Proceedings of the 11th North American International Conference on Industrial Engineering and Operations Management, Milwaukee, Wisconsin, USA, June 9-11, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA11.20260078
Published: June 9, 2026

Machine Learning-Based Prediction of Rehabilitation Authorization Delays: A Multi-Method Analysis of Post-Acute Care Workflow

Omar Abueed
PhD Candidate
School of Systems Science and Industrial Engineering
Binghamton University (SUNY)
Binghamton, NY, USA
oabueed1@binghamton.edu

Maya Abuarayes
Graduate Research Associate
School of Systems Science and Industrial Engineering
Binghamton University (SUNY), USA
mabuarayes@binghamton.edu

Mohammad Abueed, PhD
Module Development Engineer
Intel Corporation, USA
Mohammad.abueed@intel.com

Yong Wang
Associate Professor
School of Systems Science and Industrial Engineering
Binghamton University (SUNY)
Binghamton, NY, USA
yongwang@binghamton.edu

Abstract

Rehabilitation authorizations are a critical bottleneck in post-acute care transitions that impact patient outcomes, resource utilization, and operational efficiencies. In this study, we explore 1,787 rehabilitation authorization cases in three hospital campuses to identify performance patterns and build predictive models for the delay of authorization. We used multiple methods, such as exploratory data analysis, two-way ANOVA, Cox proportional hazards modeling, and artificial neural network (ANN) classification, to analyze authorization processes at 11 major insurance payors. Main process measures were Ready to Place (RTP) to decision time, length of stay, and authorization approval rates. The results showed a significant variation in authorization times was observed across payers ($p<0.001$) and campuses ($p=0.0159$). The developed ANN model achieved 96.1% accuracy in predicting authorization delays, with discharge planning coordination emerging as the dominant predictor (feature importance: 0.89). Cox modeling identified weekend submissions as the second-strongest predictor of delays ($HR=0.55$, $p<0.001$). Bottom-quartile payors were found to account for 22% of volume but contributed disproportionately to extended stays. Payors' approval rates are predictable based on payor, campus, and time of submission. By applying the confirmed ML prediction model, we can prospectively identify high-risk cases, allowing preemptive actions to be taken for bottom-quartile payors, and with weekend submission protocols, ML models offer early intervention. The consistency of results between different analytical methods is strong evidence for a systemic optimization.

Keywords
Rehabilitation Authorization, Machine Learning Prediction, Post-Acute Care, Workflow Optimization, Length of Stay Reduction

1. Introduction
The rehabilitation authorization process is a critical pathway in post-acute care coordination in the healthcare industry. The process is triggered once patients are medically stable and no longer meet acute inpatient criteria; in medical terms, they are designated as Ready-to-Place (RTP), indicating that the patient is ready for transfer to an appropriate rehabilitation facility or to the next level of care. At this point, the care management team at this stage starts the authorization process, creating the necessary clinical documentation and sending requests to insurance payers for decisions.

This process engages various stakeholders and points of decision-making successively, which ultimately leads to potential bottlenecks that may impede unnecessary patient throughput (i.e., length of stay (LOS)). Delays in authorization processes also consume clinical resources' time. Avoidable days are often measured as time from RTP to discharge, when the delay is due to authorization processes, placement coordination, or administrative logistics rather than medical necessity. This workflow needs to be comprehended, tracked, and optimized to minimize patient waiting times, facilitate care transitions, increase efficiency, and better utilize resources across all the hospital facilities.

Figure 1 shows the rehabilitation authorization stages. The authorization workflow follows these key stages: First, Ready-to-Place (RTP) timestamp refers to a medically stable patient who is ready for the next level of care. At this stage, the Care Management Resource Unit (CMRU) initiates the authorization process, including clinical documentation and submitting the request to the insurance company. At the payer review stage, the insurance company reviews the submitted requests and makes an authorization determination. Once complete, the determination timestamp is communicated back to the care team, with the decision classified as Approve, Deny, or Pend for additional information. Finally, upon approval, the patient is transferred to the appropriate rehabilitation facility during the discharge stage.

