Master of Data Science

Alexandre Aussem

LIRIS UMR 5205 CNRS Data Mining & Machine Learning Group (DM2L) University of Lyon 1 Web: perso.univ-lyon1.fr/alexandre.aussem

November 4, 2024



Présentation du cours PGM

Total de 30 heures de cours :

- 18h CM (6 séances),
- 9h TD (3 séances)
- 3h de TP (A faire chez soi)
- Travaux pratiques :
 - TP avec PyAgruM, librairie Python pour l'apprentisage de la structure et des paramètres des réseaux bayésiens dynamiques et pour l'inférence.



Outline

Independence Models

- Conditional independence
- Graphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- 4 Parameter & Structure Learning
- 5 Hidden Markov Models



L Independence Models

Outline

1

Independence Models

- Conditional independence
- Graphoids

2 PGM

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- 4 Parameter & Structure Learning
- 5 Hidden Markov Models



L Independence Models

Conditional independence

Outline

1 Independence Models

- Conditional independence
- Graphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- Parameter & Structure Learning
- Hidden Markov Models



Independence Models

Conditional independence

Modeling complex data

- To model complex data, several questions have to be answered:
 - What is the task and the loss function?
 - What are the statistical properties and assumptions and underlying the data generating process?
 - What have to be captured from the probabilistic distribution to perform the task ?
 - How to learn the model parameters and perform inference in reasonable time?
- Once the model is chosen, two more issues:
 - Learning of the parameters of the model.
 - Inference of probabilistic queries



L Independence Models

Conditional independence



- Image: In a monochromatic image, each pixel is represented by a discrete random variable. The image may be modelled using a Markov network.
- **Bioinformatics:** Consider a long sequence of ADN bases. If each base of this sequence is modelled by a discrete random variable taking values in {*A*, *C*, *G*, *T*}, the sequence may be modeled by a Markov chain.



L Independence Models

Conditional independence



- Speech processing: consider the syllables of a word represented as a random signal. To retrieve the words from the signals, we may use a hidden Markov model.
- Text: The text may be modelled by a vector whose components are the keyword appearance frequency, i.e. "bag of words" model. A *naive Bayes classifier* works well for spam detection although the order of the words and the correlation between the keywords frequencies are not taken into account.



Independence Models

Conditional independence

Complexity vs. tractability

- Poor models are usually based on simple independence assumptions among variables that are rarely met in practice but they are easy to learn.
- In contrast, rich models allow complex statistical interactions to be captured but are difficult to learn (lack of data) and computationally demanding.
- In practice, one has to achieve a trade off for the model to be able to generalize well (statistical point of view) while keeping the computational burden of training and inference as low as possible (tractable computations).



L Independence Models

Conditional independence

Basic properties

Fundamental rules of probability. Let X and Y be two random variables,

Sum rule:

$$p(X) = \sum_{Y} p(X, Y).$$

Product rule:

$$p(X,Y) = p(Y|X)p(X).$$

Independence. X and Y are independent iff

$$p(X,Y) = p(X)p(Y).$$



L Independence Models

Conditional independence

Basic properties

Conditional independence. Let X, Y, Z be random variables.

We define X and Y to be conditionally independent given Z if and only if

$$p(X, Y|Z) = p(X|Z)p(Y|Z).$$

Property: If X and Y are conditionally independent given Z, then

$$p(X|Y,Z) = p(X|Z).$$



L Independence Models

Conditional independence

Basic properties

Independent and identically distributed. A set of random variables is independent and identically distributed (*i.i.d.*) if each variable has the same probability distribution and they are jointly independent.

Bayes formula. For two random variables X, Y we have

$$p(X|Y) = \frac{p(Y|X)p(X)}{p(Y)}.$$



Independence Models

Conditional independence

Conditional independence

Let X, Y and Z denote 3 disjoint sets of random variables defined on $\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$,

Definition

X and **Y** are **conditional independent** given **Z**, denoted **X** \perp **Y** | **Z**, iff \forall (**x**, **y**, **z**) $\in \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ such that p(**z**) > 0:

 $p(\mathbf{x}, \mathbf{y} | \mathbf{z}) = p(\mathbf{x} | \mathbf{z}) p(\mathbf{y} | \mathbf{z})$

The condition is equivalent to

$$p(\mathbf{x}, \mathbf{y}, \mathbf{z})p(\mathbf{z}) = p(\mathbf{x}, \mathbf{z})p(\mathbf{y}, \mathbf{z})$$



L Independence Models

Conditional independence

Conditional independence

An alternative definition of conditional independence is:

Theorem

 $\mathbf{X} \perp \mathbf{Y} \mid \mathbf{Z}$ iff there exists two functions f and g such that

$$p(\mathbf{x}, \mathbf{y}, \mathbf{z}) = f(\mathbf{x}, \mathbf{z})g(\mathbf{y}, \mathbf{z}).$$

■ Proof: ⇒ holds trivially. To show the converse:

$$\begin{split} p(\mathbf{x}, \mathbf{y}, \mathbf{z}) p(\mathbf{z}) &= f(\mathbf{x}, \mathbf{z}) g(\mathbf{y}, \mathbf{z}) \sum_{\mathbf{x}', \mathbf{y}'} f(\mathbf{x}', \mathbf{z}) g(\mathbf{y}', \mathbf{z}), \\ p(\mathbf{x}, \mathbf{z}) p(\mathbf{y}, \mathbf{z}) &= f(\mathbf{x}, \mathbf{z}) \left(\sum_{\mathbf{y}'} g(\mathbf{y}', \mathbf{z}) \right) g(\mathbf{y}, \mathbf{z}) \left(\sum_{\mathbf{x}'} f(\mathbf{x}', \mathbf{z}) \right), \\ & \text{Université Claude Bernard (Solution Claude Bernard (Sol$$

(日) (日) (日) (日) (日) (日) (日)

Independence Models

Conditional independence

Independence models

- Conditional independences inferred from data by means of statistical independence tests can be used to learn the structure of probabilistic graphical models.
- An independence model has an axiomatic characterization or properties that allows to build formal deductive system.

Definition

An independence model *I* over a set V consists in a set of triples $\langle X, Y | Z \rangle$, called independence relations, where X, Y and Z are disjoint subsets of V.

- Equivalently, I_1 is a dependence map for I_2 (D-map), if $I_2 \subseteq I_1$. Finally,
- I I_1 is a **perfect map** for I_2 (P-map), if $I_1 = I_2$.



Independence Models

Conditional independence

Independence models

Definition

A probability distribution p defined over V is said **faithful** to an independence model I when all and only the independence relations in I hold in p, that is,

$$\langle \mathbf{X}, \mathbf{Y} \mid \mathbf{Z} \rangle \in I \iff \mathbf{X} \perp \mathbf{Y} \mid \mathbf{Z} \quad \text{w.r.t. } p.$$

An independence model *I* is said probabilistic, if there exists a probability distribution *p* which is faithful to it.



Independence Models

Graphoids



1

Independence Models

- Conditional independence
- Graphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- 4 Parameter & Structure Learning
- 5 Hidden Markov Models



L Independence Models

Graphoids

Semi-graphoids

- Consider four mutually disjoint random variables, W, X, Y and Z. The following properties hold for any probability distribution:
 - $\blacksquare \text{ Symmetry: } \langle \mathbf{X}, \mathbf{Y} \mid \mathbf{Z} \rangle \iff \langle \mathbf{Y}, \mathbf{X} \mid \mathbf{Z} \rangle.$
 - Decomposition: $\langle \mathbf{X}, \mathbf{Y} \cup \mathbf{W} \mid \mathbf{Z} \rangle \implies \langle \mathbf{X}, \mathbf{Y} \mid \mathbf{Z} \rangle.$
 - $\blacksquare \text{ Weak Union: } \langle \mathbf{X}, \mathbf{Y} \cup \mathbf{W} \mid \mathbf{Z} \rangle \implies \langle \mathbf{X}, \mathbf{Y} \mid \mathbf{Z} \cup \mathbf{W} \rangle.$
 - Contraction:
 - $\langle \mathbf{X}, \mathbf{Y} \mid \mathbf{Z} \rangle \land \langle \mathbf{X}, \mathbf{W} \mid \mathbf{Z} \cup \mathbf{Y} \rangle \implies \langle \mathbf{X}, \mathbf{Y} \cup \mathbf{W} \mid \mathbf{Z} \rangle.$
- Any independence model that respects these four properties is called a semi-graphoid.



