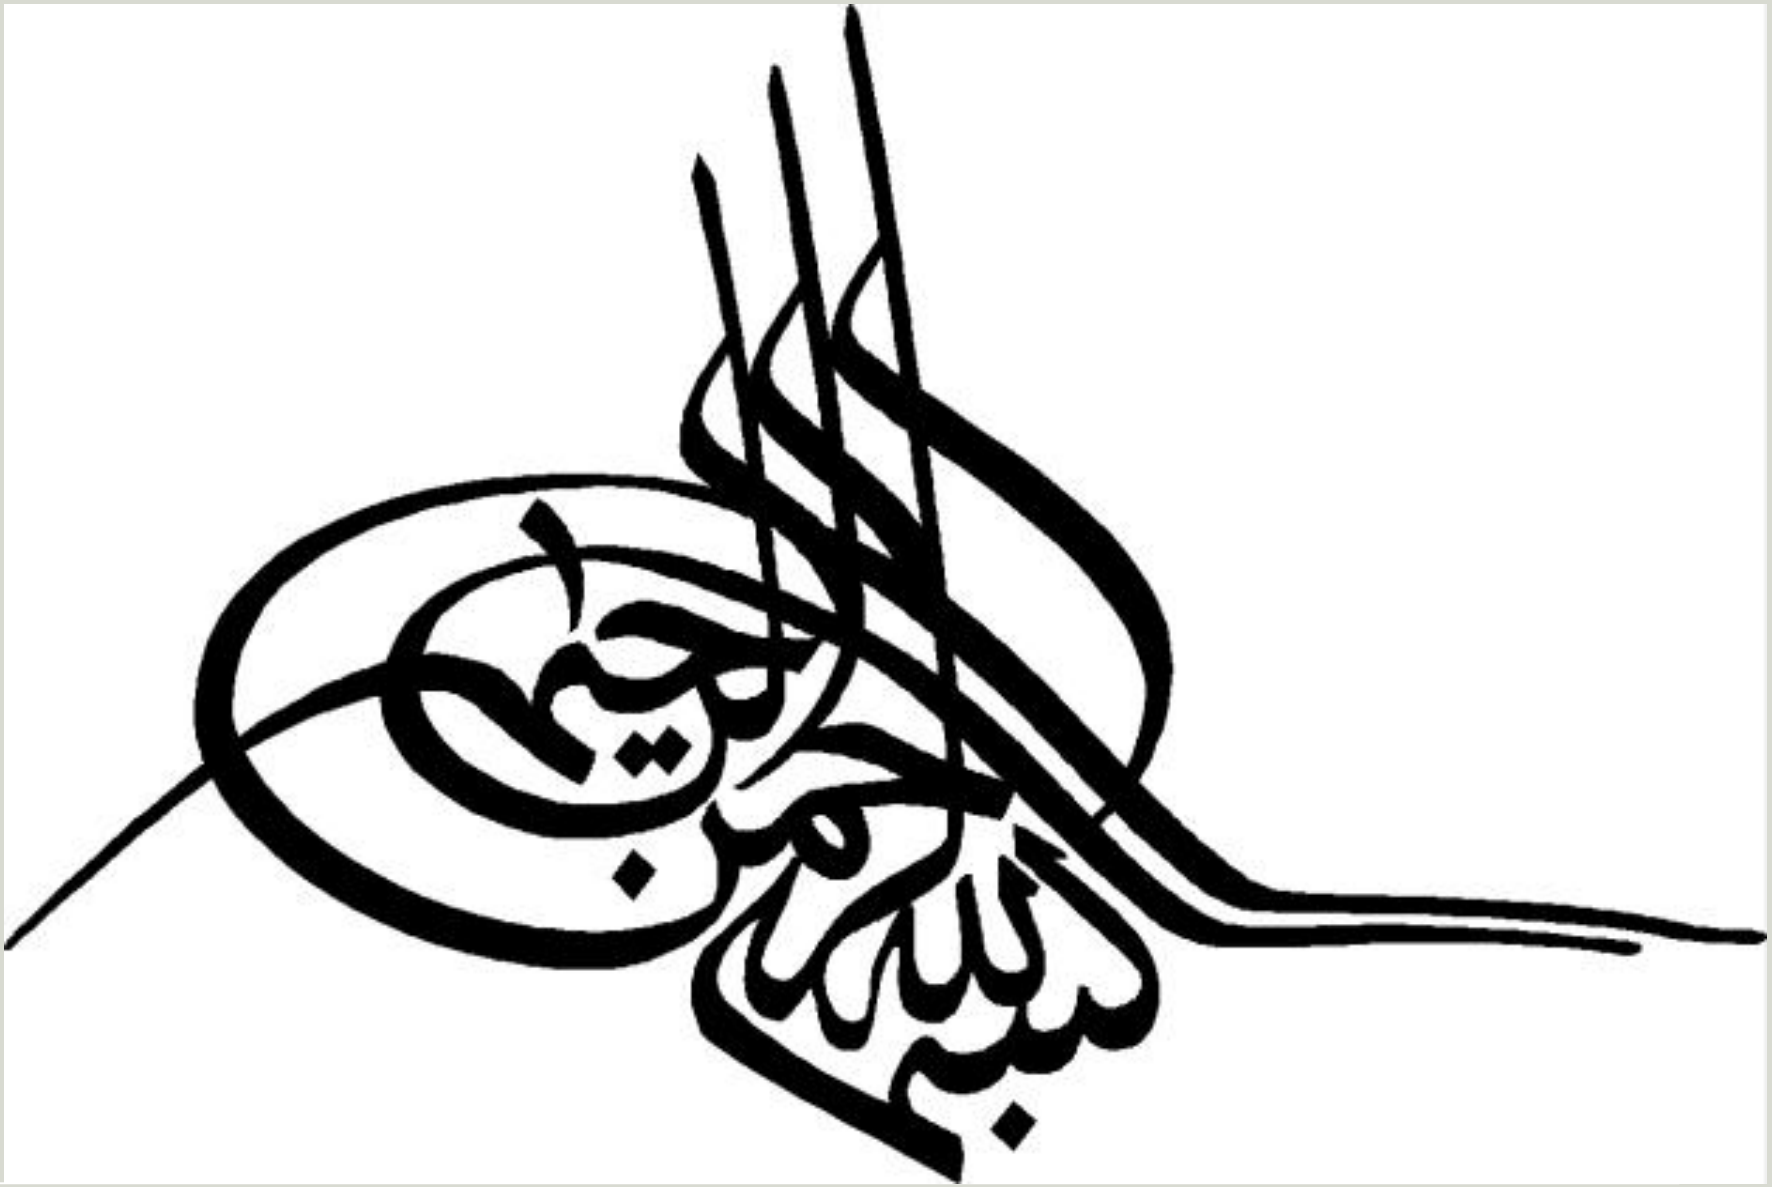
بسم الله الرحمن الرحيم

Machine Learning

M. R. Ahmadzadeh

Isfahan University of Technology

Ahmadzadeh@cc.iut.ac.ir

Textbooks

- **Introduction to Machine Learning - Ethem Alpaydin**
- **Pattern Recognition and Machine Learning, Bishop.**
- **Machine Learning, Mitchell, Tom.**
- The Elements of Statistical Learning, Hastie, T., R. Tibshirani, and J. H. Friedman.
- *Foundations of Machine Learning* by Mehryar Mohri, Afshin Rostamizadeh and Ameet Talwalkar
- Machine Learning: A Probabilistic Perspective, by Kevin P. Murphy.
- Introduction to Data Mining by Tan, Steinbach and Kumar