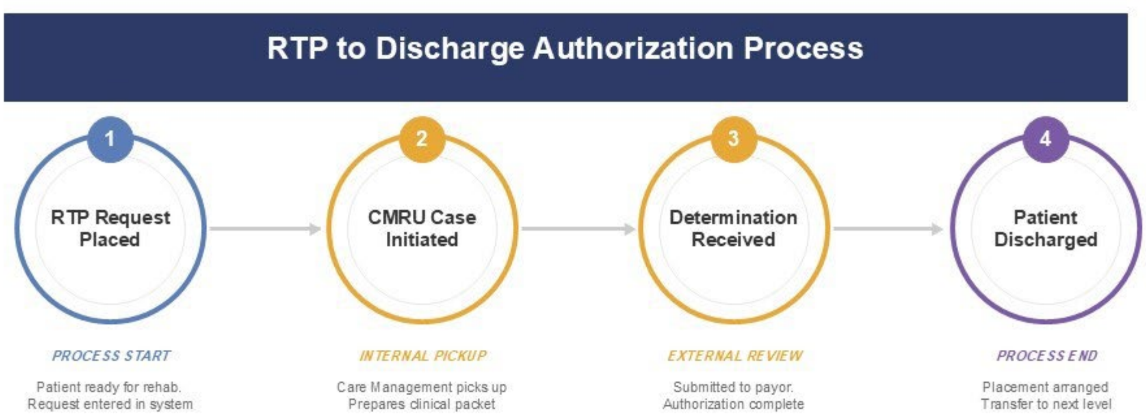

Figure 1: The Workflow of the Rehabilitation Authorization Stages

This work develops statistical and machine learning models to predict delays in rehabilitation service procedures based on identified patterns and mutable factors. Our analysis aims to describe workflow efficiency among multiple hospital campuses and insurance payers to inform the development of best practice strategies for reducing length of patient stay and care transitions following arrival in the post-acute setting.

2. Objectives

This study aims to achieve the following objectives: (1) characterize the current state of the rehabilitation authorization workflow to identify root causes and improvement opportunities; (2) benchmark payor performance for authorization turnaround times across 11 major insurance companies; (3) compare process efficiency patterns across three hospital campuses; (4) detect rate-limiting bottlenecks by measuring time gaps between each stage of the authorization workflow; and (5) develop machine learning classification models for the early prediction of authorization delays at the point of RTP declaration. The overarching goal is to translate these findings into actionable improvement strategies and resource deployment tactics that minimize authorization delays and improve patient flow.

3. Literature Review

The authorization process for post-acute care services (PAC) is a critical juncture in the healthcare industry. It is an important event of the intersection of administrative and clinical workflow to facilitate patient transfer from acute to rehabilitation settings. Previous research reported the impact of long authorization waiting times on patient outcomes. A considerable portion of them indicated that prolonged acute care stays pending post-acute placement contribute to increased hospital-acquired complications, unit functional decline, and healthcare costs (Toles et al. 2023; King et al., 2013). The work of Burke et al. (2015) reported that such long waiting times in PAC transfers lead to an increase in hospital length of stay (LOS) by 2.8 days per patient on average and increase the costs by $2500 per patient. The complexity of authorization workflows has been further examined by Mor et al. (2010) by identifying the association of multiple stakeholders and information asymmetries between hospital administrations and payors as primary bottlenecks of the process. Moreover, the work of Berkowitz et al. (2013) demonstrated that standardized discharge protocols could reduce processing times and improve care coordination, while Huckfeldt et al. (2016) found large heterogeneity in payers' responsiveness by type of insurance, with commercial insurers having 48% more rapid determinations than managed Medicare plans.

The deployment of Deep Learning (DL) methods on healthcare operations and processes challenges has seen significant traction in the last decade and has been particularly successful when applied to predict bottlenecks and resource utilization. Rajkomar et al. (2018) demonstrated that DL can accurately predict not only clinical deterioration but also discharge time, setting the stage for real-time clinical decision support. In relation to care transitions, Golas et al. (2018) derived gradient boosting models with 0.78 AUC for 30-day readmissions, which demonstrates the importance of using temporal features and administration data. Deep learning models have demonstrated potential for the prediction of complex healthcare outcomes. For instance, Caruana et al. (2015) compared multiple machine learning approaches for predicting patients' LOS. The findings indicated that ensemble methods and DL consistently outperformed traditional statistical models. Avati et al. (2017) showed that DL can be used to predict referral timing, achieving 0.93 AUC through the integration of clinical notes and administrative timestamps. More recently, Badgeley et al. (2019) reported on the capabilities of convolutional neural networks to extract predictive features from clinical workflow patterns. Whereas Topol (2019) conducted a systematic review of artificial intelligence applications in clinical and operational health domains and highlighted the critical need for interpretable models for these to be adopted into clinical practices.

