Predicting Left-Without-Being-Seen in an Emergency Department as a Dynamic Risk

Yaniv Ravid[a], Rouba Ibrahim[b], Junqi Hu[c], Kal Pasupathy[d], David M Nestler[e], Vahid Sarhangian[f], Philipp Afèche[a]

[a] : University of Toronto, Rotman School of Management
[b] : University College London, School of Management
[c] : Old Dominion University, Strome College of Business
[d] : University of Illinois Chicago, Biomedical & Health Information Sciences
[e] : Mayo Clinic, Rochester, MN
[f] : University of Toronto, Department of Mechanical & Industrial Engineering

Study objective: Accurately predicting which Emergency Department (ED) patients are at high risk of leaving without being seen (LWBS) could enable targeted interventions aimed at reducing LWBS rates. Machine Learning (ML) models that dynamically update these risk predictions as patients experience more time waiting were developed and validated, in order to improve the prediction accuracy and correctly identify more patients who LWBS.

Methods: The study was deemed quality improvement by the institutional review board, and collected all patient visits to the ED of a large academic medical campus over 24 months. The first 18 months of data were used to develop two types of classification models using XGboost: (1) a static model that uses patient and ED census information at the time of arrival to predict the risk to LWBS; and (2) a dynamic model that updates the predictions based on new information after 30 minutes for patients who are still waiting in the ED. The final six months of data were then used as an experimental period to test the accuracy of the two models.

Results: For a cohort of 150,959 ED patient arrivals, the mean patient age was 46, 51% of arrivals were female, and 2.17% of patients LWBS during their wait. The models achieved an area under the receiver operating characteristic curve (AUROC) of 0.86 when dynamically updating LWBS risk levels over time. This was in contrast to an AUROC of 0.80 for the static model from past literature that does not update the score after the moment of arrival. Over the experimental period, the dynamic model also showed the ability to reduce the number of missed LWBS cases by approximately 50% as compared to the static model, without incurring any additional false-positives.

Conclusion: Dynamically updating patients' risk to LWBS as their wait goes on significantly reduces the number of patients who LWBS that are missed by prediction models as compared to traditional static prediction approaches.

Introduction

Background

Patients who arrive at the emergency department (ED) and leave without being seen (LWBS) could face negative medical outcomes as a result [1]. In many instances, they may possibly return to the ED with a worsened condition. High rates of patients who LWBS indicate a mismatch between demand for emergency care and the hospital's ability to attend and evaluate patients in a timely fashion. Therefore, the ability to identify patients at high risk to LWBS is critical in improving patients' medical outcomes and efficiently serving all arriving patients.

Machine learning models have been developed in the literature to predict whether patients will LWBS during their wait [2]. These models have shown promise as data-driven methods to identify patients at risk to LWBS. However, these models only utilize information that is available upon patients' arrival and therefore predict a static LWBS risk level. This information, such as the patient's acuity level and the number of patients in the waiting room, is computed using the hospital's electronic medical records (EMR). Past research has shown that EMR data overestimates wait times for patients who LWBS [1]. As a result, the static risk predictions may not be actionable.

Another drawback of restricting prediction models to static information is the inability to update patients' risk to LWBS as their wait time increases. This is despite the fact that patients' waiting time was shown to be a factor which strongly impacts the decision to LWBS [1]. Accurate measurements of the time patients stay in the waiting room were computed in such studies by fitting patients with radio frequency identification (RFID) bracelets [3,4]. These bracelets were given to patients on arrival, and in-ceiling readers receive signals from the wristbands and constantly record patients' location. This novel use of RFID data uncovered the gaps between patients' true experience in the waiting room and the information available in EMR data. Past research has also shown strong associations between changes to the state of the waiting area (e.g., observing arrivals and departures) and patients' risk to LWBS [5]. However, the models rely on ex-post changes to the environment when studying patient behavior. Nonetheless, these findings suggest that information available upon patient arrival may be insufficient to accurately quantify their likelihood to LWBS. Therefore, environmental changes that can only be identified dynamically during a patient's wait, such as observing other patients getting seen by a doctor, can improve the ability to identify patients who may LWBS.

