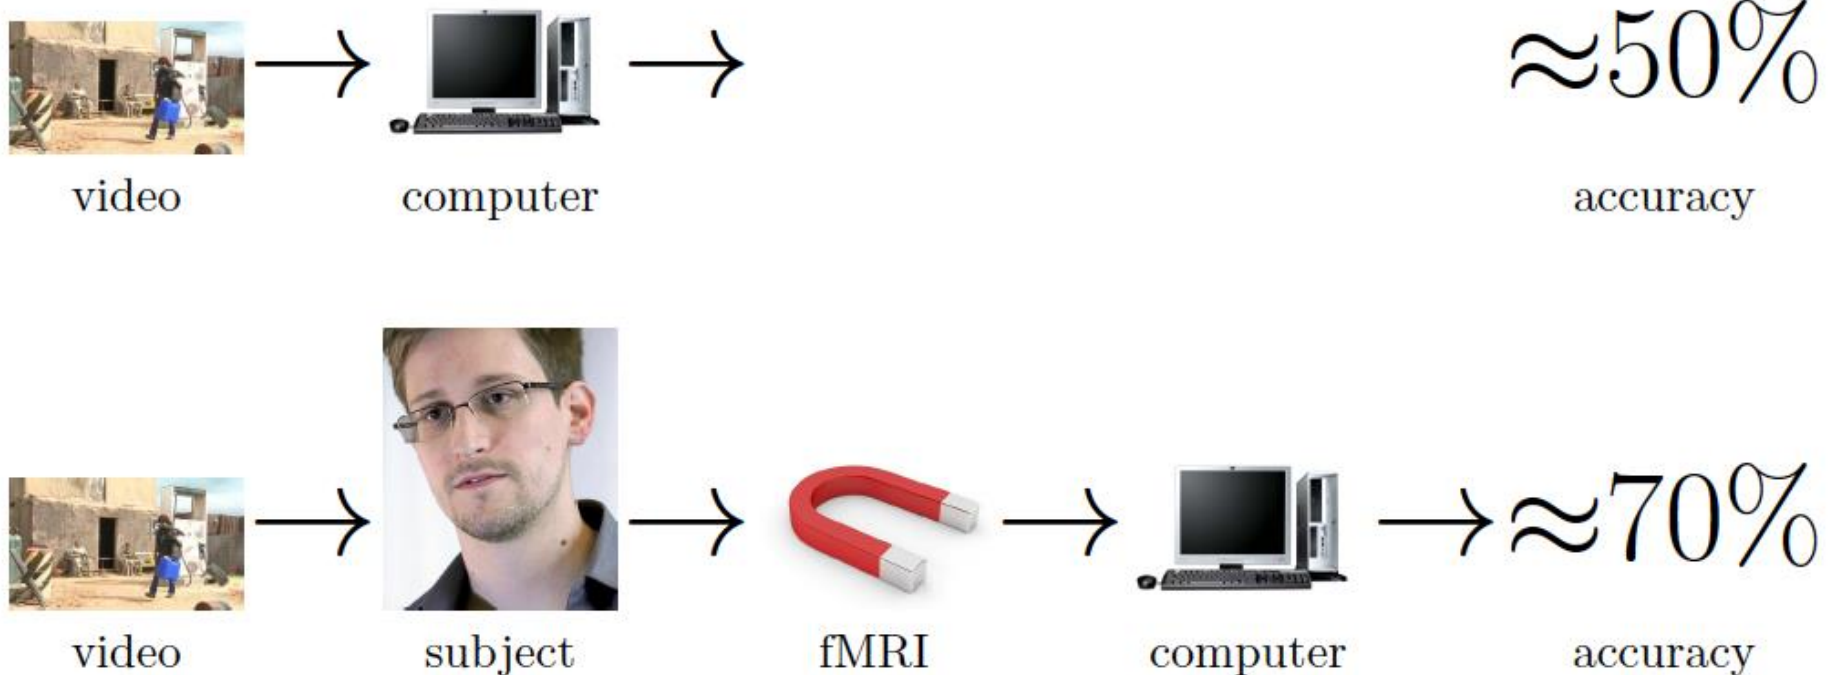


# Machine Translation



# Seeing is worse than believing

- [Barbu et al. ECCV14]



# Regression

- predicting a numeric value



# Stock market

Google Inc (NASDAQ:GOOG)