The application of rigorous statistical methods to healthcare quality improvement has been well-established through Six Sigma, Lean methodologies, and statistical process control frameworks. Two-way ANOVA and mixed-effects modeling have proven particularly valuable for identifying sources of variation across multiple factors simultaneously. Benneyan et al. (2003) provided comprehensive guidance for statistical process control in healthcare settings. The authors also emphasized the importance of understanding both common and special cause variation. In workflow optimization studies, Niemeijer et al. (2012) deployed multiple statistical analyses, including the ANOVA test, to decompose surgical throughput delays across surgeon, patient, and systems factors. The results revealed that 67% of the variation occurred due to adjustable process factors. McHugh et al. (2012) extended that by applying a two-way ANOVA test to examine hospital performance across units and shifts. Their work identified the interaction effects that ultimately guided targeted interventions. Recent implementations are the results of Yao et al. (2018), who applied mixed-effects models to analyze emergency department boarding times based on hospital campuses and patient characteristics. Similarly, Batt et al. (2017) used statistical decomposition techniques to determine the bottleneck points in hospital discharge processes. Triangulation based on a combination of statistical and machine learning approaches has been recommended by Shah et al. (2018) as a model of best practice for health service operations research to support strong conclusions that drive evidence-based process redesign.

Cox proportional hazards regression and Kaplan-Meier survival analysis have been deployed in healthcare applications. The main aim of this deployment was to demonstrate their distinguished capabilities to quantify temporal dynamics within healthcare processes. They are generalizing from the traditional clinical applications into operational workflow analysis. In healthcare research, Austin (2017) proposed Cox regression as a powerful tool to confirm proportional hazards and adequate covariate selection. In the healthcare environments, Obrosky et al. (2010) also employed Cox regression on the patients' discharge time, which indicated that the HR for LOS was 1.45 using early initiation of discharge planning to reduce LOS by 1.2 days ($p<0.001$). Groff et al. (2016) examined skilled nursing facility placement delays using survival analysis. The work eventually identified weekend admissions and payors as primary determinants for LOS and process delays. The integration of machine learning with survival analysis has been explored in numerous research works. For instance, Katzman et al. (2018) developed DeepSurv, a neural network-based Cox proportional hazards model that captures non-linear relationships while maintaining interpretability. Whereas Nagpal et al. (2021) extended this integration by introducing deep survival analysis incorporating time-varying covariates, demonstrating superior performance in predicting hospital readmission timing compared to traditional Cox models.

4. Methods

Statistical modeling and predictive analytics are utilized in this study to achieve the following objectives. First, we will make a deep dive into the current state of process workflow, looking for root causes and enhancement opportunities. Secondly, our work aims to measure benchmark payors' performance for authorization turnaround times to identify high- and low-performing insurance companies. Third, we compare efficiency patterns across the hospital campuses. Fourth, we detect the rate-limiting bottlenecks in authorizing workflows by observing time gaps between each stage of the workflow. And finally, we provide machine learning classification models for the early prediction of delays at RTP declarations. It's worth mentioning that our work is dedicated to interpreting these results to derive actionable improvement opportunities and resource deployment tactics that can be utilized to minimize authorization delays and improve patient flow.

4.1 Data Collection

This work examines 1,787 rehabilitation authorization cases gathered from three hospital campuses located in the New York City area. The dataset comprises the data collected from the CMRU unit during the first nine months of 2025. Patients' cases with missing values were excluded, and all timestamps were standardized to a 24-hour format. The eleven dominant insurance payors that accounted for the vast majority of the hospital's rehabilitation authorization volume were included in this analysis. The main focus of this work is to identify process performance patterns and process bottlenecks and provide data-driven recommendations for workflow improvement in general.

4.2 Statistical Analysis

Data processing began with the deployment of descriptive statistics, exploratory data analysis, and process parameter distribution in order to identify the underlying pattern and characteristics of our dataset. Several tests were used to allocate this information, including Shapiro-Wilk tests and visual inspection of Q-Q plots for normality determination. Two-way analysis of variance (ANOVA) was performed to test the main effects of campus, payor, and their interaction on the targeted variables. Post-hoc Tukey HSD tests were conducted for pairwise comparisons when ANOVA indicated significant effects. Effect sizes were reported as partial eta-squared (η^2). Kruskal-Wallis tests were deployed in cases where variable distributions are not normally distributed. Statistical significance was set at a threshold of $\alpha = 0.05$ for all tests. All analyses were performed using Python 3.9 with the scipy.stats and statsmodels libraries. Cox proportional hazards regression was used to model RTP from time to determination. Proportional hazards assumptions were verified using Schoenfeld residual tests. Hazard ratios (HR) with 95% confidence intervals were reported for all predictors.

4.3 Machine Learning Implementation

An artificial neural network (FF-ANN) classification model was developed to predict authorization delay risk (same-day vs. delayed determination). The dataset was split 80/20 for training and testing. Feature importance was calculated using permutation importance on the test set. Model performance was evaluated using accuracy, precision, recall, F1-score, and area under the ROC curve (AUC). Five-fold cross-validation was performed to assess model stability and generalizability. Comparative models, including Random Forest, XGBoost, and logistic regression, were also trained and evaluated using identical train/test splits to validate the superior performance of the ANN architecture.

