Probabilistic graphical models: Introduction and general information

Francis Bach, INRIA/ENS
Guillaume Obozinski, ENPC

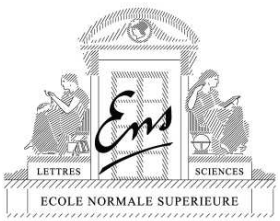

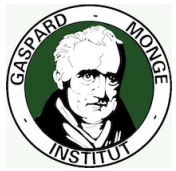

M2 MVA 2014-2015

General information

- Every Wed 9am-12pm amphi Curie until Dec 3.
- Except Nov 26 : no lecture
- **Grading** :
  - Homework 1 (20%)
  - Homework 2 (20%)
  - Take Home Exam (a longer Homework) (30%)
  - Project (30%)
- **Programming** :
  - All Hwk + Exam + Project involve programming
  - You may choose the programming language you want
  - We recommend you choose a vector oriented PL such as Python, R Matlab.
- **Polycopié** :
  - The course will be based on the book in preparation of Michael Jordan (UC Berkeley).
- **Scribes** :
  - For some lectures, the students are encouraged to volunteer as scribes to contribute to the lecture notes.

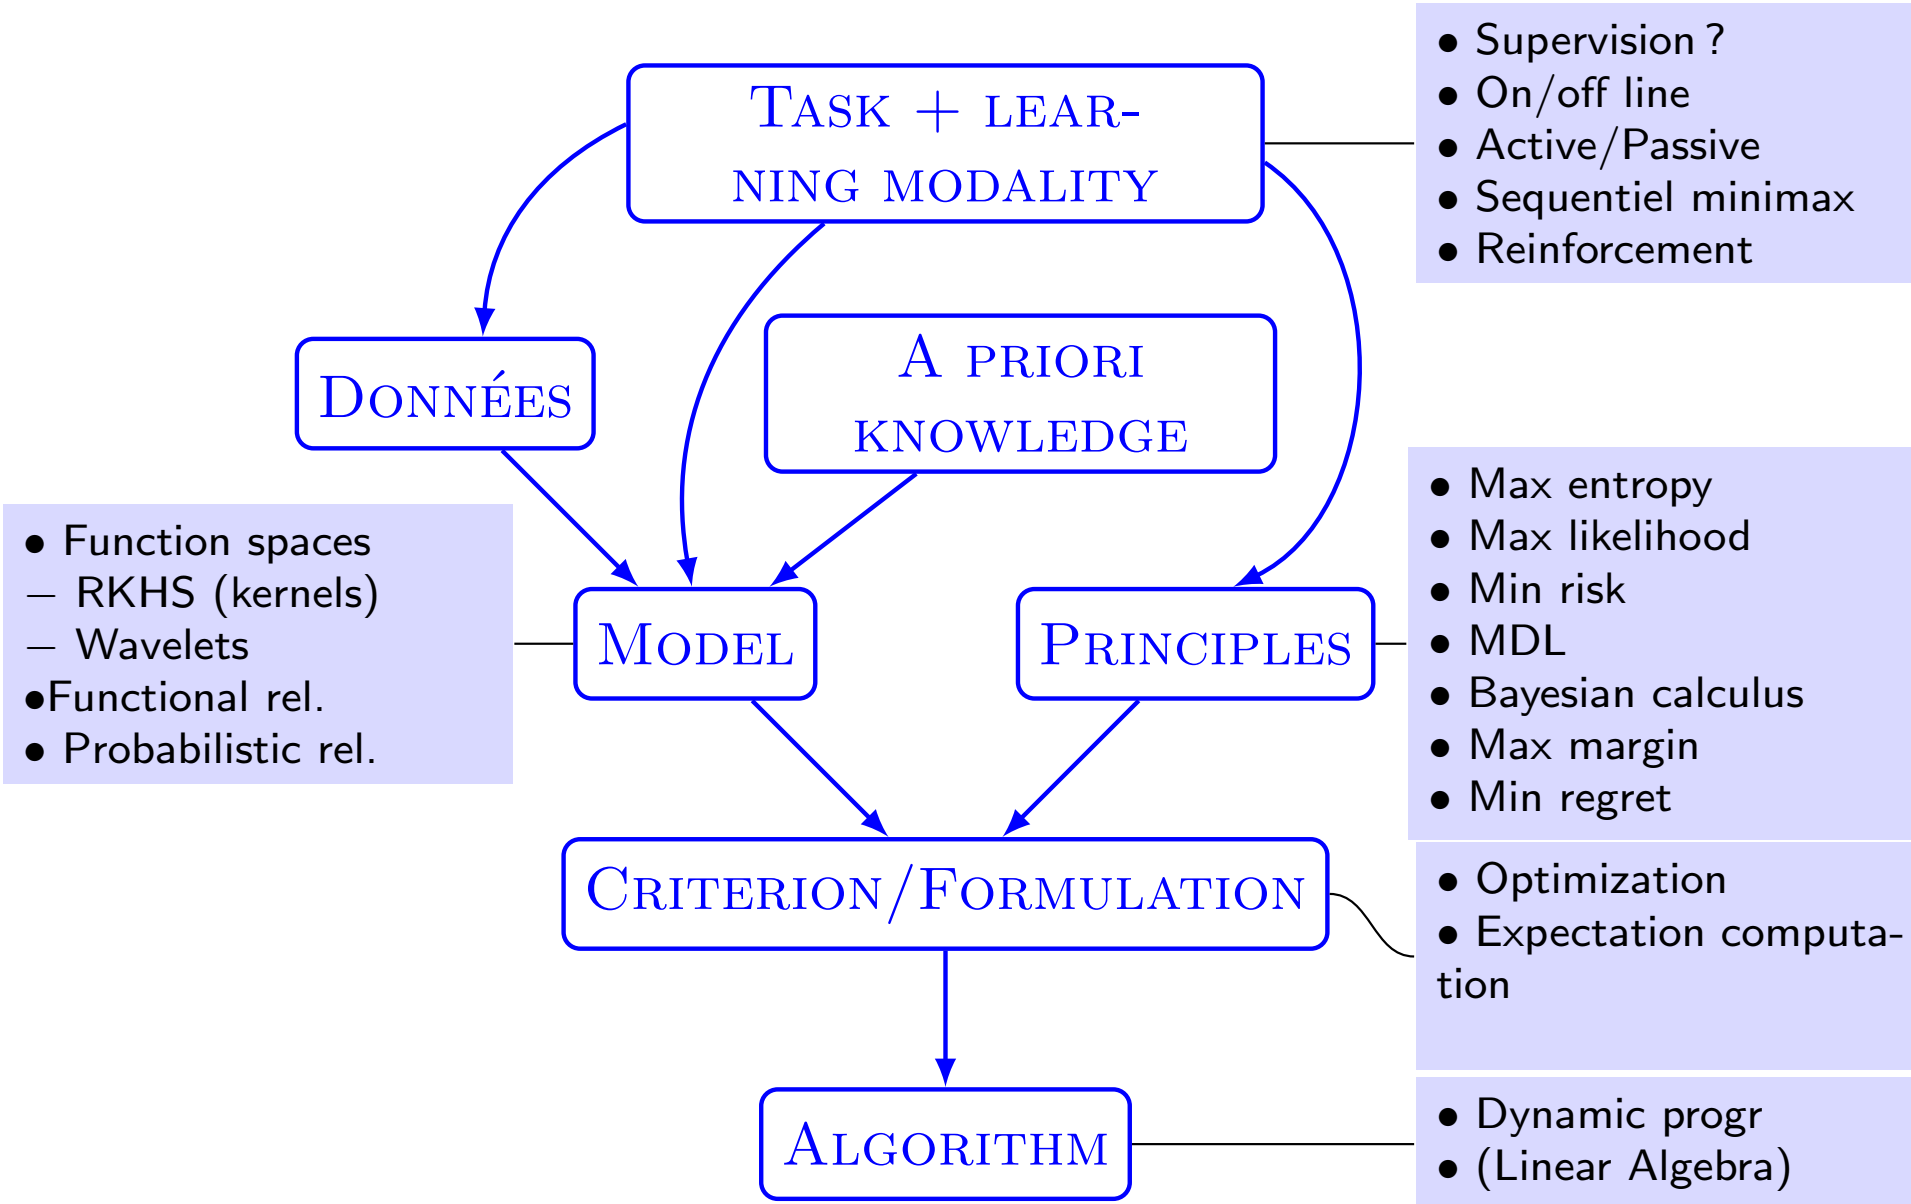
Task + learning modality
Données
A priori knowledge
Model
Principles
Criterion/Formulation
Algorithm
• Supervision ?
• On/off line
• Active/Passive
• Sequentiel minimax
• Reinforcement
• Function spaces
– RKHS (kernels)
– Wavelets
•Functional rel.
• Probabilistic rel.
• Max entropy
• Max likelihood
• Min risk
• MDL
• Bayesian calculus
• Max margin
• Min regret
• Optimization
• Expectation computation
• Dynamic progr
• (Linear Algebra)

