Norwegian University of Science and Technology
Department of Computer and Information Science

Language: English

EXAMINATION IN METHODS IN ARTIFICIAL INTELLIGENCE (TDT4171)

Tuesday 20. May 2014
Hours: 09:00 – 13:00

Read the text of each problem carefully. Make sure that you understand the problem. If you consider the information given in an problem incomplete or inaccurate, then make a note of the assumptions you find necessary to make in order to solve the problem.

Important: The following problems all have a common theme. However, they are totally independent of each other. This means that if you are having difficulties with a problem, you can safely move to the next one.

Problem 1 (30%)

We want to launch a web site that recommends content for users based on their interests. Our first task when creating a recommendation system is to filter out articles which we assume will have a low probability for the user to click on. After a lot of experimentation, we have found that there are two main indicators of this probability:

- The *interestingness* of an article as measured by similar articles the user has clicked on previously.
- The *popularity* of an article as measured by the number of clicks from all users of the system.

Further, we have found that the probability of a user clicking on an article that is both interesting and popular is `0.8`. If it is popular but not interesting, the probability is `0.6`. If it is interesting but not popular, the probability is `0.4`. The probability is `0.1` if the article is neither.

Assume our system has two articles, `a1` and `a2`. Based on previous click feedback, we have calculated their probabilities of being popular at `0.8` and `0.3`, respectively. We also have two registered users in our system, `u1` and `u2`. Based on their click history, we calculate that `u1` finds articles `a1` and `a2` interesting with a probability of `0.5` and `0.1`, respectively. The interests of `u2` are `0.3` and `0.7`, respectively.

a) Model the domain above using a Bayesian network. Make the model as simple and easy to understand as possible.

Show the graphical structure as well as the conditional probability tables. Include *users* and *articles* in the model, and assume uniform priors on their probabilities.

Answer

```
User -> Interesting
Article -> Interesting, Popular
Interesting -> Click
Popular -> Click

Users:              Articles:
u1  | 0.5           a1 | 0.5
u2  | 0.5           a2 | 0.5

Interesting:        Popular:
u1,a1 |  0.5        a1 | 0.8
u1,a2 |  0.1        a2 | 0.3
u2,a1 |  0.3
u2,a2 |  0.7

Click:
I,P   | 0.8
I,~P  | 0.4
~I,P  | 0.6
~I,~P | 0.1
```

b) Describe conditional dependency and conditional independency, and give an example of each from your model.

Answer

```
CI: popular and interesting given article
CD: article and user given click/interesting, Popular and interesting given click
```

c) Calculate the probability that user `u1` will click on article `a1`:
`P(Click=true|Article=a1,User=u1)`.

Answer

```
p(c|d1,u1) = sum(p)sum(i) (p(c|p,i) * p(p|a1) * p(i|a1,u1))

     p,i: 0.8 * 0.8 * 0.5
    p,~i: 0.6 * 0.8 * 0.5
    ~p,i: 0.4 * 0.2 * 0.5
   ~p,~i: 0.1 * 0.2 * 0.5

   = (0.8*0.8 + 0.6*0.8 + 0.4*0.2 + 0.1*0.2) * 0.5
   # 0.5 makes it easier to calculate
   = 0.61
```

Problem 2 (30%)

Now, in the previous problem we skipped over *why* a particular user would be interested in an article. Let us assume that each article consists of a main topic, and each user is primarily interested in one topic. We would like to find each user's interests to better serve articles that they are interested in. Obviously, we can't know the true interests of a user, we can only observe what she clicks on.

We can model this using a dynamic model. Each day we examine all the clicks a user has made and update the user's `Interests` based on the topic she has clicked the most on that day. Let us assume we have three topics: `Finance`, `Sport`, and `Tech`. The conditional probability table for `P(Clicked|Interests)` is:

	Clicked		
Interests	Finance	Sport	Tech
Finance	0.50	0.25	0.25
Sport	0.25	0.50	0.25
Tech	0.25	0.25	0.50

Clue

```
All interests for new users are equally likely, as we have no prior information.
For now we assume that the user's interests do not change over time.
```

a) A new user arrives at our site and clicks mostly on `Tech` on the first day. Calculate `P(Interests|Clicked=Tech)`. Assume a uniform prior on `Interests` as we have no previous information on the interests of new users.

The second day, the user clicks mostly on `Tech` again. Calculate `P(Interests|Clicked=Tech)` again, but this time use the results from the first day as a prior on `Interests`.

Answer

```
tech = 2/3, finance = 1/6, sport = 1/6.
```

b) We continue to update each user's interests every day based on their clicks and throw away previous data. What is this assumption we use called? Explain why it is reasonable here.

