CPEN 455 24W2 Tutorial

Probability and Statistics[1]

Qi Yan

University of British Columbia

January 13, 2025

[1] Based on materials from **goodfellow2016deep**; **zhang2023dive**; **pml1Book**.

Contents

Probability

Probability theory is nothing but common sense reduced to calculation. — Pierre Laplace, 1812

- Probability: quantitative degree of belief.
  - an image classifier outputs a probability distribution given an input image.
  - a large language model (e.g., chatGPT) outputs a probability distribution over the next word when making predictions.
  - a generative model (VAEs/GANs/diffusion) starts generation from known prior distributions.

PROBABILITY

Probability theory is nothing but common sense reduced to calculation. — Pierre Laplace, 1812

- Probability: quantitative degree of belief.
  - an image classifier outputs a probability distribution given an input image.
  - a large language model (e.g., chatGPT) outputs a probability distribution over the next word when making predictions.
  - a generative model (VAEs/GANs/diffusion) starts generation from known prior distributions.
- **Frequentist** perspective: probabilities represent long run frequencies of **events** that can happen multiple times.
  - if we a the coin many times, we expect it to land heads about half the time.
- **Bayesian** perspective: probability is used to quantify our **uncertainty** or ignorance about something; hence it is fundamentally related to information rather than repeated trials
  - we believe the coin is equally likely to land heads or tails on the next toss.

Random Variables

- **Random Variable**:
  - A variable with potential various random values.
  - Distinct from its possible values.

Random Variables

- **Random Variable**:
  - A variable with potential various random values.
  - Distinct from its possible values.
- **Nature of Random Variables**:
  - A description of possible states.
  - Requires a probability distribution to define likelihood of each state.

Random Variables

- **Random Variable**:
  - A variable with potential various random values.
  - Distinct from its possible values.
- **Nature of Random Variables**:
  - A description of possible states.
  - Requires a probability distribution to define likelihood of each state.
- **Types of Random Variables**:
  - *Discrete Random Variable*:
    - Finite or countably infinite states.
    - States could be numerical or non-numerical.
  - *Continuous Random Variable*:
    - Associated with real number values.

Discrete Random Variables

Described using a **Probability Mass Function (PMF)**, typically denoted by P. Associates each possible state with a probability.

PMF Characteristics

- Maps a state of a random variable to its probability of occurrence.
- PMFs can describe joint probabilities for multiple variables, e.g., $P(\mathrm{x}, \mathrm{y})$.
- Must satisfy two conditions:
  1. $0 \leq P(\mathrm{x}) \leq 1$ for all states x.
  2. $\sum_{\mathrm{x}} P(\mathrm{x}) = 1$ (Normalization).

Discrete Random Variables

Described using a **Probability Mass Function (PMF)**, typically denoted by P. Associates each possible state with a probability.

PMF Characteristics

- Maps a state of a random variable to its probability of occurrence.
- PMFs can describe joint probabilities for multiple variables, e.g., $P(\mathrm{x}, \mathrm{y})$.
- Must satisfy two conditions:
  1. $0 \leq P(\mathrm{x}) \leq 1$ for all states x.
  2. $\sum_{\mathrm{x}} P(\mathrm{x}) = 1$ (Normalization).

Example 2.1 (Uniform Distribution PMF)

Given a discrete random variable x with k states, a uniform distribution assigns:

$$P(\mathrm{x} = x_i) = \frac{1}{k}$$

This satisfies the normalization condition since $\sum_i \frac{1}{k} = 1$.

Continuous Random Variables

Probability Density Functions (PDFs) p describe probability distributions for continuous variables:

- Domain is all possible states of x.
- $p(\mathrm{x}) \geq 0$ for all x in the domain.
- $\int p(\mathrm{x})\, d\mathrm{x} = 1$ (Total probability is 1).

Understanding PDFs

- PDFs give the probability density, not the probability of specific states.
- Probability for an infinitesimal region is $p(\mathrm{x})\, d\mathrm{x}$.
- The probability that x is in a set S is the integral of $p(\mathrm{x})$ over S.

