

Introduction to Machine Learning

MSE 2400 EaLiCaRA
Dr. Tom Way

Based in part on notes from Gavin Brown, University of Manchester

What is Machine Learning?



amazon.com

facebook



Microsoft



MSE 2400 Evolution & Learning

2

Machine Learning

- a branch of artificial intelligence, is about the construction and study of systems that can learn from data.
- The ability of a computer to improve its own performance through the use of software that employs artificial intelligence techniques to mimic the ways by which humans seem to learn, such as repetition and experience.


MSE 2400 Evolution & Learning

3

Learning from Data

The world is driven by data.

- Germany's climate research centre generates 10 petabytes per year
- Google processes 24 petabytes per day
- The Large Hadron Collider produces 60 gigabytes per minute (~12 DVDs)
- There are over 50m credit card transactions a day in the US alone.




MSE 2400 Evolution & Learning

4

Learning from Data

- **Data** is recorded from some real-world phenomenon.
- What might we want to do with that data?
- **Prediction**
 - - what can we **predict** about this phenomenon?
- **Description**
 - - how can we **describe/understand** this phenomenon in a new way?



MSE 2400 Evolution & Learning

5

Learning from Data

How can we extract knowledge from data to help humans make decisions?

How can we automate decisions from data?

How can we adapt systems dynamically to enable better user experiences?

Write code to explicitly do the above tasks



Write code to make the computer **learn** how to do the tasks




MSE 2400 Evolution & Learning

6

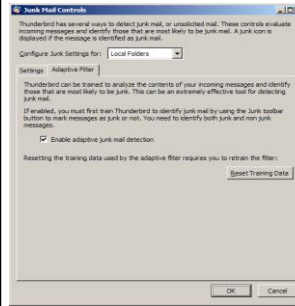
Machine Learning

- Where does it fit? What is it not?



Note that there are many definitions... this is just one way of looking at it.

- Using machine learning to detect spam emails.



To: you@gmail.com
GET YOUR DIPLOMA TODAY!
 If you are looking for a fast and cheap way to get a diploma, this is the best way out for you. Choose the desired field and degree and call us right now. For US: 1.845.709.8044 Outside US: +1.845.709.8044 "Just leave your NAME & PHONE NO. (with CountryCode)" in the voicemail. Our staff will get back to you in next few days!

ALGORITHM
 Naïve Bayes
 Rule mining

- Using machine learning to recommend books.



ALGORITHMS

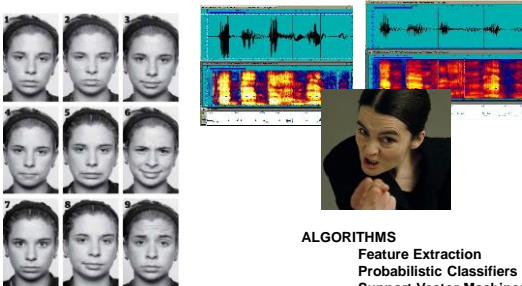
Collaborative Filtering
 Nearest Neighbour
 Clustering

- Using machine learning to identify faces and expressions.



ALGORITHMS
 Decision Trees
 Adaboost

- Using machine learning to identify vocal patterns



ALGORITHMS

Feature Extraction
 Probabilistic Classifiers
 Support Vector Machines
 + many more....

- ML for working with social network data:
 detecting fraud, predicting click-thru patterns,
 targeted advertising, etc etc etc .



facebook

amazon.com.

ebay

ALGORITHMS
 Support Vector Machines
 Collaborative filtering
 Rule mining algorithms
 Many many more....

- Driving a car
- Recognizing spam emails
- Recommending books
- Reading handwriting
- Recognizing speech, faces, etc.

What would Machine Learning have to do to make these things possible?
Even if we could, should we ?!?!?!



Many applications are immensely hard to program directly. These almost always turn out to be "pattern recognition" tasks.