Add to portfolio

More results

**744.00** +41.13 (5.85%)

Real-time: 10:43AM EST  
NASDAQ real-time data - Disclaimer  
Currency in USD

Range 735.79 - 747.99  
52 week 556.52 - 774.38  
Open 735.99  
Vol / Avg. 2.68M/2.28M  
Mkt cap 244.39B  
P/E 22.91

Div/yield -  
EPS 32.46  
Shares 328.59M  
Beta 1.08  
Inst. own 69%



Dow Jones	13,758.94	0.34%	
Nasdaq	3,151.72	0.27%	
Technology		0.33%	
GOOG	744.00	5.85%	




- A** Google Inc. (GOOG) Is Up Sharply On Q4 Results  
RTT News - 1 hour ago
  - B** Stocks to Watch: Google, Coach, Annie's  
Wall Street Journal - 1 hour ago
  - C** Google Inc (GOOG) Reports Strong Earnings, Shares Rise  
ValueWalk - 3 hours ago
  - D** Google 4th-Quarter Profits Increase as Ad Pricing Improves  
NASDAQ - 15 hours ago
  - E** Facebook Inc (FB)'s Social Graph Is a Google Inc (GOOG) Plus Killer  
Insider Monkey - 16 hours ago
- [All news for Google Inc »](#) [Subscribe](#)

Events [Add GOOG to my calendars](#)

Apr 15, 2013  
Q1 2013 Google Earnings Release

# Weather prediction



Temperature

27°C

# Pose Estimation



# Pose Estimation

- 2010: (Project Natal) Kinect
  - <http://www.youtube.com/watch?v=r5-zZDSsgFg>
- 2012: Kinect One
  - <http://youtu.be/Hi5kMNfgDS4?t=28s>
- 2013: Leap Motion
  - <http://youtu.be/gby6hGZb3ww>

# Ranking

- Comparing items



# Web search

The image is a screenshot of a Google search interface. At the top left is the Google logo. To its right is a search bar containing the text "learning to rank". Below the search bar is a dropdown menu with five suggestions: "learning to rank", "learning to rank for information retrieval", "learning to rank using gradient descent", and "learning to rank tutorial". To the right of the search bar is a blue button with a magnifying glass icon. Below the search bar is a "Search" button. On the left side of the page, there is a vertical sidebar with links to "Web", "Images", "Maps", "Videos", "News", "Shopping", and "More". Below these links, there is a section for "Manhattan, NY 10012" with a "Change location" link. At the bottom of the sidebar is a "Show search tools" link. The main content area on the right displays search results. The first result is titled "Learning to rank - Wikipedia, the free encyclopedia" with a green URL "en.wikipedia.org/wiki/Learning\_to\_rank". The text below the title says "Learning to rank or machine-learned ranking (MLR) is a type of supervised or semi-supervised machine learning problem in which the goal is to automatically ...". Below this text are four links: "Applications", "Feature vectors", "Evaluation measures", and "Approaches". The second result is titled "Yahoo! Learning to Rank Challenge" with a green URL "learningtorankchallenge.yahoo.com/". The text below the title says "Learning to Rank Challenge is closed! Close competition, innovative ideas, and fierce determination were some of the highlights of the first ever Yahoo!". The third result is titled "[PDF] Large Scale Learning to Rank" with a green URL "www.eecs.tufts.edu/~dsculley/papers/large-scale-rank.pdf". The text below the title says "File Format: PDF/Adobe Acrobat - Quick View by D Sculley - Cited by 24 - Related articles". The text below this says "Pairwise learning to rank methods such as RankSVM give good performance, ... In this paper, we are concerned with learning to rank methods that can learn on ...". The fourth result is titled "Microsoft Learning to Rank Datasets - Microsoft Research" with a green URL "research.microsoft.com/en-us/projects/mslr/". The text below the title says "We release two large scale datasets for research on learning to rank: L2R-WEB30k with more than 30000 queries and a random sampling of it L2R-WEB10K ...". The fifth result is titled "LETOR: A Benchmark Collection for Research on Learning to Rank ..." with a green URL "research.microsoft.com/~letor/". The text below the title says "This website is designed to facilitate research in LEarning TO Rank (LETOR). Much information about learning to rank can be found in the website, including ...".

