LEARNING ALGORITHMS AND SYSTEMS LABORATORY

FACULTE DES SCIENCES ET TECHNIQUES DE L'INGENIEUR

SCHOOL OF ENGINEERING / SWISS INSTITUTE OF TECHNOLOGY

Lecturer:
Prof. ***Aude Billard***
Office: *ME A3 464*
direct: +41-21-693 54 64
secretary:+41-21-693 09 39
fax: +41-21-693 78 50
E-mail: aude.billard@epfl.ch
Web: http://lasa.epfl.ch/

Assistant:
Dr. **Basilio Noris**
Office: *ME A3 395*
direct: +41-21-693 7824
E-mail: basilio.noris@epfl.ch

Exercises - III

Exercise 1: Hidden Markov Models

Imagine someone trying to deduce the weather from a piece of seaweed - folklore tells us that `soggy' seaweed means wet weather, while `dry' seaweed means sun. If it is in an intermediate state (`damp'), then we cannot be sure. However, the state of the weather is not restricted to the state of the seaweed, so we may say on the basis of an examination that the weather is probably raining or sunny. A second useful clue would be the state of the weather on the preceding day (or, at least, its probable state) - by combining knowledge about what happened yesterday with the observed seaweed state (soggish, damp, dryish, dry) we might come to a better forecast for today.

a) Design an HMM model (define its structure, number of states, hidden states, connectivity, variables to estimate, etc) that explains the observations, assuming that the hidden states (the true weather) are modelled by a simple first order Markov process. Draw the model's structure.

b) Show the unfolding in time of the observation and state sequence used, e.g. in the Forward estimation procedure.

c) Explain how you can make an estimate of what the weather was for a week given each day's seaweed observation.

d) Given a 1-week long sequence of seaweed observations, explain how you can use the model to determine whether it is winter or summer. Intuitively, if the seaweed has been dry for a while it may be summer; if it has been soggy for a while it might be winter.

e) Explain how you could adapt your model to learn transition across summer and winter based on a year-long set of seaweed observations

Exercise 2 (Typical Exam Question!):

Imagine that you have by now graduated from EPFL and work as a consultant for a car manufacturer. The car manufacturer sells 5 car models in 10 different countries. It has gathered data on its sales for the past 5 years and runs each year large pools to estimate its clients' tastes and level of satisfaction.

I) What would you recommend to your client to solve the following problems?

a) To predict the sales in each country for the next 5 years.

b) To determine which aspects of a car (speed and driving performance, aesthetic, costs, etc.) are most important to your client to help you design a new car model.

For each problem, suggest one or more Machine Learning algorithm(s) that could be used to solve the problem, and explain how it/they would solve it. That is, in each case, explain what the inputs and outputs of the algorithm would be (i.e. how the data are encoded), what cost function or bias you will use (only if applicable to your algorithm). Justify the use of this algorithm by explaining what type of computation the algorithm performs (classification, decorrelation, optimization, etc.) and why it is relevant for your problem. Draw a diagram that shows the flow of information with the variables. If the algorithm requires a preprocessing of the input, mention it explicitly.