Academic Emergency Medicine
Official Journal of the Society for Academic Emergency Medicine

ORIGINAL RESEARCH CONTRIBUTION

Predicting Emergency Department Volume Using Forecasting Methods to Create a "Surge Response" for Noncrisis Events

Valerie J. Chase, MSE, Amy E. M. Cohn, PhD, Timothy A. Peterson, MD, MBA, and Mariel S. Lavieri, PhD

Abstract

Objectives: This study investigated whether emergency department (ED) variables could be used in mathematical models to predict a future surge in ED volume based on recent levels of use of physician capacity. The models may be used to guide decisions related to on-call staffing in non–crisis-related surges of patient volume.

Methods: A retrospective analysis was conducted using information spanning July 2009 through June 2010 from a large urban teaching hospital with a Level I trauma center. A comparison of significance was used to assess the impact of multiple patient-specific variables on the state of the ED. Physician capacity was modeled based on historical physician treatment capacity and productivity. Binary logistic regression analysis was used to determine the probability that the available physician capacity would be sufficient to treat all patients forecasted to arrive in the next time period. The prediction horizons used were 15 minutes, 30 minutes, 1 hour, 2 hours, 4 hours, 8 hours, and 12 hours. Five consecutive months of patient data from July 2010 through November 2010, similar to the data used to generate the models, was used to validate the models. Positive predictive values, Type I and Type II errors, and real-time accuracy in predicting noncrisis surge events were used to evaluate the forecast accuracy of the models.

Results: The ratio of new patients requiring treatment over total physician capacity (termed the care utilization ratio [CUR]) was deemed a robust predictor of the state of the ED (with a CUR greater than 1 indicating that the physician capacity would not be sufficient to treat all patients forecasted to arrive). Prediction intervals of 30 minutes, 8 hours, and 12 hours performed best of all models analyzed, with deviances of 1.000, 0.951, and 0.864, respectively. A 95% significance was used to validate the models against the July 2010 through November 2010 data set. Positive predictive values ranged from 0.738 to 0.872, true positives ranged from 74% to 94%, and true negatives ranged from 70% to 90% depending on the threshold used to determine the state of the ED with the 30-minute prediction model.

Conclusions: The CUR is a new and robust indicator of an ED system's performance. The study was able to model the tradeoff of longer time to response versus shorter but more accurate predictions, by investigating different prediction intervals. Current practice would have been improved by using the proposed models and would have identified the surge in patient volume earlier on noncrisis days.

ACADEMIC EMERGENCY MEDICINE 2012; 19:569–576 © 2012 by the Society for Academic Emergency Medicine

Emergency department (ED) crowding affects hospitals across the country.[1,2] It is viewed as the largest safety concern from an urban physician's standpoint,[3] occurring when demand for emergent patient care exceeds the available resources, compromising the care received in hospitals across the United

From the Industrial and Operations Engineering, College of Engineering, University of Michigan, Ann Arbor (VJC, AEMC, MSL), the Center for Healthcare Engineering and Patient Safety (AEMC), and the Department of Emergency Medicine (TAP), University of Michigan, Ann Arbor, MI.

Received September 1, 2011; revisions received November 9 and December 6, 2011; accepted December 10, 2011.
Presented at the 33rd Annual Meeting of the Society of Medical Decision Making Chicago, IL, October 23–26, 2011; and The Institute for Operations Research and Management Sciences (INFORMS) 2011 annual meeting, Charlotte, NC, November 13–16, 2011.
The authors have no relevant financial information or potential conflicts of interest to disclose.
Supervising Editor: Lawrence M. Lewis, MD.

Address for correspondence and reprints: Mariel S. Lavieri, PhD; e-mail: lavieri@umich.edu.

doi: 10.1111/j.1553-2712.2012.01359.x
ISSN 1069-6563
PII ISSN 1069-6563583

States.[4,5] ED crowding has been associated with decreased clinical performance of EDs as well as decreased patient satisfaction.[6–8] It not only degrades the quality of care patients receive in the ED, but also affects the transition from the ED to an inpatient floor.[9]

Events leading to ED crowding can be related to daily operational inefficiencies, disaster events, and non–disaster-related surges in patient volume. Conditions observed on these crowding days typically include long waits for patients, full ED bed occupancy, and significant demand on all provider capacity.[10,11]

The ability to predict when ED crowding will occur remains a high priority for many departments. While initial insights have been gained, previous studies have not had sufficient success in predicting ED volume in a real-time fashion.[12–15]

The purpose of this study was to use forecasting methods with real-time data to: 1) determine which indicators could be used to accurately model the state of the system and 2) determine how far in advance a significant increase in patient volume could be predicted to adequately plan and prepare to prevent a crowding situation.