Google

learning to rank

learning to rank

learning to rank for information retrieval

learning to rank using gradient descent

learning to rank tutorial

I'm Feeling Lucky »

Search

Web

Images

Maps

Videos

News

Shopping

More

Manhattan, NY 10012

Change location

Show search tools

[Learning to rank - Wikipedia, the free encyclopedia](#)  
[en.wikipedia.org/wiki/Learning\\_to\\_rank](http://en.wikipedia.org/wiki/Learning_to_rank)  
**Learning to rank** or machine-learned ranking (MLR) is a type of supervised or semi-supervised machine learning problem in which the goal is to automatically ...  
[Applications](#) [Feature vectors](#) [Evaluation measures](#) [Approaches](#)

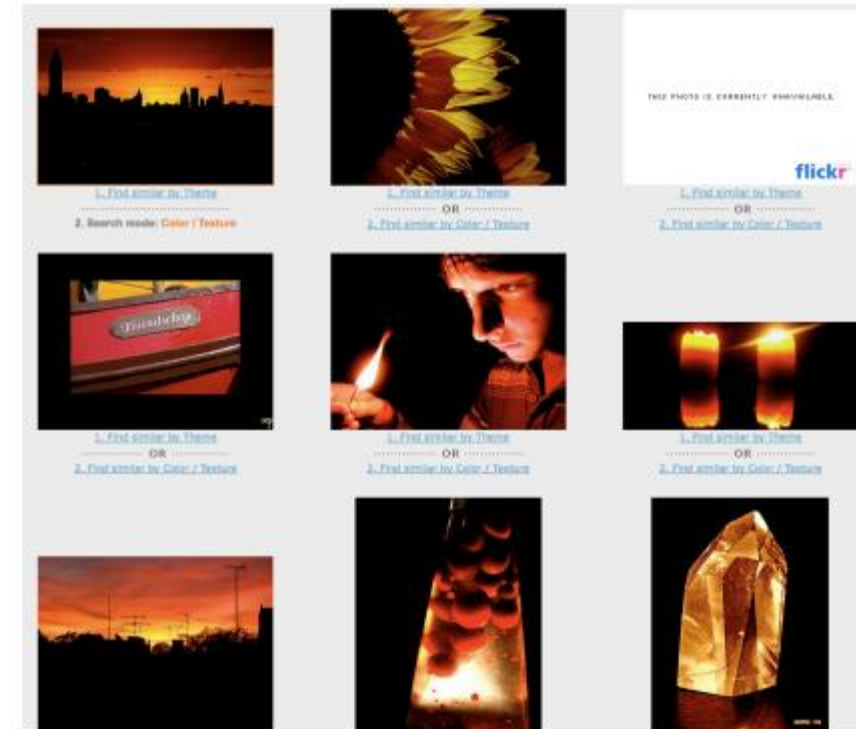
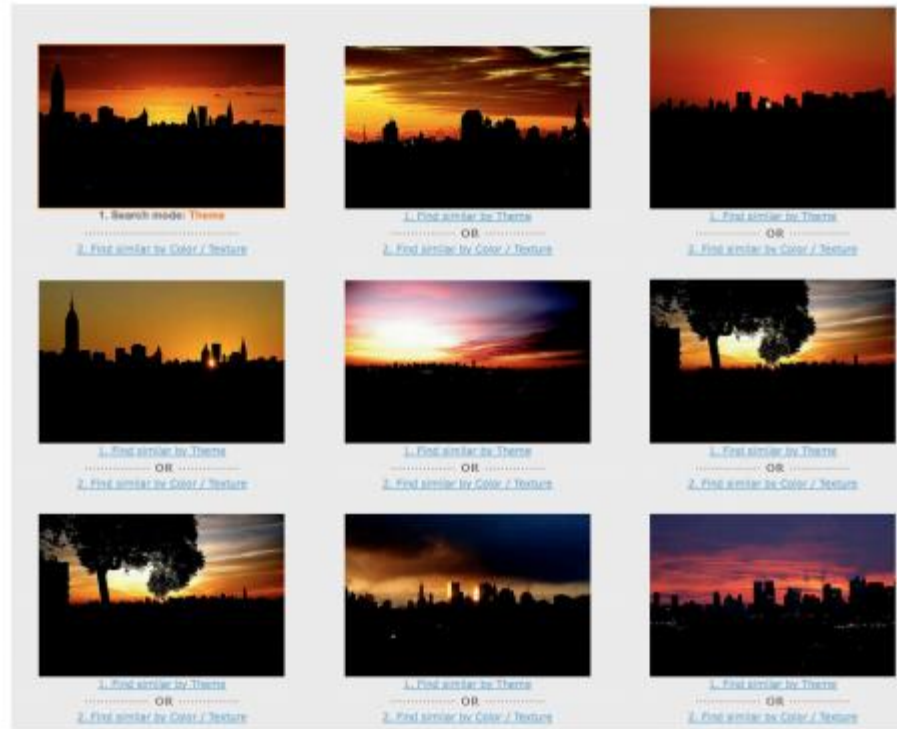
[Yahoo! Learning to Rank Challenge](#)  
[learningtorankchallenge.yahoo.com/](http://learningtorankchallenge.yahoo.com/)  
**Learning to Rank** Challenge is closed! Close competition, innovative ideas, and fierce determination were some of the highlights of the first ever Yahoo!

