

“A beginning is the time for taking the most delicate care that the balances are correct.”

Frank Herbert, *Dune*



CMPUT 365

Background review

Office hours

- Slack and eClass: Asynchronous
- Marlos: Thursday 15:00 - 16:45 in ATH 3-08 (Athabasca Hall, 3-08)
- Anna: Monday 11:00 - 13:00 in CSC 3-50
- Bryan: Wednesday 14:00 - 16:00 in CAB 3-13
- David: Tuesday 13:00 - 15:00 in CSC 3-50
- Gabor: Wednesday 9:15-11:15 in CAB 3-13
- Marcos: Friday 10:00 - 12:00 in CAB 3-13

Syllabus [[eClass](#), [Slack](#), [website](#), [Google Drive](#)]

Plan

- Probability & statistics
- Linear algebra
- Calculus

Please, interrupt me at any time!



Probability and statistics continued

Conditional probabilities with more than 2 variables

The probability of an event A given another event B is defined as:

$$\Pr(A | B) \doteq \frac{\Pr(A \cap B)}{\Pr(B)}.$$

Chain rule:

$$\Pr(A \cap B) = \Pr(A, B) = \Pr(A | B) \Pr(B)$$

What's $\Pr(A, B | C)$?


Let $D = A \cap B$. Then, $\Pr(D | C) = \Pr(D, C) / \Pr(C)$. Thus $\Pr(A, B | C) = \Pr(A, B, C) / \Pr(C)$.

Now, let $E = B \cap C$, and recall, by the chain rule, that $\Pr(A, E) = \Pr(A | E) \Pr(E)$.

We then have $\Pr(A, B, C) = \Pr(A | B, C) \Pr(B, C) = \Pr(A | B, C) \Pr(B | C) \Pr(C)$.

Putting these two together: $\Pr(A, B | C) = \Pr(A | B, C) \Pr(B | C) \Pr(C) / \Pr(C)$.

Assuming $\Pr(C) \neq 0$, $\Pr(A, B | C) = \Pr(A | B, C) \Pr(B | C)$.



Example – Probabilities with two random variables

Let X be the random variable denoting the outcome of the roll of a d20, and let Y be the random variable denoting the outcome of the roll of a d6. What's $\Pr(X + Y \geq 25)$?

$$\begin{aligned}
 \Pr(X + Y \geq 25) &= \Pr([X = 20] \cap [Y = 5]) + \Pr([X = 20] \cap [Y = 6]) + \Pr([X = 19] \cap [Y = 6]) \\
 &= \Pr(X = 20) \Pr(Y = 5) + \Pr(X = 20) \Pr(Y = 6) + \Pr(X = 19) \Pr(Y = 6) \\
 &= 1/20 \times 1/6 + 1/20 \times 1/6 + 1/20 \times 1/6 \\
 &= 1/120 + 1/120 + 1/120 \\
 &= 3/120 \\
 &= 1/40
 \end{aligned}$$


Marginalization

- The marginal probability is the probability of a single event occurring, independent of other events.
- If we have the joint distribution $\Pr(x, y)$, we can find the marginals $\Pr(x)$ and $\Pr(y)$.

$$\Pr(X = x) = \sum_{y \in Y} \Pr(X = x, Y = y)$$

$$\Pr(Y = y) = \sum_{x \in X} \Pr(X = x, Y = y)$$

		Animal's favourite activity				
		Sleep	Play			
Type of pet	Cat	0.3	0.2	$\Pr(\text{Sleep})$	$= \Pr(\text{Sleep, Cat}) + \Pr(\text{Sleep, Dog})$	$= 0.3 + 0.1 = 0.4$
	Dog	0.1	0.4	$\Pr(\text{Play})$	$= \Pr(\text{Play, Cat}) + \Pr(\text{Play, Dog})$	$= 0.2 + 0.4 = 0.6$
				$\Pr(\text{Cat})$	$= \Pr(\text{Sleep, Cat}) + \Pr(\text{Play, Cat})$	$= 0.3 + 0.2 = 0.5$
				$\Pr(\text{Dog})$	$= \Pr(\text{Sleep, Dog}) + \Pr(\text{Play, Dog})$	$= 0.1 + 0.4 = 0.5$



