

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.