# IASD: Optional Courses

masteriasd.eu

Olivier Cappé, Benjamin Negrevergne, Maxime Florens
Pierre Senellart, Étienne Decenciere

# General information

- You need to choose 6 optional courses.
- You **cannot abandon a course, or register to a new one during the semester.**
- You can follow more optional courses as "auditor only", but your grade will not be part of the final grade for the period
- Some options may not be opened if attendance is too low
  - we will do a second round in case too many options have been closed
- Check for potential schedule conflicts **before communicating your choice**
- Communicate your choices using the survey available on the website.
  **Deadline: Tuesday November 24, 2023**

# Available optional courses

- PSL Intensive weeks:
  - Digital Humanities meet Artificial Intelligence
  - Machine learning for physics and engineering
  - 'Green' Artificial Intelligence
  - Machine learning in Genomics
  - NLP for Social Sciences
- Point clouds and 3D modeling — François Goulette*
- Privacy for machine learning — Olivier Cappé, Muni Pydi
- Knowledge graphs, description logics, reasoning on data — Michaël Thomazo, Camille Bourgaux
- No SQL databases — Paul Boniol
- Deep reinforcement learning and application — Eric Benhamou
- Computational social choice — Jérôme Lang, Dominik Peters
- Incremental learning, game theory, and applications — Guillaume Vigeral
- Advanced machine learning — Yann Chevaleyre
- Monte-Carlo search and games — Tristan Cazenave
- Planning, Search and Constraint Solving — Tristan Cazenave
- Graph Analytics — Daniela Grigori*
- Machine learning on Big Data — Dario Colazzo*

# Schedule conflicts

- Graph analytics – 3D point cloud
- NoSQL Databases – Planning, search and constraint solving

# Nuages de points et modélisation 3D

# *3D Point Cloud and Modelling*

François GOULETTE
Jean-Emmanuel DESCHAUD
Tamy BOUBEKEUR

Contact : francois.goulette@ensta-paris.fr

Site Web du cours : https://www.caor.minesparis.psl.eu/presentation/cours-npm3d/

Qu’est-ce que c’est ?!?


MS Kinect

A quoi ça sert ?!?

Comment ça marche ?!?

Image couleur (RGB)

Image de profondeur (D)

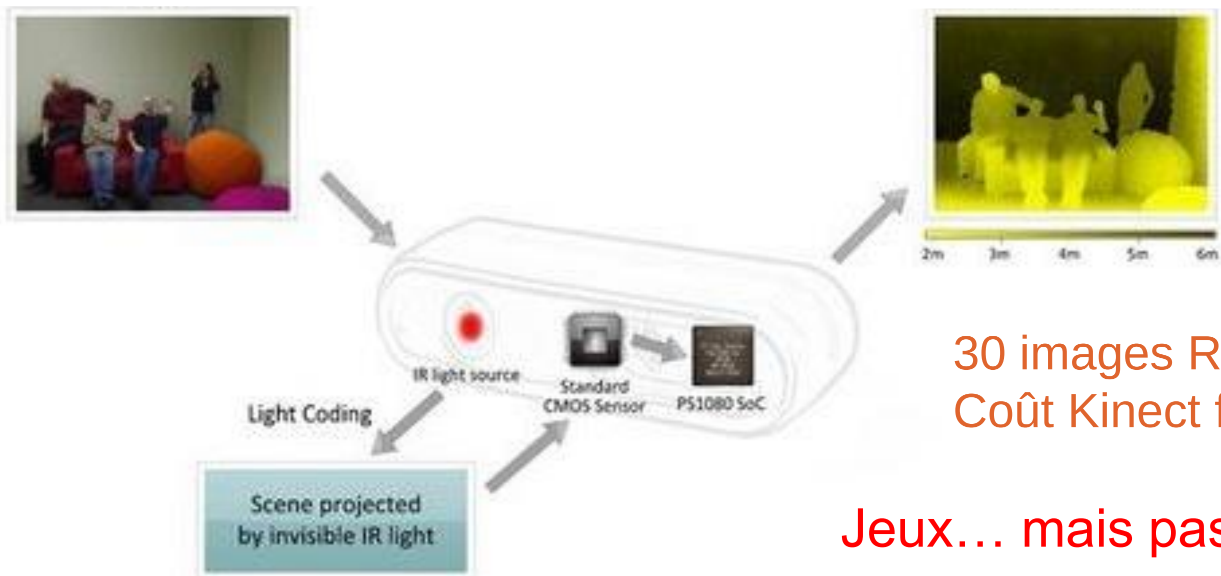


30 images RGB-D par seconde !
Coût Kinect faible ~100 €

Jeux… mais pas seulement…!

## Qu’est-ce que c’est ?!?


Scanner laser
(Faro Focus 3D)

## A quoi ça sert ?!?

Relevés 3D
à usages professionnels

Jusqu’à x100 kpts/s !
Coût scanner faible ~30 k€

## Comment ça marche ?!?

Images de profondeur
Plusieurs stations (lieux)

→Nuages de points

**Démocratisation de la 3D !**

**Recherche d’actualité et renouvelée : véhicules autonomes etc. !**

Video L3D2 Montbéliard 2013

# Et après les nuages de points ?

Rendu par point

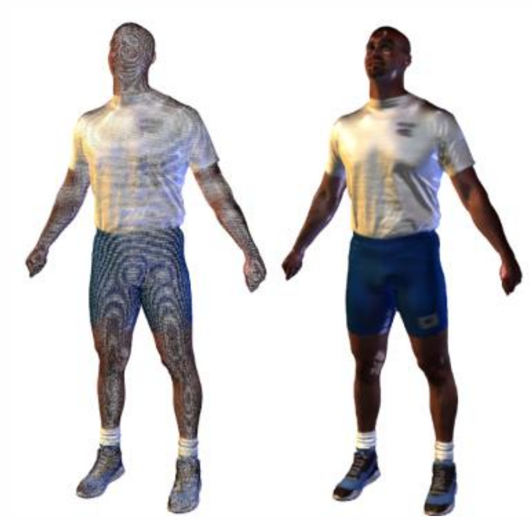

Reconstruction de surface

Modélisation

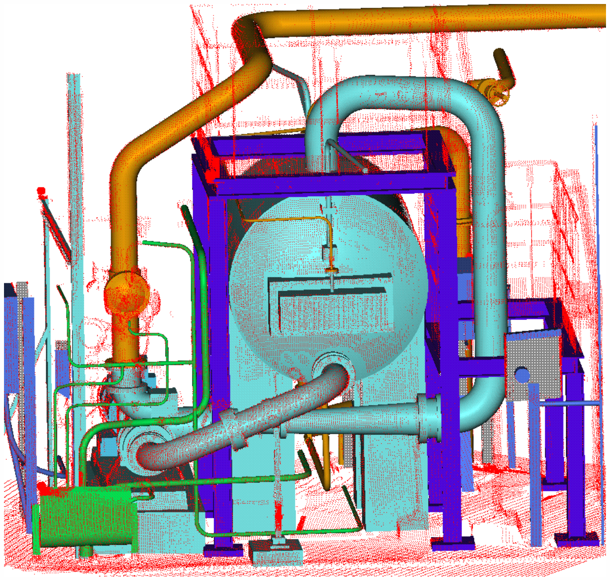

Traitements de données… sémantisation…
deep learning…de nouveaux challenges !
… calculs rapides, robustesse, interactivité…etc.

