Smart Stores: A scalable foot traffic collection and prediction system

Soheila Abrishami[1], Piyush Kumar[1], and Wickus Nienaber[2]

[1] Florida State University, Tallahassee, FL 32306, USA
abrisham,piyush@cs.fsu.edu
[2] Bloom Intelligence, Tallahassee, FL 32303, USA
wickus@bloomintelligence.com

Abstract. An accurate foot traffic prediction system can help retail businesses, physical stores, and restaurants optimize their labor schedule and costs, and reduce food wastage. In this paper, we design a large scale data collection and prediction system for store foot traffic. Our data has been collected from wireless access points deployed at over 100 businesses across the United States for a period of more than one year. This data is centrally processed and analyzed to predict the foot traffic for the next 168 hours (a week). Our current predictor is based on Support Vector Regression (SVR). There are a few other predictors that we have found that are similar in accuracy to SVR. For our collected data the average foot traffic per hour is 35 per store. Our prediction result is on average within 22% of the actual result for a 168 hour (a week) period.

Keywords: Time series forecasting, Regression algorithms, Foot traffic, Wireless access point, Bloom Intelligence

1 Introduction

Smart stores are an integral part of smart cities, which improve prosperity, save costs, reduce wastage and resource usage, and improve quality of life [15]. Demand and sales forecasting are one of the important inputs for smart stores. Based on foot traffic predictions, businesses can adjust staffing and product stock levels [14]. The further a business can accurately forecast foot traffic into the future, the more it can optimize operations management (example: labor scheduling), product management (example: stock levels) and in consequence grow profits [13, 17]. Therefore, an accurate forecast of foot traffic, by a store manager (for example a restaurant), can help increase customers' satisfaction, increase sales, and reduce food waste. Foot traffic forecasting is valuable for independent retail stores, franchises, and corporate chains alike. Accurate foot traffic information and predictions can also help reduce energy usage and improve safety in offices and buildings [8].

Compared to stock forecasting, food sales forecasting, and similar problems, the subject of foot traffic prediction has received limited research focus. The main reason for insufficient research in foot traffic forecasting is lack of real

P. Perner (Ed.): ICDM 2017, LNAI 10357, pp. 107–121, 2017.
DOI: 10.1007/978-3-319-62701-4_9

world data. To our knowledge there does not exist a public dataset that has foot traffic per minute, hour, or day, for a large number of stores. Due to lack of data, restaurants and physical stores have not had a reliable way of predicting foot traffic. To collect large scale and reliable foot traffic data, a comprehensive system with sensors is required. In this paper, we describe such a scalable data collection system that we designed and used for our experiments.

A study by Veeling [18] worked on improving foot traffic forecasting in 11 retail stores with neural networks. In the study, security gates are used to count people entering or exiting the physical stores. Beams placed at the gates are broken when people walk through them. To keep complexity low, only daily foot traffic is used in the prediction model. The study uses the Non-linear AutoRegressive with Exogenous inputs (NARX) model to predict one-step-ahead. In a different study Cortez et al. [10] used digital cameras to detect foot traffic at a sports store. Data mining methods are applied to build foot traffic forecasting models. The camera is linked with a human facial recognition system which counts the foot traffic and groups the traffic into three categories: all faces, female faces and male faces. The daily foot traffic combined with other factors like weather and special events are used to build a prediction model. They compare six forecasting methods, and their model predicts daily up to 7 days in advance. Further, we found a patent application for predicting traffic at a retail store [2], which indicates that our problem is of practical interest but not well researched.

A problem closely related to foot traffic prediction is retail sales forecasting [1]. Accurate demand forecasting is used to help retail businesses to organize and plan production. Time series forecasting models are often used for retail sales forecasting, but improving quality of forecasts still remains a challenge. Forecasting accuracy can be impacted by different factors such as time, weather conditions, economic factors, random cases, etc [15]. Several methods such as Winters exponential smoothing, the Autoregressive Integrated Moving Average (ARIMA) model [3], multiple regression, and artificial neural networks (ANNs) have been the most widely used approaches to time series forecasting because these models have the ability to capture trend and seasonal fluctuations present in aggregate retail sales. However, all these methods have shown difficulties and limitations. Therefore, it is necessary to investigate further how to improve the quality of forecasts. For sales predictions, data is usually collected when an order is placed, the transaction is automatically processed through the point of sale (POS) system, and then stored in a database.

