GHENT UNIVERSITY

ARTIFICIAL INTELLIGENCE (E016350A)
GHENT UNIVERSITY
AY 2024/2025

Aleksandra Pizurica
Asst: Yoann Arhant
E-mail: ai@lists.ugent.be

Exercises: Rational decisions

1. The Surprise Candy Company makes candy in two flavors: 70% are strawberry flavor and 30% are anchovy flavor. Each new piece of candy starts out with a round shape; as it moves along the production line, a machine randomly selects a certain percentage to be trimmed into a square; then, each piece is wrapped in a wrapper whose color is chosen randomly to be red or brown. 80% of the strawberry candies are round and 80% have a red wrapper, while 90% of the anchovy candies are square and 90% have a brown wrapper. All candies are sold individually in sealed, identical, black boxes.

   Now you, the customer, have just bought a Surprise candy at the store but have not yet opened the box. Consider the three Bayes nets in figure below.

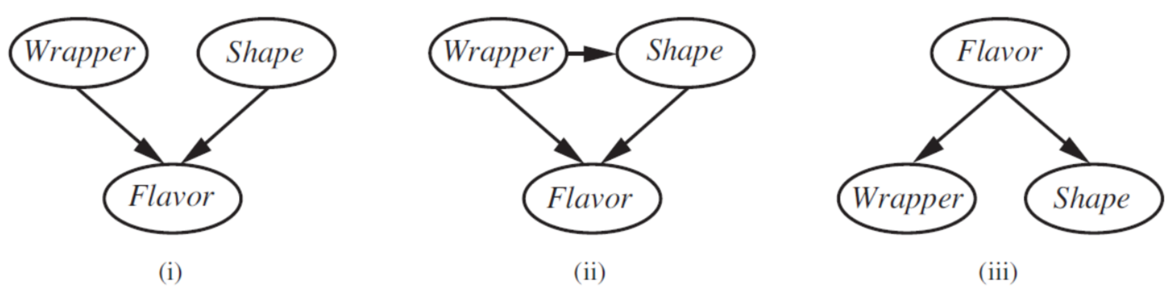

Figure 1: Three proposed Bayes nets for the Surprise Candy problem.

   (a) Which network(s) can correctly represent $\mathbf{P}(Flavor, Wrapper, Shape)$?
   (b) Which network is the best representation for this problem?
   (c) Does network (i) assert that $\mathbf{P}(Wrapper|Shape) = \mathbf{P}(Wrapper)$?
   (d) What is the probability that your candy has a red wrapper?
   (e) In the box is a round candy with a red wrapper. What is the probability that its flavor is strawberry?
   (f) A unwrapped strawberry candy is worth s on the open market and an unwrapped anchovy candy is worth a . Write an expression for the value of an unopened candy box.
   (g) A new law prohibits trading of unwrapped candies, but it is still legal to trade wrapped candies (out of the box). Is an unopened candy box now worth more than, less than, or the same as before?

2. In 1713, Nicolas Bernoulli stated a puzzle, now called the St. Petersburg paradox, which works as follows. You have the opportunity to play a game in which a fair coin is tossed repeatedly until it comes up heads. If the first heads appears on the n-th toss, you win $\$2^n$.

   (a) Show that the expected monetary value of this game is infinite.

   (b) How much would you, personally, pay to play the game?

   (c) Nicolas's cousin Daniel Bernoulli resolved the paradox in 1738 by finding that the utility of money is logarithmic (i.e., $U(S_n) = a \log_2 n + b$ where S_n is the state of having $\$n$). Assume $k = c$ for simplicity, i.e., we paid all we have to play the game. What is the expected utility of the game under this assumption?

   (d) What is the maximum amount it would be rational to pay to play the game, assuming that one's initial wealth is $\$k$?

