ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

AI-HealthCloud: A Computational Feasibility Study of Influenza Admission Forecasting, Simulated Healthcare Resource Allocation and Public Health Surge Planning

Oluwasanya Ogunsakin[1]; Temitope Iyamu Fadairo[2]; Misheck Malangano[3]; Kevin Mukasa[4]

[1]Bournemouth University, Bournemouth, Dorset, United Kingdom.
[2]DePaul University, Chicago, Illinois, United States of America.
[3]University of West Georgia, Carrollton, Georgia, United States of America.
[4]Maharishi International University, Fairfield, Iowa, United States of America

Publication Date: 2026/11/25

Abstract: Background: Linking forecasts of respiratory-disease-related hospital demand with resource-allocation decisions may support preparedness, but operational feasibility cannot be inferred from surveillance counts alone. Objective: To evaluate the computational feasibility of an auditable forecast-to-allocation workflow, termed AI-HealthCloud, using a limited influenza admissions extract and explicitly hypothetical resource constraints. Methods: Ten U.S. national weekly observations and four weeks of preserved admissions observations for 52 jurisdictions were used to evaluate deterministic one-week forecasting rules. Two jurisdictional target weeks (104 scored jurisdiction–target pairs, but only two time origins) and seven national rolling targets were compared. A synthetic capacity threshold of 115 percent of October 24 admissions and a hypothetical 1,200-unit resource budget supported retrospective policy comparisons, with additional surge and budget sensitivity tests. Results: The half-trend rule yielded 18.38 percent weighted absolute percentage error (WAPE) across the two jurisdictional target weeks versus 19.93 percent for persistence; an exploratory jurisdiction-cluster bootstrap interval for its mean absolute-error advantage included zero. On seven national rolling targets, nonnegative linear extrapolation produced the lowest WAPE (8.29 percent). In one illustrative allocation scenario, forecast-excess weighting left 1,208 modeled unmet units versus 1,572 for equal allocation (23.1 percent fewer); a constrained linear program left 1,401. These are outcomes of assumed capacity arithmetic, not observed hospital shortages. Conclusions: The calculations establish a reproducible, limited-sample feasibility demonstration, not evidence of superior AI prediction, effective hospital resource allocation, improved patient outcomes, or a deployed cloud service. Manual transcription, unresolved aggregation differences, short time coverage, assumed capacities, and absent independent seasonal and operational validation limit interpretation.

Keywords: *Influenza Hospitalization; Computational Feasibility; Public-Health Surveillance; Deterministic Forecasting; Resource-Allocation Simulation; Demand Surge; Decision Support; Reproducibility.*

How to Cite: Oluwasanya Ogunsakin; Temitope Iyamu Fadairo; Misheck Malangano; Kevin Mukasa (2026) AI-HealthCloud: A Computational Feasibility Study of Influenza Admission Forecasting, Simulated Healthcare Resource Allocation and Public Health Surge Planning. *International Journal of Innovative Science and Research Technology*, 11(11), 1521-1532. https://doi.org/10.38124/ijisrt/26nov961

I. INTRODUCTION

➢ *Background and Motivation*

Respiratory disease surveillance can indicate changes in healthcare demand, but anticipation and action are different scientific problems. A forecast of hospital admissions does not automatically establish how many beds, staff members, respiratory devices, or transfer slots are available. Similarly, an optimization algorithm can produce an apparently efficient recommendation while relying on inputs that are too incomplete to justify operational use. AI-HealthCloud addresses this interface as a proposed decision-support architecture rather than treating data ingestion, statistical prediction, or cloud storage alone as an innovation.

In the United States, the CDC National Healthcare Safety Network (NHSN) publishes aggregated weekly Hospital Respiratory Data (HRD), including influenza and

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

other respiratory-virus admissions, hospital capacity and occupancy, and associated reporting coverage [1]. CDC's FluSight program has also established a sophisticated comparison environment for influenza forecasting. Its September 2026 evaluation reports performance across numerous competing models and evaluates prediction intervals, calibration, and forecast skill rather than relying only on model complexity [2]. Consequently, the existence of AI predictions or cloud-hosted reporting should not be presented as new in itself.

A central research challenge is to connect temporally valid forecasts to transparently constrained decisions while distinguishing observed inputs from simulation assumptions. The present report evaluates whether such a computational workflow can be executed and audited with a limited sample; it does not test whether the framework improves decisions under realistic clinical constraints.

➢ *Problem Statement*

A surveillance-only process identifies a possible increase in admissions but does not quantify the consequences of competing preparedness decisions. A decision system that allocates resources according to raw case counts may ignore per-capita burden, existing capacity, reporting changes, and local delivery constraints. More elaborate decision algorithms can perform worse if their objectives or assumptions are poorly aligned with the evaluation outcome. A unified research approach must therefore assess (a) data quality, (b) out-of-sample forecasting reliability, (c) feasibility of resource constraints, and (d) simulated decision performance separately before any integrated claim is made.

This study documents a traceable computational protocol through deterministic prediction benchmarks, a stylized resource-allocation model, separated observed and assumed inputs, and executable analyses. Its purpose is to establish technical feasibility and identify the limits of inference from a restricted public-data extract, not to demonstrate clinical effectiveness or validate a trained artificial-intelligence system.

➢ *Objectives, Questions, and Scope*

The objectives are to (1) specify an integrated yet modular forecasting-to-decision architecture; (2) document the provenance and limits of a preserved public-health data extract; (3) compare transparent one-week statistical forecasting baselines without look-ahead; (4) evaluate hypothetical allocation rules and a constrained linear program; and (5) examine sensitivity to assumed capacities, allocation budgets, and surge magnitudes. The research questions concern execution and auditability, variation of baseline accuracy by aggregation, and sensitivity of simulated outcomes to explicitly defined assumptions.

The analysis concerns reported influenza-associated admissions rather than infections or individual patient outcomes. The working extract comprises ten national weeks ending September 5–November 7, 2026, and four jurisdictional weeks ending October 17–November 7, 2026. No patient-level records, verified staffed-bed counts, hospital transfer data, operational cloud deployment, or clinical intervention evaluations are included. The AI-HealthCloud label denotes a proposed future architecture; no trained AI model is evaluated in the current experiment.

II. RELATED RESEARCH, TECHNICAL CONTEXT AND GAP

➢ *Surveillance Infrastructure and Reporting Definitions*

The CDC NHSN Hospital Respiratory Data program provides geographically aggregated hospitalization observations and capacity snapshots, together with measures of reporting completeness [1]. The Jan. 2026 NHSN HRD protocol distinguishes weekly new admissions, accumulated from Sunday through Saturday, from single-day bed-capacity and occupancy snapshots for Wednesday of the reporting week [3]. These two measurement classes must not be treated as interchangeable. A week with many new admissions does not mean the same number of beds were occupied simultaneously. Historical data also cross different reporting regimes: the period from May through October 2024 was voluntary, with the current mandated reporting period beginning in November 2024 [1].

These distinctions are central to this research. A backtest that ignores changing reporting coverage risks learning changes in data submission rather than changes in underlying demand. A resource model that treats admission flow as the stock of hospital beds risks systematically misstating shortages. An adequate empirical study must measure completeness by jurisdiction and week; check whether the national and component observations reconcile; restrict or stratify reporting regimes; and separate patient arrivals, length of stay, existing occupancy, staffing, and physical bed capacity.

