Probabilistic Graphical Models

Cyrill Stachniss

Wolfram Burgard

Henrik Kretzschmar

Probabilistic Graphical Models

- Marriage between probability theory and graph theory
- Tool for dealing with uncertainty, independence, and complexity
- Notion of modularity – a complex system that consists of simpler parts
- Probability theory is the “glue” for the individual parts
- Play an increasingly important role in robotics, vision, and machine learning

Important Questions

- **Representation:** How can a graphical model compactly represent a joint probability distribution?
- **Inference:** How to efficiently infer states given observations?
- **Learning:** How to estimate the parameters (and structure) of the model?
- **Applications:** What has this machinery been used for?

A Typical Example for a PGM – Bayes Network

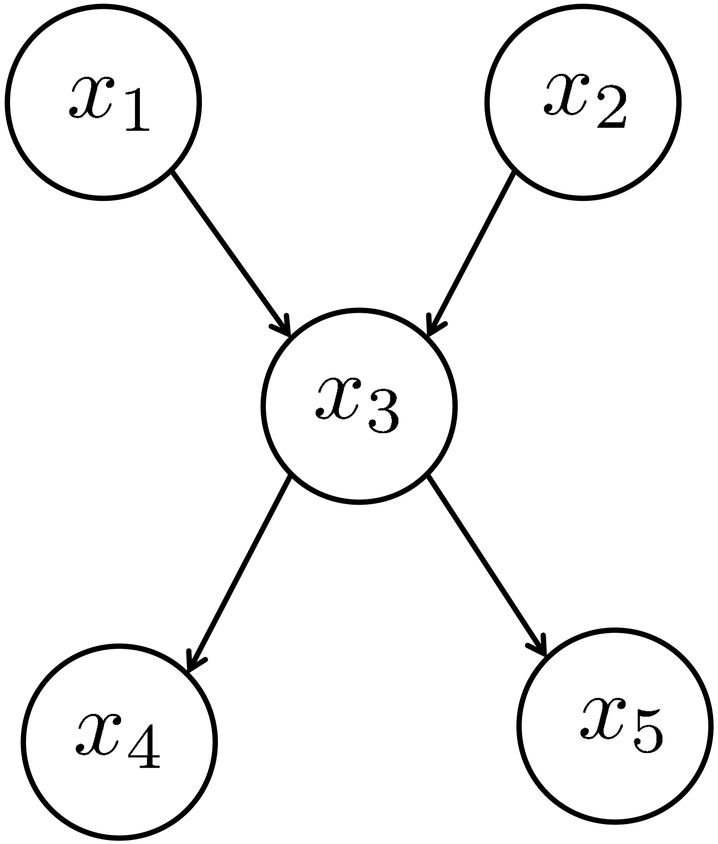

Bayes Network

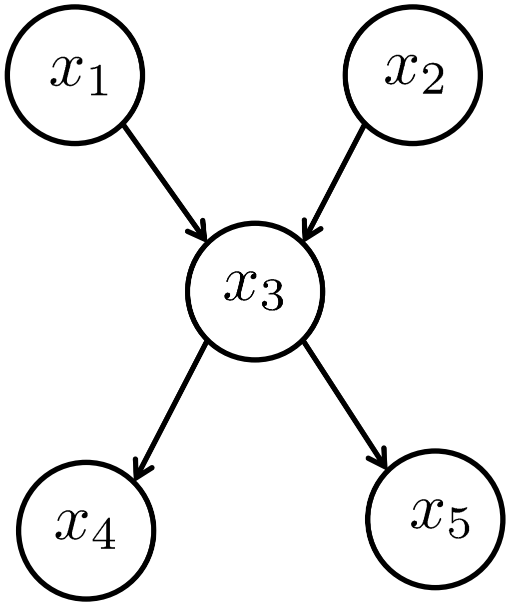

- Expresses dependencies between the variables
- Arrow can be read as "source influences target"
- E.g., x_1 influences x_3, i.e. x_3 depends on x_1
- Cond. independencies allow for simplifying the joint posterior:

$$p(x_1, \ldots, x_5) = p(x_1)p(x_2)p(x_3 \mid x_1, x_2)p(x_4 \mid x_3)p(x_5 \mid x_3)$$

Topics of this Seminar

- Key concepts in graphical models
- Different representations
- Inference algorithms
- Parameter learning
- State sequences and hidden states
- Sampling techniques
- Dimensionality reduction
- Model comparison
- Continuous latent variables