Structured problems in HD

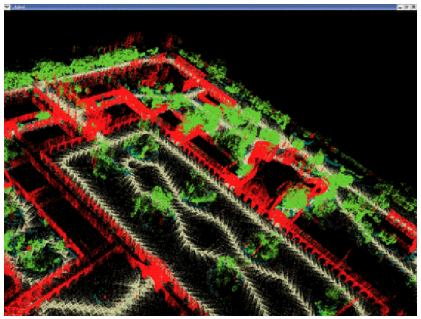

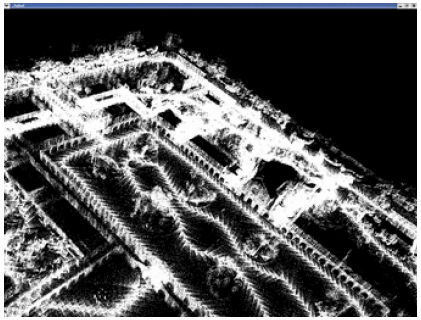

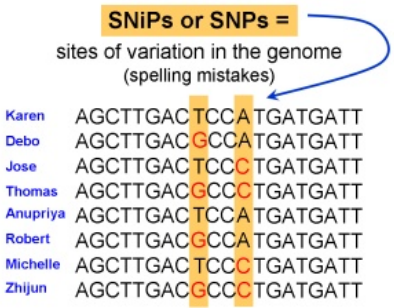

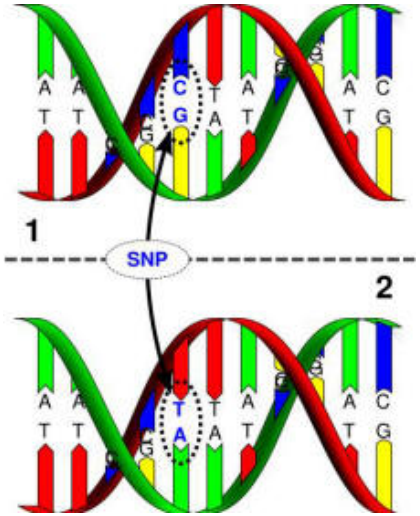

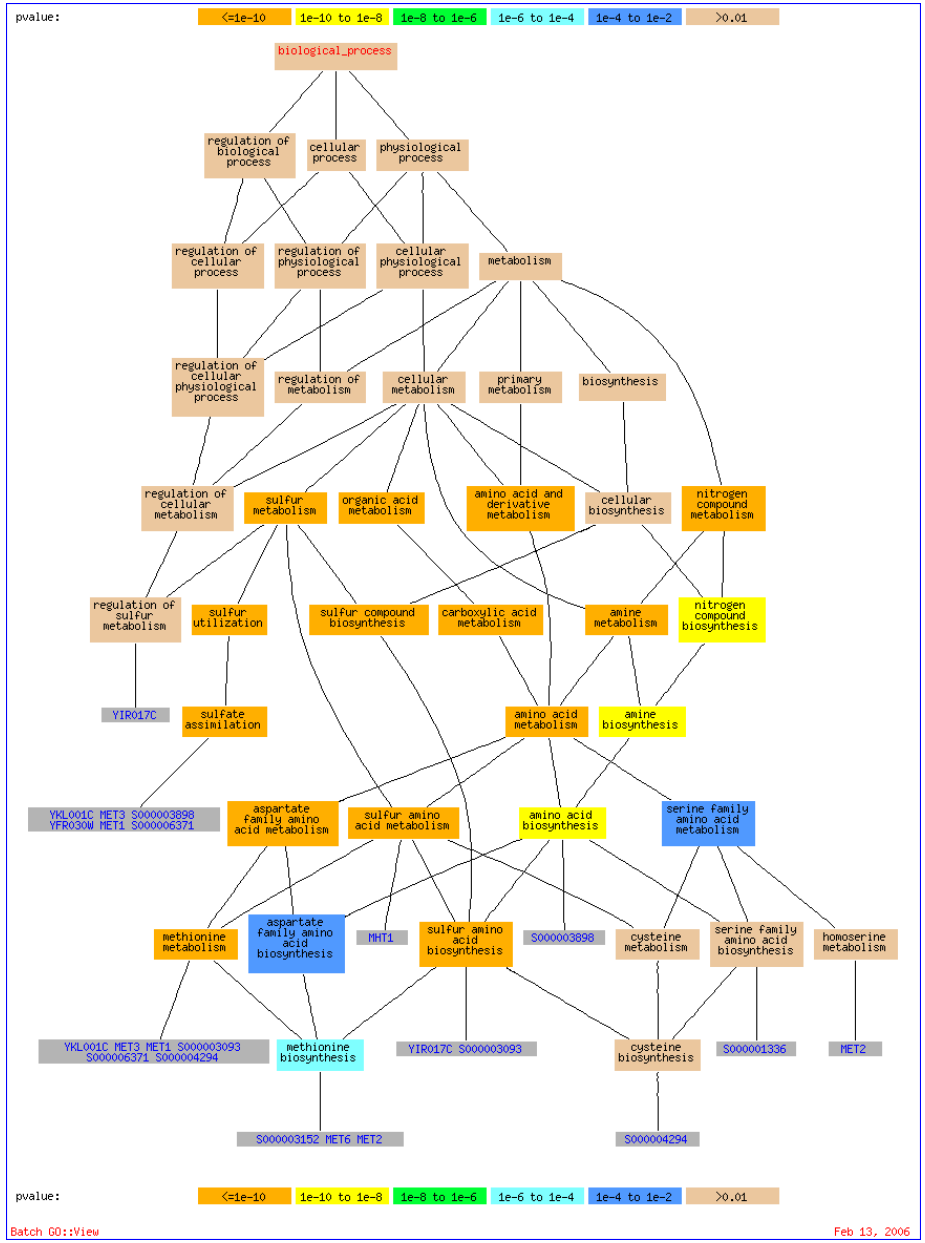

Sequence modelling

How to model the distribution of DNA sequences of length k ?