In this research, we aim to predict hourly foot traffic based on historical data that contain foot traffic of stores for every hour. Data for our work was gathered by Cisco Meraki wireless access points that are installed in different stores, such as gyms, restaurants and bars. WiFi access points installed in businesses detect smartphones with WiFi turned on, whether or not a user connects his or her device to the wireless network. However, the device's presence can only be detected while the device is within range of the network [7]. After preprocessing the raw data, we built a model for prediction. We choose traditional machine learning approaches rather than time series analysis such as ARIMA. Time series models with multi-step-ahead forecast have high error rate [6], and would reduce the effectiveness of the prediction model beyond one step-ahead (one hour-ahead). Our goal is to predict foot traffic each hour of the next 7 days; we are predicting $h = 168$ steps ahead (hour).

We are able to predict hourly foot traffic for the next hour, day and week by using historical data and regression algorithms such as Random Forest Regression [4] and Support Vector Regression [9]. The variation in foot traffic can be caused by different factors such as weather and holidays. Hence, we consider these influencing factors in the forecasting model to improve forecast accuracy. The prediction models learn from the collected data for different stores to see how the prediction models work for each different type of business and location.

The remainder of this paper is organized as follows. Section 2 presents the description of data collection system in the study. Section 3 explains the prediction model and implementation of the developed model. Experimental design and results are presented in Section 4. Finally, conclusions and future work are described in Section 5.

2 Data Collection System

For this research raw data is collected from Cisco Meraki Wireless Access Points (AP) by Bloom Intelligence. The AP's are installed at customers of Bloom Intelligence such as gyms, bars, restaurants, etc. The raw data is processed and classified to extract foot traffic and other engagement metrics for the location. This data is then made available to Bloom Intelligence customers through an analytics dashboard.

2.1 Data Collection

The IEEE 802.11 specification for wireless communication provide a mechanism for devices to discover other compatible 802.11 devices. An 802.11 enabled WiFi device, a mobile phone for instance, would broadcast a probing requests that can be received by any other compatible device, a Cisco Meraki Wireless Access Points. The AP captures these probing request for each device when they are in range and is able to collect information about the device. Some of the data collected and utilized in this research is highlighted in table 1.

Table 1. Information collected from access points

AP MAC	MAC address of the observing AP
Client MAC	MAC address of the probing device
Seen Epoch	Observation time in seconds since the UNIX epoch
RSSI	Device receiver signal strength indication (RSSI) as seen by the AP. This determines the proximity of the device to the AP.

Devices create probing events at different intervals and can be affected by many factors such as the operating system in use, applications that are installed, etc. Cisco Meraki with their analytics partners have found that the request interval can vary greatly. From our collected data we have seen similar results [7].

The data collected by the AP are relayed to the Bloom Intelligence (BI) big data processing servers in an aggregated JSON format that defines each probe as an observation. At the BI servers the data is validated and normalized followed by classified and analytics generation. In the final analytics step the foot traffic per hour is aggregated and made available for the prediction model.

2.2 Preprocessing

The observations, collected by AP's, are analyzed using a proprietary heuristic to discover patterns. The heuristic looks for two types of patterns that classify the foot traffic as either a visit or passerby for a given location. The visit patterns are then analyzed to aggregate hourly foot traffic per hour and per location.

We will now briefly describe the preprocessing (classifying heuristic) shown in algorithm 1. Using the data defined in table 1, a passerby is defined as any device that is detected by the AP at least once. A visitor is a device that is considered present at the location for 5 minutes or more within a 20-minute window. A 20-minute window is initiated by an observation with an RSSI of at least 15. The window is maintained (device is present) when the observations has an RSSI of 10 or more. Once a window is classified as a visit the duration of the visit can be determined. The 20- minute window period is increased by checking for continued activity at the AP. If there is no activity for 20 minutes from the last observations time, the visit window is deemed terminated. Any new observations that occur after this 20-minute period would result in a new pattern analysis. The implication of this procedure is that a device can have multiple visits to a location within a time period like a day or an hour and, visits can also span multiple consecutive hours or days. Figure 1 shows seen devices that are classified as visitors.

