

Machine Learning

Introduction

Marc Toussaint
U Stuttgart

What is Machine Learning?

- 1) Machine Learning is an approach to understand learning by building learning systems
(“synthetic” approach to a science of learning)

What is Machine Learning?

- 1) Machine Learning is an approach to understand learning by building learning systems
(“synthetic” approach to a science of learning)
- 2) A long list of methods/algorithms for different data analysis problems
 - in sciences
 - in commerce

What is Machine Learning?

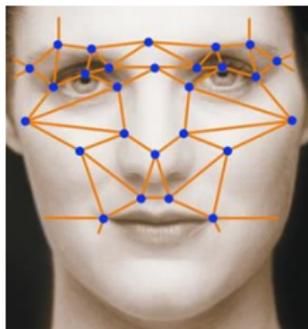
- 1) Machine Learning is an approach to understand learning by building learning systems
(“synthetic” approach to a science of learning)
- 2) A long list of methods/algorithms for different data analysis problems
 - in sciences
 - in commerce
- 3) A framework to develop your own learning algorithms/methods

What is Machine Learning?

- 1) Machine Learning is an approach to understand learning by building learning systems
(“synthetic” approach to a science of learning)
- 2) A long list of methods/algorithms for different data analysis problems
 - in sciences
 - in commerce
- 3) A framework to develop your own learning algorithms/methods
- 4) Machine Learning = information theory/statistics + computer science

Examples for ML applications...

Face recognition



keypoints



eigenfaces

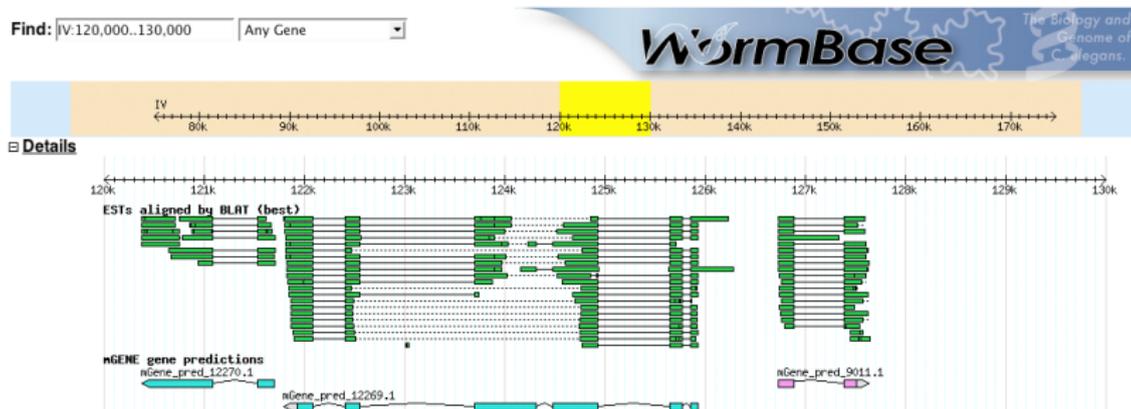
(e.g., Viola & Jones)

Hand-written digit recognition (US postal data)



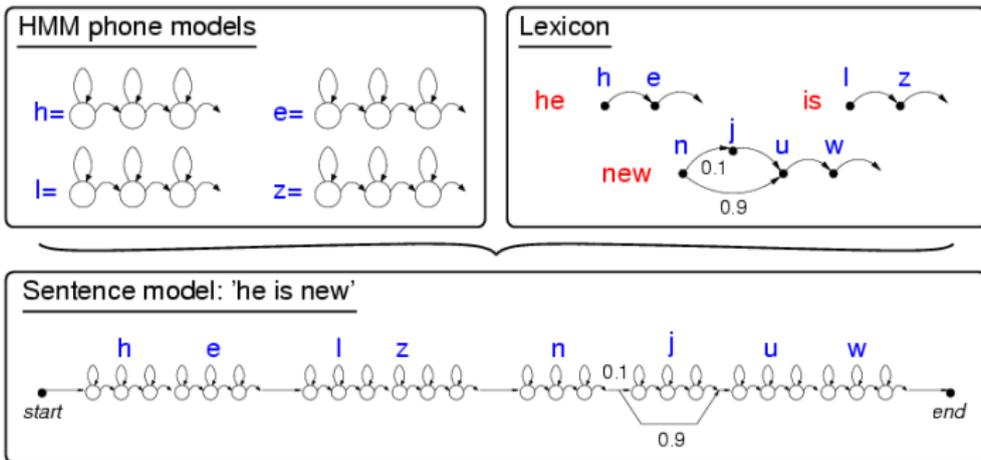
(e.g., Yann LeCun)

