Strategy | Industrial automation | Performance management | Supervised Learning

Data-Driven Key Performance Indicator Modeling for Robotic Mobile Fulfillment Systems

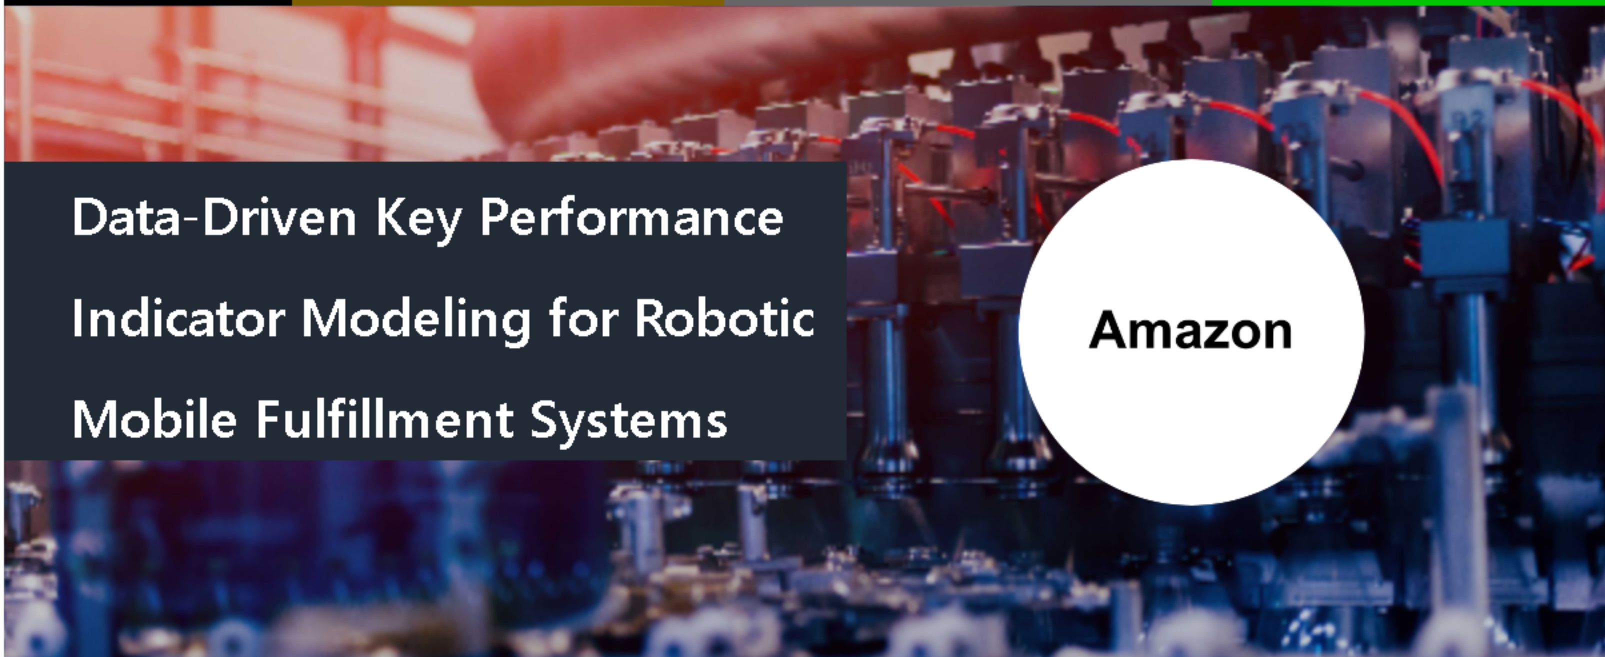

Amazon

BUSINESS PROBLEM

E-commerce companies operating robotic mobile fulfillment systems (RMFS) face significant challenges in predicting and optimizing brownfield warehouse performance across complex networks. Traditional modeling methods excel in providing design stage insights but struggle to anticipate the dynamic nature of these systems, which constantly evolve due to rapid market changes, continuous experimentation, and shifting customer demands. Companies with large, digitized networks can benefit from data-driven approaches that leverage this operational data to provide accurate predictions and actionable insights for decision-making.

