Forecasting Models of Emergency Department Crowding

Lisa M. Schweigler, MD, MPH, MS, Jeffrey S. Desmond, MD, Melissa L. McCarthy, ScD, Kyle J. Bukowski, MBA, BSIE, Edward L. Ionides, PhD, and John G. Younger, MD, MS

Abstract

Objectives: The authors investigated whether models using time series methods can generate accurate short-term forecasts of emergency department (ED) bed occupancy, using traditional historical averages models as comparison.

Methods: From July 2005 through June 2006, retrospective hourly ED bed occupancy values were collected from three tertiary care hospitals. Three models of ED bed occupancy were developed for each site: 1) hourly historical average, 2) seasonal autoregressive integrated moving average (ARIMA), and 3) sinusoidal with an autoregression (AR)-structured error term. Goodness of fits were compared using log likelihood and Akaike's Information Criterion (AIC). The accuracies of 4- and 12-hour forecasts were evaluated by comparing model forecasts to actual observed bed occupancy with root mean square (RMS) error. Sensitivity of prediction errors to model training time was evaluated, as well.

Results: The seasonal ARIMA outperformed the historical average in complexity adjusted goodness of fit (AIC). Both AR-based models had significantly better forecast accuracy for the 4- and the 12-hour forecasts of ED bed occupancy (analysis of variance [ANOVA] $p < 0.01$), compared to the historical average. The AR-based models did not differ significantly from each other in their performance. Model prediction errors did not show appreciable sensitivity to model training times greater than 7 days.

Conclusions: Both a sinusoidal model with AR-structured error term and a seasonal ARIMA model were found to robustly forecast ED bed occupancy 4 and 12 hours in advance at three different EDs, without needing data input beyond bed occupancy in the preceding hours.

ACADEMIC EMERGENCY MEDICINE 2009; 16:301–308 © 2009 by the Society for Academic Emergency Medicine

Keywords: crowding, forecasting, emergency service, hospital, operations research

Emergency department (ED) overcrowding has become a significant problem throughout the United States, leading to possible increased health care costs, causing raised stress levels among staff and patients in EDs, and most importantly, adversely affecting patient outcomes.[1–9] One aspect of the problem is the difficulty of anticipating the timing and magnitude of overcrowded conditions. The ability to predict crowded conditions, especially hour by hour, could substantially impact ED operations. To this end, we evaluate how time series–based models perform in short-term forecasting of ED occupancy.

Traditionally, ED operations directors have found historical averages to be reliable and accurate for long-term forecasts of ED behavior. For example, a director might use the average ED bed occupancy on Monday evenings at 21:00 over the past 2 years to determine how many staff should be working in the ED at that time. However, short-term forecasting of ED bed occupancy, such as might be useful for calling in additional staff or opening up hospital beds, is likely to need more accurate forecasting techniques.

Several authors have looked to time series techniques, such as autoregression (AR) models, as potentially useful tools in forecasting ED behavior (e.g., patient volume or arrivals, length of stay, or patient acuity) without needing the input of many different

From the Departments of Emergency Medicine (LMS, JSD, JGY) and Statistics (ELI), University of Michigan, Ann Arbor, MI; the Department of Emergency Medicine, The Johns Hopkins University School of Medicine (MLM), Baltimore, MD; and Administrative Consulting, William Beaumont Hospital (KJB), Royal Oak, MI.

Presented at the Society of Academic Emergency Medicine Annual Meeting, Chicago, IL, 2007.

This work was supported in part by a pilot grant from the University of Michigan Center for Computational Medicine and Biology (JGY) and by the Robert Wood Johnson Foundation, with whom LMS is a clinical scholar.

Received July 8, 2008; revisions received September 25 and November 11, 2008; accepted November 12, 2008.

Address for correspondence and reprints: Lisa M. Schweigler, MD, MPH, MS; e-mail: lschweig@umich.edu.

doi: 10.1111/j.1553-2712.2009.00356.x
ISSN 1069-6563
PII ISSN 1069-6563583

predictor variables.[10–14] The premise of these models is straightforward: an ED's level of activity in the near future is strongly correlated to its activity now. These studies show that, in general, time series methods provide better statistical fit than more traditional approaches such as multivariate linear regression or historical experience. However, most time series studies of ED behavior have remarked on the ability of time series models to closely fit past events; performance against future behavior has not typically been included. Furthermore, time series approaches have not yet been used to directly investigate ED *crowding*, instead modeling related behaviors such as patient arrivals per hour. To our knowledge, only one group has studied the effectiveness of time series methods for predicting future behavior, and in that work they focused on total *daily* occupancy, rather than hourly forecasts such as might be useful for "higher resolution" real-time operations management.[10]