If there are n visitors (client MAC address) and m stores (access points), time complexity for calculation of hourly visitor numbers will be $\mathcal{O}(nm)$. Since the number of observations of a particular client MAC address is small, we assume it to be $\mathcal{O}(1)$ in this analysis. In the DAM and cache-oblivious model, our algorithm runs in $\mathcal{O}(nm/B)$ I/Os [11].

This study uses hourly foot traffic of different retails locations that are collected from August 2015 to October 2016. The preprocessed statistics have precomputed the foot traffic for every hour for each calendar day in this period.

3 Our Prediction System

In this section we present our prediction model and the metrics for measuring the forecast error.

Algorithm 1: Calculate visitor numbers per hour for one access point

```
Input  : observations time O and rssi values corresponding to observations
         time R
Output: # of visitors per hour
  function CALCULATE_VISITOR_NUMBERS(O, R)
      total number of visitors = 0
      foreach client mac address do
          create a list of sorted observations time (already sorted in DB)
          create a list of rssi values corresponding with sorted observations
          while list of observations time and rssi are not empty do
              slides window till first observation with rssi ≥ 15
              ANALYSIS_THE_WINDOW_FOR_THE_PRESENCE(window)
              if window is visit then
                  extend the window till to 20 minutes gap
                  total number of visitors ++
                  reset the window
              end
          end
      end
      return total number of visitors
  end function
  function ANALYSIS_THE_WINDOW_FOR_THE_PRESENCE(window)
      if presence in window ≥ 5 minutes then
          return window is visit
      end
  end function
```

3.1 Forecast error measures

For forecasting models there is no general applicable accuracy measure because there are a variety of forecasting objectives, and also data scales and patterns are different [12]. Thus, in order to reduce possible bias generated by one single accuracy measure, in this study we use three measures, including root mean square error (RMSE), mean absolute percentage error (MAPE), and mean absolute error (MAE). If $\bar{Y}$ is a vector of n predictions, and Y is the vector of actual observation then RMSE, MAPE and MAE are defined as follows:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(\bar{Y}_i - Y_i\right)^2} \tag{1}$$

$$MAE = \sum_{i=1}^{n}\left|\bar{Y}_i - Y_i\right| \tag{2}$$

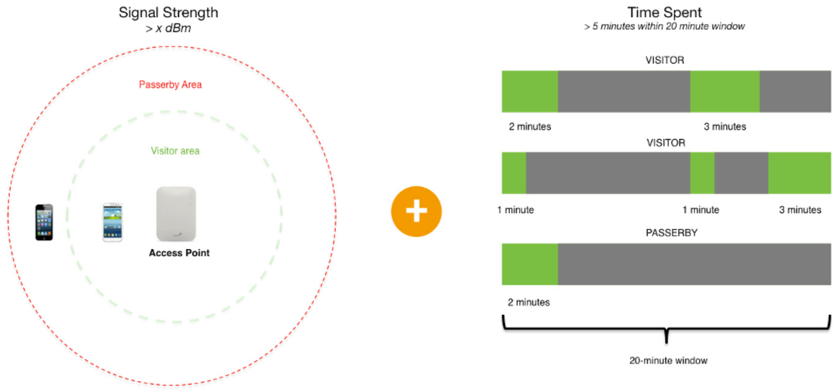

Fig. 1. Computing visitor state [7].

$$MAPE = \sum_{i=1}^{n} \left| \frac{\bar{Y}_i - Y_i}{Y_i} \right| * 100 \tag{3}$$

These measure have some advantages and disadvantages. For example the MAE is easy to understand, but it is not appropriate for comparing forecast performance between different data sets. MAPE is scale-independent and suitable for comparison of prediction accuracy of multiple data sets, however it has large value when any y_i is close to zero. The MAPE will be infinite if $y_i = 0$.

3.2 Prediction Model

After collecting and preprocessing data we have datasets which contain foot traffic per hour. In other words, for each retail store we have a dataset that shows us how many people visit the store every hour. Now we want to build a prediction model to forecast foot traffic of each store for every hour of the next seven days.