Video Niessner 2013

# Déroulement

1/ **Perception 3D ; capteurs et étalonnage** **(FG)**
2/ **Recalage et consolidation** **(FG)**
3/ **Description locale des courbes et surfaces** **(FG)**
4/ **Rendu de nuages de points et maillages** **(TB)**
5/ **Reconstruction de courbes et surfaces** **(JED)**
6/ **Modélisation et segmentation** **(FG)**
7/ **Apprentissage profond et nuage de points 3D** **(JED)**

- **Séminaire de recherche** (chercheurs, doctorants)

# Organisation

- **Jeudis après-midi, 13h45-18h**
  - Cours + TP informatique
  - Venir avec ordinateur portable
    - Logiciels : Python, CloudCompare (installés à l'avance)
- **Lieu**
  - Paris Santé Campus OU Mines Paris (A CONFIRMER)
- **Language**
  - Courses in French, educational documents in English.
  - Practical courses : documents in English, accompanied in French and English
- **Evaluation**
  - Comptes-rendus de TP (1/3) et projets sur articles (2/3)

**INSCRIPTION** (pour être tenu informé) :
Site Web du cours https://www.caor.minesparis.psl.eu/presentation/cours-npm3d/

# Questions ?...

Rue Soufflot, Paris

Points 3D couleur : Données LARA-3D (MINES ParisTech)

Gris : MNS (IGN)

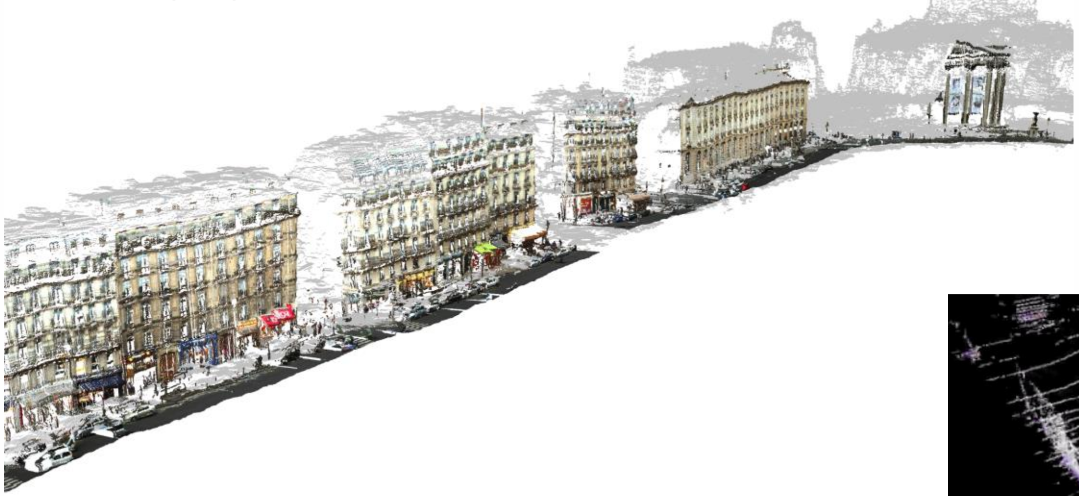

3D - SLAM

Localisation de véhicule autonome

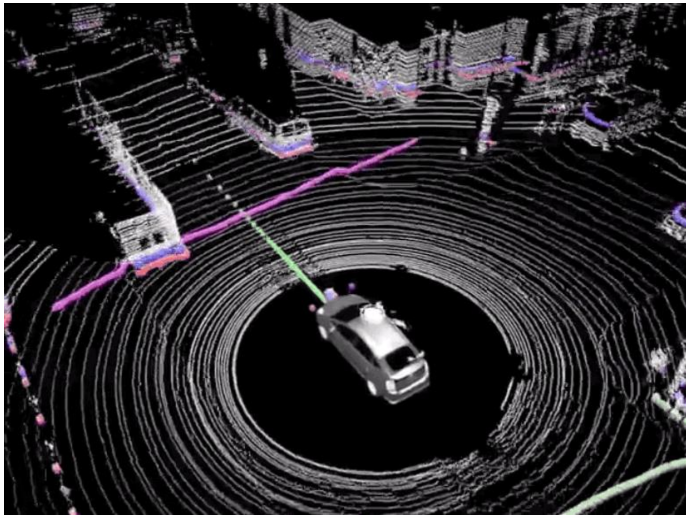

# IASD 2022–2023 : Privacy for Machine Learning

Classes on Tuesday afternoon (PSC): $8 \times 3$h + defense
Contact: Olivier Cappé (`olivier.cappe@cnrs.fr`)

`https://github.com/ftramer/LM_Memorization`

[...] Reuters 5/11 Tim Hortons releases 'Java Daddy' tv ad where actor plays non-binary character called 'Java' and challenges Michael Jackson to an Apple Watch video choosing between the two A man was killed early Monday in a drive-by shooting on his front porch in the Englewood neighborhood on the South Side, police said. The shooting happened about 1:30 a.m. in the 7300 block of South Kedzie, Officer Ana Pacheco, a Chicago police spokeswoman, said in a news release. The victim, who had his back to the gunman when the shooting occurred, was struck in the chest by gunfire, [...]

`https://thispersondoesnotexist.com`


really?

This courses covers the basics of Differential Privacy (DP), a framework that has become, in the last ten years, a de facto standard for enforcing user privacy in data processing pipelines. DP methods seek to reach a proper trade-off between protecting the characteristics of individuals and guaranteeing that the outcomes of the data analysis stays meaningful.

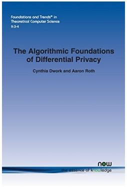


- The first part of the course is devoted the basic notion of epsilon-DP and understanding the trade-off between privacy and accuracy, both from the empirical and statistical points of view.
- The second half of the course will cover more advanced aspects, including the different variants of DP and the their use to allow for privacy-preserving training of large and/or distributed machine learning models.

Keywords: Randomized response, differential privacy (epsilon-DP, Rényi DP, ...), Laplace mechanism, DP-SGD, federated learning

# In Practice

## Lectures

- Jamal Atif (LAMSADE)
- Olivier Cappé (DI ENS)
- Muni Sreenivas Pydi (LAMSADE)

## Grades

Practicals (please bring you laptop in class) / homeworks (Python/Colab) + project (in goups) on research papers with final project defense

Prerequisite: Basic probability and statistics, Pyhton + first semester courses on machine learning and optimization

# Knowledge Graphs, Description Logics and Reasoning on Data

C. Bourgaux, M. Thomazo

# Context

Data is at the core of many applications, but:

- data is *heterogeneous* in several ways:
  - models: relational, textual, ...
  - vocabulary: different languages, attribute names,...
- *semantics* of the data is important, but often implicit;
- final users may not be IT experts.

**Challenge**

How to allow a user to efficiently access the relevant data?

# Knowledge Representation

## General Goal

"Develop formalisms for providing high-level descriptions of the world that can be effectively used to build intelligent applications" (Nardi and Brachman, 2003)

# Knowledge Representation

## General Goal

"Develop formalisms for providing high-level descriptions of the world that can be effectively used to build intelligent applications" (Nardi and Brachman, 2003)

- formalism: well-defined syntax and formal, unambiguous semantics;

# Knowledge Representation

## General Goal

"Develop formalisms for providing high-level descriptions of the world that can be effectively used to build intelligent applications" (Nardi and Brachman, 2003)

- formalism: well-defined syntax and formal, unambiguous semantics;
- high-level description: only relevant aspects are represented;

# Knowledge Representation

## General Goal

"Develop formalisms for providing high-level descriptions of the world that can be effectively used to build intelligent applications" (Nardi and Brachman, 2003)

- formalism: well-defined syntax and formal, unambiguous semantics;
- high-level description: only relevant aspects are represented;
- intelligent applications: inferring the implicit from the explicit;

# Knowledge Representation

## General Goal

"Develop formalisms for providing high-level descriptions of the world that can be effectively used to build intelligent applications" (Nardi and Brachman, 2003)

- formalism: well-defined syntax and formal, unambiguous semantics;
- high-level description: only relevant aspects are represented;
- intelligent applications: inferring the implicit from the explicit;
- effectively used: practical reasoning tools and efficient implementations.

