Machine Learning for Robotics
Intelligent Systems Series

Georg Martius

MPI for Intelligent Systems, Tübingen, Germany

April 19, 2017

Organizational structure of the lecture

- Teaching language is English, although you can ask in German
- Mondays 12 c.t.–14:00 Lectures
- Thursdays 12 c.t.–14:00 Recitations
- Exercises:
  - exercise sheets have to be returned in the following week
  - Need 50% passed sheets to be eligible for passing the course
  - Later in the course we will have projects
  - final exam will most likely be a presentation of the final project
- Lecture notes: mostly black board, but there will be background material to read
- Webpage: `georg.playfulmachines.com/course-machine-learning-for-robotics`
- Next week Monday (24th) is canceled (moved to today)

Machine Learning Overview

Machine learning is not voodoo,
it is about automatically finding a function that best solves a given task.

Three different classes of tasks:

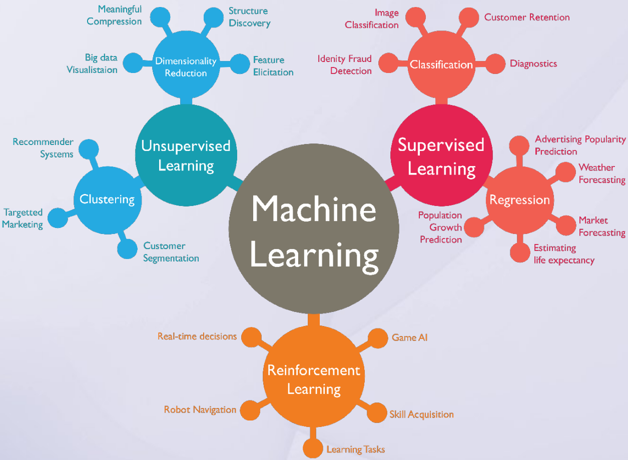

Machine Learning Overview

Supervised Learning

given: $\{x, y\}_i \sim \mathcal{D}$ with data point $x \in \mathbb{R}^n$ and label $y \in \mathcal{Y}$ and $\mathcal{D}$ the data distribution.
What to find function $h(\cdot)$ such that

$$h(x) = y \quad \forall (x, y) \sim D$$

To measure quality of h and to be able to optimize something: Define loss function

$$J(h) = \mathbb{E}_{\mathcal{D}}[\text{dist}(y, h(x))]$$

(distance between true label y and predicted label $f(x)$)

Task: find function that minimized loss: $h^* = \arg\min_h J(h)$

Math can be so easy ;-)

We will see why this is not so easy in practice.

Supervised Learning – Examples

Classification: $\mathcal{Y}$ is discrete

Examples:

Recognize handwritten digits:

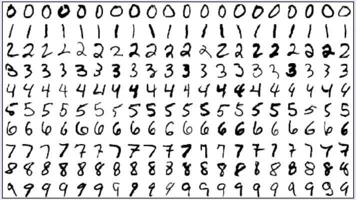

(MNIST)

Classify pathology images:

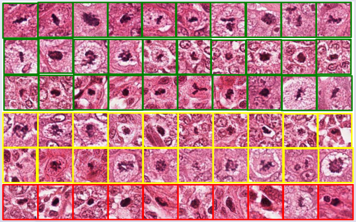

(Mitosis in breast cancer)

Regression: $\mathcal{Y}$ is continuous

Examples:

Predicting Ozon levels

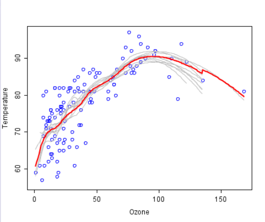

Predicting torques

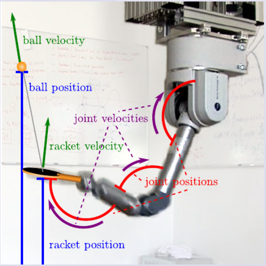

Machine Learning Overview

Unsupervised Learning

given: $\{x\}_i$ with $x \in \mathbb{R}^n$
What to find function $f(\cdot)$ such that $f(x) = y$ where y low dimensional, e.g. a cluster number

- Much less clear what is the objective.
- Many algorithms but no unifying theory.

Unsupervised Learning – Examples

Clustering: discrete y

Examples:

Genome comparison:

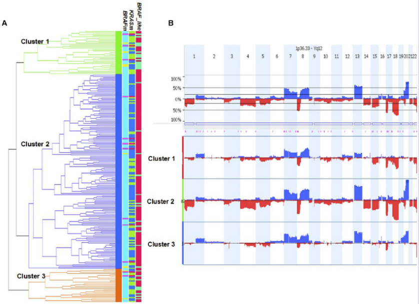

(by Tao Xie)

Dim. reduction: continuous y

Examples:

Finding descriptors for face expressions

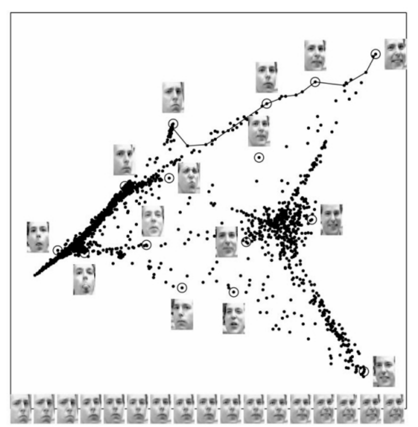

(by Sam T Rowels)

Both cases are expecially useful for high-dimensional data

Machine Learning Overview

Reinforcement Learning

given:

- system to interact with: $s_{t+1} = S(a_t, s_t)$ where s_t is the state and a_t is the action.
- reward/utility function: $r_t = U(a_t, s_t)$

What to find function $f(\cdot)$ (policy) such that $a = f(s)$ and $\mathbb{E}[r]$ is maximized.

In general: stochastic systems formulated as Markov Decision Processes.

- Need to simultaneously learn f and potentially models of S and U.
- Reward can be sparse (e.g. only at the end of an long action sequence)

Reiforcement Learning – Examples

Robot Control

(by MPI-IS)

Deepmind AlphaGo

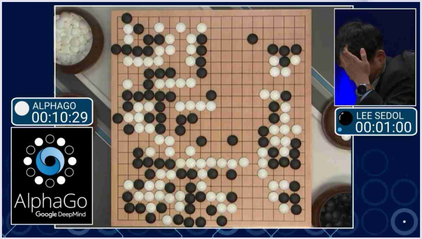

(go-baduk-weiqi.de)

Improve performance by learning from experience
and exploring new strategies.

Rough plan of the course

- Supervised learning
  - linear regression, regularization, model selection, . . .
  - neural networks
- Unsupervised learning
  - Clustering: k-means, spectral, DBSCAN?, . . .
  - Dimensionality reduction: PCA, ICA, LLE, ISOMAP?, Autoencoder, sparse coding and learning representations
- Reinforcement Learning
  - Markov Decision Processes (MDPs) and background
  - Bellman equations and TD learning, Q-Learning, . . .
  - Continuous Spaces:
    - Actor-Critic
    - Reinforcement Learning with parametrized policies
    - Episodic RL as parametrized optimization problem
  - Bayesian optimization for RL?
- if there is time: Artificial Curiosity, . . .