Importance

We proposed that dynamically updating the predicted risk of each patient's tendency to LWBS results in better risk stratification that can be acted on in real time. This improved ability can help design interventions and practices that improve service. Introducing wait-time as a feature in the dynamic risk calculation can also help create guidelines and targets for how quickly patients should be evaluated. For example, the order in which patients are seen may be adjusted based on accurate predictions of their imminent tendency to LWBS. Such priority decisions are typically based on whether patients have been medically screened and how long they have been waiting. The addition of a dynamic LWBS risk dimension would prioritize patients that are likely to abandon, reduce inefficiencies and avoid costly outcomes.

Goals of this Investigation

This study aimed to improve upon traditional static prediction models that classify patients based on their risk to LWBS. The study uses RFID data to give more accurate data of when patients exit the waiting room. The availability of such precise measurements allowed us to track patients' experience in the waiting room and dynamically update patients' predicted risk to LWBS. The dynamic models then predicted patients' risk to LWBS at specific times during their wait. An experimental period is then used to evaluate whether these dynamic prediction models can be more accurate than traditional static models. In addition to AUROC, another criteria used to compare predictive performance is the number of true-positives and false-positives generated by each model.

Materials and Methods

Setting and EMR Data Collection

A retrospective cohort study was conducted among all patients who visited the ED at a large academic medical center. Patient characteristics used in the analysis were age, sex and Emergency Severity Index (ESI) score. These were documented for each patient in an EMR system, which also included the time in which the patient entered the ED. This *entering time* was recorded into the EMR system at the patient's moment of triage. All patients wait in the same area and can observe all other waiting patients. The time at which each patient was assigned an ED treatment area was also included in the EMR data. At the time of treatment room assignment, a nurse calls for the patient and if present, the patient then exits the waiting room into the treatment area. This is defined as the patient's *exit time* (from the waiting room). However, those patients who LWBS are missing when called on by the nurse. The nurse then records in the EMR system that the patient has LWBS.

RFID Exit Times and Congestion Computation

RFID bracelets were used to record the exit time for patients who LWBS. RFID bracelets were placed upon entering the ED. When patients left the waiting room while wearing the RFID bracelet, the precise exit time was captured in the RFID data. Equipped with these RFID timestamps in addition to the EMR timestamps, the number of patients present in the waiting room at any point in time t can be computed. This measurement is referred to as the waiting room *congestion* at time t. This congestion equals the number of patients whose entering time precedes t, and whose exit time (whether through bed assignment or LWBS) is after t.

Study Cohort

The studied dataset included 150,959 patient-level entries of all ED visits over a two-year period. A total of 3,270 (2.17%) patients left without being seen. Of these patients, exit time was available through RFID data for 3,022 (92%) patients. There is no discernible pattern that describes the group of patients who LWBS with missing RFID data. It is assumed that these patients never had an RFID bracelet placed on their arrival. Such patients whose LWBS time was missing were removed from the analysis. Additionally, 2 records with missing sex data were removed. All 2,126 patients of ESI score 1 were also removed for the purpose of analysis. This is because over the two years of data collection, no patient with an ESI score of 1 was observed to LWBS, as expected, since these are usually patients with life-threatening conditions.

Lastly, pediatric patients (≤19 years) were served by nurses, doctors, and beds that were designated for pediatric care. Therefore, pediatric patients from the dataset were also removed for the purpose of analysis. The final cohort consisted of 119,326 patient visits to the ED.

Panel Data Creation

To study how the risk of LWBS evolved with patients' wait time, each patient visit was split into intervals, and patients' tendency to LWBS in each interval was studied. The number of intervals chosen impacts predictive accuracy. A large number of intervals would result in inaccurate predictions. Therefore, patient visits are split into two intervals. The first consists of the initial 30 minutes of the patient's waiting time. The second interval begins at the 30-minute mark of waiting and ends at the patient's exit time. The choice of 30 minutes for the length of the first interval balances the number of patients who LWBS in the first and second interval. If the first interval were too short, then a very small number of patients who LWBS will be observed in it. If the first interval were too long, then a very small number of patients who LWBS will be observed during the second waiting intervals.