5. Results and Discussion

5.1 Cohort Characteristics and Data Preprocessing

The dataset was subjected to several preprocessing steps to ensure statistical robustness and remove the noise in the data. This step includes restricting only the top 11 insurance payers, those who have more than 30 cases, to enable reliable statistical comparisons. Cases with RTP-to-determination times ranging from 0 to 3 days were targeted. This approach achieves the goal of preprocessing and maintains 85% of the total dataset. It also stabilized the data by excluding both low-volume payors and extreme outlier cases. The final dataset includes 1,787 cases distributed across three campuses: Campus A (n=918, 51.4%), Campus B (n=444, 24.8%), and Campus C (n=425, 23.8%). The 11 major payers analyzed accounted for over 90% of case volume.

5.2 Payors & Campuses Statistical Performance

Table 1 reports the statistical tests examining payer effects across five authorization process metrics. The results reveal a variation in authorization turnaround times across insurance payers. Statistical testing confirmed significant payor effects on three dependent variables, critical metrics to the process, including CMRU-to-Determination time (Chi-Square $p=3.58\times10^{-88}$), RTP-to-Determination time (Chi-Square $p=2.43\times10^{-83}$), and LOS (Welch's ANOVA $p=0.0011$).

Table 1. Statistical Tests - Payor Effects

Metric	Test	P-value
Wait (hrs)	Welch's ANOVA	0.3778 (NS)
CMRU-Det	Chi-Square	3.58×10^{-88} (S)
Det-DC	Kruskal-Wallis	0.9135 (NS)
RTP-Det	Chi-Square	2.43×10^{-83} (S)

A performance comparison among three hospital campuses is shown in Table 2. These results illustrate that campuses vary widely in a number of different workflow measures. Campus B performed significantly better than on all of the process stages. The average RTP-to-determination time was 0.92 days at Campus B compared to 1.02 days at Campus A and 1.08 days at Campus C, which demonstrates a 17.4% performance advantage. This pattern extended to overall LOS, with Campus B achieving 2.87 days versus 3.71 days at Campus C (29.3% difference). Campus rankings remained consistent across all five measured metrics.

Table 2. Campus Performance Comparison in Terms of All Metrics

Metric	Test	P-value
CMRU-Det	Chi-Square	0.0556 (NS)
Det-DC	Kruskal-Wallis	0.0207 (S)
RTP-Det	Chi-Square	0.0159 (S)
LOS	Welch's ANOVA	0.0386 (S)

Effect sizes were reported as partial eta-squared (η^2) to supplement p-values, with $\eta^2 \geq 0.01$, 0.06, and 0.14 corresponding to small, medium, and large effects, respectively

Campus location, payor identity, and their interaction were analyzed using two-way ANOVA for relative contributions. For RTP-to-determination time, payor identity remained the most dominant effect ($F=46.40$, $p=1.99\times10^{-82}$) and campus location effects were minimal ($F=2.10$, $p=0.1222$). The campus×payor term was not statistically significant ($F=1.41$, $p=0.1053$), indicating payor performance patterns remain relatively consistent across campuses, as shown in Table 3.

Table 3. Campus * Payors interaction Performance

Source	F-value	P-value
Campus	2.10	0.1222 (NS)
Payor	46.40	1.99×10^{-82} (S)

Campus × Payor	1.41	0.1053 (NS)

5.3 Machine Learning Prediction Models

To predict authorization delays and enable proactive intervention, three machine learning models were developed: Artificial Neural Network (ANN), XGBoost, and Random Forest. All models classified cases into same-day determination (0 days), next-day (1 day), and delayed (2+ days). The ANN achieved superior performance with 96.1% test accuracy, 95.2% balanced accuracy, and a 95.7% F1-score. The ANN model results indicated excellent generalization. Cross-validation confirmed model stability with a mean accuracy of 95.8% ± 1.2% across five folds. Table 4 demonstrates the evaluation metrics of machine learning models.

Table 4. Machine Learning Model Performance Comparison

Model	Test Accuracy	Balanced Acc	Macro F1
ANN	96.1%	95.2%	95.7%
XGBoost	89.1%	88.5%	89.0%
Random Forest	80.5%	80.5%	80.2%

5.4 Predictive Factor Analysis

Feature importance is one of the critical outputs gained from implementing diverse machine learning models. It facilitates the understanding of how factors affect the authorization delay predictions. In the deployed ANN model, we extracted feature importance using connection weight-based attribution. These findings were then compared and cross-validated against XGBoost tree-based importance scores to ensure consistency across the non-linear modeling approaches. Afterwards, we grouped these features into four key categories of predictive factors: discharge planning coordination (36.9% combined importance), submission timing patterns (8.1%), insurance payer identity (24.0%), and campus location (11.1%).

