

Optimizing Hospital Resource Allocation Using Artificial Intelligence and Predictive Analytics

Executive Presentation Summary

Sara Seri
Master of Science in Business Analytics
Grand Canyon University
2026

Executive Overview

This project develops an end-to-end hospital operations decision-support framework. CDC National Healthcare Safety Network weekly data were used to forecast next-week inpatient occupancy. The selected forecasting model was then connected to a linear optimization model that recommends surge-bed activation and supplemental nursing coverage under different demand conditions.

Data Foundation

Measure	Result
Source	CDC/NHSN Weekly Hospital Respiratory Data by Jurisdiction
Original observations	20,770 weekly records
Original variables	322 variables
Modeling panel	17,360 state/territory-week records
Time coverage	August 2020 through July 2026
Primary target	Next-week inpatient occupancy percentage

Forecasting Results

Four forecasting approaches were compared using chronological 70/30, 80/20, and 95/5 train-test partitions. XGBoost achieved the lowest average forecast error and was selected as the leading model.

Model	Mean RMSE	Mean MAE	Mean R ²	Within 3 Points
Persistence Baseline	3.114	1.900	0.891	82.74%
Linear Regression	2.922	1.831	0.905	83.35%
Random Forest	2.899	1.811	0.906	83.81%
XGBoost	2.888	1.805	0.907	83.75%




Figure 1. Forecast RMSE across chronological train-test splits.

On the 80/20 test split, XGBoost produced an RMSE of 2.593 occupancy percentage points, an MAE of 1.655, an R² of 0.924, and 85.28% of forecasts within three occupancy points of the observed value.




Figure 2. Actual and predicted next-week inpatient occupancy.

Key Predictive Drivers

The most influential signals were current and recent inpatient occupancy, four-week occupancy trends, ICU occupancy, the Operational Stress Index, and respiratory admissions. These findings show that short-term hospital demand is persistent but is also affected by changing clinical and respiratory conditions.




Figure 3. Most influential forecasting features.

Resource Optimization Results

The selected forecast was translated into prescriptive recommendations using a transparent linear-programming model. Because the public CDC dataset does not contain individual hospital staffing costs or licensed-bed details, the

optimization uses a standardized 500-bed reference hospital. These assumptions are illustrative and should be replaced with facility-specific inputs before deployment.

Scenario	Forecast Occupancy	Surge Beds	Additional Nurses/Shift	Unmet Demand
Typical Demand	72.09%	0	0	0
Elevated Demand	76.82%	4	1	0
High Demand	80.63%	24	5	0
Critical Demand	83.07%	37	8	0




Figure 4. Forecast demand and optimized bed capacity.

Executive Recommendations

- Run the occupancy forecast weekly and distribute results to operations, nursing, and capacity-management leadership.
- Use forecast thresholds to activate flexible staffing and surge capacity before occupancy reaches critical levels.
- Monitor MAE, RMSE, forecast bias, and predictions within three occupancy points to detect model drift.
- Replace the illustrative optimization assumptions with hospital-specific bed, staffing, cost, and safety constraints.
- Keep clinical and operational leadership in the approval loop; the framework should support, not replace, management judgment.

Important Limitations

The source data are aggregated by jurisdiction rather than by individual hospital. The project therefore demonstrates a scalable analytical framework, not a facility-ready staffing standard. Hospital-level deployment requires local census, unit-level capacity, staffing ratios, labor costs, transfer agreements, and governance requirements.

Included Submission Materials

- Phase 1 report and workbook: data understanding, quality assessment, variables, and feature engineering.
- Phase 2 report and workbook: forecasting methods, model comparison, predictions, and feature importance.
- Phase 3 report and workbook: optimization assumptions, resource recommendations, and sensitivity analysis.
- Python scripts: reproducible data preparation, forecasting, and optimization workflow.
- Original and cleaned datasets, supporting CSV tables, and fourteen publication-ready figures.

Optimizing Hospital Resource Allocation Using Artificial Intelligence and Predictive Analytics

Phase 1: Data Understanding, Acquisition, and Feature Design

Dataset Overview

The analysis uses the CDC Weekly Hospital Respiratory Data (HRD) Metrics by Jurisdiction dataset. The uploaded file contains 20,770 weekly observations and 333 variables spanning 2020-08-08 through 2026-07-11. The file contains 67 geographic series, including states and territories, 10 HHS regions, and a national series.

Purpose and Suitability

The dataset was selected because it contains direct measures of inpatient and ICU occupancy, respiratory admissions, hospitalizations, and reporting coverage. These measures support one-week-ahead forecasting of hospital capacity pressure and can be connected to an optimization model for proactive bed and staffing allocation.