➢ *Epidemiological Forecasting Benchmarks*

The CDC FluSight program demonstrates that forecasting severe respiratory disease is already an active empirical research field [2], [4]. Its 2025–2026 evaluation compares forecasts for short-horizon influenza hospital admissions and places substantial emphasis on weighted interval score and interval coverage. The report also describes reduced reliability of some forecasts during rapid increases or changing seasonal trends. These observations favor robust baseline models, explicit forecast horizons, and probabilistic evaluation. For this proof of concept, only deterministic point forecasts are possible because the accessible sample is too small for credible uncertainty calibration or seasonal model training.

The persistence method holds the most recently observed value constant. Linear trend extrapolation projects the latest difference forward. A half-trend rule dampens that increase. All three are transparent benchmark estimators, not artificial-intelligence models. They are useful precisely because a future machine-learning implementation should demonstrate benefits over simple, reproducible alternatives. Model complexity without comparative evidence is not an appropriate claim of advancement.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

In a three-season retrospective analysis of CDC FluSight hospitalization forecasts, Frutos et al. [5] compared performance-weighted and unweighted probabilistic ensembles; the best-performing strategy varied by season and geographic aggregation. This is an important counterpoint to selecting a model on the basis of a few autumn weeks. The present study does not reproduce FluSight ensemble evaluation and cannot claim comparable forecast skill, but it follows the principle of comparing forecasts with transparent baselines and explicitly stated scoring periods.

Although it examines Medicare financial risk rather than respiratory disease, Mukasa [6] reviews the use of AI for predictive risk adjustment and emphasizes model transparency, fairness, and auditability. Its contribution to AI-HealthCloud is a transferable governance lesson: predictive sophistication should be evaluated alongside interpretability and the consequences of using a forecast in a regulated healthcare setting. The article does not supply evidence of influenza forecast accuracy or hospital capacity performance.

➢ *Resource-Allocation Research Considerations*

Forecasting and allocation are connected but not equivalent. An optimal decision requires an objective function, decision variables, resource constraints, feasibility conditions, and an explicit representation of uncertainty. Hospitals need usable staffed capacity, not merely the number of beds recorded in an aggregate table. Geographic redistribution may face transportation, medical compatibility, professional licensure, jurisdictional authority, and timing restrictions. Because the available extract contains no actual staffing, transfer, or length-of-stay observations, the model here uses abstract preparedness service units. The unit is designed to demonstrate algorithmic consequences and is not intended as a substitute for beds or nurses.

In a genuine deployment, the objective might minimize expected unmet bed-days and transfer burden while penalizing severe inequality of access. Such a model would require forecasts of admissions, patient duration and case mix, baseline occupied capacity, flexible staffing, clinical constraints, and allowable interventions. The current exercise instead evaluates whether alternative policies behave as their mathematical definitions imply when applied to an observed subsequent week under identical hypothetical assumptions.

Healthcare-specific optimization literature provides a much closer precedent. Fattahi et al. [7] examined resource sharing and patient transfers using a multi-stage stochastic formulation for pandemic planning. Jiang et al. [8] expressly combined epidemic prediction with healthcare resource allocation in a rolling predict-then-optimize framework. Johnson et al. [9] examined hospital ward-level bed requirements, while Begen et al. [10] described a Monte Carlo tool used to support actual hospital transfer planning. Together, these works show that predictive demand and constrained allocation have already been integrated in the literature; they also expose the operational information missing from the current exercise: patient flows, bed census, length of stay, transfer eligibility, and uncertainties. AI-HealthCloud is therefore not presented as the first integrated prediction-and-allocation framework.

Onaji et al. [11] investigate multi-agent deep reinforcement learning for dynamic network-resource orchestration under service-quality constraints. Their telecommunications study illustrates an alternative adaptive resource-allocation paradigm, but its simulated network results cannot be transferred directly to hospital beds, clinical staffing, or patient transfers. The present AI-HealthCloud experiment therefore retains explicitly defined allocation policies and a linear program; healthcare-specific reinforcement learning remains an untested extension.

➢ *Governance, Transparency and Reporting*

Prediction systems used in health settings must identify what the model predicts, how predictions are evaluated, whether performance changes by group or location, and who is authorized to act on the output. The TRIPOD+AI statement addresses transparent reporting of clinical prediction-model research, and its general commitments to clear model specification, evaluation, and limitations are useful reference principles here, although this study predicts aggregate population-level admissions rather than individual patient prognosis [12]. CDC's own methodology emphasizes reporting coverage and careful interpretation of respiratory hospitalization measures [13].

A cloud-ready design also needs provenance, dataset versioning, access controls, reproducible deployments, monitoring, and clear human review. Claims of scalability, latency, reliability, data protection, or interoperability must be measured in the implemented environment and cannot be established merely by drawing a systems diagram.

The World Health Organization [14] emphasizes human autonomy, safety, transparency, accountability, and equity in AI for health. Those considerations inform the proposed review and auditing layers but do not establish regulatory compliance. This proof of concept uses aggregated public observations; any future patient-level integration would raise materially different privacy, security, and governance obligations.

Ogunsakin and Mukasa [15] review ethical and regulatory challenges surrounding AI-enabled health cloud systems in the United States, highlighting privacy, patient-data control, stakeholder trust, and accountability. These themes support making governance and human review explicit requirements of AI-HealthCloud rather than assuming that cloud deployment itself ensures responsible use. Because the current prototype uses public aggregate observations, its data-protection exposure differs from that of patient-level clinical systems; future integration of confidential records would require separate legal, privacy, and security assessment.

A related publication record attributed to Manyura [16] is titled “Artificial Intelligence for Secure Health Information Exchange: Advancing Healthcare Operations Through Intelligent Interoperability.” Its ResearchGate listing suggests

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

topical proximity to secure health information exchange, but a publisher-hosted full text could not be independently examined, and its listed July 2026 date is not reconciled with the supplied volume-coded DOI. Accordingly, that bibliographic listing is cited here only to acknowledge related work; no specific methodological design choice, experimental result, or interoperability performance claim in this study is derived from it.

➢ *Research Gap and Precise Contribution*

Prior studies have examined influenza forecasting ensembles and prediction-informed resource optimization [7], [5], [8]. Neither the use of cloud computing nor the general connection of forecasts and allocation is claimed here as a novel invention. The narrower contribution is an auditable, reproducible feasibility demonstration that exposes differences among simple forecasting baselines, makes synthetic allocation assumptions explicit, examines sensitivity to those assumptions, and documents the evidence required for future clinical or operational assessment. Whether the present design constitutes a methodological advance beyond prior systems has not been established.

III. METHODS AND PROPOSED SYSTEM ARCHITECTURE

➢ *Study Design and Evidence Classification*

This computational feasibility study separates four categories of information: (O) preserved CDC-derived surveillance observations; (F) retrospective deterministic point forecasts scored against later observations; (S) model-generated resource-allocation and surge scenarios based on hypothetical capacities and budgets; and (P) proposed, unimplemented cloud components. Observed epidemiological quantities, predicted values, and simulated resource outcomes are not treated as interchangeable evidence.

The complete original NHSN data extract was not locally acquired, so this study cannot establish a full historical data-quality profile or perform a multi-season predictive evaluation. The local experiment uses a preserved, manually assembled FluSight-derived working subset. A separate unexecuted script is provided for future bulk source acquisition and auditing. No claim of training machine-learning models, auditing the full source archive, or benchmarking a deployed cloud application is made.