For building the prediction model we choose traditional machine learning approaches instead of using time series analysis such as ARIMA. We prefer using regression algorithms such as Random Forest Regression and Support Vector Regression because time series models have high error for multi-step-ahead forecast [6]. This work aims to predict foot traffic per hour for next week which means $h = 168$ hour-ahead (24 hours $\times$ 7 days $= 168$).

Function Learning. In order to use the regression algorithms for building the prediction model we block up the data into overlapping windows of time, then use it to predict the next days (or weeks). Moreover, we believe that future traffic will resemble past traffic and also foot traffic will be different depending on the

Table 2. Factors that may impact the foot traffic

Factor	**Range or an example of the factor**
Weather	Temperature, rainfall level, snowfall level
Holidays	Public holidays, school holidays
Special events	Happy hour, sport games, local concerts, conferences, other events
Location	Close to schools, tourist cities

time of year, day of week, and time of day. For instance, in a restaurant, traffic may be higher on weekends, or at specific times such as during periods of lunch or dinner. Therefore, we can conclude foot traffic on hour h and day d, where $0 \leq h \leq 23$ and $0 \leq d \leq 6$, is close to foot traffic on hour h and day d of last weeks. For example, if we want to predict foot traffic on Monday at 8 pm, we should look at the foot traffic on the previously observed Mondays at 8 pm.

Thus, based on this logic, we can define a function learning problem. Let foot traffic at timestamp t be equal to v_t. Our function f takes N inputs, and will output the foot traffic for an hour. The N inputs to f correspond to the foot traffic from the same hour and day in previous N weeks. Formally:

$$f(v_{t-w}, v_{t-2w}, v_{t-3w}, \ldots, v_{t-Nw}) = v_t$$

where $w = 24 \times 7 = 168$. Now we can easily use our data to solve for f using regressions of different types.

Building a Prediction Model Considering Factors. Once function learning is created for the datasets, we can use regression algorithms and build a model that can predict how traffic moved in our historic datasets. Historical data alone is not sufficient in producing accurate prediction values. The foot traffic can be affected by factors such as weather, special events and public holidays. In table 2 we list the factors that we have considered for our prediction model.

Weather: Before adding this feature to our prediction model, we performed some analysis to explore the potential relationships between the weather features and foot traffic, and no significant correlation observed. Moreover, because weather forecasts are inaccurate and they can cause error propagation, we do not use these features in our prediction model.

Holiday: We classify holidays into two categories: regular holidays and festival holidays. The regular holidays include the official four holidays in Florida (Martin Luther King Day, Memorial Day, Independence Day, and Labor Day). The festival holidays include New Year's day, Thanksgiving day, and Christmas day which the stores are usually closed. We believe based on the business type of stores the holiday effect needs to be considered for the day before and after the holiday as well. For example, for a bar, foot traffic is impacted before, on, and after the holiday.

In order to incorporate the effect of holidays into the model, for festival holidays we only need to return zeros for all 24 hours because the stores are

Algorithm 2: Prediction with holiday consideration

function PREDICTION_HOLIDAY($Data$)
  PREDICTION_WITH_RANDOM_FOREST_REGRESSION($Data$)
  DETECT_HOLIDAY_BEHAVIOR($Data$)
  for *the days that are holiday* **do**
    $rate \times (prediction\ results\ for\ the\ holiday)$
  return $updated\ prediction\ results\ for\ a\ week$
end function

function PREDICTION_WITH_RANDOM_FOREST_REGRESSION($Data$)
  return prediction results for a week
end function

function DETECT_HOLIDAY_BEHAVIOR($Data$)
  if *holiday behavior is close to Sundays* **then**
    $rate = Average(\frac{foot\ traffic\ on\ last\ Sundays}{foot\ traffic\ on\ last\ days\ that\ have\ same\ weekday\ as\ holiday})$
  else if *holiday behavior is close to last holidays* **then**
    $rate = \frac{foot\ traffic\ on\ the\ last\ holiday}{foot\ traffic\ on\ last\ day\ that\ has\ same\ weekday\ as\ holiday}$
  else
    $rate = 1$
  return $rate$
end function

usually closed. However, for regular holidays we need an approach to consider holiday effect. The simplest way for including holidays into the model is using dummy variables. For instance, normal days are coded as 0, regular holidays as 1 and festival days as 2. This simple way is effective with large data sets, but we have only around a year of data, and there are only a few holidays in one year. Instead, we use an alternative approach to improve the forecast accuracy for holidays.