Data Quality Findings

No duplicate geography-week records were identified. Occupancy percentage variables provide substantially better continuity than many raw bed-count fields, so the initial forecasting design uses percentage occupancy measures as the primary operational targets. Historical count variables containing the value 999 were treated as unavailable in the analytical file, and percentage values outside the 0–100 range were excluded from modeling while retained in the audit documentation. Reporting coverage variables will remain available as data-quality controls.

Initial Feature Engineering

- Respiratory Admissions Total: combined COVID-19, influenza, and RSV weekly admissions.
- Respiratory Hospitalized Total: combined currently hospitalized respiratory patients.
- Operational Stress Index: weighted combination of inpatient occupancy, ICU occupancy, and respiratory admissions burden.
- Capacity Pressure: Low (<70%), Moderate (70–<80%), High (80–<90%), or Critical (≥90%) based on inpatient occupancy.
- Lag Features: one-week and two-week prior occupancy and admission values within each jurisdiction.
- Rolling Features: four-week historical averages using prior observations only.
- Forecast Targets: next-week inpatient occupancy percentage and next-week capacity-pressure category.

Modeling Panel Design

The clean modeling panel contains 17,360 state/territory-week observations. National and HHS regional rows are retained for descriptive benchmarking but excluded from the initial modeling panel to reduce

aggregation overlap. Each jurisdiction is ordered chronologically, and all lag and rolling features are calculated within jurisdiction to prevent cross-jurisdiction contamination.

Planned Forecasting Objective

The principal continuous target is next-week inpatient occupancy. A complementary classification target groups next-week capacity pressure into Low, Moderate, High, and Critical categories. The forecasting stage will compare a seasonal baseline, linear regression, Random Forest, XGBoost, and an appropriate time-series model. The optimization stage will translate the selected forecast into recommended beds, surge capacity, and staffing requirements under operational constraints.




Figure 1. United States inpatient and ICU occupancy trends.




Figure 2. Distribution of jurisdiction-week capacity pressure.




Figure 3. United States weekly respiratory admissions.




Figure 4. Missingness among selected analytical variables.




Figure 5. Jurisdictions with the highest median inpatient occupancy.

Phase 1 Conclusion

The dataset is suitable for forecasting and hospital resource-allocation research, but its reporting structure requires careful handling. The initial design prioritizes occupancy percentages, reporting-completeness controls, jurisdiction-specific lags, and transparent capacity-pressure thresholds. This Phase 1 package establishes the audited data foundation needed for exploratory analysis, forecasting, optimization, stress testing, and deployment planning.

Optimizing Hospital Resource Allocation Using Artificial Intelligence and Predictive Analytics

Phase 2: Forecasting Model Development and Assessment

Forecasting Objective

The forecasting phase estimates next-week inpatient occupancy at the jurisdiction level. Chronological train-test partitions were used so that earlier weeks trained the models and later weeks tested them. This design preserves temporal order and avoids using future observations to predict the past.

Models Evaluated

- Persistence Baseline: assumes next week's occupancy will equal the current week's occupancy.
- Linear Regression: provides an interpretable linear benchmark.
- Random Forest: captures nonlinear relationships and interactions.
- XGBoost: sequentially combines weak decision trees and controls overfitting through regularization.

Evaluation Strategy

The models were evaluated with 70/30, 80/20, and 95/5 chronological splits using MAE, RMSE, MAPE, R-squared, and the percentage of predictions within three and five occupancy points of the observed value.

Model Results

Split	Model	MAE	RMSE	MAPE	R2	Within 3 pts
70/30	Persistence Baseline	2.212	3.811	3.37%	0.846	79.2%
70/30	Linear Regression	2.060	3.476	3.14%	0.872	80.3%
70/30	Random Forest	2.045	3.444	3.14%	0.874	81.0%
70/30	XGBoost	2.051	3.461	3.15%	0.873	80.6%
80/20	Persistence Baseline	1.811	2.911	2.72%	0.904	83.6%
80/20	Linear Regression	1.675	2.634	2.52%	0.921	84.9%
80/20	Random Forest	1.662	2.606	2.51%	0.923	85.7%
80/20	XGBoost	1.655	2.593	2.51%	0.924	85.3%
95/5	Persistence	1.677	2.619	2.50%	0.924	85.4%

Baseline						
95/5	Linear Regression	1.759	2.656	2.62%	0.922	84.9%
95/5	Random Forest	1.726	2.648	2.58%	0.922	84.7%
95/5	XGBoost	1.710	2.611	2.56%	0.924	85.4%

Across the three chronological splits, XGBoost produced the lowest average RMSE and was selected as the leading forecasting model. Its average RMSE was 2.888 occupancy points, average MAE was 1.805, and average R-squared was 0.907. The close performance of Random Forest and XGBoost indicates that recent occupancy and lagged operational conditions contain a strong and stable forecast signal.