➢ *Data Provenance, Geography and Cleaning*

The source is the CDC FluSight Forecast Hub's public target-hospital-admissions.csv, a curated target file derived from hospital surveillance. The preserved sample includes national values for ten consecutive weeks ending September 5 through November 7, 2026. For the jurisdictions, four consecutive Saturdays ending October 17, October 24, October 31, and November 7 provide 52 units: the fifty states, District of Columbia, and Puerto Rico. The data therefore contain 208 jurisdiction-week observations and ten national weekly observations, with the national series evaluated separately rather than pooled with component geographies.

The November 7 jurisdictional observations were preserved during the initial review. Observations for October 17, October 24, and October 31 were manually transcribed into a local CSV from the public target series. Local checks found no duplicate jurisdiction–week keys, negative values, or missing admissions values within the assembled four-week panel; these checks cannot establish that every transcribed value matches the original source file. The complete upstream source was not frozen locally, and reported totals remain subject to revision and source-level reconciliation.

For the four evaluated weeks, the locally assembled sum across 52 jurisdictions differs from the separately preserved U.S. aggregate by 8, 7, 1, and 2 admissions, respectively (Table 1). These differences have not been reconciled against a complete source download. They could arise from transcription, aggregation conventions, or source revisions; no cause can be inferred from the available files. National and jurisdictional evaluations are therefore analyzed separately, and results must be treated as exploratory pending source verification.

Akinyelure [17] synthesizes how community-generated behavioural-health insights and participatory information can complement conventional surveillance for policy design. This informed a prospective provenance requirement for AI-HealthCloud: a future ingestion layer should distinguish official epidemiological counts from qualitative community inputs, record source reliability and consent conditions, and forbid unvalidated contextual variables from silently changing quantitative forecasts. The present experiment uses only the preserved aggregate influenza admission series; no behavioural-health, social-platform, or community-participation data were collected or analyzed. Consequently, no performance benefit from community-data integration is asserted.

Table 1 National Row Versus Component Jurisdictions (Unresolved Differences)

Week Ending	US Row	52-Juris. Sum	Difference
2026-10-17	5090	5082	8
2026-10-24	6451	6444	7
2026-10-31	8583	8582	1
2026-11-07	9395	9393	2

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

➢ *Deterministic Forecast Rules and Chronological Evaluation*

Let A[i,t] denote reported weekly influenza admissions in jurisdiction i for reporting week t. The models evaluated are:

$$\text{Persistence: } F_{persist}[i,t+1] = A[i,t].$$

$$\text{Nonnegative linear extrapolation: } F_{linear}[i,t+1] = \max(0, 2A[i,t] - A[i,t-1]).$$

$$\text{Damped trend: } F_{half}[i,t+1] = \max(0, A[i,t] + 0.5(A[i,t] - A[i,t-1])).$$

The non-negativity floor prevents negative predicted admissions. For jurisdictions, the two chronological target weeks are October 31 and November 7. The preceding two weeks are used at each origin to generate predictions. This yields 104 jurisdiction–target evaluations per model; the shared target weeks are strongly correlated and are not 104 independent temporal tests. A separate latest-week comparison reports the 52 jurisdictional forecasts for November 7. For the U.S. series, seven successive one-week rolling predictions are available after a minimum three-week initial history, and a trailing-three-week mean is included as a fourth baseline. No future target observation enters its own prediction.

All forecast rules are investigator-defined deterministic statistical baselines; none is a fitted or trained AI model. The 0.5 trend attenuation was not optimized on independent historical training data. The small number of chronological target weeks and correlated jurisdictional observations preclude a robust model-selection or predictive-superiority claim.

For observed y[j] and predictions f[j], mean absolute error is $MAE = \text{mean}(|y[j]-f[j]|)$, root mean squared error is $RMSE = \text{sqrt}(\text{mean}((y[j]-f[j])^2))$, and $WAPE = 100\times\text{sum}(|y[j]-f[j]|)/\text{sum}(y[j])$. WAPE describes aggregate absolute error relative to observed volume, and can be driven by high-count jurisdictions. It is not an individualized error rate. Model ranking over this short autumn rising period should not be extrapolated to winter peaks or declining phases.

Methodological equity safeguards were also specified by analogy to Akinyelure's review of predictive analytics in marginalized mental-health populations [18]. That work emphasizes the risk of unrepresentative samples and context-insensitive predictions. For a later AI-HealthCloud implementation, appropriate evaluation would separately report forecast errors and calibration, where applicable, across geographical and, only when lawfully supported by suitable data, demographic subgroups; it would also document differential reporting coverage and involve affected users in design review. None of these subgroup analyses was executable on the current short, aggregated dataset, so the present results do not establish fairness or cultural generalizability.

Akinyelure's earlier review of AI-enabled mental-health diagnostics [19] highlights interpretability, privacy, data bias, and participatory design as prerequisites for equitable use. These principles informed the proposed decision-control checklist—transparent forecast rules, documented variable provenance, human authorization before allocation, audit logs, and planned subgroup evaluation. They did not supply a trained prediction model to this experiment. Although the implemented baseline formulas are readily interpretable, interpretability by itself neither demonstrates unbiased decisions nor substitutes for performance or privacy validation.

➢ *Synthetic Service Capacity and Allocation Model*

In the absence of validated staffed-bed or staffing data, let $C[i] = 1.15\times A[i,Oct24]$ define an illustrative service-capacity threshold in abstract weekly-admission-equivalent units. This is not an observed NHSN bed inventory and cannot be converted directly into beds or clinicians. Let estimated additional need be $N_hat[i] = \max(0,F_half[i,Nov7]-C[i])$. A hypothetical divisible allocation budget B takes the values 600, 1,200, or 1,800 abstract units. Nonnegative assignments x[i] must satisfy $\text{sum}_i\ x[i] = B$.

Four deployable-in-principle decision rules were compared. An equal rule assigns B/52 to each jurisdiction. A previous-volume rule allocates B in proportion to A[i,Oct31]. A forecast-excess rule allocates B in proportion to N_hat[i]. Finally, an actual linear program minimizes a weighted sum of forecast unmet demand, subject to the budget, hypothetical minimum-allocation floor, and per-jurisdiction cap. The weights are based on recent relative admission growth; the floor assigns at least 10% of the budget in equal minimum shares, and the cap limits a jurisdiction to at most 25% of the budget. These policy parameters are illustrative research choices, not accepted hospital planning standards.

For the linear program, use nonnegative variables x[i] and u[i], and minimize $\text{sum}_i\ w[i]u[i]$, subject to $u[i] \geq N_hat[i]-x[i]$, $\text{sum}_i\ x[i] = B$, $x[i] \geq 0.10B/52$, and $x[i] \leq \min(0.25B, N_hat[i]+0.10B/52)$ where this upper bound is no smaller than the floor. The numerical implementation uses SciPy's HiGHS linear-programming solver. This is an executed mathematical optimization rather than a drawing of one. A perfect-hindsight oracle allocates in proportion to future observed excess; it is an unattainable reference bound and is never characterized as a deployable model.