[\[PDF\] Large Scale Learning to Rank](#)  
[www.eecs.tufts.edu/~dsculley/papers/large-scale-rank.pdf](http://www.eecs.tufts.edu/~dsculley/papers/large-scale-rank.pdf)  
File Format: PDF/Adobe Acrobat - [Quick View](#)  
by D Sculley - [Cited by 24](#) - [Related articles](#)  
Pairwise **learning to rank** methods such as RankSVM give good performance, ... In this paper, we are concerned with **learning to rank** methods that can learn on ...

[Microsoft Learning to Rank Datasets - Microsoft Research](#)  
[research.microsoft.com/en-us/projects/mslr/](http://research.microsoft.com/en-us/projects/mslr/)  
We release two large scale datasets for research on **learning to rank**: L2R-WEB30k with more than 30000 queries and a random sampling of it L2R-WEB10K ...

[LETOR: A Benchmark Collection for Research on Learning to Rank ...](#)  
[research.microsoft.com/~letor/](http://research.microsoft.com/~letor/)  
This website is designed to facilitate research in **LEarning TO Rank** (LETOR). Much information about **learning to rank** can be found in the website, including ...

# Given image, find similar images



<http://www.tiltomo.com/>

Slide Credit: David Sontag

# Recommendation systems

amazon [Try Prime](#) [David's Amazon.com](#) [Today's Deals](#) [Gift Cards](#) [Sell](#) [Help](#) [Daily Lightning Deals](#) [Back-to-School Savings](#) [Shop now](#)

Shop by Department [▼](#) Search [Books](#) [Go](#) Hello, [David](#) [Your Account](#) [Try Prime](#) [Cart](#) [Wish List](#)

[Your Amazon.com](#) [Your Browsing History](#) [Recommended For You](#) [Amazon Beterizer](#) [Improve Your Recommendations](#) [Your Profile](#) [Learn More](#)

[Your Amazon.com](#) > [Recommended for You](#) > [Books](#) > [Subjects](#) > [Science & Math](#) > [History & Philosophy](#)

**Just For Today**  
[Browse Recommended](#)

**Recommendations**  
**History & Philosophy**  
[History of Science](#)  
[Philosophy of Biology](#)  
[Philosophy of Medicine](#)

These recommendations are based on [items you own](#) and more.

view: [All](#) | [New Releases](#) | [Coming Soon](#)

- 

**Causality: Models, Reasoning and Inference**  
by Judea Pearl (September 14, 2009)  
Average Customer Review: [★★★★★](#) (18)  
In Stock

**List Price:** \$50.00  
**Price:** ~~\$32.49~~  
[61 used & new](#) from \$28.00

[Add to Cart](#) [Add to Wish List](#)

☐ I own it ☐ Not interested ☒ [★★★★★](#) Rate this item  
Recommended because you purchased [Probabilistic Graphical Models](#) and more (Fix this)
- 

**The Lady Tasting Tea: How Statistics Revolutionized Science in the Twentieth Century**  
by David Salsburg (May 1, 2002)  
Average Customer Review: [★★★★★](#) (75)  
In Stock

**List Price:** \$18.99  
**Price:** ~~\$13.88~~  
[81 used & new](#) from \$9.00

[Add to Cart](#) [Add to Wish List](#)

