

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

Frank Herbert, *Dune*



CMPUT 365 Background review

Reminder

You **should be enrolled in the private session** we created in Coursera for CMPUT 365.
Don't leave it to the last minute!

I **cannot** use marks from the public repository for your course marks.

You **need** to **check, every time**, if you are in the private session and if you are submitting quizzes and assignments to the private section.

The deadlines in the public session **do not align** with the deadlines in Coursera.

If you have any questions or concerns, **talk with the TAs** or email us
`cmput365@ualberta.ca`.

Reminder

You **should be enrolled in the private session** we created in Coursera for CMPUT 365.

Don't leave it to the last minute!

I **cannot** use marks from the public repository for your course marks.

You **need** to **check, every time**, if you are in the private session and if you are submitting quizzes and assignments to the private section.

The deadlines in the public session **do not align** with the deadlines in Coursera.

If you have any questions or concerns, **talk with the TAs** or email us

`cmput365@ualberta.ca`.

Notes

- The deadlines in Coursera and Canvas should now be aligned with the schedule
- Office hours are mostly settled:
 - Lucas Monday 16:00 – 18:00 @ UCOMM 2-138
 - Yuyang Tuesday 13:00 – 15:00 @ UCOMM 2-138
 - Shashank Wednesday 10:30 – 12:30 @ UCOMM 2-138
 - Aaron Wednesday 14:00 – 16:00 @ UCOMM
 - Siddarth Thursday 11:00 – 13:00 @ UCOMM 3-138
 - Dasha Thursday 13:00 – 15:00 @ UCOMM 2-138
 - Diego Friday 10:00 – 12:00 @ UCOMM 3-138
 - Parham Friday 15:00 – 17:00 @ UCOMM 2-138

Please, interrupt me at any time!



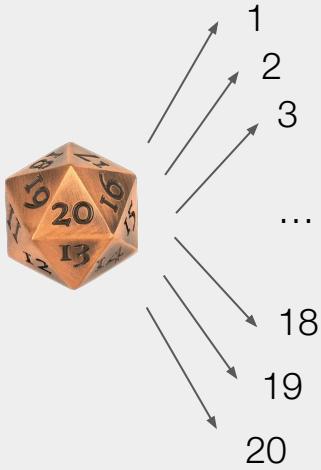
Probability

Probability – The basics

A probability is a function that associates a number between 0 and 1 to an event, with this number being a measure of the likelihood of that set of outcomes.

Example

Dungeons & Dragons!



$$\mathbf{Pr}(\text{rolling 20}) = 1/20 = 5\%$$

$$\begin{aligned}\mathbf{Pr}(\text{rolling 19 or 20}) &= \\ \mathbf{Pr}(\text{rolling 19}) + \mathbf{Pr}(\text{rolling 20}) &= \\ 1/20 + 1/20 &= 10\%\end{aligned}$$

$$\begin{aligned}\mathbf{Pr}(\text{rolling 20 and 20}) &= \\ \mathbf{Pr}(\text{rolling 20}) \times \mathbf{Pr}(\text{rolling 20}) &= \\ 1/20 \times 1/20 &= 1/400 = 0.25\%\end{aligned}$$

Probability – Somewhat more formally

*A probability is a function that associates a number between 0 and 1 to an event, with this number being a **measure of the likelihood** of that set of outcomes.*

- A probability distribution is defines how the probability is distributed among the outcomes.

Example



For an unbiased dice, each number is equally likely (i.e., uniform probability distribution). Thus, for each outcome $e \in S$, $\Pr(e) = 1/|S|$.

Probability – Somewhat more formally

*A probability is a function that associates a number between 0 and 1 to an event, with this number being a **measure of the likelihood** of that set of outcomes.*

- A probability distribution is defines how the probability is distributed among the outcomes.
- A way of calculating the probability of a specific event is a matter of identifying the sample space (set of all possible outcomes) and the probability distribution.

Example 1



For an unbiased dice, the probability of rolling a 20 is $\Pr(\text{rolling 20}) = 1/20$.

Example 2



For an unbiased dice, the probability of rolling higher than 18 is $\Pr(\text{rolling 19 or 20}) = 1/20 + 1/20 = 1/10$.

Probability – Properties

A probability is a function that associates a number between 0 and 1 to an event, with this number being a measure of the likelihood of that set of outcomes.