➢ *Synthetic Scenario Outcomes*

The subsequent November 7 admissions are the observed check for the 0% surge case. Additional stress cases multiply those observed admissions by factors 1.20, 1.40, and 1.60. For an assumed scenario multiplier m, define $D[i,m] = m\times A[i,Nov7]$. The modeled excess before allocation is $E[i,m] = \max(0,D[i,m]-C[i])$, and modeled unmet demand after an assigned quantity x[i] is $U[i,m] = \max(0,E[i,m]-x[i])$. Total unmet demand is $\text{sum}_i\ U[i,m]$. Excess assigned units that exceed a jurisdiction's modeled requirement are categorized as wasted allocation; no real-world equipment or labor waste is inferred.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

The policy inputs for forecasting and allocation are restricted to October 24 and October 31 observations. The later November 7 observations are used to score the policies retrospectively, and hypothetical surge multipliers are applied only at the scoring stage. The oracle is explicitly excluded from prospective policy comparisons. The calculations are repeated at three budgets and four surge multipliers, providing a compact sensitivity grid. The framework currently optimizes only an abstract decision, not an operational patient transfer or staff roster.

➢ *Proposed Cloud-Ready Systems Design*

The proposed architecture consists of (1) public surveillance and authorized operational inputs; (2) versioned ingestion and schema validation; (3) quality checks covering reporting completeness, inconsistent totals, and regime changes; (4) explicitly timestamped forecasting; (5) a configurable allocation and optimization layer; (6) a simulator that records assumptions separately from observations; and (7) a human-reviewed reporting and auditing interface. These are design components, not tested cloud capabilities.

A future implementation might store immutable source objects in secured cloud object storage, use scheduled or event-triggered jobs for validation, maintain a versioned feature store, train models in containerized jobs, and serve forecast and planning outputs through controlled APIs. Such components are implementation options, not proof of interoperability or operational performance. In particular, proposed cloud-based deployment has not been completed here; no cloud-specific throughput, recovery-time, or cost claims are supported by these experiments. The local research scripts are intentionally portable but are not equivalent to a production system.

For the proposed interoperability layer, Manyura et al. [20] provide a directly relevant methodological comparator: their 2025 Medicaid health information-exchange review develops a staged sociotechnical design addressing readiness, standardization, data governance, AI analysis, privacy, equity, and evaluation, and discusses HL7 FHIR and API-oriented exchange. AI-HealthCloud uses these dimensions as requirements to be tested if future deployments ingest authorized hospital or Medicaid operational data. The current CDC-derived CSV analysis did not connect to Medicaid systems, exchange patient records, execute FHIR transactions, implement federated learning, or measure cross-institution interoperability; no such outcome is claimed.

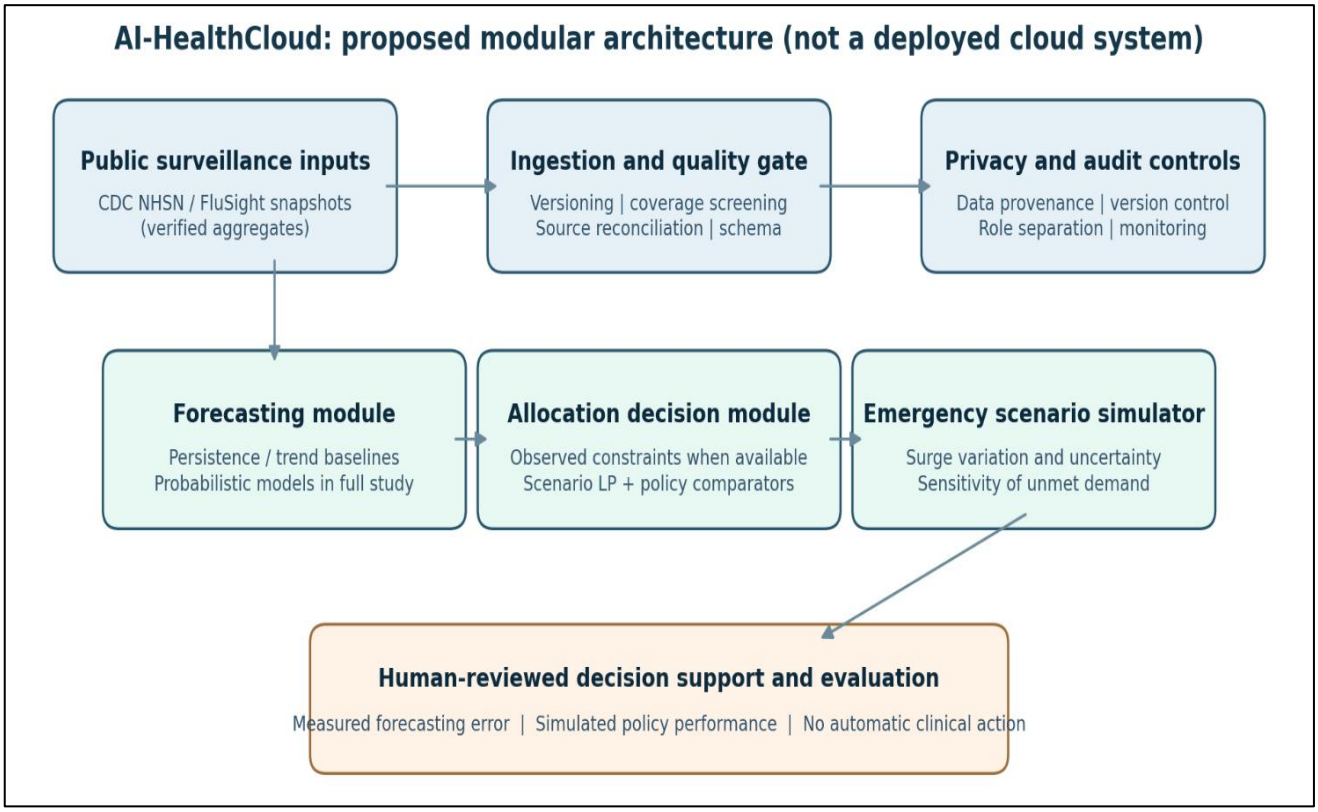


Fig 1 Proposed Architecture; Not a Deployed or Benchmarked Cloud System.

In a different critical-infrastructure context, Onaji et al. [21] propose integrating predictive monitoring, high-availability engineering, and trust assurance to support resilient communication systems. Their architecture offers a conceptual comparison for separating monitoring, decision support, and service-continuity controls in AI-HealthCloud. This citation is not evidence that the present healthcare prototype is fault-tolerant or secure: no blockchain, automated recovery, or cloud-availability benchmark was implemented in this study.

➢ *Reproducibility and Integrity Controls*

The accompanying supplementary archive contains the preserved working inputs, executable baseline and simulation

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

scripts, generated numerical outputs, SHA-256 input checksums, and figures. Running python run_proof_of_concept.py reproduces the deterministic forecast and allocation results; running extended_analysis.py reproduces exploratory uncertainty and sensitivity calculations. The full NHSN data audit script is included solely as a future-work implementation aid and was not executed against a complete source extraction. All figures should be read with their observed, modeled, simulated, or proposed status in mind.

IV. RESULTS AND DISCUSSION

➢ *Observed National and Geographic Patterns*

Within the preserved ten-week national series, reported influenza-associated hospital admissions rise from 1,256 for the week ending September 5 to 9,395 for the week ending November 7, 2026, an increase of 8,139 (approximately 648%). This describes counts in the available working extract, not population incidence, causal effects, or the full influenza season. Counts may be revised by the source. Because this sequence largely captures a rising period, it is unsuitable on its own for selecting a general-purpose seasonal forecasting model.