Our alternative method for prediction with holiday consideration is using a rate, in a way that the prediction results are updated with a rate. Once, the prediction model returns results for holidays, the results are updated with a rate. Details regarding the rate calculation and prediction model with holiday consideration are shown in algorithm 2.

Special events: The impact of special events depends on the type of event. Events like a happy hour occur on a regular, predictable time, like a Friday from 5pm to 6pm, that can be captured by the prediction model. However, events like sport games, local concerts, and conferences can have irregular or one-off occurrences need to be considered as special days and their behavior could be considered similar to holidays. To incorporate their effects into the prediction model the same as holidays, dummy variables can be used.

Location: The location of some stores induces specific irregularities in foot traffic. For instance, those stores that are close to universities or located in tourist cities show complex behavior. See experiment 5 for more details.

Table 3. Data collected from different stores

Group #	Business type	Avg. foot traffic per hour	# of stores
1	Gym	72	4
2	Coffee shop	26	6
3	Restaurant	28	26
4	Bar	40	19
5	Barbershop	16	1

4 Experimental Results

This section presents several experiments that were conducted to evaluate the forecasting performance of the multiple regression algorithms using real-world store foot traffic data. In order to validate the general forecasting performance of our proposed prediction models, multiple experiments were performed to explore foot traffic prediction in different types of businesses and locations with a variety of foot traffic. Further, we do some experiments to show how holidays can affect foot traffic, and how our prediction model with holiday considerations can manage this irregular behavior.

We have access to data from over 100 different stores. However, in some cases data collection has only been active for 6 months or less. At the results we have chosen 56 different stores that their data collection has been performed for more than one year. The stores are categorized as: gym, restaurant, coffee shop, bar and barbershop. The collected data is from August 2015 to October 2016, and contain the aggregated foot traffic counts for each hour of the day in this period.

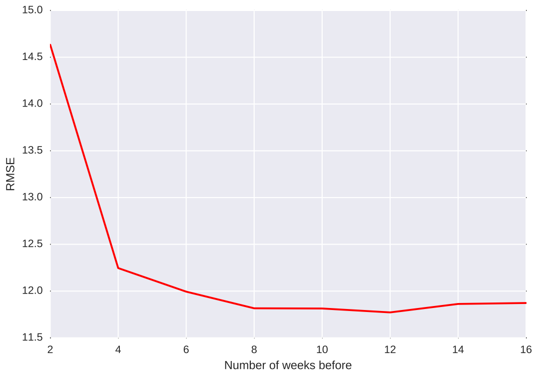

Fig. 2. Correlation between number of weeks before (N) and RMSE.

As mentioned in section 3.2, we had to compute a constant value of N (number of weeks before) for our function learning problem. There are 9504 (396 days $\times$ 24 hours) rows in our dataset. However, after creating the learning function described in section 3.2, number of rows in our dataset will be reduced. For choosing N we run an experiment to see the quality of the prediction models with different values of N. As figure 2 shows, after $N = 8$ there is no significant improvement in error. There is a trade-off between increasing N and computation time for function learning. Increasing N decreases our total dataset size for learning. Hence we chose $N = 8$ for our experiments. To create the learning function, we process each row, for which we need data from the 8 previous weeks. For the first 8 weeks there is no data. Therefore, the number of rows available for function learning, will be 9504 - (8 weeks $\times$ 7 days $\times$ 24 hours) = 8160. Details regarding the datasets are shown in table 3.

For all the experiments the datasets are partitioned into training and test datasets. We use Random Forest Regression from Python package sklearn [16], Support Vector Regression from Libsvm [5], and Google cloud predictor [1].

Now we explain 5 experiments that show the quality of our prediction model:

Experiment 1: In this experiment we compare the forecast measures for three different prediction models of the five random stores: a gym, restaurant, coffee shop, bar and barbershop. The values of RMSE, MAE and MAPE of the one-week forecasts are presented in table 4. Although support vector regression has the least error compared to other models for most stores, this difference is not significant and can be ignored.