3. Consider a student who has the choice to buy or not buy a textbook for a course. We'll model this as a decision problem with one Boolean decision node, B, indicating whether the agent chooses to buy the book, and two Boolean chance nodes, M, indicating whether the student has mastered the material in the book, and P, indicating whether the student passes the course. Of course, there is also a utility node, U. A certain student, Sam, has an additive utility function: 0 for not buying the book and $-\$100$ for buying it; and \$2000 for passing the course and 0 for not passing. Sam's conditional probability estimates are as follows:

   - $\mathrm{P}(p \mid b, m) = 0.9$
   - $\mathrm{P}(p \mid b, \neg m) = 0.5$
   - $\mathrm{P}(p \mid \neg b, m) = 0.8$
   - $\mathrm{P}(m \mid b) = 0.9$
   - $\mathrm{P}(m \mid \neg b) = 0.7$
   - $\mathrm{P}(p \mid \neg b, \neg m) = 0.3$

   You might think that P would be independent of B given M, but this course has an openbook final-so having the book helps.

   (a) Draw the decision network for this problem.

   (b) Compute the expected utility of buying the book and of not buying it.

   (c) What should Sam do?

4. **(Old exam question)** It is Sunday evening and Sam is finishing up preparing for the AI exam that is coming up on Monday. Sam has already mastered all the topics except one: Decision Networks. He is contemplating whether to spend the remainder of his evening learning that topic ($A = \textit{learn}$), or just go to sleep ($A = \textit{sleep}$). Decision Networks are going to be on the exam ($D = d$) or not ($D = \neg d$) with equal *a-priori* probabilities. The student's utility of satisfaction is only affected by these two random variables D and A, as shown below:

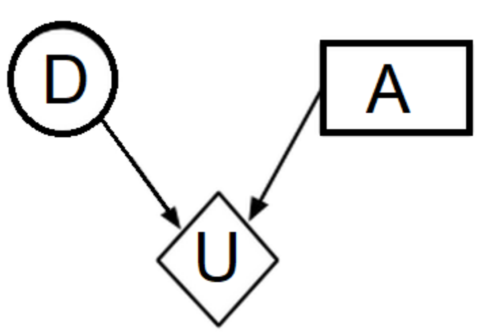

D	$\mathbf{P}(D)$
d	0.5
$\neg d$	0.5

D	A	$U(D, A)$
d	$learn$	1400
$\neg d$	$learn$	800
d	$sleep$	0
$\neg d$	$sleep$	2000

(a) Compute the following, where EU is expected utility and MEU is maximum expected utility:

(i) $EU(learn) =$

(ii) $EU(sleep) =$

(iii) $MEU(\{\}) =$

(iv) Action that achieves $MEU(\{\}) =$

(b) The teaching assistant's happiness (H) on the same day is affected by whether decision networks are going to be on the exam. The happiness (H) influences the chance that the assistant posts on Facebook (f) or doesn't post on Facebook ($\neg f$): if happy, he posts with probability 0.6 and posts with probability 0.2 if unhappy. The prior on D and utility tables remain unchanged. Assume the prior probabilities $\mathrm{P}(h) = 0.6$ and $\mathrm{P}(f) = 0.44$. Further on, the assistant is happy with probability 0.95 if the decision networks are on the exam and is happy with probability 0.25 if they are not on the exam. Some tables that are computed from the data above and that might be useful in answering the questions below are:

D	F	$\mathbf{P}(D \mid F)$
d	f	0.666
$\neg d$	f	0.334
d	$\neg f$	0.370
$\neg d$	$\neg f$	0.630

F	D	$\mathbf{P}(F \mid D)$
f	d	0.586
$\neg f$	d	0.414
f	$\neg d$	0.300
$\neg f$	$\neg d$	0.700

(i) Draw the decision network that reflects this extended description of the problem.

(ii) Compute the expected utilities of the two actions in cases when the assistant posts or not on Facebook:

* $EU(learn \mid f) =$
* $EU(sleep \mid f) =$
* $EU(learn \mid \neg f) =$
* $EU(sleep \mid \neg f) =$

(iii) What are the maximum expected utilities and what are the corresponding optimal actions?

* $MEU(\{f\}) =$
* Optimal Action($\{f\}$) $=$
* $MEU(\{\neg f\}) =$
* Optimal Action($\{\neg f\}$) $=$