Random variable - takes on values w/ certain probabilities

ex. Coin Flip

Random variable X

$$X = \begin{cases} \text{Heads} & \text{wp } 0.5 \\ \text{Tails} & \text{wp } 0.5 \end{cases}$$

ex. $X \sim$ Uniform $(0, 1)$

$\rightarrow X$ will take on values btwn. 0 & 1 w/ equal prob.

Expected Value - weighted mean of the random variable

$X \sim$ Uniform $(0, 1)$

$\mathbb{E}[X] = 0.5$

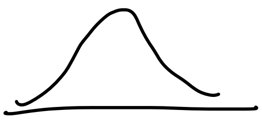

$Y \sim \text{Normal}(\underset{\text{mean}}{0}, \underset{\text{variance}}{1})$

$\mathbb{E}\{Y\} = 0$

- Linearity of Expectation

$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$

$$\mathbb{E}[cX] = c\,\mathbb{E}[X]$$

Indicator Random Variables

- Combination of events
- Break them down

N coin flips

$X_i \quad i \in [1, N] \longrightarrow$ Indicator for result of each coin flips

C106B Discussion 7: Probability and CV

1 Introduction

In anticipation for the upcoming unit on environment feedback, localization, and mapping, today we'll talk about:

1. Linearity of Expectation
2. Multivariate Random Variables
3. Hidden Markov Models
4. Low-Level Computer Vision Review

2 Linearity of Expectation

Random variables are commonly used to perform probabilistic calculations. They take on values corresponding to some distribution. The average value of some random variable X is known as the *expectation* of X.

Linearity of expectation allows us to compute the average value of a combination of multiple random variables. The theorem states:

$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$
$$\mathbb{E}[cX] = c\mathbb{E}[X]$$

Problem 1: Suppose I toss all 10 (unique) pairs of my socks in the washing machine, but when I collect them from my dryer, I only have 16 socks remaining. In expectation, how many pairs can I expect to see?

$X_i \rightarrow$ the odds that sock pair i survives
$i \in 1 \dots 10$

$$X_i = \begin{cases} 1 & \text{if sock pair } i \text{ survives} \\ 0 & \text{if it's eaten up} \end{cases}$$

$$P[X_i = 1] = \frac{\cancel{16}^{\,4}}{\cancel{20}_{\,\cancel{8}}} \times \frac{\cancel{15}^{\,3}}{19} = \frac{12}{19}$$

$$\mathbb{E}[X_1 + X_2 + \dots + X_{10}] = \text{expected number of surviving pairs}$$
$$= 10 \times \frac{12}{19} \approx \boxed{6.3 \text{ pairs}}$$

- Multivariate R.V. (random vectors)
  - Puts information about our system into one vector

$$X = \begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix}$$

<u>ex.</u> x, y, z location of our robot

- CDF : Cumulative Distribution Function

$$\underset{\substack{\uparrow \\ \text{R.V.}}}{F_X}\Big(\underset{\substack{\uparrow \\ \text{value}}}{x}\Big) = P(X \le x)$$

- Variance: squared avg distance from mean

$$Var(X) = \sigma^2(X) = \frac{\sum (x_i - \mu)^2}{N} = \mathbb{E}[X^2] - (\mathbb{E}[X])^2$$

- Covariance: extent to which they correspond in value

$$cov(X, Y) = \mathbb{E}\big[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])\big]$$

$$= \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$$

3 Multivariate Random Variables

A *multivariate random variable*, also known as a *random vector*, can be thought of as a group of random variables that are associated with one another in a single mathematical system. For example, a robot's predicted (x, y, z) location might be the output of some probabilistic function of observed environment variables, and the coordinates are grouped together to represent the state.