Characterizing each patient by these two intervals, the dataset is transformed into a panel dataset, where each observation is a patient-interval pair. For patients whose exit time is prior to the 30-minute mark of waiting, only one interval exists in the panel dataset. For those who wait longer than 30 minutes, two intervals exist in the panel dataset. If a large number of intervals were used for patient waits, the panel dataset would grow in size, while the number of observed LWBS cases in the dataset would not change. Therefore, the task of predicting LWBS behavior becomes more difficult.

Variable Definition

Patient features used for prediction were age, sex and ESI level. The entering timestamp of each patient was used to determine the day of the week and the hour of day in which they entered the waiting room. Congestion upon arrival was also used as a feature for prediction. Note that these features are static and do not change as a patient's wait continues.

Each patient-interval pair is also characterized by the *sojourn* time the patient has been waiting (0 and 30 minutes for the first and second interval, respectively). Patient-interval pairs representing the second waiting interval are also characterized by the congestion at the onset of the second interval. Additionally, one of the study's goals is to understand how events that patients saw in the past dynamically affect their future decisions. Therefore, the number of events that patients saw during their first waiting interval is also used to describe their second waiting interval. Specifically, these event counts include the number of arrivals, departures into service, and abandonments. Table 1 lists all of the independent variables used in the analysis.

The dependent variable associated with each patient-interval pair is a binary indicator which took value 1 if the patient decided to LWBS during the interval. If the patient was otherwise served during the interval, the indicator variable took value 0. Alternatively, for the first interval of waiting, the indicator variable took value 0 if the patient continued to wait into a second waiting interval.

Model Setup and Tuning

The two years during which data were collected were split into a model *estimation period*, consisting of the first 18 months of data, and an *experimentation period* consisting of the final six months of data. Three predictive models described below were trained and tested using only the first 18 months of collected data. Patients who arrived during this 18-month estimation period were randomly divided into training (80%) and testing (20%) datasets. This testing set was used to validate the models and avoid overfitting. This was done by early-stopping the training process when performance on the testing set began to worsen. The hyperparameter selection process, as described below, was also guided by performance on the testing population. During these procedures, the final six months of data were not used in any way. The final six months of data were instead used for the *experimentation* period during which the estimated models were applied. During this period, the models' ability to identify who among patients that were held out from the estimation process chose to LWBS was measured.

Model 1: First, a static benchmark model that resembles past LWBS prediction models [2] is fit. The training and testing panel datasets were therefore reduced to only represent each patient encounter once. The model only used information that was available upon the patient's arrival at the waiting room. Any dynamic features such as sojourn time and congestion after 30 minutes were therefore not used for prediction. The exact list of static features is shown in Table 1. The binary outcome these static features were used to predict took value 1 if the patient LWBS at any point during their wait, and 0 otherwise.

Model 2: Next, the full panel dataset is utilized and the task of predicting whether a patient will LWBS during each specific waiting interval is considered. The same static features used in Model 1 are relied on, with the addition of sojourn time. Sojourn time equaled 0 for the first intervals in the panel dataset, and 30 minutes for the second intervals. The binary outcome was now interval-specific and equaled 1 for patient-interval pairs whenever the patient chose to LWBS during the interval. The binary outcome equaled 0 otherwise. Unlike the static model, this model was able to dynamically update patients' risk to LWBS as their wait went on.

Model 3: The final statistical model, in addition to the stipulations in Model 2, further allowed for time-varying congestion effects and introduced more dynamic features. For example, this model included the number of arrivals that a patient saw during their first 30 minutes of waiting to predict their risk to LWBS during the second interval. Table 1 indicates the exact variables used in the model. The binary outcome variable was identical to that used in Model 2.

