

Supervised Learning for Integrated Forecasting and Inventory Control

Joost van der Haar

We explore using supervised learning with custom loss functions for multi-period inventory control with feature-driven demand. This method directly considers feature information such as promotions and trends to make periodic order decisions, does not require distributional assumptions on demand, and is sample efficient. The application of supervised learning in inventory control has thus far been limited to problems for which the optimal policy structure is known and takes the form of a simple decision rule, such as the newsvendor problem. We present an approximation approach to expand its use to inventory problems where the optimal policy structure is unknown. We test our approach on lost sales, perishable goods, and dual-sourcing inventory problems. It performs on par with state-of-the-art heuristics under stationary demand. It outperforms them for non-stationary perishable goods settings where demand is driven by features, and for non-stationary lost sales and dual-sourcing settings where demand is smooth and feature-driven.