We contend that three chief requirements of a useful ED crowding forecasting model are 1) that it can be used at different EDs with varying operations environments, 2) that it performs significantly better than the hourly historical average, and 3) that it requires the smallest amount of information possible with which to make predictions. Although large multivariate models no longer constitute a computationally challenging problem, for short-term forecasting purposes they would require a continuous supply of high-fidelity data, often streaming from different administrative units (e.g., hospital bed control, operating rooms). At present, departments with access to such dense operational informatic resources are rare. Our requirement of parsimony of information led us to choose ED occupancy as our crowding metric over other more complex metrics such as the National Emergency Department Overcrowding Scale (NEDOCS) or the Emergency Department Work Index (EDWIN).[15,16]

In this study, we addressed the following questions: How do AR models perform compared to the hourly historical average in forecasting (up to 12 hours into the future) the occupancy of an ED? Furthermore, do some models perform better at one institution than another? Is there a model that performs sufficiently well regardless of the ED toward which it is applied such that it might constitute a standard? Finally, how far back in time should a model look to generate the best forecasts? Too short of a training time may impair performance by generating imprecise parameter estimates. Too long of a training time may prevent a model from adapting to very recent changes in occupancy behavior. The models were evaluated for the accuracy of their 4- and 12-hour forecasts for a year's worth of Monday evenings at three large teaching hospitals, which tend to be the most crowded times for many EDs.[17–20]

METHODS

Study Design and Setting

We conducted a multicenter retrospective analysis of hourly clinical activity at three adult EDs (Site 1 annual census, 98,199; Site 2 annual census, 59,344; and Site 3 annual census, 55,757). No patient- or provider-level identifying information was included, and therefore the study was considered exempt from informed consent requirements by the institutional review boards at all three sites.

Study Protocol

Hourly occupancy was defined as the number of patients within each adult ED and its waiting room divided by the number of permanent beds (excluding makeshift hallway beds, chairs, etc.) in that ED available during the hour in question. Patients who ultimately left before being evaluated were included in the counts while they were still registered as being in the ED. The hourly denominator was corrected for circumstances when greater or lesser numbers of beds were available in each ED (e.g., when ED-adjacent clinic space became available after normal clinic hours). Occupancy values for the adult EDs were collected retrospectively from each institution's clinical information system for the period beginning midnight, July 1, 2005, and ending 11:00 PM, June 30, 2006, resulting in 8,760 sequential hourly occupancy values for each center.

Evaluation of different statistical models was directed at their goodness of fit and their ability to make forecasts from 15:00 Monday through 02:00 Tuesday for 51 of the 52 Mondays included in the data set. These were times when all three sites frequently experienced occupancy levels that were higher and less predictable than at other times during the week and thus represented a stringent test platform.

Data Analysis

Analysis was conducted in R 2.7.1 (Comprehensive R Archive Network, http://cran.r-project.org) and Matlab R2008a (The Mathworks, Inc., Natick, MA). Prior to building AR-based models, diagnostic time domain analyses were performed to identify dominant frequencies within each site's occupancy behavior (data not shown). As discussed under Results, 24-hour periodicity was the primary mode at each site, and subsequent time domain models were limited to this frequency.

Following these preliminary model checks, we evaluated in detail three models of ED crowding, including the hourly historical average and two autocorrelation models. The specifications of the models are included in Data Supplement S1 (available as supporting information in the online version of this paper). In brief, they were the historical average, which is the mean occupancy for each site for each hour of the day: a 24-hour seasonal model (seasonal autoregressive integrated moving average [ARIMA] (1,0,1)/(0,1,1)), where occupancy at any time is a function of occupancy both 1 and 24 hours prior, and a sinusoidal model with an AR-structured error term, where occupancy at any time is a function of a 24-hour period sine wave fit to each ED's diurnal pattern and combined with 1-hour AR. The standard descriptive notation for ARIMA models is ARIMA(p,d,q), where p denotes the number of autoregressive parameters, d is the number of differencing passes, and q is the number of moving average parameters. A seasonal ARIMA is described by

ARIMA(p,d,q)/(sp,sd,sq) in which sp, sd, and sq provide the additional information on the seasonal autoregressive, differencing, and moving average components of the model, respectively. The two AR-based models were specifically chosen because they account most parsimoniously for both the 24-hour periodicity of ED occupancy behavior and the strong predictiveness of a previous hour's occupancy on the next hour's occupancy.

Each model was evaluated in two ways, as summarized in Figure 1. Goodness of fit was evaluated retrospectively using log-likelihood values across the ensemble of 51 Monday evenings in the data set by maximum likelihood regression to the 168 hours (7 days) prior to Monday, t = 15:00. Details of these calculations are also included in Data Supplement S1.

The second means of evaluating model performance was to consider prospective accuracy. As represented graphically in Figure 1, each model was trained on a defined number of hours of prior ED occupancy (annotated as goodness-of-fit domain in the legend to Figure 1) and then allowed to generate a forecast of ED occupancy for a subsequent number of hours (forecast domain). To build the models, we only used observed ED occupancy from the training period; no data from the subsequent prediction period were used in building the predictions. Thus, the forecasting performance was prospectively evaluated in a virtual manner from previously collected data.