- Nonnegativity: $\mathbf{Pr}(A) \geq 0$.
- Normalization: $\sum_{e \in S} \mathbf{Pr}(e) = 1$.
- Additivity: $\mathbf{Pr}(A \cup B) = \mathbf{Pr}(A) + \mathbf{Pr}(B)$; $A \cap B = \{\}$.



Random variables and expectations

Random variables

Random variables are ways to map outcomes of random processes to real numbers.

They are not a traditional variable, nor random 😅

Capital letter!

Example 1

$$X = \left\{ \begin{array}{l} 1 \text{ if roll 1} \\ 2 \text{ if roll 2} \\ \dots \\ 19 \text{ if roll 19} \\ 20 \text{ if roll 20} \end{array} \right\}$$



Example 2

$$Y = \left\{ \begin{array}{l} 1 \text{ if heads} \\ 0 \text{ if tails} \end{array} \right\}$$



Example 3

$$Z = \left\{ \text{sum of 2 dice} \right\}$$



We can write $\Pr(X = 20)$ to represent $\Pr(\text{rolling 20})$.



We can write $\Pr(X \geq 19)$ to represent $\Pr(\text{rolling 19 or 20})$.



Examples

When rolling a d20 dice, let X be the random variable denoting the outcome of the roll.

$$\mathbf{Pr}(1 \leq X \leq 20) = 1$$

$$\mathbf{Pr}(X = 15) = 1/20$$

$$\mathbf{Pr}(X = 0) = 0$$

$$\mathbf{Pr}(2X = 1) = 0$$

Conditional probabilities

Chain rule:

$$\mathbf{Pr}(A \cap B) = \mathbf{Pr}(A, B) = \mathbf{Pr}(A \mid B) \mathbf{Pr}(B)$$

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

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

In a classroom with 100 students, out of those 100, 20 students play tabletop RPG, and 30 students have read *The Lord of the Rings* books. There are 15 students who play tabletop RPG who have read LOTR. What is the probability that a student has read LOTR given that the student plays tabletop RPG?



Let X be the random variable denoting the probability that a student plays tabletop RPG, and let Y be the random variable denoting the probability that a student has read LOTR.

$$\mathbf{Pr}(X) = 0.2$$

$$\mathbf{Pr}(Y) = 0.3$$

$$\mathbf{Pr}(X \cap Y) = 0.15$$

$$\mathbf{Pr}(Y \mid X) = 0.15/0.2 = 0.75$$

Conditional probabilities

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

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

When playing D&D, Tristan needs to roll 17 or higher on a d20 to successfully hit the troll. Tristan gets a critical hit when they roll a 20. Knowing that Tristan has successfully hit the target, what's the likelihood that Tristan got a critical hit?

Let X be the random variable denoting the number Tristan rolled on a d20, and Y a binary random variable denoting whether Tristan rolled a 20 ($Y=1$) or not ($Y=0$).

$$\mathbf{Pr}(X \geq 17) = 1/5 \quad \mathbf{Pr}(Y = 1 \cap X \geq 17) = 1/20$$

$$\frac{\mathbf{Pr}(Y = 1 \cap X \geq 17)}{\mathbf{Pr}(X \geq 17)} = \frac{1/20}{1/5} = \frac{5}{20} = 25\%$$



Independence

Two events are *independent* when the likelihood of an event does not change after knowing the other event. A is independent of B if and only if

$$\mathbf{Pr}(A \mid B) = \mathbf{Pr}(A).$$

$$\mathbf{Pr}(A \mid B) = \mathbf{Pr}(A \cap B) / \mathbf{Pr}(B)$$

$$\mathbf{Pr}(A \cap B) = \mathbf{Pr}(A \mid B)$$

$$\mathbf{Pr}(B)$$

$$\mathbf{Pr}(A \cap B) = \mathbf{Pr}(A) \mathbf{Pr}(B)$$

$$\begin{aligned} \mathbf{Pr}(B \mid A) &= \mathbf{Pr}(B \cap A) / \mathbf{Pr}(A) \\ &= \mathbf{Pr}(B) \mathbf{Pr}(A) / \mathbf{Pr}(A) \\ &= \mathbf{Pr}(B) \end{aligned}$$

Example



Tristan rolls two d20 dice. Given that they rolled a 1 on the first dice, what's the likelihood of them running a 20 on the second dice?