For all three models, the XGBoost classifier [6] was used to predict the risk to LWBS using the corresponding features. This ensemble model combines multiple non-linear classification algorithms, and as a result captures complex relationships between features and the decision to LWBS. Past literature has also shown that XGBoost specifically outperforms other learning algorithms in various ED contexts [2, 7]. The XGBoost Python package was used in the analysis. The objective used in the estimation period was the logistic probability of observing the LWBS decisions in the dataset. A high-performing model would therefore predict high probabilities for observations in the data that chose to LWBS. The metric for evaluating predictive performance and preventing overfitting was the AUROC.

For Model 1, the binary classification was made once for every patient and represented learning each patient's likelihood to LWBS at any point during their wait. For Models 2 and 3, the binary classification was made for each patient-interval pair, and represented learning the dynamic likelihood that the patient will LWBS during the interval.

The training population was used to estimate each model, and the testing population was used to ensure this estimation did not result in overfitting. A tree-specific booster was used for all models. The model hyperparameters were tuned for each model separately in order to maximize the testing performance. These hyperparameters included the maximum tree depth, a regularization parameter, the step size shrinkage between boosting steps, and the subsampling ratio of training instances. For each model, the set of hyperparameters that gave it the best AUROC on the testing dataset was chosen.

Table 1. Definitions of the features used, and their usage in the three predictive models.

Feature	Type	Description	Model 1	Model 2	Model 3
Age	Binary	Focal patient's age	X	X	X
Sex	Numerical	Focal patient's sex, 1 (0) for male (female)	X	X	X
Day of week	Categorical	Six dummy variables for day of the week in which the patient arrived (Friday as reference class)	X	X	X
Hour of day	Categorical	Five dummy variables for the four-hour window in which the patient arrived (0am to 4am, etc., with noon to 4pm as the reference class)	X	X	X
ESI	Categorical	Three dummy variables for the ESI level of the patient (ESI 2 as reference class, ESI 1 not included in analysis)	X	X	X
Congestion at arrival	Numerical	Number of non-pediatric patients who are waiting at the onset of the focal patient's wait (i.e., initial congestion fixed effect)	X	X	X
Sojourn	Numerical	Number of hours the focal patient has been waiting for at the beginning of the waiting interval		X	X
Congestion_30*	Numerical	Number of non-pediatric patients who are waiting at the onset of the focal patient's 2nd waiting interval.			X
Departures_30*	Numerical	Number of non-pediatric departures into service observed by the focal patients in the 30 minutes prior to the onset of the 2nd waiting interval.			X
Arrivals_30*	Numerical	Number of non-pediatric arrivals observed by the focal patients in the 30 minutes prior to the onset of the 2nd waiting interval			X
Abands_30*	Numerical	Number of non-pediatric LWBS observed by the focal patients in the 30 minutes prior to the onset of the 2nd waiting interval			X

*These patient-flow variables are only defined for the second waiting intervals in the panel dataset, and equal 0 for the first intervals.

Experimental Validation

The final six months of data representing the experimentation period were then used to apply the models and measure their performance. The number of patients who LWBS that were caught and missed by each model in this period was measured. The number of false-positives that were raised by each model was also calculated. A false-positive indicates a patient who is predicted to LWBS but in actuality did not LWBS. Specifically, two levels of prediction thresholds that result in different false-positive rates were considered. By choosing a high (low) threshold for prediction, the models predicted fewer (more) patients as likely to LWBS. If patient risk scores vary from 0 to 1, a high (low) threshold for prediction could mean only classifying patients with a risk score above 0.85 (0.25) as patients who will LWBS. As a result, a lower (higher) rate of false-positives will be achieved. The number of patients who LWBS that were missed by each model when only 10% of all panel-dataset observations were allowed to be falsely predicted to LWBS was measured. Then, these measures were computed when this rate is increased to 50%. By fixing the model's false-positive rates to these values, the benefit of added features in each model can be seen through the reduction in the occurrences of missed LWBS cases.