The November 7 jurisdictional sample has larger absolute counts for Florida, California, and Texas, while Puerto Rico and Louisiana display high reported admissions per 100,000 residents in the original rate extract. Counts, per-capita rates, available staffed capacity, and actual medical need measure different things; none should serve as an automatic allocation priority without explicit governance and operational information.

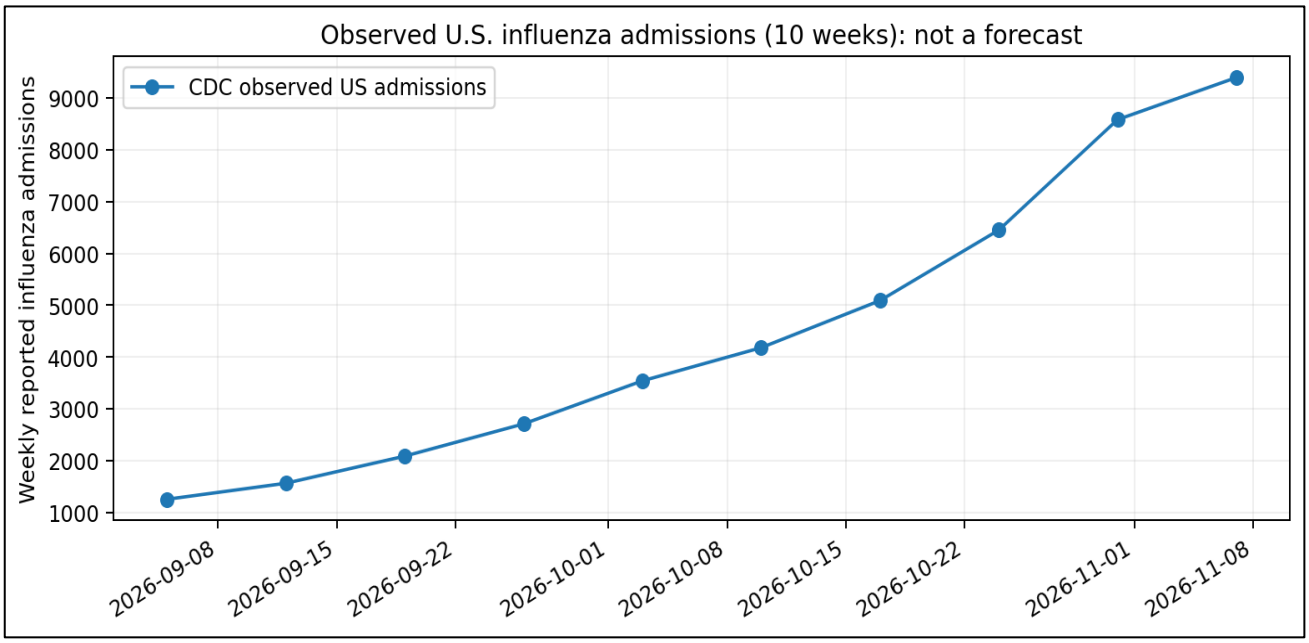


Fig 2 Ten-Week U.S. Published Influenza Admission Observations.

➢ *Jurisdictional Forecasting Comparisons*

For the latest forecast target (November 7, 2026), the damped-trend baseline records MAE 24.87, RMSE 40.45, and WAPE 13.77% across the 52 jurisdictions. Persistence records MAE 25.83, RMSE 46.00, and WAPE 14.30%; nonnegative linear extrapolation records MAE 34.48, RMSE 61.50, and WAPE 19.09%. These descriptive differences come from one forecast origin and are not proof of predictive superiority.

Pooling October 31 and November 7 targets produces 104 jurisdiction–target scored pairs but only two temporal origins. The damped-trend baseline has WAPE 18.38% and MAE 31.77, compared with WAPE 19.93% and MAE 34.45 for persistence; nonnegative linear extrapolation yields WAPE 20.61% and MAE 35.63. The forecast comparisons are exploratory, and jurisdictional observations share calendar shocks and reporting conditions.

An exploratory bootstrap resampling 52 jurisdiction clusters (20,000 replicates, fixed seed 20261114) estimated the damped-trend minus-persistence mean-absolute-error advantage at 2.68 admissions, with a 95% bootstrap percentile interval of −2.94 to +8.46 admissions. This interval crosses zero. Because all jurisdictions share only two evaluation weeks and common epidemic conditions, the interval does not account adequately for temporal uncertainty; it should not be treated as formal evidence of better future forecast performance.

Table 2 Two-Origin Evaluation Over 52 Jurisdictions (104 Scored Pairs Per Model)

Forecast Rule	MAE	RMSE	WAPE (%)
half trend	31.77	53.91	18.38
nonnegative linear trend	35.62	67.83	20.61
persistence	34.45	60.69	19.93

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

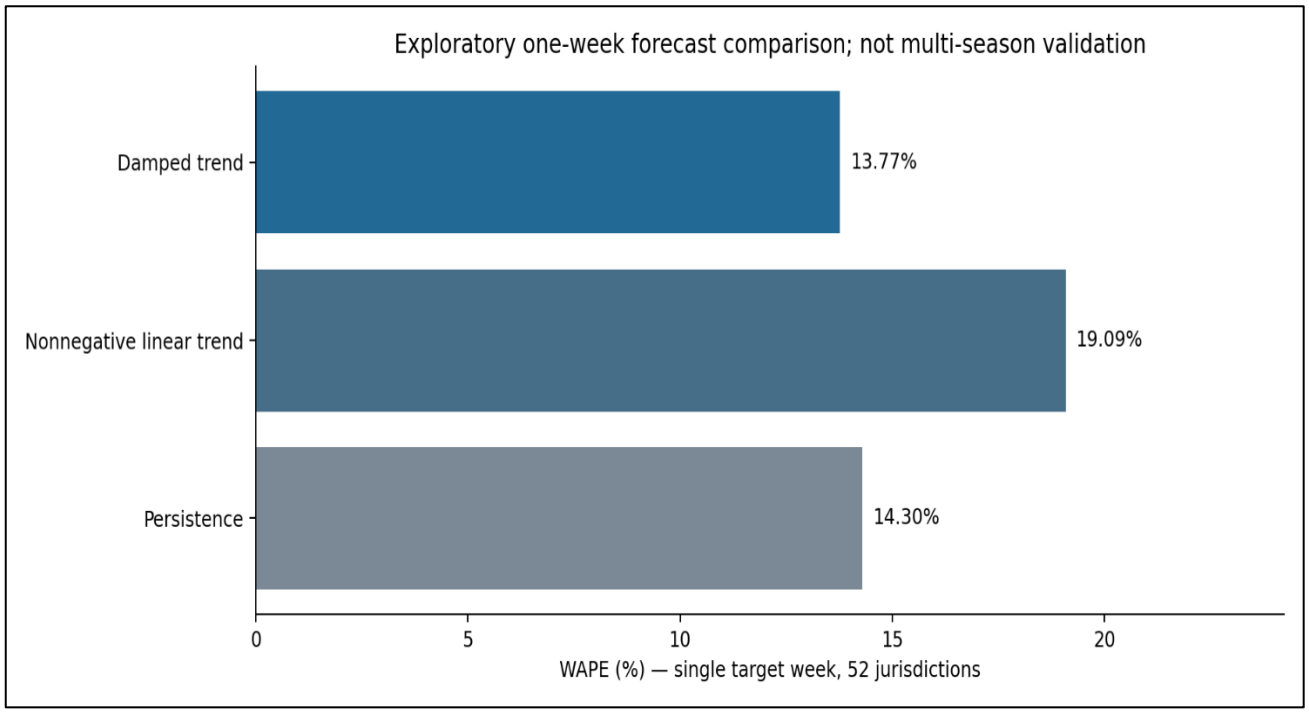


Fig 3 Latest Target-Week Comparison. This Single-Week Figure Differs from the Two-Origin Table.