Let X be the random variable denoting the roll on the first dice, and Y be the equivalent for the second dice.

$$\mathbf{Pr}(X = 1) = 1/20 \quad \mathbf{Pr}(Y = 20) = 1/20 \quad \mathbf{Pr}(X = 1 \cap Y = 20) = 1/400$$

$$\mathbf{Pr}(Y = 20 \mid X = 1) = (1/400)/(1/20) = 1/20$$

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 $\Pr(X + Y \geq 25)$?

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

Conditional probabilities with more than 2 variables

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

$$\mathbf{Pr}(A \mid B) \doteq \frac{\mathbf{Pr}(A \cap B)}{\mathbf{Pr}(B)}$$

Chain rule:

$$\mathbf{Pr}(A \cap B) = \mathbf{Pr}(A, B) = \mathbf{Pr}(A \mid B) \mathbf{Pr}(B)$$

What's $\mathbf{Pr}(A, B \mid C)$?

Let $D = A \cap B$. Then, $\mathbf{Pr}(D \mid C) = \mathbf{Pr}(D, C) / \mathbf{Pr}(C)$. Thus $\mathbf{Pr}(A, B \mid C) = \mathbf{Pr}(A, B, C) / \mathbf{Pr}(C)$.

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

We then have $\mathbf{Pr}(A, B, C) = \mathbf{Pr}(A \mid B, C) \mathbf{Pr}(B, C) = \mathbf{Pr}(A \mid B, C) \mathbf{Pr}(B \mid C) \mathbf{Pr}(C)$.

Putting these two together: $\mathbf{Pr}(A, B \mid C) = \mathbf{Pr}(A \mid B, C) \mathbf{Pr}(B \mid C) \mathbf{Pr}(C) / \mathbf{Pr}(C)$.

Assuming $\mathbf{Pr}(C) \neq 0$, $\mathbf{Pr}(A, B \mid C) = \mathbf{Pr}(A \mid B, C) \mathbf{Pr}(B \mid C)$.

Marginalization

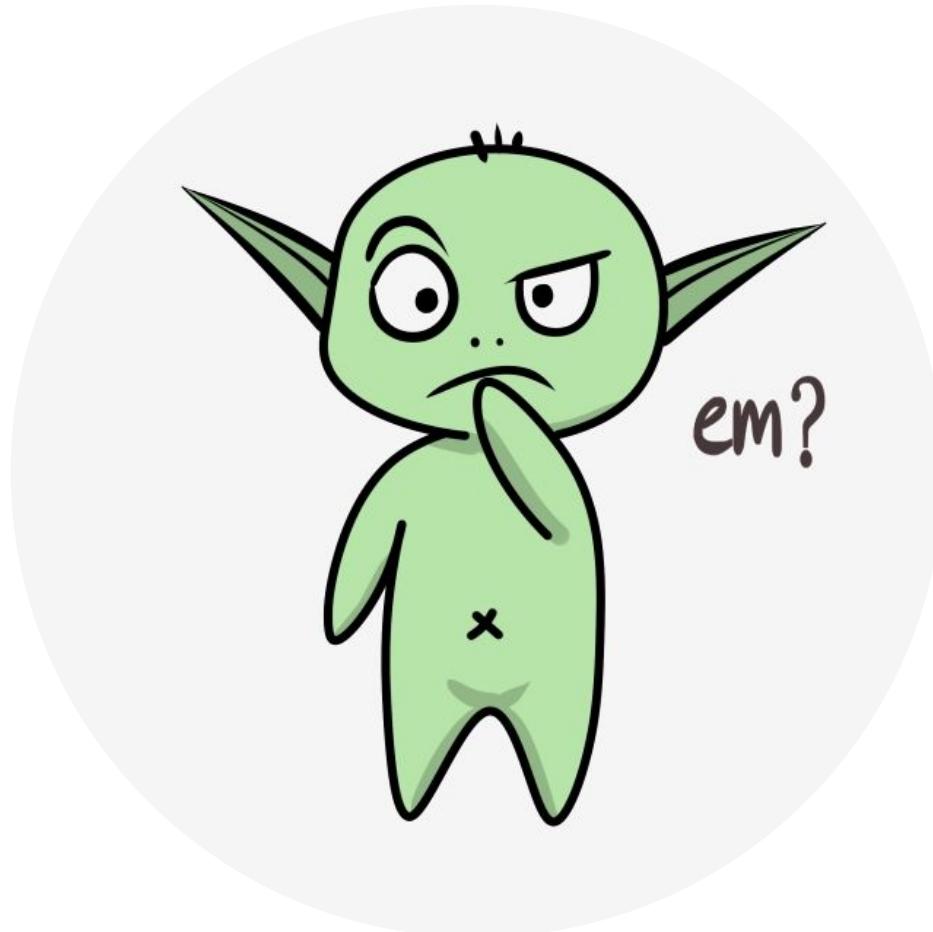
- The marginal probability is the probability of a single event occurring, independent of other events.
- If we have the joint distribution $\mathbf{Pr}(x, y)$, we can find the marginals $\mathbf{Pr}(x)$ and $\mathbf{Pr}(y)$.