3.1 Mean, Covariance, and Cross-Covariance

The CDF takes a vector as input with the number of entries corresponding to the number of random variables:

$$F_X(x) = P(X_1 \leq x_1, ..., X_n < x_n)$$

As a result, the mean, or expected value, of a multivariate random variable is a vector as well:

$$\mathbb{E}[X] = (\mathbb{E}[X_1], ..., \mathbb{E}[X_n])^T$$

The *variance* of a single random variable equals the average distance from the mean of the values that it can take. It can be computed in two different ways:

$$Var(X) = \sigma^2(X) = \frac{\Sigma(x_i - \bar{x})^2}{N} = \mathbb{E}[X^2] - \mathbb{E}[X]^2$$

The *covariance* of multiple random variables represents the extent to which they correspond in values. If they both increase and decrease together, this value will be positive; whereas if they move in opposite directions, the value will be negative. Independent random variables will have 0 covariance (although the converse does not hold true).

$$cov(X, Y) = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])] = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$$

A single multivariate random variable will have an associated covariance matrix to represent pairwise covariance. Covariance matrices are symmetric positive semidefinite. A cross-covariance matrix can be calculated to represent the covariance between two different multivariate random variables. The element in the i, j position represents the covariance between the i-th value in the first vector and the j-th value in the second.

$$cov(\mathbf{X}, \mathbf{Y}) = \mathbb{E}[(\mathbf{X} - \mathbb{E}[\mathbf{X}])(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])^T] = \mathbb{E}[\mathbf{X}\mathbf{Y}^T] - \mathbb{E}[\mathbf{X}]\mathbb{E}[\mathbf{Y}]^T$$

Problem 2: Interpret the following covariance matrix. Which of the values is/are invalid?

$$\begin{bmatrix} 3 & 5 & 0 \\ 5 & 12 & -7 \\ 6 & -7 & -4 \end{bmatrix}$$

$$X = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \end{bmatrix}$$

Covariance matrix should be symmetric

$cov(X, Y) = cov(Y, X)$

Variances should be ≥ 0

$Var(X_3) \geq 0$

-4 is invalid

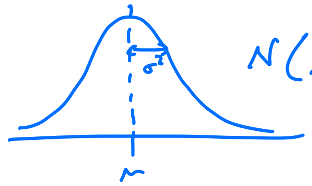

$N(\mu, \sigma^2)$

3.2 Multivariate Gaussians

A multivariate Gaussian random variable, or normal distribution, follows the following distribution:

$$p(x; \mu, \Sigma) = \frac{1}{(2\pi)^{(\frac{n}{2})}|\Sigma|^{(\frac{n}{2})}} exp(-\frac{1}{2}(x-\mu)^T\Sigma^{-1}(x-\mu))$$

means

cov matrix

The sum of Gaussians is a Gaussian:

$$X \sim N(\mu_X, \sigma_X^2)$$
$$Y \sim N(\mu_Y, \sigma_Y^2)$$
$$Z = X + Y$$
$$Z \sim N(\mu_X + \mu_Y, \sigma_X^2 + \sigma_Y^2)$$

The PDF of a multivariate Gaussian with a diagonal covariance matrix will be the same as that of n independent Gaussians (uncorrelated implies independence).

Problem 3: The isocontours of a function $f : \mathbf{R}^2 \longrightarrow \mathbf{R}$ are of the form $x \in \mathbf{R}^2 : f(x) = c$. Find the isocontours of a multivariate Gaussian, both with and without a diagonal covariance matrix. What kind of intuition does this give you about Gaussians?