Forecast accuracy was examined over 51 consecutive Monday evenings for all three sites over the study year. A forecast was defined as a prediction of ED occupancy either 4 or 12 hours beyond the available data, which in each case was artificially cut off at t = 14:00 for each study day (15:00 was therefore the first hour of forecast). Accuracy was quantified by comparing the predicted occupancy to the actual occupancy during the forecast and calculating the error as the root mean forecast sum of squares,

$$\varepsilon_{RMS} = \left(\frac{1}{K}\sum_{k=1}^{K}(x_k - \mu_k)^2\right)^{1/2}$$

where k is each hour of a K-hour forecast, x_k is the actual occupancy, and μ_k is the model-predicted occupancy.

To determine the impact of duration of training time (i.e., the number of hours of occupancy behavior provided to a model to allow predictions), a series of 4- or 12-hour forecasts of occupancy from 15:00 Monday to 02:00 Tuesday were made with an increasingly greater number of training hours, from 168 hours (7 days) to 336 hours (14 days). Forecast root mean square (RMS) error was calculated as described previously, and the mean RMS errors for each site were determined for each training period.

RESULTS

Table 1 summarizes key operational characteristics of the three study sites during the study period July 2005–June 2006, both at the ED and at the hospital level. Appreciable differences are seen in total ED volume, number of ED patients per ED bed per year, number of available inpatient beds, average weekday adult inpatient bed occupancy, percentage of days in study period with inpatient bed occupancy greater than 95%, attending and resident staffing hours, left-before-evaluation rates, size of observation unit, and percentage of patients seen in a minor care area.

A summary of the occupancy data at the three sites is shown in Figure 2. The clinical activity at each site is depicted as a heat map scaled over each day or over each week of the study frame. These images show that 1) the occupancy patterns differ between the three institutions and 2) all three institutions show diurnal variation in bed occupancy, with the busiest times occurring later in the day.

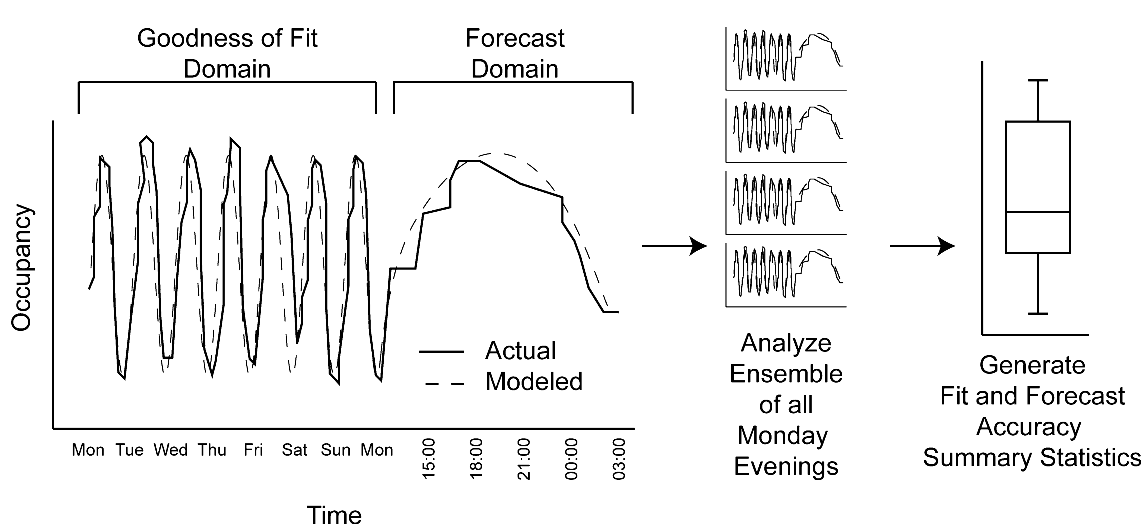

Figure 1. General analytical strategy. For each study center, data were segmented into 1-week segments. For each segment, three statistical models were examined for goodness of fit and for 4- and 12-hour forecast accuracy. From this ensemble of model fits and model forecasts, summary statistics including log likelihood, Akaike's Information Criteria (AIC), and root mean square (RMS) accuracy forecasts were generated.

Table 1
Operational Characteristics of the Adult EDs/Hospitals at the Three Study Sites during the July 1, 2005–June 30, 2006, Study Period

Operational Characteristics	Site 1	Site 2	Site 3
Total adult ED census	98,199	59,344	55,757
Number of adult ED patients/ED bed/year	1,488	1,091	1,742
Number of adult inpatient beds	901	946	502
Average weekday adult inpatient bed occupancy	85%	84%	96%
Number of days in study period with adult inpatient bed occupancy >95%	22	4	184
Percentage of adult ED patients admitted	32%	23%	29%
Hours of adult ED faculty staffing*	83	55	51
Hours of adult ED resident staffing*	39	101	90
Adult left without being seen rate (yearly)	1%	3%	4%
ED-based observation unit?	Yes	Yes	Yes
Number of beds in ED observation unit	21	14	16
Dedicated pediatric† ED space?	Yes	Yes‡	Yes
Percentage of overall ED patients seen in pediatric ED	14%	NA	25%
Percentage of pediatric patients seen in adult ED	6%	2%	2%
Minor care area?§	Yes	Yes	Yes
Percentage of total patients treated in minor care area	20%	25%	7%

ED = emergency department; NA = not applicable.
*Total number of hours per day, e.g., two physicians in ED at all times = 48 hours/day, observation unit staffing not included.
†Under age 18 = pediatric.
‡Site 2 has a pediatric ED run and staffed entirely by the Department of Pediatrics; it is operationally separate from the Department of Emergency Medicine; therefore, pediatric ED–specific data are not reported.
§For example, fast track or urgent care.