Continuous Random Variables

Probability Density Functions (PDFs) p describe probability distributions for continuous variables:

- Domain is all possible states of x.
- $p(\mathrm{x}) \geq 0$ for all x in the domain.
- $\int p(\mathrm{x})\, d\mathrm{x} = 1$ (Total probability is 1).

Understanding PDFs

- PDFs give the probability density, not the probability of specific states.
- Probability for an infinitesimal region is $p(\mathrm{x})\, d\mathrm{x}$.
- The probability that x is in a set S is the integral of $p(\mathrm{x})$ over S.

Example 2.2 (Uniform Distribution PDF)

Consider $u(\mathrm{x}; a, b)$ for a uniform distribution over $[a, b]$, where $a < b$:

$$u(\mathrm{x}; a, b) = \begin{cases} 0 & \text{for } \mathrm{x} \notin [a, b] \\ \frac{1}{b-a} & \text{for } \mathrm{x} \in [a, b] \end{cases}$$

This function is always nonnegative and integrates to 1, representing the uniform distribution $\mathrm{x} \sim U(a, b)$.

Probability Distributions

Marginal Probability

Concept of Marginal Probability

The probability distribution over a subset of variables, derived from a joint distribution of multiple variables.

Discrete Random Variables

Given discrete random variables x and y, and the joint distribution $P(\mathrm{x}, \mathrm{y})$, the marginal probability $P(\mathrm{x})$ is calculated as:

$$\forall \mathrm{x} \in X, P(\mathrm{x} = x) = \sum_{y \in Y} P(\mathrm{x} = x, \mathrm{y} = y)$$

Continuous Random Variables

For continuous variables, marginal probability is found using integration:

$$p(\mathrm{x}) = \int p(\mathrm{x}, \mathrm{y}) \, d\mathrm{y}$$

Origin of the Term

The term "marginal" refers to the practice of summing probabilities in a table and writing the totals in the margins.

Definition

The probability of an event given that another event has occurred, denoted as $P(\mathrm{y} = y \mid \mathrm{x} = x)$.

Computation Formula

$$P(\mathrm{y} = y \mid \mathrm{x} = x) = \frac{P(\mathrm{y} = y, \mathrm{x} = x)}{P(\mathrm{x} = x)}$$

Note: Defined only if $P(\mathrm{x} = x) > 0$.

Understanding Conditional Probability

- Not to be confused with the consequences of actions or interventions.
- The conditional probability of an event y given x is different from the probability that y would happen if x were to be caused by some action (called intervention query for causal modeling).

Chain Rule

A joint probability distribution can be decomposed into conditional probabilities:

$$P(\mathrm{x}^{(1)}, \ldots, \mathrm{x}^{(n)}) = P(\mathrm{x}^{(1)}) \prod_{i=2}^{n} P(\mathrm{x}^{(i)} | \mathrm{x}^{(1)}, \ldots, \mathrm{x}^{(i-1)})$$

Application of the Chain Rule

- Simplifies the computation of joint distributions.
- Derived from the definition of conditional probability.

Chain Rule

A joint probability distribution can be decomposed into conditional probabilities:

$$P(\mathrm{x}^{(1)}, \ldots, \mathrm{x}^{(n)}) = P(\mathrm{x}^{(1)}) \prod_{i=2}^{n} P(\mathrm{x}^{(i)} | \mathrm{x}^{(1)}, \ldots, \mathrm{x}^{(i-1)})$$

Application of the Chain Rule

- Simplifies the computation of joint distributions.
- Derived from the definition of conditional probability.