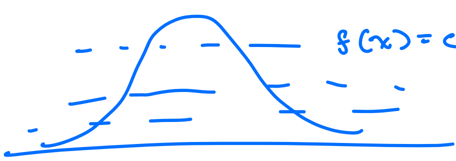

$f(x) = c$

$$\Sigma = \begin{bmatrix} 25 & 0 \\ 0 & 9 \end{bmatrix}$$

$$\mu = \begin{bmatrix} 3 \\ 2 \end{bmatrix}$$

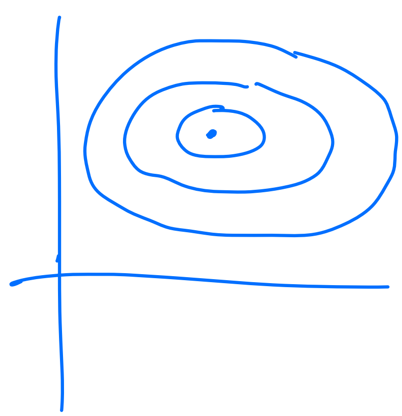

* Multivariate Gaussians have singular means

– Allows us to estimate multivariate R.V.s well

$$\Sigma = \begin{bmatrix} 10 & 5 \\ 5 & 5 \end{bmatrix}$$

$$\mu = \begin{bmatrix} 3 \\ 2 \end{bmatrix}$$

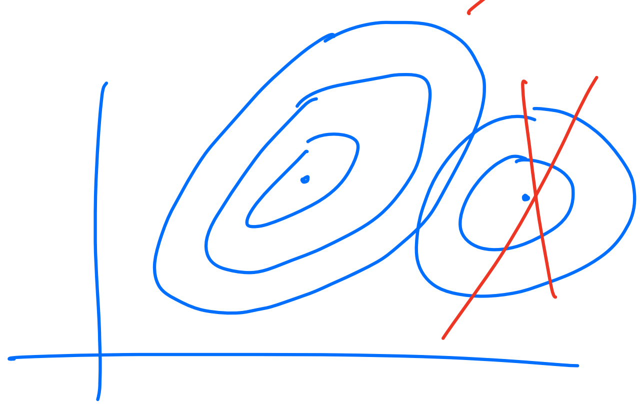

$S_1 \to S_2 \to S_3 \to \cdots$

- Next state depends <u>only</u> on current state

  * Markov property

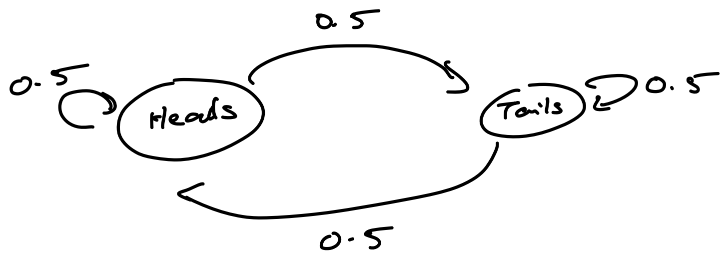

Evolution of robot states

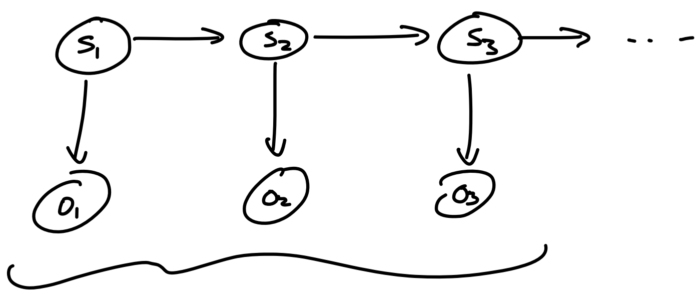

While we can't observe state, we <u>can</u> see observations

↳ <u>ex.</u> sensor readings

We use observations to measure state

4 Hidden Markov Models

A *Markov Chain* possesses the Markov property - the next state depends only on the current state and is entirely independent of the past. For example, a coin toss can be expressed as a Markov Model:

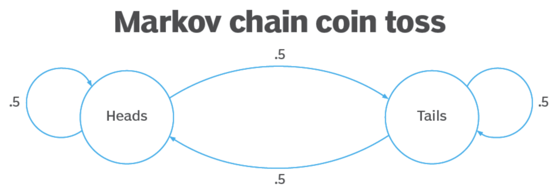

A *Hidden Markov Model* is often used in systems where data is continuously fed in over time. HMMs assume that states themselves have the Markov property. However, they are unknown - *observations* are used to form a probabilistic distribution of your current state. The observations could be sensor readings, for example, and the state might be the (unknown) location of your robot in the space. This is quite useful for localization!

