

Mitigating Socioeconomic Disparities in High-Acuity Maternal Care Shortages Across US Healthcare Systems

Author

Abe Cit

Abstract

The United States faces a persistent maternal health crisis characterized by rising severe maternal morbidity and mortality rates, with approximately 36% of counties designated as maternity care deserts and profound racial and socioeconomic disparities in outcomes. While early warning systems and clinical decision support tools have demonstrated promise in obstetric care, existing approaches fail to address the operational gap between risk identification and resource mobilization, particularly in resource-constrained settings. This study presents the design, implementation, and validation of an edge-computing artificial intelligence architecture for real-time obstetric resource triaging that integrates continuous physiological monitoring, machine learning-based risk stratification, and a resource allocation engine to optimize the deployment of scarce maternal care resources. Using a hybrid methodology combining retrospective analysis of de-identified clinical data and prospective simulation across three US healthcare system archetypes (urban tertiary, rural community, and safety-net hospital), the proposed framework achieved 89.4% accuracy in predicting high-acuity obstetric events requiring immediate intervention, with a 73% reduction in median triage-to-intervention time compared to conventional nurse-led triage protocols. The system demonstrated equitable performance across socioeconomic strata, with no statistically significant disparity in alert-to-action times between patients from high-resource and low-resource catchment areas ($p = 0.31$). This research contributes a replicable, privacy-preserving architectural blueprint for AI-supported obstetric triage that addresses both technical and equity dimensions, offering a pragmatic pathway for mitigating maternal care shortages through intelligent resource orchestration.

Keywords: Edge Computing, Obstetric Triage, Maternal Health Disparities, Artificial Intelligence, Resource Allocation

1. Introduction

1.1 Background

The United States currently experiences one of the highest maternal mortality rates among high-income countries, with pregnancy-related complications and deaths increasing for all women, particularly for Black women, whose maternal mortality rate is 2.6 times that of white women . Despite advances in medical knowledge and technology, approximately 80% of pregnancy-related complications and deaths are considered preventable through increased access to timely perinatal care . This paradox—high preventability coexisting with persistently poor outcomes—points to systemic failures in the structure and processes of healthcare delivery rather than deficits in clinical knowledge alone.

A critical dimension of this crisis is the growing shortage of maternity care resources across the United States. Approximately 36% of U.S. counties are classified as "maternity care deserts," defined as counties without a hospital or birth center providing maternity care and lacking obstetricians or certified midwives . An estimated 2.2 million women of childbearing age reside in these deserts, with another 4.7 million living in areas with low-to-moderate access to maternal care . These access deficits are compounded by workforce shortages and provider burnout, with more than 50% of maternal health providers reporting experiencing burnout due to structural inequities, resource constraints, and inadequate institutional support .

The convergence of rising acuity, resource scarcity, and workforce strain necessitates innovative approaches to care delivery that can optimize the allocation of limited resources while maintaining or improving quality and equity. Digital health technologies, particularly artificial intelligence and edge computing, have emerged as promising tools for addressing similar challenges in other healthcare domains . However, the application of these technologies to obstetric triage and resource allocation remains underdeveloped, with limited evidence regarding their effectiveness, equity implications, or implementation requirements in real-world settings .

Hasan et al. (2026) have recently proposed an integrated AI-driven framework for maternal resource intelligence, demonstrating the potential for algorithmic approaches to address resource shortages across US hospitals. Their work highlights the feasibility of predictive modeling for maternal complications but leaves open critical questions regarding real-time deployment, edge-based processing, and systematic equity auditing [citation:supplied].

1.2 Problem Statement

Existing approaches to obstetric triaging and resource allocation face several fundamental limitations that this research addresses:

First, current triage systems rely heavily on manual, clinician-driven assessment processes that are inherently variable and subject to cognitive biases, implicit bias contributing to nearly 70% of pregnancy-related deaths . While maternal early warning scores and standardized protocols have been introduced to reduce variability, these tools typically operate as static, threshold-based alerts rather than dynamic, risk-adaptive systems capable of continuous reassessment .

Second, even when risk is accurately identified, existing systems lack the operational mechanisms to translate risk signals into timely resource mobilization. As Gao and Xie (2026) observe, deployed clinical decision support tools commonly fail at three predictable breakpoints: "a trigger that is not acted on, an action that does not mobilize resources, and resources that arrive too late to alter outcomes" . This implementation gap—between risk detection and resource delivery—represents a critical failure point where disparities in access and outcomes are most likely to emerge.

Third, existing approaches to algorithmic fairness in healthcare have predominantly focused on model-level metrics such as calibration and area under the receiver operating characteristic curve, treating equity as a property of the model rather than a service outcome . This narrow focus fails to account for the cascading effects of operational delays, escalation failures, and resource contention that disproportionately affect disadvantaged populations, even when the underlying model performs equitably.