- Pattern Classification (2nd ed.) by Richard O. Duda, Peter E. Hart and David G. Stork
- Pattern Recognition, 4th Ed., Theodoridis and Koutroumbas

Grading Criteria

- Midterm Exam ≈ 25%
- HW, Comp. Assignments and projects: ≈ 30%
- Final exam ≈ 45%

- **Course Website:**
- http://ivut.iut.ac.ir or http://elearning.iut.ac.ir/
- **Email: Ahmadzadeh@cc.iut.ac.ir**
- EBooks…

Contents

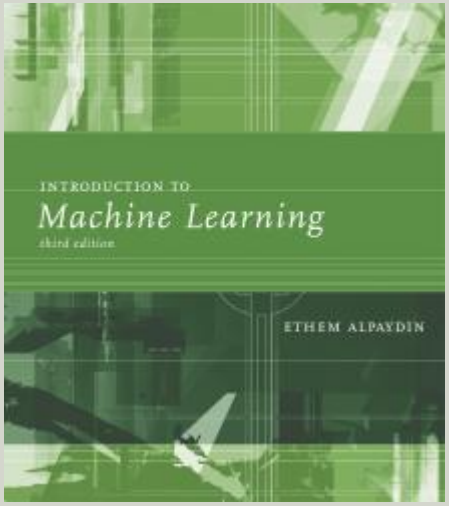

Lecture Slides for

INTRODUCTION TO MACHINE LEARNING

3RD EDITION

ETHEM ALPAYDIN

alpaydin@boun.edu.tr
http://www.cmpe.boun.edu.tr/~ethem/i2ml3e

CHAPTER 1:

INTRODUCTION

Big Data

- Widespread use of personal computers and wireless communication leads to “big data”
- We are both producers and consumers of data
- Data is not random, it has structure, e.g., customer behavior
- We need “big theory” to extract that structure from data for

  (a) Understanding the process

  (b) Making predictions for the future

Why "Learn"?

- Machine learning is programming computers to optimize a performance criterion using example data or past experience.
- There is no need to "learn" to calculate payroll
- Learning is used when:
  - Human expertise does not exist (navigating on Mars),
  - Humans are unable to explain their expertise (speech recognition)
  - Solution changes in time (routing on a computer network)
  - Solution needs to be adapted to particular cases (user biometrics)

What We Talk About When We Talk About "Learning"

- Learning general models from a data of particular examples
- Data is cheap and abundant (data warehouses, data marts); knowledge is expensive and scarce.
- Example in retail: Customer transactions to consumer behavior:

  People who bought "Blink" also bought "Outliers" (www.amazon.com)

- Build a model that is *a good and useful approximation* to the data.

Data Mining

- Retail: Market basket analysis, Customer relationship management (CRM)
- Finance: Credit scoring, fraud detection
- Manufacturing: Control, robotics, troubleshooting
- Medicine: Medical diagnosis
- Telecommunications: Spam filters, intrusion detection
- Bioinformatics: Motifs, alignment
- Web mining: Search engines
- …

What is Machine Learning?

- Optimize a performance criterion using example data or past experience.
- Role of Statistics: Inference from a sample
- Role of Computer science: Efficient algorithms to
  - Solve the optimization problem
  - Representing and evaluating the model for inference

Machine Learning vs Pattern Recognition

- Pattern Recognition: automatic discovery of regularities in data and the use of these regularities to take actions – classifying the data into different categories. Example: handwritten recognition. Input: a vector $\mathbf{x}$ of pixel values. Output: A digit from 0 to 9.
- Machine Learning: a large set of input vectors $\mathbf{x}_1, \ldots, \mathbf{x}_N$, or a training set is used to tune the parameters of an adaptive model. The category of an input vector is expressed using a target vector $\mathbf{t}$. The result of a machine learning algorithm: $y(\mathbf{x})$ where the output y is encoded as the target vectors.

Applications

- Association
- Supervised Learning
  - Classification
  - Regression
- Unsupervised Learning
- Reinforcement Learning

Learning Associations

- Basket analysis:

  $P(Y \mid X)$ probability that somebody who buys X also buys Y where X and Y are products/services.

  Example: $P(\text{Chips} \mid \text{Yogurt}) = 0.7$

Classification

- Example: Credit scoring
- Differentiating between low-risk and high-risk customers from their *income* and *savings*