☐ I own it ☐ Not interested ☒ [★★★★★](#) Rate this item  
Recommended because you added [The Theory That Would Not Die](#) to your Wish List (Fix this)
- 

**The Eighth Day of Creation: Makers of the Revolution in Biology, 25th Anniversary Edition**  
by Horace Freeland Judson (November 1, 1996)  
Average Customer Review: [★★★★★](#) (10)  
In stock on September 4, 2013

**List Price:** \$56.00  
**Price:** ~~\$36.09~~  
[59 used & new](#) from \$26.95

[Add to Cart](#) [Add to Wish List](#)

☐ I own it ☐ Not interested ☒ [★★★★★](#) Rate this item  
Recommended because you purchased [Molecular Biology of the Cell](#) (Fix this)
- 

**The Machinery of Life**  
by David S. Goodsell (April 28, 2009)  
Average Customer Review: [★★★★★](#) (41)  
In Stock

**List Price:** \$25.00  
**Price:** ~~\$17.49~~  
[92 used & new](#) from \$12.00

[Add to Cart](#) [Add to Wish List](#)



# Recommendation systems

Machine learning competition with a \$1 million prize

## Leaderboard

Display top 20 leaders.

Rank	Team Name	Best Score	% Improvement	Last Submit Time
1	<a href="#">The Ensemble</a>	0.8553	10.10	2009-07-26 18:38:22
2	<a href="#">BellKor's Pragmatic Chaos</a>	0.8554	10.09	2009-07-26 18:18:28
<b>Grand Prize - RMSE &lt;= 0.8563</b>				
3	<a href="#">Grand Prize Team</a>	0.8571	9.91	2009-07-24 13:07:49
4	<a href="#">Opera Solutions and Vandelay United</a>	0.8573	9.89	2009-07-25 20:05:52
5	<a href="#">Vandelay Industries I</a>	0.8579	9.83	2009-07-26 02:49:53
6	<a href="#">PragmaticTheory</a>	0.8582	9.80	2009-07-12 15:09:53
7	<a href="#">BellKor in BigChaos</a>	0.8590	9.71	2009-07-26 12:57:25
8	<a href="#">Dace</a>	0.8603	9.58	2009-07-24 17:18:43
9	<a href="#">Opera Solutions</a>	0.8611	9.49	2009-07-26 18:02:08
10	<a href="#">BellKor</a>	0.8612	9.48	2009-07-26 17:19:11
11	<a href="#">BigChaos</a>	0.8613	9.47	2009-09-23 23:06:52
12	<a href="#">Feeds2</a>	0.8613	9.47	2009-07-24 20:06:46
<b>Progress Prize 2008 - RMSE = 0.8616 - Winning Team: BellKor in BigChaos</b>				
13	<a href="#">xianqiang</a>	0.8633	9.26	2009-07-21 02:04:40
14	<a href="#">Granty</a>	0.8634	9.25	2009-07-26 15:58:34
15	<a href="#">Ces</a>	0.8642	9.17	2009-07-25 17:42:38
16	<a href="#">Invisible Ideas</a>	0.8644	9.14	2009-07-20 03:26:12
17	<a href="#">Just a guy in a garage</a>	0.8650	9.08	2009-07-22 14:10:42
18	<a href="#">Craig Carmichael</a>	0.8656	9.02	2009-07-25 16:00:54
19	<a href="#">J Dennis Su</a>	0.8658	9.00	2009-03-11 09:41:54
20	<a href="#">acmehill</a>	0.8659	8.99	2009-04-16 06:29:35
<b>Progress Prize 2002 - RMSE = 0.8712 - Winning Team: KorBell</b>				
<b>Cinematch score on quiz subset - RMSE = 0.9514</b>				

NETFLIX

Browse Recommendations Friends Queue Buy DVDs

Home Genres New Releases Previews Netflix Top 100 Cn

### Movies For You

Randy, the following movies were chosen based on your interest in:

[Bowling for Columbine](#)  
[Carnivale: Season 1](#)  
[Eisenstein 95.2](#)

**The Big One**  
★ ★ ★ ★ ★  
or subversive  
y from  
n /  
angel

**You really liked it...**  
All Discs Guaranteed!  
Now only for just \$5.99  
Shop as low  
uses  
Original art  
LIGHT

**movies**

**users**

1		?	3	5	?
?	1				2
	4		4	5	?