Additionally, the Shapley values (SHAP) of important features [8] were reported to indicate which features strongly impact the prediction output of each model. These values measure the contribution that each feature makes toward the predicted risk to LWBS. When the Shapley value of a specific patient feature is very positive (negative), the feature then strongly increases (decreases) the predicted risk to LWBS. The SHAP package in Python was used for this purpose, with a prediction explainer specifically built for tree-based models.

Results

Cohort Characteristics

The final study cohort consisted of 119,326 patients. Table 2 describes this final cohort used in the analysis, as well as the estimation period cohort used to develop the three predictive models, and the experimentation period cohort used to measure the model's performance.

Over the entire study cohort, 24.1% of the 2,345 patients who LWBS did so within 30 minutes of arrival. Of those who LWBS in the first 30 minutes, 29.5% were detected as LWBS by a nurse within 30 minutes of entering the waiting room. Therefore, a panel dataset built without RFID data would incorrectly include a second waiting interval for 399 patients.

On average, patients who LWBS are first noticed to be gone after 3.08 hours. However, accurate exit time measurements using RFID data show that the average true wait time for these patients is 2.29 hours. Therefore, the average inflation in observed waiting time for LWBS patients is 47 minutes.

Predictive Performance

The three models were used to predict the risk to LWBS for every patient who visited the ED in the six-month experimental period. Using the panel dataset defined by these visits, all three models were used to predict the risk to LWBS for every patient-interval pair in the dataset. Figure 1 presents the receiver operating characteristic (ROC) and precision-recall (PR) curves of the three predictive models on this experimentation group. Model 1 achieved an AUROC of 0.803, while Models 2 and 3 achieved an AUROC of 0.863 and 0.864, respectively. The PR curves show that at the same precision level, the dynamic predictors can achieve a recall of 60% where the static predictions would only correctly recall 10% of all LWBS cases.

The three models' precisions on the experimental cohort are illustrated in Table 3. Model 1 gave the highest number of missed LWBS cases both at 10% and 50% false-positive rate. Model 2 gave the lowest number of missed LWBS cases under the same false-positive rate. At the 50% false-positive rate specifically, Model 1 missed more than twice as many LWBS cases as Model 2 (57 missed LWBS cases compared to 25, over a period of six months).

Table 2. Cohort statistics.

	Entire Cohort	Model Estimation Cohort	Model Experimentation Cohort
Unique Visits (% LWBS)	119,326 (2.4%)	88,533 (2.4%)	30,793 (2.6%)
ESI 2	28,248 (0.7%)	21,530 (0.8%)	6,718 (0.6%)
ESI 3	71,714 (2.3%)	53,458 (2.3%)	18,256 (2.6%)
ESI 4	18,609 (5.2%)	12,988 (5.3%)	5,621 (5.0%)
ESI 5	755 (7.2%)	557 (7.0%)	198 (7.6%)
Age	54.89 (20.05)	54.83 (20.05)	55.07 (20.03)
% Male	48%	48%	48%
Exit time (hours)	0.70 (1.10)	0.67 (1.08)	0.76 (1.16)
ESI 2	0.27 (0.58)	0.27 (0.59)	0.27 (0.55)
ESI 3	0.76 (1.14)	0.74 (1.12)	0.82 (1.21)
ESI 4	1.09 (1.32)	1.06 (1.32)	1.16 (1.33)
ESI 5	0.91 (1.28)	0.96 (1.36)	0.79 (1.00)
Congestion upon arrival	6.18 (6.96)	5.91 (6.78)	6.94 (7.39)
Hours to LWBS	1.58 (1.15)	1.56 (1.14)	1.63 (1.16)
Offered wait time for patients who LWBS (hours)	2.50 (1.52)	2.50 (1.54)	2.47 (1.45)

Data is reported as n (percent), mean (SD), or percent.