Gene annotation



(Gunnar Rättsch, Tübingen, mGene Project)

Speech recognition



(This is the basis of all commercial products)

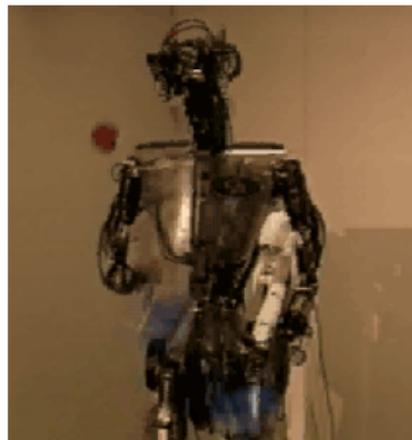
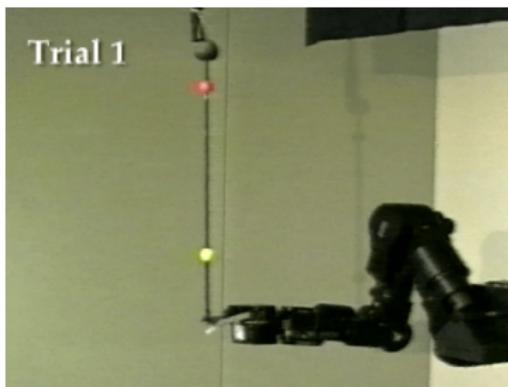
Spam filters

	george	you	your	hp	free	hpl	!	our	re	edu	remove
spam	0.00	2.26	1.38	0.02	0.52	0.01	0.51	0.51	0.13	0.01	0.28
email	1.27	1.27	0.44	0.90	0.07	0.43	0.11	0.18	0.42	0.29	0.01

- More examples:
 - Google's (and many others') analysis of user preferences
 - Medical diagnosis
- **Machine Learning became an important technology in science as well as commerce**

Examples of ML *for behavior...*

Learning motor skills



(around 2000, by Schaal, Atkeson, Vijayakumar)

Learning to walk

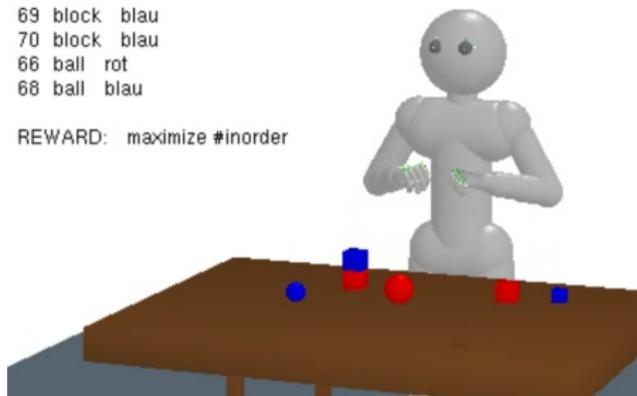


(Rus Tedrake et al.)

Learning effects of actions

OBJECTS:
60 table braun
65 block rot
67 block rot
69 block blau
70 block blau
66 ball rot
68 ball blau

REWARD: maximize #inorder



(Tobias Lang & M Toussaint)

Types of ML

Types of ML

- *Supervised* learning: learn from “labelled” data $\{(x_i, y_i)\}_{i=1}^N$
Unsupervised learning: learn from “unlabelled” data $\{x_i\}_{i=0}^N$ only
Semi-supervised learning: many unlabelled data, few labelled data

Types of ML

- *Supervised* learning: learn from “labelled” data $\{(x_i, y_i)\}_{i=1}^N$
Unsupervised learning: learn from “unlabelled” data $\{x_i\}_{i=0}^N$ only
Semi-supervised learning: many unlabelled data, few labelled data
- *Reinforcement* learning: learn from data $\{(s_t, a_t, r_t, s_{t+1})\}$
 - learn a predictive model $(s, a) \mapsto s'$
 - learn to predict reward $(s, a) \mapsto r$
 - learn a behavior $s \mapsto a$ that maximizes reward