L Independence Models

Graphoids



- Another property holds in strictly positive distributions, that is when p > 0:
 - Intersection: $\langle \mathbf{X}, \mathbf{Y} \mid \mathbf{Z} \cup \mathbf{W} \rangle \land \langle \mathbf{X}, \mathbf{W} \mid \mathbf{Z} \cup \mathbf{Y} \rangle \implies \langle \mathbf{X}, \mathbf{Y} \cup \mathbf{W} \mid \mathbf{Z} \rangle.$
- Any independence model that respects these five properties is called a **graphoid**.



Independence Models

Graphoids

The characterization problem

- It is possible to detect contradictory conditional independence relations, by checking if they respect the semi-graphoid properties.
- Do the semi-graphoid properties provide a sufficient condition to characterize a probabilistic independence model? No!
- Uncompleteness: the graphoid axioms are insufficient to characterize probabilistic independence models.



L Independence Models

Graphoids

The characterization problem

The following set of probabistic CI relations (and their symetric counterparts) satisfies the graphoid axioms, yet does not have any faithful probability distribution:

 $\langle A,B \mid \{C,D\} \rangle \land \langle C,D \mid A \rangle \land \langle C,D \mid B \rangle \land \langle A,B \mid \emptyset \rangle.$

- In fact, no finite set of CI properties characterizes the probabilistic independence models.
- Probabilistic independencies have no finite complete axiomatic characterization.



Independence Models

Graphoids

The characterization problem



Figure: p denotes CI of probability distribution, and p > 0 stands formation strictly positive probability distribution.

Outline

Independence Models

Conditional independenceGraphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- Parameter & Structure Learning
- Hidden Markov Models



- PGMs

Probabilistic Graphical models

- A probabilistic graphical model (PGM) represents graphically a joint distribution.
- The nodes in the graph represent random variables, and the (lack of) edges represent conditional independence (CI) assumptions.
- Several useful properties:
 - Provide a simple way to visualize the probabilistic structure of a joint probability distribution.
 - Insights into the CI properties can be obtained by inspection of the graph.
 - Complex computations, required to perform inference and learning can be expressed in terms of graphical manipulations.

Université Claude Bernard

Several kinds of PGMs: directed, undirected, mixed etc.

Graphs

- A graph \mathcal{G} is an ordered pair of sets $(\mathbf{V}, \mathcal{E})$. $\mathbf{V} = \{V_1, \dots, V_n\}$, are the *nodes* (or vertices), \mathcal{E} represents the *edges*.
- A clique is a set of nodes such that each node is adjacent to every other node in the set. A maximal clique is a clique that does not accept any other clique as a proper superset.
- A walk between two nodes V_1 and V_k is a sequence of adjacent nodes in the form V_1, \ldots, V_k . A walk with only distinct nodes is called a *path*. A path with $V_1 = V_k$ is called a *cycle*.
- A *complete graph* is a graph that has only one maximal clique.
- A *chordal graph* is a graph in which every cycle with more than 3 distinct nodes admits a smaller cycle as a proper subset.



Graphical models



 Directed, undirected and mixed graphs edges. The expressiveness of these models differ.

(a)

Probabilistic graphical models

- A PGM always consists in a set of parameters Θ and a graphical structure *G*.
- *G* encodes a set of conditional independence relations between the variables and induces an independence model denoted *I*(*G*). By definition, *I*(*G*) is an I-map for *p*, that is,

$$\langle \mathbf{X}, \mathbf{Y} \mid \mathbf{Z} \rangle \in I(\mathcal{G}) \implies \mathbf{X} \perp \mathbf{Y} \mid \mathbf{Z} \quad \text{w.r.t. } p.$$

- G allows explicit modelling of expert knowledge in the form of conditional independencies.
- And the edges provide a convenient way of communicating the investigator's beliefs of the causal influences among variables



Probabilistic graphical models

- A PGM is a compact graphical model for a joint distribution.
- The relationship between factorization, conditional independence, and graph structure comprises much of the power of the graphical modeling framework:
- The conditional independence viewpoint is most useful for designing models.
- The factorization viewpoint is most useful for designing inference algorithms.
- **Problems**: structure learning (\mathcal{G}), parameter learning (Θ), and inference using the model (e.g. $P(\mathbf{X} \mid \mathbf{Y})$).



Undirected graphical models

Outline

Independence Models

Conditional independenceGraphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- Parameter & Structure Learning
- 6 Hidden Markov Models



Undirected graphical models

Markov Networks

Markov networks (MNs), also called Markov random fields, are the most popular graphical models based on undirected graphs

Factorization

The probability distribution p factorizes as

$$p(\mathbf{x}) = \frac{1}{Z} \prod_{\mathbf{C} \in \mathcal{C}l_{\mathcal{G}}} \phi_C(\mathbf{x}_C), \text{ with } Z = \sum_{\mathbf{x}} \prod_{\mathbf{C} \in \mathcal{C}l_{\mathcal{G}}} \phi_C(\mathbf{x}_C)$$

Cl_G is the set of all cliques in G. Z is called the *partition function*. Each ϕ_i function is called a **factor**, a potential function, or a clique potential.

Université Claude Bernard

-

・ ロ ト ・ 雪 ト ・ 雪 ト ・ 日 ト

• $\phi_C(x_C) \ge 0$ ensures that $p(\mathbf{x}) \ge 0$.

PGMs

Undirected graphical models



• $\langle \mathbf{X}, \mathbf{Y} | \mathbf{Z} \rangle$ belongs to $I(\mathcal{G})$ iff \mathbf{Z} *u*-separates \mathbf{X} and \mathbf{Y} in \mathcal{G} , that is, every path between a node in \mathbf{X} and a node in \mathbf{Y} contains a node in \mathbf{Z} .



From the graph, we see that $\{A, B\} \perp \{D, E, F\} \mid C$ holds but not $\{A, B\} \perp F \mid E$.



Undirected graphical models



- Let *G* be an undirected graph over the random variables **V**, and *p* a probability distribution over **V**.
- \mathcal{G} is an **I-map** for p if for all $\mathbf{X}, \mathbf{Y}, \mathbf{Z} \in \mathbf{V}$,

$$\mathbf{X} \perp\!\!\!\perp_{\mathcal{G}} \mathbf{Y} \mid \mathbf{Z} \implies \mathbf{X} \perp\!\!\!\perp_{P} \mathbf{Y} \mid \mathbf{Z}.$$

Theorem

 $I(\mathcal{G})$ is an *I*-map for p if p factorizes into a product of potentials over the **cliques** in \mathcal{G} . The converse holds only if p > 0.



Undirected graphical models

Two undirected graphs





Undirected graphical models



The factorization over the maximal cliques are for each Markov network:

(Graph 1)
$$p(\mathbf{v}) = \phi_1(a, b)\phi_2(b, c)\phi_3(c, d)\phi_4(d, a)$$

(Graph 2) $p(\mathbf{v}) = \phi_1(a, b, d)\phi_2(d, b, c)$

In the case of binary variables, we may define the clique potentials in the form of numerical tables.



Undirected graphical models

Clique potentials

	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
(a) $\phi_1(a, b)$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
	A
0 1	0 1
$\begin{array}{c c c} C & 0 & 3/3 & 2/3 \\ 1 & 2/3 & 1/3 \end{array}$	$\begin{array}{c ccccc} D & 0 & 3/10 & 1/10 \\ 1 & 1/10 & 2/10 \end{array}$
(c) $\phi_3(c,d)$	(d) $\phi_4(d,a)$
	Université Claude Bernar

Lyon 1

æ

・ロト ・聞ト ・ヨト ・ヨト

L PGMs

Lundirected graphical models

Clique potentials

		B			
A	D	0	1		
0	0	1/8	3/8		
	1	1/8	1/8		
1	0	2/8	1/8		
	1	2/8	2/8		
(e) $\phi_1(a,b,d)$					

		B			
C	D	0	1		
0	0	1/6	4/6		
	1	1/6	1/6		
1	0	1/6	2/6		
	1	2/6	3/6		
(f) $\phi_2(b, c, d)$					


Undirected graphical models

Clique potentials

- These potentials are valid, i.e. $\sum_{a,b,c,d} p(a,b,c,d) = 1$.
- However individual clique potentials do not necessarily sum to 1, and therefore do not necessarily correspond to marginal or conditional probability distributions.
- Potential functions in Markov network do not lend to an intuitive probabilistic interpretation. One must go through a factorization to obtain a proper probability measure.
- Potential functions are often expressed as exponential parametric functions for practical reasons.