- Naive model $\rightarrow 4^n - 1$ parameters
- Indépendant model $\rightarrow 3n$ parameters

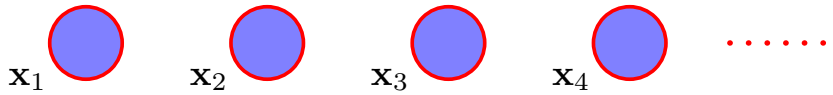

First order Markov chain :

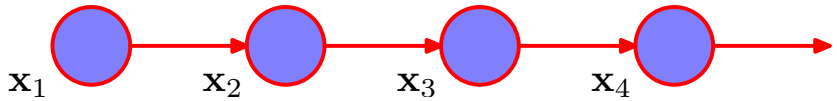

Second order Markov chain :

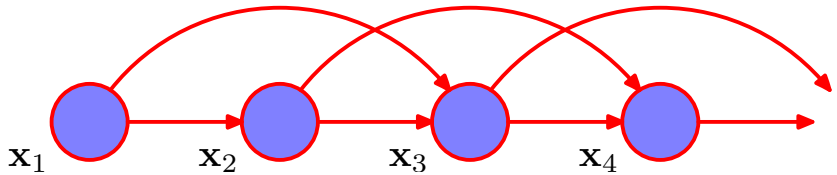

Number of parameters $\mathcal{O}(n)$ for chains of length n.

Models for speech processing

- Speech modelled by a sequence of unobserved phonemes
- For each phoneme a random sound is produced following a distribution which characterizes the phoneme

Hidden Markov Model : HMM (Modèle de Markov caché)

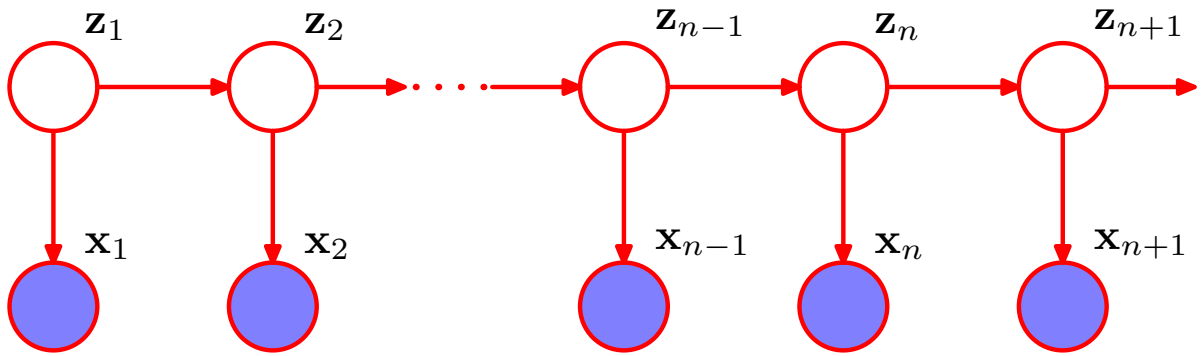

$\rightarrow$ **Latent** variable models

Modelling image structures

Markov Random Field

(Champ de Markov caché)