Feature Importance

Lastly, Figure 2 illustrates the Shapley values of the features in the dynamic Models 2 and 3 among the experimentation cohort. Congestion observed upon arrival has the largest average absolute impact on the predicted risk in both models. Specifically, high congestion generally increases the predicted risk to LWBS. In Model 2, sojourn time is the second most impactful feature, with patients who waited for 30 minutes having greater predicted risk to LWBS. However, sojourn time is only the fourth most impactful feature in Model 3. Instead, the congestion level at the onset of the second interval has the second largest impact. According to this model, patients who saw greater congestion after waiting for 30 minutes had a larger predicted risk to LWBS.

Figure 1. ROC and PR curves of the three models.

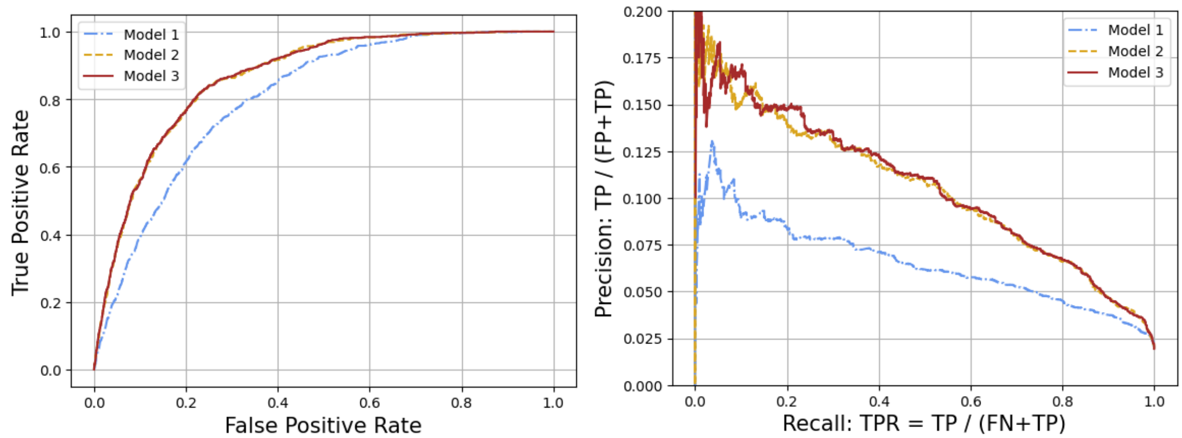

Table 3. Comparing the three models during the six-month experimentation period. The middle (rightmost) column corresponds to a high (low) prediction threshold which results in lower (higher) number of false-positives and true-positives. The total number of patients who LWBS in the six-month period was 810.

	LWBS missed at 10% false-positive rate	LWBS missed at 50% false-positive rate
Model 1	488	57
Model 2	352	25
Model 3	355	23

Note: since many observations shared predicted risk values, exact false-positive rates cannot always be achieved. For each model, the lowest number of missed cases that achieves the desired rate is therefore reported.

Discussion

Patients' decision to LWBS while waiting in the ED presents a major operational hurdle to many hospitals. Recent research has suggested the capability to identify patients with high tendency to LWBS. However, this study shows that more cases can be identified and fewer cases can be missed by also predicting when during the patients' wait is the tendency to LWBS the highest. Instead of treating LWBS as a single binary outcome for every patient, it was treated as a dynamic risk level that changed with patients' experience in the waiting room. The results showed that a static binary prediction performed poorly and missed more LWBS cases than a dynamic risk level can otherwise identify.

The ROC curves show an improvement in AUROC from the static score of Model 1 to the dynamic score of Model 2. This dynamic score was only based on the amount of time patients have waited. The PR curves also illustrate this result, as the dynamic model relying only on sojourn time can achieve significantly better precision and recall than the static model. During the experimental cohort representing six months of ED arrivals, the dynamic predictor always found more LWBS cases than its static counterpart. This result highlighted the importance of dynamically updating LWBS risk predictions over time.