Undirected graphical models

Clique potentials

- Not every probability distribution is UG-faithful.
- The probabilistic IC relation $X \perp Y$ and $X \not \perp Y \mid Z$ cannot be faithfully represented an UG model because $X \perp Y \mid \emptyset$ necessarily implies $X \perp Y \mid Z$ (strong union property) in UG models.
- As a result, the only undirected graph that is an I-map for p is the complete graph, which necessarily results in 7 free parameters instead of 6, as p(x, y, z) = p(x)p(y)p(z|x, y).
- A Markov network is not perfectly suited to encode *p* in this situation.



Undirected graphical models



The notions of Markov blanket and Markov boundary are essential in feature selection..

Definition

A *Markov blanket* of X in V is a subset $\mathbf{M} \subseteq (\mathbf{V} \setminus \mathbf{X})$ such that $\mathbf{X} \perp \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{M}) \mid \mathbf{M}$. A *Markov boundary* is an inclusion-optimal Markov blanket, i.e., none of its proper subsets is a Markov blanket.

In a faithful UG the Markov boundary of a variable X is unique and is given by neighbors of X.



Undirected graphical models

Image denoising using MRF

- An UG model representing a Markov random field for image de-noising,
- x_i is a binary variable denoting the state of pixel i in the unknown noise-free image,
- y_i denotes the corresponding value of pixel i in the observed noisy image.





Undirected graphical models

Image denoising using MRF

Because a potential function is an arbitrary, non-negative function over a maximal clique, we may define a joint distribution over x and y by

$$p(\mathbf{x}, \mathbf{y}) = \frac{1}{Z} \exp\{-E(\mathbf{x}, \mathbf{y})\}\$$

- Neighbouring pixels are correlated and only a small percentage of the pixels are corrupted.
- We want the energy to be lower when $\{x_i, x_j\}$ and $\{x_i, y_i\}$ have the same sign than when they have the opposite sign.
- The complete energy function takes the form

$$E(\mathbf{x}, \mathbf{y}) = h \sum_{i} x_i - \beta \sum_{\{i,j\} \in E} x_i x_j - \eta \sum_{i} x_i y_i$$

Université Claude Bernard (Golyon 1)

Undirected graphical models

Image denoising using MRF

Given y are the (observed) pixels of the noisy image, one has to solve the MAP:

$$\arg\max_{x} p(\mathbf{x}, \mathbf{y}) = \arg\max_{x} \exp\{-E(\mathbf{x}, \mathbf{y})\}$$

- A local maximum can be easily obtained by simple coordinate-wise gradient ascent methods.
- This is an example of the Ising model which has been widely studied in statistical physics.



Undirected graphical models

Image denoising using MRF

- Illustration of image de-noising using a Markov random field (Besag, 1974).
- On the top, the corrupted image after randomly changing 10% of the pixels. On the bottom, the restored images obtained using iterated conditional models (ICM)



L Directed graphical models

Outline

Independence Models

Conditional independenceGraphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

3 Inference and MAP Estimation

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- Parameter & Structure Learning
- 6 Hidden Markov Models



L Directed graphical models

Bayesian networks

Definition

A **Bayesian network** consists in a set of random variables $\mathbf{V} = \{V_1, \ldots, V_n\}$, a simple directed acyclic graph $\mathcal{G} = (\mathbf{V}, \mathcal{E})$, and a set of parameters Θ . Together, \mathcal{G} and Θ define a probability distribution p over \mathbf{V} which factorizes as:

$$p(\mathbf{v}) = \prod_{V_i \in \mathbf{V}} p(v_i | \mathbf{pa}_{V_i}).$$

- **p** \mathbf{p}_{V_i} denotes the *parents* of node V_i in \mathcal{G} .
- Θ are the probabilities $p(v_i | \mathbf{pa}_{V_i})$.



Probabilistic Graphical Models

L PGMs

L Directed graphical models





The corresponding factorization is

$$p(\mathbf{v}) = p(a)p(d|a)p(b|a)p(c|b,d)$$



Directed graphical models

Conditional probability tables

Λ						
	0 1			0	1	
	0.4 0.6	A	0	0.6	0.4	-
	(g) <i>p</i> (<i>a</i>)		(h) $p(d a)$			
						7
	В		B	D	0	1
	0 1	-	0	0	0.8	0.2
A	0 0.3 0.7 1 0.1 0.9		0	1	0.7	0.3
			4	0	0.5	0.5
	(i) $p(b a)$		I	1	0.7	0.3
			(j) $p(c b,d)$			

Table: A set of conditional probability tables that define a valid set of parameters Θ for the Bayesian network structure.

・ ロ ト ・ 雪 ト ・ 目 ト ・

э

Directed graphical models

Conditional probability tables

- These tables define valid conditional probability distributions that can be intuitively interpreted.
- Each of the factors $p(v_i | \mathbf{pa}_{V_i})$ can be seen as a potential function $\phi_i(v_i, \mathbf{pa}_{V_i})$ in a Markov network.
- In a Bayesian network, each factor defines a conditional probability distribution for V_i , and thus respects the normalization constraint $\sum_{v_i} \phi_i(v_i, \mathbf{pa}_{V_i}) = 1.$



L Directed graphical models

Parametric conditional distributions

- The number of parameters required to specify a PCT grows exponentially with *M* the number of parents.
- A more parsimonious form uses a logistic sigmoid function acting on a linear combination of the parents. Consider a graph comprising M parents x₁,..., x_M and a single child y,

$$p(y = 1 | x_1, \dots, x_M) = \sigma(w_0 + \sum_{i=1}^M w_i x_i)$$

- σ(a) = (1 + exp(-a))⁻¹ is the sigmoid function and w = (w₀, w₁..., w_M)^T is a vector of M + 1 parameters.
- The conditional distribution is now governed by a number of parameters that **grows linearly** with *M*.



Directed graphical models

Linear-Gaussian models

- A multivariate Gaussian can be expressed as a directed graph corresponding to a linear-Gaussian model.
- Examples of linear-Gaussian models: probabilistic principal component analysis, factor analysis, and linear dynamical systems.
- If node *i* represents a continuous random variable X_i having a Gaussian distribution of the form,

$$p(x_i \mid pa_i) = \mathcal{N}(x_i \mid \sum_{j \in pa_i} w_{ij} x_j + b_i, \sigma_i^2)$$

- w_{ij} and b_i are parameters governing the mean, and σ²_i is the variance of the conditional distribution.
- The mean and covariance of the joint distribution are determined recursively.



Directed graphical models



- Every DAG G induces a formal independence model I(G) over V, by means of a graphical separation criterion called d-separation
- Within a path V_1, \ldots, V_k , an intermediate node V_i is said to be a *collider* is an intermediate node V_i in the form $V_{i-1} \rightarrow V_i \leftarrow V_{i+1}$ called a *v*-structure.
- *d*-separation is equivalent to *u*-separation when *G* contains no *v*-structure



Directed graphical models



• Let $\mathbf{X} \perp _{\mathcal{G}} \mathbf{Y} \mid \mathbf{Z}$ denotes a CI relation encoded in a DAG \mathcal{G} .

Definition

For any disjoint set of random variables X, Y and Z, $X \perp _{\mathcal{G}} Y \mid Z$ *iff* Z *d*-separates X and Y in \mathcal{G} , that is, every path between X and Y contains

- a non-collider that belongs to Z,
- or a collider that does not belong to $\mathbf{Z} \cup \mathbf{AN}_{\mathbf{Z}}$.
- AN_Z are the ancestors of nodes Z.



Probabilistic Graphical Models

PGMs

Directed graphical models





We have $A \perp C \mid \{D, B\}$ and $D \perp B \mid A$ because $D \rightarrow C \leftarrow B$ is a closed path.



Directed graphical models



- A friendly interpretation of *d*-separation is to consider a path as an information flow.
- Consider a path between X and Y, and a conditioning set \mathbf{Z} .
- When Z is empty, each intermediate node that is not a collider is open, that is, it lets the flow go through. Conversely, each intermediate node that is a collider is closed, and blocks the flow.
- The variables in Z change the state of the nodes, i.e. from open to closed and vice-versa.
- If \mathbf{Z} *d*-separates *X* and *Y*, all the paths between *X* and *Y* are closed.