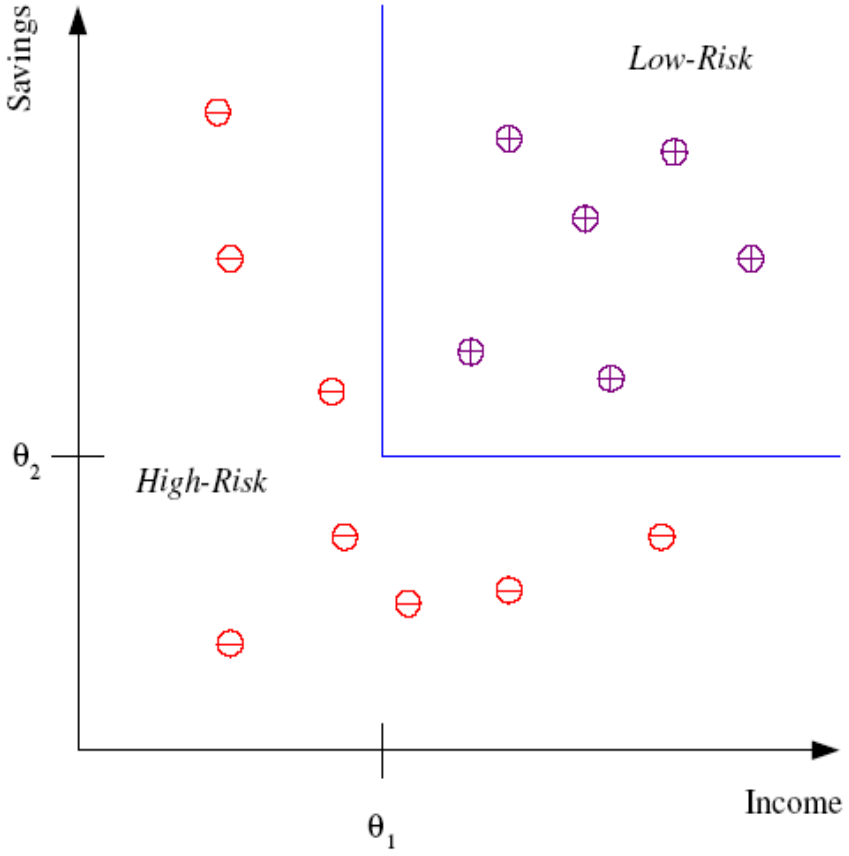

Discriminant: IF *income* > θ_1 AND *savings* > θ_2

THEN low-risk ELSE high-risk

Classification: Applications

- Face recognition: Pose, lighting, occlusion (glasses, beard), make-up, hair style
- Character recognition: Different handwriting styles.
- Speech recognition: Temporal dependency.
- Medical diagnosis: From symptoms to illnesses
- Biometrics: Recognition/authentication using physical and/or behavioral characteristics: Face, iris, signature, etc
- Outlier/novelty detection:

Face Recognition

Training examples of a person

Test images

ORL dataset,
AT&T Laboratories, Cambridge UK

A classic example of a task that requires machine learning:

It is very hard to say what makes a 2

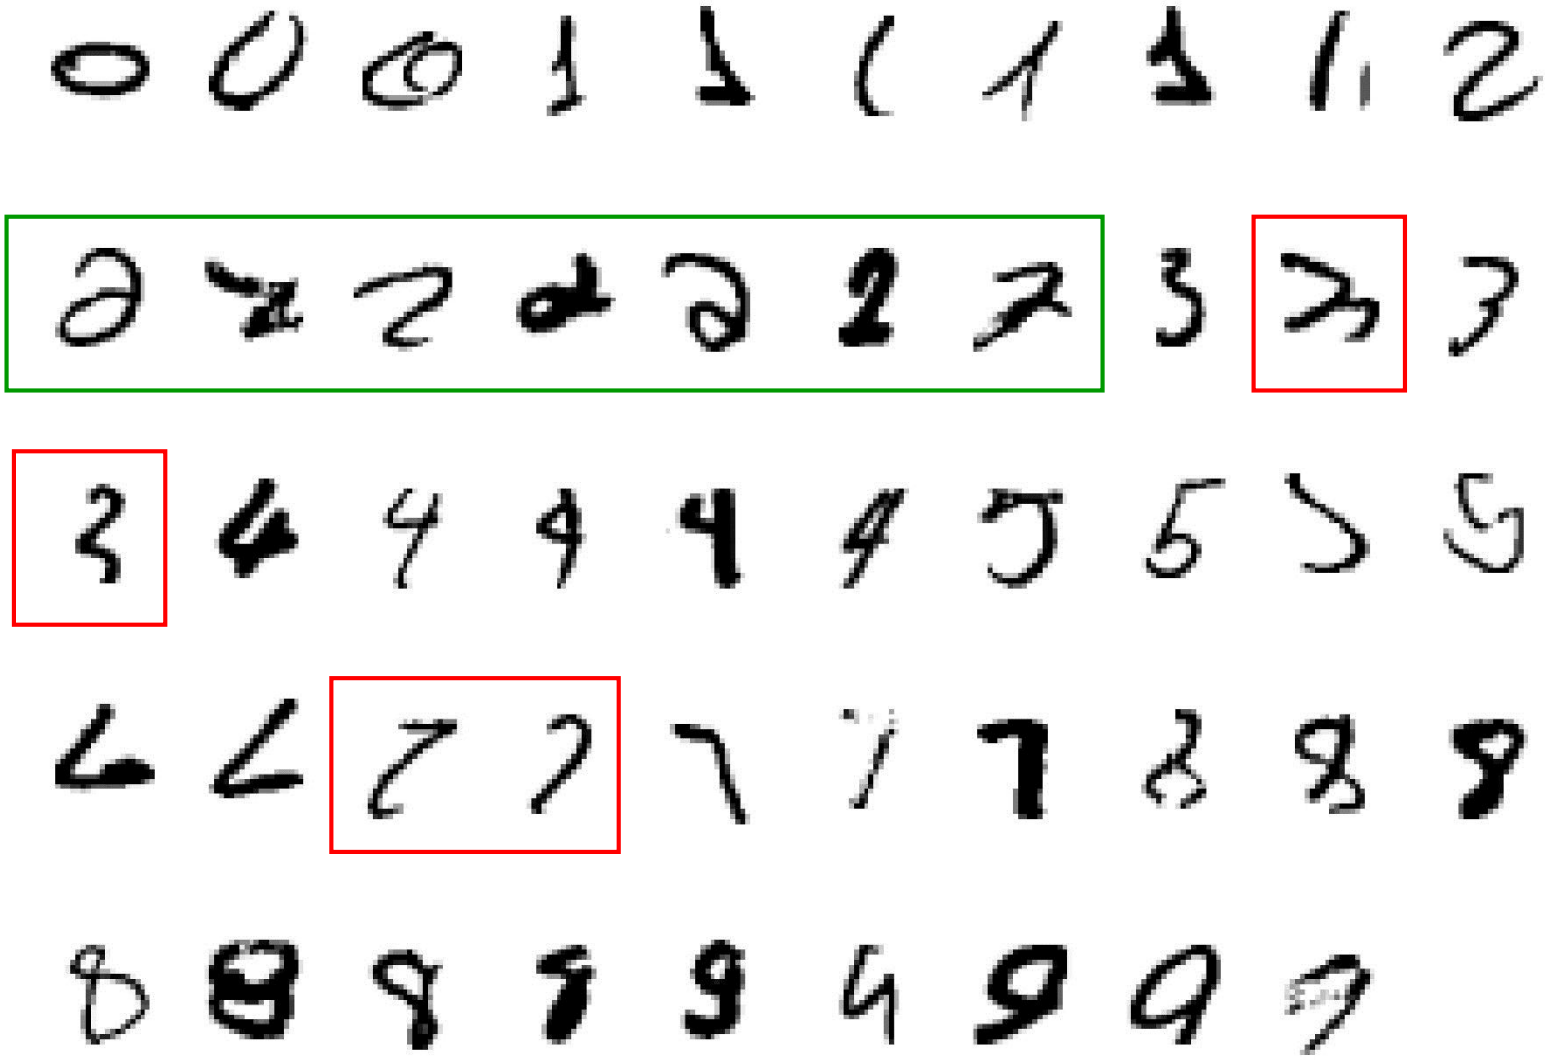

Regression

- Example: Price of a used car
- x : car attributes

  y : price

  $$y = g(x \mid \theta)$$

  $g(\ )$ model,

  θ parameters

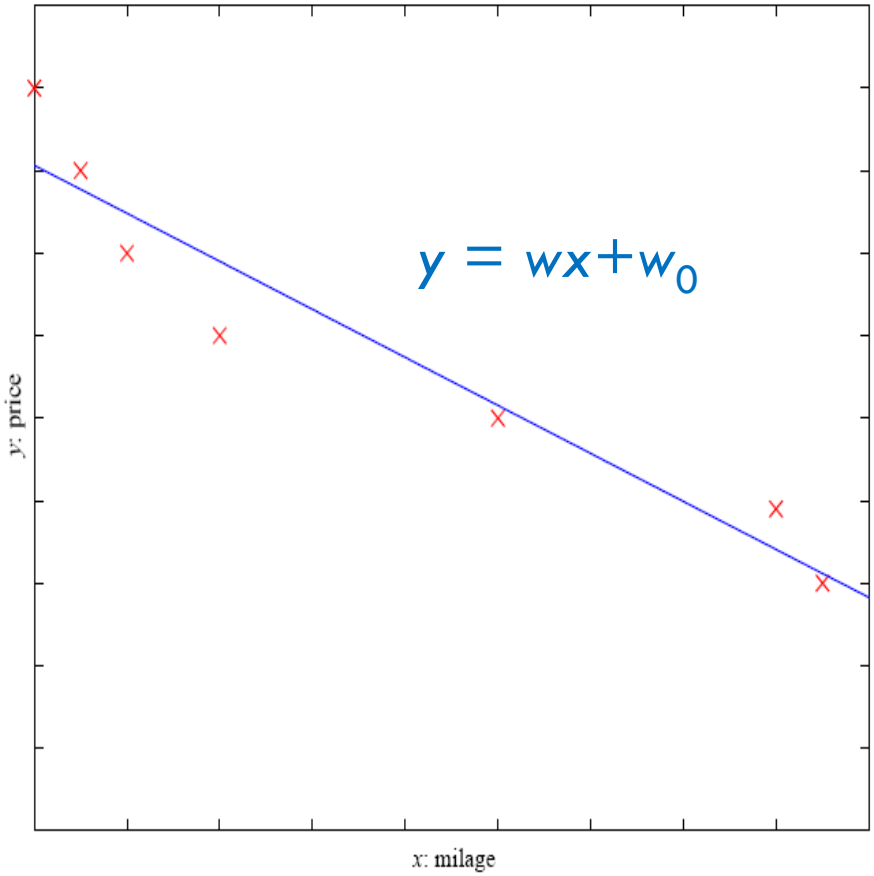

Regression Applications

- Navigating a car: Angle of the steering
- Kinematics of a robot arm

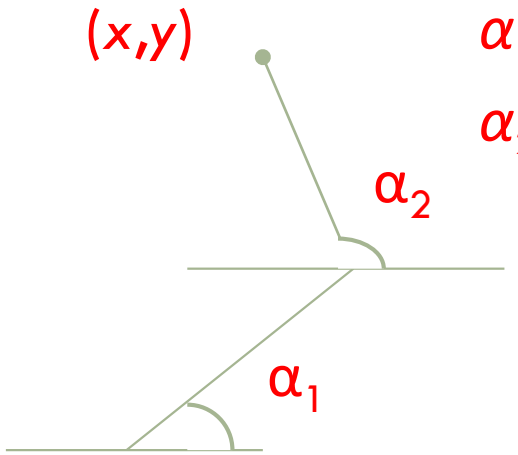

$\alpha_1 = g_1(x,y)$

$\alpha_2 = g_2(x,y)$

- Response surface design

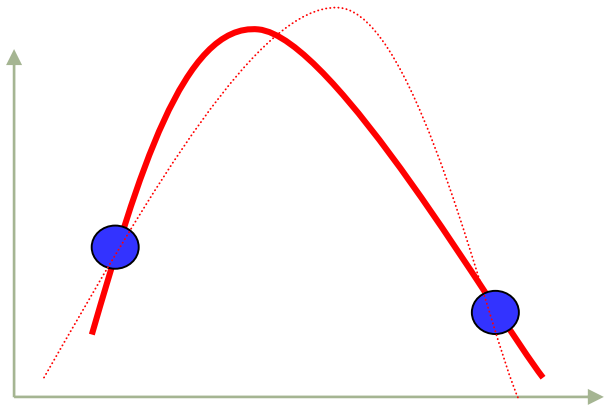

Supervised Learning: Uses

- Prediction of future cases: Use the rule to predict the output for future inputs
- Knowledge extraction: The rule is easy to understand
- Compression: The rule is simpler than the data it explains
- Outlier detection: Exceptions that are not covered by the rule, e.g., fraud

Unsupervised Learning

- Learning “what normally happens”
- No output
- Clustering: Grouping similar instances
- Example applications
  - Customer segmentation in customer relationship management (CRM)
  - Image compression: Color quantization
  - Bioinformatics: Learning motifs

Reinforcement Learning

- Learning a policy: A sequence of outputs
- No supervised output but delayed reward
- Credit assignment problem
- Game playing
- Robot in a maze
- Multiple agents, partial observability, ...

Resources: Datasets - Journals

- UCI Repository: http://www.ics.uci.edu/~mlearn/MLRepository.html
- Statlib: http://lib.stat.cmu.edu/

- Journal of Machine Learning Research www.jmlr.org
- Machine Learning
- Neural Computation
- Neural Networks
- IEEE Trans on Neural Networks and Learning Systems
- IEEE Trans on Pattern Analysis and Machine Intelligence
- Journals on Statistics/Data Mining/Signal Processing /Natural Language Processing/ Bioinformatics/ ...

Resources: Conferences

- International Conference on Machine Learning (ICML)
- European Conference on Machine Learning (ECML)
- Neural Information Processing Systems (NIPS)
- Uncertainty in Artificial Intelligence (UAI)
- Computational Learning Theory (COLT)
- International Conference on Artificial Neural Networks (ICANN)
- International Conference on AI & Statistics (AISTATS)
- International Conference on Pattern Recognition (ICPR)
- …