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

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

Example

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



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:

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

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:

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

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?

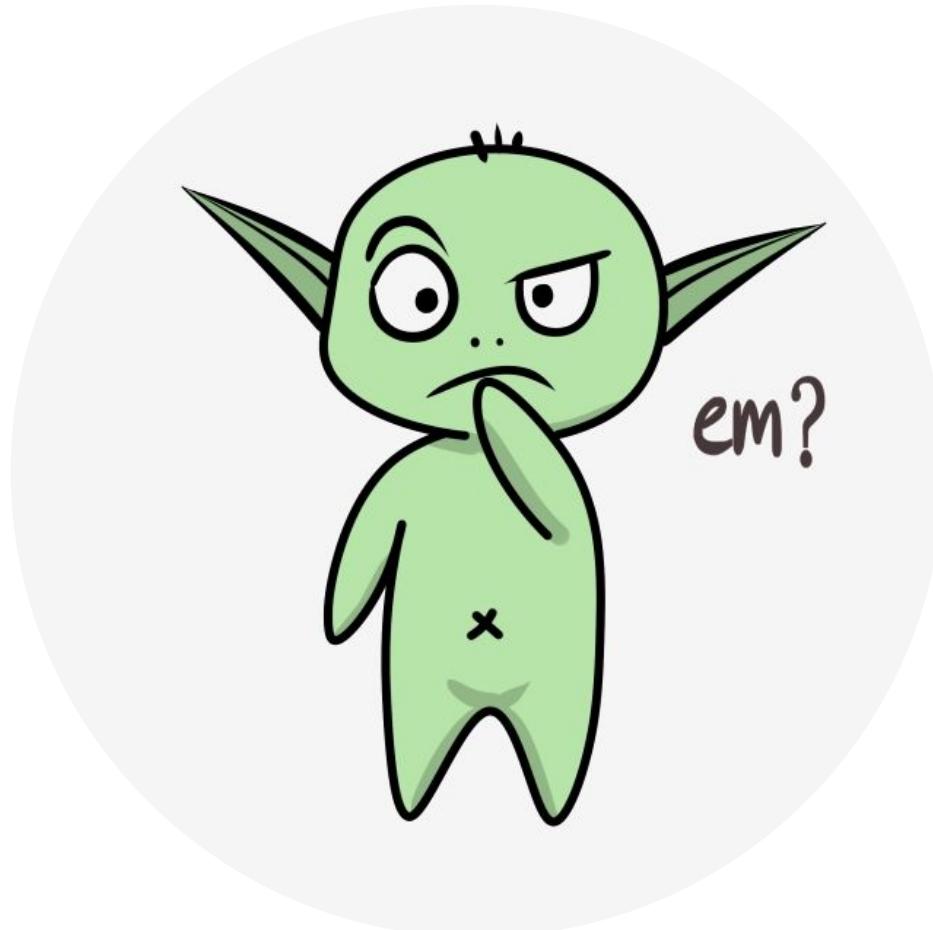
Let X be the random variable denoting the 1d8 damage roll, and Y be the r.v. denoting the d20 roll.

$$\mathbb{E}[X | Y < 15] = 0$$

$$\mathbb{E}[X | Y \geq 16] = 4.5$$

$$\mathbb{E}[X | Y = 16] = 1 \Pr(X = 1 | Y = 16) + \dots + 8 \Pr(X = 8 | Y = 16) = 36/8 = 4.5$$

$$\mathbb{E}[X] = 15/20 \times 0 + 5/20 \times 4.5 = \frac{3}{4} \times 0 + \frac{1}{4} \times 4.5 = 1.125.$$



