



Advanced Topics in Machine Learning

A. LAZARIC (*INRIA-Lille*)

DEI, Politecnico di Milano

SequeL – INRIA Lille

Outline

Elements of Statistical Learning Theory

Online Learning (a.k.a. Sequence Prediction)

Multi-armed Bandit Problem

Administrative Stuff

Outline


Elements of Statistical Learning Theory

Online Learning (a.k.a. Sequence Prediction)

Multi-armed Bandit Problem

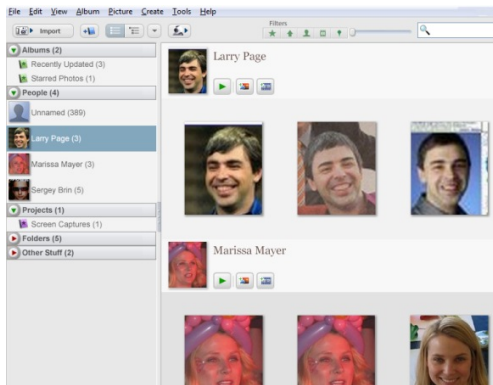
Administrative Stuff

Machine Learning



Handwriting recognition

Machine Learning



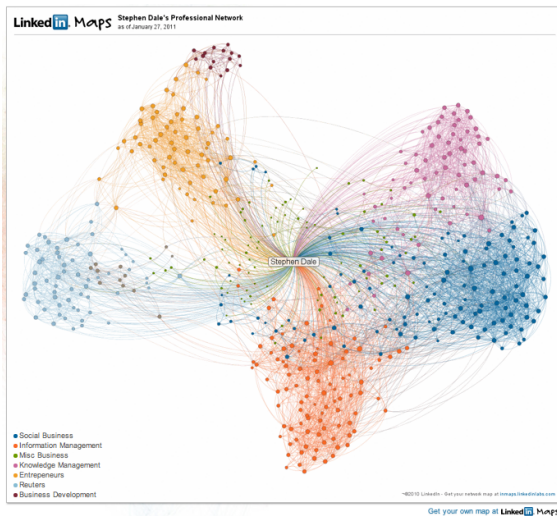
Face detection/recognition

Machine Learning




Page ranking

Machine Learning



Machine Learning



Sentiment classification

Machine Learning



Recommendation systems

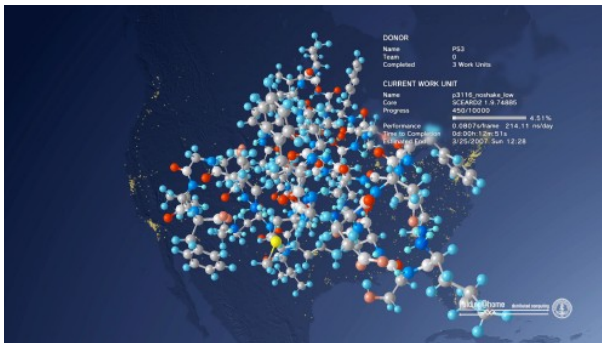
Machine Learning

The image shows the Netflix logo, which consists of the word "NETFLIX" in a bold, white, sans-serif font with a 3D effect, set against a solid red background. The letters are slightly shadowed to give them a three-dimensional appearance.

NETFLIX

Recommendation systems

Machine Learning




Bioinformatics

Machine Learning



Robotics

Machine Learning




Robotics

Machine Learning




Natural Language Processing

Machine Learning



Natural Language Processing

Machine Learning



Entertainment

Elements of Statistical Learning Theory

What: Theoretical analysis of the *behavior* of a learning algorithm.

Elements of Statistical Learning Theory

What: Theoretical analysis of the *behavior* of a learning algorithm.

Why: Learning algorithms work on *noisy* data (i.e., noisy data) and thus their behavior is *non-deterministic* (but *predictable!*).

Elements of Statistical Learning Theory

What: Theoretical analysis of the *behavior* of a learning algorithm.

Why: Learning algorithms work on *noisy* data (i.e., noisy data) and thus their behavior is *non-deterministic* (but *predictable!*).

