Discrete Inference and Learning

Karteek Alahari (Thoth team, INRIA Grenoble)
& Guillaume Charpiat (TAU team, INRIA Saclay)

(formerly Yuliya Tarabalka (Titane team, INRIA Sophia-Antipolis))

Overview

- Course summary and organization
- Chapters overview

Context

Graphical models (or probabilistic graphical models) provide a powerful paradigm to jointly exploit probability theory and graph theory for solving complex real-world problems. They form an indispensable component in several research areas, such as statistics, machine learning, computer vision, where a graph expresses the conditional (probabilistic) dependence among random variables.
This course will focus on discrete models, that is, cases where the random variables of the graphical models are discrete. After an introduction to the basics of graphical models, the course will then focus on problems in representation, inference, and learning of graphical models. We will cover classical as well as state of the art algorithms used for these problems. Several applications in machine learning and computer vision will be studied as part of the course.

TL;DR:

- modeling computer vision problems with graphs
- algorithms for discrete optimization on graphs

Interactive Binary Segmentation

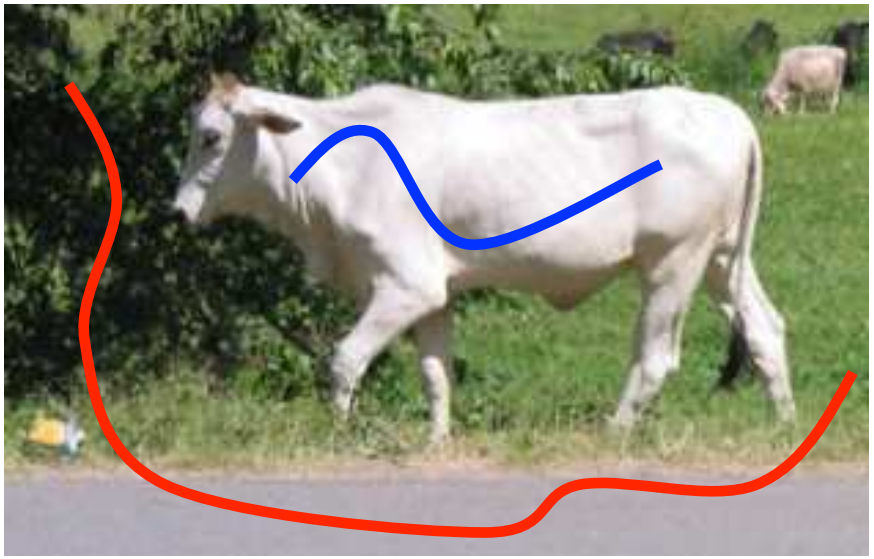

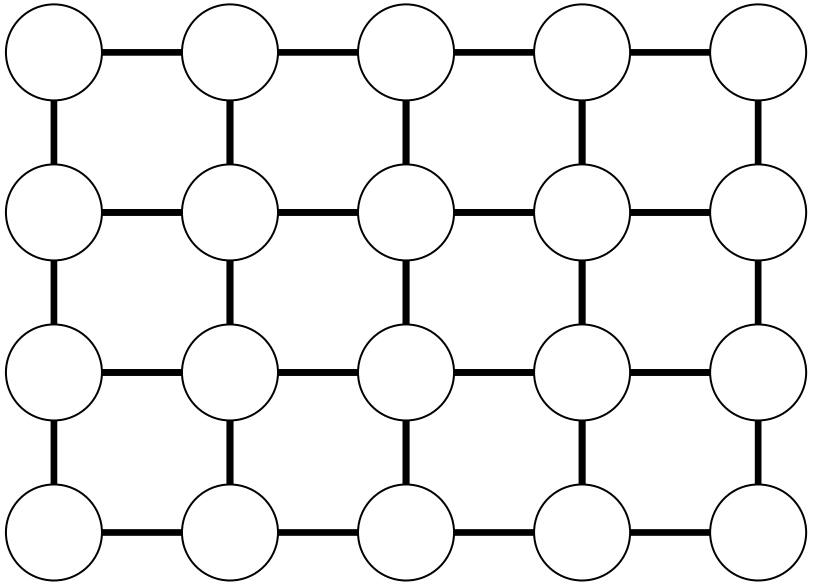

Foreground histogram of RGB values FG

Background histogram of RGB values BG

‘1’ indicates foreground and ‘0’ indicates background

Results

Object Detection

Felzenszwalb and Huttenlocher, 2004

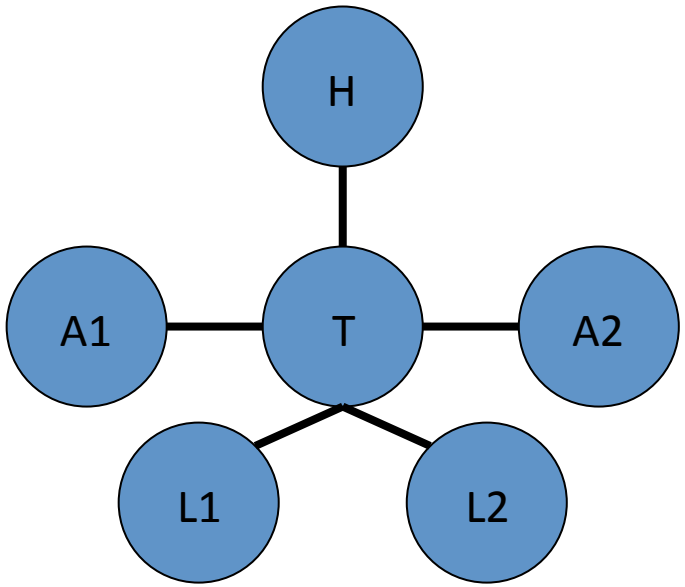

Labels - Poses of parts

Unary Potentials:
Fraction of foreground pixels

Pairwise Potentials:
Favour Valid Configurations

Results

Object Detection

Felzenszwalb and Huttenlocher, 2004

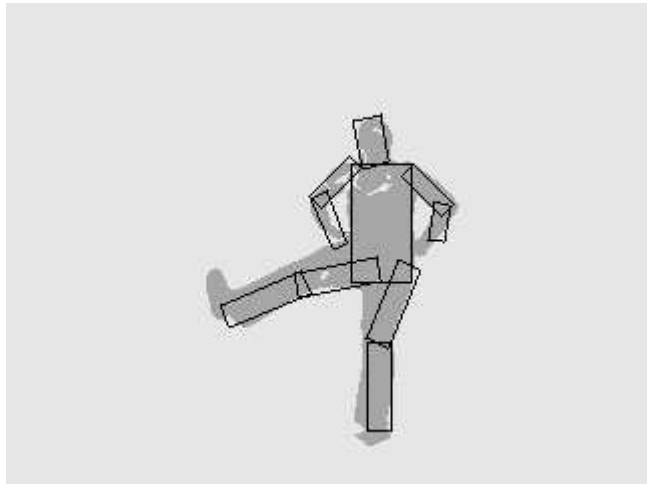