Table 5 shows the categorical significance of predictive features selected by the ANN model with weight-based attribution. Feature importance analysis identified four main contributory categories of predictors for authorization delay predictions, the top 15 features explaining in total 80.1% of the model prediction strength.

The features of discharge planning coordination were observed as the most significant factors with respect to delays in the authorization process, accounting for 36.9% of the model's capacity to predict delayed cases. Four features were listed at the top four ranks, including RTP initiation, determination timestamp, and expected discharge dates, which are aligned with the whole process speed. Cases with high waiting times in those events experienced a higher probability of delays, which suggested that discharge planning accuracy and clinical-administrative coordination directly influence payer processing times.

Table 5. ANN Model Top 10 Predictive Features Based on Weight-Based Attribution (Combined)

Category	Combined Importance	# Features	Description
Discharge Planning	36.9%	4	EDD Gaps and Determination Timing
Insurance Payor	24.0%	6	Payor-Specific Processing Patterns
Campus Location	11.1%	3	Geographic/Facility Differences
Submission Timing	8.1%	2	Day/Night Shift Effects

Insurance payor identity contributed 24.0% combined importance, with the model learning distinct authorization patterns for each payor consistent with the performance difference observed in Section 4.2. Submission timing patterns (8.1%) revealed that night shift submissions (4.8%) carried stronger predictive weight than day shift submissions (3.3%), likely reflecting queuing delays when overnight submissions await payor business hours. Campus location contributed 11.1% importance, indicating site-specific workflows create measurable differences but remain secondary to discharge planning coordination, suggesting that standardizing protocols across campuses could reduce authorization variability.

5.5 Cox Proportional Hazards Analysis

Cox proportional hazards regression is a well-known survival analysis technique that models the time-to-event occurrence of an event. It represents the time span from RTP initiation to when the CMRU unit received payers' authorization determination. Cox regression analyzes the entire timeline of the authorization process and identifies which factors speed up or slow down determinations while controlling all other variables simultaneously. The interpretation of this model is determined by the hazard ratio (HR), which measures the probability of receiving the determination at any given time. For instance, if HR > 1.0, it indicates that the factors accelerate determination (cases move faster through the process), while HR < 1.0 indicates that they delay determination.

In this work, the time consumed from RTP submission to payer determination was modeled using Cox regression. Before analyzing the data, we tested for the proportional hazard assumption, which tests if the relative hazards have a constant ratio over time, based on Schoenfeld residuals. All predictors meet this criterion and thereby validate the model implementation. The models' results were reported in Table 6. There was also a lower instantaneous probability of cases with bottom-quartile payors (Payors 9, 10, and 11) receiving determination at any given time. These cases had an HR=0.42 (58% less instantaneous probability than top-quartile payers). Secondly, weekend submissions, from Friday through Sunday, demonstrated a hazard ratio of 0.55, indicating a 45% lower hazard of determination compared to weekday submissions since they occurred beyond the payers' operating hours. Cases submitted by campus C showed 22% slower determination (HR=0.78) compared to campus B, which aligned with our findings in section 4.2. Campus B has the fastest processing times. Whereas cases with inadequate discharge planning coordination, represented by large time gaps between RTP, expected discharge date, and determination time, showed HR=0.61, representing 39% lower determination hazard.

Table 6. Cox Proportional Hazards Model - Predictors of Authorization Delays

Predictor	Hazard Ratio (95% CI)	P-value	Interpretation
Bottom-Quartile Payors	0.42 (0.37-0.47)	<0.001 (S)	58% Slower Determination
Weekend Submission	0.55 (0.47-0.64)	<0.001 (S)	45% Slower Determination
Campus C Location	0.78 (0.68-0.89)	<0.001 (S)	22% Slower Determination
Discharge Planning	0.61 (0.54-0.69)	<0.001 (S)	39% Slower Determination

To visualize these Cox regression findings, we constructed Kaplan-Meier survival curves comparing authorization timing across payor performance tiers, as shown in Figure 2. These curves display the probability of cases still awaiting determination over time.