Experiment 2: In order to compare error for different types of businesses, MAPE is compared for stores described in table 3. As table 5 illustrates, the gyms have around 13% errors, while coffee shops, restaurants, bars and barbershops have 20% , 24%, 22%, and 17% errors respectively. This is due to the types of stores (datasets); we believe the reason is that gyms have members that go to gyms routinely, therefore foot traffic patterns are less random.

Experiment 3: This experiment describes the comparisons of actual and predicted outputs of three prediction models for a gym and coffee shop. Figure 3 shows the forecasting outputs, generated by different models, of one-week in advance for gym. The three regression models work mostly the same, and there is not a significant difference between the outputs. In addition, as we can see in Figure 4, outputs generated using the three regression models are very close for coffee shop as well. For the sake of brevity, the plots for other types of businesses in experiments 3, 4 and 5 are not included.

Experiment 4: As we discussed before, foot traffic are affected by holidays. In this experiment we observe how prediction with holiday consideration improves the forecast accuracy. Figure 5 presents prediction results of a gym for the week of 9/4/2016 to 9/11/2016, of which the September 5, 2016 is a holiday (Labor day). As the plot shows, prediction results of the holiday that are updated with the rate are more accurate than forecasted outputs with no holiday consideration.

[1] https://cloud.google.com/prediction/

Table 4. Comparison of different prediction model

Business Type	Model	RMSE	MAE	MAPE
Gym	Random Forest Regression	16.622	12.22	0.122
	Support Vector Regression	14.856	10.5	0.101
	Google cloud predictor	14.929	10.910	0.114
Coffee shop	Random Forest Regression	6.80	4.66	0.178
	Support Vector Regression	6.988	4.75	0.178
	Google cloud predictor	6.892	4.821	0.179
Restaurant	Random Forest Regression	7.341	5.059	0.1965
	Support Vector Regression	7.057	4.744	0.173
	Google cloud predictor	6.879	4.654	0.163
Bar	Random Forest Regression	10.133	6.863	0.197
	Support Vector Regression	8.344	5.684	0.1657
	Google cloud predictor	8.593	6.702	0.202
Barbershop	Random Forest Regression	4.237	2.869	0.175
	Support Vector Regression	4.182	2.934	0.179
	Google cloud predictor	4.624	3.13	0.175

Table 5. Error comparison for different type of business

Group #	**Business type**	**MAPE**
1	Gym	13%
2	Coffee shop	21%
3	Restaurant	24%
4	Bar	22%
5	Barbershop	17%

For the non-holidays, solid red and blue line are overlapping and only the red one can be seen. The comparison of predicted values with and without holiday consideration for a restaurant are shown in figure 6. The prediction results are for the week of 7/2/2016 to 7/9/2016, of which 4th of July is a holiday.

Experiment 5: In this experiment we see how proximity to universities and locating in tourist cities can affect the prediction results. For a bar in downtown

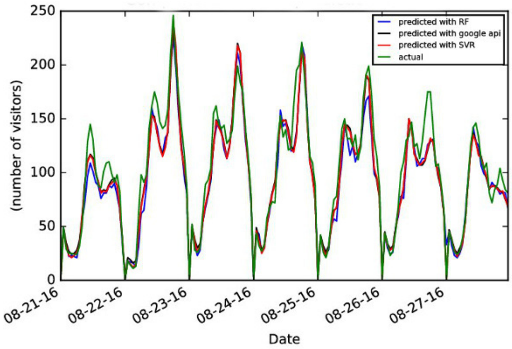

Fig. 3. Comparison of the predicted results using the three different models for a gym.

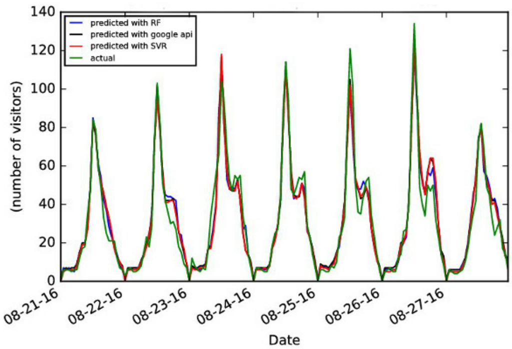

Fig. 4. Comparison of the predicted results using the three different models for a coffee shop.