# Unsupervised Learning

## Discovering structure in data




Unsupervised Learning  
Y not provided

# Clustering Data: Group similar things



# Clustering images

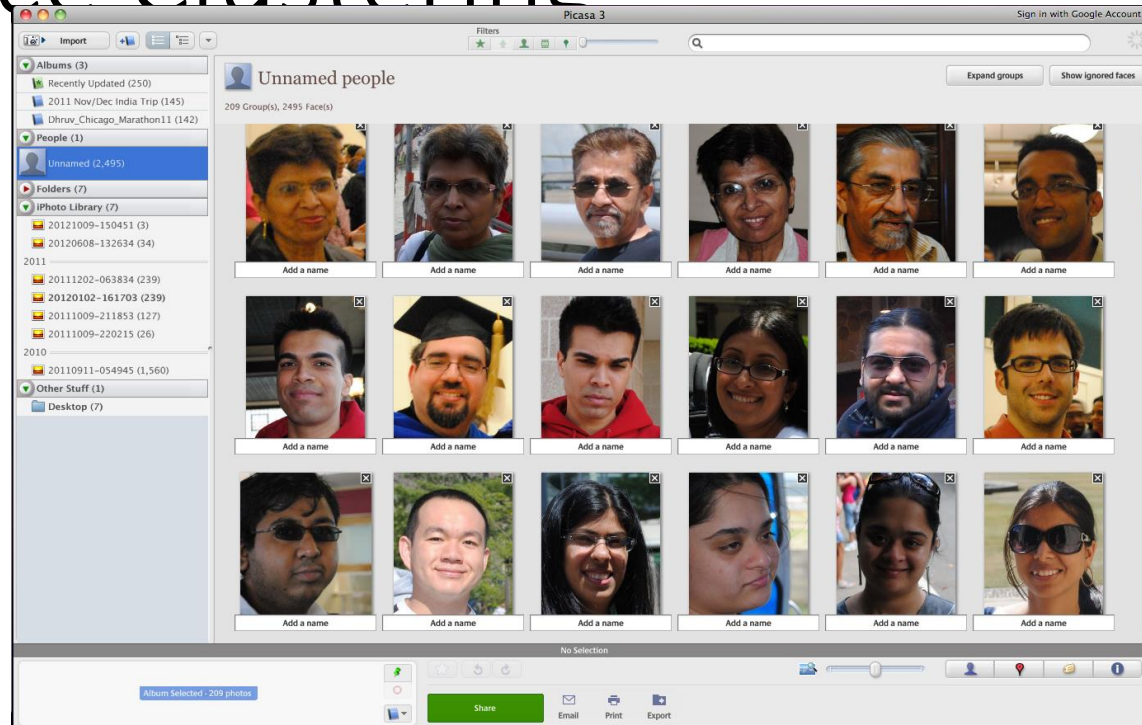


[Goldberger et al.]

Slide Credit: David Sontag



# Face Clustering



Picassa

iPhoto





# Clustering web search results

The screenshot shows the Clusty search engine interface. At the top, there's a navigation bar with links for web, news, images, wikipedia, blogs, jobs, and more. A search bar contains the word "race", and there are buttons for "Search" and "advanced preferences". Below the navigation bar, the left sidebar shows a list of clusters. The "Human" cluster is selected, showing sub-clusters like "Classification Of Human", "Statement, Evolved", and "Other Topics". The main content area displays a list of 8 documents related to the "Human" cluster. Each document entry includes a title, a brief description, and a URL. The documents are: 1. "Race (classification of human beings) - Wikipedia, the free ...", 2. "Race - Wikipedia, the free encyclopedia", 3. "Publications | Human Rights Watch", 4. "Amazon.com: Race: The Reality Of Human Differences: Vincent Sarich ...", 5. "AAPA Statement on Biological Aspects of Race", 6. "race: Definition from Answers.com", 7. "Dopefish.com", and 8. "Site for newbies as well as experienced Dopefish followers, chronicling the birth of the Dopefish, its numerous appearances in several computer games, and its eventual take-over of the human race. Maintained by Mr. Dopefish himself, Joe Siegler of Apogee Software."

web news images wikipedia blogs jobs more »

Clusty

race

Search advanced preferences

clusters sources sites

All Results (238) remix

- Car (28)
- Race cars (7)
- Photos, Races Scheduled (5)
- Game (4)
- Track (3)
- Nascar (2)
- Equipment And Safety (2)
- Other Topics (7)
- Photos (22)
- Game (14)
- Definition (13)
- Team (18)
- Human (8)**
  - Classification Of Human (2)
  - Statement, Evolved (2)
  - Other Topics (4)
- Weekend (8)
- Ethnicity And Race (7)
- Race for the Cure (8)
- Race Information (8)

more | all clusters

find in clusters:  Find

Cluster Human contains 8 documents.

