Probabilistic Robotics, BAIPR6, Fall 2011

Partial Exam, Wednesday November 23th, 13:00 - 15:00, room G4.15

Arnoud Visser

December 13, 2011

Question 1

This question is based on the environment for knowledge-based agents introduced in 'Artificial Intelligence - A Modern Approach'[1]. The **wumpus world** is a cave consisted of rooms connected by passageways. Lurking in the cave is the terrible wumpus, but we ignore the beast for this question. An agent would enter this cave in search for a piece of gold, but also this reward is not in scope for this question. Only relevant is the knowledge that some rooms contain bottomless pits that will trap anyone who wanders in these rooms. A sample wumpus world of 4x4 is given in Figure 1.

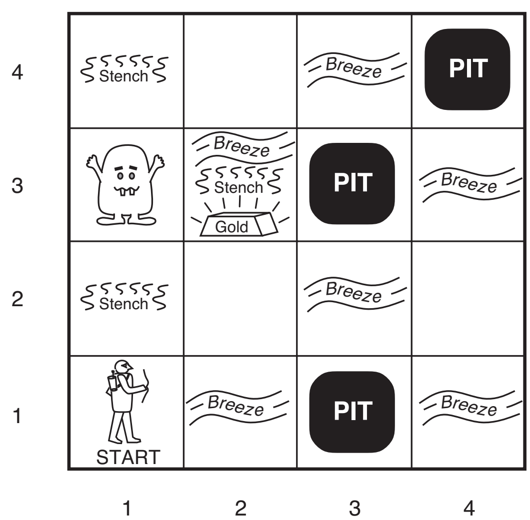

Figure 1: A typical wumpus world. The agent is in the bottom left corner. Further there are three pits, a wumpus and a piece of gold. (Courtesy Russell and Norvig [1]). Note: in this problem you may forget the wumpus, the stench and the piece of gold. Concentrate on the pits and the occurrence of a breeze.

The agent will perceive a *Breeze* in a room directly (not diagonally) adjacent to a pit. The world is partially observable, the agent give only partial information about the world. For example, Figure 2 shows a situation in which each of the three reachable squares - [1,3], [2,2], and [3,1] - might contain a pit. Pure logical inference can conclude nothing about which square is most likely to be safe, so a logical agent might have to choose randomly. We will see that a probabilistic agent can do much better than the logical agent.

The relevant properties of the wumpus world of Figure 2 is the prior information that (1) a pit causes breezes in all neighboring squares, and (2) each square other than [1,1] contains a pit with probability $p_{i,j} = 0.2$. The agent has moved around in the environment and collected the observations $z = \neg b_{1,1} \wedge b_{1,2} \wedge b_{2,1}$ about the presence of breeze $b_{i,j}$ and the evidence $e = \neg p_{1,1} \wedge \neg p_{2,1} \wedge \neg p_{1,2}$ about the presence of a pit $p_{i,j}$. To make a decision for the next square to go to, we are interested in the probabilities of $\mathbf{P}(P_{1,3}|e,z)$, $\mathbf{P}(P_{2,2}|e,z)$, and $\mathbf{P}(P_{3,1}|e,z)$. If we concentrate on calculating $P_{1,3}$ and call this the $query$-square, the other two probabilities $P_{2,2}, P_{3,1}$ can be calculated equivalently. The wumpus world can than be divided up in 4 regions: the $query$-square we like to know the probability from, the $known$-squares were we have

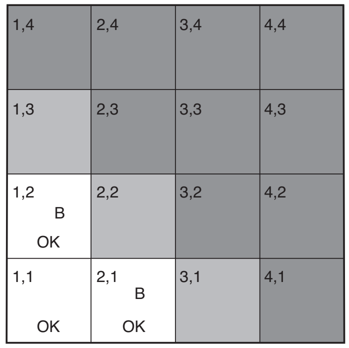

Figure 2: After perceiving a breeze in both [1,2] and [2,1], the logical agent is stuck - there is no save place to explore (Courtesy Russell and Norvig [1]). Note: probability can be used to show that some places are less safe than others.

evidence e from, the frontier the probabilities are conditioned by observations z from and the $other$-squares behind the frontier. One can now argue that the other squares cannot cause the detection of a breeze in the $known$-region; the observation z is independent of the state of $other$ iff the state of $known$, $query$ and $frontier$ is given.

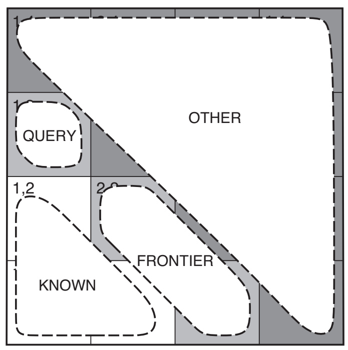

Figure 3: Division of the state space in the regions $known$, $frontier$, $query$ and $other$ (Courtesy Russell and Norvig [1]).

