

Background

In the hyper-competitive third-party logistics industry, every minute counts. Third party logistics (3PL) brokers operate in the trucking spot market, where agents match one-off shipments with truckers (carriers) willing to transport them. Due to national trucking shortages and increasing demand, brokers must work quickly to contact carriers likely to accept a particular shipment while maintaining profitability. Adding to the challenge, brokers have often agreed on a price with the shipper before contracting a carrier, so minimizing carrier costs can make the difference between profit and loss. A leading 3PL firm approached Mosaic Data Science with the goal of using analytics to prioritize carriers for brokers to call, thereby helping the agents to source low-cost carriers with the fewest possible calls.

Analysis

First, Mosaic's data science team met with project stakeholders and future analytics end-users to discuss the project requirements and broker workflows. The team decided to rank carriers using an expected value framework that weights the predicted profit margin from using a carrier for a shipment by the probability that the carrier would accept the load. Mosaic used exploratory data analysis of historical shipment data and consultation with subject matter experts to determine how these values would be calculated.

Exploratory analysis of the historical shipment data revealed that carriers who had above-average profit margins (i.e., took shipments for a lower than normal cost) in the past were likely to yield above-average profit margins on similar current shipments. Similar shipments were shown to be those of a similar distance and near the same geographic origin markets. As a result, model features were calculated for each carrier based on the carrier's past performance on shipments of each of four different distance bands (local, short, medium, or long) and out of specific regional markets. Markets had previously been defined by the client as groups of zip codes in specific metropolitan regions. To control for seasonal variation in shipping demand and carrier costs (e.g., diesel prices), all features used a median-centered average profit margin centered on the profit margin across all carriers for the shipment type.

The exploratory analysis also showed that carriers tended to operate on just a few of the many possible origin-destination market pairs. For example, a carrier might tend to take shipments out of Los Angeles, CA to a few specific destinations. Therefore, carriers were determined to be likely to accept a shipment if they had accepted jobs on the same market-to-market shipping lane, or out of the same origin market, in the past 90 days. In addition, carriers often post future availability in specific locations to online "load boards;" carriers with posted availability near an origin market were assigned a higher probability of acceptance.

The software package Mosaic created was split into two components: a computationally-heavy batch process that runs overnight to train a machine learning model for carrier profitability analysis, and a lightweight real-time process that can be accessed with a custom API to deliver instantaneous carrier recommendations.

All scripts were written in Python so that they could easily integrate with the company's existing custom-built software platform.

Results

Mosaic helped create an analytics toolbox with a set of models and data processing scripts that provide advanced analytics-based insights and recommendations based on internal and external data sources. The high-level architecture for the sourcing recommendation toolbox capabilities are pictured in Figure 1. Mosaic designed and built all components shown in the grey box; components outside of the box were part of the existing infrastructure maintained by the client.