METHODS

Study Design

This was a retrospective study that used patient data abstracted from an ED administrative database to build a model for forecasting ED crowding. This study was approved by the institutional review board through a waiver of informed consent, as no identifiable patient information was reported.

Study Setting and Population

This study involved data from a large, academic, tertiary care, urban, Level I trauma-verified ED with an annual adult patient census of approximately 69,000. The ED is divided into multiple treatment spaces; 46 beds are available for the primary assessment of new ED patients, including three resuscitation bays for critically ill patients. A nine-bed unit is used for holding admitted patients and for ED-based observation protocols.

All adult patients (≥18 years old) triaged to one of three main ED treatment areas from 12:00 AM July 1, 2009, through 11:45 PM November 30, 2010, were included in this analysis. Minor care, a physically separate treatment space, was excluded from this analysis.

Study Protocol

The main data sources for this study were: 1) the ED patient-specific report from the ED administrative database (Centricity; GE Healthcare, Waukesha, WI) and 2) the database of physician capacity scheduling. For all patients who arrived during the study time period, the following were retrieved: 1) date and time of check-in, 2) date and time a room was assigned to the patient, 3) date and time a physician was assigned to a patient, 4) date and time of patient disposition, 5) disposition category, and 6) date and time a provider accepted the patient on an inpatient ward (only for patients being admitted to an inpatient floor).

A database was created from average long-run historical information on resident and faculty productivity including number of residents and faculty in the ED throughout the day, and the rate at which physicians were able to see patients per hour, to develop a measure of physician capacity for use in constructing the models. This database was initially stratified by varying levels of experience (resident, physician assistant, nurse practitioner, faculty, etc.) and was associated with a long-run historical average treatment rate for each level of experience. The total number of providers at any given time was then multiplied by the respective treatment productivity for that provider's level of experience. The summation of all weighted averages was then used to represent the total physician capacity at any given point in time.

All of the aforementioned data were collected for a retrospective analysis to determine if overcrowding would occur and further, if it could be predicted. The models were developed with the July 2009 through June 2010 data (model creation data set) and were validated against the remaining July 2010 through November 2010 data (model validation data set).

Data Analysis

The data analysis was conducted in four phases. First, a database evaluation was completed to ensure all desired aspects affecting overcrowding were addressed using the observed data sets. Since the data did not naturally include information directly describing physician capacity with respect to other ED status-describing variables, the need for such a variable was identified. Therefore, to relate both data sets, a new variable termed ''care utilization ratio'' (CUR) was developed to explicitly convey the ratio of the number of new patients to be treated (i.e., new arrivals minus patients triaged) to the total estimated physician capacity for that period.

Data collected in the physician capacity database were aggregated into a total physician capacity for each time interval. This capacity was then used in conjunction with the new arrivals and number of patients triaged for the respective interval to generate the CUR for that time interval with the equation

$$\text{CUR} = \frac{\text{new arrivals - patients triaged}}{\text{total physician capacity}}.$$

This ratio indicates, at a given point in time, whether or not the physician capacity of the ED was capable of treating all patients currently waiting for a provider. The initially calculated CUR values were a lagging indicator of the system or, rather, an indication of the state of the system in the period ending at the point in time the summary was captured from the ED patient-specific database. When comparing CUR to other predictors or measures of crowding, such as length of stay, patients leaving before evaluation, waiting room wait, etc., we found that CUR was essentially insensitive to weekly and seasonal changes.

Second, a graphical analysis was used to identify any patterns in the CUR dictated by month, day, and time of day. The results of these graphs were intended to show the CUR's capability in conveying the possibility of overcrowding and the state of the ED in any given time interval. They were also intended to guide the next steps of analysis in generating predictions of the state of the ED.

Third, after identifying patterns in the stratified CURs, the variability by time of day was identified as a robust indicator of the state of the system to eventually suggest when the system should be inspected to determine whether an additional physician should be called to add to the provider capacity. To predict the probability of a CUR greater than 1, a binary logistic regression analysis was performed using a stepwise approach to identify significant predictors to include in the model.

