Quantitative Analysis
Probability and Distribution
Independent and Dependent Events
A event that has no effect on the probability of an event happening
Mutually Exclusive Events
Two events that cannot happen at the same time
Conditional Probability
Probability of an event happening given that another event has happened
Bayes' Theorem
Random Variables
A variable whose value is a numerical outcome of a random phenomenon
Discrete Random Variables
Random variables that can only take specific, countable values
Continuous Random Variables
Random variables that can take any value within a range
Probability Mass Function (PMF)
Probability that a discrete random variable equals a specific value
Probability Density Function (PDF)
Probability that a continuous random variable falls within a range
Cumulative Distribution Function (CDF)
Probability that a random variable is less than or equal to a specific value
Expectation
Average value of a random variable over many trials
Variance
Measure of how spread out the values are from the mean
Standard Deviation
Square root of variance, measure of dispersion
Discrete Uniform Distribution
Equal probability for all outcomes
Bernoulli Distribution
Two outcomes: success (1) and failure (0)
Binomial Distribution
Number of successes in n independent Bernoulli trials
Hypergeometric Distribution
Probability of k successes in n draws without replacement
Poisson Distribution
Probability of events occurring in a fixed interval
Geometric Distribution
Number of trials until first success
Negative Binomial Distribution
Number of trials until r successes
Continuous Uniform Distribution
Equal probability over a range
Normal Distribution
Bell curve, most common distribution
Standard Normal Distribution
Normal distribution with μ = 0, σ = 1
Exponential Distribution
Time between events in a Poisson process
Gamma Distribution
Generalization of exponential distribution
Beta Distribution
Distribution of probabilities, used in Bayesian statistics
Chi-Squared Distribution
Sum of squares of standard normal variables
Student's t-Distribution
Used when sample size is small
F-Distribution
Ratio of two chi-squared distributions
Multivariate Distributions
Joint Distribution
Probability of two or more random variables
Marginal Distribution
Distribution of one variable from a joint distribution
Conditional Distribution
Distribution of one variable given another
Independence
Variables do not affect each other
Covariance
Measure of how two variables change together
Correlation
Standardized covariance, ranges from -1 to 1
Multivariate Normal Distribution
Generalization of normal distribution to multiple variables
Descriptive Statistics and Inferential Statistics
Descriptive Statistics
Summarizing data using measures and visuals
Measures of Central Tendency
Mean, median, mode
Measures of Dispersion
Range, variance, standard deviation, IQR
Shape of Distribution
Skewness, kurtosis
Visualization
Histogram, boxplot, scatterplot, bar chart
Inferential Statistics
Drawing conclusions about a population from a sample
Sampling Methods
Random, stratified, cluster, systematic, convenience
Sampling Distribution
Distribution of a statistic across samples
Central Limit Theorem
Sample means approach normal distribution as n → ∞
Confidence Intervals
Range of values likely to contain the true population parameter
Hypothesis Testing
Null hypothesis (H₀) vs alternative hypothesis (H₁)
Type I Error
Reject H₀ when it is true (false positive)
Type II Error
Fail to reject H₀ when it is false (false negative)
Significance Level (α)
Probability of Type I error, typically 0.05
Power
1 - β, probability of correctly rejecting H₀
p-Value
Probability of observing data as extreme as the sample, assuming H₀ is true
Confidence Interval
Range of values likely to contain the true population parameter
Z-Test
Used when σ is known and n ≥ 30
T-Test
Used when σ is unknown or n < 30
Chi-Squared Test
Tests independence or goodness of fit for categorical data
ANOVA
Compares means across three or more groups
Regression Analysis
Simple Regression
y = β₀ + β₁x + ε
Multiple Regression
y = β₀ + β₁x₁ + β₂x₂ + ... + βₖxₖ + ε
Assumptions
Linearity, independence, homoscedasticity, normality of residuals
Coefficients
Interpretation: holding other variables constant, a unit change in xᵢ changes y by βᵢ
Adjusted R²
Accounts for number of predictors, penalizes adding irrelevant variables
Multicollinearity
Predictors are correlated, inflates coefficient variance
Heteroscedasticity
Residuals have unequal variance, violates assumptions
Residuals
Difference between observed and predicted values
Modeling Cycle
Data Collection
Sources, sampling, quality
Exploratory Data Analysis
Visualize, summarize, find patterns
Feature Engineering
Create, transform, select variables
Model Selection
Choose algorithm based on problem type
Training
Fit model to training data
Validation
Tune hyperparameters, prevent overfitting
Evaluation
Test on unseen data, measure performance
Deployment
Put model into production
Monitoring and Improving Models
Cross-Validation
Rotate through multiple train/test splits for robust evaluation