Organization of this lecture

- **Part 1:** *The Basis*
 - Basic regression & classification

- **Part 2:** *The Breadth of ML ideas*
 - PCA, PLS
 - Local & lazy learning
 - Combining weak learners: boosting, decision trees & stumps
 - Other loss functions & Sparse approximations: SVMs
 - Deep learning

- **Part 3:** *In Depth Topics*
 - Bayesian Learning, Bayesian Ridge/Logistic Regression
Gaussian Processes GP classification
 - Active Learning
Recommender Systems

- Missing:
 - Neural Networks
 - Graphical Models & structure learning
 - un-, semi-supervised learning

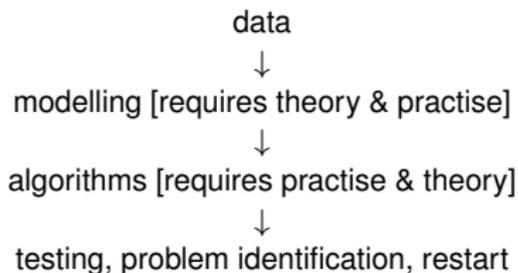
- Is this a theoretical or practical course?

Neither alone.

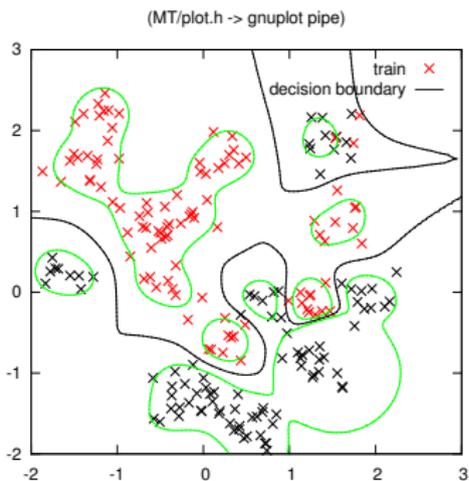
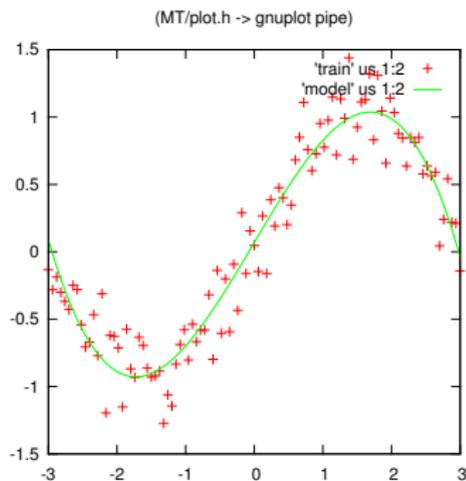
- Is this a theoretical or practical course?

Neither alone.

- The goal is to teach how to design good learning algorithms



Basic regression & classification



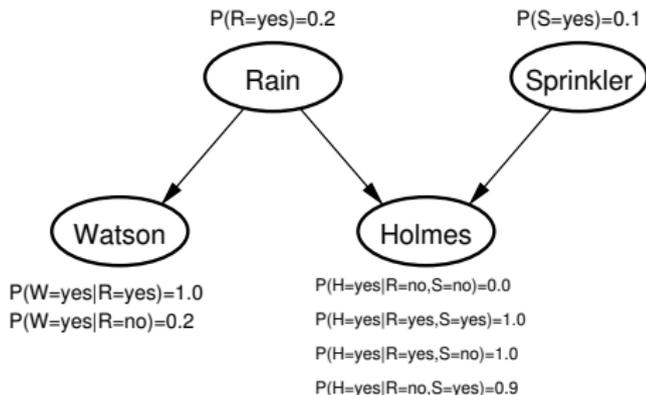
- Regression: map input x to continuous value $y \in \mathbb{R}$
Classification: map input x to one of M classes $y \in \{1, 2, \dots, M\}$

Basic regression & classification

- A must-know!
- High practical relevance for applications
- Focus on linear methods on non-linear features, regularization, cross-validation
“linear|polynomial|Kernel Ridge|Lasso Regression|Classification”
- Relations to SVM, GPs, feature selection