Probabilistic Graphical Models

L PGMs

L Directed graphical models

d-separation (again)





L Directed graphical models



- $A \perp B \mid C$ because the only path $A \rightarrow C \rightarrow B$ is closed by C that is observed.
- $A \not\perp F \mid \emptyset$, because of the open path $A \to C \to E \to F$.
- Conditioning on *E* does not *d*-separate *A* and *F* either, it closes the previous path but opens a new one with *A* → *C* → *E* ← *D*.
- To close all paths, it is sufficient to condition on $\{C, D\}$, *E* it no longer necessary in the conditioning set.



L Directed graphical models



• The notions of *Markov blanket* and *Markov boundary* are essential in feature selection..

Definition

A *Markov blanket* of X in V is a subset $\mathbf{M} \subseteq (\mathbf{V} \setminus \mathbf{X})$ such that $\mathbf{X} \perp \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{M}) \mid \mathbf{M}$. A *Markov boundary* is an inclusion-optimal Markov blanket, i.e., none of its proper subsets is a Markov blanket.

In a faithful DAG the Markov boundary of a variable X is unique and is given by $\mathbf{MB}_x = \mathbf{PC}_X \cup \mathbf{SP}_X$, that is, the parents, children and spouses of X.



L Directed graphical models



 As with Markov networks, a Bayesian network structure always defines an I-map of the underlying probability distribution.

Theorem

Let \mathcal{G} be a DAG, $I(\mathcal{G})$ is an I-map for p iff p factorizes recursively over \mathcal{G} .



L Directed graphical models

Local Markov property

- From the *d*-separation, every node is independent of its non-descendants given its parents (a.k.a. local Markov property), that is, V_i ⊥ ND_{Vi} \ PA_{Vi} | PA_{Vi}.
- Because G is a DAG, we may arrange its nodes in a topological ordering V₁,..., V_n according to G, that is, i < j if V_i → V_j is in G.
- From the chain rule of probabilities, we show that

$$p(\mathbf{v}) = \prod_{i=1}^{n} p(v_i | v_1, \dots, v_{i-1})$$
$$= \prod_{i=1}^{n} p(v_i | \mathbf{pa}_{V_i})$$



Directed graphical models

Famous networks used as benchmarks



L Directed graphical models

Graphs with three nodes

- A Markov chain is a particular DAG.
- We have $X \not\perp Y \mid \emptyset$:

$$p(x,y) = \sum_{z} p(x)p(z|x)p(y|z)$$
$$= p(x)\sum_{z} p(z|x)p(y|z)$$
$$= p(x)p(y|x)$$
$$\neq p(x)p(y)$$





Directed graphical models

Graphs with three nodes

- A Markov chain is a special DAG.
- We verify that $X \perp Y \mid Z$:

$$p(y|z,x) = \frac{p(x,y,z)}{p(x,z)}$$

$$= \frac{p(x,y,z)}{\sum_{y'} p(y',x,z)}$$

$$= \frac{p(x)p(z|x)p(y|z)}{\sum_{y'} p(x)p(z|x)p(y'|z)}$$

$$= p(y|z)$$





Directed graphical models

Graphs with three nodes

Z is a latent cause.

• We verify that $X \not\perp Y \mid \emptyset$:

$$\begin{array}{lcl} p(x,y) & = & \displaystyle \sum_{z} p(z) p(x|z) p(y|z) \\ & \neq & \displaystyle p(x) p(y) \end{array}$$





Directed graphical models

Graphs with three nodes

Z is a latent cause.

• We verify that $X \perp Y \mid Z$:

$$p(x, y|z) = \frac{p(x, y, z)}{p(z)}$$
$$= \frac{p(z)p(y|z)p(x|z)}{p(z)}$$
$$= p(x|z)p(y|z)$$





Directed graphical models

Graphs with three nodes

Explaining away or V-structure.

• We verify that $X \perp Y \mid \emptyset$:

$$p(x,y) = \sum_{z} p(x,y,z)$$
$$= p(x)p(y)\sum_{z} p(z|x,y)$$
$$= p(x)p(y)$$





Directed graphical models

Graphs with three nodes

Explaining away or V-structure.

• We verify that $X \not\perp Y \mid Z$:

$$p(x, y|z) = \frac{p(x, y, z)}{p(z)}$$
$$= \frac{p(x)p(y)p(z|x, y)}{p(z)}$$
$$\neq p(x|z)p(y|z)$$





Directed graphical models

Hidden Markov chain

- A Hidden Markov Model is a dynamic Bayesian network. Often used because we only have a noisy observation of the random process.
- Y_t are the visible variables, and X_t the hidden variables.
- We have:

$$\begin{array}{cccc} X_{t+1} & \bot & X_{t-1} \mid X_t \\ Y_{t+1} & \bot & Y_t \mid X_t \end{array}$$





Probabilistic	Graphical	Models
---------------	-----------	--------

L Illustration



Independence Models

Conditional independenceGraphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

3 Inference and MAP Estimation

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- 4 Parameter & Structure Learning
- 5 Hidden Markov Models



- Illustration

Various inferences on Asia



L Illustration

Toy problem 1

Consider a bag containing the following tokens:



- Show that Value ⊥ Form | Color.
- Build all the faithful DAGs of p(V, C, F).
- Learn the parameters.
- Compute P(V = 1 | F =square) using the model.



L Illustration

Toy problem 2

- Consider a car with Battery (0=flat, 1=fully charged), Fuel tank (0=empty, 1=full) and Fuel Gauge reading (0=empty, 1=full).
- Assume that $B \perp F \mid \emptyset$.



- $P(G = 1 \mid B = 1, F = 1) = 0.8$
- $P(G = 1 \mid B = 1, F = 0) = 0.2$
- $P(G=1 \mid B=0, F=1) = -0.2$
- $P(G = 1 \mid B = 0, F = 0) = 0.1$
 - P(B=1) = 0.9
 - P(F=1) = 0.9

Université Claude Bernard (G) Lyon 1

Compute:

- $p(F = 0 \mid G = 0)$
- $p(F = 0 \mid G = 0, B = 0)$

・ コット (雪) (小田) (コット 日)

PGM's expressiveness

Outline

Independence Models

Conditional independenceGraphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

Inference and MAP Estimation

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- 4 Parameter & Structure Learning
- 5 Hidden Markov Models


PGMs

PGM's expressiveness

Bayesian curve fitting and prediction

Predictive distribution: $p(\hat{t}|\hat{x}, \mathbf{x}, \mathbf{t}, \alpha, \sigma^2) \propto \int p(\hat{t}, \mathbf{t}, \mathbf{w}|\hat{x}, \mathbf{x}, \alpha, \sigma^2) \, \mathrm{d}\mathbf{w}$



(日)

PGMs

PGM's expressiveness

Relation between DAG and UG





PGMs

PGM's expressiveness

Relation between DAG and UG

$$p(x) = p(x_1)p(x_2 \mid x_1)p(x_3 \mid x_2) \dots p(x_N \mid x_{N-1})$$

= $\frac{1}{Z}\psi_{1,2}(x_1, x_2)\psi_{2,3}(x_2, x_3) \dots \psi_{N-1,N}(x_{N-1}, x_N).$

This is easily done by identifying,

$$\psi_{1,2}(x_1, x_2) = p(x_1)p(x_2 \mid x_1)$$

$$\psi_{2,3}(x_2, x_3) = p(x_3 \mid x_2)$$

$$\vdots$$

$$\psi_{N-1,N}(x_{N-1}, x_N) = p(x_N \mid x_{N-1})$$

The maximal cliques in the UG are the pairs of neighbouring nodes in the DAG. In this case, Z = 1.
Université Claude Bernard (Guide Bernard (Guide

イロト イロト イヨト

э

PGMs

PGM's expressiveness

Relation between DAG and UG



The process of 'marrying the parents' is known as moralization, and the resulting undirected graph, after dropping the arrows, is called the "moral graph".