Expectations

The expectation of a numeric random variable is the weighted average of its possible numeric outcomes, where the weights are the prob. of the outcome occurring:

$$\mathbb{E}[Y] \doteq \sum_{y \in Y} y \mathbf{Pr}(Y = y).$$

Example



$$\begin{aligned} \mathbb{E}[Y] &= \frac{1}{4} \times 1 + \frac{1}{4} \times 2 + \frac{1}{4} \times 3 + \frac{1}{4} \times 4 \\ &= 10/4 = 2.5. \end{aligned}$$

We can also compute the expectation of a function of a random variable:

$$\mathbb{E}[f(Y)] \doteq \sum_{y \in Y} f(y) \mathbf{Pr}(Y = y).$$

Properties of expectations

The expectation of a numeric random variable is the weighted average of its possible numeric outcomes, where the weights are the prob. of the outcome occurring:

$$\mathbb{E}[Y] \doteq \sum_{y \in Y} y \mathbf{Pr}(Y = y).$$

- $\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$, if X and Y are independent in nature.
- $\mathbb{E}[X Y] = \mathbb{E}[X] \mathbb{E}[Y]$, if X and Y are independent in nature.
- $\mathbb{E}[X + c] = \mathbb{E}[X] + c$, where c is not a random variable of the model.
- $\mathbb{E}[cX] = c \mathbb{E}[X]$, where c is not a random variable of the model.

Conditional expectations

Law of total expectation:

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

$$\mathbb{E}[X] = \sum_{y \in Y} \mathbb{E}[X | Y = y] \Pr(Y = y)$$

A conditional expectation of a random variable is the expected value of the variable given that an event is already known to have happened.

$$\mathbb{E}[X | Y = y] \doteq \sum_{x \in X} x \Pr(X = x | Y = y).$$

Example



Consider a D&D player who needs to roll 16 or higher to hit the target. When they hit the target, they cause 1d8 of damage. What's the expected damage this player will cause during such a battle?

Let X be the random variable denoting the 1d8 damage roll, and Y be the r.v. denoting the d20 roll.

$$\begin{aligned} \mathbb{E}[X | Y < 15] &= 0 & \mathbb{E}[X | Y = 16] &= 1 \Pr(X = 1 | Y = 16) + \dots + 8 \Pr(X = 8 | Y = 16) = 36/8 = 4.5 \\ \mathbb{E}[X | Y > 16] &= 4.5 & \mathbb{E}[X] &= 15/20 \times 0 + 5/20 \times 4.5 = \frac{3}{4} \times 0 + \frac{1}{4} \times 4.5 = 1.125. \end{aligned}$$



Linear algebra

Vectors and matrices

A vector can be thought as a list of numbers.

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}$$

A matrix can be thought as a table of numbers.

$$\mathbf{M} = \begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix}$$

Products

A vector is a matrix with one of its dimensions 1. Same rules apply.

A dot product between two vectors, \mathbf{v} and $\mathbf{w} \in \mathbb{R}^d$, is defined as:

$$\langle \mathbf{v}, \mathbf{w} \rangle = \mathbf{v}^\top \mathbf{w} = \sum_i v_i w_i.$$

A product between a matrix $\mathbf{M} \in \mathbb{R}^{n \times d}$ and a matrix $\mathbf{P} \in \mathbb{R}^{d \times p}$ is defined such that:

$$\mathbf{MP} = \mathbf{R},$$

$$\text{where } r_{ij} = m_{i1} p_{1j} + m_{i2} p_{2j} + \dots + m_{id} p_{dj} = \sum_{k=1}^d m_{ik} p_{kj}$$

$$\begin{bmatrix} m_{11} & m_{12} & m_{13} \\ p_{32} & & \\ m_{21} & m_{22} & m_{23} \\ p_{32} & & \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \\ p_{31} & p_{32} \end{bmatrix} = \begin{bmatrix} m_{11} p_{11} + m_{12} p_{21} + m_{13} p_{31} & m_{11} p_{12} + m_{12} p_{22} + m_{13} p_{32} \\ m_{21} p_{11} + m_{22} p_{21} + m_{23} p_{31} & m_{21} p_{12} + m_{22} p_{22} + m_{23} p_{32} \\ m_{31} p_{11} + m_{32} p_{21} + m_{33} p_{31} & m_{31} p_{12} + m_{32} p_{22} + m_{33} p_{32} \end{bmatrix}$$

Expectations in vector form

When dealing with more than one random variable, sometimes it is useful to use vector and/or matrix notation.

