CSE 571 - Robotics
Homework 1 - Bayes Filters and Fully Connected Networks

Due Tuesday, April 20th @ 11:59pm

This homework involves three math problems, a programming assignment in Python, and, for extra credit, an ethical discussion on your work this week. The goal of the math problems is to help you understand conditional distributions and Bayesian inference. For the programming portion, you will be implementing a simple neural network in Python using NumPy without using any existing deep learning frameworks. Our aim is to teach you what these frameworks do, so you'll have a better understanding when it's time to actually use them. The code for this week's homework can be found on Github at `https://github.com/fishbotics/uw-robotics-571-sp21`

Useful reading material: Lecture notes, **Chapter 2 of Probabilistic Robotics, Thrun, Burgard and Fox (pdf shared with class)** and Chapter 2 of Gaussian Processes for Machine Learning, Rasmussen and Williams (Available online at: `http://www.gaussianprocess.org/gpml/chapters/`)

Collaboration: Students can discuss questions and work together to solve the problems, but each student MUST write up their own answers and code their own software. We will be checking both the code and PDFs for plagiarism. If we find that you have shared your solutions or code online, you will get an automatic 50% on the assignment. If we find that you have plagiarized solutions from another student or the internet, you will get an automatic 0% on this assignment.

Late Policy: You are allowed 6 late days for the entire quarter. After using these up, you will incur a penalty of 20% of the maximum grade per day. Please plan ahead!

1 Writing assignments [40 points]

1.1 Conditional Independence [5 points]

Let X,Y and Z be three random variables. Assuming that X and Y are conditionally independent given Z:

$$p(x, y|z) = p(x|z)p(y|z)$$

show that:

$$p(x|z) = p(x|z, y)$$
$$p(y|z) = p(y|z, x)$$

1.2 Bayes Filter [15 points]

A special-purpose robot is equipped with a vacuum unit to clean the floor. The robot has a binary sensor to detect whether a floor tile is clean or dirty. However, neither the cleaning unit nor the sensor are perfect. From previous experience, you know that the robot succeeds in cleaning a dirty floor tile with a probability of

$$\mathbb{P}(x_{t+1} = \text{clean}|x_t = \text{dirty}) = 0.6,$$

where x_t is the state of the floor tile at time t and x_{t+1} is the resulting state after the action has been applied. Activating the cleaning unit when the tile is clean will never make it dirty. Assume the robot always cleans at every time t (i.e. the transition probabilities model the fact that the robot is cleaning the floor tile).

The probability that the sensor indicates that the floor tile is clean although it is dirty is given by $p(z_t = \text{clean}|x_t = \text{dirty}) = 0.2$, and the probability that the sensor correctly detects a clean tile is given by $p(z_t = \text{clean}|x_t = \text{clean}) = 0.6$.

Unfortunately, you have no knowledge about the current state of the floor tile. However, after cleaning the tile, the robot's sensor indicates that it is clean. Compute $p(x_{t+1} = \text{clean}|z_{t+1} = \text{clean})$, *i.e.*, the probability that at time $t + 1$, the floor tile is now clean given that the sensor indicates it is clean. Assume a prior distribution at time t as $p(x_t = \text{clean}) = c$, where $0 \leq c \leq 1$. Then, using a software package of your choice, plot $p(x_{t+1} = \text{clean}|z_{t+1} = \text{clean})$ for $0 \leq c \leq 1$. You do not need to submit your code for plotting.

1.3 Gaussian Conditioning [20 points]

Let X and Y denote two scalar random variables that are jointly Gaussian:

$$\begin{aligned} p(x, y) &= \mathcal{N}(\mu, \Sigma) \\ &= \frac{1}{2\pi\sqrt{|\Sigma|}} \exp\left\{-\frac{1}{2}\left(\begin{bmatrix}x\\y\end{bmatrix} - \begin{bmatrix}\mu_X\\\mu_Y\end{bmatrix}\right)^\intercal \Sigma^{-1}\left(\begin{bmatrix}x\\y\end{bmatrix} - \begin{bmatrix}\mu_X\\\mu_Y\end{bmatrix}\right)\right\}, \end{aligned}$$

where $\mu = \begin{bmatrix}\mu_X & \mu_Y\end{bmatrix}^\intercal$ and $\Sigma = \begin{bmatrix}\sigma_X^2 & \sigma_{XY}\\ \sigma_{XY} & \sigma_Y^2\end{bmatrix}$ are the mean and covariance, respectively. Show that conditioning on Y results in a Gaussian over X:

$$\begin{aligned} p(x \mid y) &= \mathcal{N}(\mu_{X|Y}, \sigma_{X|Y}^2) \\ &= \frac{1}{\sqrt{2\pi}\sigma_{X|Y}} \exp\left\{-\frac{1}{2}\frac{(x - \mu_{X|Y})^2}{\sigma_{X|Y}^2}\right\} \end{aligned}$$

with $\mu_{X|Y} = \mu_X + \frac{\sigma_{XY}}{\sigma_Y^2}(y - \mu_Y)$ and $\sigma_{X|Y}^2 = \sigma_X^2 - \frac{\sigma_{XY}^2}{\sigma_Y^2}$.

1.3.1 Hints

- Use the definition of the conditional distribution and "complete the square" to get the answer

- Given $p(x, y)$ is jointly Gaussian, the marginal distribution of Y is also Gaussian: $p(y) = \mathcal{N}(\mu_Y, \sigma_Y^2)$
- For a 2x2 matrix positive definite matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, $A^{-1} = \frac{1}{|A|} \begin{bmatrix} d & -c \\ -b & a \end{bmatrix}$, $|A| = ad - bc$

2 Programming problems [60 points]

See `https://github.com/fishbotics/uw-robotics-571-sp21` for the programming assignment.

3 Extra credit [20 points]

Robotics, AI, and computer vision have the power to reshape our world and have the potential for both great outcomes and terrible consequences. As researchers, it is important to consider the potential outcomes of our work. The goal of this extra credit is to encourage you to think critically about the power of the technical material in this course. Grades for this section will be based on the quality of your argument and the effort you put into your answers–not on your specific opinions. Because these topics are inherently controversial, these answers will not be published anywhere without individual permission and will only be viewed by your instructors.

1. In the programming section, you implemented a neural network that does simple image detection. Describe why such a network may have an "opinion" (as in, is not perfectly fair) and in a paragraph or two, describe an incident in the real-world where a neural network's partiality had negative consequences (for example, social, political, or economic consequences) and discuss how these negative outcomes could have been mitigated.
2. Who is currently responsible for preventing bias in AI (for example, researchers, business leaders, politicians, etc)? Why? What is a downside of this current system? In a paragraph or two, describe who you think should be responsible and how would you design a fair and efficient system?

4 Submission

You will be using Gradescope `https://www.gradescope.com/` to submit the homework. Please submit the written assignment answers as a PDF. For the code, use the submit script in the homework directory and submit the compressed directory that it produces.