Figure 2. Shapley values of the features in Models 2 (top) and Model 3 (bottom), ranked by average absolute impact on the final predicted value. Each row illustrates the distribution of feature values among the experimentation cohort, with blue (red) sample points illustrating a low (high) feature value. Vertical stacking of points in each row indicates Shapley value density.

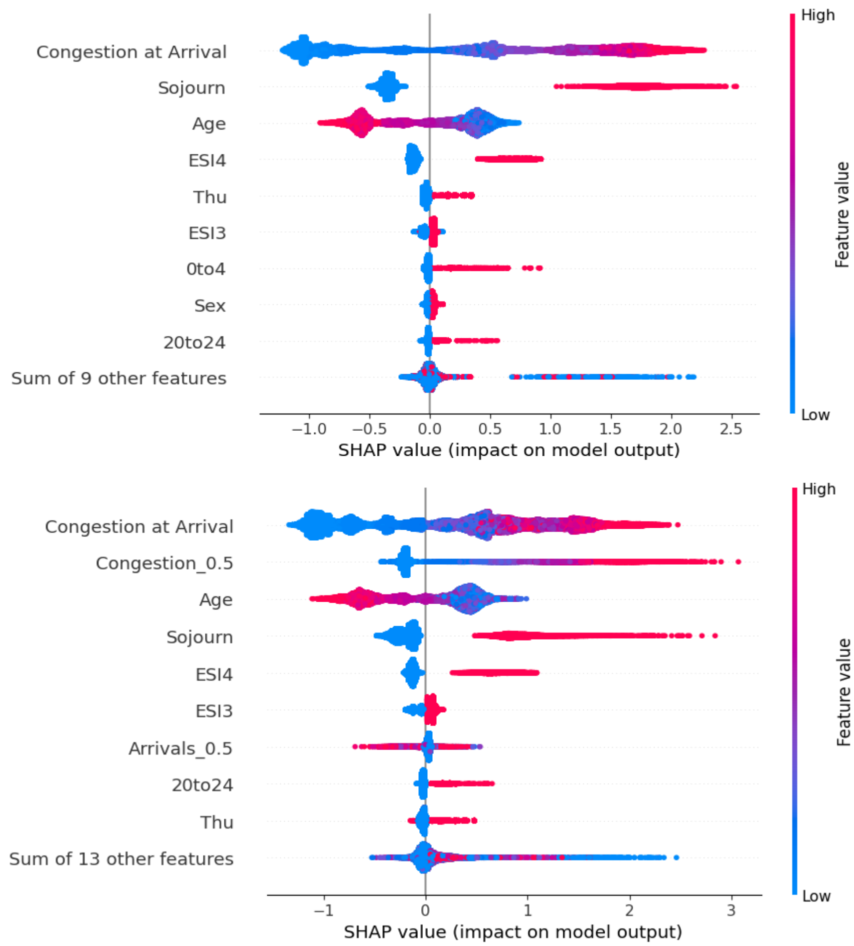

An extension to the dynamic model which accounted for events that occurred during patients' waits in addition to the time they have spent waiting was studied in Model 3. These events were used when computing patients' risk to LWBS. However, this extension did not significantly outperform the basic dynamic predictor, despite the additional features. Including more granular information about what patients observe during their wait does not improve our ability to predict their tendency to LWBS. A simple model that only tracks how long the patient has been waiting can perform just as well without the additional dynamic features, despite their sometimes significant impact to predicted risk. Among these features, the Shapley value of the congestion present at the 30 minute mark significantly impacts patients' risk to LWBS in their second waiting interval. This could be due to the fact that the risk to LWBS is only predicted twice, and so both sojourn time and congestion at 30 minutes take value zero for patients in their first waiting interval. As a result, both features indicate that the patient is in their second waiting interval, which could be the underlying source to LWBS risk. Across all models, higher initial congestion generally increased the predicted risk while greater age reduced the risk. In the dynamic models, increased sojourn time also increased the predicted risk to LWBS, yet the degree to which the predicted risk grew depended on patient characteristics.