A D > A P > A D > A D >

 Graph moralization plays an important role in exact inference techniques such as the junction tree algorithm. PGMs

PGM's expressiveness

Relation between DAG and UG

$$p(x) = p(x_1)p(x_2)p(x_3)p(x_4 \mid x_1, x_2, x_3)$$
$$= \frac{1}{Z}\psi_{1,2,3,4}(x_1, x_2, x_3, x_4).$$

- In going from a directed to an undirected representation we had to discard some CI properties from the graph (e.g. X₁ ⊥ X₂ | X₄).
- It turns out that the two types of graph can express different CI properties.



PGMs

PGM's expressiveness

PGM's expressiveness





PGMs

PGM's expressiveness



\square p(a, b, c) for the noisy XOR (exclusive OR) relationship

$$P(A = B \oplus C) = 1 - \epsilon$$

 $\label{eq:posterior} {\scriptstyle \blacksquare} \ p>0 \ \text{for any} \ \epsilon\in]0,1/2[\cup]1/2,0[.$



PGMs

PGM's expressiveness



- We have $A \perp B$, $B \perp C$ and $C \perp A$.
- Due to the strong union property, $A \perp B \implies A \perp B \mid C$, no undirected graph that can encode any of the independence relations in p
- *p* is not UG-faithful. The complete graph is the I-map.
- The Markov network model requires 7 free parameters to encode *p*.



PGMs

PGM's expressiveness



- *p* is not DAG-faithful either. Due to the composition property, any DAG that encodes two of the independence relations in *p* necessarily breaks a dependence relation as well $(A \perp B \land A \perp C \implies A \perp \{B, C\})$.
- The DAG $A \to C \leftarrow B$ encodes only one of the independence relation. This BN structure results in the factorization p(a,b,c) = p(a)p(b)p(c|a,b), which encodes p with 6 free parameters.
- In this example p is neither UG-faithful nor DAG-faithful, so both Markov networks and Bayesian networks are not well-suited models to encode p efficiently.
- Yet, *p* can be encoded efficiently with only 4 parameters.



PGMs

PGM's expressiveness

Extensions

- Classical PGMs have a limited expressive power as independence models.
- Over the years, many alternative PGMs have been proposed to overcome these limitations, by extending and unifying UGs and DAGs.
 - Ancestral graphs, Anterial graphs, LWF chain graphs, AMP chain graphs...
 - Four types of edges are allowed: directed edges and three types of undirected edges.
- Increased expressive power comes at the expense of an increased complexity.
- Factorization of p? Practical parametrization of the model? Learning and inference?



Outline

Independence Models

Conditional independenceGraphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- 3 Inference and MAP Estimation
 - Inference in a chain
 - Sum-product algorithm
 - Max-sum algorithm
 - Parameter & Structure Learning
 - 5 Hidden Markov Models



The 4 basic problems with PGMs

There are 4 basic problems to be solved for the model to be useful in real-world applications:

- 1 Problem 1: **Inference**. Given some observation, compute the conditional distribution of the remaining variables (NP-hard if loops in the graph).
- 2 Problem 2: MAP Inference. Find the MAP over this conditional distribution.
- 3 Problem 3: Learning. Given a sequence of observations, estimate the MAP of the parameters (Easy problem with a complete data set).
- 4 Problem 4: Learning. Given a sequence of observations, learn the topological structure of the PGM (NP-hard).



Problem 1: Inference

- Suppose we have a set of correlated random variables with joint distribution $p(x_1, \ldots, x_N | \theta)$.
- Let us partition this vector into the visible variables X_v, which are observed, and the hidden variables, X_h, which are unobserved.
- Inference refers to computing the posterior distribution of the unknowns given the evidence:

$$p(\mathbf{x}_h | \mathbf{x}_v, \theta) = \frac{p(\mathbf{x}_h, \mathbf{x}_v | \theta)}{p(\mathbf{x}_v | \theta)} = \frac{p(\mathbf{x}_h, \mathbf{x}_v | \theta)}{\sum_{\mathbf{x}_h'} p(\mathbf{x}_h', \mathbf{x}_v | \theta)}$$



Problem 1: Inference

- Sometimes only some of the hidden variables are of interest.
- Let's partition the hidden variables into **query variables**, *X_q*, whose value we wish to know, and the remaining **nuisance variables**, *X_n*, which we are not interested in.
- We can compute what we are interested in by marginalizing out the nuisance variables:

$$p(\mathbf{x}_q | \mathbf{x}_v, \theta) = \sum_{\mathbf{x}_n} p(\mathbf{x}_q, \mathbf{x}_n | \mathbf{x}_v, \theta)$$



Inference in a chain



Independence Models

- Conditional independence
- Graphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- 3 Inference and MAP Estimation
 - Inference in a chain
 - Sum-product algorithm
 - Max-sum algorithm
 - 4 Parameter & Structure Learning
 - 5 Hidden Markov Models



Inference in a chain

Problem 1: Inference in a chain



Exact inference on a graph comprising a chain of nodes can be performed efficiently in time that is linear in the number of nodes.

$$p(x) = \frac{1}{Z}\psi_{1,2}(x_1, x_2)\psi_{2,3}(x_2, x_3)\dots\psi_{N-1,N}(x_{N-1}, x_N).$$

The algorithm that can be interpreted in terms of messages passed along the chain.



Inference in a chain

Problem 1: Inference in a chain

Consider the inference problem of finding the marginal distribution $p(x_n)$

$$p(x_n) = \sum_{x_1} \dots \sum_{x_{n-1}} \sum_{x_{n+1}} \dots \sum_{x_N} p(\mathbf{x})$$

= $\frac{1}{Z} \left[\sum_{x_{n-1}} \psi_{n-1,n}(x_{n-1}, x_n) \left[\dots \left[\sum_{x_1} \psi_{1,2}(x_1, x_2) \right] \right] \dots \right]$
 $\times \left[\sum_{x_{n+1}} \psi_{n,n+1}(x_n, x_{n+1}) \left[\dots \left[\sum_{x_N} \psi_{N-1,N}(x_{N-1}, x_N) \right] \right] \dots \right]$
= $\frac{1}{Z} \mu_{\alpha}(x_n) \mu_{\beta}(x_n).$



Inference in a chain

Problem 1: Inference in a chain

The algorithm that can be interpreted in terms of messages passed along the chain.



- With *N* discrete variables each having *K* states, the messages $\mu_{\alpha}(x_n)$ and $\mu_{\beta}(x_n)$ can be evaluated recursively in $O(NK^2)$. by exploiting the IC properties of this simple graph in order to obtain an efficient calculation.
- This is linear in the length of the chain, in contrast to the exponential cost of a naive approach.



Inference and MAP Estimation

Inference in a chain

Exact inference in a (poly)tree

- As for chains, inference can be performed efficiently using local message passing in trees and polytrees.
- A polytree is a directed acyclic graph whose underlying undirected graph is a tree.
- The message passing formalism is also applicable to undirected and directed trees and to polytrees. It is called the sum-product algorithm.
- It requires a graphical construction called a factor graph.



Inference and MAP Estimation

Inference in a chain



First transform the PGM into a factor graph:



L Sum-product algorithm

Outline

Independence Models

- Conditional independence
- Graphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

3 Inference and MAP Estimation

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- Parameter & Structure Learning
- 5 Hidden Markov Models



Inference and MAP Estimation

Sum-product algorithm

Sum-product algorithm

Start form the leaves:







L Sum-product algorithm

Sum-product algorithm

Consider a simple example to illustrate the operation of the sum-product algorithm:



Université Claude Bernard

э

A D > A P > A D > A D >

L Sum-product algorithm

Sum-product algorithm

Say node x₃ is the root node. Start from the leaf nodes x₁ and x₄ towards the root x₃ and perform the following sequence of messages:

$$\mu_{x_1 \to f_a}(x_1) = 1$$

$$\mu_{f_a \to x_2}(x_2) = \sum_{x_1} f_a(x_1, x_2)$$

$$\mu_{x_4 \to f_c}(x_4) = 1$$

$$\mu_{f_c \to x_2}(x_2) = \sum_{x_4} f_c(x_2, x_4)$$

$$\mu_{x_2 \to f_b}(x_2) = \mu_{f_a \to x_2}(x_2) \mu_{f_c \to x_2}(x_2)$$

$$\mu_{f_b \to x_3}(x_3) = \sum_{x_2} f_b(x_2, x_3) \mu_{x_2 \to f_b}(x_2)$$

Université Claude Bernard

・ロット (雪) (日) (日)