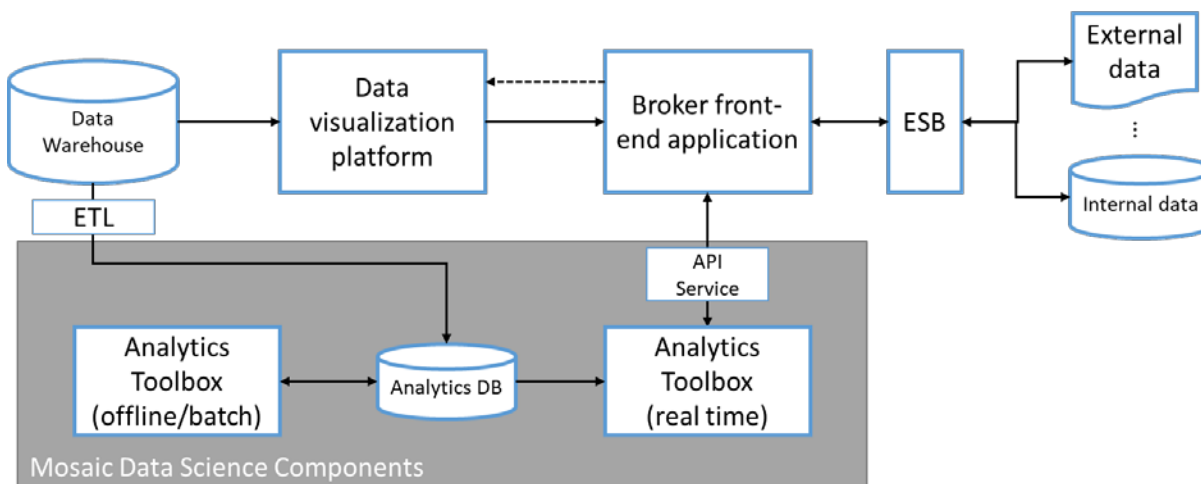


Figure 1. High-level architecture diagram for analytics toolbox to support sourcing recommendation capability.

Mosaic created a set of Python scripts to run a nightly batch process and an API to provide real-time ranked carrier lists to brokers throughout the workday. Figure 2 shows a diagram of the dynamic workflow of the Analytics Toolbox API.

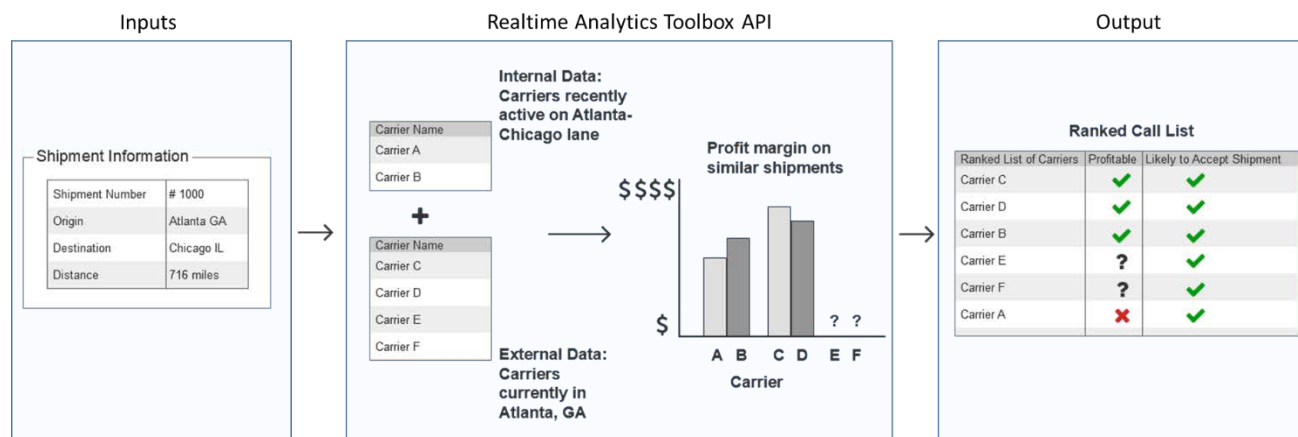


Figure 2. Overview of real-time recommendations generated by the Analytics Toolbox.

When a new shipment has been accepted by the broker and needs to be sourced (i.e., a carrier needs to be found who will accept the load), the script will seamlessly use information from each shipment that comes in to supplement a list of available carriers based on internal data. The analytics toolbox then pulls estimated profitability information for the relevant carriers based on the previous night's batch

processing. Next, the toolbox uses the number of recent shipments carriers have taken on the relevant origin-destination market pair (a shipping lane), or near the relevant origin market, to assign an estimated probability that each carrier would accept the current shipment if it were offered. Finally, the acceptance probabilities are combined with the expected profit margin for each carrier, ranked by this combination, and returned as a JSON-formatted list. Because all the processor-intensive computations are run at night and staged for the API, this entire real-time process can be completed for hundreds of users at a time, taking just milliseconds to process each request.

Instead of manually deciding which carriers to call and in which order, agents will be able to simply work their way down an already-sorted list. Even new employees at the firm with little to no experience sourcing shipments can systematically contact carriers that have the best expected value to the firm, given the dual goals of maximizing profit and minimizing time to book.

With this new system, agents will be able to quickly call the right carriers at the right time to maximize corporate profitability and customer satisfaction.

For More Information

Want to learn more? Please contact info@mosaicdatascience.com