To perform the binary regression analysis, the real-time CUR variables were categorized into binary categories with 0 indicating that the capacity ratio was below 1 and the physicians were able to clear the current ED queue and 1 indicating the capacity ratio was above 1 and the physicians were unable to clear the current queue. This strategy was applied to seven different time intervals: 15 minutes, 30 minutes, 1 hour, 2 hours, 4 hours, 8 hours, and 12 hours to compare how using various intervals affected the predictions.

Time periods used in the initial analysis included: 1) every previous period equivalent to the period being predicted up to 1 day prior, 2) every day beyond 1 day up to 7 days prior, and 3) every week after 7 days up to 4 weeks prior (see Figure 1 for a generic illustration). For example, to predict the probability that CUR would be greater than 1 in the next 8 hours, the following variables were analyzed for significance with respect to their contribution to the overall resulting probability: 8 hours ago, 16 hours ago, 24 hours ago, 2 days ago, 3 days ago, 4 days ago, 5 days ago, 6 days ago, 7 days ago, 2 weeks ago, 3 weeks ago, and 4 weeks ago.

A 95% significance test was conducted on the p-values for the resulting equation. Insignificant variables were eliminated one by one, starting with the highest p-value exceeding 0.05 and iterating through, recalculating the equation with the new set of variables until the equation contained only significant variables (i.e., variables with p-values less than 0.05). This test was performed for each of the time intervals until each prediction equation contained only significant variables.

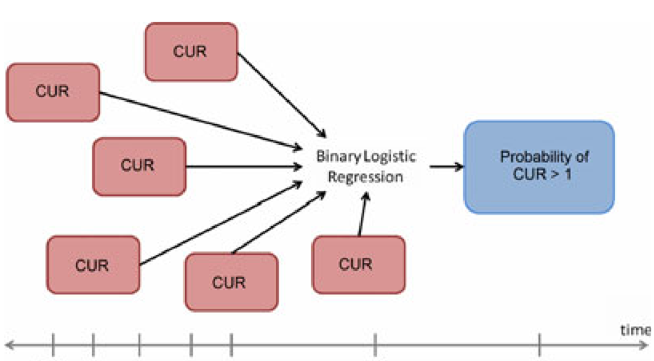

Figure 1. Relationship between CUR values in previous periods to prediction of CUR larger than 1 in the next period. CUR = care utilization ratio.

Finally, the models created were validated against the model validation data set and three separate occasions on which the patient arrivals doubled the normal volume observed in the ED, indicating that the system was overcrowded. This validation allowed us to determine the effect of changing the planning horizon within our models.

RESULTS

We recorded 60,155 ED visits for the model creation data set and 26,383 additional visits for the model validation data set. The two data sets are compared in Table 1.

Additionally, the acuity of patients who arrived during the model creation data set and each of the days where patient arrivals were twice the normal volume is summarized in Table 2. Based on historical data, total physician capacity ranged from 3.2 to 12.3 patients per hour as the number and type (faculty, intern, etc.) of physicians increased and decreased throughout the day. Hourly CURs ranged from 0 (in periods where no new patients arrived) to 3.6.

Graphical Analysis

The CURs were stratified according to month, day, and time of day using Minitab version 16 (Minitab Inc., State College, PA) to identify any patterns. It was found that the trend was relatively consistent with little variation of the average or variance when stratified by month or day, which indicated both that average provider productivity assumptions held and that the behavior of CUR could be expected regardless of the present month or day. However, when the data were stratified by time of day using box plots, a pattern regarding CUR arose (Figure 2).

Table 1
Comparison of Data Used for Model Generation to Data Used for Model Validation

Average Daily Rate	Model Creation Data Set	Model Validation Data Set
Arrivals	157	161
Hospital admissions	59	60
ED discharges	107	107

Table 2
Acuity Distribution of Surge Days Compared to Nonsurge Days During Model Creation

	Acuity Level				
	1	2	3	4	5
Model creation data set	2	42	47	8	1
Surge Day 1	4	41	48	5	2
Surge Day 2	0	43	51	6	0
Surge Day 3	1	38	52	7	1

Values are reported in percentages.