1. Program the computer to do the pattern recognition task directly. ❌


1. Program the computer to be able to **learn** from examples.
2. Provide "training" data. ✅

Definition of Machine Learning

- self-configuring data structures that allow a computer to do things that would be called "intelligent" if a human did it
- "making computers behave like they do in the movies"

A Bit of History

- Arthur Samuel (1959) wrote a program that **learned** to play the game of Checkers.



1940s

Human reasoning / logic first studied as a formal subject within mathematics (Claude Shannon, Kurt Godel et al).

1950s

The "Turing Test" is proposed: a test for true machine intelligence, expected to be passed by year 2000. Various game-playing programs built. 1956 "Dartmouth conference" coins the phrase "artificial intelligence".

1960s

A.I. funding increased (mainly military). Famous quote: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."

- Ax. 1. $P(x) \wedge \square \forall y [\phi(x) \rightarrow \psi(y)] \rightarrow P(y)$
- Ax. 2. $P(\neg x) \rightarrow \neg P(x)$
- Th. 1. $P(\phi) \rightarrow \square \exists x [\phi(x)]$
- DE. 1. $G(x) \iff \forall y [P(y) \rightarrow \psi(x)]$
- Ax. 3. $P(G)$
- Th. 2. $\square \exists x G(x)$
- DE. 2. $\phi \text{ ens } x \iff \phi(x) \wedge \forall y [\phi(x) \rightarrow \square \forall y [\phi(x) \rightarrow \psi(y)]]$
- Ax. 4. $P(\phi) \rightarrow \square P(\phi)$
- Th. 3. $G(x) \rightarrow G \text{ ens } x$
- DE. 3. $E(x) \iff \forall y [\phi \text{ ens } x \rightarrow \square \exists x \phi(x)]$
- Ax. 5. $P(E)$
- Th. 4. $\square \exists x G(x)$



1970s

A.I. "winter". Funding dries up as people realize it's hard. Limited computing power and dead-end frameworks.

1980s

Revival through bio-inspired algorithms: Neural networks, Genetic Algorithms. A.I. promises the world – lots of commercial investment – mostly fails. Rule based "expert systems" used in medical / legal professions.

1990s

AI diverges into separate fields: Computer Vision, Automated Reasoning, Planning systems, Natural Language processing, **Machine Learning**...


...*Machine Learning begins to overlap with statistics / probability theory.*



$$P(A|B) = \frac{P(B|A)P(A)}{P(B)^3}$$

2000s



ML merging with statistics continues. Other subfields continue in parallel. First commercial-strength applications: Google, Amazon, computer games, route-finding, credit card fraud detection, etc... Tools adopted as standard by other fields e.g. biology



2010s.... ??????

MSE 2400 Evolution & Learning

The future?



MSE 2400 Evolution & Learning

20

Microsoft

http://www.youtube.com/watch?v=NS_L3Yyv2RI

Chris Bishop
Chief Research Scientist
Microsoft Research Cambridge
Microsoft Corp.


Chris Bishop is a Chief Research Scientist at Microsoft Research Cambridge, where he helps with the strategic direction and planning for the lab, and jointly leads the Machine Learning and Perception Group. His research interests include probabilistic approaches to machine learning, as well as their application to fields such as biomedical sciences and healthcare.

Chris is also Professor of Computer Science at the University of Edinburgh, where he is a member of the Institute for Adaptive and Neural Computation in the School of Informatics. Chris is a Fellow of Queen's College, Cambridge and a Fellow of the British Computer Society. In 2004 he was elected Fellow of the Royal Academy of Engineering, and in 2007 he was elected Fellow of the Royal Society of Edinburgh. In 2008 he was selected as the Royal Institution **MSE 2400 Evolution & Learning**

Microsoft has a MAJOR worldwide investment in Machine Learning

21

Programming language : "Matlab"



MATrix LABoratory

- Interactive scripting language
- Interpreted (i.e. no compiling)
- Objects possible, not compulsory
- Dynamically typed
- Flexible GUI / plotting framework
- Large libraries of tools
- Highly optimized for maths

MSE 2400 Evolution & Learning

22