

Avoiding the Wumpus

CS 580

Intro to Artificial Intelligence

Wumpus world

Problem setup


Environment

4x4 grid, containing

- 1 pile of gold
- Several pits
- Wumpus (carnivorous)

Goal

Reach the gold without falling in a pit or being eaten by the wumpus




Wumpus world simplified

For the moment,

- Multiple Wumpuses, but they do not move
- Gold is in (4,4), agent in (1,1)
- No pits
- We have a “wumpus adjacency sensor”
- We know the marginal probability that a wumpus is in any single cell

We can use these facts to build a model of the probability that a wumpus is any specific cell, and use this to decide which cells are safest to move to.



Setting up the probabilities

Define random variables

- Let W_{ij} be the RV for the event that a wumpus is in cell (i,j)
- Let D_{kl} be the RV for the event that a wumpus is detected to be **adjacent** to cell (k,l)

Define probabilities and independence relationships


Sensor is noise free, but we still can't measure wumpus direction directly

$$p(W_{ij} = \text{TRUE}) = 0.2, \quad W_{ij} \perp W_{kl} \text{ for } (i,j) \neq (k,l)$$

$$p(D_{ij} = \text{TRUE} \mid W_{ij+1} = \text{TRUE}) = p(d_{ij} \mid w_{i+1j}) = p(d_{ij} \mid w_{ij-1}) = p(d_{ij} \mid w_{i-1j}) = 1$$

$$p(d_{ij} \mid \neg w_{ij+1}, \neg w_{i+1j}, \neg w_{ij-1}, \neg w_{i-1j}) = 0$$

$$\underline{D_{ij}} \perp \underline{W_{mn}} \mid \underline{W_{kl}} \text{ where } (i,j) \text{ adj to } (k,l)$$



Setting up the query

“What is the probability that a wumpus is in (1,3) given no detection in (1,1) and detections in (2,1) and (1,2)?”


We'll repeat this query for (3,1) and (2,2).

Note: if $D_{11} = \text{False}$, we know $W_{11} = W_{21} = W_{12} = \text{False}$ because of the sensor model

Step 1: Write down query as a probability statement

Step 2: Put into joint distribution form

Step 3: Rearrange, marginalize, plug-in



Step 1: writing the query as a probability statement

“What is the probability a wumpus is in (1,3) given no detection in (1,1), and detections in both (1,2) and (2,1)?”

$$p(W_{13} = \text{TRUE} \mid D_{11} = \text{FALSE}, D_{21} = \text{TRUE}, D_{12} = \text{TRUE}, \\ W_{11} = \text{FALSE}, W_{21} = \text{FALSE}, W_{12} = \text{FALSE})$$

Abbreviated version

$$p(w_{13} \mid \underline{\neg d_{11}}, \underline{d_{21}}, \underline{d_{12}}, \underline{\neg w_{11}}, \underline{\neg w_{21}}, \underline{\neg w_{12}})$$
$$p(w_{13} \mid \underline{\text{sensors}}, \underline{\text{cleared}})$$

Let “sensors” and “cleared” stand in for the collection of RV settings we started with

Step 2: Joint distribution form

“What is the probability a wumpus is in (1,3) given no detection in (1,1), and detections in both (1,2) and (2,1)?”

$$p(w_{13} \mid \text{sensors, cleared}) = \alpha \sum p(w_{13}, \neg d_{11}, d_{12}, d_{21}, \underline{D_{13} = \hat{d}_1, \dots, D_{44} = \hat{d}_k}, \neg w_{11}, \neg w_{21}, \neg w_{12}, \underline{W_{31} = \hat{w}_1, \dots, W_{44} = \hat{w}_k})$$

$$p(w_{13} \mid \text{sensors, cleared}) = \alpha \sum_h p(w_{13}, \text{sensors, cleared}, \underline{\text{hidden} = h})$$

Here, we're using the definition of conditional probability, normalizing, and marginalizing over all the variables not mentioned in the query

