Predicting Future Arrivals and Occupancy Levels in an Emergency Department

Xiaopei Zhang
(with Ward Whitt)

Columbia University
Department of Industrial Engineering & Operations Research

Overview

The Data and the Problem

- We have a dataset which contains the records of each customer that visited an Emergency Department (ED) of an Israeli Hospital located in Haifa, Israel.
- We know the **arrival and departure times of each patient**, and the admission decision (release or admitted to a deparment of the hospital).
- We **don't know** any detailed information about each customer (such as age, sex, etc.) and what happens during the visit.

The Data and the Problem

- The ED has approximately 40 beds and has independent resources.
- We have 46 months' data, during which $180,000+$ patients visited the ED.

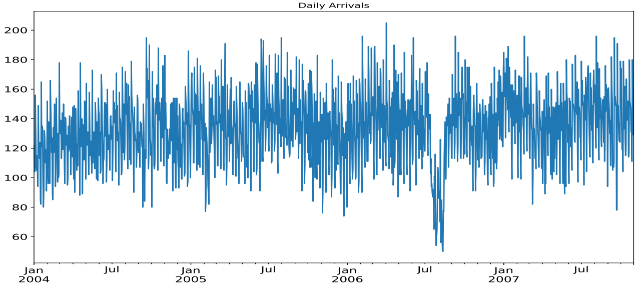

Figure: Daily Total Arrivals from 01/01/2004 to 10/31/2007.

The Data and the Problem

- **Question**: How can we make real-time prediction for arrivals and occupancy levels based on recent and historical data up to present time?
- We want to explore several techniques that could be used to predict the future arrivals and occupancy levels.
- The work is based on one of our previous work (Whitt and Zhang, 2017), where we
  - conducted data analysis of the patient flow, and
  - proposed a stochastic model to characterize the system.

Previous Work: Develop a Stochastic Model

- In Whitt and Zhang (2017), we proposed an $M_t^T/GI_t/\infty$ queueing model for the ED data by analyzing part of the data:
  - Daily total arrivals are (periodic) time series.
  - Given the daily total arrivals, the arrivals are a non-homogenuous Poisson process within a day (but with constant rate for each hour).
  - The length of stay (LoS) of each customer only depends on the arrival time and is independent of other LoSs and the arrival process.

The Model Features

- The model we proposed describes the system well:
  - It captures the **time-varying** structure (both the arrival rates and the LoS distributions) with period length 1 week.
  - It properly deals with the **over-dispersion** of the arrival process. (More random than a Poisson process.)

The Over-Dispersion of the Arrival Process

- $I(t) = \dfrac{\mathrm{Var}(A(t))}{\mathbb{E}(A(t))}$, Index of dispersion for counts, where $A(t)$ is the arrival counting process.
- For Poisson process, $I(t) \equiv 1$.

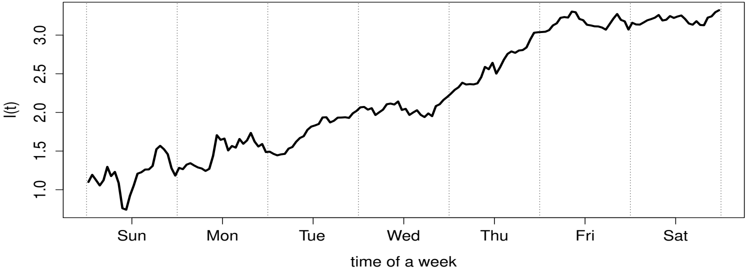

Figure: Index of dispersion for counts over a week.

The Time-Varying Nature of Arrivals and Length of Stay (LoS)

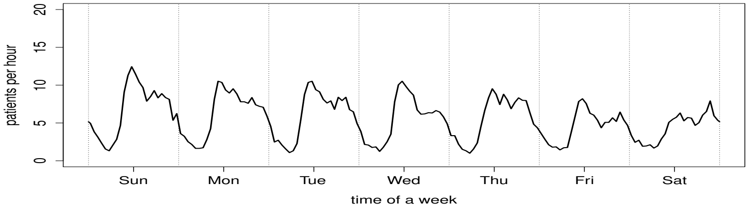

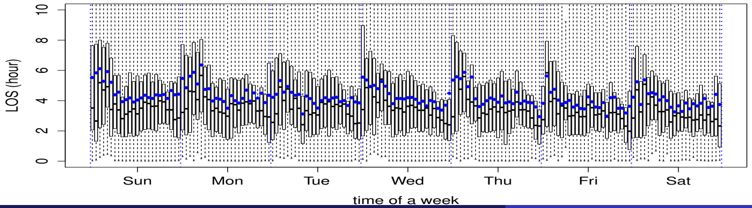

Making Prediction

- Now we turn to our new study, i.e., how to predict future arrivals and occupancy levels?
- We use the full data set and remove the war period.
- We split the data set into a training set (before war) and a test set (after war).