Orlando (a tourist city) which has an average of 55 hourly foot traffic, MAPE by SVR is 30% which is more than some other bars that have MAPE of 22%. Also, for a bar close to Florida State University in Tallahassee with an average of 21 hourly foot traffic, MAPE by SVR is 31%. As figure 7 shows the traffic behavior is very random which causes poor prediction. Furthermore, for a coffee shop with an average of 21 hourly foot traffic which is close to FSU, the MAPE is 33%.

5 Conclusion and future work

In this paper, we presented a scalable data collection and prediction system to forecast hourly foot traffic, one-week in advance. Raw data for more than one year was gathered by using wireless access points installed in more than 100 stores. We preprocessed the raw data in order to calculate the foot traffic

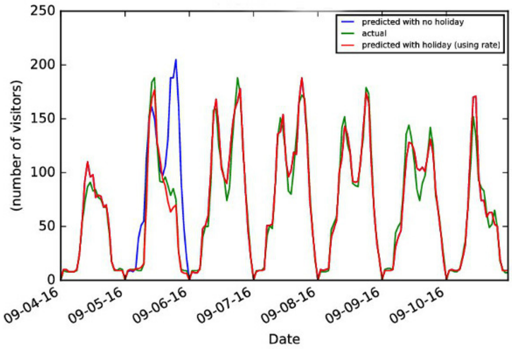

Fig. 5. Comparison of the predicted results with and without holiday consideration for a gym

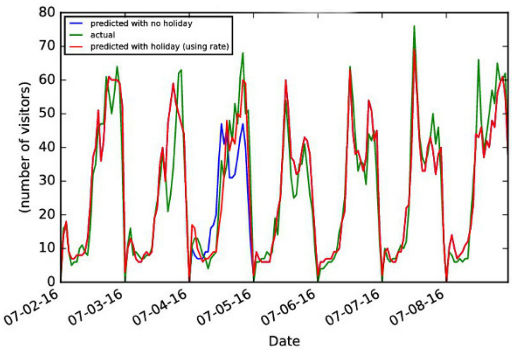

Fig. 6. Comparison of the predicted results with and without holiday consideration for a restaurant.

per hour. After preprocessing data, we used a regression algorithm to build a prediction model, predicting the foot traffic for the next 168 hours. Experiments show that the best results are for SVR, however a few other regression models exist that have accuracies close to SVR. Average error for hourly prediction for one-week in advance is 22%.

Future research will focus on three different aspects: handling the location impact on the forecast model, on top of the prediction model for the next 168 hours; separately build a forecast model only for the next day to improve the accuracy of the next 24 hours; and finally, develop the prediction model to a real-time system and integrate the model into the store environment.

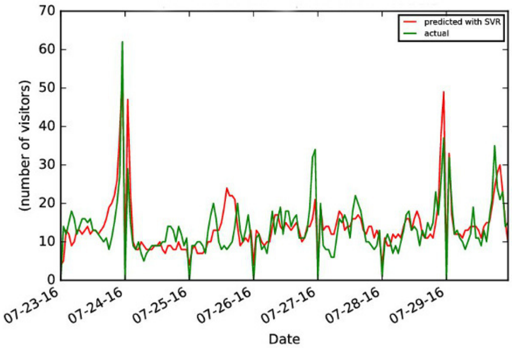

Fig. 7. Random traffic behavior which causes poor prediction.

6 Acknowledgments

The research for this paper was financially supported by Bloom Intelligence and Florida State University. William Wilson and Robin Johnston have been particularly helpful on this project from Bloom Intelligence. Multiple departments at Florida State University have helped with the administration of this research project, particularly the Dean's office at College of Arts and Sciences and the FSU Foundation.