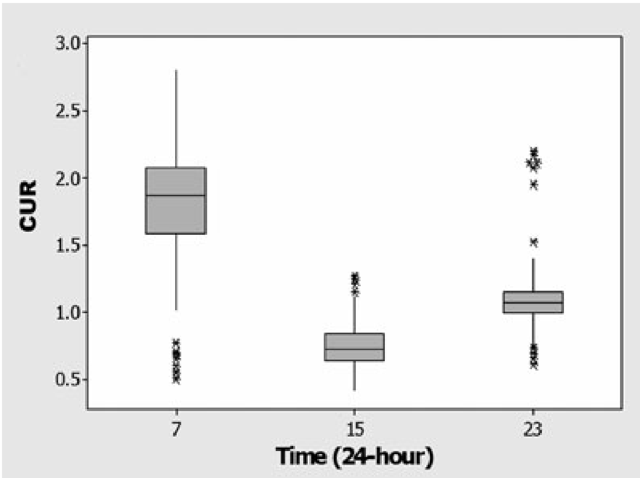

Figure 2. Box plots of CURs by 8-hour interval for the 2009/2010 period. CUR = care utilization ratio.

From 23:00 to 07:00 the following morning, CURs ranged from 1.6 to 2.1 with large variance, preventing providers from treating all patients in a timely manner and adding to the queue of patients waiting. The system is able to regain control over the 07:00 to 15:00 time period, with CURs ranging from 0.6 to 0.8 with a relatively small variance and few occurrences of CUR greater than 1, allowing the providers to eliminate any queue that may have built up throughout the early morning and treat any new patients who may arrive. Finally, the day concludes with CURs ranging from 0.9 to 1.1 and a small variance with a few extreme occurrences of CUR greater than 1.1, sometimes providing the opportunity to either clear the queue if CUR happens to be less than 1, but sometimes building up another queue into the morning hours compromising the state of the system as this cyclic pattern begins again.

This same pattern was observed with all analyzed time intervals. Since we identified time of day as the most significant variation in CUR, and intend CUR to predict the need for additional staff, any seasonal variation in ED use will only result in a seasonal pattern of calling in additional staff. Any emerging pattern in calling in staff would ideally be considered by management when developing the ED staffing schedule, but the seasonal pattern would not affect the usefulness or accuracy of CUR. A binary logistic regression analysis followed to predict the probability the CUR was greater than 1 for each of the time intervals.

Binary Logistic Regression Analysis

The individual patient visits were stratified, again using Minitab version 16, according to the same time intervals (15 minutes, 30 minutes, 1 hour, 2 hours, 4 hours, 8 hours, and 12 hours), and were then analyzed according to a 95% confidence test to identify the resulting significant equations. Table 3 summarizes the resulting equations.

Since the p-values indicate all of the equations are significant, the three best models were identified using deviance values, which are the result of a hypothesis test evaluating the null hypothesis (assuming the model fits the data) compared to the alternative hypothesis (assuming the model does not fit the data). Any deviance value greater than our chosen significance of 0.05 would imply that the null hypothesis would be accepted, indicating that the models are, in fact, significant predictors of the probability the CUR will be greater than 1 in the next period. The resulting deviance values of the 30-minute, 8-hour, and 12-hour models are 1.000, 0.951, and 0.864, respectively. Deviance values of 0.05 or larger indicate that these models would be robust predictors of the probability that CUR would be observed to be greater than 1 in the next period.

Table 3
Significance of Binary Regression Equations

Time Interval*	p-value†	Number of Variables‡
15 minutes	<0.001	19
30 minutes	<0.001	28
1 hour	<0.001	14
2 hours	<0.001	4
4 hours	<0.001	7
8 hours	<0.001	4
12 hours	<0.001	3

CUR = care utilization ratio.
*Prediction horizon for the model.
†Significance of the model with only significant inputs included.
‡Number of independent variables (i.e., CUR values in previous periods) used in prediction horizon.

Model Validation

To determine which of the significant binary logistic regression equations (30 minutes, 8 hours, or 12 hours) would best predict the behavior of the CUR, the equations were validated against the model validation data set and three days on which the ED observed significant surges in patient volume.

First, the positive predictive values of all equations were analyzed to test each model's ability to correctly diagnose positive results (i.e., correctly predict when the ED will experience a CUR greater than 1). Positive predictive values were calculated for a range of percentage thresholds that identify various levels at which the result from a binary regression model would be chosen to predict the state of the ED, depending on the required sensitivity. Table 4 summarizes the resulting comparison of all models based on these predictions.

The thresholds used to evaluate the positive predictive values of the three models illustrate a range of

Table 4
Comparison of Models Using Positive Predictive Values

% Threshold	30 Minutes	8 Hours	12 Hours
20	0.738	0.472	0.226
30	0.783	0.376	0.226
40	0.814	0.191	0.111
50	0.840	0.089	0
60	0.872	0.016	0

probabilities used to gauge the result of the binary regression equation. The threshold chosen to gauge the system is based on a user-required sensitivity to reliably predict the state of the ED.