# Ontology-Based Data Access

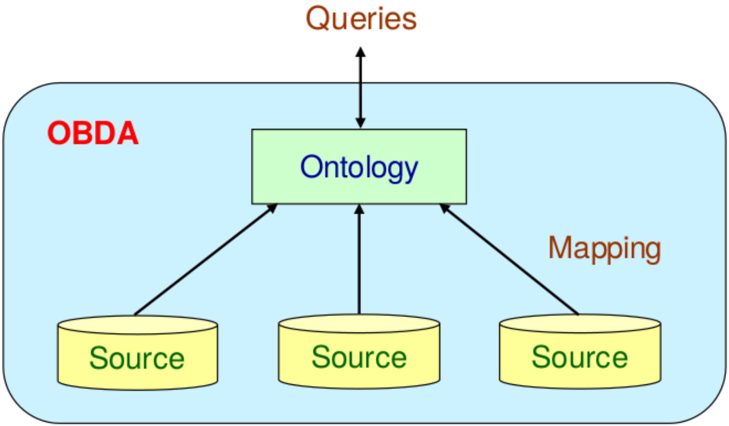

# Practical Matters

Curriculum:

- Intoduction to Knowledge Graphs and Logic ($2 \times 3$ hours)
- Reasoning with Description Logics ($2 \times 3$ hours)
- Using Ontologies to Query Data ($2 \times 3$ hours)
- Opening Topics ($2 \times 3$ hours)

Type of courses: Lectures, Hands-On sessions, Tutorials

Evaluation: written exam.

# NoSQL

Paul Boniol

Contact: boniol.paul@gmail.com

# Why NoSQL?

More and More data…

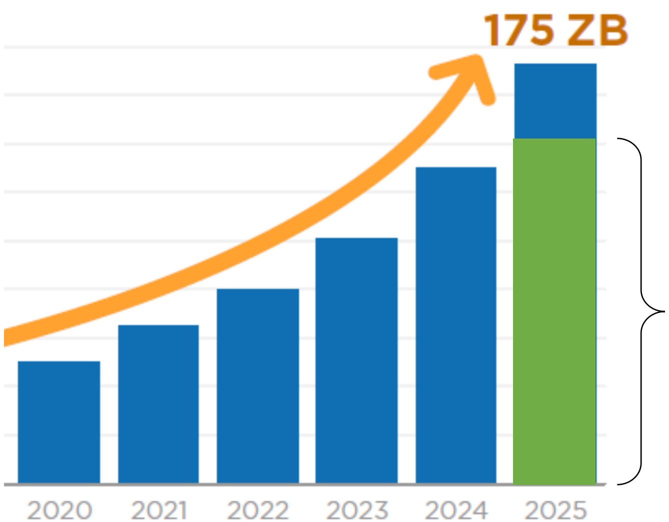


80% are complex multi-dimensional data

(e.g., time series, text, audio, images, videos, logs…)

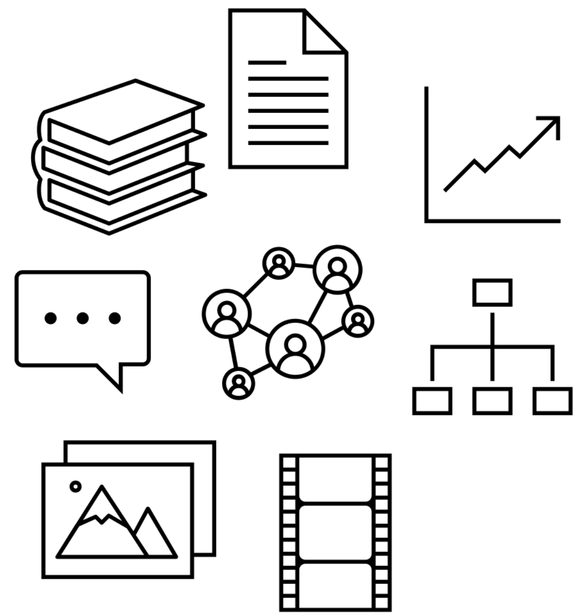

# What is NoSQL?

How to represent and store data outside traditional formats?

How to search efficiently?

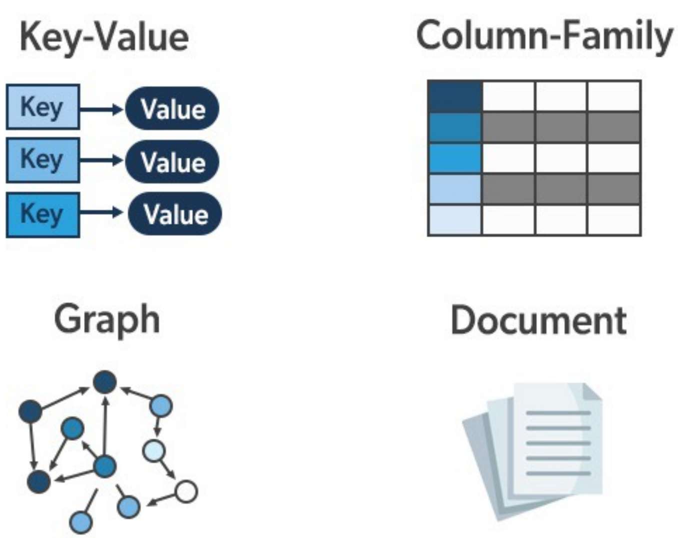


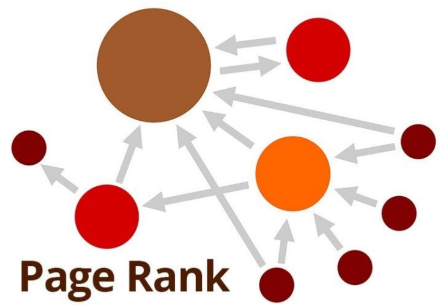


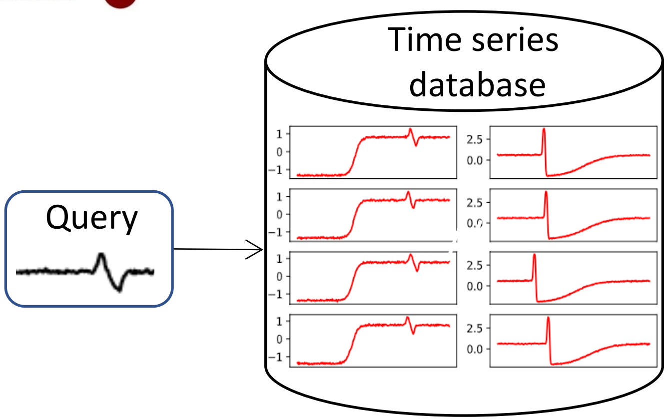

# Curriculum and provisional schedule

8 sessions (of 3 hours) + project defenses

## General Introduction

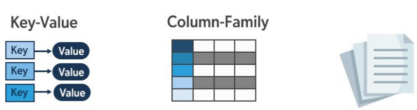


We will explore the following topics:

- SQL vs NoSQL
- ACID and BASE
- CAP theorem
- Types of NoSQL databases
    - Key-value store
    - Document-store
    - Column-store

## Graph Database

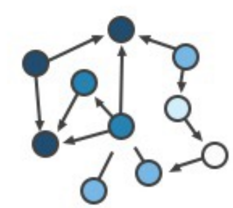

We will explore the following topics:

- Basic graph theory
- Graph structure
- Graph data modeling
- Labeled-property graph
- Real system: Neo4j

## Search and indexing

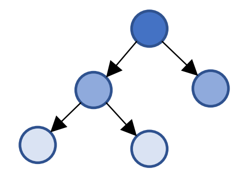

We will explore the following topics:

- Information retrieval
- Indexing
- PageRank

### Special focus on time series

- Time series similarity
- Time series symbolization
- Time series indexing
- Deep learning for indexing

# Evaluation

## One homework

- 40% of the total grade
- Topic:
  - **Select one research paper** (among a pre-selected list of papers).
  - **Summarize** it
  - **Explain** the method and the results in your own words
  - **Comment** on its **strengths** and **limitations**

## One Project

- 60% of the total grade
- Topic:
  - **Select one research paper** (among a pre-selected list of papers).
  - **replicate** it
    - re-implement the method
    - Reproduce the experimental results
    - *[optional] Evaluate it on a new case*
  - **Present** your results in a **defense**