We can list n random variables into a vector $\mathbf{x} \in \mathbb{R}^n$, getting a vector of random variables. We call such a vector a *random vector*.

$$\mathbb{E}[\mathbf{x}] = \begin{bmatrix} \mathbb{E}[x_1] \\ \dots \\ \mathbb{E}[x_n] \end{bmatrix}$$


Several properties still apply, such as $\mathbb{E}[\mathbf{A} + \mathbf{B}] = \mathbb{E}[\mathbf{A}] + \mathbb{E}[\mathbf{B}]$.



Calculus


Derivatives

The derivative $df(a)/dx$ of a function f is the instantaneous rate of change of $y = f(x)$ with respect to x when $x = a$.




Derivatives

The derivative $df(a)/dx$ of a function f is the instantaneous rate of change of $y = f(x)$ with respect to x when $x = a$.



Derivatives

The derivative $df(a)/dx$ of a function f is the instantaneous rate of change of $y = f(x)$ with respect to x when $x = a$.




Derivatives

The derivative $df(a)/dx$ of a function f is the instantaneous rate of change of $y = f(x)$ with respect to x when $x = a$.

Useful property


The derivative of a function is zero at its local minima and its local maxima.



Useful property


We can sample from f and we can use its gradient to find a local minimum or a local maximum. That's stochastic gradient descent / ascent:

$$x' \leftarrow x \pm \alpha \nabla_x f(x).$$



Example – Stochastic gradient descent


Say we have a function $f(z) = z^2$, and we want to find the z that minimizes its value.



$$z' \leftarrow z \pm \alpha \nabla_z f(z)$$

Example – Stochastic gradient descent

Say we have a function $f(z) = z^2$, and we want to find the z that minimizes its value.




$$\frac{df(z)}{dz} =$$

$$z' \leftarrow z \pm \alpha \nabla_z f(z)$$

Example – Stochastic gradient descent

Say we have a function $f(z) = z^2$, and we want to find the z that minimizes its value.




$$\frac{df(z)}{dz} = 2z$$

$$z' \leftarrow z \pm \alpha \nabla_z f(z)$$

Example – Stochastic gradient descent (intuition)

Say we have a function $f(z) = z^2$, and we want to find the z that minimizes its value.



$$\frac{df(z)}{dz} = 2z$$

$$\alpha = 0.4$$

$$z' \leftarrow z \pm \alpha \nabla_z f(z)$$

$$\nabla f(4) = 2 \times 4 = 8$$

$$z' \leftarrow 4 - 0.4 \times 8$$

$$z' = 0.8$$

$$\nabla f(0.8) = 2 \times 0.8 = 1.6$$

$$z'' \leftarrow 0.8 - 0.4 \times 1.6$$

$$z'' = 0.16$$

The gradient vector

The gradient of f , denoted by ∇f , is a generalization of derivatives to a multi-dimensional function (the collection of all of its partial derivatives).


$$\nabla f(x_0, y_0, \dots) = \left[\frac{\partial f(x_0, y_0, \dots)}{\partial x}, \frac{\partial f(x_0, y_0, \dots)}{\partial y}, \dots \right]^T$$

Example

If $f(x, y) = x^2 + x \ln y$, which one is the right ∇f ?

- a. $\begin{bmatrix} 2x + \ln y \\ x/y \end{bmatrix}$ b. $\begin{bmatrix} 2x + x \ln y \\ x^2 + x/y \end{bmatrix}$

∇f outputs a vector with all possible partial derivatives of f .



Next class

Reminder: Practice quiz for Coursera's *Fundamentals of RL: Sequential decision-making* is due today at midnight.

- What **I** plan to do: Fundamentals of RL: An introduction to sequential decision-making (Bandits)
 - Discuss, more in depth, things related to bandits (Chapter 2 of the textbook).
- What I recommend **YOU** to do for next class:
 - Continue collecting (and posting) questions in eClass/Slack about Chapter 2 of the textbook.
 - Wrap up Coursera's week Fundamentals of RL: An introduction to sequential decision-making.
 - Submit programming assignment for Week 1 of Fundamentals of RL in Coursera.