➢ *Separate National Rolling-Baseline Comparison*

For seven national rolling one-week forecasts from September 26 through November 7, 2026, nonnegative linear extrapolation has WAPE 8.29%; damped trend has 10.78%, persistence 18.29%, and a trailing-three-week mean 34.34%. The ordering differs from the two-origin jurisdictional comparison, illustrating dependence on aggregation and the limited rising-period sample. No confidence interval or multi-season generalization is asserted for these national results.

Table 3 National Rolling One-Week Forecasts (Seven Origins; Limited Rising-Season Sample)

Forecast Rule	N Origins	MAE	WAPE (%)
nonnegative linear trend	7	473.0	8.29
half trend	7	615.1	10.78
persistence	7	1043.9	18.29
trailing three mean	7	1959.6	34.34

➢ *Hypothetical Allocation Comparison*

Under the illustrative capacity threshold C[i] = 1.15×October 24 admissions and the 1,200-unit hypothetical budget, the November 7 admissions yield 2,254.65 modeled excess units before allocation. Equal distribution leaves 1,572.08 modeled unmet units, prior-volume weighting 1,544.00, forecast-excess weighting 1,208.27, and the specified constrained linear program 1,400.57. A non-deployable hindsight oracle leaves 1,054.65. These figures describe arithmetic over assumed service-unit thresholds; they are not observed hospital shortfalls.

In this single defined scenario, forecast-excess weighting generates 23.14% fewer modeled unmet units than equal distribution. This difference is conditional on the chosen inputs, model objective, and allocation constraints. It cannot be interpreted as prevented hospital admissions, avoided patient harm, released beds, improved clinical outcomes, or financial savings. The constrained linear program performs worse on the reported unweighted unmet-unit metric because its floor, cap, and weighting choices change the optimization objective; this does not establish that either policy is clinically preferable.

Table 4 Conditional Simulation at 1,200 Synthetic Allocation Units

Policy	Unmet Proxy Units	Unused Proxy Units
equal	1,572	517
prior admissions weighted	1,544	489
forecast excess weighted	1,208	154
forecast lp priority	1,401	346
oracle hindsight upper bound	1,055	0

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

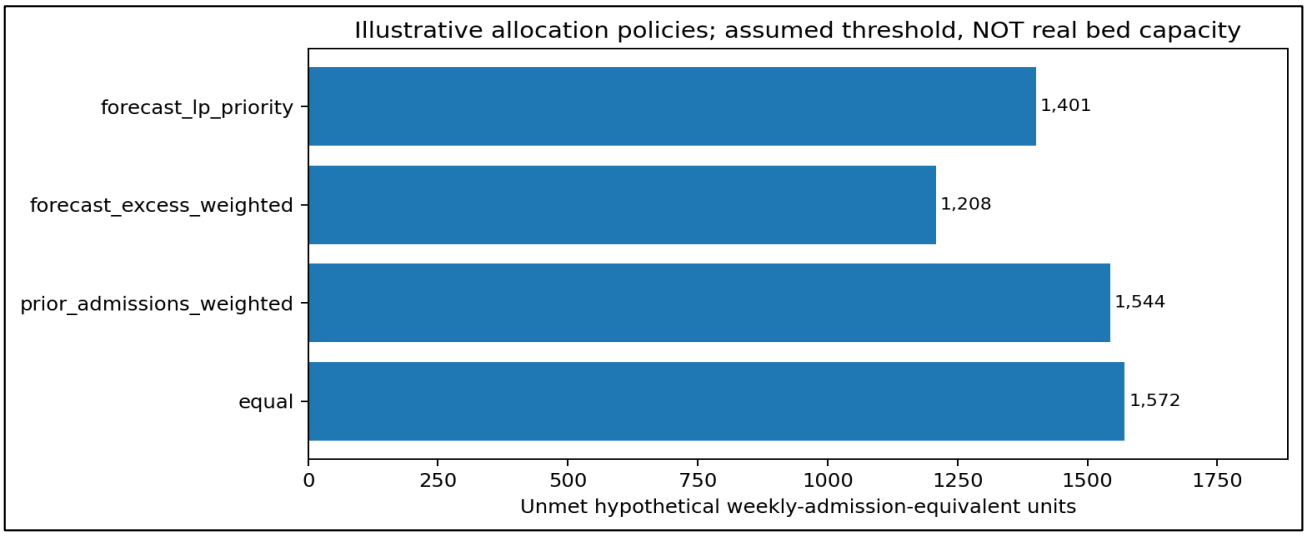


Fig 4 Simulated Unmet Abstract Service Units, not Real Clinical Shortages.

➢ *Stress Testing and Sensitivity to Assumptions*

For the 1,200-unit resource budget and assumed additional demands of 0%, 20%, 40%, and 60%, equal allocation leaves approximately 1,572, 3,064, 4,793, and 6,585 modeled unmet units, whereas forecast-excess weighting leaves 1,208, 2,735, 4,563, and 6,429. The relative modeled difference declines from approximately 23.1% to 2.4% as simulated demand rises. These values are deterministic sensitivity outputs for synthetic resource units, not measured public-health emergency outcomes.

Further exploratory tests vary the synthetic capacity multiplier over 0.90, 1.00, 1.15, 1.30, and 1.50; assumed additional demand over 0%, 20%, 40%, and 60%; and budgets over 600, 1,200, and 1,800 units, giving 60 scenarios. Forecast-excess weighting leaves fewer modeled unmet units than equal allocation in all 60 scenarios, but is not the lowest-unmet policy against every alternative rule in every scenario. These outcomes are mathematically contingent on the selected capacity proxies, budget, service-unit definition, forecasts, and demand realizations, and are not 60 independent validation trials. Different operational constraints could reverse the policy ranking.