How: Overview on

- ▶ Concentration of measures inequalities
- ▶ VC–Theory for binary classification
- ▶ L1-Regularized least–squares regression

Outline

Elements of Statistical Learning Theory

Online Learning (a.k.a. Sequence Prediction)

Multi-armed Bandit Problem


Administrative Stuff

Online Prediction



Binary prediction

Online Prediction



Online data compression

Online Prediction




Market prediction

Online Prediction



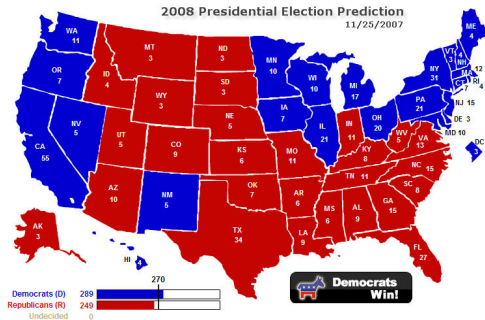
Betting systems

Online Prediction



Weather forecast


Online Prediction



Elections prediction



Online Prediction



Traffic forecast

Online Learning (a.k.a. Sequence Prediction)

What: Prediction of *any sequence* of elements in an *online* fashion (e.g., stock market, binary sequences, weather, ...).

Online Learning (a.k.a. Sequence Prediction)

What: Prediction of *any sequence* of elements in an *online* fashion (e.g., stock market, binary sequences, weather, ...).

Why: Most of the time data are not available in batches but are provided in a *stream* and their behavior might *change over time*.

Online Learning (a.k.a. Sequence Prediction)

What: Prediction of *any sequence* of elements in an *online* fashion (e.g., stock market, binary sequences, weather, ...).

Why: Most of the time data are not available in batches but are provided in a *stream* and their behavior might *change over time*.

How: Overview on

- ▶ Learning with expert advice for continuous prediction
- ▶ Learning with expert advice for discrete prediction
- ▶ Efficient forecasters

Outline

Elements of Statistical Learning Theory

Online Learning (a.k.a. Sequence Prediction)

Multi-armed Bandit Problem


Administrative Stuff

Online Prediction



Online web advertising

Online Prediction




Adaptive routing

Online Prediction



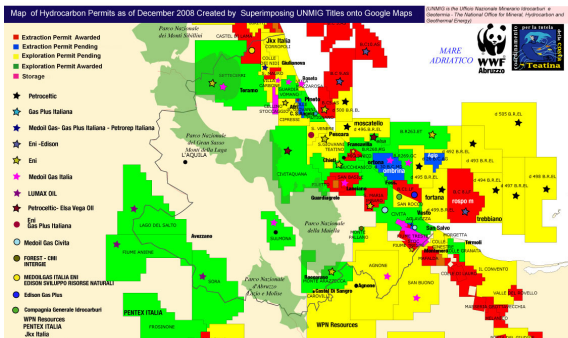
Clinical trials

Online Prediction



Computer games

Online Prediction



Oil/gas mining

Online Prediction



Adaptive marketing

Multi-armed Bandit Problem

What: Study of learning algorithms that are *responsible* for the decisions and data acquisition at the same time.

Multi-armed Bandit Problem

What: Study of learning algorithms that are *responsible* for the decisions and data acquisition at the same time.

Why: The observed data often depend on the learning system itself. Thus, it should trade-off between collection of *good* and *informative* data.

Multi-armed Bandit Problem

What: Study of learning algorithms that are *responsible* for the decisions and data acquisition at the same time.

Why: The observed data often depend on the learning system itself. Thus, it should trade-off between collection of *good* and *informative* data.

How: Overview on

- ▶ The exploration–exploitation dilemma
- ▶ Stochastic bandits (UCB)
- ▶ Non–stochastic bandits (Exp3)
- ▶ Extensions to large set of arms
- ▶ Connections to game theory

Outline

Elements of Statistical Learning Theory

Online Learning (a.k.a. Sequence Prediction)

Multi-armed Bandit Problem

Administrative Stuff



Schedule

<i>Topic</i>	<i>Day/Time</i>	<i>Location</i>
Statistical Learning Theory	Apr. 2nd, 9:00-13:00	Sala Seminari
Statistical Learning Theory	Apr. 3rd, 9:00-13:00	Sala Seminari
Online Learning	Apr. 4th, 9:00-13:00	Sala Seminari
Online Learning	Apr. 11th, 9:00-13:00	Sala Seminari
Multi-armed Bandits	Apr. 12th, 9:00-13:00	Sala Seminari
Multi-armed Bandits	Apr. 13th, 9:00-13:00	Sala Seminari

Assignments

- ▶ **Option 1:** Report on 2-3 papers (better if you propose them)
- ▶ **Option 2:** Presentation on 2-3 papers (on the last day)
- ▶ **Option 3:** Short project (if some interesting research questions come out)

Advanced Topics in Machine Learning



Alessandro Lazaric

alessandro.lazaric@inria.fr

sequel.lille.inria.fr