

MS-E2191 - Graduate Seminar on Operations Research

An Approximate Dynamic Programming Algorithm for Large-Scale Fleet Management: A Case Application

Homework solution

Possible answers that could be found from the chapter 5 of the paper.

How well the model estimates marginal values of drivers (5.1)?

- The values of the drivers are obtained from the value function of last iteration round N at time $t = 0$.
- Aggregation by driver home domicile and capacity type was used
- To get some idea how the marginal values actually behave when more drivers are added, the simulation is run by
 1. adding 10, 20, 30, 40 and 50 drivers of four different driver classes. This resulted relatively linear behavior with up to 20 new drivers.
 2. Adding 20 new drivers with 20 different types. This simulation was run 10 times. For each driver type, 95% confidence intervals of the value function were calculated.
- The estimate of marginal value from value function fell to the simulated confidence interval in 18 of 20 drivers.
- With 95% confidence interval, it is expected that only 1 driver would be outside of the confidence intervals
- Confidence intervals were really wide and only 20 driver types were tested. Simulation was only run 10 times.
- For small changes in number of drivers the model can likely give good estimates

How the Driver Remix Experiment was done (5.2.)?

- The goal was to increase the value of the objective function by changing the number of drivers belonging to each driver class using the drivers' marginal values.
- Stochastic gradient algorithm was used to determine the number of drivers to be added or removed. It has step for the number of adding new drivers of
$$R_a^{n+1} = \max \{0, R_a^n + \beta(\bar{v}_a^n - \bar{v}_*^n)\}$$
where the new number is maximum of 0 and the old number plus difference between average marginal value of all attributes a and marginal value of an specific attribute type a. This difference is scaled by β to keep the changes minimal in each iteration.
- After the update we scale the number of drivers to stay constant and continue to next iteration
- Driver Remix was applied at iteration 400 so that the value functions stabilized before that.
- The experiment resulted in increase of 0.1 million in the objective function and change of 5% percentage points.