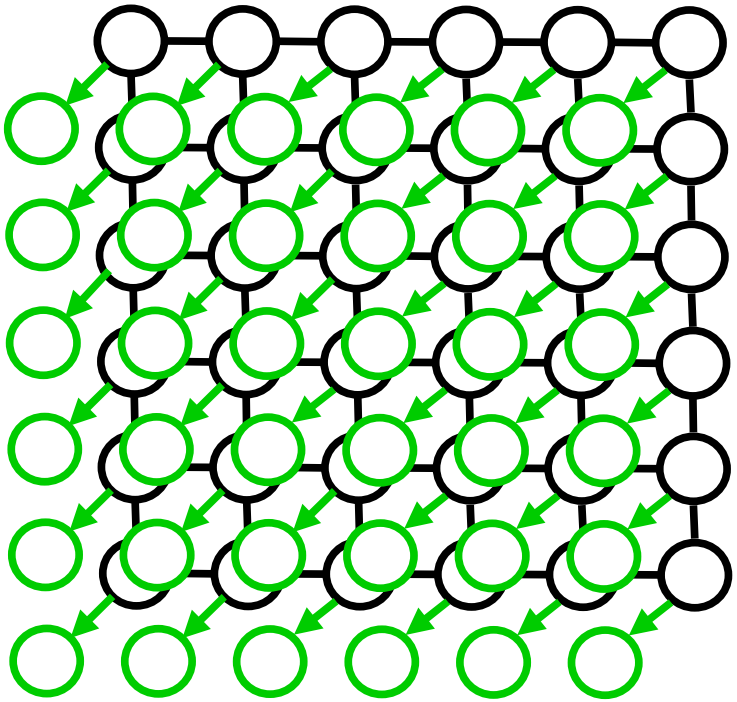

Original image

Segmentation

$\rightarrow$ *oriented graphical model* vs *non oriented*

Anaesthesia alarm (Beinlich et al., 1989)

"The ALARM Monitoring system"