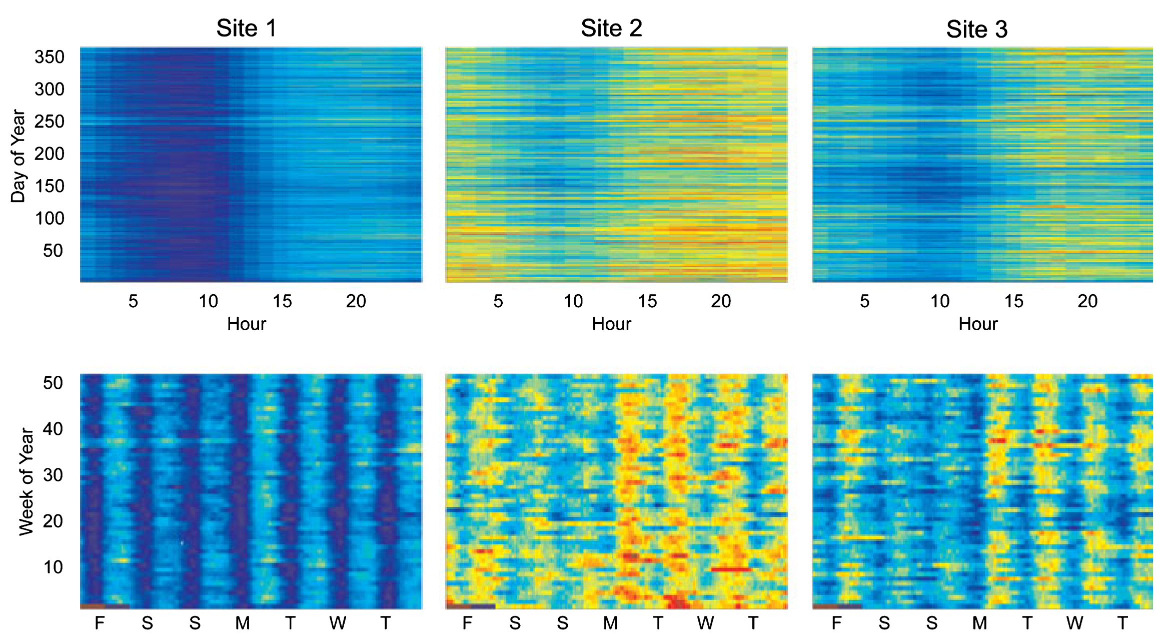

Figure 2. Heat map representations of the study period, midnight, July 1, 2005, through 23:00, June 30, 2006. (Top panel) Hour of the day is shown on the horizontal axis and days of the study year advance from bottom to top along the vertical axis. Occupancy is depicted by color, with dark blue showing least occupied and bright red showing most occupied. All three sites share the same color scale. (Bottom panel) Similar representation showing occupancy trends by day of week.

Table 2 shows the parameter estimates for the AR-based models. Occupancy means calculated for the historical averages model are not reported.

Goodness-of-fit measures are shown in Table 3. The historical average model consistently produced the best goodness of fit. However, the seasonal ARIMA (1,0,1)/(0,1,1) performed best according to Akaike's Information Criterion (AIC), which penalizes for increased model complexity. In this case, the historical average model included 24 model parameters (1 for each hour of the day), while the sinusoidal model with autocorrelated error required only 4 to make its

Table 2
Parameters of the AR-based Models

Model	Site 1	Site 2	Site 3
Seasonal ARIMA (1,0,1)/(0,1,1)			
AR term	0.907 ± 0.031	0.910 ± 0.033	0.910 ± 0.038
MA term	0.336 ± 0.082	0.042 ± 0.083	0.053 ± 0.080
24-hr seasonal term	−0.840 ± 0.141	−0.772 ± 0.127	−0.831 ± 0.127
Sinusoidal with AR error			
AR term	0.927 ± 0.021	0.903 ± 0.104	0.898 ± 0.041
Intercept	0.601 ± 0.052	1.235 ± 0.104	0.982 ± 0.107
Sine component	0.257 ± 0.026	0.192 ± 0.419	0.221 ± 0.033
Cosine component	0.160 ± 0.025	0.064 ± 0.044	0.114 ± 0.036

AR = autoregression; ARIMA = autoregressive integrated moving average; MA = moving average.