Bayesian Modelling

- Mr. Holmes lives in Los Angeles. One morning when Holmes leaves his house, he realizes that his grass is wet. Is it due to rain, or has he forgotten to turn off his sprinkler?
- Holmes checks Watsons grass, and finds it is also wet. What does that imply on rain vs. sprinkler?



$$\iff P(H, W, S, R) = P(H|S, R) P(W|R) P(S) P(R)$$

Bayesian modelling

- Fundamental view on information processing and learning
- Provides general tools for formulating structured probabilistic models (e.g., latent variables, mixtures, hierarchical, deep models)
→ a framework for formulating novel learning algorithms
- Bayesian view on linear models + regularization
 - regularization \leftrightarrow prior, “error” \leftrightarrow likelihood

Reinforcement Learning (another course..)

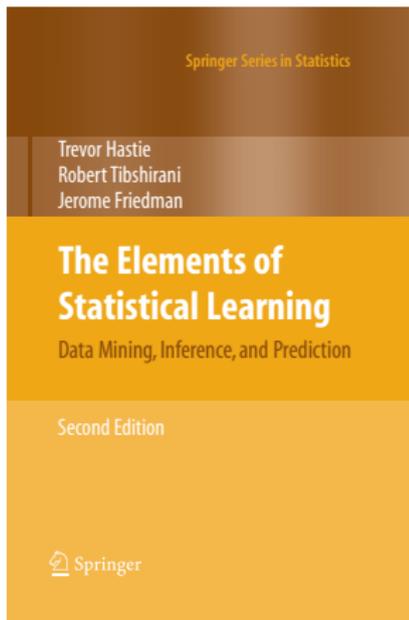
[PacMan]

I Szita, A Lorincz: *Learning to Play Using Low-Complexity Rule-Based Policies: Illustrations through Ms. Pac-Man*. JAIR 2007.

Reinforcement Learning (another course..)

- Behavior!, learning to act
- Basic RL methods (Temporal Difference, Q-learning, traces)
- Regression in RL, Bayesian methods in RL
- Applications

Books



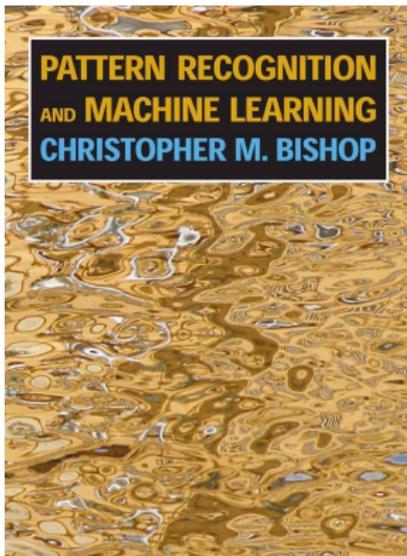
Trevor Hastie, Robert Tibshirani and Jerome Friedman: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* Springer, Second Edition, 2009.

<http://www-stat.stanford.edu/~tibs/ElemStatLearn/>

(recommended: read introductory chapter)

(this course will not go to the full depth in math of Hastie et al.)

Books



Bishop, C. M.: *Pattern Recognition and Machine Learning*.

Springer, 2006

<http://research.microsoft.com/en-us/um/people/cmbishop/prml/>
(some chapters are fully online)

Organisation

- Vorlesungs-Webpage:

<http://ipvs.informatik.uni-stuttgart.de/mlr/marc/teaching/13-MachineLearning/>

- Slides, Übungen & Software (C++)
 - Links zu Büchern und anderen Ressourcen
-
- Sekretariat/Organisatorische Fragen:
Carola Stahl, Carola.Stahl@ipvs.uni-stuttgart.de, Raum 2.217
-
- 2 geplante Übungen: Dienstag 14:00-15:30 & 15:45-17:15, 0.453
-
- Regelung zu Übungen:
 - Bearbeitung der Übungen ist wichtig!
 - Zu Beginn jeder Übung in Liste eintragen:
 - Teilnahme
 - Welche Aufgaben wurden bearbeitet
 - Zufällige Auswahl zur Präsentation der Lösung
 - 50% bearbeitete Aufgaben notwendig für *aktive Teilnahme*