Table 4 shows that the 30-minute model performs consistently better than the 8- or 12-hour models when comparing positive predictive values for the three models. The input variables for the 30-minute model are summarized in Table 5.

To further analyze how well the 30-minute model predicts the probability the CUR will be greater than 1 in the next period, Type I and Type II errors were calculated and compared across the previously identified thresholds (20% to 60%). Table 6 summarizes the results of this analysis.

Table 5
Significant Input Variables for 30-Minute Model

Input Variable*	Coefficient	p-value
Constant	−1.7814	<0.001
30 minutes	0.2156	<0.001
1 hour	0.1326	<0.001
3 hours	0.1033	<0.001
4 hours	0.0580	0.038
7 hours	−0.0973	0.001
8 hours	−0.0669	0.026
9 hours	−0.1153	<0.001
10 hours	−0.1147	0.001
11 hours	−0.0905	0.008
12 hours	−0.1023	0.003
13 hours	−0.0833	0.014
14 hours	−0.0968	0.003
15 hours	−0.1028	0.001
16 hours	−0.1491	<0.001
19 hours	0.1240	<0.001
20 hours	0.0671	0.020
21 hours	0.0664	0.024
22 hours	0.0655	0.032
23 hours	0.1423	<0.001
1 day	0.2396	<0.001
2 days	0.1472	<0.001
3 days	0.1378	<0.001
4 days	0.1489	<0.001
5 days	0.1484	<0.001
1 week	0.1135	0.001
2 weeks	0.1725	<0.001
3 weeks	0.1096	0.001
1 month	0.0841	0.007

CUR = care utilization ratio.
*Value of CUR at the indicated previous time interval.

Table 6
Type I and Type II Errors for the 30-Minute Prediction Model

% Threshold	True Positive	False Positive	True Negative	False Negative
20	94	30	70	6
30	90	22	78	10
40	85	17	83	15
50	79	13	87	21
60	74	10	90	26

Values are reported as percentages.

The resulting false-positive rates shown in Table 6 (30% to 10% for thresholds 20% to 60%, respectively) identify, using the 30-minute model as a stand-alone predictor, the worst-case scenario for false-positive rates. Combining the 30-minute model with other potential predictor variables, or combining this model with patient-level variables (such as patient acuity or waiting patients), may subsequently decrease these false-positive rates.

Since this model provided significant predictions and low Type I and Type II errors at the given thresholds, the 30-minute model was tested against 3 days where the ED observed significant surges in patient volume. The model was compared to each of these days to determine the time at which the surge could have been predicted for each day to allow the ED staff to prepare appropriately and avoid delays in patient care.

Figures 3 through 5 illustrate the observed CUR for every 30-minute period in the day prior to and during the observed surge (i.e., first and second 0–24 hours, respectively). They also show the predicted probability in each of these periods that CUR would exceed 1, indicating that physician capacity would be exceeded by the demand of new patients to be treated. Figure 3 illustrates the first day a surge in patient volume was observed. Figures 4 and 5 illustrate the second and third surge days in chronological order following the first.

Using this prediction model on each of the three days that experienced a surge in patient volume, we were able to determine with the various thresholds when the ED could have potentially predicted the surge and aligned its resources to accommodate the forecasted increase in patient volume. The times that would have triggered this preparation for the various thresholds during each of the days experiencing a surge are summarized in Table 7.

DISCUSSION

Variable Indicators

In an attempt to directly address the issue of treatment capacity,[16] a ratio of new patients requiring treatment over the physician capacity at that point in time was initialized as a robust predictor of the state of the ED. This CUR allowed for an analysis of several models constructed from various time intervals to determine which time interval best predicted the state of the system. In determining the most robust model, several conditions were considered: the significance of each model in a comparison among all of the models, as well as an individual comparison of significance at various thresholds dictating the probability acceptable to make a decision based on the binary regression output; the time at which the models would have predicted a surge during 3 days when twice the normal patient volume was observed; and finally the time to prepare for a surge allowed by various predictions.