DATA SOURCES

The research utilized extensive operational data from Amazon Robotics' Robotic Mobile Fulfillment Systems (RMFS) and was ultimately obfuscated in publishing. Original data included hourly time-series data, as well as site characteristics such as zone configurations and equipment details. This data spanned multiple years and sites across North America.

Data Types and Format

Data utilized was tabular time-series sampled at various frequencies as well as unstructured semi-static site characteristics and classes.

Author: Steffan Sowards

APPROACH

Machine learning techniques were employed to develop predictive models for RMFS performance metrics. Gradient boosted tree ensembles were selected across model types. Template-specific models were created to address unique facility configurations. Interpretability techniques, including SHAP analysis and permutation feature importance, provided insights into system behavior and KPI drivers.

IMPACT

The developed predictive framework demonstrated the viability of this form of brownfield data-driven modeling in warehouse automation. Models achieved a mean R^2 value of 0.7838 across all templates and KPIs, with particularly strong performance in the most predictive metrics (mean R^2 of 0.9660). The framework was used to generate immediate practical artifacts such as lookup tables that visualize performance trade-offs between different parameters. These tables allow operators to quickly assess the impact of proposed changes and identify optimal parameter combinations that maintain desired performance levels. The model's scenario modeling capability also supported proactive planning by simulating "what-if" scenarios for inventory layouts, robot fleet management, and peak season preparation. This work also enhanced system understanding through interpretability techniques that identified key performance drivers, non-linear performance behavior, variable interactions, and system tipping points. This knowledge helps has been captured internally to help prioritize operational adjustments and system design modifications, focusing resources on factors with the greatest impact. By bridging the gap between basic analytics and complex simulation tools, this framework enables more efficient resource allocation, supports strategic planning, and facilitates rapid adaptation to changing market conditions.

DRIVERS	The e-commerce industry's rapid growth and technological advancements enables this type of work. The catalysts for this solution were the need for rapid data-driven insights in RMFS operations, dynamic market conditions, complex system interactions, and the availability of rich operational data.
BARRIERS	Barriers included the scope and complexity of RMFS systems and the need to balance model interpretability with predictive accuracy. The rich availability of data required careful handling, engineering, and selection to maintain the utility and control of model inputs at the expense of perceived model performance.
ENABLERS	Amazon's massive active brownfield Robotic Mobile Fulfillment Systems network was central to this work. Key enablers were the collaboration with Amazon Robotics, access to extensive operational data, and the use of advanced machine learning techniques like CatBoost and SHAP analysis.
ACTIONS	Initial project actions included data collection and preprocessing, feature engineering, model development using AutoML and CatBoost, performance evaluation, and model interpretation using SHAP and permutation feature importance. Enduring impact was enabled through findings reports, look-up table generation, Sagemaker code documentation, and technical handoff to Amazon Robotics internal analytics and data-science personnel.
INNOVATION	Innovative aspects of this work included making use of large operational datasets in RMFS system modeling, the use of machine learning to replace traditional simulation methods, the development of a novel and generalizable feature engineering framework, the generation of practical operations artifacts, and the application of SHAP analysis for system performance behavior assessment.
IMPROVEMENT	The solution achieved a mean R^2 value of 0.7838 across all templates and KPIs, with particularly strong performance in the top metric (mean R^2 of 0.9660). This performance enables reliable forecasting and operational optimization in a format sensitive to non-linear system behavior without the time and cost associated with high fidelity simulation.
BEST PRACTICES	Best practices for replicating this solution include leveraging real-world operational data, assessing a broad range of advanced machine learning techniques, and focusing only on actionable inputs for model interpretability and utility. Ongoing use additionally required continued model performance monitoring, retraining, and validation against evolving system and market behavior.
OTHER APPLICATIONS	Potential applications extend beyond RMFS to other complex warehouse automation systems, such as Automated Storage and Retrieval Systems (AS/RS) and Autonomous Mobile Robots (AMRs). The approach could also be applied in other industries requiring predictive modeling of complex systems, like general manufacturing and logistics.