Step 3: rearrange, marginalize, normalize, solve! (1)



$$p(w_{13} \mid \text{sensors}, \text{cleared}) =$$

$$\alpha \sum_h p(w_{13}, \text{sensors}, \underline{\text{cleared}}, \underline{\text{hidden}} = h)$$

If we split *hidden* into two pieces, *adjacent* and *far* we can leverage conditional independence to simplify

$$p(\text{sensors} \mid \underline{\text{cleared}}, \underline{\text{adjacent}}, \underline{\text{far}}) = p(\text{sensors} \mid \underline{\text{cleared}}, \underline{\text{adjacent}})$$

We're going to have to do some rearranging first to be able to use this fact



Step 3: rearrange, marginalize, normalize, solve! (2)

$$p(w_{13} \mid \text{sensors, cleared}) = \alpha \sum_h p(w_{13}, \text{sensors, cleared}, \text{hidden} = h)$$

Break *hidden* into
adj and *far*

$$= \alpha \sum_{h_1} \sum_{h_2} p(w_{13}, \text{sensors, cleared}, \text{adj} = h_1, \text{far} = h_2)$$

Use product rule

$$= \alpha \sum_{h_1} \sum_{h_2} [p(\text{sensors} \mid w_{13}, \text{cleared}, \text{adj} = h_1, \text{far} = h_2) \\ \cdot p(w_{13}, \text{cleared}, \text{adj} = h_1, \text{far} = h_2)]$$

Use conditional
independence of
the sensors

$$= \alpha \sum_{h_1} \sum_{h_2} \underbrace{p(\text{sensors} \mid w_{13}, \text{cleared}, \text{adj} = h_1)}_{\substack{\uparrow \\ \text{Note: the first term does not depend on } h_2}} \cdot p(w_{13}, \text{cleared}, \text{adj} = h_1, \text{far} = h_2)$$

Note: the first term does not depend on h_2

Step 3: rearrange, marginalize, normalize, solve! (3)

$$p(w_{13} \mid sensors, cleared) =$$

w_{13} , $cleared$, $adjacent$, and far are all W_{ij} random variables, which are independent!

$$\alpha \sum_{h_1} p(sensors \mid w_{13}, cleared, adj = h_1) \left[\sum_{h_2} p(w_{13}, cleared, adj = h_1, far = h_2) \right]$$

Don't involve h_1 or h_2

Doesn't involve h_2

$$p(w_{13} \mid sensors, cleared) =$$

$$\alpha \sum_{h_1} p(sensors \mid w_{13}, cleared, adj = h_1) \left[\sum_{h_2} p(w_{13}) \cdot p(cleared) \cdot p(adj = h_1) \cdot p(far = h_2) \right]$$

$$= \alpha p(w_{13}) \cdot p(cleared) \sum_{h_1} p(sensors \mid w_{13}, cleared, adj = h_1) \cdot p(adj = h_1) \sum_{h_2} p(far = h_2)$$

Step 3: rearrange, marginalize, normalize, solve! (4)

We can simplify further

$$p(w_{13} \mid \text{sensors}, \text{cleared}) =$$




$$\alpha p(w_{13}) \cdot \overbrace{p(\text{cleared})}^{\text{Combine } \alpha \text{ and } p(\text{cleared})} \sum_{h_1} \overbrace{p(\text{sensors} \mid w_{13}, \text{cleared}, \text{adj} = h_1) \cdot p(\text{adj} = h_1)}^{\text{Recall the sensor model. This is either 0 or 1 depending on the values for adj}} \sum_{h_2} \overbrace{p(\text{far} = h_2)}^{\text{Sums to 1}}$$

$$p(w_{13} \mid \text{sensors}, \text{cleared}) = \alpha' p(w_{13}) \sum_{h_1 \in \text{cons}(\text{sensors}, w_{13})} p(\text{adjacent} = h_1)$$

$$\text{cons}(\text{sensors}, w_{13}) = \{W_{ij} : \text{the wumpus could have generated } \text{sensors} \text{ (and } w_{13})\}$$