Search Results

- Race (classification of human beings) - Wikipedia, the free ...**  
The term **race** or racial group usually refers to the concept of dividing **humans** into populations or groups on the basis of various sets of characteristics. The most widely used **human** racial categories are based on visible traits (especially skin color, cranial or facial features and hair texture), and self-identification. Conceptions of **race**, as well as specific ways of grouping **rac**es, vary by culture and over time, and are often controversial for scientific as well as social and political reasons. History · Modern debates · Political and ...  
[en.wikipedia.org/wiki/Race\\_\(classification\\_of\\_human\\_beings\)](http://en.wikipedia.org/wiki/Race_(classification_of_human_beings)) - [cache] - Live, Ask
- Race - Wikipedia, the free encyclopedia**  
General. **Racing** competitions The **Race** (yachting **race**), or La course du millénaire, a no-rules round-the-world sailing event; **Race** (biology), classification of flora and fauna; **Race** (classification of human beings) **Race** and ethnicity in the United States Census, official definitions of "**race**" used by the US Census Bureau; **Race** and genetics, notion of racial classifications based on genetics. Historical definitions of **race**; **Race** (bearing), the inner and outer rings of a rolling-element bearing. **RACE** in molecular biology "Rapid ... General · Surnames · Television · Music · Literature · Video games  
[en.wikipedia.org/wiki/Race](http://en.wikipedia.org/wiki/Race) - [cache] - Live, Ask
- Publications | Human Rights Watch**  
The use of torture, unlawful rendition, secret prisons, unfair trials, ... Risks to Migrants, Refugees, and Asylum Seekers in Egypt and Israel ... In the run-up to the Beijing Olympics in August 2008, ...  
[www.hrw.org/background/usa/race](http://www.hrw.org/background/usa/race) - [cache] - Ask
- Amazon.com: Race: The Reality Of Human Differences: Vincent Sarich ...**  
Amazon.com: **Race: The Reality Of Human Differences: Vincent Sarich, Frank Miele: Books ...** From Publishers Weekly Sarich, a Berkeley emeritus anthropologist, and Miele, an editor ...  
[www.amazon.com/Race-Reality-Differences-Vincent-Sarich/dp/0813340861](http://www.amazon.com/Race-Reality-Differences-Vincent-Sarich/dp/0813340861) - [cache] - Live
- AAPA Statement on Biological Aspects of Race**  
AAPA Statement on Biological Aspects of **Race** ... Published in the American Journal of Physical Anthropology, vol. 101, pp 569-570, 1996 ... PREAMBLE As scientists who study **human** evolution and variation, ...  
[www.physanth.org/positions/race.html](http://www.physanth.org/positions/race.html) - [cache] - Ask
- race: Definition from Answers.com**  
**race** n. A local geographic or global **human** population distinguished as a more or less distinct group by genetically transmitted physical  
[www.answers.com/topic/race-1](http://www.answers.com/topic/race-1) - [cache] - Live
- Dopefish.com**  
Site for newbies as well as experienced Dopefish followers, chronicling the birth of the Dopefish, its numerous appearances in several computer games, and its eventual take-over of the **human race**. Maintained by Mr. Dopefish himself, Joe Siegler of Apogee Software.  
[www.dopefish.com](http://www.dopefish.com) - [cache] - Open Directory

# Unsupervised Learning

## Dimensionality Reduction / Embedding




Unsupervised Learning  
Y not provided


# Embedding images

Images have thousands or millions of pixels.

Can we give each image a coordinate, such that similar images are near each other?