Interpretation

The strongest predictors were recent inpatient occupancy, rolling occupancy history, ICU conditions, and the operational stress index. This indicates that near-term hospital demand is persistent but also influenced by short-run changes in ICU utilization and respiratory burden. The improvement over the persistence baseline demonstrates that machine learning adds measurable value beyond simply carrying the current occupancy level forward.

Operational Stress Testing

Four operational scenarios were developed to connect forecasts with management decisions. Normal weeks maintain standard coverage, elevated respiratory weeks activate flexible staffing and limited surge capacity, high-strain weeks open surge units and coordinate transfers, and critical scenarios activate emergency operations and regional load balancing. These scenarios will become formal constraints in the optimization phase.




Figure 6. Forecast RMSE across chronological splits.




Figure 7. Actual and predicted next-week occupancy.




Figure 8. Actual versus predicted occupancy.




Figure 9. Top forecasting features.




Figure 10. Forecast residual distribution.

Phase 2 Conclusion

The forecasting analysis supports the hypothesis that historical hospital operating measures can predict next-week inpatient occupancy. XGBoost demonstrated the strongest average out-of-sample performance and will feed the Phase 3 optimization model. The next phase will translate predicted occupancy into recommended bed, staffing, and surge-capacity allocations under safety and cost constraints.

Optimizing Hospital Resource Allocation Using Artificial Intelligence and Predictive Analytics

Phase 3: Bed, Staffing, and Surge-Capacity Optimization

Optimization Objective

Phase 3 translates the selected XGBoost occupancy forecast into operational resource recommendations. The optimization minimizes the combined cost of surge-bed activation, overtime nursing coverage, and unmet demand, while requiring sufficient bed and nursing capacity to cover forecast demand plus a safety buffer.

Scope and Assumptions

Because the CDC dataset reports jurisdiction-level occupancy percentages rather than hospital-specific staffed-bed and labor-cost data, the optimization is presented as a transparent scenario model for a standardized 500-bed reference hospital. The assumptions are illustrative planning inputs and should be replaced with an organization's actual capacity, labor, and cost data before deployment.

- Licensed beds: 500.0
- Regular staffed beds: 400.0
- Maximum surge beds: 100.0
- Regular nurse capacity (beds per shift): 400.0
- Nurse-to-occupied-bed planning ratio: 0.2
- Maximum overtime nurses per shift: 25.0
- Safety buffer percent: 0.05
- Surge bed cost per week: 2500.0
- Overtime nurse cost per week: 4500.0
- Unmet bed-demand penalty per bed-week: 20000.0

Optimization Formulation

Decision variables include the number of surge beds to activate, the number of additional nurses required per shift, and any remaining unmet bed demand. The model minimizes weekly incremental cost. Constraints ensure that bed capacity and nurse-supported capacity meet buffered forecast demand, subject to maximum surge and overtime limits.

Scenario Results

Scenario	Forecast %	Occupied Beds	Buffered Demand	Surge Beds	OT Nurses	Weekly Cost
Typical Demand	72.1%	360.4	378.5	0	0	\$0
Elevated Demand	76.8%	384.1	403.3	4	1	\$11,158
High	80.6%	403.1	423.3	24	5	\$79,168

Demand						
Critical Demand	83.1%	415.3	436.1	37	8	\$122,762

Interpretation

The optimization recommends little or no incremental capacity during typical weeks, then progressively activates surge beds and overtime nursing coverage as predicted occupancy rises. Under high and critical scenarios, proactive allocation eliminates or sharply reduces modeled shortages. The comparison with a reactive no-surge approach shows the operational value of linking forecasts to advance capacity decisions.

Sensitivity and Risk

Sensitivity analysis varied the safety buffer and maximum surge-bed limit. Higher safety buffers improve resilience but increase planned capacity cost. Restrictive surge limits may create unmet demand during critical weeks. This tradeoff demonstrates why hospital leaders should calibrate optimization assumptions to local clinical standards, staffing contracts, transfer agreements, and financial constraints.




Figure 11. Forecast demand and optimized bed capacity.




Figure 12. Recommended surge beds by scenario.




Figure 13. Recommended additional nursing coverage.




Figure 14. Illustrative reactive and optimized planning costs.

Phase 3 Conclusion

The optimization phase demonstrates how next-week occupancy forecasts can be converted into specific bed and staffing recommendations. The model provides a practical decision-support framework rather than a universal staffing standard. Deployment would require replacing the illustrative assumptions with facility-level licensed beds, staffed beds, unit-specific nurse ratios, labor costs, and surge constraints. The next phase will evaluate implementation risk, external evidence, cost-benefit logic, dashboard deployment, and lifecycle monitoring.