L Sum-product algorithm

Sum-product algorithm

Then, from the root node towards the leaf nodes:

$$\begin{aligned} \mu_{x_3 \to f_b}(x_3) &= 1 \\ \mu_{f_b \to x_2}(x_2) &= \sum_{x_3} f_b(x_2, x_3) \\ \mu_{x_2 \to f_a}(x_2) &= \mu_{f_b \to x_2}(x_2) \mu_{f_c \to x_2}(x_2) \\ \mu_{f_a \to x_1}(x_1) &= \sum_{x_2} f_a(x_1, x_2) \mu_{x_2 \to f_a}(x_2) \\ \mu_{x_2 \to f_c}(x_2) &= \mu_{f_a \to x_2}(x_2) \mu_{f_b \to x_2}(x_2) \\ \mu_{f_c \to x_4}(x_4) &= \sum_{x_2} f_c(x_2, x_4) \mu_{x_2 \to f_c}(x_2) \end{aligned}$$



L Sum-product algorithm

Sum-product algorithm

- One message has now passed in each direction across each link,
- To evaluate the marginals:

$$p(x_2) = \frac{1}{Z} \mu_{f_a \to x_2}(x_2) \mu_{f_b \to x_2}(x_2) \mu_{f_c \to x_2}(x_2)$$

= $\frac{1}{Z} \left[\sum_{x_1} f_a(x_1, x_2) \right] \left[\sum_{x_3} f_b(x_2, x_3) \right] \left[\sum_{x_4} f_c(x_2, x_4) \right]$
= $\frac{1}{Z} \sum_{x_1} \sum_{x_3} \sum_{x_4} f_a(x_1, x_2) f_b(x_2, x_3) f_c(x_2, x_4)$
= $\frac{1}{Z} \sum_{x_1} \sum_{x_3} \sum_{x_4} p(\mathbf{x})$



Max-sum algorithm



Independence Models

- Conditional independence
- Graphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

3 Inference and MAP Estimation

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- Parameter & Structure Learning
- 5 Hidden Markov Models



Max-sum algorithm

Problem 2: MAP Inference

- Suppose we have a set of correlated random variables with joint distribution $p(x_1, \ldots, x_N | \theta)$.
- Let us partition this vector into the visible variables X_v, which are observed, and the hidden variables, X_h, which are unobserved.
- MAP Inference refers to computing the MAP of the posterior distribution:

$$\mathbf{x}_{h}^{\star} = \arg \max_{\mathbf{x}_{h}} p(\mathbf{x}_{h} | \mathbf{x}_{v}, \theta))$$



Max-sum algorithm

Problem 2: The max-sum algorithm

- The sum-product algorithm takes a joint distribution p(x) expressed as a factor graph and efficiently find marginals over the component variables.
- MAP inference: find a setting of the variables that has the largest probability and give the probability.
- This can be addressed through a closely related algorithm called max-sum algorithm, which can be viewed as an application of dynamic programming in the context of graphical models



Max-sum algorithm

Problem 2: MAP inference in a chain



 MAP inference on a graph comprising a chain of nodes can be performed efficiently in time that is linear in the number of nodes.

$$p(x) = \frac{1}{Z} \max_{x_1,\dots,x_N} \left[\psi_{1,2}(x_1, x_2) \psi_{2,3}(x_2, x_3) \dots \psi_{N-1,N}(x_{N-1}, x_N) \right]$$
$$= \frac{1}{Z} \max_{x_1} \left[\max_{x_2} \psi_{1,2}(x_1, x_2) \left[\dots \max_{x_N} \psi_{N-1,N}(x_{N-1}, x_N) \right] \right].$$

- The structure of this calculation is identical to that of the sum-product algorithm,
- Application: find the most probable sequence of hidden states in a HMM, known as the Viterbi algorithm.

・ロット (雪) (日) (日)

Max-sum algorithm

Exact inference in general graphs

- The message passing framework (i.e. sum-product and max-sum algorithms) can be generalized to graphs having loops, using the junction tree algorithm (Lauritzen et al., 1988).
- A DAG is first converted to an UG by moralization, (not required for an UG).
- Next the graph is triangulated, i.e. adding extra links to eliminate chord-less cycles containing four or more nodes.



Max-sum algorithm

The junction tree algorithm

- Then, construct a tree-structured undirected graph called a join tree, whose nodes correspond to the maximal cliques of the triangulated graph, and whose links connect pairs of cliques that have variables in common.
- The selection of which pairs of cliques to connect in this way is important and is done so as to give a maximal spanning tree.
- If the number of variables in the largest clique is high, the junction tree algorithm becomes impractical.



Max-sum algorithm

Approximate inference in general graphs

- For many problems of practical interest, it is not be feasible to use exact inference, effective approximation methods are needed.
- A simple idea to approximate inference in graphs with loops is to apply the sum-product algorithm as it is.
- This approach is known as loopy belief propagation (Frey and MacKay, 1998) and is possible because the message passing rules are purely local, even though there is no guarantee that it will yield good results.



Parameter & Structure Learning

Outline

Independence Models

- Conditional independence
- Graphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

- 3 Inference and MAP Estimation
 - Inference in a chain
 - Sum-product algorithm
 - Max-sum algorithm
- 4 Parameter & Structure Learning
 - Hidden Markov Models



Parameter & Structure Learning

The 4 basic problems with PGMs

There are 4 basic problems to be solved for the model to be useful in real-world applications:

- 1 Problem 1: **Inference**. Given some observation, compute the conditional distribution of the remaining variables (NP-hard if loops in the graph).
- 2 Problem 2: MAP Inference. Find the MAP over this conditional distribution (NP-hard).
- 3 Problem 3: **Learning**. Given a sequence of observations, estimate the MAP of the **parameters** (Easy problem with a complete data set).
- 4 Problem 4: Learning. Given a sequence of observations, learn the Markov boundary of some variable or the complete topological structure of the PGM (NP-hard).



Parameter & Structure Learning

Problem 3: Parameter learning

Find the MAP estimate for the parameters:

$$\hat{\theta} = \arg \max_{\theta} \sum_{i=1}^{N} \log p(\mathbf{x}_i | \theta) + \log p(\theta)$$

• $p(\theta)$ is the prior on the parameters.


Problem 3: Learning from complete data

- If all the variables are fully observed (i.e. no missing data and no hidden variables), the data is complete.
- For a DGM with complete data, the likelihood is given by

$$p(\mathbf{x}|\theta) = \prod_{i=1}^{N} p(x_i|\theta)$$
$$= \prod_{i=1}^{N} \prod_{t=1}^{V} p(x_{it}|x_{i,\text{pa}(t)}, \theta_t)$$
$$= \prod_{t=1}^{V} p(\mathcal{D}_t|\theta_t)$$

 \square \mathcal{D}_t is the data associated with node t and its parents.

Problem 3: Learning from complete data

Now suppose that the prior factorizes as well:

$$p(\theta) = \prod_{t=1}^{V} p(\theta_t)$$

Then clearly the posterior also factorizes:

$$p(\theta|\mathcal{D}) \propto p(\mathcal{D}|\theta)p(\theta) = \prod_{t=1}^{V} p(\mathcal{D}_t|\theta_t)p(\theta_t)$$



Probabilistic Graphical Models

Parameter & Structure Learning

Problem 3: Learning with missing and/or latent variables

- If we have missing data and/or hidden variables, the likelihood no longer factorizes, and indeed it is no longer convex.
- This means we will usually can only compute a locally optimal ML or MAP estimate.
- Bayesian inference of the parameters is even harder and requires suitable approximate inference techniques.



Problem 4: structure learning

- Given a sequence of observations, learn the topological structure of the PGM (NP-hard). The problem of learning a BN structure has attracted much attention.
- Problem: the number of possible DAGs with n variables is *superexponential* w.r.t n. For instance, NS(5) = 29281 and NS(10) = 4.2×10^{18} .
- Search-and-score methods search over a space of structures employing a scoring function to guide the search. The most prominent algorithm in this class is the Greedy Equivalent Search (GES).
- Constraint-based algorithms use statistical independence tests to impose constraints on the network structure and infer the final DAG. PC is prototypical constraint-based algorithm.