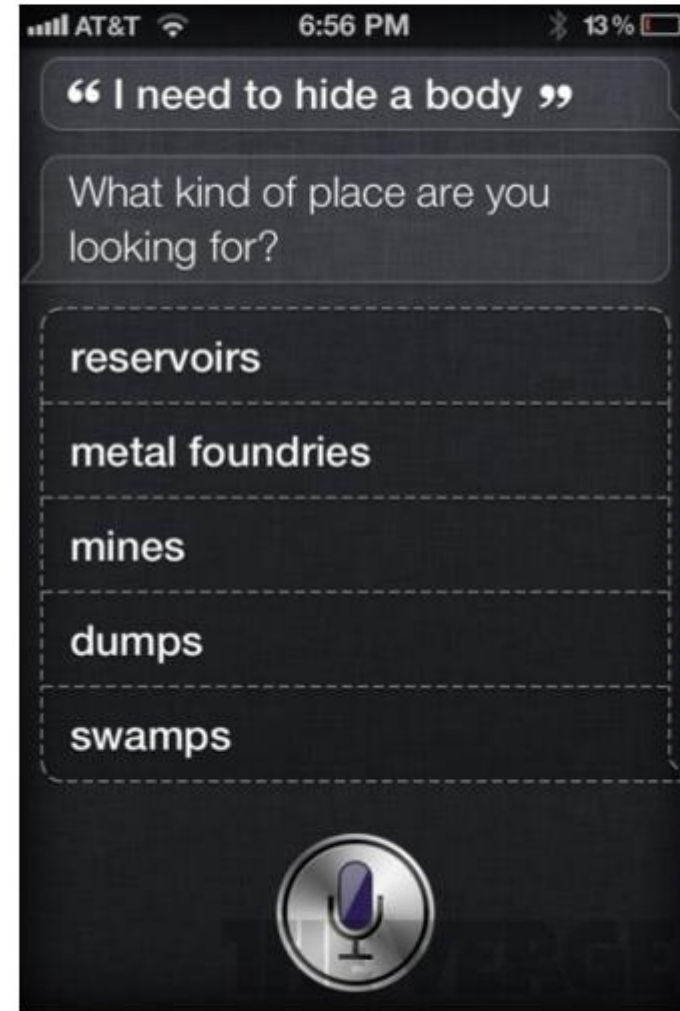

# Embedding words




# Structured prediction

- From data to discrete classes

# Speech recognition



# Reinforcement Learning



Learning from feedback



# Reinforcement Learning: Learning to act

- There is only one “supervised” signal at the end of the game.
- But you need to make a move at every step
- RL deals with “credit assignment”





# Learning to act

- Reinforcement learning
- An agent
  - Makes sensor observations
  - Must select action
  - Receives rewards
    - positive for “good” states
    - negative for “bad” states
- Towel Folding
  - <http://youtu.be/gy5g33S0Gzo>

# Growth of Machine Learning

- Machine learning is preferred approach to
  - Speech recognition, Natural language processing
  - Computer vision
  - Medical outcomes analysis
  - Robot control
  - Computational biology
  - Sensor networks
  - ...
- This trend is accelerating
  - Big data
  - Improved machine learning algorithms
  - Faster computers
  - Good open-source software

# Syllabus

- Basics of Statistical Learning
  - Loss functions, MLE, MAP, Bayesian estimation, bias-variance tradeoff, overfitting, regularization, cross-validation
- Supervised Learning
  - Nearest Neighbour, Naïve Bayes, Logistic Regression, Support Vector Machines, Kernels, Neural Networks, Decision Trees
  - Ensemble Methods: Bagging, Boosting
- Unsupervised Learning
  - Clustering: k-means, Gaussian mixture models, EM
  - Dimensionality reduction: PCA, SVD, LDA
- Advanced Topics
  - Weakly-supervised and semi-supervised learning
  - Reinforcement learning
  - Probabilistic Graphical Models: Bayes Nets, HMM
  - Applications to Vision, Natural Language Processing

# Reference Books

- A Course in Machine Learning, Hal Daumé III (**online version (v.0.99) available**), 2017
- Artificial Intelligence: A Modern Approach (3rd Edition), Russell and Norvig. Prentice Hall, 2009
- Bayesian Reasoning and Machine Learning, Barber, Cambridge University Press, 2012 (**online version available**)
- Introduction to Machine Learning (2nd Edition), Alpaydin, MIT Press, 2010
- [Pattern Recognition and Machine Learning](#), Bishop, Springer, 2006
- Machine Learning: A Probabilistic Perspective, Murphy, MIT Press, 2012

