Class 1

MAS.500 Hands on Foundations in Media Technology

Feb 5, 2014

Survey of Applied Machine Learning

People

Prof. Joseph Paradiso

joep@media.mit.edu

TA's

Artem Dementyev

artemd@media.mit.edu

Spencer Russell

sfr@media.mit.edu

Responsive Environments E14-548

Goals

- Learn how to apply machine learning
- Learn the tools
- Some theory

Schedule

About 1 month - 8 Classes

Topic 1 - Linear Classification

Topic 2 - Non linear classification

Topic 3 - Regression, ensemble methods

Toolkits: Python (mostly), Weka, Matlab,

Assignments

Weekly readings and mini projects

Class website: http://excedrin.media.mit.edu/responsive/machine-learning-module/

Machine Learning

How computers learn from data

Finding generalized functions from examples

Why machine learning?

Some problems are difficult to be explicitly programmed

Formal Definition

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E

Tom Mitchell, 1998

Machine Learning is Multidisciplinary

- Artificial Intelligence
- Statistics
- Mathematics
- Information Theory
- Philosophy
- Psychology

History

I.—COMPUTING MACHINERY AND INTELLIGENCE

By A. M. Turing

Before 1950s - The Turk

1950s - Alan Turing

1960s - Neural Networks (perceptron) - Minsky

1970s - Dormant

1980s - Revival with decision trees

1990s - today: Internet. Rapid growth

Types of Machine Learning

- Supervised
- Unsupervised
- Semi-supervised
- Reinforcement learning

Classification vs. regression

Offline vs. online

Supervised Learning

Learning from labeled examples

Algorithms:

- Artificial Neural Networks
- Support Vector Machines
- Many more...

Supervised Learning: Classification

Predicting discrete variable (Class)

Training Data

[class, height, weight, location]

e.g. [kangaroo, 1m, 40kg, Australia] ….

Input

[???, 5m, 500 kg, Africa]

Output

elephant

Examples

- Handwriting recognition
- Credit card fraud
- Gesture recognition

https://www.youtube.com/watch?v=uEyDRwG4b48

Supervised Learning: Regression

Predicting continuous variable

Input:

[Height, weight, location]

Output:

Age

Example - Brain Computer Interface

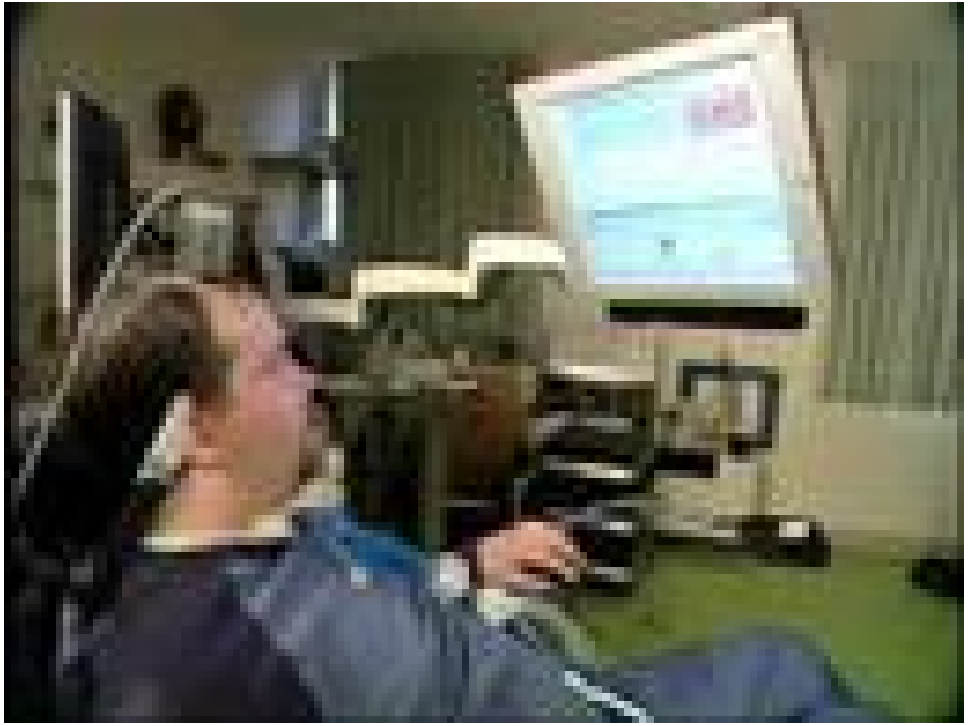

https://www.youtube.com/watch?v=TJJPbpHoPWo

Unsupervised Learning

No labeled examples

Finding structure in the data

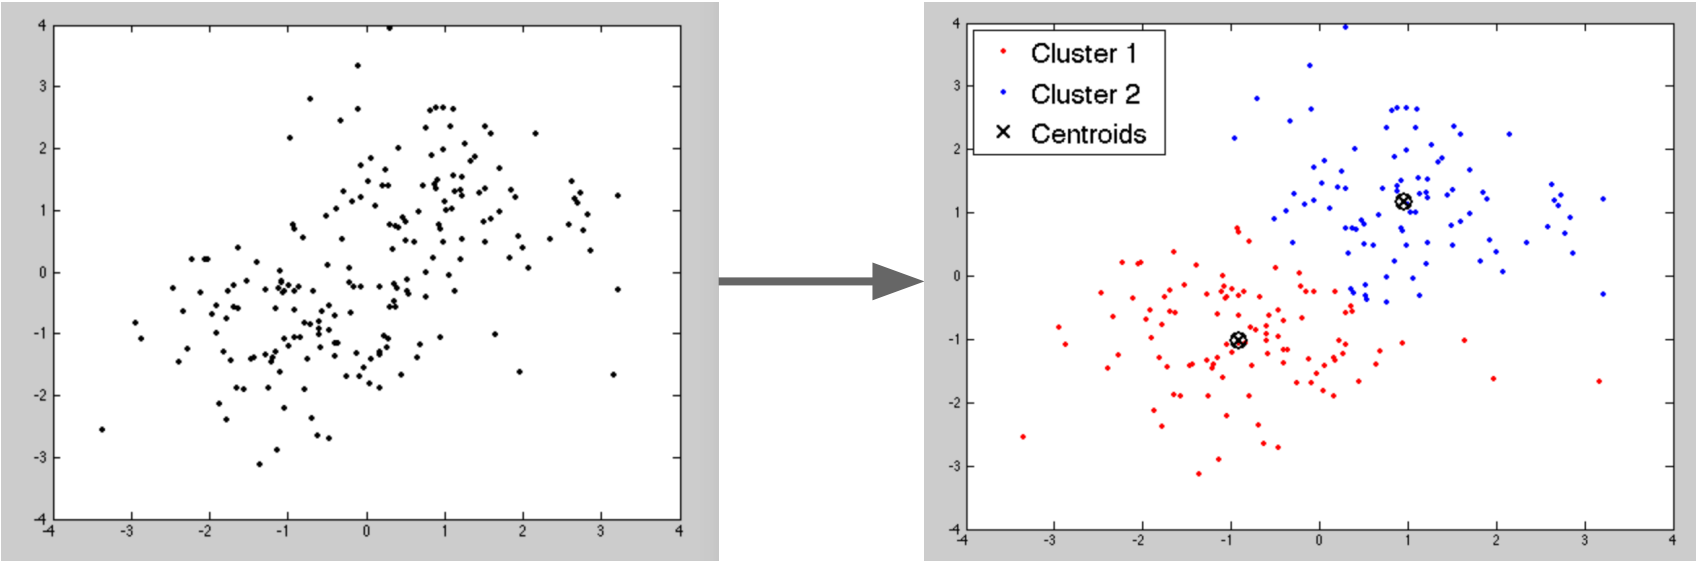

Example

Data mining applications

- Financial data
- Cancer genes
- Advertisments
- News article categorization

Semi-supervised learning

- Small number of labeled data
- Mostly unlabeled data
- Using unlabeled data to improve learning

Used where it is hard to label data

Examples

Speech Recognition

- Lots of recordings but human transcription is costly

Webpage classification

Reinforcement Learning

- Used a lot in robotics
- Learning with good and bad rewards
- Actions to maximizing good rewards

https://www.youtube.com/watch?v=W_gxLKSsSIE

Assignments

Instructions on website: http://excedrin.media.mit.edu/responsive/machine-learning-module/

1. Reading
2. Mini Project