Example 3.1

Applying the rule to three variables a, b, and c:

$$\begin{aligned} P(\mathrm{a}, \mathrm{b}, \mathrm{c}) &= P(\mathrm{a}|\mathrm{b}, \mathrm{c}) P(\mathrm{b}, \mathrm{c}) \\ P(\mathrm{b}, \mathrm{c}) &= P(\mathrm{b}|\mathrm{c}) P(\mathrm{c}) \\ P(\mathrm{a}, \mathrm{b}, \mathrm{c}) &= P(\mathrm{a}|\mathrm{b}, \mathrm{c}) P(\mathrm{b}|\mathrm{c}) P(\mathrm{c}) \end{aligned}$$

Independence of Random Variables

Two random variables x and y are **independent** if:

$$\forall x \in X, y \in Y, \quad p(\mathrm{x}=x, \mathrm{y}=y) = p(\mathrm{x}=x)p(\mathrm{y}=y)$$

Conditional Independence

x and y are **conditionally independent** given z if:

$$\forall x \in X, y \in Y, z \in Z, \quad p(\mathrm{x}=x, \mathrm{y}=y|\mathrm{z}=z) = p(\mathrm{x}=x|\mathrm{z}=z)p(\mathrm{y}=y|\mathrm{z}=z)$$

Notation

Independence is denoted by $\mathrm{x} \perp \mathrm{y}$, while conditional independence given z is denoted by $\mathrm{x} \perp \mathrm{y}|\mathrm{z}$.

Expectation, Variance and Covariance

Expectation and Expected Value

Definition

The expectation or expected value of a function $f(\mathrm{x})$ with respect to a distribution $P(\mathrm{x})$ is the mean value f takes when x is drawn from P.

Computation

For discrete variables:

$$\mathbb{E}_{\mathrm{x}\sim P}[f(\mathrm{x})] = \sum_{\mathrm{x}} P(\mathrm{x})f(\mathrm{x}),$$

For continuous variables:

$$\mathbb{E}_{\mathrm{x}\sim P}[f(\mathrm{x})] = \int p(\mathrm{x})f(\mathrm{x})d\mathrm{x}.$$

Properties of Expectation

- Expectations are linear:

  $$\mathbb{E}_{\mathrm{x}}[\alpha f(\mathrm{x}) + \beta g(\mathrm{x})] = \alpha\mathbb{E}_{\mathrm{x}}[f(\mathrm{x})] + \beta\mathbb{E}_{\mathrm{x}}[g(\mathrm{x})],$$

  where α and β are constants.
- Notation can be simplified to $\mathbb{E}[f(\mathrm{x})]$ when the context is clear.

Variance

The measure of the spread of a function $f(\mathrm{x})$ of a random variable:

$$\mathrm{Var}(f(\mathrm{x})) = \mathbb{E}[(f(\mathrm{x}) - \mathbb{E}[f(\mathrm{x})])^2]$$

Standard deviation is the square root of the variance.

Covariance

Indicates how two variables $f(\mathrm{x})$ and $g(\mathrm{y})$ linearly relate to each other:

$$\mathrm{Cov}(f(\mathrm{x}), g(\mathrm{y})) = \mathbb{E}[(f(\mathrm{x}) - \mathbb{E}[f(\mathrm{x})])(g(\mathrm{y}) - \mathbb{E}[g(\mathrm{y})])]$$

Positive (negative) covariance implies a positive (negative) linear relationship.

Variance

The measure of the spread of a function $f(\mathrm{x})$ of a random variable:

$$\mathrm{Var}(f(\mathrm{x})) = \mathbb{E}[(f(\mathrm{x}) - \mathbb{E}[f(\mathrm{x})])^2]$$

Standard deviation is the square root of the variance.

Covariance

Indicates how two variables $f(\mathrm{x})$ and $g(\mathrm{y})$ linearly relate to each other:

$$\mathrm{Cov}(f(\mathrm{x}), g(\mathrm{y})) = \mathbb{E}[(f(\mathrm{x}) - \mathbb{E}[f(\mathrm{x})])(g(\mathrm{y}) - \mathbb{E}[g(\mathrm{y})])]$$

Positive (negative) covariance implies a positive (negative) linear relationship.

Correlation

Correlation normalizes covariance to measure the strength of the linear relationship (scale-invariant).

Independence v.s. Covariance

Independence implies zero covariance but not vice versa. Zero covariance does not imply independence.