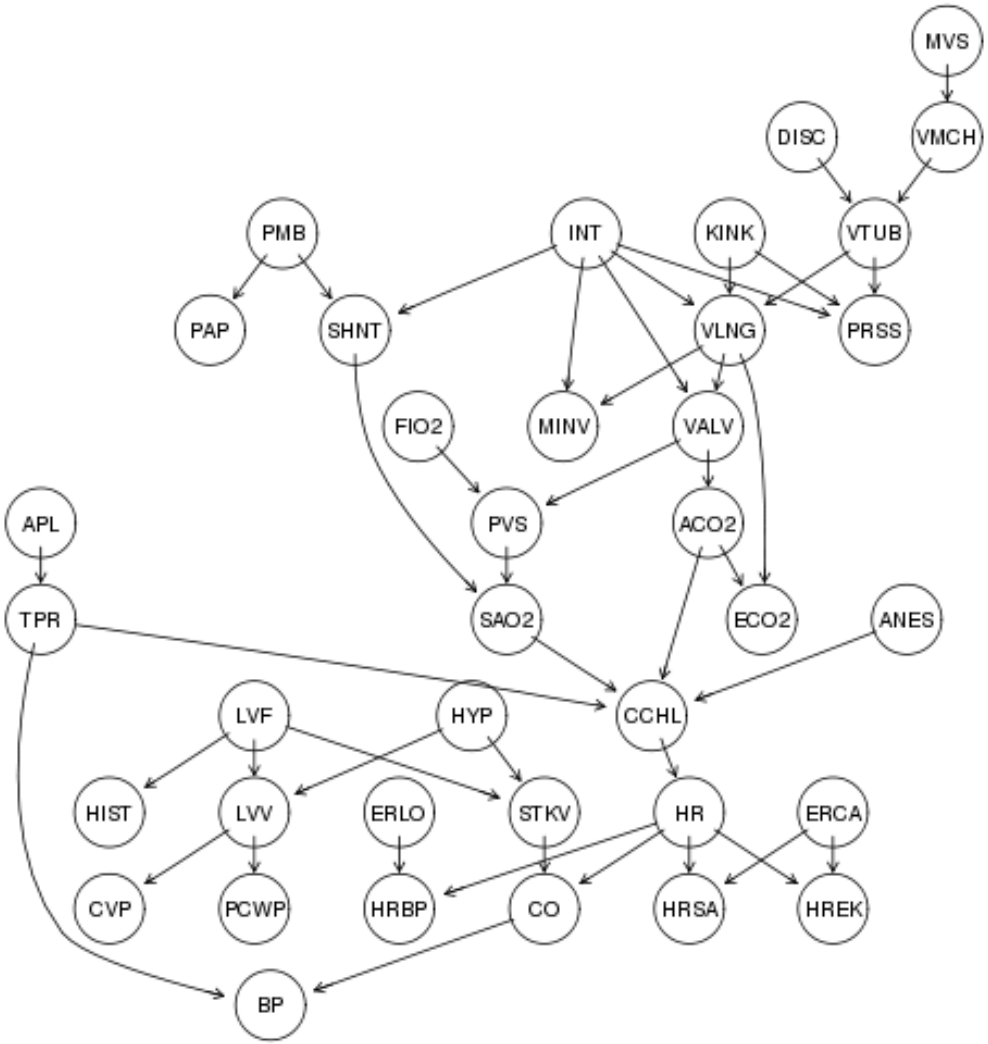

CVP	central venous pressure
PCWP	pulmonary capillary wedge pressure
HIST	history
TPR	total peripheral resistance
BP	blood pressure
CO	cardiac output
HRBP	heart rate / blood pressure.
HREK	heart rate measured by an EKG monitor
HRSA	heart rate / oxygen saturation.
PAP	pulmonary artery pressure.
SAO2	arterial oxygen saturation.
FIO2	fraction of inspired oxygen.
PRSS	breathing pressure.
ECO2	expelled CO2.
MINV	minimum volume.
MVS	minimum volume set
HYP	hypovolemia
LVF	left ventricular failure
APL	anaphylaxis
ANES	insufficient anesthesia/analgesia.
PMB	pulmonary embolus
INT	intubation