References

1. Alon, I., Qi, M., Sadowski, R.J.: Forecasting aggregate retail sales:: a comparison of artificial neural networks and traditional methods. Journal of Retailing and Consumer Services 8(3), 147 – 156 (2001), http://www.sciencedirect.com/science/article/pii/S0969698900000114
2. Anand, J., Young, K., Choudhary, R., Howell, D.: Predicting shopper traffic at a retail store (Aug 9 2011), https://www.google.com/patents/US7996256, uS Patent 7,996,256
3. Box, G.E., Jenkins, G.M., Reinsel, G.C.: Time Series Analysis: Forecasting and Control. Wiley-Blackwell (2008)
4. Breiman, L.: Random forests. Mach. Learn. 45(1), 5–32 (Oct 2001), http://dx.doi.org/10.1023/A:1010933404324
5. Chang, C.C., Lin, C.J.: Libsvm: A library for support vector machines. ACM Trans. Intell. Syst. Technol. 2(3), 27:1–27:27 (May 2011), http://doi.acm.org/10.1145/1961189.1961199
6. Cheng, H., Tan, P.N., Gao, J., Scripps, J.: Multistep-ahead time series prediction. In: Ng, W.K., Kitsuregawa, M., Li, J., Chang, K. (eds.) Advances in Knowledge Discovery and Data Mining: 10th Pacific-Asia Conference, PAKDD 2006, Singapore, April 9-12, 2006. Proceedings, pp. 765–774. Springer Berlin Heidelberg (2006), http://dx.doi.org/10.1007/11731139_89
7. Cisco: Location Analytics (CMX). https://documentation.meraki.com (2015)

8. Clark, G., Mehta, P.: Artificial intelligence and networking in integrated building management systems. Automation in Construction 6(5), 481–498 (1997)
9. Cortes, C., Vapnik, V.: Support-vector networks. Mach. Learn. 20(3), 273–297 (Sep 1995), http://dx.doi.org/10.1023/A:1022627411411
10. Cortez, P., Matos, L.M., Pereira, P.J., Santos, N., Duque, D.: Forecasting store foot traffic using facial recognition, time series and support vector machines. In: International Joint Conference SOCO'16-CISIS'16-ICEUTE'16: San Sebastián, Spain, October 19th-21st, 2016 Proceedings, pp. 267–276. Springer International Publishing (2017), http://dx.doi.org/10.1007/978-3-319-47364-2_26
11. Frigo, M., Leiserson, C.E., Prokop, H., Ramachandran, S.: Cache-oblivious algorithms. In: Proceedings of the 40th Annual Symposium on Foundations of Computer Science. pp. 285–. FOCS '99, IEEE Computer Society, Washington, DC, USA (1999), http://dl.acm.org/citation.cfm?id=795665.796479
12. Hyndman, R.J., Koehler, A.B.: Another look at measures of forecast accuracy. International Journal of Forecasting 22(4), 679 – 688 (2006), http://www.sciencedirect.com/science/article/pii/S0169207006000239
13. Özgür Kabak, Ülengin, F., Aktaş, E., Şule Önsel, Topcu, Y.I.: Efficient shift scheduling in the retail sector through two-stage optimization. European Journal of Operational Research 184(1), 76 – 90 (2008), http://www.sciencedirect.com/science/article/pii/S0377221706010277
14. Lam, S., Vandenbosch, M., Pearce, M.: Retail sales force scheduling based on store traffic forecasting. Journal of Retailing 74(1), 61 – 88 (1998), http://www.sciencedirect.com/science/article/pii/S0022435999800888
15. Lasek, A., Cercone, N., Saunders, J.: Restaurant sales and customer demand forecasting: Literature survey and categorization of methods. In: Leon-Garcia, A., Lenort, R., Holman, D., Staš, D., Krutilova, V., Wicher, P., Cagáňová, D., Špirková, D., Golej, J., Nguyen, K. (eds.) Smart City 360°: First EAI International Summit, Smart City 360°, Bratislava, Slovakia and Toronto, Canada, October 13-16, 2015. Revised Selected Papers, pp. 479–491. Springer International Publishing (2016), http://dx.doi.org/10.1007/978-3-319-33681-7_40
16. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in python. J. Mach. Learn. Res. 12, 2825–2830 (Nov 2011), http://dl.acm.org/citation.cfm?id=1953048.2078195
17. Perdikaki, O., Kesavan, S., Swaminathan, J.M.: Effect of traffic on sales and conversion rates of retail stores. Manufacturing & Service Operations Management 14(1), 145–162 (2012)
18. Veeling, B.: Improving Visitor Traffic Forecasting in Brick-and-Mortar Retail Stores with Neural Networks. Bachelor thesis, Universiteit Twente (2014)