Bias-Variance Tradeoff
Bias: underfitting, high error on both train and test
Variance: overfitting, low error on train, high error on test
Regularization
L1 (Lasso): sparse, feature selection
L2 (Ridge): shrinks coefficients, keeps all features
Elastic Net: combines L1 and L2
Model Selection
Akaike Information Criterion (AIC)
Bayesian Information Criterion (BIC)
Cross-Validation
Estimate performance on unseen data
Learning Curves
Training error vs validation error as function of training set size
Overfitting
Model memorizes training data, poor generalization
Underfitting
Model too simple, high bias
Feature Importance
Which features contribute most to predictions
Simulations
Monte Carlo Simulation
Use random sampling to estimate probabilities or expectations
Sensitivity Analysis
How does output change when inputs change
Scenario Analysis
What-if analysis, exploring different outcomes
Stochastic Processes
Random processes evolving over time
Markov Chains
Next state depends only on current state
Poisson Processes
Events occurring randomly at a constant rate
Brownian Motion
Random walk in continuous time
Geometric Brownian Motion
Used in finance to model stock prices
Monte Carlo Methods
Use random sampling to estimate probabilities or expectations
Law of Large Numbers
Sample means converge to true population mean as n → ∞
Central Limit Theorem
Sum of many independent random variables approaches normal distribution
Bootstrapping
Resampling with replacement to estimate sampling distribution
Jackknife
Leave-one-out resampling for bias and variance estimation
Permutation Tests
Shuffle labels to test hypothesis without assumptions
Bayesian Methods
Prior × Likelihood = Posterior
Conjugate Priors
Prior and posterior same family, simplifies calculation
MCMC
Markov Chain Monte Carlo, sampling from complex distributions
Gibbs Sampling
Special case of MCMC, sampling one variable at a time
Metropolis-Hastings
General MCMC algorithm, accept/reject proposals
Variational Inference
Approximate posterior with simpler distribution
Probabilistic Programming
Stan, PyMC, Pyro, TensorFlow Probability
Hierarchical Models
Parameters themselves have distributions, pooling information across groups
Time Series
Components
Trend, seasonality, cyclical, irregular
Stationarity
Mean, variance, autocorrelation constant over time
Autocorrelation
Correlation with past values
Partial Autocorrelation
Correlation after removing intermediate effects
ARIMA
AutoRegressive Integrated Moving Average
AR
Autoregressive, p lags
I
Integrated, d differences
MA
Moving Average, q lagged errors
SARIMA
Seasonal ARIMA, adds seasonal components
Exponential Smoothing
Weighted average, recent observations weighted more
Simple Exponential Smoothing
Level only
Holt's Method
Level + trend
Holt-Winters
Level + trend + seasonality
Forecast Evaluation
Rolling window cross-validation, MAE, RMSE, MAPE
Stationarity Tests
Augmented Dickey-Fuller (ADF)
Null hypothesis: unit root (non-stationary)
KPSS
Null hypothesis: stationarity
Differencing
Subtract consecutive observations to remove trend
Log Transform
Stabilize variance, convert multiplicative to additive
Seasonal Decomposition
Separate trend, seasonal, residual components
Autocorrelation Function (ACF)
Correlation at lag k
Partial Autocorrelation Function (PFAC)
Correlation at lag k after removing intermediate lags
Forecasting
Point forecast
Single value prediction
Interval forecast
Prediction interval, uncertainty quantification
Rolling Forecast
Retrain model as new data arrives
Expanding Window
Use all historical data up to current point
Sliding Window
Use fixed recent window
Multivariate Time Series
Vector Autoregression (VAR)
Multivariate ARIMA
Cointegration
Long-run equilibrium between non-stationary series
Error Correction Models
Short-run deviations from long-run equilibrium
Deep Learning
RNN
Recurrent Neural Network
LSTM
Long Short-Term Memory
GRU
Gated Recurrent Unit
Transformer
Attention mechanism, parallel processing
Attention
Weighs importance of different parts of input
Encoder-Decoder
Compress input, generate output
Pre-training
Learn from large unlabeled data
Fine-tuning
Adapt pre-trained model to specific task
Transfer Learning
Use knowledge from one domain for another
Evaluation Metrics
Classification
Accuracy, precision, recall, F1, AUC-ROC
Regression
MAE, MSE, RMSE, R², MAPE
Ranking
Precision@k, MAP, NDCG
Calibration
Predicted probabilities match actual frequencies
Cross-Entropy
Measure of difference between predicted and actual distributions
Log Loss
Penalizes confident wrong predictions
Bias-Variance Decomposition
Error = Bias² + Variance + Irreducible noise
Overfitting
Model memorizes training data, poor generalization
Underfitting
Model too simple, high bias
Regularization
L1 (Lasso): sparse, feature selection
L2 (Ridge): shrinks coefficients, keeps all features
Elastic Net: combines L1 and L2
Cross-Validation
Estimate performance on unseen data
Learning Curves
Training error vs validation error as function of training set size
Feature Importance
Which features contribute most to predictions