Probabilistic Graphical Models, Spring 2012

Problem Set 3: Exact inference

Due: Monday, March 19, 2012 at 10am (electronically)

For the following questions, you may use the programming language of your choice. You are allowed to use basic graph packages (e.g., for representing and working with directed or undirected graphs), but are **not** permitted to use any graphical model or probabilistic inference packages.

E-mail a PDF of your full assignment, including all code written, to Chris.

1. The *Alarm* Bayesian network [1], shown in Figure 1, was an alarm message system for patient monitoring that was designed in 1989 as a case study for applying Bayesian networks to medical diagnosis. The Alarm Bayesian network is provided in the file "alarm.bif" (the format should be self-explanatory). This question will explore variable elimination as applied to the Alarm BN.

   Implement the sum-product variable elimination algorithm from class (also described in Section 9.3 of Koller & Friedman). Do not implement pruning of inactive variables. Use the min-fill heuristic (see page 314) to choose an elimination ordering. Note that you do **not** need to make this query-specific, i.e. it should be performed only once on the whole graph and then this elimination order should be used for all queries.

   (a) What is the elimination ordering found by the min-fill heuristic? How many fill edges were added?

   (b) What is the induced width of the graph with respect to the ordering found by min-fill? List the variables in the largest clique.

   Compute the value of the following queries (report up to 4 significant digits):

   (c) $p(\text{StrokeVolume} = \text{High} \mid \text{Hypovolemia} = \text{True}, \text{ErrCauter} = \text{True}, \text{PVSat} = \text{Normal}, \text{Disconnect} = \text{True}, \text{MinVolSet} = \text{Low})$

   (d) $p(\text{HRBP} = \text{Normal} \mid \text{LVEDVolume} = \text{Normal}, \text{Anaphylaxis} = \text{False}, \text{Press} = \text{Zero}, \text{VentTube} = \text{Zero}, \text{BP} = \text{High})$

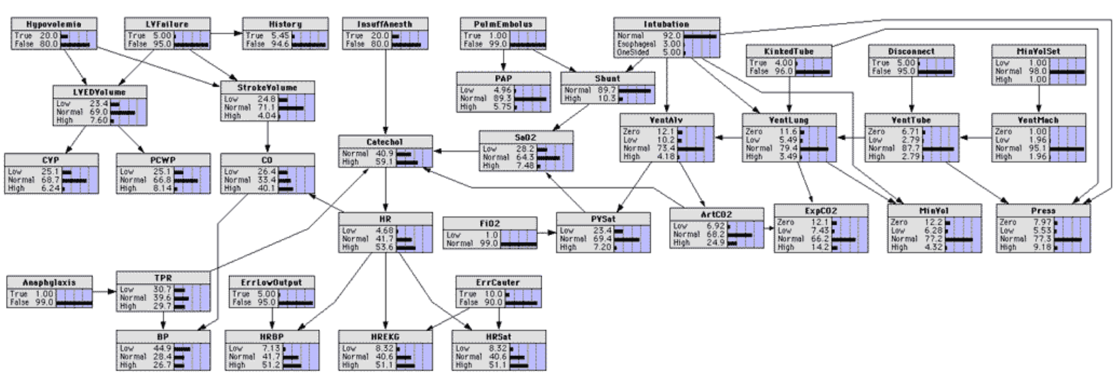

Figure 1: ALARM ("A Logical Alarm Reduction Mechanism") Bayesian network. [1]

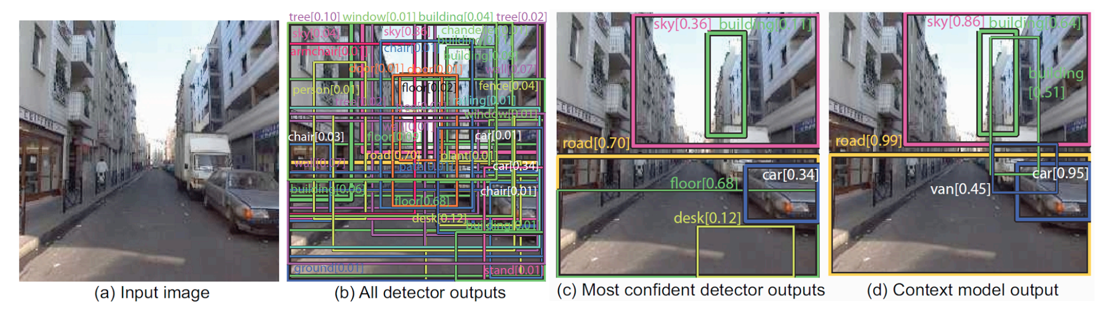

Figure 2: Using context within object detection for computer vision. [2]

(e) $p(\text{LVFailure} = \text{False} \mid \text{Hypovolemia} = \text{True}, \text{MinVolSet} = \text{Low}, \text{VentLung} = \text{Normal}, \text{BP} = \text{Normal})$

(f) $p(\text{PVSAT} = \text{Normal}, \text{CVP} = \text{Normal} \mid \text{LVEDVolume} = \text{High}, \text{Anaphylaxis} = \text{False}, \text{Press} = \text{Zero})$

2. When trying to do object detection from computer images, *context* can be very helpful. For example, if "car" and "road" are present in an image, then it is likely that "building" and "sky" are present as well (see Figure 2). In recent work, a tree-structured Markov random field (see Figure 3) was shown to be particularly useful for modeling the prior distribution of what objects are present in images and using this to improve object detection [2].