Table 3
Goodness-of-fit Results for the Three Models Studied

Study Site	Historical Average	Seasonal ARIMA (1,0,1)/(0,1,1)	Sinusoidal with AR Error
Goodness of Fit by Log Likelihood			
1	292 ± 33	261 ± 8	281 ± 11
2	175 ± 31	144 ± 13	158 ± 14
3	191 ± 40	156 ± 12	175 ± 12
Goodness of Fit by AIC			
1	−536 ± 66	−514 ± 17	−551 ± 21
2	−302 ± 62	−281 ± 27	−305 ± 27
3	−335 ± 80	−306 ± 24	−340 ± 25

AIC = Akaike's Information Criterion; AR = autoregression; ARIMA = autoregressive integrated moving average.

predictions (an AR term, a sine coefficient, a cosine coefficient, and an intercept).

Forecast performance is summarized in Figure 3. The box plots depict accuracy, for each site and each model, over either 4 or 12 hours of prediction as the RMS of the summed residuals between observed and predicted occupancy values. These plots show the increase in accuracy achieved by moving from a simple historical average system to more sophisticated models. For each site, the accuracies of the three methods were compared with one-way analysis of variance (ANOVA), followed by post hoc comparisons with Tukey-Kramer statistics. For each site, the two AR-based models outperformed the historical average. In post hoc testing, no differences were noted between the sinusoidal-AR model and the seasonal ARIMA model at each site. Examination of the effect of training time on forecast accuracy revealed no significant benefit beyond 168 hour (1 week) training periods (data not shown).

DISCUSSION

In our study, we show that AR models with seasonal/sinusoidal adjustment consistently outperform the historical average in short-term forecasting of ED occupancy up to 12 hours in advance and do so at several different institutions. Although there is variability in model performance for any given Monday, the reductions in error are potentially important operationally. For example, in moving from a historical average to either the seasonal ARIMA model or the sinusoidal model with AR-structured error, Site 2 would see a roughly 33% improvement in its 12-hour forecasts of ED crowding. We therefore posit that AR-based models should constitute the standard for predictive models, using time-series approaches and ED occupancy as the crowding metric. We found that a training time of 1 week (168 hours) was sufficient to build a model with excellent reliability.

It is not surprising that the two autoregressive models outperformed the historical average: while the historical average is an easy-to-understand approach to predicting long-term future ED volume and occupancy and has a well-founded theoretical basis in queuing theory,[21] it cannot be expected to perform well in situations where there is frequent irregularity in short-term behavior, e.g., unusually high ED occupancy on a Monday night, or increased demand on ED services during a virulent cold and flu season.

The ED literature describes many possible ED crowding metrics, such as staff opinion, leaving-before-evaluation rates, amount of time on ambulance diversion, ED length of stay, or calculated scores such as NEDOCS or EDWIN.[15,16,22–25] However, many of these measures of ED crowding are not easily obtained from EDs that do not systematically make an effort to collect such information. Most EDs do keep records on when patients present to the department and when they depart. From this operational information it is straightforward to calculate ED bed occupancy, defined as number of patients in the ED over number of permanent treatment bays available to that ED. It has been shown that the measure of ED bed occupancy performs no worse than more complex scores such as EDWIN in identifying ED crowding.[25–27]

Another important consideration in the development of this study was how to interpret the ED bed occupancy metric: should it be treated as a continuous metric or should a threshold approach be used, in which either an ED is crowded or it is not? While some important ED performance metrics may become abnormal at an easily discerned threshold occupancy level,

