

Data Science Case Study

Forecasting Hotel Room Demand

Mosaic was engaged by a leading hotel chain to assess the best way to predict future demand for hotel rooms across their various properties.



Industry

Hotel



Use Case

Demand Forecasting,
Open Source vs License



Techniques

Prophet, Time Series
Modeling, Classification,
Analytics Adoption & Scale



Outcome

Production-grade models
provide accurate forecasts
to room occupancy over
global properties & various
timeframes.

BACKGROUND

A prominent publicly traded hotel chain that operates global properties across multiple brands had been investing heavily in developing advanced analytics capabilities and capacity to bring value now, and into the future, for the business. Efforts were underway to bring data together in ways not previously explored, with a focus on enabling analytics across the enterprise.

The business had been using an existing demand forecasting model from an enterprise analytics software company, but were dissatisfied with its level of accuracy which hindered the business in appropriately planning and executing resource allocation. The hotel chain needed an analytics consulting partner who could provide predictive analytical capabilities to improve the accuracy of future demand estimates.

Mosaic, a leading data science consultancy, was engaged by the hotel chain to assess the best way to predict future demand for hotel rooms across their various properties. The ultimate objective was maximizing revenue from a resource with constrained supply (i.e. limited number of rooms) and fluctuating demand over time (i.e. night(s) of stay). This is a critical analytics task for hotel chains, as unoccupied rooms on a given night earn zero revenue, while demand in excess of room capacity carries a cost in terms of lost revenue.

[Therefore, forecast of future demand helps the hotel industry make key decisions in revenue management¹.](#) One can assume a generally negative correlation between price and quantity of demand, and determining how this relationship applies to a given hotel property can inform decisions on room rates offered by that property. Obviously, the appropriateness of such decisions depends on the accuracy of demand forecasting.

FORECASTING HOTEL ROOM DEMAND

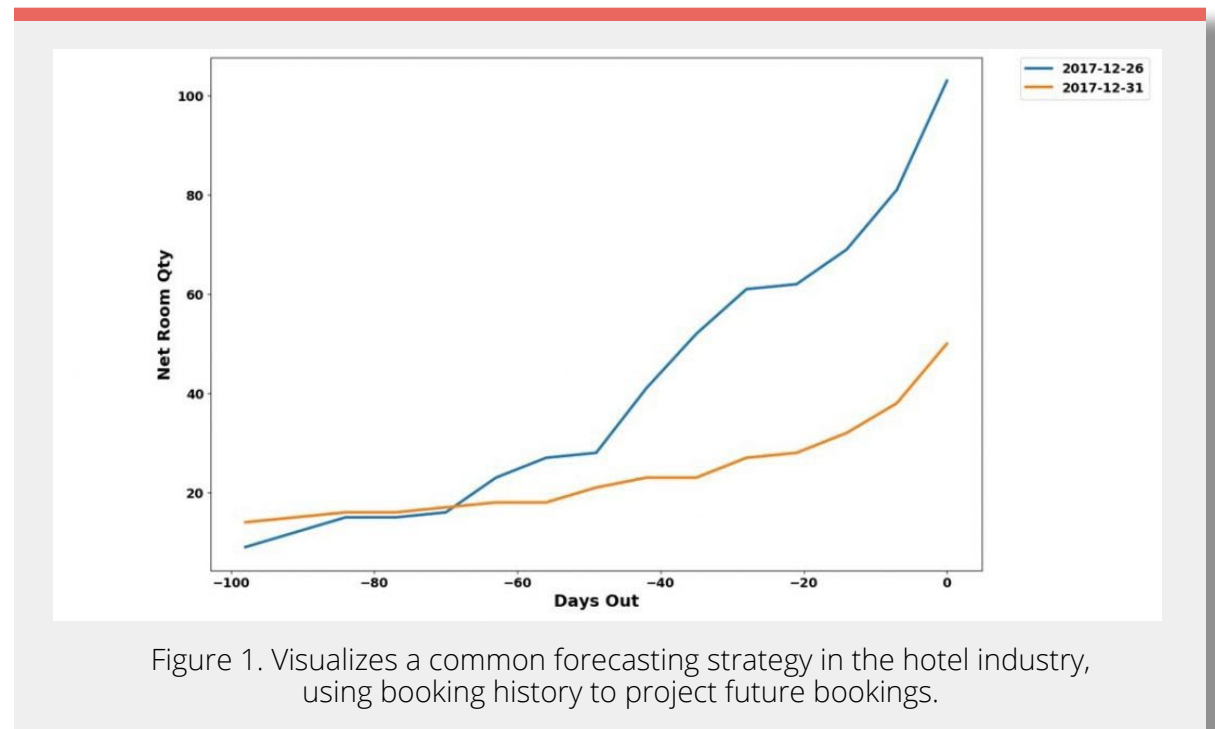


Figure 1. Visualizes a common forecasting strategy in the hotel industry, using booking history to project future bookings.



Mosaic needed to develop forecasts that outperformed the current analytics tool used by the hotel chain, providing the business with an accurate picture of demand. The first step was becoming familiar with traditional approaches to demand forecasting in the hotel industry. Typically, this type of problem is viewed from two angles: an historical time-series modeling approach and an advanced booking approach. The time-series approach models future demand day-by-day by using historical data to fit a parameterized model, and then extrapolating the model into the future. The advanced booking approach uses historic booking data for a given day to extrapolate future bookings given current bookings on-hand. After spinning up quickly on these approaches, the Mosaic data science consultants began to implement these analytical methods using an open-source toolset.

Mosaic attacked the historical booking model similarly to any other machine learning (ML) modeling problem: by testing various features and ML algorithms. The challenge in this case was that almost all the information came from time-series features (day of week, month, week of year, holidays, etc.). Although Mosaic was able to get improved results this way, experimentation showed that one could get comparable results with decreased computation time using time-series forecasting, so that was the approach ultimately adopted.

For the advanced booking approach, Mosaic decided to fit a model to the bookings themselves (rather than just using the past data in a lookup table) and obtained good results, especially in the few days leading up to the forecast date. These projections were then combined with the time-series model for an overall demand forecast.

RESULTS

Mosaic was able to outperform the current analytical forecasting tool across multiple properties and timeframes. Mosaic's data scientists were able to achieve this result using open-source software, which could save the hotel chain significant licensing costs. Now, the hotel chain is able to allocate resources more effectively, leading to a number of downstream positive effects on metrics and bottom line net income. Not only are the business decision makers using data more efficiently, the analytics team at this hotel chain gets another highly visible project 'win', inspiring more confidence and more projects for the team.

Endnotes

1. <https://www.sciencedirect.com/science/article/pii/S0020025500000827>