Covariance Matrix

For a random vector $\mathbf{x} \in \mathbb{R}^n$, the covariance matrix is $n \times n$ with:

$$\mathrm{Cov}(\mathrm{x}_i, \mathrm{x}_j) = \mathrm{Var}(\mathrm{x}_i, \mathrm{x}_j)$$

Diagonal elements represent the variance.

A Bit of Information Theory

Entropy

Definition

Entropy, denoted as $H(\mathrm{X})$, is a measure of the uncertainty or unpredictability in a random variable X.

Shannon Entropy

For a discrete random variable X with possible values $\{x_1, x_2, \ldots, x_n\}$ and probability mass function $P(\mathrm{X})$, entropy is defined as:

$$H(\mathrm{X}) = -\mathbb{E}_{\mathrm{x}\sim P}[\log P(\mathrm{x})] = -\sum_{i=1}^{n} P(x_i) \log P(x_i)$$

where the logarithm is base 2 for bits.

Definition

Entropy, denoted as $H(\mathrm{X})$, is a measure of the uncertainty or unpredictability in a random variable X.

Shannon Entropy

For a discrete random variable X with possible values $\{x_1, x_2, \ldots, x_n\}$ and probability mass function $P(\mathrm{X})$, entropy is defined as:

$$H(\mathrm{X}) = -\mathbb{E}_{\mathrm{x}\sim P}[\log P(\mathrm{x})] = -\sum_{i=1}^{n} P(x_i) \log P(x_i)$$

where the logarithm is base 2 for bits.

Interpretation

Entropy quantifies the expected amount of information conveyed by identifying the outcome of X. Higher entropy implies a more uncertain outcome, while lower entropy implies a more predictable outcome.

Properties of Entropy

- $H(\mathrm{X}) \geq 0$ for all X.
- $H(\mathrm{X}) = 0$ if and only if one outcome has a probability of 1 (no uncertainty).

Kullback-Leibler (KL) Divergence

Measures how one probability distribution P diverges from a second, reference probability distribution Q:

$$D_{\mathrm{KL}}(P \| Q) = \mathbb{E}_{\mathrm{x} \sim P} \left[\log \frac{P(\mathrm{x})}{Q(\mathrm{x})} \right]$$

It represents the extra amount of information needed to code samples from P using a code optimized for Q.

Properties of KL Divergence

- Non-negative: $D_{\mathrm{KL}}(P \| Q) \geq 0$
- Zero if and only if P and Q are the same distribution.
- Non-symmetric: $D_{\mathrm{KL}}(P \| Q) \neq D_{\mathrm{KL}}(Q \| P)$

Cross-Entropy

Related to KL divergence, cross-entropy combines the entropy of P with the KL divergence between P and Q:

$$H(P, Q) = H(P) + D_{\mathrm{KL}}(P\|Q)$$

Interpretation

Minimizing cross-entropy with respect to Q equates to minimizing the KL divergence, often used in optimization problems such as training machine learning models.

Example 5.1 (Cross-Entropy and KL Divergence in Image Classification)

In image classification, a model predicts a probability distribution Q over classes for a given image, while the true distribution P is typically one-hot encoded (1 for the correct class, 0 for others).
Cross-entropy in this context measures the difference between the distributions P and Q:

$$H(P, Q) = -\sum_{\mathrm{c}} P(\mathrm{c}) \log Q(\mathrm{c})$$

where c indexes over the classes.

Connection to KL Divergence

Cross-entropy decomposes into the sum of the true distribution's entropy and the KL divergence:

$$H(P, Q) = H(P) + D_{\mathrm{KL}}(P \| Q)$$

Since $H(P)$ is constant (the true label is fixed), minimizing cross-entropy $H(P, Q)$ with respect to Q is equivalent to minimizing $D_{\mathrm{KL}}(P \| Q)$.
During training, optimizing cross-entropy encourages the model to make predictions Q that match the true distribution P, effectively reducing the divergence from the true label distribution.

REFERENCES