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:

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

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

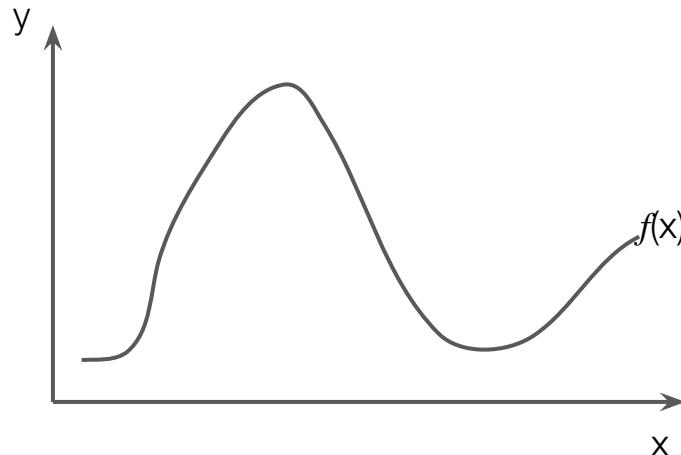
$$\begin{bmatrix} m_{11} & m_{12} & m_{13} \\ p_{32} \\ m_{21} & m_{22} & m_{23} \\ p_{32} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \\ p_{31} & p_{32} \end{bmatrix} = \begin{bmatrix} m_{11}p_{11} + m_{12}p_{21} + m_{13}p_{31} & m_{11}p_{12} + m_{12}p_{22}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



Calculus

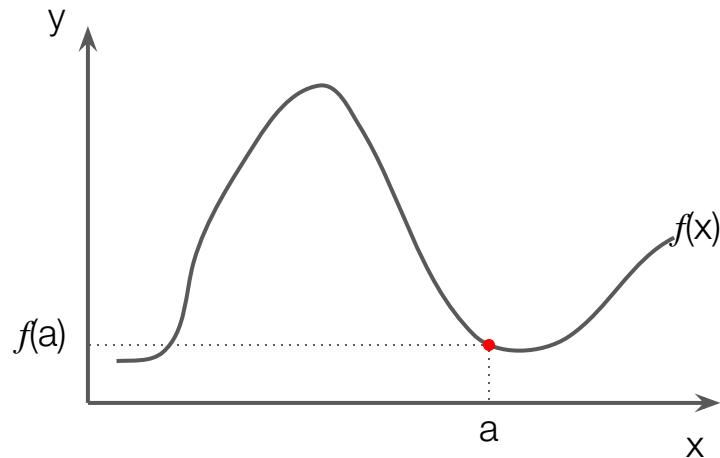
Derivatives

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



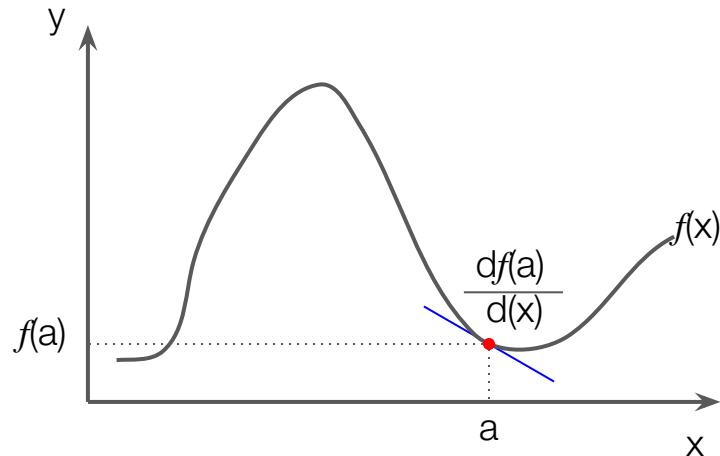
Derivatives

The derivative $d\mathit{f}(a)/dx$ 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 $x = a$.