# Deep Reinforcement Learning

Master IASD: Artificial Intelligence & Systems, Data

E.Benhamou, D. Saltiel

2021-2022

# What you will learn in this class?

- ✓ Intro and Course Overview
- ✓ Supervised Learning behaviors
- ✓ Intro to Reinforcement Learning
- ✓ Policy Gradients
- ✓ Actor-Critic Algorithms (A2C, A3C and Soft AC)
- ✓ Value Function Methods
- ✓ Deep RL with Q-functions
- ✓ Advanced Policy Gradient (DDPG, Twin Delayed DDPG)
- ✓ Trust Region & Proximal Policy Optimization (TRPO, PPO)
- ✓ Optimal Control and Planning
- ✓ Model-Based Reinforcement Learning
- ✓ Model-Based Policy Learning
- ✓ Exploration and Stochastic Bandit in RL
- ✓ Exploration with Curiosity and Imagination
- ✓ Offline RL and Generalization issues
- ✓ Offline RL and Policy constraints

# Why DRL?

- ✓ Is a very promising type of learning as it does not need to know the solution
- ✓ Only needs the rules and good rewards
- ✓ Combines best aspects of deep learning and reinforcement learning.
- ✓ Can lead to impressive results in games, robotic, finance

# References


- ✓ Goodfellow, Bengio, Deep Learning
- ✓ Sutton & Barto, Reinforcement Learning: An Introduction
- ✓ Szepesvari, Algorithms for Reinforcement Learning
- ✓ Bertsekas, Dynamic Programming and Optimal Control, Vols I and II
- ✓ Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming
- ✓ Powell, Approximate Dynamic Programming

# Computational Social Choice

- Social choice:
  *designing and analysing methods for collective decision making*
- Topics: Voting rules    Fair division    Matching
- Lecturers:
  - 4 lectures: Dominik Peters (dominik.peters@lamsade.dauphine.fr)
  - 4 lectures: Jérôme Lang (lang@lamsade.dauphine.fr)
- Wednesdays 13:45-17:00
- Intersection of computer science / AI and economics

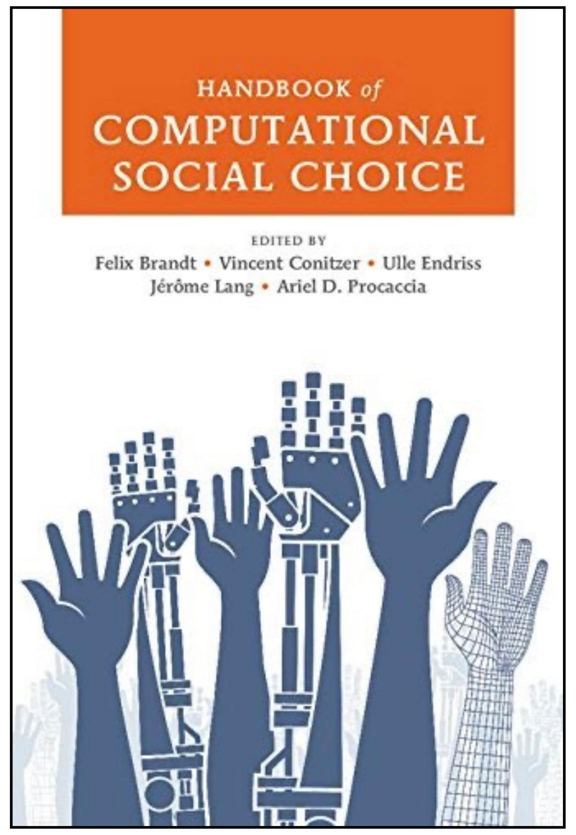

# Plan of the course

## Allocation
(how to decide who gets what)

- Fair cake cutting
  (proportionality and envy-freeness, protocols, query complexity)
- Rent division
  (quasi-linear utilities, maximin solution, linear programming)
- Indivisible goods
  (relaxations of envy-freeness, maximising Nash welfare, NP-hardness, approximations)
- Random assignment
  (fairness via randomness, strategyproofness, impossibility theorem)
- Apportionment
- Stable matching

## Voting and collective decisions
(how to decide what to do)

- Voting rules
  (the good, the bad, and the ugly, and how to tell which is which; axioms, input formats, information, computation)
- Strategic voting
  (famous impossibility theorems of Arrow and Gibbard-Satterthwaite, escape routes)
- Multiwinner voting
  (designing objective functions, algorithms and complexity, proportional representation)
- Public goods & participatory budgeting
  (portioning, public decision making, the core, the method of equal shares)
- Communication issues
- Applications to moral AI

# Computational Social Choice

## Reasons to take the course

- Mathematics with societal applications
- Learning rigorous tools for evaluating decision making procedures
- Learning patterns for designing good methods
- Excellent field to get started doing research
- Interdisciplinary


No prerequisites (the course is self-contained) but
a basic level in discrete maths and algorithmics will help.

Dauphine UNIVERSITÉ PARIS | PSL LAMSADE UMR 7243

# Online Learning in Games


Rida Laraki

CNRS, PSL


IASD, January-April, 2020

# Outline of the course 1/2

- 0) **C1** : Introduction.
- 1) **C2** : Zero-Sum Games (finite case) : minmax, maxmin, value, mixed strategies, von-Neuman minmax theorem and its link with linear programming. Two learning procedures : fictitious play (follow the leader) / and better-reply (Blackwell approachability).
- 2) **C3** : Zero-Sum Game (general case) : Sion minmax theorem (proofs : by separation, by discretization, by learning -fictitious play-). Extensive Form Games (Zermelo's, Gale Stewart, Kuhn's theorems).
- 3) **C4** : Vector payoff games : Blackwell approachability and the equivalence with no-regret and calibration. Application to zero-sum games. Link with online optimisation (Online Gradient Descent, Follow the leader, Online Mirror Decent).

# Outline of the course 2/2

- 4) **C5** : *N*-player games : Rationalizability, Nash equilibrium, potential games, monotone games. Existence and variational characterisation of NE. Best reply dynamics and fictitious play : convergence in potential/acyclic games, non convergence in general games (Shapley triangle). Fictitious play with finite memory.
- 5) **C6** : Smooth fictitious play, link with regret learning (follow the perturbed leader) and convergence to coarse equilibria. Internal regret and prediction : convergence to correlated equilibria.
- 6) **C7** : Repeated Games : cooperation and folk theorems. Bayesian Learning, merging, grain of truth, convergence and impossibilities. Hypothesis Testing and convergence to Nash Equilibria.
- 7) **C8** : Continuous time learning dynamics, link with discrete time, local stability, ESS, variational stability (if we have time !)
- 8) **C9** : Examination : oral presentation by all student of their article (20mn+5mn questions). The final grade is the sum with the oral and the written report grades.