Table 5 Modeled Unmet Units by Demand-Surge Scenario, B=1,200

Extra surge	Equal	Forecast-weighted	Constrained LP
0%	1,572	1,208	1,401
20%	3,064	2,735	2,829
40%	4,793	4,563	4,599
60%	6,585	6,429	6,444

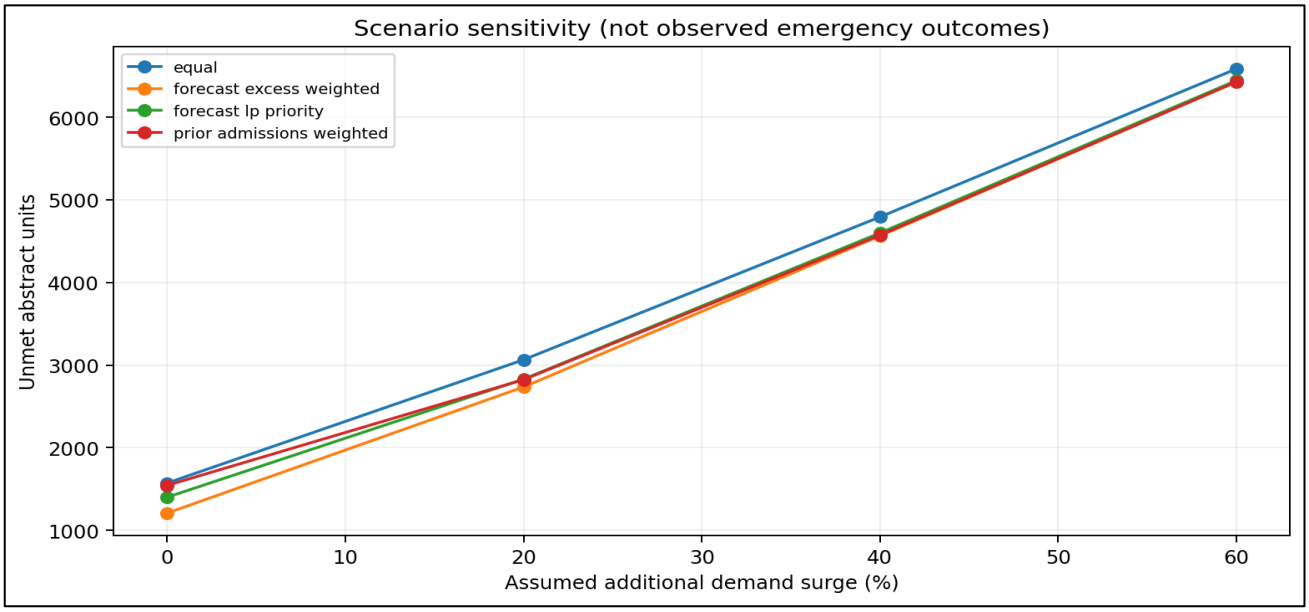


Fig 5 Stress Sensitivity; 20–60% Demand Increases are Hypothetical Scenarios.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

➢ *Evidence Limits and Interpretation*

The prototype demonstrates execution of transparent statistical baselines, chronological forecast scoring, deterministic allocation rules, a constrained linear program, and sensitivity calculations. It does not demonstrate trained-AI accuracy, outbreak detection, actual hospital capacity optimization, enhanced patient care, economic savings, secure cloud operation, or measured emergency readiness. The source panel is incomplete, four national-to-jurisdiction discrepancies remain unresolved, only two jurisdictional evaluation dates are available, and staffing, occupancy, length-of-stay, and transfer feasibility are unmeasured. These are substantive limits on inference.

V. CONCLUSIONS AND FURTHER RESEARCH

➢ *Summary of Principal Findings*

This proof of concept establishes that simple statistical admission forecasts can be linked to transparent hypothetical allocation policies and audited with a reproducible workflow. Results differ between the seven national and two jurisdictional forecast origins; no forecast rule is validated as superior across influenza seasons. Under specified assumptions, forecast-excess weighting reduces simulated unmet units relative to equal allocation; an additional constrained linear program does not reduce the reported unweighted unmet-unit outcome as much. The proposed cloud architecture is a design, not an implemented system.

➢ *Contribution and Relationship to Prior Work*

The study contributes an explicit feasibility protocol that separates observed admission series, retrospective deterministic forecasts, hypothetical capacity and surge assumptions, and proposed system components. Because prior research already integrates epidemic forecasting and resource planning [7], [8], the present analysis is not positioned as the first such integration or as evidence of a novel validated algorithm. Its value is methodological transparency and a reproducible starting point for independent evaluation.

➢ *Requirements for a Larger Empirical Evaluation*

Future empirical work should acquire and archive a complete, versioned NHSN/FluSight extraction; quantify completeness by jurisdiction and reporting regime; reconcile geographic totals; and evaluate forecasting rules over multiple complete respiratory seasons using prespecified chronological training and testing. Seasonal baselines, appropriate learned models, and calibrated uncertainty forecasts should be compared with MAE, WAPE, interval coverage, and relevant proper scoring rules.

Operationally meaningful allocation research would require observed or defensibly sourced staffed capacity, bed occupancy, patient length of stay, admissions-to-census conversion, transfer feasibility, and distribution or redeployment costs. Optimization should target clinically meaningful outcomes and test fairness constraints, rather than assuming that lower abstract unmet units imply better patient care. Without those inputs, simulation findings should remain explicitly hypothetical.

A further implementation study could deploy the proposed pipeline in an appropriately governed cloud environment and evaluate ingestion reliability, data quality, processing latency, access control, computing costs, failure recovery, and audited user decisions. Such metrics are not available from a systems diagram or local scripts. Future access to identifiable clinical data would require a separate privacy, regulatory, and institutional evaluation.

➢ *Limitations*

The study is limited by a ten-week national working series, four manually assembled jurisdictional weeks, unresolved component-to-national discrepancies, two jurisdictional forecast origins, lack of a complete source-file archive, investigator-selected coefficients, absent predictive intervals, and synthetic service-capacity and surge assumptions. There is no deployment, staffing, census, transfer, cost, patient outcome, or causal evidence. A synthetic weekly-admission-equivalent resource unit is not interchangeable with one hospital bed, one clinical shift, one avoided admission, or any monetary amount. The findings are preliminary, context-sensitive computational demonstrations.

➢ *Conclusion*

The AI-HealthCloud architecture can be described and its forecast-to-decision logic can be executed and reproduced using a restricted influenza admission dataset. Baseline results and synthetic allocation outcomes demonstrate feasibility of the analytical workflow, not superior clinical prediction or validated healthcare resource optimization.

The limited observations, manually assembled source panel, and unobserved operational resource constraints prevent claims of improved hospital preparedness, lower healthcare expenditure, reduced patient harm, or a production-grade AI/cloud implementation. The architecture should therefore be interpreted as a proposed decision-support framework supported by an exploratory computational feasibility study.

A larger study with audited multi-season data, independently justified resource measures, and external validation would be needed to determine whether any forecast-guided policy produces robust operational benefits beyond existing forecasting and optimization approaches.

DATA AND CODE AVAILABILITY

The observations originated from the publicly accessible CDC FluSight hospitalization target series. The four-week jurisdictional panel was locally preserved, with three jurisdictional weeks manually transcribed; the complete original source archive was not captured locally. The study supplies the assembled input files, baseline and scenario scripts, source-reconciliation output, uncertainty and sensitivity calculations, and result figures in an accompanying supplementary archive. No stable public code repository has yet been deposited. National/jurisdiction differences are disclosed and remain unresolved.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

ETHICS AND REPORTING DECLARATIONS

Only aggregate, publicly accessible surveillance values were analyzed; no patient-level records or human participants were directly recruited. Numerical results are classified as (O) observed working-extract values, (F) retrospectively scored forecasts, (S) hypothetical scenario outputs, or (P) proposed but unimplemented system components. Author affiliations, contributions, funding, competing-interest disclosures, and any journal-required ethics declarations must be supplied and verified by the authors for a specific submission; they are not inferred in this report.

