Exercises Algorithmical & Statistical Modelling, Fall 2012, Sheet 6

Please return on Thu, Dec 6, in the lecture

Problem 1 (20 points) (Model some domain of your choice by a BN, with a diagnostic task in mind (e.g., some medical area, or car engine, or an "intelligent" computer user interface that tries to diagnose the state of mind of its current user, or nuclear power plant, or whatever). Think of at least 10 relevant observational variables, some of the functioning as "symptoms" (should be physically measurable), some as "causes", some as mediatory quantities (for instance, blood pressure might be mediating between an original cause "diabetes" to a symptom "kidney pain" – the role of cause vs. intermediary vs. symptom is not always clear cut). Specify the value range for each of your variables. Arrange them in a plausible BN and specify reasonably thought-out tables for the conditional distributions of three of your variables.

Problem 2 (50 points) Below you see a picture by Pop Art artist Roy Lichtenstein, titled "HIM". Graphit and touche on paper, 22 x 17 inches, 1964. Saint Louis Art Museum. Taken from http://www.lichtenstein-foundation.org). Lichtenstein's hallmark is his imitation of raster points (as known from printing) in his paintings – the dots you see on this image are hand-painted... Your task: design a "Lichtenstein dot detector" in the form of Markov random field (MRF), which gets black and white images like this HIM picture as input (= activation pattern on visible units, which correspond to pixels) and develops, through its stochastic update dynamics, a segmentation indicator for the "Lichtenstein-dotted" ares. More precisely, over image areas with Lichtenstein dots, the hidden MRF layer units should develop an activation of +1, while in non-dotted areas, their activation should develop toward –1. – Notice that the Lichtenstein dots are larger than the image pixels, that is, one Lichtenstein dot corresponds to a small cluster of image pixels. Let's say, for simplicity, that each Lichtenstein dot covers 2 x 2 image pixels, with blank pixels in between according to the pattern shown in the schema below (we assume that the Lichtenstein dots all have the same size across different Lichtenstein paintings – I hope you don't mind that we pretend that the world is simple).

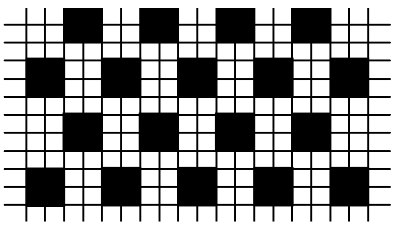

Deliverables: A description of the MRF topology and RVs, plus an energy function for the MRF, plus an explanation in words of the energy component(s) that you choose. You may assume that a local pattern matching algorithm P is given, which gets a 5x5 binary pixel image as input and returns 1 if the input image corresponds to some 5x5 subsection of a Lichtenstein dot pattern, and returns 0 else.

Note: this problem is taken from the final exam 2007 of this lecture.

Problem 3 (30 points). Besides exact inference algorithms like the join-tree algorithm, there exist numerous approximate algorithms for inference in BNs. One common approach is to use sampling for computing approximate inferences in Bayesian networks. Figure out how. Concretely, let $(X_i)_{i=1,...,N}$ be the node variables of a (directed, acyclical) BN with vertices i (may be identified with the X_i) and edges $E \subseteq N \times N$. Assume all X_i have discrete distributions and the conditional distributions $P(X_i \mid pa(X_i))$ are provided as tables. The task is to compute an estimate for $P(Y \mid Z_1 = a_1, ..., Z_k = a_k)$, where $Y, Z_1,..., Z_k$ are among the X_i. (Note: $P(Y \mid Z_1 = a_1, ..., Z_k = a_k)$ is the conditional distribution of Y given one-point "evidence" on the $Z_1,..., Z_k$.) Invent and describe a sampling-based estimation algorithm for $P(Y \mid Z_1 = a_1, ..., Z_k = a_k)$. Note: don't hunt for MCMC sampling, that's not easy. Instead, look out for a kind of primitive brute-force approach, which generates samples $(x_1, ..., x_N)$ for all of the RVs, each sample freshly produced from tabula rasa without the local modification mechanisms typical for MCMC. Your sampling algorithm can be utterly inefficient but it should be correct.