KINK	kinked tube.
DISC	disconnection
LVV	left ventricular end-diastolic volume
STKV	stroke volume
CCHL	catecholamine
ERLO	error low output
HR	heart rate.
ERCA	electrocauter
SHNT	shunt
PVS	pulmonary venous oxygen saturation
ACO2	arterial CO2
VALV	pulmonary alveoli ventilation
VLNG	lung ventilation
VTUB	ventilation tube
VMCH	ventilation machine

http://www.bnlearn.com/documentation/networks/

Probabilistic model

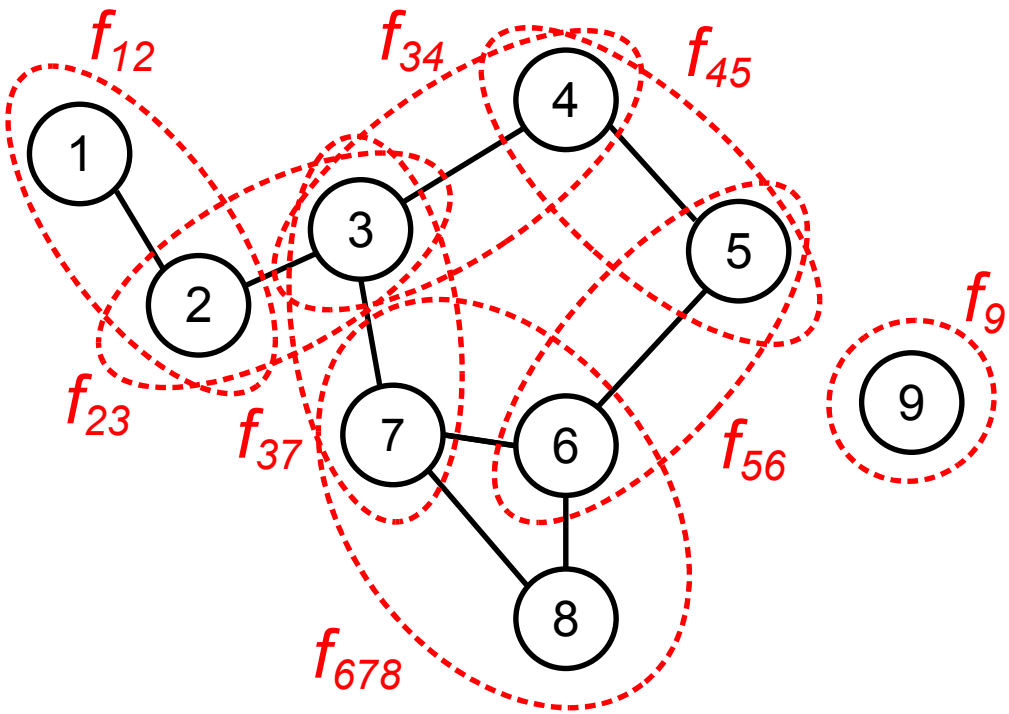

$$p(x_1, x_2, \ldots, x_9) = f_{12}(x_1, x_2)\, f_{23}(x_2, x_3)\, f_{34}(x_3, x_4)\, f_{45}(x_4, x_5) \ldots$$
$$f_{56}(x_5, x_6)\, f_{37}(x_3, x_7)\, f_{678}(x_6, x_7, x_8)\, f_9(x_9)$$