# Apprentissage Automatique Avancé

# Objectif du cours

- **Partie 1: Apprentissage Bayésien:**
  - Dans le paradigme dominant de l'apprentissage artificiel vu en cours est le « max vraissemblance »
  - L'apprentissage Bayésien permet d'aggréger l'ensemble des modèles probables (plutot que de n'en prendre qu'un) et fournit des intervalles de confiance

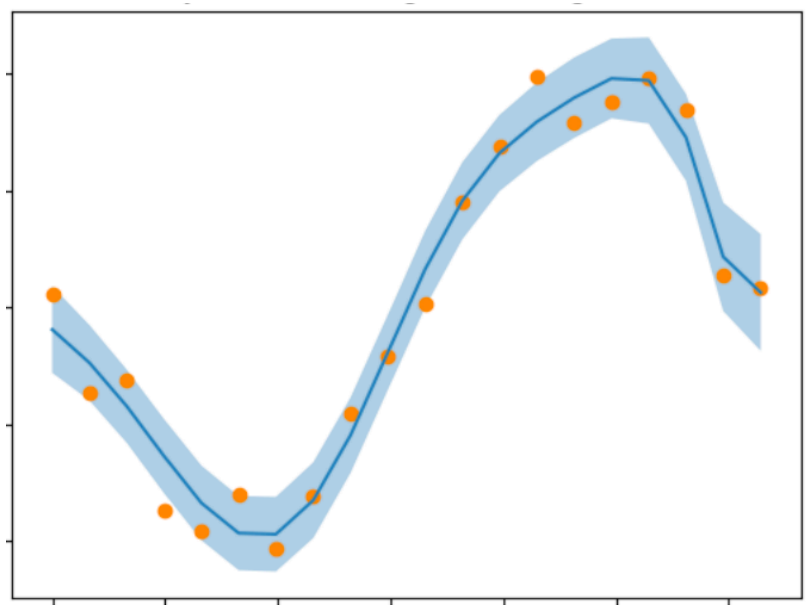

- **Partie 2 : Recommandation et Ranking**
  - Apprentissage de moteur de recherche, apprendre à partir de labels ordinaux…

# Intervenants/Contenu

- **Moez Draif** (VP & chief scientist chez Capgemini, prof à Imperial college). **6h**
  **Methodes bayesiennes en Apprentissage.**
    - Intro : présentation de la structure du cours
    - Approche bayésienne, différences avec la statistique fréquentiste
    - Régression linéaire bayésienne
    - Lois conjugées
    - Qu'est ce que le Topic Modeling ?
    - Algorithme LDA (latent dirichelet allocation)
    - Processus Gaussiens
    - Optimisation bayésienne
    - Exemples d'applications
- **Julian Arbel (CR INRIA). 6h**
    - Méthodes variationnelles
    - deep learning bayésien (variationnel et MCMC)
- **Clément Calauzene (Criteo).**
  **Learning to Rank and Recommender Systems. 6h**
    - Ce cours va explorer les architectures possibles pour construire un système de recommandation. En particulier,
    - Comment découpler la phase de pre-selection (retrieval) de la recommandation elle-même (ranking) pour des raisons de passage à l'échelle ?
    - Comment construire, apprendre et évaluer un modèle ranking ?
    - Comment intégrer la nature implicite du feedback observé (click, rating...) qui n'est pas aussi riche qu'une supervision ?

# Evaluation

- Controle continu:
  Lecture et présentation (en 15 ou 20 minutes) d’un papier de recherche dans une thématique liée au cours

# Monte Carlo Search and Games

Tristan Cazenave

- Monte Carlo Tree Search
- Nested Monte Carlo Search
- Nested Rollout Policy Adaptation

AlphaGo


Lee Sedol

Ke Jie

# Monte Carlo Tree Search

- UCB (Upper Confidence Bounds)
- UCT (Upper Confidence bounds applied to Trees)
- AMAF (All Moves As First)
- RAVE (Rapid Action Value Estimation)
- GRAVE (Generalized RAVE)
- Sequential Halving
- SHUSS (Sequential Halving Using ScoreS)
- PUCT (Prior UCT)

# Nested Monte Carlo Search

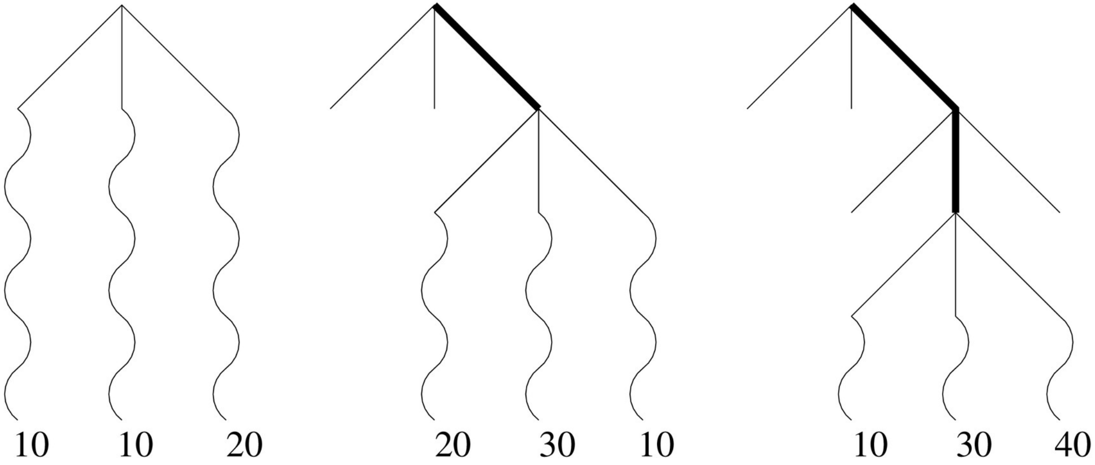


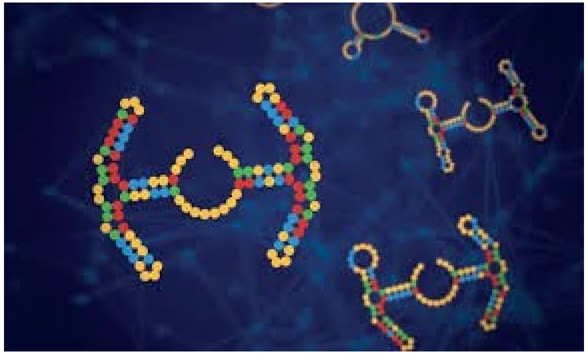

# Nested Monte Carlo Search

- Theoretical Analysis
- Applications
- Discovery of Mathematical Expressions
- Two Player Games

# Nested Rollout Policy Adaptation

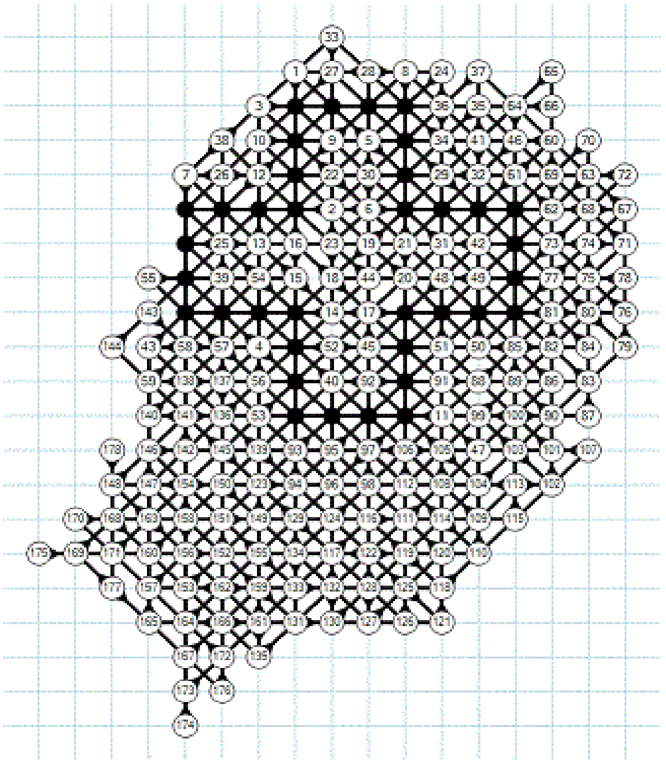

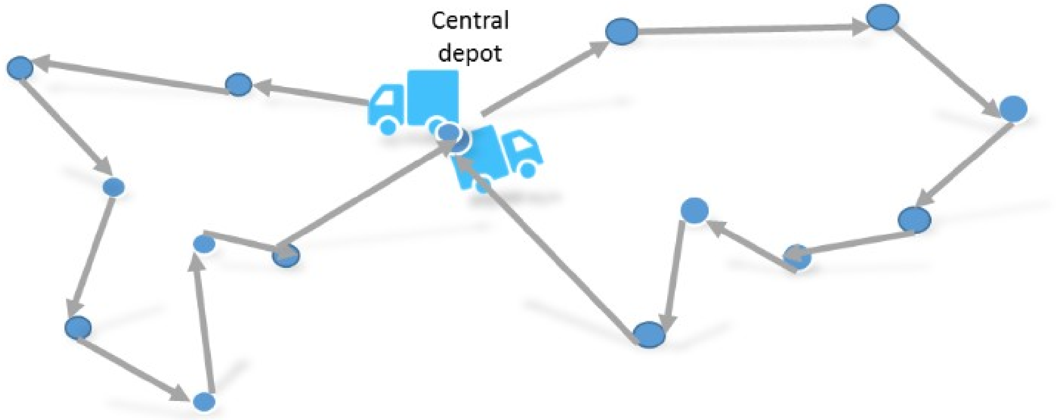

# Nested Rollout Policy Adaptation

- Presentation of the Algorithm
- Applications
- Selective Policies
- Weak Schur Numbers
- Theoretical Analysis
- Generalized NRPA
- Warm Starting
- Bias Weights Learning
- Playout Policy Learning