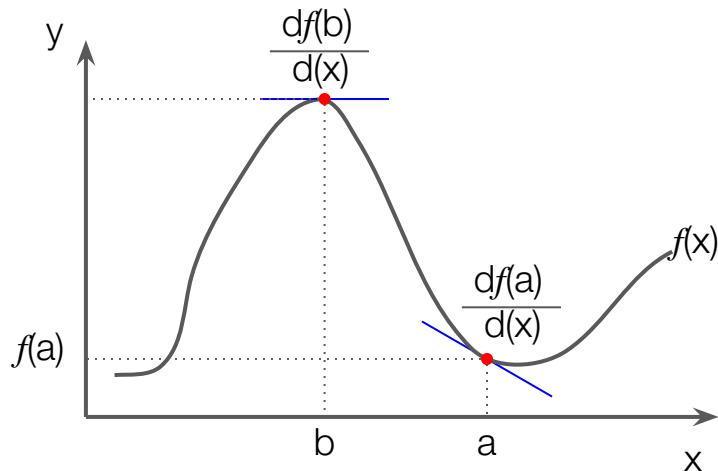


Derivatives

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:

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

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).

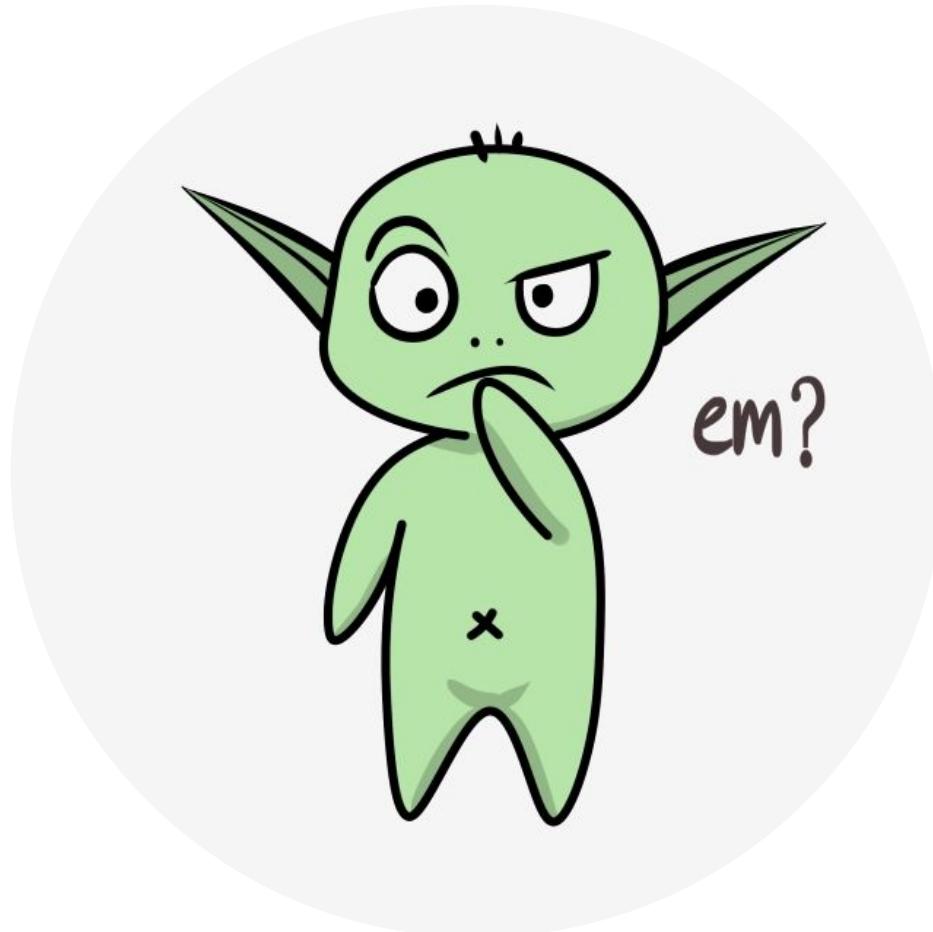
$$\nabla f(x_0, y_0, \dots) = \left[\frac{\partial f(x_0, y_0, \dots)}{\partial x}, \frac{\partial f(x_0, y_0, \dots)}{\partial y}, \dots \right]^T$$

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

- What **I** plan to do
 - Start Fundamentals of RL: An introduction to sequential decision-making (Bandits)
- What I recommend **YOU** to do for next class:
 - Finish the recommended reading for Coursera's M1W2.
 - Start/Submit practice quiz and programming assignment for Coursera's Fundamentals of RL: Sequential decision-making (M1 W2).