W_{ij} consistent with the sensors (1)

What settings for W_{ij} in *adjacent* are **consistent** with sensors $\neg d_{11}, d_{12}, d_{21}$ and w_{13} ?



$$p(\text{adj}) = p(w_{22}, w_{31}) = 0.2 * 0.2 = 0.04$$


$$p(\text{adj}) = p(w_{22}, \neg w_{31}) = 0.2 * 0.8 = 0.16$$

$$p(\text{adj}) = p(\neg w_{22}, w_{31}) = 0.8 * 0.2 = 0.16$$

$$p(w_{13}) = 0.2$$

W_{ij} consistent with the sensors

What settings for W_{ij} in *adjacent* are **consistent** with sensors $\neg d_{11}$, d_{12} , d_{21} and $\neg w_{13}$?



$$p(\text{adj}) = p(w_{22}, \neg w_{31}) = 0.2 * 0.8 = 0.16 \quad p(\text{adj}) = p(w_{22}, w_{31}) = 0.2 * 0.2 = 0.04$$

$$p(w_{13}) = 0.8$$

$$p(w_{13} \mid \text{sensors, cleared}) = \alpha' p(w_{13}) \sum_{h_1 \in \text{cons}(\text{sensors}, w_{13})} p(\text{adjacent} = h_1)$$

$$\text{cons}(\text{sensors}, w_{13}) = \{W_{ij} : \text{the wumpus could have generated sensors}\}$$

Solving for α

Plugging in,

$$p(w_{13} \mid \text{sense, cleared})$$

$$= \alpha p(w_{13}) [p(w_{22}, w_{31}) + p(\neg w_{22}, w_{31}) + p(w_{22}, \neg w_{31})]$$

$$= \alpha (0.2) [(0.04) + (0.16) + (0.16)] = \alpha 0.072$$

$$p(\neg w_{13} \mid \text{sense, cleared})$$



$$= \alpha p(\neg w_{13}) [p(w_{22}, \neg w_{31}) + p(w_{22}, w_{31})]$$

$$= \alpha (0.8) [(0.16) + (0.04)] = \alpha 0.16$$

$$\alpha (0.072 + 0.16) = 1, \alpha = 4.31\dots$$

$$p(w_{13} \mid \text{sense, cleared}) = 0.310$$

$$p(\neg w_{13} \mid \text{sense, cleared}) = 0.689$$



Solving for w_{31} and w_{22}

W_{31} is symmetric with W_{13} , so $p(w_{31}|\dots)=p(w_{13}|\dots)=0.31$



W_{22} is not symmetric, but has the same “consistent configurations” as before, just different probabilities

$$\begin{aligned} p(w_{22}|\text{sense,cleared}) &= \alpha p(w_{22}) [p(w_{13},w_{31}) + p(w_{13},\neg w_{31}) + p(\neg w_{13},w_{31}) + p(\neg w_{13},\neg w_{31})] \\ &= \alpha (0.2) [(0.04) + (0.16) + (0.16) + (0.64)] = \alpha 0.2 \end{aligned}$$

$$\begin{aligned} p(\neg w_{22}|\text{sense,cleared}) &= \alpha p(\neg w_{22}) [p(w_{13},w_{31})] \\ &= \alpha (0.8) [(0.04)] = \alpha 0.032 \end{aligned}$$

$$\alpha = 4.31, p(w_{22}|\text{sense,cleared}) = 0.86$$

So the agent should avoid (2,2) for either (3,1) or (1,3)



Summary and preview

Wrapping up

- We can apply all the rules of probability we've discussed to answer **queries** about the probability of specific **states** of the environment
- After developing a model, we can **compute** a desired probability with a simple three step approach
- Simplifying the model can help save significant computation overhead

Next time

- Graphically modeling independence structure using Bayes Nets