Bayesian Networks

- Example

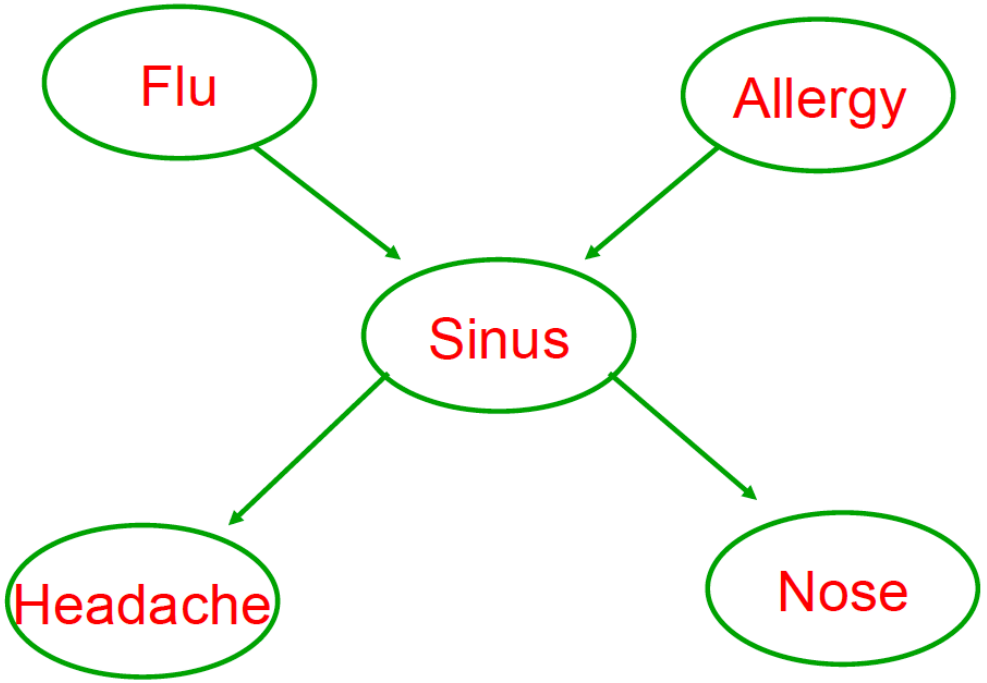

Slide courtesy: Dhruv Batra

Outline

Schedule

- 14/01 : Introduction and basics
  - graphical models
  - Markov Random Fields (MRF)
  - dynamic programming
- 04/02 : Message passing & Belief propagation
- 11/02 : Graph cuts
  - binary energy minimization
  - extension to multi-label energies
  - application to computer vision (image grid graph: segmentation)
- 18/02 : Further optimization on graphs
  - move-making algorithms
  - Ishikawa construction
  - Tree-ReWeighted message passing (TRW)

...

- 25/02 : Extracting features & Primal-dual formulations
  - separability
  - from k-nearest neighbors to decision trees and random forests
  - primal-dual formulations
  - Support Vector Machines (SVM)
- 03/03 : Extracting features (II) : neural networks
  - multi-layer perceptron
  - Convolution Neural Networks (CNN)
  - Designing deep network architectures to tackle image resolution issues, example with remote sensing imagery
- 10/03 : CRF in dual form
  - dual decomposition
  - Conditional Random Fields (CRF)
- 12/03 : Bayesian networks
  - parameter learning
  - structure learning
  - inference
- 17/03 : Project defenses

Organisation and evaluation

- Courses: lessons
- Validation: project

Schedule

8 classes of 3 hours, most often on Tuesday afternoons (13h45 – 17h with a break) at CentraleSupelec (not every week, check the webpage for details).

Webpage & mailing-list:

`https://lear.inrialpes.fr/people/alahari/disinflearn/index.html`

Prerequisite

- General maths (probabilities, Bayesian statistics, analysis, differential calculus...)

To attend the course

- All information will be available on the website soon
- See you on Tuesday at 13h45, room EB.114, Eiffel building (CentraleSupelec)