Problem 4: Suppose we observe $O_1 = a$ and $O_2 = b$. Compute the probability distribution $P(W_2|O_1 = a, O_2 = b)$. [Source: CS 188 Fa22 Discussion 9]

↳ second state given observations

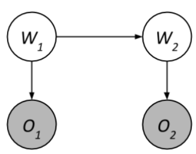

W_1	$P(W_1)$
0	0.3
1	0.7

prob. dist. of our start state

W_t	W_{t+1}	$P(W_{t+1}\|W_t)$
0	0	0.4
0	1	0.6
1	0	0.8
1	1	0.2

Transition probabilities

W_t	O_t	$P(O_t\|W_t)$
0	a	0.9
0	b	0.1
1	a	0.5
1	b	0.5

Observation probabilities conditioned on state

1. Joint probability 1st state & observation
2. Transition to the next state (prob. of second state w/ only O_1)
3. Update (add info from O_2)
4. Normalize

Start: $\rightarrow$ End:

W_1	$P(W_1)$
0	0.3
1	0.7

W_2	$P(W_2)$
0	0.25
1	0.75

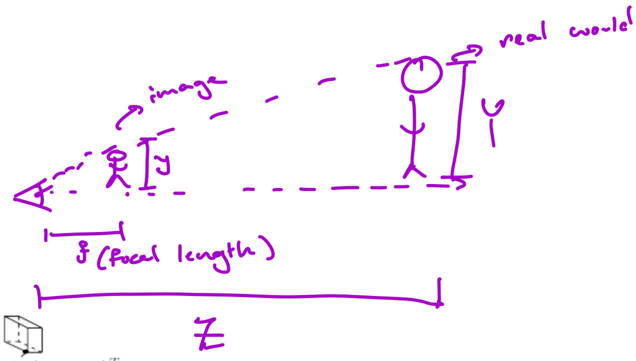

5 Computer Vision

5.1 The Pinhole Camera Model

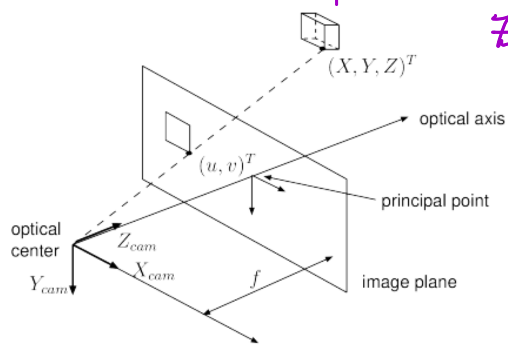

Problem 5: Using the image above, find the relationship between the 3D point $(X, Y, Z)^T$ to its corresponding 2D projection (u, v) onto the imaging plane (assume the focal length is 1).

Similar triangles: $\frac{y}{f} = \frac{Y}{Z}$ $\quad y = f\left(\frac{Y}{Z}\right)$

5.2 Homography

The pinhole camera model has a particular center of projection, the point from which we view the world. The image plane is offset a certain focal length in some direction from that point. If we rotate the camera without moving it around, we maintain the same center of projection - we're just looking a different way! This is known as a *homography transformation* and can be thought of as a series of unprojection, rotation, and then reprojection. It's quite a useful function to have, especially when your robot is moving around. The homography transformation can be used to straighten images - if we've taken a picture with the camera pointed sideways, we can rotate it so that it looks as though the camera is pointed straight!

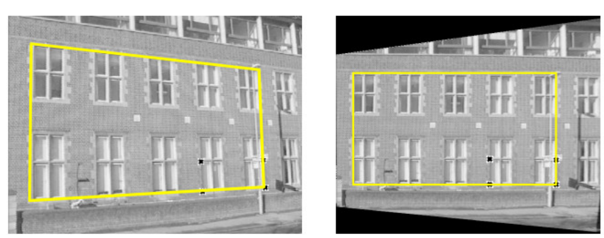

from Hartley & Zisserman

Problem 6: Let p correspond to a point on one image and let p' correspond to the same point in the scene, but projected onto another image. Write a general equation for how a homography matrix H maps points from one image to another. How would H be restricted if it must describe an affine transformation?

$p' = Hp$

Problem 7: How can you compute a homography matrix with real-world points?

– Perform feature mapping
 - Identify matching points in each picture
 - Perform least squares → most likely homography matrix