# Planning, search, and constraint solving

Arnaud Lallouet & Tristan Cazenave

- SAT
- CSP
- A*
- IDA*
- Design of heuristics: Rubik’s Cube
- Retrograde Analysis: 15-Puzzle
- Planning: Multiagent Pathfinding (Amazon Warehouses)
- Partial Moves: Multiple Sequence Alignment (Biology)

# Graph analytics

Daniela Grigori
Paris-Dauphine University, PSL

# Many types of data are graphs

Social networks

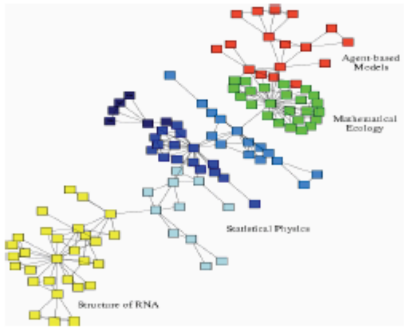


Economic networks

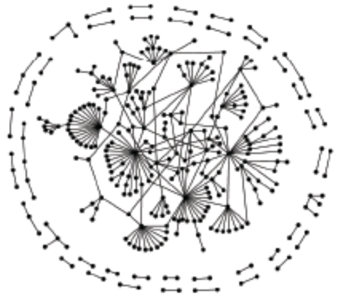

Communication networks

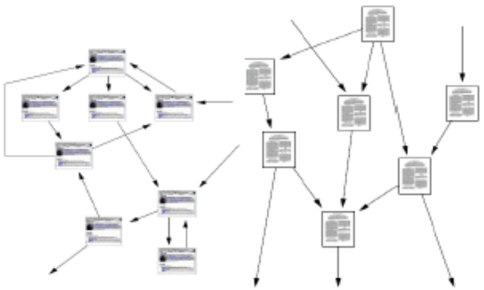

Information networks:
Web & citations

Internet

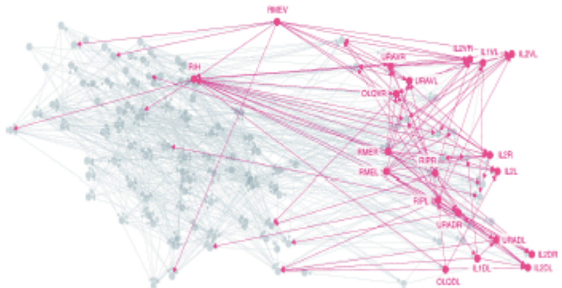

Networks of neurons

# Graph Data: Social Networks


**Facebook social graph**

**4-degrees of separation [Backstrom-Boldi-Rosa-Ugander-Vigna, 2011]**

# Graph Data: Media Networks

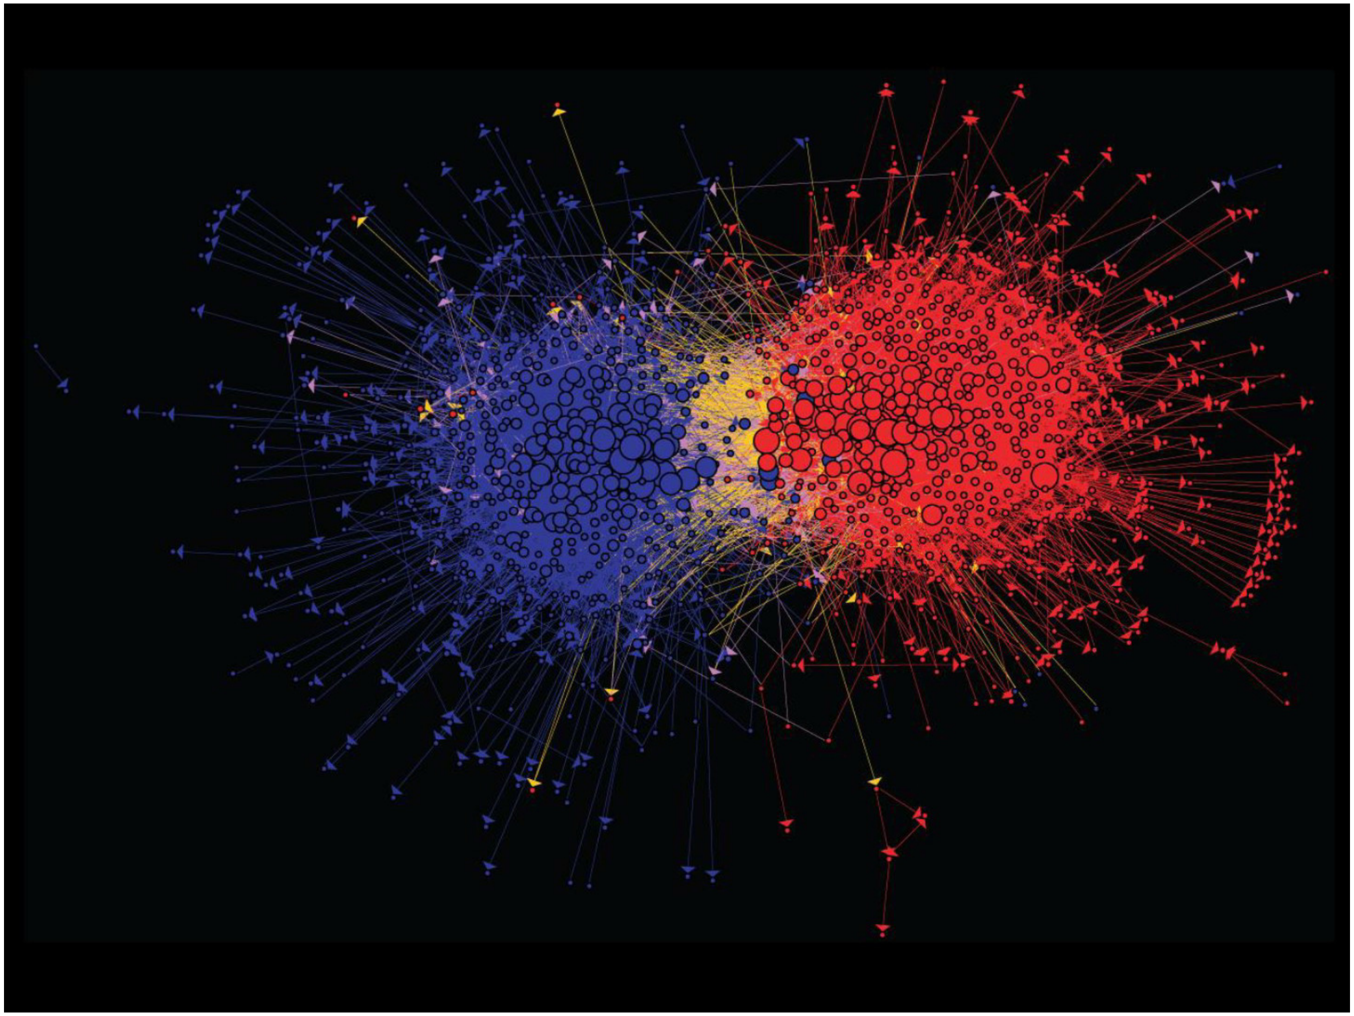

**Connections between political blogs**

Polarization of the network [Adamic-Glance, 2005]

