

Assignment for CSC 2506. Due February 7. Worth 15% of the mark

Part I

Design a directed graphical model for a person's choice of clothing for a day (an arbitrary person within some general class, not one person in particular). You should produce both a DAG representing independence relationships, and a quantitative specification of the conditional probabilities that define the joint distribution based on this DAG. You should try to make the model simple (ie, having few arcs in the DAG), to the extent that this can be done while keeping it reasonably realistic.

There is no single correct answer, and, if you like, you may define the exact problem in whatever way seems most interesting, as long as the situation you model is at least as rich as the one described below. (For instance, you shouldn't produce a network with no edges, having decided that all the variables you are considering are independent!)

For a model of men's dress, variables describing the following aspects of dress could be appropriate:

- Whether he is wearing sandals, casual shoes, dress shoes, or cowboy boots.
- Whether he is wearing dress pants, jeans, or shorts.
- Whether he is wearing a T-shirt, a plaid shirt, or a white dress shirt.
- Whether he is wearing a tie.
- Whether he is wearing a suit coat.

His choice of dress could depend on the activities he plans for the day, such as

- Whether he will be attending a business meeting.
- Whether he will be attending a funeral.
- Whether he will be attending a rodeo.

His choice might also be influenced by a personal preference for formal or casual dress, and by whether the weather is expected to be hot. You may find that to get a simple model you need to invent additional variables as well.

You should check out the qualitative aspects of your model by seeing whether various of the conditional independence relationships implied by the d -separation criterion seem reasonable.

Part II

Write a computer program (in whatever language you prefer) that randomly samples from the joint distribution defined by the model you designed in Part I. The program need not be general; it has only to work for this single model.

Use the program to check out the quantitative aspects of your model, by generating about fifty examples of a person's activities and dress for a day (plus any other variables in the model), and deciding whether you think these are reasonable examples of what real people wear when they do these things. If you find flaws in your model, try to fix them, or explain why fixing them is difficult.

Finally, produce a program that samples from the conditional distribution for the person's activities (business meeting, funeral, etc.) given that they are wearing dress shoes, with other clothing unspecified (perhaps you can see only their feet). Do this by sampling from the joint distribution and discarding realizations in which the person isn't wearing dress shoes.