Learning the Markov Boundary

```
Algorithm 1 IAMB
Require: T: target; D = data set
Ensure: MB : Markov boundary of T
    Phase I: Add true positives to MB
 1: MB = \emptyset
 2: repeat
 3:
       [assoc, Y] = \max_{X \in (U \setminus MB \setminus T)} dep(T, X | MB)
   if assoc \neq 0 then
 4:
 5:
         MB = MB \cup Y
      end if
 6:
 7: until MB has not changed
    Phase II: Remove false positives from MB
 8: for all X \in MB do
 9:
      if dep(T, X | MB \setminus X) = 0 then
         MB = MB \setminus X
10:
    end if
11:
12: end for
```

Université Claude Bernard (GB) Lyon 1

・ コット (雪) ・ (目) ・ (目)

Aliferis et al. 2002

Learning the Markov Boundary



・ロト ・ 同ト ・ ヨト ・ ヨト

э

Constraint-based structure learning

Input : \mathcal{D} : dataset; **U**: set of random variables. S_{XY} ensemble de variables d-sépatant X et Y. **Output** : Graph pattern (*gp*) such that $X \perp_{\mathcal{G}} Y | \mathbf{Z} \Rightarrow X \perp_{P} Y | \mathbf{Z}.$

1. Step 1:

- Forall [pair of nodes $X, Y \in \mathbf{U}$]
- Search for a subset $S_{XY} \subseteq \mathbf{U}$ such that $X \perp_P Y | S_{XY}$;
- If [no such set can be found]
- ► Create the link X Y in gp;

2. Step 2:

- ► Forall [uncoupled meeting *X* − *Z* − *Y*]
- If $[Z \notin S_{XY}]$
- Orient X Z Y as $X \to Z \leftarrow Y$;



Skeleton further orientation

 $\label{eq:input} \textbf{Input}: An essential graph G (partially directed acyclic graph pattern)$

 \mathbf{Output} : DAG pattern (gp) such that

$$X\perp_{\mathcal{G}} Y|\mathbf{Z}\Rightarrow X\perp_{P} Y|\mathbf{Z}.$$

- 1. Step 1:
 - ▶ Forall [link *X* − *Y* such that there is a path from *X* to *Y*]
 - Orient X Y as $X \to Y$;
- 2. Step 2:
 - ▶ **Forall** [uncoupled meeting X Z Y such that $X \to W$, $Y \to W$ and Z W]

人口 医水黄 医水黄 医水黄素 计目录

• Orient Z - W as $Z \rightarrow W$;

PC algorithm on ASIA



 $\chi^2: \ \textit{S} \bot \textit{A} \ \textit{L} \bot \textit{A} \ \textit{B} \bot \textit{A} \ \textit{O} \bot \textit{A} \ \textit{X} \bot \textit{A} \ \textit{D} \bot \textit{A} \ \textit{T} \bot \textit{S} \ \textit{L} \bot \textit{T} \ \textit{O} \bot \textit{B} \ \textit{X} \bot \textit{B}$



PC algorithm on ASIA



 $\chi^2: D \perp S | \{L, B\} \quad X \perp O | \{T, L\} \quad D \perp O | \{T, L\}$



Search for V-structures and further orientation





Probabilistic Graphical Models

Parameter & Structure Learning

ALARM



Probabilistic Graphical Models

Parameter & Structure Learning

ALARM



・ロト ・ 理 ト ・ 理 ト ・ 理 ト

Gene networks





æ

Epidemiological studies



▲□▶ ▲□▶ ▲三▶ ▲三▶ 三三 のへで

A causal graph: Hip fracture risk factor analysis



Outline

Independence Models

- Conditional independenceGraphoids

2 PGMs

- Undirected graphical models
- Directed graphical models
- Illustration

PGM's expressiveness

3 Inference and MAP Estimation

- Inference in a chain
- Sum-product algorithm
- Max-sum algorithm
- Parameter & Structure Learning
- 5 Hidden Markov Models



Markov Models

- A Markov model is a stochastic model used to model randomly changing systems where it is assumed that future states depend only on the current state not on the events that occurred before it.
- This assumption is called the:

Markov property

$$P(X_{n+1} = j | X_1 = i_1, \dots, X_n = i_n) = P(X_{n+1} = j | X_n = i_n)$$



Markov chains

- The Markov property enables reasoning and computation with the model that would otherwise be intractable.
- If $P(X_{n+1} = j | X_n = i_n) = p_{ij}$ does not depend on *n* then the Markov model is **homogeneous**.
- The simplest Markov model is the Markov chain. It models the state of a system with a random variable that changes through time.



Markov chains

Markov chain



Transition matrix

$$P = \begin{pmatrix} 0.3 & 0.2 & 0.5 & 0\\ 0.1 & 0.2 & 0.3 & 0.4\\ 0 & 0.8 & 0.2 & 0\\ 0.2 & 0.1 & 0.3 & 0.4 \end{pmatrix}$$

Such that
$$\forall i, \sum_{j} P_{ij} = 1$$



Coin toss Models

- Someone is performing coin tosses in a room. He tells you the result of the coin flips, nothing else (e.g. probability of heads, number of coins, transition probabilities).
- We only observe a sequence of heads (H) and tails (T).
- Which model (1 or 2 coins) best matches the observations?

Two scenarios:

Single coin





(日)

 $P(H) = p_2$

 $P(H) = p_1$

Hidden Markov models

- A hidden Markov model (HMM) is a Markov model in which the system is assumed to be a Markov process with unobserved (hidden) states.
- In simpler Markov models (like a Markov chain), the state is directly visible to the observer, and therefore the state transition probabilities are the only parameters.
- In a HMM, the output, dependent on the state, is visible. Each state has a probability distribution over the possible **outputs**. Therefore, the sequence of outputs generated by an HMM gives some information about the sequence of state.
- Many applications in temporal pattern recognition such as speech, handwriting, gesture recognition, and bioinformatics.



Hidden Markov models

- The random variable *q_t* is the hidden state at time *t*. which is assumed to consist of one of N possible values {*s*₁,...,*s_n*}, modeled as a categorical distribution.
- The random variable O_t is the observation at time t (with $y(t) \in \{y_1, y_2, y_3, y_4\}$). O_t is typically a letter from an alphabet of M symbols $V = \{v_1, \ldots, v_M\}$.
- In the standard HMM, the state space is discrete, while the observations themselves can either be discrete or continuous (e.g. Gaussian distribution).
- The parameters of a hidden Markov model are of two types: N² transition probabilities and NM emission probabilities (also known as output probabilities).



Example: Urn and ball model





Discrete symbol HMM



Notations

A complete specification of an HMM is given its paremeters $\Lambda = (A,B,\pi)$ and is defined by:

- Its *n* hidden states $S = \{s_1, \ldots, s_n\}$.
- The *M* observable symbols $V = \{v_1, \ldots, v_M\}$. O_t denotes the symbol at time *t*.
- The state transition matrix $a_{ij} = A(i, j)$
- The observation symbol probability distribution *B*: $b_j(k) = P(O_t = v_k | q_t = s_j)$ with $\sum_{k=1}^{M} b_j(k) = 1$
- The initial state distribution $\pi = {\pi_j}_{j=1,...,n}$ where $\pi_j = P(q_1 = s_j)$ and $\sum_{j=1}^n \pi_j = 1$.



The 3 basic problems with HMMs

There are 3 basic problems to be solved for the model to be useful in real-world applications:

- **1** Problem 1: **Evaluation**. Compute of the probability $P(O|\Lambda)$ of the observation sequence $\{O_1, \ldots, O_T\}$ given an HMM $\Lambda = (A, B, \pi)$.
- **2** Problem 2: **Inference**. Given a sequence $\{O_1, \ldots, O_T\}$ and the model Λ , chose a state sequence $Q = q_1, \ldots, q_T$ which is meaningful (i.e. that best explains the observations) in some sense ? Several optimality criteria to be imposed.
- 3 Problem 3: **Training**. Given a sequence $\{O_1, \ldots, O_T\}$, how do we adjust the model $\Lambda = (A, B, \pi)$ to maximize $P(\mathcal{O}|\Lambda)$?