Model Performance

Our results show that there are three potentially significant models to predict the probability that CUR will be larger than 1 in the next period, indicating that patient

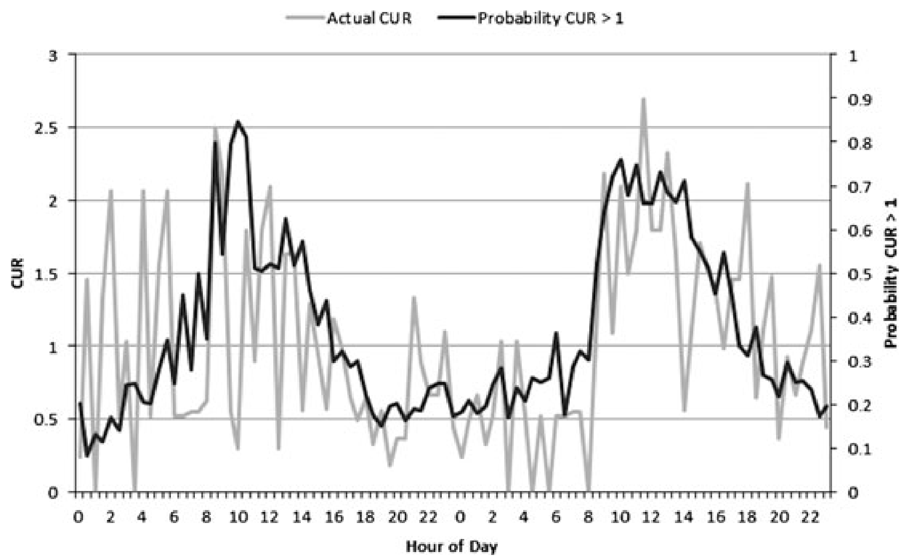

Figure 3. Comparison of predicted probability of CUR larger than 1 and observed (actual data) of the ED during the first day experiencing volume surge. CUR = care utilization ratio.

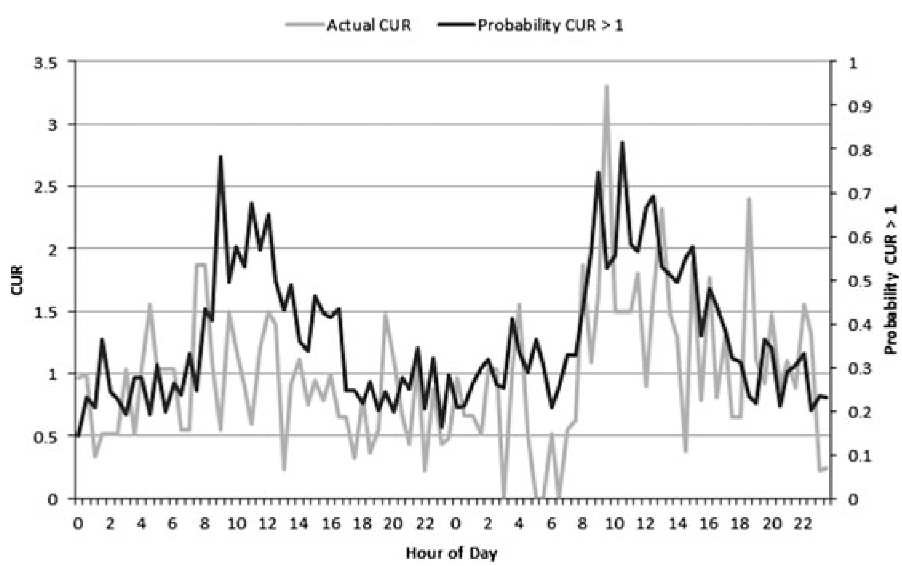

Figure 4. Comparison of predicted probability of CUR larger than 1 and observed (actual data) of the ED during the second day experiencing volume surge. CUR = care utilization ratio.

demand will exceed the scheduled physician capacity. Of these three models, based on the comparison of positive predictive values and Type I and Type II errors, the 30-minute prediction provides the most robust prediction of all the models when used as the sole method to predict the behavior of the system in the next period.

Tables 4 and 6 in conjunction outline the capability of the 30-minute prediction model; this model's deviance value of 1.000 suggests that it has the best goodness-of-fit to the data when compared against the other two strongest models (8 and 12 hours). When these three models are compared using positive predictive values at varying thresholds, the 30-minute model allows for a higher correct prediction percentage at all thresholds compared to the other two models. When the model is categorized by threshold and analyzed based on Type I and Type II errors, the model is capable of correctly predicting when the CUR will be greater than 1 or less than 1 in the next period 94% to 74% and 70% to 90% of the time, respectively, in thresholds ranging from 20% to 60%.