Answer

```
Clumsy question. We're not "throwing away" data!
First-order markov assumption
Previous measures are conditionally independent.
```

After a while we find that it was unreasonable to assume that a user's interests do not change over time. For instance, many people are interested in Soccer during the World Cup, but less so at other times. These short term interests introduce some uncertainty of what a user really is interested in.

c) Suggest a transition model that updates a user's interests so they over time return to equal probability if no other input from the user is given.

Answer

```
to same state: 0.9 to other state 0.05
```

d) Temporal models such as the one we have built here have four basic inference tasks: `filtering`, `prediction`, `smoothing` and `most likely explanation`. Describe briefly what these inference tasks compute.

Answer

```
Filtering: Computes the belief (posterior) given all evidence to the present

Prediction: Computes the posterior over a future state given all evidence to the
present

Smoothing: Computes the posterior over a past state, given all evidence to the
present

Most likely: Most probable squence of states that generated the evidence
```

e) After a long period of observations, we can better estimate the user's true long term interests as they tend not to change much. Of the four inference tasks, which one is suitable for this task?

Answer

```
smoothing
```

Problem 3 (20%)

To improve our recommendation system further, we'd like to keep track of some more detailed information of what users actually are clicking on. For instance, we can calculate the number of clicks per region, age range, gender, time of day, and in total for any article. We can also calculate this for all topics. For any given user, we can calculate when they tend to click on which topics during the day. There are many other statistics we also can calculate.

We would like to use these statistics to train a system to predict whether or not a user will click on a given article.

a) In what cases does it make sense for intelligent systems to learn from data? Point out reasons relevant to our system.

Answer

```
unknown environment
laxy designers: system construction method
undergoes changes over time
```

b) Highlight some differences between decision trees and neural networks. Which approach would you use for this problem and why?

Answer

```
Might be a bit unclear here: what was meant that the system should predict for _any_
given article, meaning that the model would need to use information about both the
user and article as inputs. This is in contrast to building a model for a _single_
article, where the model would only need information about the user. Of course, you
could build as many models as you have articles, but then you need a large amount of
data.

So, the inputs would be statistics of the document applicable for the given user.
I.e. first retrieve relevant statistics for the article given what we know about the
user: region, age, gender etc. These statistics are then put into the model, and a
probability is calculated based on these.

These statistics are real-valued, and the probability is real-valued, so it makes
sense to use neural networks.

However, since the problem text is a bit unclear (it does not specifically say to
calculate the probability), one could very well interpret this as a series of
decisions based on data for each article whether or not a user would click on it.
This requires discretization of continuous data, which has not been curriculum.
But it can't be ruled out.

Based on this, the grade given should be based on how cognizant the student is
of these matters.
```

c) Each time a user is presented with an article, we record the the corresponding statistics at that point in time and whether or not she clicked. Given historical data, how would you set up a learning task for this problem?

Answer

```
Supervised learning, use the statistics as input values, the click or skip
as target value. Minimize the error between predicted and target values
(for instance through gradient descent / backprop).
```

d) A common problem in machine learning is overfitting. Explain this problem and what we can do to avoid it.

Answer

```
Does not generalize well to unseen data, overfits training data. Use validation sets
to avoid it for NN, for DT post pruning etc.
```

Problem 4 (20%)

After putting in all the work above in our site, we've become pretty successful. However, with increased costs due to higher traffic, we need to generate some income. We will do that by using our recommendation system to present relevant advertisements to each user.

We start by grouping users together by region, age, gender and interests. Each time we want to present the user with an ad, we first find the group that best matches what we know about the user.

For each group we then need to select a relevant ad, but there are many unknowns here. For instance, we don't know which ads will appeal to which types of users. We would still like to select the best advertisement as revenue will be generated only if the user clicks on it. Our goal is to maximize the total revenue over time.

a) Briefly describe what reinforcement learning is, and why it is useful here.

Answer

```
Reinforcement learning only gives feedback intermittently.
We only get rewards (revenue) when an ad is clicked.
```

b) When selecting an ad to show we need to strike a balance between exploitation and exploration. Describe this balance and suggest a method suitable for our problem.

Answer

```
Bandits! We can't be greedy or random. Some formula that mixes more exploitation
the more we know would be good. Factors here: value increases quicker for higher
paid ads etc.
```

c) Explain how we can solve the problem of maximizing total revenue over time using Q-learning. How do we handle the non-deterministic outcomes of choosing an advertisement?

Answer

```
Action/state table is ad/group with some extra reward states for each ad.
No gamma value here (0) as this is really a simple decision.

Each time we want to present an ad, select the state and choose the action
(ad) according to the action selection in c).

Another table for visits as described in the lecture takes care of non-determinism.

Difficult question?
```