   In this question, you will replicate some of the results from [2]. In particular, you will implement both the sum-product and max-product belief propagation algorithms for exact inference in tree-structured Markov random fields. Note, it is not necessary to read this paper to complete this assignment.[1]

   You will use your algorithms to do inference in two CRFs corresponding to two images, one of a kitchen scene (see Figure 4) and the other of an office scene (see Figure 5). The two input files ("kitchen.uai" and "office.uai") are in the format described here:

   `http://www.cs.huji.ac.il/project/PASCAL/fileFormat.php`

   You are allowed to special-case your code for (tree-structured) binary pairwise MRFs.

[1]That said, please see `http://people.csail.mit.edu/myungjin/HContext.html` if you are curious for more details. We omit the spatial prior and the global image features, and use only the co-occurences prior and the local detector outputs (b_i, c_{ik}, and s_{ik} from [2]'s Figure 3).

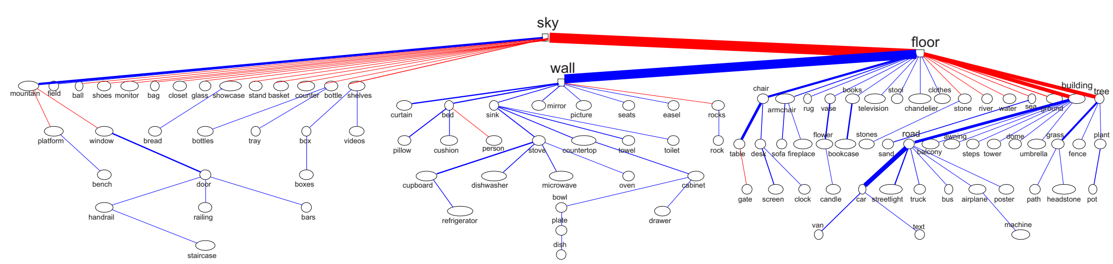

Figure 3: Pairwise MRF of object class presences in images. Red edges denote negative correlations between classes. The thickness of each edge represents the strength of the link. [2]

Figure 4: Kitchen

Figure 5: Office

Each CRF describes the conditional distribution $p(b_1, \ldots, b_{111}, \mathbf{c} \mid \mathbf{s})$, where $b_i \in \{0, 1\}$ denotes the presence or absence of objects of type i in the image (the corresponding object names are given in the file "names.txt"), the variables $\mathbf{c} = \{c_{ik}\}$ where $c_{ik} \in \{0, 1\}$ specify whether location k in the image contains object i, and $\mathbf{s}$ are features of the image. The evidence (i.e. the $\mathbf{s}$ variables) is already subsumed into the edge and node potentials, and so only the $\mathbf{b}$ and $\mathbf{c}$ variables are represented in the provided MRFs.

For each of the below questions, report the answer only for variables 1 through 111 (the object presence variables $\mathbf{b}$), and use the names (e.g., "wall") that are provided.

(a) For each of the two images, what is the MAP assignment, i.e. $\arg\max_{\mathbf{b},\mathbf{c}} p(\mathbf{b}, \mathbf{c} \mid \mathbf{s})$? Just report which objects are present in the image (i.e., names(i) for which $b_i = 1$) according to the MAP assignment.

(b) Use the sum-product algorithm to compute the single-node marginals. For each of the two images, what objects are present with probability greater than 0.8 (i.e., names(i) for which $p(b_i = 1 \mid \mathbf{s}) \geq 0.8$)?

(c) For the two images, what objects are present with probability greater than 0.6?

References

[1] I.A. Beinlich, H.J. Suermondt, R.M. Chavez, and G.F. Cooper. The alarm monitoring system: a case study with two probabilistic inference techniques for belief networks. Technical Report KSL-88-84, Stanford University. Computer Science Dept., Knowledge Systems Laboratory, 1989.

[2] Myung Jin Choi, Joseph J. Lim, Antonio Torralba, and Alan S. Willsky. Exploiting hierarchical context on a large database of object categories. In *IEEE Conference on Computer VIsion and Pattern Recognition (CVPR)*, 2010.