Based on this division of regions, the following inference can be made:

$$
\begin{aligned}
\mathbf{P}(P_{1,3}|e,z) &\equiv \mathbf{P}(query|known,z) \\
&= \eta \sum_{frontier} \sum_{other} \mathbf{P}(z|query,known,frontier,other)\mathbf{P}(query,known,frontier,other) \\
&= \eta \sum_{frontier} \mathbf{P}(z|query,known,frontier) \sum_{other} \mathbf{P}(query,known,frontier,other) \\
&= \eta \sum_{frontier} \mathbf{P}(z|query,known,frontier) \sum_{other} \mathbf{P}(query)P(known)P(frontier)P(other) \\
&= \eta P(known) \sum_{frontier} \mathbf{P}(z|query,known,frontier)P(frontier) \sum_{other} P(other) \\
&= \eta^{'} \sum_{frontier} \mathbf{P}(z|query,known,frontier)P(frontier)
\end{aligned}
$$

Notice that the expression $\mathbf{P}(z|query,known,frontier)$ is a filter, with value 1 when the observation $z = \neg b_{1,1} \wedge b_{1,2} \wedge b_{2,1}$ can be explained by enough pits in the frontier-query combination (between 1 and 3 pits), and value 0 otherwise. There are in total 5 configurations of the frontier consistent with this observation (see Figure 4 and only for those 5 configuration $P(frontier)$ should be calculated.

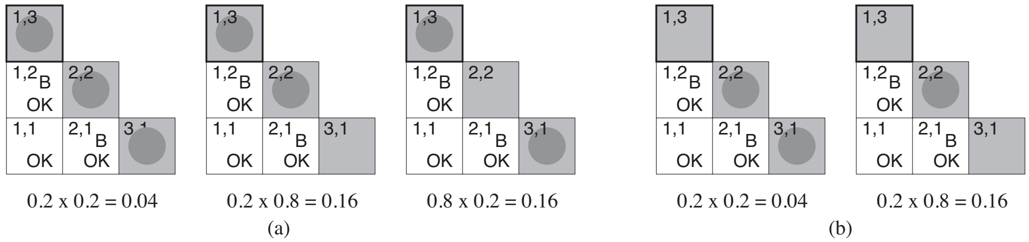

Figure 4: Configuration of the frontier variables $P_{2,2}$ and $P3,1$ which are consistent with the observation $z = \neg b_{1,1} \land b_{1,2} \land b_{2,1}$ (Courtesy Russell and Norvig [1]): (a) three configurations with $P_{1,3} = true$ and (b) two configurations with $P_{1,3} = false$

The resulting probability indicates that based on the observations and evidence the change on a pit for square [1,3] is 31%:

$$\mathbf{P}(query|known, z) = \eta^{'} \langle 0.2(0.04 + 0.16 + 0.16), 0.8(0.04 + 0.16) \rangle \approx \langle 0.31, 0.69 \rangle \tag{1}$$

(a) Show that the square [2,2] contains a pit with a probability of roughly 86%.

(b) In the previous calculation the prior probability of a pit was estimated with probability 0.2, independently of the presence of pits of the other squares. Suppose now instead that $N/5$ pits are scattered at random among the N squares other than [1,1]. Are the variables $P_{i,j}$ and $P_{k,l}$ still independent? What is the joint distribution $\mathbf{P}(P_{1,1}, \ldots, P_{4,4})$ now? Redo the calculation for the probabilities of pits in [1,3] and [2,2].

Question 2

For this question you have to rely on a distance-only sensor. You try to locate your friend using her cell phone signals. Suppose that on the map of Amsterdam, the Science Park is located at $m_0 = (10, 8)^T$, and your friend's home is situated at $m_1 = (6, 3)^T$. You have access to the data received by two cell towers. You have access to the data received by two cell towers, which are located at the positions $x_0 = (12, 4)^T$ and $x_1 = (5, 7)^T$, respectively. The distance between your friend's phone and the the towers can be computed from the intensities of your friend's cell phone signals. The distance measurements are distributed by white Gaussian noise with variances $\sigma_0^2 = 1$ for tower 0 and $\sigma_1^2 = 1.5$ for tower 1. You receive the distance measurements $d_0 = 3.9$ and $d_1 = 4.5$ from the two towers.

(a) Make a drawing of the situation.

(b) At which of the two places is your friend more likely to be? Explain your calculations.

(c) Now, suppose you have prior knowledge about your friend's habits which suggest that your friend is currently is at home with probability $P(at\_home) = 0.7$, at the university with $P(at\_university) = 0.3$ and at any other place with $P(other) = 0$. Use this prior knowledge to recalculate the likelihoods of b).

Question 3

A common drive model for indoor robots is *holonomic*. A holonomic robot has many controllable degrees of freedom as the dimension of its configuration (or pose) space. Here, you are asked to generalize the velocity model of Section 5.3 of the 'Probabilistic Robotics' book [2] to a holonomic robot operating in the plane. Assume the robot can control its forward velocity, an orthogonal sidewards velocity, and a rotational velocity. Let us arbitrarily give sidewards motion to the left positive values, and sideward motions to the right negative values.

(a) State a mathematical model for such robot, assuming its controls are subject to independent Gaussian noise.

(b) Provide a procedure for calculating $p(x_t|u_t, x_{t-1})$.

(c) Provide a sampling procedure for sampling $x_t \approx p(x_t|u_t, x_{t-1})$.

Success!

Acknowledgements

One question is based on an assignment from Russell and Norvig's book "Artificial Intelligence - A Modern Approach"[1]. Another question is based on an assignment from the Albert-Ludwigs-Universität Freiburg, written by Wolfram Burgard. The last question is directly from the "Probabilistic Robotics" book [2].

References

[1] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, 3rd edition, 2009.

[2] S. Thrun, W. Burgard and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*, The MIT Press, September 2005, ISBN 0-262-20162-3.