Things Need to Predict

- According to our model, we need to predict three things: 1) the daily total arrivals, 2) the arrival rate within a day, and 3) LoS distributions.
- Daily total arrivals: We tried several models including SARIMAX (Seasonal AutoRegressive Integrated Moving Average with eXogenous regressors) model, MLP (Multi-Layer Perceptron) model, etc.
- The arrival rate within a day: Moving average estimation.
- Distribution of LoS for each patient: Rolling empirical distribution.
- Occupancy level: Based on the prediction of the above quantities and current state of the ED.

Daily Total Arrivals Prediction

- We tried both highly structured traditional time series models as well as neural-network-based machine learning models.

Method	Training MSE	Test MSE
(i) Regression w/ DoW	248.9	264.6
(ii) Dynamic regression w/ DoW	248.3	269.9
(iii) SARIMA	221.5	263.4
(iv) Regression w/ calendar and weather var.	206.0	234.3
(v) SARIMAX (iii)+(iv)	181.6	191.8
(vi) MLP w/ everything	205.7	265.8

Table: Summary of the training and test results for the five methods to predict the daily arrival totals.

Daily Total Arrivals Prediction

- We finally choose SARIMAX$(6,1,0)(0,0,2)_7$ as our model for daily total arrivals.
-
$$A_t = x_t^T\beta + A_{t-1} + \sum_{i=1}^{6} \phi_i(A_{t-i} - A_{t-1-i}) + \epsilon_t + \Theta_1\epsilon_{t-7} + \Theta_2\epsilon_{t-14},$$

  where x_t includes calender indicators and min/max temporature on day t.

Hourly Occupancy Level Prediction

- We want to utilize as much currently available information as we can to accurately predict the occupancy level for the next few hours.
- Future Occupancy Level = Current Customers Remaining + Future Arrivals
- We have answered the second part (Future Arrivals) by the previous slide.
- For the first part (Current Customers Remaining), the idea is estimating the 'survival' probability

$$P(W > j + k | W > j), \quad j, k \geq 0,$$

by the historical data, where W denotes the LoS of a patient.

Test the Estimator for 1-Hour-Ahead Prediction

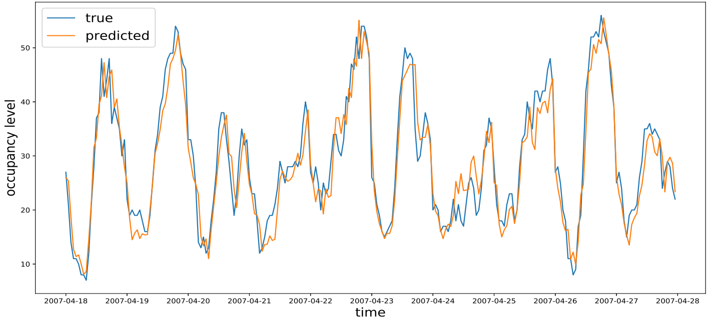

Figure: The predicted 1-hour-ahead occupancy level.

- The MSE is 14.65. (Compared to 23.33 if we use $\hat{Q}_{k+1} = Q_k$ and 51.91 if we use moving average.)

Predict the Occupancy for More Than 1 Hour Ahead

- Our occupancy predictor can be easily applied to predict for several hours in the future, but as we are using more and more future arrivals prediction, we expect the prediction error increases over time.

# of hours ahead	1	2	3	4	5	6
MSE	14.65	25.21	35.05	66.33	113.59	160.16

Table: MSE for predicting the occupancy several hours ahead.

Summary

- It is helpful to use exogenous variables such as calendar and temperature data when we make predictions.
- Good time series methods may do better than machine learning such as MLP on such highly structured low-dimensional problem.
- The data are critical, but the model also can help in prediction problems.
- With longer time period, we see a trend. So for prediction and model building for the short-term future, we should emphasize recent data.

References

References

1 Whitt, W. and Zhang, X. (2017). A Data-Driven Model of an Emergency Department. Operations Research for Health Care, vol. 12, No. 1, 2017, pp. 1-15.

2 Whitt, W. and Zhang, X. (2018). Forecasting Arrivals and Occupancy Levels in an Emergency Department. Submitted to Operations Research for Health Care.

3 Whitt, W. (1999). Predicting Queueing Delays. Management Science, vol. 45, No. 6, June 1999, pp. 970-888.

4 Ibrahim, R. and Whitt, W. (1999a). Real-Time Delay Estimation Based on Delay History. Manufacturing and Service Operations Management, vol. 11, No. 3, Summer 2009, pp. 397-415.

5 Ibrahim, R. and Whitt, W. (1999b). Real-Time Delay Estimation in Overloaded Multiserver Queues with Abandonment. Management Science, vol. 55, No. 10, October 2009, pp. 1729-1742.

6 Brockwell, P.J. and Davis, R.A. (2016). Introduction to Time Series and Forecasting. Springer.

7 Friedman, J., Hastie, T. and Tibshirani, R. (2001). The Elements of Statistical Learning. New York: Springer series in statistics.