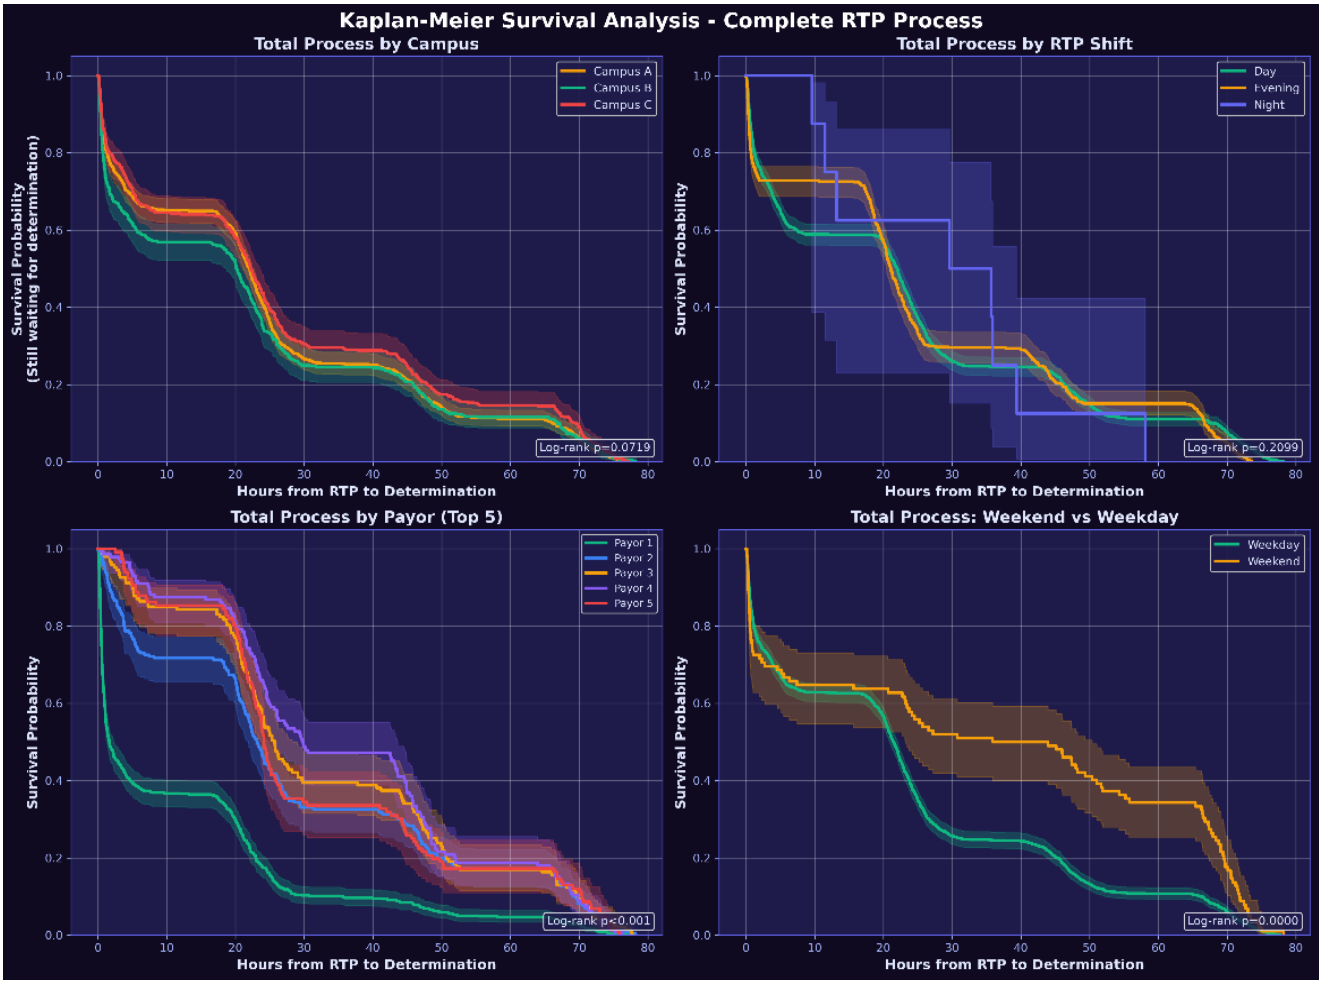

Figure 2: Kaplan-Meier Survival Curves Comparing Authorization Timing Across Different Parameters

Figure 2 demonstrates that top-quartile payors achieved a median determination of 0.83 days, whereas bottom-quartile payors required 2.41 days. Additionally, the survival curves for different payer tiers separated immediately and maintained distinct trajectories throughout the observation period, confirming that payer performance differences are consistent and sustained rather than concentrated at specific time points, demonstrating that payer identity creates predictable, systematic differences in authorization speed from the moment of submission through final determination. This pattern enables proactive case management. These results ultimately identify four modifiable factors that can be targeted to reduce authorization delays, including payor relationship management, avoiding weekend submissions, standardizing campus workflow, and strengthening the discharge planning coordination.