This model, if it had been used on three separate occasions when the ED observed a non–crisis-related surge in patients, would have allowed for a 30-minute warning early in the day, providing the department with an opportunity to prepare for the surge. This 30-minute warning was deemed acceptable under the condition that current policy requires that a physician who is on call at any given time is expected to be within 20 minutes of the hospital throughout his or her on-call shift to allow for prompt arrival and substitution of the departing physician.

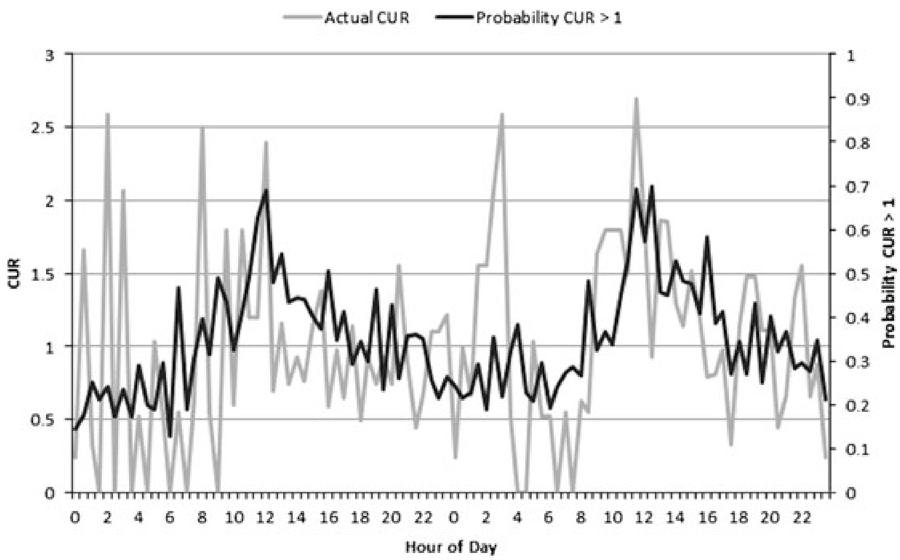

Figure 5. Comparison of predicted probability of CUR larger than 1 and observed (actual data) of the ED during the third day experiencing volume surge. CUR = care utilization ratio.

Table 7
Time Surge in Patient Volume Detected During Each of the Days Experiencing an Abnormal Increase In Patient Volume

% Threshold	First Surge Day	Second Surge Day	Third Surge Day
20	00:30	00:00	00:00
30	06:00	01:30	02:00
40	08:30	03:00	08:00
50	08:30	08:00	10:30
60	09:00	08:30	11:00

The 8- and 12-hour prediction models would provide an initial assessment of the expected CUR values, but do not hold enough power as standalone predictors when compared to the 30-minute model using positive predictor values and Type I and Type II errors. In addition to using the 8- and 12-hour prediction models to develop an initial assessment of the future state of the ED with respect to CUR, our preliminary analysis of other factors, such as wait times, total length of stay, patients in the waiting room, etc., shows potential for further validation and support of our model. Preliminary analysis of the length of stay of the waiting room, length of stay of patients waiting for a physician, and total length of stay identified moderate increases in average wait times on the surge days compared to the nonsurge days. Therefore, analyzing the relationships between the CUR and other factors provides a natural extension of this work for future research.

LIMITATIONS

We only developed our models independently as standalone predictors, but it would be interesting to look at the possibility of using multiple models in combination to obtain a range of predictions to allow for a longer planning period in the event of a predicted surge. In addition, since physician capacity and patient visits are only two aspects of ED crowding, it might be useful to incorporate the prediction value into a larger model incorporating other variables describing the state of the ED (patient wait times, total number of patients in the ED, acuity, number of patients boarding, bed capacity, length of stay, number of patients in the waiting room, etc.). However, although these factors could prove to strengthen our model in future analyses, we found the 30-minute model to be a significant stand-alone predictor of surges in patient volume with the variables included.