# Graph Data: Information Nets

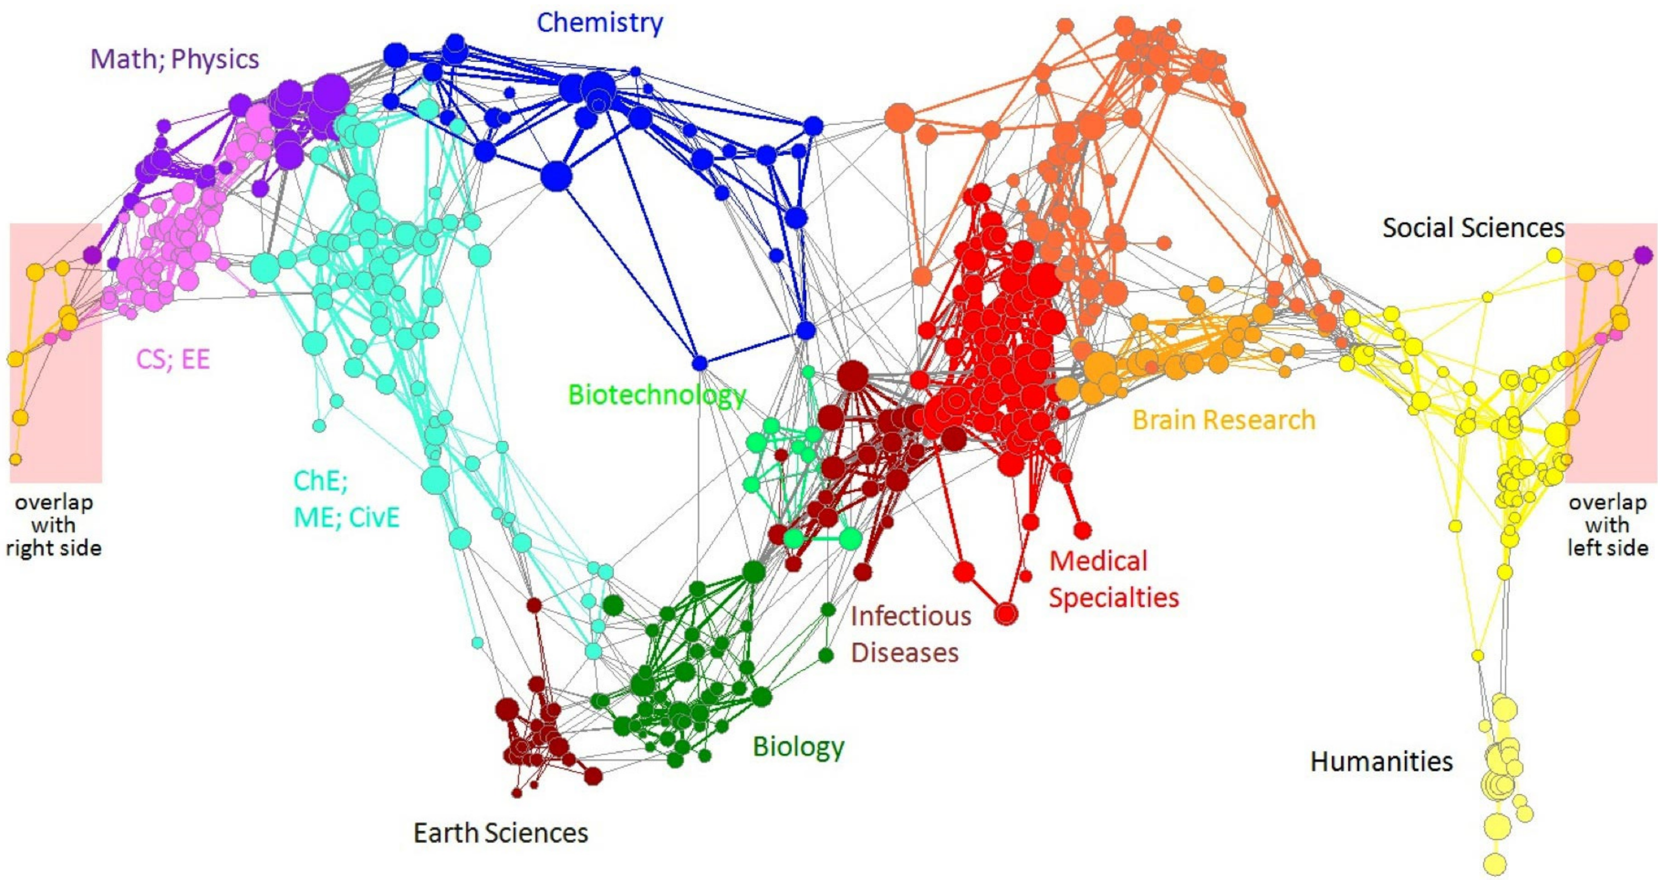


**Citation networks and Maps of science**

**[Börner et al., 2012]**

# Graph Data: Communication Nets

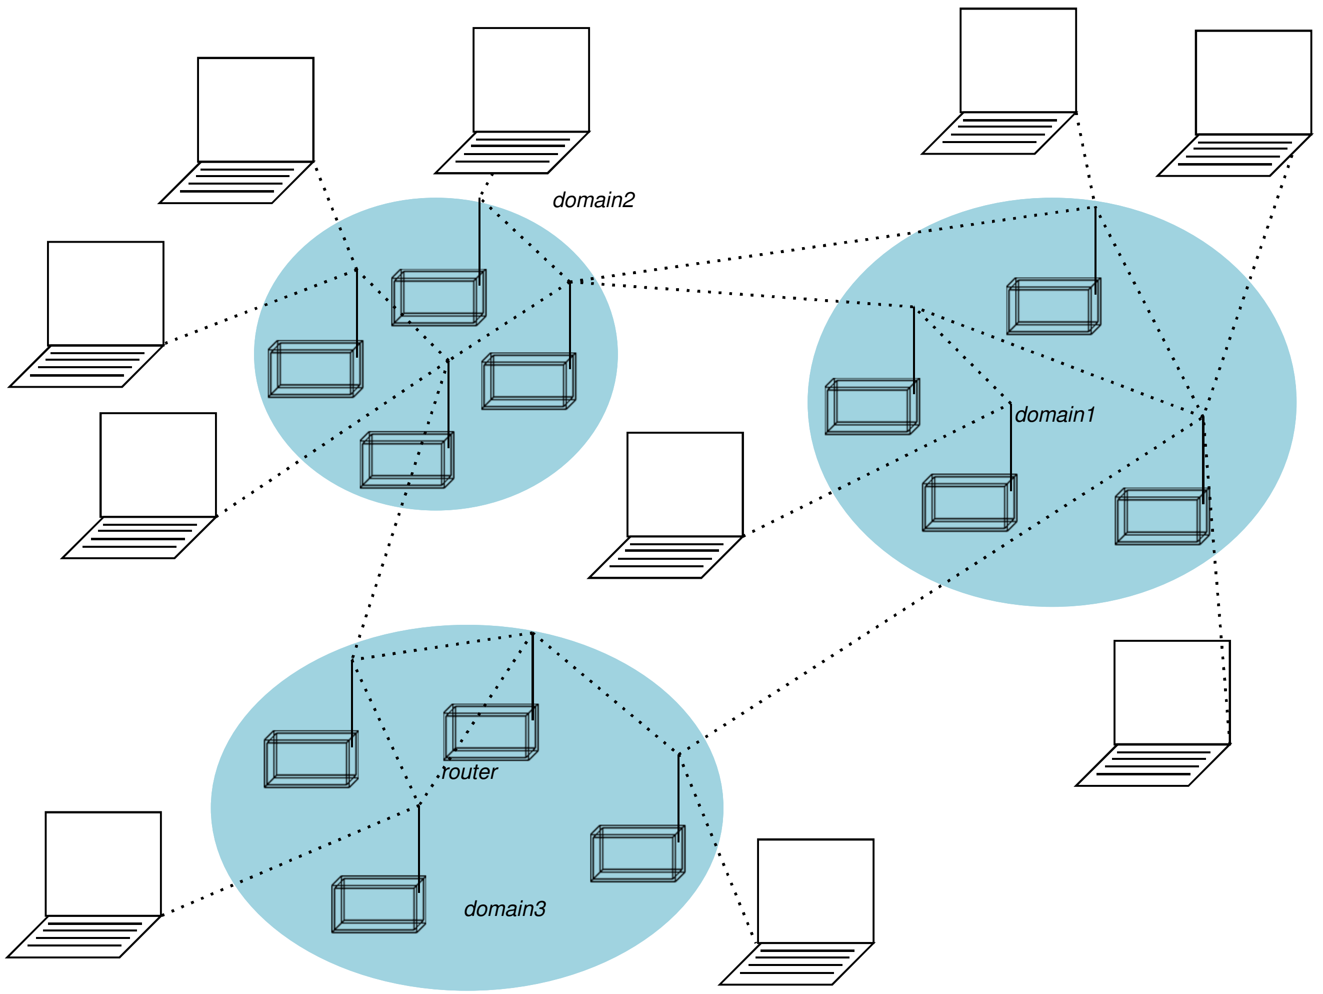


**Internet**

# (5) Networks: Biomedicine

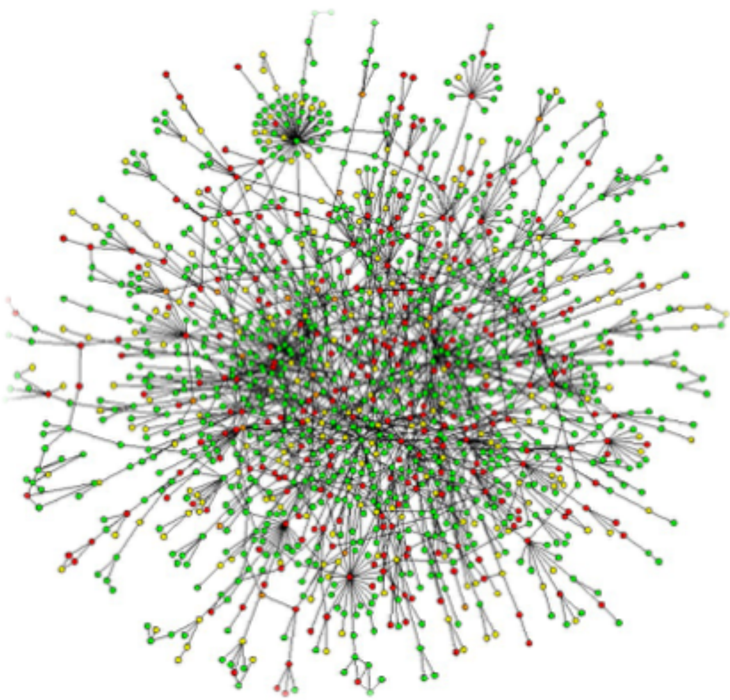

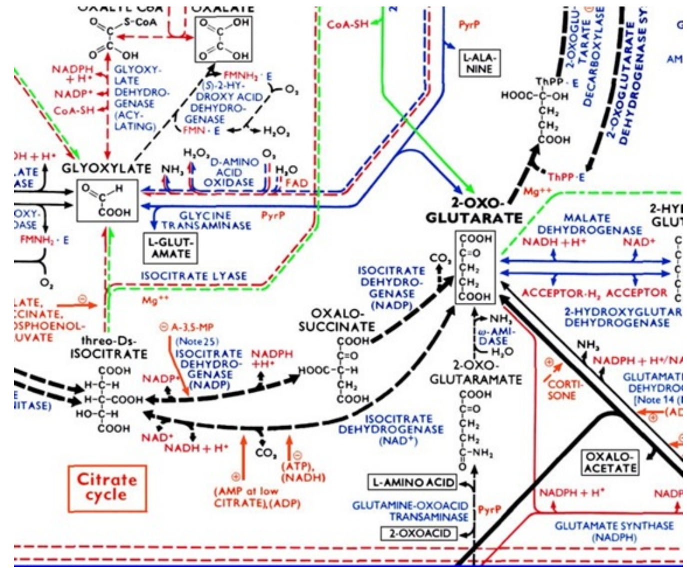


**Protein-protein interaction (PPI) networks:**

Nodes: Proteins

Edges: 'Physical' interactions

**Metabolic networks:**

Nodes: Metabolites and enzymes

Edges: Chemical reactions

# Graphs : Machine learning

- Complex domains (knwoledge, text, images, etc) have rich relational structure, which can be represented as a relational graph

- By explicitly modeling relationships we achieve better performance

# Different kinds of graph analysis

- **Path analysis:** for route optimization that is particularly applicable to logistics, supply and distribution chains and traffic optimization for smart cities.
- **Connectivity analysis:** tor determining weaknesses in networks such as a utility power grid, comparing connectivity across networks.
- **Community analysis:** Distance and density–based analysis is used to find groups of interacting people in a social network, for example, and identifying whether they are transient and predicting if the network will grow.