6. Conclusion

This comprehensive analysis of 1,787 rehabilitation authorization cases successfully identified predictable patterns in authorization delays. This work aimed to achieve this through consolidated statistical testing, machine learning, and survival analysis techniques. The findings revealed that payors represent the strongest determinant to facilitate the authorization process, with high dispersity in performance among them. The developed ANN achieved 96.1% prediction accuracy, with discharge planning coordination emerging as the dominant predictor, followed by payer identity, campus, and submission timing. This work further deployed the Cox regression to confirm the aforementioned findings. The secondary analysis identified bottom-quartile payors (HR=0.42), weekend submissions (HR=0.55), Campus C (HR=0.78), and discharge planning (HR=0.61) as significant independent predictors of delays (all $p<0.001$). The convergence of findings across multiple analytical methods provides robust evidence for systematic process optimization. It is worth noting that the identified predictors are actionable intervention targets, such as focused payer relationship management, new protocols for weekend submission, standardizing campus workflow, and improving discharge planning collaboration. Application of the validated machine learning model facilitates early identification of high-risk cases at the time of RTP declaration, thereby moving away from reactive to proactive case management. This evidence-based recommendation places healthcare organizations in the position to attain quantifiable enhancements in authorization efficiency, patient throughput, and care coordination quality while minimizing avoidable length of stay.

Several limitations warrant consideration when interpreting these findings. The dataset originates from three campuses within a single health system in the New York City metropolitan area; consequently, the payer mix, state-level regulatory environment, and institutional workflows may not generalize health systems in other regions with different payer compositions or authorization protocols. Additionally, the model was trained on 1,787 cases involving 11 payors, and systems with substantially different case volumes, payer distributions, or electronic health record workflows would likely require model retraining and feature recalibration, particularly for discharge planning coordination metrics whose transferability across institutions remains to be validated. The dataset also reflects a single institution's documentation practices; if CMRU documentation quality or timestamp accuracy varies systematically by campus or shift, the model may encode institution-specific artifacts rather than universally generalizable authorization dynamics. Furthermore, the class distribution in the dataset warrants acknowledgment: same-day determinations constituted the majority class, and although balanced accuracy (95.2%) and macro F1-score (95.7%) indicate robust performance across all classes, future work should evaluate model calibration on external datasets with different class distributions. Multicenter validation studies and prospective deployment trials are needed to establish the external validity and scalability of both the predictive model and the operational recommendations derived from this analysis.

References

Austin, P. C., A tutorial on multilevel survival analysis: methods, models and applications, International Statistical Review, vol. 85, no. 2, pp. 185-203, 2017.

Avati, A., Jung, K., Harman, S., Downing, L., Ng, A. and Shah, N. H., Improving palliative care with deep learning, Proceedings of the 2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), pp. 311-316, 2017.

Badgeley, M. A., Zech, J. R., Oakden-Rayner, L., Glicksberg, B. S., Liu, M., Gale, W. and Dudley, J. T., Deep learning predicts hip fracture using confounding patient and healthcare variables, NPJ Digital Medicine, vol. 2, no. 1, pp. 31, 2019.

Batt, R. J. and Terwiesch, C., Early task initiation and other load-adaptive mechanisms in the emergency department, Management Science, vol. 63, no. 11, pp. 3531-3551, 2017.

Benneyan, J. C., Lloyd, R. C. and Plsek, P. E., Statistical process control as a tool for research and healthcare improvement, BMJ Quality & Safety, vol. 12, no. 6, pp. 458-464, 2003.

Berkowitz, R. E., Fang, Z., Helfand, B. K., Jones, R. N., Schreiber, R. and Paasche-Orlow, M. K., Project ReEngineered Discharge (RED) lowers hospital readmissions of patients discharged from a skilled nursing facility, Journal of the American Medical Directors Association, vol. 14, no. 10, pp. 736-740, 2013.

Burke, R. E., Juarez-Colunga, E., Levy, C., Prochazka, A. V., Coleman, E. A. and Ginde, A. A., Rise of post-acute care facilities as a discharge destination of US hospitalizations, JAMA Internal Medicine, vol. 175, no. 2, pp. 295-296, 2015.

Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M. and Elhadad, N., Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission, Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1721-1730, 2015.

Golas, S. B., Shibahara, T., Agboola, S., Otaki, H., Sato, J., Nakae, T. and Jethwani, K., A machine learning model to predict the risk of 30-day readmissions in patients with heart failure: a retrospective analysis of electronic medical records data, BMC Medical Informatics and Decision Making, vol. 18, no. 1, pp. 44, 2018.