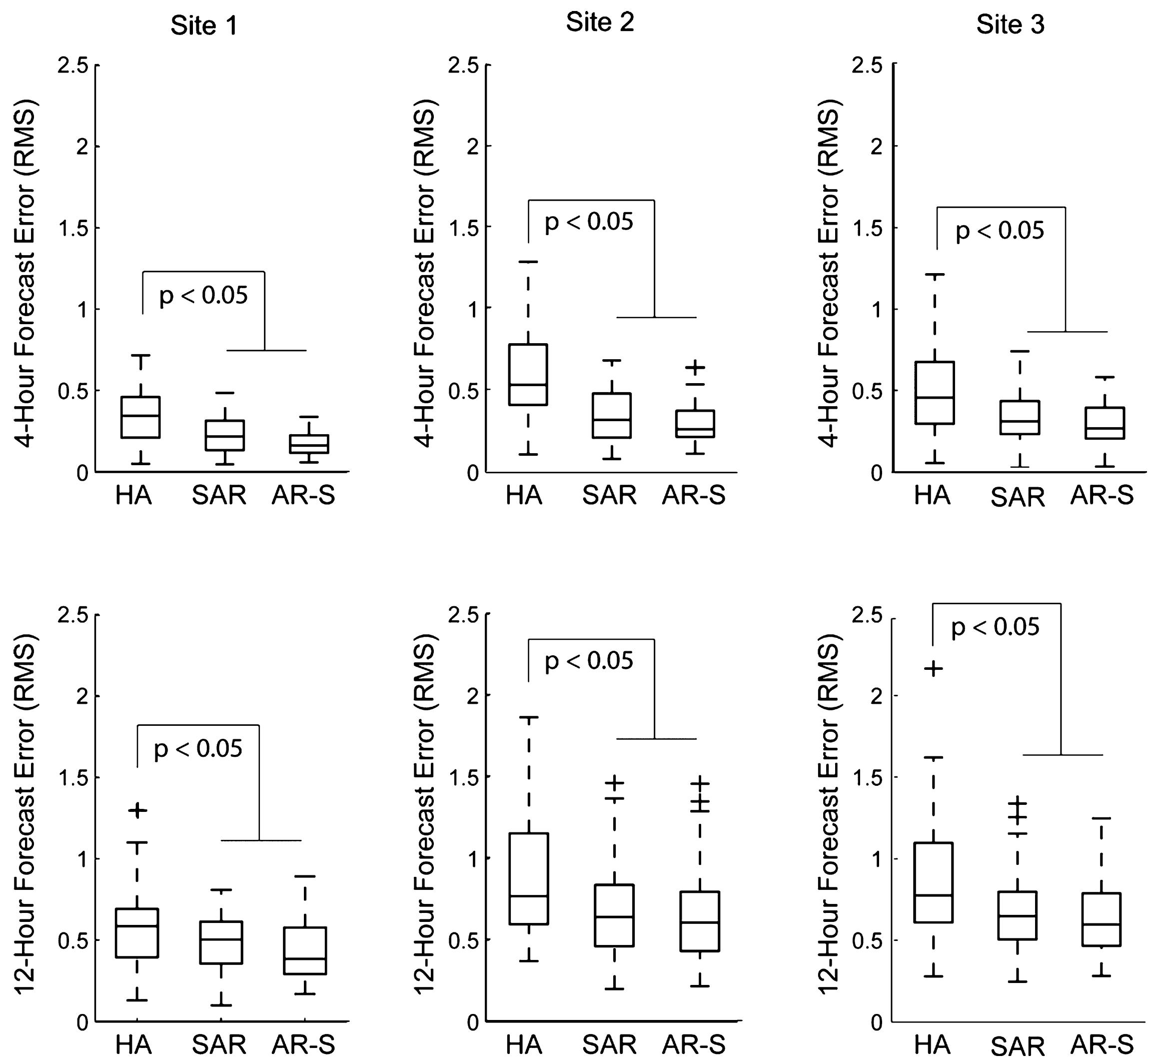

Figure 3. Twelve-hour forecast accuracy for the three models, by study site. The differences between predicted and actual ED occupancy from 15:00 Monday to 02:00 Tuesday for 51 weeks were quantified using the root-mean sum of squares. Models include HA (historical average), SAR (seasonal autoregressive integrated moving average [ARIMA] (1,0,1)/(0,1,1)), and sinusoidal with AR-structured error term (AR-S). One-way analysis of variance (ANOVA) was used to compare the different methods at each site and returned a p-value of < 0.01 for each instance. Post hoc comparisons were performed with the Tukey-Kramer method. The p-values for these results are shown in the figures.

a crowding metric that can supply a universally applicable threshold of "this ED is now crowded" remains to be developed. The value of picking a dichotomous outcome (e.g., crowded or not, on ambulance diversion or not) is attractive in that it permits evaluating forecast strategies with receiver operating characteristic approaches, but risks limiting generalizability of the forecasting scheme. The goal of our study was not to predict when a participating ED would reach a specific threshold of crowding, but rather to formulate predictions of different occupancy levels. It remains up to the individual institutions using any of the ED crowding metrics to interpret the meaning of the values obtained. Once such a level is established, our results indicate that an AR-based forecasting rule would perform well.

We considered possible applications of short-term forecasts of ED occupancy. For example, some institutions have successfully implemented a "dashboard" approach, in which ED and hospital administrators can make immediate patient flow and resource allocation decisions based on real-time ED and hospital operations data displayed on a computer interface (the dashboard), such as ED volume and hospital bed occupancy.[16,28] The addition of accurate and frequent short-term forecasts of ED crowding as a dashboard tool could be invaluable in helping administrators mitigate crowded conditions. Short-term forecasts of ED crowding could also prove valuable in regionalizing ambulance traffic. In its examination of the current state and the future of emergency care in the United States, the 2006 Institute of Medicine Report "Future of Emergency Care" called for refinement of methods to enable regional coordination of patient flow between different EDs to help alleviate crowding.[29] Complementary cornerstones of effective regionalization would be up-to-the-minute knowledge of crowding across EDs and a reliable means of predicting their status in the near future. The latter point is critical; delivering patients to an ED that is currently less than fully occupied but is likely to become so in the near future may not be the most effective triage choice.

Investigators have studied approaches other than historical averages and time-series analysis to forecasting ED behavior. The recent literature discusses numerous methods ranging from multivariable regression analysis to nonlinear techniques, discrete event simulation, and neural networks.[11,19,27,30–35] As reported, all of these approaches function reasonably well in providing short-term forecasts of various lengths for a variety of ED operational characteristics. However, many of these models may use proprietary software and often require input of many operational variables, some from outside of the ED, to generate their forecasts. The AR-based models shown to perform well in our analysis do not have these problems—we demonstrated that they require only one input variable, and they use widely available open-source software (R 2.7.1, Comprehensive R Archive Network, http://cran.r-project.org).