Also, our models only incorporate data from one hospital, and therefore it would be useful to assess the capability of the models using data from other hospitals, although our models use generic data sources that could be reasonably collected from other ED databases. Our calculations of an aggregate treatment rate also provide an opportunity for deviation among differently structured EDs, and the values are long-run historical averages, when in practice, individual productivity is highly variable. This aggregate measure assumes, since schedulers for each role do not coordinate their efforts to balance ''slow'' and ''fast'' providers, in the long run the ED productivity would be stable due to the randomness of assignment on any given day. Aggregate productivity assumes that a balanced ED functions at a certain productivity rate. This assumption is necessary in this type of a model because considering individual provider capacities would be complex and therefore impractical in this situation. Therefore, potential for under- or overestimation a provider's unique capacity at any given point in time still exists, but the aggregate measure provides a good starting point for this type of analysis. The evaluation of a single site also provides a unique mix of ED faculty, residents, and physician assistants, which is unlikely to be identical at other institutions, but the models we generated show the need to measure capacity well.

Finally, patient dynamics are not accounted for in this model. A very sick patient occupies much more time and consumes any individual provider's capacity for new patients until that patient is stabilized and then also requires a large amount of follow-up work to maintain stability and make disposition. In reality, therefore, the capacity of the ED is variable, whereas we used a static measure in our models.

CONCLUSIONS

The care utilization ratio is a new and robust indicator of the system's performance and can aid in further development of quantitative models to predict ED overcrowding situations. The study was able to model the tradeoff of longer response time versus shorter, but more accurate predictions, by investigating different prediction intervals. Current practice would have been improved by using the proposed models and would have identified the surge in patient volume earlier on noncrisis days. ED overcrowding warrants continued attention from a qualitative and quantitative perspective to provide hospitals with implementable tools and knowledge to prepare for and reduce the likelihood of these types of situations to provide safe and timely care for all patients.

The authors acknowledge the support of Jeffrey Desmond, MD, Yedan Yin, MSE, Jon Lemus, the University of Michigan Center for Healthcare Engineering and Patient Safety, and the editors for their helpful feedback and suggestions.

References

1. 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.
2. Bernstein SL, Asplin BR. Emergency department crowding: old problem, new solutions. Emerg Med Clin North Am. 2006; 24:821–37.
3. Sklar DP, Crandall CS, Zola T, Cunningham R. Emergency physician perceptions of patient safety risks. Ann Emerg Med. 2010; 55:336–40.
4. Derlet R, Richards J, Kravitz R. Frequent overcrowding in U.S. emergency departments. Acad Emerg Med. 2001; 8:151–5.
5. Pines JM, Garson C, Baxt WG, Rhodes KV, Shofer FS, Hollander JE. ED crowding is associated with variable perceptions of care compromise. Acad Emerg Med. 2007; 14:1176–81.
6. Cowan RM, Trzeciak S. Clinical review: emergency department overcrowding and the potential impact on the critically ill. Crit Care. 2005; 9:291–5.
7. Hwang U, Richardson L, Livote E, Harris B, Spencer N, Morrison RS. Emergency department crowding and decreased quality of pain care. Acad Emerg Med. 2008; 15:1248–55.
8. Pines JM, Pollack CV, Diercks DB, Chang AM, Shofer FS, Hollander JE. The association between emergency department crowding and adverse cardiovascular outcomes in patients with chest pain. Acad Emerg Med. 2009; 16:617–25.
9. Thompson S, Nunez M, Garfinkel R, Dean MD. Efficient short-term allocation and reallocation of patients to floors of a hospital during demand surges. Operations Res. 2009; 57:261–73.
10. Schneider S, Zwemer F, Doniger A, Dick R, Czapranski T, Davis E. Rochester, New York: a decade of emergency department overcrowding. Acad Emerg Med. 2001; 8:1044–50.
11. Derlet RW. Overcrowding in emergency departments: increased demand and decreased capacity. Ann Emerg Med. 2002; 39:430–2.
12. Hoot NR, Epstein SK, Allen TL, et al. Forecasting emergency department crowding: an external, multicenter evaluation. Ann Emerg Med. 2009; 54:514–22.
13. Hoot NR, LeBlanc LJ, Jones I, et al. Forecasting emergency department crowding: a discreet event simulation. Ann Emerg Med. 2007; 52:116–25.
14. Schweigler LM, Desmond JS, McCarthy mL, Bukowski KJ, Ionides EL, Younger JG. Forecasting models of emergency department crowding. Acad Emerg Med. 2009; 16:301–8.
15. Hoot NR, Zhou C, Jones I, Aronsky D. Measuring and forecasting emergency department crowding in real time. Ann Emerg Med. 2007; 49:747–55.
16. Howard PK. Overcrowding: not just an emergency department issue. J Emerg Nurs. 2005; 31:227–8.