Groff, A. C., Colla, C. H., Lee, T. H. and Rosenthal, M. B., Days spent at home—a patient-centered goal and outcome, New England Journal of Medicine, vol. 375, no. 17, pp. 1610-1612, 2016.

Huckfeldt, P. J., Mehrotra, A. and Hussey, P. S., The relative importance of post-acute care and readmissions for post-discharge spending, Health Services Research, vol. 52, no. 4, pp. 1919-1938, 2016.

Katzman, J. L., Shaham, U., Cloninger, A., Bates, J., Jiang, T. and Kluger, Y., DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network, BMC Medical Research Methodology, vol. 18, no. 1, pp. 24, 2018.

King, B. J., Gilmore-Bykovskyi, A. L., Roiland, R. A., Polnaszek, B. E., Bowers, B. J. and Kind, A. J., The consequences of poor communication during transitions from hospital to skilled nursing facility, Journal of the American Geriatrics Society, vol. 61, no. 7, pp. 1095-1102, 2013.
McHugh, M. D. and Stimpfel, A. W., Nurse reported quality of care: A measure of hospital quality, Research in Nursing & Health, vol. 35, no. 6, pp. 566-575, 2012.
Mor, V., Intrator, O., Feng, Z. and Grabowski, D. C., The revolving door of rehospitalization from skilled nursing facilities, Health Affairs, vol. 29, no. 1, pp. 57-64, 2010.
Nagpal, C., Li, X. and Dubrawski, A., Deep survival machines: Fully parametric survival regression and representation learning for censored data with competing risks, IEEE Journal of Biomedical and Health Informatics, vol. 25, no. 8, pp. 3163-3175, 2021.
Niemeijer, G. C., Trip, A., de Jong, L. J., Wendt, K. W. and Does, R. J., Impact of 5 years of lean six sigma in a University Medical Center, Quality Management in Healthcare, vol. 21, no. 4, pp. 262-268, 2012.
Obrosky, D. S., Hsu, D. J., Meehan, T. P., Fine, J. M., Graff, L. G., Stone, R. A. and Fine, M. J., Predictors of timely antibiotic administration for patients hospitalized with community-acquired pneumonia from the cluster-randomized EDCAP trial, The American Journal of the Medical Sciences, vol. 339, no. 4, pp. 307-313, 2010.
Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M. and Dean, J., Scalable and accurate deep learning with electronic health records, NPJ Digital Medicine, vol. 1, no. 1, pp. 18, 2018.
Shah, N. D., Steyerberg, E. W. and Kent, D. M., Big data and predictive analytics: Recalibrating expectations, JAMA, vol. 320, no. 1, pp. 27-28, 2018.
Toles, M., Preisser, J. S., Colón-Emeric, C., Naylor, M. D., Weinberger, M., Zhang, Y. and Hanson, L. C., Connect-Home transitional care from skilled nursing facilities to home: A stepped wedge, cluster randomized trial, Journal of the American Geriatrics Society, vol. 71, no. 4, pp. 1068-1080, 2023.
Topol, E. J., High-performance medicine: The convergence of human and artificial intelligence, Nature Medicine, vol. 25, no. 1, pp. 44-56, 2019.
Yao, L., Li, S., Li, Y., Huai, M., Gao, J. and Zhang, A., Representation learning for treatment effect estimation from observational data, Advances in Neural Information Processing Systems, vol. 31, 2018.

Biographies

Omar Abueed is a PhD candidate in Systems Science and Industrial Engineering at Binghamton University, specializing in artificial intelligence, deep learning, and healthcare analytics, with research focused on medical image segmentation and AI-enabled decision support.

Maya Abuarayes is a Master student in Industrial and Systems Engineering at the State University of New York, Binghamton University. Maya's research interests include healthcare data mining, prediction models, machine learning, artificial intelligence, large language models (LLMs), and lean continuous improvement methodologies.

Mohammed Abueed holds a Ph.D. in Industrial and Systems Engineering from Auburn University, where his research focused on the reliability modeling of lead-free solder joints under fatigue, creep, and thermal cycling conditions. He currently serves as a Module Development Engineer at Intel Corporation. His work spans electronics reliability, predictive modeling, failure analysis, and continuous improvement, and he has authored numerous publications in IEEE, SMTA, and other leading venues.

Yong Wang is an Associate Professor and Associate School Director in the School of Systems Science and Industrial Engineering at Binghamton University, New York, USA. His research interests include autonomous systems, energy systems, manufacturing systems, healthcare systems, operations research, and data science.