Improving Long-Range 3D Object Detection Methods for Autonomous Box Trucks using Sensor Fusion

Jeremy Flics (Honors Capstone), Nisarg Polra (Honors Capstone), Claire Wan, Claire Tinker, Kristen Worden, Herbert Li, Kayvan Najarian (Faculty Mentor), Collin Peterson (Sponsor Mentor)

Background

One of the major goals of JP Morgan is to build a follow the sun scheduling model that leverages their global workforce to meet demand as it comes. However, supervisors in charge of staffing certain tasks do not have the necessary line of sight into short-term changes in request volume to schedule flexibly. As such, scheduling is done statically across multiple offices during US working hours. In particular, Our scope was restricted to considering three offices in Chicago, Manilla, and Bengaluru all working during Central Time standard working hours.

Proposed Solution

Our solution is two pronged. We first conducted a one-time statistical analysis of historical data to recommend a new baseline schedule that adheres to the follow the sun model. We then built and ML driven tool to predict changes in request volume and recommend adjustments to the base schedule up to two weeks in advance.

Follow the Sun Schedule and Work Volume

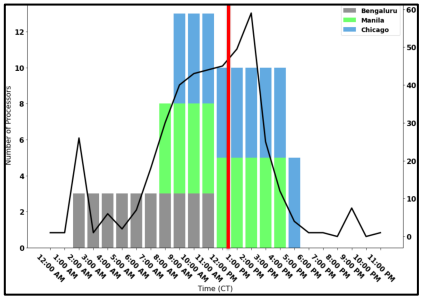

Volume Prediction

In this project we had to make multi-step ahead time series predictions using relatively limited data. In our early testing, two models showed promising results. Long Short Term Memory (LSTM), and Gradient Boosted Regression Trees (GBRT). Eventually we went with the latter, because it had lower error metrics.

Our GBRT based model had a Mean Absolute Error of **4.6** which was about **26.7%** of the percentage of hourly work volume. Since our target to be production ready was 20%, this was a viable proof of concept.

User Interface Design

Once volumes were predicted, we had to serve our insights to stakeholders in an interpretable manner. To do this, we got feedback from UI/UX designers at the firm and came up with the following design. This frontend was implement using React/Typescript and deployed to AWS on a Kubernetes cluster.

System Design

Since our product is targeting only a few key users, and only needs to be accessed a few times a day, paying for dedicated server uptime did not seem worth it. As such, we leveraged AWS services to design an event driven system.

Throughout designing our system, we emphasized principles of strong separation of concern between tasks to improve interoperability. This was done because many legacy systems in the firm are in the process of being modernized, and we did not want to add complexity to that process. Below is a, partially redacted, systems diagram of our architecture.