Abstact models vs concrete ones

Abstracts models

- Linear regression
- Logistic regression
- Mixture model
- Principal Component Analysis
- Canonical Correlation Analysis
- Independent Component analysis
- LDA (Multinomiale PCA)
- Naive Bayes Classifier
- Mixture of experts

Concrete Models

- Markov chains
- HMM
- Tree-structured models
- Double HMMs
- Oriented acyclic models
- Markov Random Fields
- Star models
- Constellation Model

Operations on graphical models

Probabilistic inference

Computing a marginal distr. $p(x_i)$ ou $p(x_i|x_1 = 3, x_7 = 0)$

Decoding (MAP inference)

What is the most likely instance ?

$$\text{argmax}_z p(z|x)$$

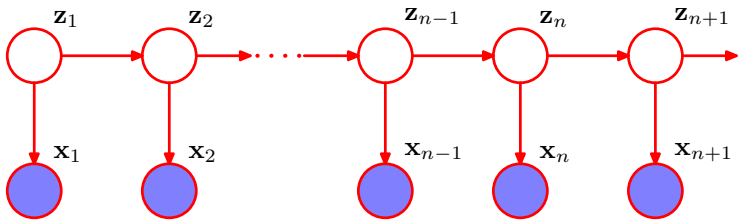

Learning (or Estimation)

Soit $p(x;\boldsymbol{\theta}) = \frac{1}{Z(\boldsymbol{\theta})} \prod_C \psi(x_C, \theta_C)$, we want to find

$$\text{argmax}_{\boldsymbol{\theta}} \prod_{i=1}^{n} p(x^{(i)};\boldsymbol{\theta}) = \text{argmax}_{\boldsymbol{\theta}} \frac{1}{Z(\boldsymbol{\theta})} \prod_{i=1}^{n} \prod_C \psi(x_C^{(i)}, \theta_C)$$

Course outline

- **Course 1**
  Introduction
  Maximum likelihood
  Models with a single node
- **Course 2**
  Linear regression
  Logistic regression
  Generative classification (Fisher discriminant)
- **Cours 3**
  K-means
  EM
  Gaussian mixtures
  Graph Theoretic aspects
- **Cours 4**
  Unoriented graphical models
  Oriented graphical models
- **Cours 5**
  Exponential families
  Information Theory
- **Cours 6**
  Gaussian Variables
  Factorial Analysis
- **Cours 7**
  Sum-product algorithm
- **Cours 8**
  Approximate inférence
- **Cours 9**
  Bayesian methods

To start : models with 1 and 2 nodes...

Regression and classification

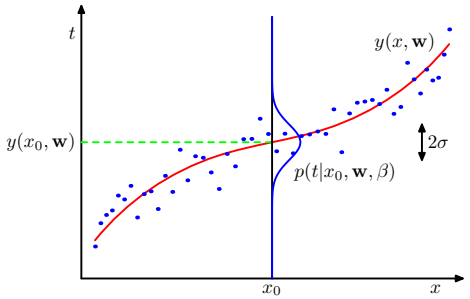

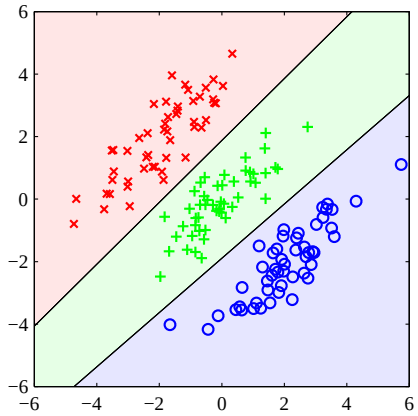

Mixture models