- **Centrality analysis:** find the most influential people in a social network, for example, or to find most highly accessed web pages—such as by using the PageRank algorithm.

# Graph-Parallel Systems

Pregel oogle

APACHE GIRAPH

GraphLab

*Exploit graph structure to achieve orders-of-magnitude performance gains over more general data-parallel systems.*

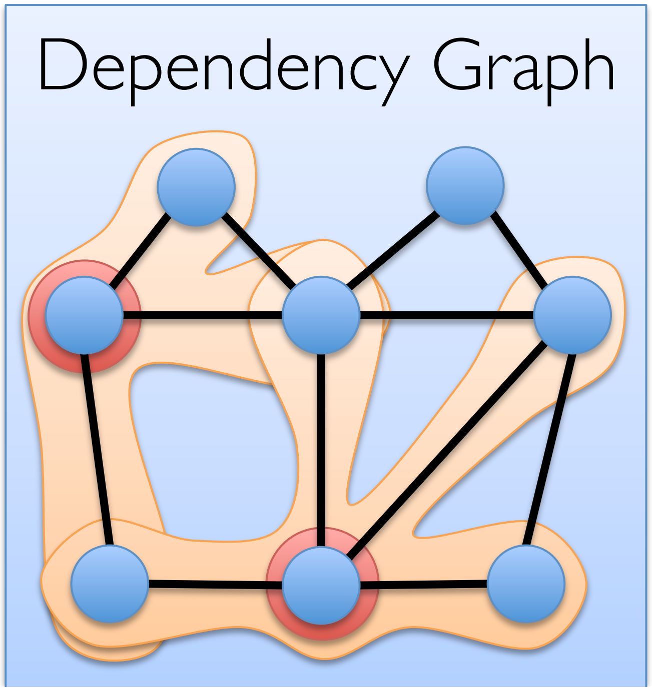

# Agenda

- Graph analytics
  - Link Analysis
  - Community detection
  - Graph similarity, Graph clustering, ..
- Machine learning with graphs : Graph Neural Networks
- Frameworks for parallel graph analytics
  - Pregel − a model for parallel-graph computing
  - GraphX Spark − unifying graph- and data −parallel computing
- Practical work : graph analytics with GraphX

# Systems, paradigms and languages for Big Data analytics and machine learning

- Follwow up of ADB course
- Focus on
  - large scale, map-reduce based data processing via Spark
  - Principles and techniques behind RDD, Dataframes, Datasets (in Scala)
  - Partitioning and shuffle-and-sort tuning
  - Large scale machine learning in SparkML
    - end-to-end pipelines for regression and classification
    - from-scratch map-reduce implementation in Spark of various gradient-descendant techniques, from batch to AdaGrad.
- Lab sessions
  - >= 50%
- Evaluation : written exam