



AI in Supply Chain: Machine Learning Demand Forecasting and Inventory Optimisation



AI in Supply Chain: Machine Learning Demand Forecasting and Inventory Optimisation

Technical depth: Practitioner · **Practical mode:** Lab

Introduction

Many supply chains hold too much of the wrong stock and too little of the right stock because forecasts are built in spreadsheets, rarely measured and disconnected from inventory policy. AI and machine learning can improve accuracy, but only when data is prepared properly, models are tested against simple benchmarks and planners understand what drives the numbers. This Core Concept course equips planners and analysts to build, test and govern AI forecasts and to convert forecast error into stock decisions. Working hands on with realistic planning data, participants produce a Forecast-to-Inventory Model and Deployment Brief for one product group.

Course Objectives

- Assess forecasting maturity and data readiness and select AI use cases with a clear value case
- Build and compare statistical and machine learning demand forecasts using backtesting and standard accuracy metrics
- Translate forecast error into safety stock, reorder points and differentiated service levels
- Interpret model outputs and drivers so that planners can explain and challenge a forecast
- Apply risk, monitoring and override controls to forecasting and inventory models in live use
- Produce a Forecast-to-Inventory Model and Deployment Brief for one product group

Target Audience

- Demand planners and forecasting analysts responsible for statistical forecasts
- Inventory and replenishment analysts setting stock policies and service levels
- Supply chain analysts building planning reports and models
- Planning system and ERP key users responsible for forecasting modules
- Data analysts supporting supply chain, retail or distribution functions
- Materials and spare parts planners handling intermittent demand

Course Outline

Day 1: AI in Supply Chain Planning and Data Readiness

- Forecasting Maturity Assessment Against SCOR Digital Standard Plan Processes
- AI Use-Case Screening Matrix: Value, Data Availability and Feasibility
- Demand Profiling with Syntetos-Boylan Classification
- Data Readiness Audit: Sales, Promotions, Stock-Outs and External Drivers
- CRISP-DM Project Framing for a Forecasting Initiative

Day 2: Forecasting and Inventory Models

- Statistical Baselines: Exponential Smoothing, ARIMA and Croston's Method
- Machine Learning Forecasting with Gradient-Boosted Trees and Feature Engineering
- Probabilistic Forecasting with Quantile Regression and Prediction Intervals
- Forecast Accuracy Metrics: MAPE, WAPE, MASE and Bias
- Inventory Policy Models: EOQ, Reorder Point and Newsvendor Service Levels

Day 3: Building Forecasts and Stock Policies in the Lab

- Time Series Cross-Validation and Backtesting in Python
- Forecast Value Added Analysis Against a Naive Benchmark
- Safety Stock Calculation from Forecast Error Distributions
- ABC-XYZ Segmentation for Differentiated Service Level Targets
- Forecast Driver Explanation Using SHAP Values

Day 4: Advanced Optimisation, Risk and Governance

- Multi-Echelon Inventory Optimisation Concepts and Network Trade-Offs
- Demand Sensing and Disruption Scenarios in Forecast Models
- Model Drift Monitoring and Retraining Triggers
- AI Risk Assessment Using ISO/IEC 23894 and the NIST AI RMF
- Planner Override Governance and Human-in-the-Loop Controls

Day 5: Lab Build and the Forecast-to-Inventory Model

- Retail Replenishment Case: Promotions and Stock-Out Censoring
 - Spare Parts Case: Intermittent Demand and Service Level Trade-Offs
 - Forecast-to-Inventory Model Build for an Own Product Group
 - Deployment Brief: Business Case, KPIs and ISO/IEC 42001 Controls
 - Peer Review of Models and Deployment Briefs
-

Skills You Will Gain

- Time Series Forecasting
- Machine Learning Feature Engineering
- Forecast Accuracy Measurement
- Safety Stock Modelling
- Inventory Segmentation
- Model Explainability
- Model Performance Monitoring
- AI Risk Assessment

Why Attend This Course

- Return with a working Forecast-to-Inventory Model for a real product group, supported by a deployment brief
- Judge whether a vendor or in-house AI forecast actually beats a simple benchmark before it is trusted
- Explain to managers why stock targets move when forecast error changes
- Work through retail, spare parts and distribution data alongside planners from different sectors

Conclusion

Better forecasts matter only when they change stock decisions and planners trust them. This course moves from data readiness and use-case selection, through statistical, machine learning and probabilistic forecasting methods, to safety stock, segmentation, explainability and the controls that keep models reliable after go-live. The final day turns that material into a Forecast-to-Inventory Model and Deployment Brief, giving participants a tested and explainable way to bring AI forecasting into daily planning without losing control of service levels or working capital.
