Bayesian Inference about Everything

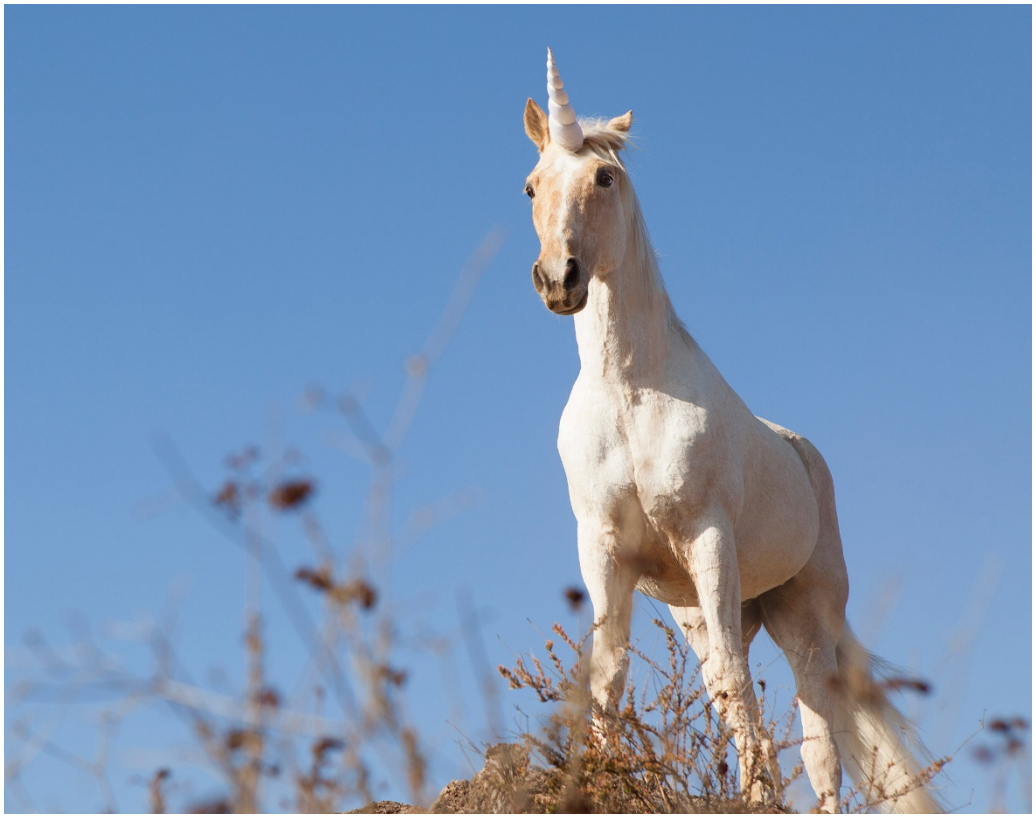

CSC411/2515: Machine Learning and Data Mining, Winter 2018

Michael Guerzhoy and Lisa Zhang

Model Everything

- What is the probability of me (ever) observing a horse that has a horn if unicorns exist?
  - $P(I|U) = 0.9$ (say)
- What is the probability of me (ever) observing a horse that has a horn if unicorns do not exist?
  - $P(I|\neg U) = 0.0001$ (say)
- How to compute those probabilities?
  - Model everything in the world: the probability that a quantum fluctuation would create an image of a unicorn on my retina (a physical process), the probability that someone would be able to play this prank on me (a sociological process), the probability that I an hallucinating a unicorn (a psychological process)…
  - In practice, just guess

Prior Beliefs about the World

- What is the probability that unicorns exist?
  - My *prior beliefs* about the world, before observing the data
  - $P(U) = 10^{-10}$
- Now, having observed a unicorn, what do I think about unicorns?

$$P(U|I) = \frac{P(I|U)P(U)}{P(I)} = \frac{P(I|U)P(U)}{P(I|U)P(U)+P(I|\neg U)P(\neg U)}$$

$$= \frac{0.9\times10^{-10}}{0.9\times10^{-10} + 0.0001\times(1-10^{-10})} = 9\times10^{-7}$$

Reminder: Unicorns and Stats

- We can encode out beliefs about what the values of the parameters could be using

$$P(\theta)$$

likelihood

Prior

- Using Bayes' rule, we have

$$P(\theta = \theta_0|\text{data}) = \frac{\text{P}(\theta=\theta_0, data)}{P(data)} = \frac{P(data|\theta = \theta_0)P(\theta=\theta_0)}{P(data)}$$

$$= \sum_{\theta_1} P(data|\, \theta = \theta_1)\, P(\theta = \theta_1)$$