SUPPLEMENTARY METHODS AND REPRODUCIBILITY NOTES

The supplementary archive contains run_proof_of_concept.py, extended_analysis.py, preserved national and jurisdictional input CSVs, national and jurisdictional forecasting outputs, source reconciliation, allocation and sensitivity outputs, calculation manifests, and Figures 1–5. The file labeled full_CDC_NHSN_audit_ready_not_run.py is a future-source-audit script and was not executed against the complete NHSN archive.

Run python run_proof_of_concept.py and python extended_analysis.py in the extracted archive after installing the listed Python dependencies (NumPy, pandas, SciPy, and Matplotlib). The results were reproduced using the locally preserved working extract, not a full externally downloaded CDC archive. The provenance and reconciliation fields must accompany any reuse of the data tables.

The original source-derived input values, code, and generated outputs are retained separately. The synthetic scenario budgets, capacity multipliers, minimum-allocation rules, and demand surges are specified in the scripts and should not be interpreted as published clinical standards or measured physical resources.

Interpretation rule: Do not combine observed admissions (O), retrospective forecast scores (F), hypothetical allocation outcomes (S), and untested cloud-design elements (P) into a claim of demonstrated healthcare-system effectiveness.

REFERENCES

[1]. Centers for Disease Control and Prevention. (2026a, November 13). Weekly hospital respiratory data (HRD) metrics by jurisdiction, National Healthcare Safety Network (NHSN) [Data set]. https://data.cdc.gov/d/ua7e-t2fy

[2]. Centers for Disease Control and Prevention. (2026b, September 30). FluSight 2025–2026 evaluation. https://www.cdc.gov/flu-forecasting/evaluation/2025-2026-report.html

[3]. Centers for Disease Control and Prevention. (2026c). NHSN hospital respiratory data (HRD) reporting protocol. https://www.cdc.gov/nhsn/pdfs/pscmanual/HRD-Protocol-Final.pdf

[4]. CDC FluSight Forecast Hub. (2026). Target hospital admissions [Data file]. https://raw.githubusercontent.com/cdcepi/FluSight-forecast-hub/main/target-data/target-hospital-admissions.csv

[5]. Frutos, A. M., Hines, A., Shandross, L., Ray, E. L., Reich, N. G., Borchering, R. K., & Biggerstaff, M. (2026). Retrospective evaluation of trained and untrained probabilistic ensemble forecasts for influenza hospital admissions—United States, 2022–2025. Infectious Disease Modelling, 11(3), 1227–1233. https://doi.org/10.1016/j.idm.2026.04.002

[6]. Mukasa, K. (2026). AI-driven forecasting for Medicarerisk score modernization. International Journal for Multidisciplinary Research, 8(1). https://doi.org/10.36948/ijfmr.2026.v08i01.66658

[7]. Fattahi, M., Keyvanshokooh, E., Kannan, D., & Govindan, K. (2023). Resource planning strategies for healthcare systems during a pandemic. European Journal of Operational Research, 304(1), 192–206. https://doi.org/10.1016/j.ejor.2022.01.023

[8]. Jiang, J., Sheng, D., Chen, X., Tian, Q., Li, F., & Yang, P. (2024). Data-driven collaborative healthcare resource allocation in pandemics. Transportation Research Part E: Logistics and Transportation Review, 192, Article 103828. https://doi.org/10.1016/j.tre.2024.103828

[9]. Johnson, M. R., Naik, H., Chan, W. S., Greiner, J., Michaleski, M., Liu, D., Silvestre, B., & McCarthy, I. P. (2023). Forecasting ward-level bed requirements to aid pandemic resource planning: Lessons learned and future directions. Health Care Management Science, 26, 477–500. https://doi.org/10.1007/s10729-023-09639-2

[10]. Begen, M. A., Rodrigues, F. F., Rice, T., & Zaric, G. S. (2024). A forecasting tool for a hospital to plan inbound transfers of COVID-19 patients from other regions. BMC Public Health, 24, Article 505. https://doi.org/10.1186/s12889-024-18038-3

[11]. Onaji, V., Adediji, E., Zeyeum, J. N., Ayano, K., & Olufemi, D. (2025). Deep reinforcement learning for dynamic network slicing and resource orchestration in software-defined critical telecom infrastructure. International Journal of Computer Applications Technology and Research, 14(11), 53–73. https://doi.org/10.7753/IJCATR1411.1006

[12]. Collins, G. S., Moons, K. G. M., Dhiman, P., Riley, R. D., Beam, A. L., Van Calster, B., Ghassemi, M., Liu, X., Reitsma, J. B., van Smeden, M., Boulesteix, A.-L., Camaradou, J. C., Celi, L. A., Denaxas, S., Denniston, A. K., Glocker, B., Golub, R. M., Harvey, H., Heinze, G., ... Logullo, P. (2024). TRIPOD+AI statement: Updated guidance for reporting clinical prediction models that use regression or machine learning methods. BMJ, 385, e078378. https://doi.org/10.1136/bmj-2023-078378

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov961

[13]. Centers for Disease Control and Prevention. (2026d). CDC's hospital respiratory data (HRD) methodology. https://www.cdc.gov/nhsn/psc/hospital-respiratory-reporting/data-methods.html
[14]. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. https://www.who.int/publications/i/item/9789240029200
[15]. Ogunsakin, O. L., & Mukasa, K. (2026). Navigating ethics and regulation in AI health cloud: Challenges and opportunities for United States healthcare. International Journal of Innovative Science and Research Technology, 11(1), 942–952. https://doi.org/10.38124/ijisrt/26jan367
[16]. Manyura, S. (2026). Artificial intelligence for secure health information exchange: Advancing healthcare operations through intelligent interoperability. International Journal of Computer Applications Technology and Research [bibliographic listing; publisher full text and publication metadata not independently verified]. ResearchGate. https://www.researchgate.net/publication/410754414_Artificial_Intelligence_for_Secure_Health_Information_Exchange_Advancing_Healthcare_Operations_Through_Intelligent_Interoperability
[17]. Akinyelure, F. (2025). Leveraging behavioural health data for policy innovation: Closing the loop between community insights and public health decision-making. International Journal of Innovative Science and Research Technology, 10(7), 3458–3466. https://doi.org/10.38124/ijisrt/25jul1532
[18]. Akinyelure, F. M. (2023). Bridging the gap: Integrating predictive analytics with culturally competent mental health care delivery in marginalized populations. International Journal of Research in Psychiatry, 3(2), 12–17. https://doi.org/10.22271/27891623.2023.v3.i2a.76
[19]. Akinyelure, F. M. (2021). AI in mental health diagnostics: Ethical imperatives and design strategies for equitable implementation. International Journal of Research in Medical Science, 3(2), 14–19. https://doi.org/10.33545/26648733.2021.v3.i2a.167
[20]. Manyura, S., Alo, A., Saaka, D. K., & Kabera, N. R. (2025). A methodological framework for developing AI-enabled health information interoperability solutions to improve secure health information exchange across Medicaid ecosystems. World Journal of Advanced Research and Reviews, 28(2), 2667–2679. https://doi.org/10.30574/wjarr.2025.28.2.3712
[21]. Onaji, V., Kangethe, L., & Usoh, R. (2026). AI-enabled ICT resilience architecture for high-availability, secure, and blockchain-assured communication systems. International Journal of Innovative Science and Research Technology, 11(1), 1669–1683. https://doi.org/10.38124/ijisrt/26jan696