LIMITATIONS

It is important to note that the evaluation of the models in this study may be limited by the operational similarities between the three sites studied. All were relatively large, tertiary academic referral institutions. However, the hospitals under consideration are located in varied socioeconomic surroundings and therefore are likely to have different demands placed on them at different times in terms of patient presentations and bed availability. Table 1 shows that there are large differences in some operational characteristics between the three study sites. As seen in Figure 2, the three sites were found to have clear differences in the magnitude of both occupancy and variability of occupancy. Despite these differences in operational characteristics of the three EDs, the models performed similarly relative to each other at all three sites, adding weight to the argument that they might also work in a similar way at other comparable institutions. We would emphasize, however, that extrapolating these models to very different departments, and in particular low-volume sites where departmental occupancy may be zero over several sequential hours, should be done with care.

The premise of this study was to develop AR-based models of ED occupancy that are parsimonious and applicable to any ED. As with any modeling problem, the modeler faces a trade-off between model complexity, practicality, and generalizability. In this study, it is possible that higher-order ARIMA models may have provided even further reduction in the forecasting error than achieved with our approach. Indeed, a department wishing to undertake a systematic statistical consideration of this problem might well explore a large ensemble of related models, possibly incorporating locally available real-time operational covariates to determine which is optimal in their setting. The models explored here are not the final answer, but a reasonable platform from which to move forward.

We evaluated the performance of our models in two ways: model goodness of fit via log likelihood and the AIC and actual forecast performance via ANOVA of the RMS error of the different approaches. The goodness-of-fit measurements did not conclusively favor the AR-based models, but in this instance their interpretation is subtle. Specifically, while the historical average appears to be a very simple model (in that one could easily calculate it by hand), in actuality it is one including 24 separate parameters that must be fit to observed data. As a result, it is significantly advantaged in the calculation of log likelihood and similarly disadvantaged in calculation of AIC, which rewards model simplicity. However, this discussion is in part academic: the AR-based models performed clearly better than the historical average in their *forecast* accuracy, which is arguably a much more meaningful metric to individuals responsible for clinical operations than goodness of fit.

The models developed in this study were designed to be tools for operations managers that might help them decide when to institute interventions to mitigate ED crowding in their individual institution or regional ambulance network. The model predicts ED occupancy but does not provide insight into causes or consequences, nor would it be expected to shed light on site-specific solutions. However, as the model provides a parameterization of occupancy trends over time, it could readily be implemented as a statistical instrument for evaluating operational changes whose effects might be more complex than simply reducing absolute occupancy.

CONCLUSIONS

Using only preceding ED bed occupancy as input, AR-based models with seasonal or sinusoidal adjustment generated robust short-term forecasts of subsequent ED bed occupancy. This forecasting method was found to work equally well at three different institutions with differing operational characteristics, without having to adjust any of the model input variables. The simplicity of this approach makes it attractive for implementation in various applications such as regional out-of-hospital, ED, and hospital operations.

References

1. Institute of Medicine (U.S.). Committee on the Future of Emergency Care in the United States Health System. Hospital-based Emergency Care: At the Breaking Point. Washington, DC: National Academies Press, 2007.
2. Rondeau KV, Francescutti LH. Emergency department overcrowding: the impact of resource scarcity on physician job satisfaction. J Healthc Manag. 2005; 50:327–40.
3. Schull MJ, Vermeulen M, Slaughter G, Morrison L, Daly P. Emergency department crowding and thrombolysis delays in acute myocardial infarction. Ann Emerg Med. 2004; 44:577–85.
4. Pines JM, Hollander JE, Localio AR, Metlay JP. The association between emergency department crowding and hospital performance on antibiotic timing for pneumonia and percutaneous intervention for myocardial infarction. Acad Emerg Med. 2006; 13:873–8.
5. Fee C, Weber EJ, Maak CA, Bacchetti P. Effect of emergency department crowding on time to