The XGboost model has been used to predict static LWBS risks in the past [2], but rarely do these studies report model precision [9]. These precision measurements indicate the level of incorrect noise in the form of false-positive predictions that a model outputs. While identifying LWBS cases is a critical objective, highly noisy predictions would make it difficult to separate the true LWBS cases from those that are falsely predicted to LWBS. To shed light onto the precision at which LWBS cases can be identified, the analysis illustrates the precision of both the traditional static binary predictor and the dynamic risk models. The results showed that under similar false-positive rates, the dynamic predictors always outperform the static ones. When the models are allowed to falsely predict as many as half of all observations in the experimental period as likely to LWBS, Models 2 and 3 correctly identified 96.9% and 97.2% of all patients who LWBS, respectively. In effect, the dynamic prediction models can be used to separate the population of waiting patients into two equal groups, those with high risk to LWBS and those of low LWBS risk. Based on the experimentation cohort, nearly all of patients who LWBS will be correctly placed in the high-risk group. Across a period of six months, with more than 100 daily arrivals, this equates to only 23 missed LWBS cases, less than one per week.

The ability to find patients who will abandon during their wait, as well as the intervals during their wait in which they will likely LWBS, also allows ED staff to intervene and respond to the risk. Possible interventions may take the form of adjusted scheduling decisions, and prioritizing high LWBS risk patients. The design of such interventions could be the focus of future research. Another limitation of the study is the lack of patient-level clinical variables (e.g., chief complaint, medical history, lab and imaging results) that may further improve the ability to identify LWBS behavior. Lastly, predictive models cannot reliably evaluate counterfactual scenarios. If, for instance, the medical staff was increased to reduce congestion, it cannot be claimed that LWBS rates will decline. Therefore, designing interventions based on predicted risk to LWBS should be done carefully, accounting for such counterfactual scenarios.

References

[1] Geers J, Pasupathy K, Lovik K, Finley J, Hellmich T, Marisamy G, Nestler D, Sadosty A, Sir M, Heaton H (2020) Characterization of emergency department abandonment using a real-time location system. The American Journal of Emergency Medicine. 38 (4) 759-762.

[2] Casey P, Zolfaghar K, Eckert C, Waters L, Sonntag H, McKelvey T, Mark NM (2018) Predicting Patients at Risk for Leaving Without Being Seen Using Machine Learning. Annals of Emergency Medicine. 72 (4) S5-S6.

[3] Koenig KR, Pasupathy KS, Hellmich TR, Hawthorne HJ, Karalius VP, Sir M, Das D, Heaton HA, Nestler DM (2021) Measuring Sensitivity and Precision of Real-Time Location Systems (RTLS): Definition, Protocol and Demonstration for Clinical Relevance. Journal of Medical Systems. 45 (1) 15.

[4] Griffin JM, Hellmich TR, Pasupathy KS, Funni SA, Pagel SM, Srinivasan SS, Heaton HA Sir MY, Nestler DM, Blocker RC, Hawthorne HJ, Koenig KR, Herbst KM, Hallbeck MS (2020) Attitudes and Behavior of Health Care Workers Before, During, and After Implementation of Real-Time Location System Technology. Mayo Clinic Proceedings: Innovations, Quality & Outcomes. 4 (1) 90-98.

[5] Batt RJ, Terwiesch C (2015) Waiting Patiently: An Empirical Study of Queue Abandonment in an Emergency Department. Management Science. 61(1) 39-59.

[6] Chen T, Guestrin, C. (2016). Xgboost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. (785-794).

[7] Morey J, Winters R, Jones D (2025) Artificial Intelligence to Predict Billing Code Levels of Emergency Department Encounters. Annals of Emergency Medicine. 85 (1) 63 - 73.

[8] Lundberg SM, Lee, SI. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems, 30.

[9] Hunter D, Carr B, Morey J, Jones D (2024) 242 Machine Learning Model to Predict Emergency Department Patients Who Left Without Being Seen. Annals of Emergency Medicine. 84 (4). S112 - S113.