Markov Models

November 9, 2008

1 Modeling

1. Consider modeling the change of weather. Any given day, the weather can be in one of two possible states: there can be either *rain* or *sun*. If it rained yesterday, the probability for *rain* today will be $p_{r \to r} = 0.7$, and the probability for *sun* will be $1 - p_{r \to r} = 0.3$. If it was sunny yesterday, the probabilities for *rain* and *sun* today are $p_{s \to r} = 0.4$ and $1 - p_{s \to r} = 0.6$. This situation can be modeled by the Markov model depicted below:

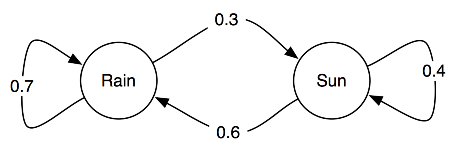

Model this domain in your language.

2. Now consider you are a prisoner in jail and you're currently planning a jailbreak. As you have a strong interest in not getting wet after escaping you need to know the current weather. However, there is no window in your cell and the only indication of the weather you get is the observation of the guard carrying an umbrella or not. Obviously, the probability that the guard carries an umbrella depends on the weather. On rainy days, the guard carries an umbrella with a probability of $p_{r \to u} = 0.8$, and on sunny days, with probability $p_{s \to u} = 0.1$. The probabilities $p_{r \to s}, p_{s \to s}$ for weather changes are as in Task 1. This situation is nicely modeled by the hidden Markov model shown below:

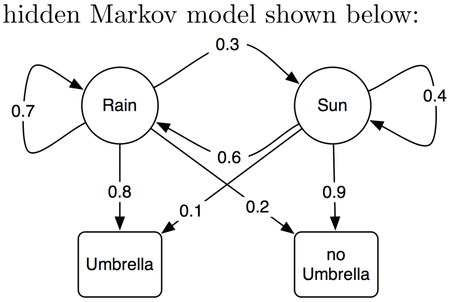

3. As everybody knows, an alternative theory is that the weather actually depends on whether or not we carry an umbrella, as it always seems to rain on the days that we don't. Specifically, lets assume that:
If it rained yesterday and the guard carries an umbrella today, it will rain with probability $p_{ru \to r} = 0.3$;
If it rained yesterday and the guard carries no umbrella today, it will rain with probability $p_{r\bar{u} \to r} = 0.9$;
If it was sunny yesterday and the guard carries an umbrella today, it will rain with probability $p_{su \to r} = 0.4$;
If it was sunny yesterday and the guard carries no umbrella today, it will rain with probability $p_{s\bar{u} \to r} = 0.8$.

   Model this situation in your language.

4. Let us go back to the model where the fact that you carry an umbrella or not depends on the weather. But now you have not a single guard but more than one. Additionally you know that every guard behaves the same, that is, the probability of carrying an umbrella or not is the same for each guard. Probability values are $p_{r \to s}, p_{s \to s}$ and $p_{r \to u}, p_{s \to u}$ as in Task 2.

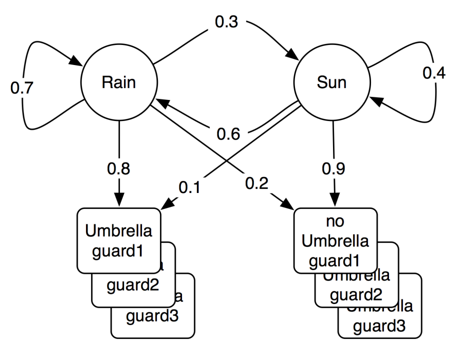

   Model this situation in your language. In particular, make sure that you enforce the constraint that all guards share the same probability distribution.

5. Consider extending the model developed in Task 3 to include an (arbitrary) number of guards (as in Task 4). That is, the current weather should be influenced by the guards carrying umbrellas or not. Specifically, we assume that if at least one guard carries an umbrella it will probably rain.

   How can such a situation be modeled in general in statistical relational learning systems? Give a model for this situation in your system if possible, or argue why this is not possible.

2 Learning

We will consider the model for multiple guards developed in modeling Task 4) for all learning tasks given below.

1. Let us first consider parameter learning from fully observable data. The dataset *seq_2guards_summer.txt* contains 10 sequences of observations (umbrella or not) for three guards, and the corresponding hidden states (sun or rain).

   Use the learning algorithm built into your system to learn the parameters of the model from data (note you probably need to bring the file in an appropriate format for your system). Start with a model that has uniform initial state, transition and observation probabilities. We denote the learned model by M_{summer}.

   Report the parameters learned by the model, and the final likelihood of the training set given the learned model.

2. Alternatively, one can learn a model from the observations only. The dataset *seq_3guards_winter.txt* contains sequences in which the hidden state is universally *bars*, meaning that it was not observed.

   Again learn a model from this partially observable data, and report the model parameters and training set likelihood. Denote the learned model by M_{winter}.

3. We have now found a sequence of observations, *seq_test_inv.txt*, and would like to know whether the sequence was recorded in summer or in winter. To decide this, compute the likelihood of this sequence under the two models M_{winter} and M_{summer} trained above. Report the likelihoods.

4. The quality of a learned model typically depends on the amount of training data available. Take the dataset from the season matching the test data *seq_test_inv.txt*, which can be *seq_3guards_winter.txt* or *seq_2guards_summer.txt* depending on the result of the previous task.

   Learn models M_3, M_6, M_{10} taking only the first 3,6, and 10 sequences from this dataset. Report the likelihood of *seq_test_inv.txt* for these three models.