Direct evaluation of $P(O|\Lambda)$

The most straightforward way to compute of $P(O|\Lambda)$ is through enumerating all every state sequence q_1, \ldots, q_T :

$$\begin{split} P(O|\Lambda) &= \sum_{Q} P(O,Q|\Lambda) = \sum_{Q} P(O|Q,\Lambda) P(Q|\Lambda) \\ &= \sum_{Q} P(q_1|\Lambda) \prod_{t=1}^{T} P(O_t|q_t,\Lambda) \prod_{t=2}^{T} P(q_t|q_{t-1},\Lambda) \\ &= \sum_{Q} \pi_{q_1} \prod_{t=2}^{T} b_{q_t}(O_t) a_{q_{t-1},q_t} \end{split}$$

■ The calculation of *P*(*O*|Λ) involves *O*(2*T* · *n*^{*T*}) calculations. A more efficient procedure is needed.



Forward Approach

• Let
$$\alpha_t(i) = P(O_1, \dots, O_t, q_t = s_i | \Lambda) = P(O_1^t, q_t = s_i | \Lambda)$$

$$\begin{aligned} \alpha_1(i) &= P(O_1, q_1 = s_i | \Lambda) = \pi_i b_i(O_1) \\ \alpha_{t+1}(j) &= P(O_1, \dots, O_t, O_{t+1}, q_{t+1} = s_j | \Lambda) \\ &= \sum_{i=1}^n P(O_1^t, O_{t+1}, q_t = s_i, q_{t+1} = s_j | \Lambda) \\ &= \sum_{i=1}^n P(O_{t+1} | q_{t+1} = s_i, \Lambda) P(O_1^t, q_t = s_i | \Lambda) a_{ij} \\ &= \sum_{i=1}^n \alpha_t(i) a_{ij}] b_j(O_{t+1}) \end{aligned}$$

• Finally : $P(O|\Lambda) = \sum_{i=1}^{n} \alpha_T(i)$



Backward Approach

Likewise, let
$$\beta_t(i) = P(O_{t+1}, \dots, O_T | q_t = s_i, \Lambda)$$

$$\beta_T(i) = 1$$

As previously

$$\beta_t(i) = \sum_{j=1}^n a_{ij} b_j(O_{t+1}) \beta_{t+1}(j)$$

- Finally : $P(O|\Lambda) = \sum_{i=1}^{n} \pi_i b_i(O_1) \beta_1(i)$.
- In both cases, the complexity is $O(n^2T)$



Which state q_t is the most likely?

There are several ways of finding the **optimal state sequence**.

• Which state q_t is the most likely?

• Let
$$\gamma_t(i) = P(q_t = s_i | O_1^T)$$

$$\begin{aligned} \gamma_t(i) &= P(q_t = s_i | O_1^T) \\ &= P(q_t = s_i | O_1, \dots, O_t, O_{t+1}, \dots, O_T) \\ &= \frac{P(O_1, \dots, O_t, q_t = s_i | \Lambda) P(O_{t+1}, \dots, O_T | q_t = s_i, \Lambda)}{P(O_1^T | \Lambda)} \\ &= \frac{\alpha_t(i) \beta_t(i)}{\sum_{j=1}^n \alpha_t(j) \beta_t(j)} \end{aligned}$$

Then we solve $q_t = argmax_i[\gamma_t(i)]$

One problem is that the state sequence may not even be valid, for instance if state transitions have zero probability.
Université Claude Bernard Output

・ロット (雪) (日) (日)

Optimal state sequence: Viterbi algorithm

The single best path sequence is given by $\max_Q P(O, Q|\Lambda)$. Define

$$\delta_t(i) = \max_{q_1,\ldots,q_T} P(q_1,\ldots,q_{t-1},q_t=s_i,O_1,\ldots,O_T|\Lambda)$$

By induction (dynamic programming), we have the recursion:

$$\delta_{t+1}(j) = [\max_i \delta_t(i)a_{ij}]b_j(O_{t+1})$$

Hence the complete recursive procedure,

1
$$\delta_1(i) = \pi_i b_i(O_1)$$

2 $\delta_{t+1}(j) = [\max_i \delta_t(i) a_{ij}] b_j(O_{t+1})$
3 $\max_Q P(O, Q|\Lambda) = \max_i \delta_T(i)$



Training

- The most difficult problem of HMMs is to adjust the model parameters to maximize the likelihood.
- Suppose we have the sequence $\mathcal{O} = \{O^1, \dots, O^n\}$, the goal is to find the parameters $\Lambda = (A, B, \pi)$ such that $P(\mathcal{O}|\Lambda) = \prod_{k=1}^n P(O^k|\Lambda)$ is locally maximum using gradient or EM techniques.
- We compute Λ_{k+1} from Λ_k such that $P(\mathcal{O}|\Lambda_{k+1}) \ge P(\mathcal{O}|\Lambda_k)$.
- Eventually, the likelihood function converges to a critical point.
- We define next the Baum-Welch iterative procedure for choosing model parameters.



Baum-Welch Algorithm

Let
$$\xi_t^k(i, j) = P(q_t = s_i, q_{t+1} = s_j | O^k, \Lambda)$$

$$= \frac{P(q_t = s_i, q_{t+1} = s_j, O^k | \Lambda)}{P(O^k | \Lambda)}$$

$$= \frac{\alpha_t^k(i) a_{ij} b_j(O_{t+1}^k) \beta_{t+1}^k(j)}{P(O^k | \Lambda)}$$

We have
$$\gamma_t(i) = P(q_t = s_i | O_1^T)$$

$$= \sum_{j=1}^n P(q_t = s_i, q_{t+1} = s_j | O^T, \Lambda)$$

$$= \sum_{j=1}^n \xi_t^k(i, j) = \frac{\alpha_t(i)\beta_t(i)}{\sum_{j=1}^n \alpha_t(j)\beta_t(j)}$$
Université Claude Bernarcie de la 2000

Baum-Welch Algorithm

The parameters of new HMM model Λ_{p+1} are re-estimated from the previous one Λ_p :

$$\bar{a}_{ij} = \frac{\sum_{k=1}^{m} \sum_{t} \xi_{t}^{k}(i,j)}{\sum_{k=1}^{m} \sum_{t} \gamma_{t}^{k}(i)}$$
$$\bar{b}_{j}(l) = \frac{\sum_{k=1}^{m} \sum_{t/O_{t}^{k}=v_{l}\}} \gamma_{t}^{k}(j)}{\sum_{k=1}^{m} \sum_{t} \gamma_{t}^{k}(i)}$$
$$\bar{\pi}_{i} = \frac{1}{m} \sum_{k=1}^{m} \gamma_{1}^{k}(i)$$



Baum-Welch Algorithm

1 Given
$$\Lambda_0 = (A, B, \pi)$$
 et $p = 0$

- **2** Do: Compute $\xi_t^k(i, j)$ with $\gamma_1^k(i), \forall 1 \le i, j \le n$ with $1 \le t \le T 1$ and Λ_p
- 3 Estimate $\bar{a}_{ij}, \bar{b}_j(l), \bar{\pi}_i$
- 4 Let $\Lambda_p = (\bar{A}, \bar{B}, \bar{\pi})$
- 5 p ← p+1
- 6 Until convergence


Hidden Markov Models

Extensions

- HMMs are generative models: they model a joint distribution of observations and hidden states.
- A discriminative model can be used in place of the generative model of standard HMMs. This type of model directly models the conditional distribution of the hidden states given the observations X.
- HMM can also be generalized to allow continuous observations and/or state spaces (typically Gaussian), however, in general, exact inference in HMMs with continuous latent variables is infeasible.
- A uniform prior distribution over the transition probabilities was implicitly assumed. Another prior candidate is the Dirichlet distribution.



Hidden Markov Models

References I



D. Koller and N. Friedman.

Probabilistic Graphical Models: Principles and Techniques. MIT Press, 2009.



Christopher M. Bishop.

Pattern Recognition and Machine Learning. Springer, 2006.



G. Bontempi and S. Ben Taieb. Statistical foundations of machine learning

OTexts: Melbourne, Australia, 2016.



Maxime Gasse.

Probabilistic Graphical Model Structure Learning: Application to Multi-Label Classification.

Ph.D. thesis, Université de Lyon, 2017.



Hidden Markov Models

References II

Lawrence R. Rabiner.

A tutorial on Hidden Markov Models and selected applications in speech recognition, Proceedings of the IEEE, vol. 77, no 2, 1989.

G. Obozinski and F. Bach

Introduction to Graphical Models, Master "Mathématiques Appliquées", Parcours "Mathematiques, Vision et Apprentissage", ENS Cachan 2017.