antibiotics in patients admitted with community-acquired pneumonia. Ann Emerg Med. 2007; 50:501–9.
6. Sprivulis PC, Da Silva JA, Jacobs IG, Frazer AR, Jelinek GA. The association between hospital overcrowding and mortality among patients admitted via Western Australian emergency departments. Med J Aust. 2006; 184:208–12.
7. Pines JM, Hollander JE. Emergency department crowding is associated with poor care for patients with severe pain. Ann Emerg Med. 2008; 51:1–5.
8. Krochmal P, Riley TA. Increased health care costs associated with ED overcrowding. Am J Emerg Med. 1994; 12:265–6.
9. Bayley MD, Schwartz JS, Shofer FS, et al. The financial burden of emergency department congestion and hospital crowding for chest pain patients awaiting admission. Ann Emerg Med. 2005; 45:110–7.
10. Tandberg D, Qualls C. Time series forecasts of emergency department patient volume, length of stay, and acuity. Ann Emerg Med. 1994; 23:299–306.
11. Jones SS, Thomas A, Evans RS, Welch SJ, Haug PJ, Snow GL. Forecasting daily patient volumes in the emergency department. Acad Emerg Med. 2008; 15:159–70.
12. Champion R, Kinsman LD, Lee GA, et al. Forecasting emergency department presentations. Aust Health Rev. 2007; 31:83–90.
13. Milner PC. Forecasting the demand on accident and emergency departments in health districts in the Trent region. Stat Med. 1988; 7:1061–72.
14. Milner PC. Ten-year follow-up of ARIMA forecasts of attendances at accident and emergency departments in the Trent region. Stat Med. 1997; 16:2117–25.
15. Weiss SJ, Derlet R, Arndahl J, et al. Estimating the degree of emergency department overcrowding in academic medical centers: results of the National ED Overcrowding Study (NEDOCS). Acad Emerg Med. 2004; 11:38–50.
16. Bernstein SL, Verghese V, Leung W, Lunney AT, Perez I. Development and validation of a new index to measure emergency department crowding. Acad Emerg Med. 2003; 10:938–42.
17. National Center for Health Statistics. National Hospital Ambulatory Medical Care Survey. Centers for Disease Control, 2005. Available at: http://www.cdc.gov/nchs/about/major/ahcd/ahcd1.htm. Accessed Sep 23, 2008.
18. Asaro PV, Lewis LM, Boxerman SB. Emergency department overcrowding: analysis of the factors of renege rate. Acad Emerg Med. 2007; 14:157–62.
19. Hoot NR, Leblanc LJ, Jones I, et al. Forecasting emergency department crowding: a discrete event simulation. Ann Emerg Med. 2008; 52:116–25.
20. McCaig L, Burt C. National Hospital Ambulatory Medical Care Survey: 2002 emergency department summary. Adv Data. 2004; 340:1–34.
21. Little JD. A proof of the queueing formula $L = \lambda W$. Oper Res. 1961; 9:383–7.
22. Weiss SJ, Ernst AA, Derlet R, King R, Bair A, Nick TG. Relationship between the National ED Overcrowding Scale and the number of patients who leave without being seen in an academic ED. Am J Emerg Med. 2005; 23:288–94.
23. Andrulis DP, Kellermann A, Hintz EA, Hackman BB, Weslowski VB. Emergency departments and crowding in United States teaching hospitals. Ann Emerg Med. 1991; 20:980–6.
24. Hwang U, Concato J. Care in the emergency department: how crowded is overcrowded? Acad Emerg Med. 2004; 11:1097–101.
25. Jones SS, Allen TL, Flottemesch TJ, Welch SJ. An independent evaluation of four quantitative emergency department crowding scales. Acad Emerg Med. 2006; 13:1204–11.
26. McCarthy ML, Aronsky D, Jones ID, et al. The emergency department occupancy rate: a simple measure of emergency department crowding? Ann Emerg Med. 2008; 51:15–24.
27. Hoot NR, Zhou C, Jones I, Aronsky D. Measuring and forecasting emergency department crowding in real time. Ann Emerg Med. 2007; 49:747–55.
28. Spektor M, Ramsay C, Sommer B, Likourezos A, Davidson SJ. Electronic dashboard and a multidisciplinary hospital-wide team decrease patient throughput intervals and reduce number of admitted patients held in the emergency department [abstract]. Ann Emerg Med. 2008; 52:S87.
29. Institute of Medicine (U.S.). Committee on the Future of Emergency Care in the United States Health System. Emergency Medical Services at the Crossroads. Washington, DC: National Academies Press, 2007.
30. McCarthy ML, Zeger SL, Ding R, Aronsky D, Hoot NR, Kelen GD. The challenge of predicting demand for emergency department services. Acad Emerg Med. 2008; 15:337–46.
31. Hung GR, Whitehouse SR, O'Neill C, Gray AP, Kissoon N. Computer modeling of patient flow in a pediatric emergency department using discrete event simulation. Pediatr Emerg Care. 2007; 23:5–10.
32. Flottemesch TJ, Gordon BD, Jones SS. Advanced statistics: developing a formal model of emergency department census and defining operational efficiency. Acad Emerg Med. 2007; 14:799–809.
33. Asplin BR, Flottemesch TJ, Gordon BD. Developing models for patient flow and daily surge capacity research. Acad Emerg Med. 2006; 13:1109–13.
34. Hoot N, Aronsky D. An early warning system for overcrowding in the emergency department. AMIA Annu Symp Proc. 2006: 339–43.
35. Jones SA, Joy MP, Pearson J. Forecasting demand of emergency care. Health Care Manag Sci. 2002; 5:297–305.

Supporting Information

The following supporting information is available in the online version of this paper:

Data Supplement S1. Supplemental material.

The document is in PDF format.

Please note: Wiley Periodicals Inc. is not responsible for the content or functionality of any supporting information supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.