

CONTENT

Chapter 1: Introduction to Real-Time Data Analytics 1-13

- 1.1 Evolution of Data Analytics
- 1.2 Batch vs Near Real-Time vs Real-Time Analytics
- 1.3 Characteristics of Real-Time Systems
- 1.4 Role of AI in Modern Analytics
- 1.5 Industry Use Cases (IoT, Finance, Smart Cities, Healthcare)
- 1.6 Challenges: Latency, Scalability, Reliability

Chapter 2: Fundamentals of Data Streams 14-28

- 2.1 Data Stream Models
- 2.2 Event Time vs Processing Time
- 2.3 Stream Ordering and Late Arriving Data
- 2.4 Windowing Concepts (Tumbling, Sliding, Session)
- 2.5 Stream Processing Semantics (Exactly-once, At-least-once)
- 2.6 Data Quality in Streams

Chapter 3: Data Collection and Ingestion 29-52

- 3.1 Real-Time Data Sources (Sensors, Logs, Clickstreams)
- 3.2 APIs, Webhooks and Message Queues
- 3.3 Apache Kafka Architecture
- 3.4 Data Serialization (JSON, Avro, Protobuf)
- 3.5 Fault Tolerance and Backpressure Handling
- 3.6 Data Validation and Preprocessing

Chapter 4: Real-Time Data Processing Architectures 53-72

- 4.1 Lambda Architecture
- 4.2 Kappa Architecture
- 4.3 Event-Driven Architecture
- 4.4 Microservices-based Streaming Systems
- 4.5 Stateful vs Stateless Processing
- 4.6 Design Trade-offs and Best Practices

Chapter 5: Streaming Frameworks and Tools	73-96
5.1 Apache Kafka Streams	
5.2 Apache Spark Structured Streaming	
5.3 Apache Flink Concepts	
5.4 Stream Joins and Aggregations	
5.5 State Management and Checkpointing	
5.6 Performance Optimization Techniques	
Chapter 6: AI and Machine Learning for Real-Time Analytics	97-116
6.1 Introduction to Online and Incremental Learning	
6.2 Real-Time Feature Engineering	
6.3 Streaming Classification and Regression	
6.4 Concept Drift Detection	
6.5 Real-Time Anomaly Detection	
6.6 Model Serving and Inference Pipelines	
Chapter 7: Real-Time Analytics, Visualization and Monitoring	117-137
7.1 Time-Series Data Analytics	
7.2 Real-Time Dashboards and Visualization Tools	
7.3 Alerting and Notification Systems	
7.4 KPI Definition and Tracking	
7.5 Log Analytics and Metrics Collection	
7.6 Observability in Streaming Systems	
Chapter 8: Data Governance, Security and Ethics	138-154
8.1 Data Privacy and Compliance (GDPR basics)	
8.2 Security in Streaming Pipelines	
8.3 Encryption and Access Control	
8.4 Ethical Issues in AI-driven Analytics	
8.5 Bias and Fairness in Real-Time AI	
8.6 Responsible AI Practices	

**Chapter 9: Deployment, Scaling and Cloud
Integration**

155-172

9.1 Containerization with Docker

9.2 Kubernetes for Streaming Applications

9.3 Auto-scaling and Load Balancing

9.4 Cloud Streaming Services (AWS, Azure, GCP overview)

9.5 CI/CD for Data Pipelines

9.6 Cost Optimization Strategies

Chapter 10: Case Studies and Project

173-190

10.1 Industry Case Studies

10.2 Design of End-to-End Real-Time Analytics Pipeline

10.3 Integration of AI Models with Streaming Data

10.4 Performance Evaluation and Optimization

10.5 Mini/Capstone Project

10.6 Documentation and Presentation