

AUTOMATED DEMAND FORECASTING FOR A MAJOR LOGISTICS COMPANY

Elder Research implemented an automated framework for time-series forecasting at a major logistics company. Our system, combining R and Apache Spark™, produces 35 million forecasts in under one hour, and selects the optimal time-series forecast algorithm in each of three forecasting windows. Forecast results from our framework were 88% accurate at a four-week horizon.

INDUSTRY

- » Logistics

BUSINESS NEED

- » Modernize operations through analytics
- » Automate production forecasting to support next phases of software development

SOLUTION

- » Implement production framework for time-series forecasting as an R package
- » Selects from 11 time-series algorithms at multiple forecast horizons to produce the most accurate forecasts

BENEFIT

- » High-grade production-scale solution for demand forecasting
- » Flexible/extensible architecture for changing algorithms
- » 35 million forecasts in under an hour with median accuracy of 88% four weeks out

THE CHALLENGE

Logistics is a mature, technologically-advanced, and analytically-sophisticated industry. Still, even after decades of improvements coming from the Industrial Engineering and Operations Research fields, major efficiencies can still be realized by applying advanced analytics, data infrastructure, and computing power. All business processes in logistics rely on accurate demand forecasting in the short, medium, and long-term to inform resourcing, planning, and staffing to support future needs. Our client was three months into a highly-visible, strategic analytics project and with an urgent need to have forecast results in their production system. Given the strategic importance of this project, they needed to quickly scale a prototype forecast model into their automated production system that interfaces with a new platform for planners.

THE SOLUTION

Elder Research was hired to provide a bridge between technical experts and the application development team responsible for time-series forecast implementation. We worked collaboratively with the prototype model authors, software developers and architects, database administrators, and business stakeholders to ensure that our production solution would meet requirements, interface with existing systems, and provide the flexibility required for future development. We also provided a valued perspective on statistical and optimization methods and techniques for the operations research team that created the prototype model.

RESULTS

In three weeks we delivered a functioning production time-series forecasting framework using R and Spark. After six months we had scaled to a refined framework that produces 35 million forecasts in under one hour on over 2000 locations in our client's network. This framework features automated execution and algorithm selection at short, medium, and long-term horizons. At a four-week interval, our forecasts had a median accuracy of 88%, despite high variance in the characteristics of the entities being forecast. We also developed a flexible forecasting model API to enable easy inclusion/exclusion of time-series algorithms as better techniques are identified or existing algorithms are replaced.

Contact Us

www.elderresearch.com
 contact@elderresearch.com
 (434) 973-7673
 © 2020 Elder Research Inc.



ELDER RESEARCH
 — DATA SCIENCE · AI · MACHINE LEARNING —

Office Locations

Charlottesville, VA
 Washington, DC
 Baltimore, MD
 Raleigh, NC
 London, UK