Material

- **See seminar website for online material**
- C. Bishop: Pattern Recognition and Machine Learning (online: Chapter 8). The book is available in the **library** in bldg. 101
- D. MacKay: Information Theory, Inference, and Learning Algorithms (online)
- D. Barber: Bayesian Reasoning and Machine Learning (online)
- Kollar and Friedman: Probabilistic Graphical Models. The book is available in the **library** in bldg. 101
- Klinger & Tomanek: Classical Probabilistic Models and Conditional Random Fields (online)
- Sutton & McCallum: An Introduction to Conditional Random Fields for Relational Learning (online)
- K. Murphy: An Introduction to GMs (online)
- M. Jordan: An Introduction to GMs (online)

Topics Overview

1. Inference and the Sum-Product Algorithm
2. Junction Tree Algorithm
3. Loopy Belief Propagation
   & Loopy Intersection Propagation
4. Hidden Markov Models
5. Mixture Models and EM
6. Sampling Methods
7. Conditional Random Fields
8. Model Comparison and Occam's Razor
9. Continuous Latent Variables

Topics in Detail

1. Inference and the Sum-Product Algorithm

- Inference = Inferring states given potentially noisy observations/information
- Techniques for exact inference given cycle-free graphs

- **Material:** C. Bishop, Chapter 8.4 - 8.4.4 & 8.4.6

Supervisor: Henrik Kretzschmar

Topics in Detail

2. Junction Tree Algorithm

- Tool for inference given a graph with (some) loops/cycles
- Idea: eliminates cycles by clustering variables
- Exact inference
- Inefficient in the general case

- **Material:** D. Barber, Chapter 6

Supervisor: Cyrill Stachniss

Topics in Detail

3. Loopy Intersection Propagation

- Belief Propagation is only exact on cycle-free graphs
- Loopy Intersection Propagation is a technique for efficiently dealing with Gaussian distributions and loopy graphs
- **Material:** D. Tipaldi (Loopy Intersection Propagation) and C. Bishop, Chapter 8.4ff (Belief Propagation)

Student: Philipp

Supervisor: Cyrill Stachniss

Topics in Detail

4. Hidden Markov Models

- Simplest form of dynamic Bayes networks
- Considers state sequences
- Introducing hidden (unobserved) states often allows for simplifications
- Forward-Backward algorithm
- Sum-Product Algorithm for HMMs

- **Material:** C. Bishop, Chapter 13-13.2.3

Student: Oleg

Supervisor: Wolfram Burgard

Topics in Detail

5. Mixture Models and EM

- Parameter Learning
- Known structure, but partial observability
- Algorithm to find a locally optimal MLE

- **Material:** C. Bishop, Chapter 9 - 9.3.2

Student: Nichola
Supervisor: Henrik Kretzschmar

Topics in Detail

6. Sampling Methods

- Often, distributions cannot be represented in closed form
- Sampled distributions are an efficient way for representing posteriors
- Also used for generating samples from a distribution given by the PGM

- **Material:** C. Bishop, Chapter 11-11.2

Student: Stefanie
Supervisor: Henrik Kretzschmar

Topics in Detail

7. Conditional Random Fields

- Graphical model that represent a conditional distribution, i.e. $p(x \mid y)$ instead of $p(x)$
- Often used for labeling or parsing of sequential data

- **Material:** CRF Introduction by Klinger & Tomanek or Sutton & McCallum

Student: Daniel
Supervisor: Wolfram Burgard

Topics in Detail

8. Model Comparison and Occam's Razor

- How to compare different models
- General problem in probabilistic modeling (not only in PGM)
- Approximation method of Laplace
- Occam's Razor
- Model fitting and comparison

- **Material:** D. MacKay, Chapter 27 & 28

Student: Alejandro
Supervisor: Cyrill Stachniss

Topics in Detail

9. Continuous Latent Variables

- Dealing with continuous variables
- Dimensionality reduction
- PCA
- Probabilistic PCA
- Factor Analysis

- **Material:** C. Bishop, Chapter 12-12.2.2 & 12.2.4

Supervisor: Wolfram Burgard

Organization

- "Blockseminar"
- 27.07.2011 9-17 (+28.07.2011 9-13)
- 9 topics for 9 students
- Registration via the internet portal
- First come first serve policy
- Your tasks
  - Give a talk of 30 min (+5-10 min questions)
  - Write a summary (style will be provided)

Important Dates

- **TBD:** lecture on PGM basics
- **July 10:** deadline for show your presentation to your supervisor
- **July 26:** deadline for submitting your summary and your slides via email to stachnis@informatik.uni-freiburg.de
- **July 27 & 28:** seminar, attending both days is mandatory
- Discuss questions and your presentation **early in time** with your supervisor

Lecture on PGM Basics

- Doodle to indicate possible dates: http://www.doodle.com/tedixtic5n28p93h
- Enter your availability until May 9

Choose Your Topic!