

Turo Case Study



Jérôme Selles

Director of Data Science and Analytics



Turo offers car owners and travelers with a service to share vehicles by renting them.

Use Case

Turo uses ScienceOps to deploy the data science team's dynamic pricing model and recommender engine into their web and mobile app.

Results

Yhat removed Turo's data science team's dependence on other engineering teams so that they can realize the value of their work almost immediately.

Turo wants to put the world's billion cars to better use.

Much like AirBNB revolutionized the hotel industry, Turo is transforming the car rental industry, one adventure at a time. Turo offers more than 800 makes and models of vehicles. Owners offer their cars, sometimes including delivery, to renters in more than 4,500 cities and 400 airports across the U.S. Turo also started its international expansion last year, launching in both Canada and the UK. An undertaking of this size is no small endeavor.

Early on, the [Turo team recognized the importance of data](#), and began capturing and storing information on how users navigated the site to lend and borrow vehicles in a central Amazon Redshift cluster. Jerome Selles, Director of Data Science and Analytics, defines the team's mission as "turning data into an asset for the company" and divides Turo's approach to leveraging data into two major categories, "decision science and data features."

"Data features are algorithms we build into the app that directly impact and personalize a user's experience."

"Decision science is making the right decisions using data, like what aspect of the product to focus on next or what advertisement to run. Data features, on the other hand, are algorithms we build into the app that directly impact and personalize a user's experience," explains Jerome.

"Initially, we focused on the decision science side. We were doing a lot of ad hoc analyses and reporting, but we hadn't built predictive models. The app was still relatively simple. We were providing a list of available cars, without any kind of data magic to show you the right one or help you along in the process."

"...we knew that it would be really hard for users to predict fluctuations in demand...so we decided to build a dynamic pricing algorithm."

"Our first pricing scheme was static--listers chose one price that was fixed

Turo Case Study

all year. In reality though, car demand varies, just like hotel demand, so prices should change accordingly. We wanted to give our users the flexibility to vary rates by day, but we knew that it would also be really hard for owners to predict fluctuations in demand. So we decided to build a dynamic pricing algorithm.”

“The algorithm evaluates a few straightforward variables like car make, model, and year, as well as some less obvious factors like demand and competition...”

“Our team built a model that could suggest a reasonable price based on all the data we had at our disposal. The algorithm evaluates a few straightforward variables like car make, model, and year, as well as some less obvious factors like demand and competition, both intra- and extra- Turo.”

“Our data science teams works predominantly in Python but our engineering team develops in Java and Javascript, so there was no clear path to production.”

“Our data science teams works predominantly in Python but our engineering team develops in Java and Javascript, so there was no clear path to production. I actually came across Yhat’s platform at a data science meetup in San Francisco. I thought the idea was really interesting. What stood out to me was the vision of separating the development tracks, uncoupling data science modeling from the user facing product.”

“To be honest, my first inclination was to see if we could build something in house. With enough time, our team was able to port the pricing model into production but we struggled with refreshing the model without any downtime on the API. Yhat provides an elegant solution to this problem. When we thought about the opportunity cost of building and maintaining a “good enough” solution versus purchasing reliable enterprise software with support, it didn’t make sense to try to recreate the wheel. We looked at a couple of competing products as well, but Yhat had the most robust deployment solution and the best support.”

“Yhat had the most robust deployment solution and the best support. It took less than a day to migrate our pricing model to the Yhat platform and deploy it into production as REST API endpoints.

“It took less than a day to migrate our pricing model to the Yhat platform and deploy it into production as REST API endpoints. Since then we’ve also written and deployed a complex recommender engine using the Yhat platform. We can update our models in seconds, rather than in weeks or months like it would’ve taken without Yhat. Yhat removed our dependence on other engineering teams so that we can realize the value of our work almost immediately.”



Turo is a car rental marketplace where travelers can rent any car they want, wherever they want it, from a nationwide community of local car owners. Travelers choose from a unique selection of nearby cars, while car owners earn extra money and help fuel the adventures of travelers they meet along the way.

www.turo.com



Yhat (pronounced Y-hat) is an end-to-end data science lifecycle management suite. Yhat helps equip and organize teams to analyze/model data collaboratively in the cloud, manage their own compute resources without IT, and deploy predictive models to production faster, more reliably and without writing any custom deployment code.

www.yhat.com

Address

45 Main Street
Suite #707
Brooklyn, New York
11201

Email

info@yhathq.com
support@yhathq.com

Phone

(718) 855-2107