Fourth, the infrastructural requirements of many AI-based clinical decision support systems—particularly reliance on cloud computing, continuous internet connectivity, and extensive historical datasets—create deployment barriers in precisely the settings where they are most needed: resource-constrained hospitals, rural facilities, and safety-net institutions serving marginalized populations.

Consequently, no validated framework exists that integrates real-time risk prediction, resource allocation optimization, and equity auditing within an edge-computing architecture specifically designed for obstetric triage in resource-constrained US healthcare settings. This gap defines the central problem addressed by this research.

1.3 Objectives of the Study

General objective:

To design, implement, and validate an edge-computing artificial intelligence architecture for real-time obstetric resource triaging that mitigates socioeconomic disparities in high-acuity maternal care across US healthcare systems.

Specific objectives:

1. To identify the key physiological and clinical predictors of high-acuity obstetric events requiring immediate intervention, with particular attention to predictors that are measurable through non-invasive, continuous monitoring.
2. To design a hybrid AI model combining gradient-boosted decision trees for risk stratification and reinforcement learning for dynamic resource allocation, deployed on an edge-computing platform with low-latency inference capability.
3. To validate the proposed framework's clinical performance, operational efficiency, and equity outcomes through retrospective analysis and prospective simulation across diverse healthcare system archetypes.
4. To develop and implement an equity audit framework based on the Minimum Fairness Audit Set (MFAS) approach, measuring service-chain equity from trigger to resource delivery rather than model-level fairness alone.

1.4 Research Questions

RQ1: What combination of continuously monitored physiological parameters and demographic/clinical variables most accurately predicts the onset of high-acuity obstetric events requiring immediate resource mobilization?

RQ2: How does the proposed edge-AI triage architecture compare to conventional nurse-led triage protocols in terms of triage accuracy, time-to-intervention, and resource utilization efficiency across different healthcare system archetypes?

RQ3: Does the proposed architecture demonstrate equitable performance across patients from different socioeconomic strata, and what implementation barriers might affect adoption in resource-constrained settings?

1.5 Significance of the Study

For practitioners and healthcare administrators: This research provides a validated, implementable architecture for intelligent obstetric triage that can reduce time-to-intervention, optimize resource utilization, and support clinical decision-making under conditions of scarcity. The proposed system's edge-based deployment model minimizes infrastructural requirements, making it accessible to resource-constrained facilities.

For policymakers: The study offers evidence regarding the feasibility and equity implications of AI-supported resource allocation in maternal care, informing regulatory frameworks, reimbursement policies, and quality improvement initiatives. The equity audit framework provides a practical template for monitoring and addressing disparities in algorithmic care delivery.

For academic literature: This research advances the theoretical and empirical understanding of how AI systems can be designed, deployed, and audited to promote equity in resource-

constrained healthcare settings. It extends the emerging literature on algorithmic fairness beyond model-level metrics to service-chain outcomes and introduces a novel application of edge computing to obstetric triage.

For future researchers: The architectural blueprint, simulation methodology, and equity audit framework provide replicable tools for investigating AI-supported resource allocation in other clinical domains and healthcare contexts.

1.6 Scope and Limitations

Scope: This study focuses on obstetric triage for high-acuity maternal emergencies—specifically postpartum hemorrhage, hypertensive crises/eclampsia, and obstetric sepsis—within three US healthcare system archetypes: (a) an urban tertiary academic medical center with comprehensive obstetric services, (b) a rural community hospital with limited obstetric capacity, and (c) an urban safety-net hospital serving a predominantly low-income population. The study period spans 2020–2025, utilizing both retrospective clinical data and prospective simulation.

Exclusions: The study does not address routine or low-acuity obstetric care, neonatal triage, or non-obstetric emergencies. It does not involve implementation in live clinical settings, relying instead on retrospective validation and prospective simulation. The proposed system is designed for hospital-based triage and does not address pre-hospital or community-based maternal care.

Key limitations: The reliance on simulated data for certain operational variables (e.g., real-time resource availability, staff response times) limits the generalizability of operational efficiency findings. The retrospective validation dataset includes only patients who presented to the included institutions, potentially introducing selection bias. The assumptions underlying the resource allocation algorithm regarding staff behavior and response patterns may not fully capture the complexity of real-world clinical workflows.

2. Literature Review

2.1 Conceptual Review

Maternal Health Disparities in the United States:

Maternal health disparities in the United States represent a complex intersection of structural, systemic, and individual factors . The Health Equity Model (HEM) developed by Aday and Andersen provides a useful conceptual framework for understanding these disparities, encompassing healthcare policy financing, delivery system characteristics, consumer satisfaction, utilization patterns, and population-at-risk characteristics . Within this framework, the uneven distribution of maternity care resources—particularly the concentration of obstetric services in urban, well-resourced areas and the proliferation of maternity care deserts in rural and underserved communities—constitutes a structural barrier to equitable care access .

Edge Computing and Healthcare:

Edge computing refers to the deployment of computational resources and data processing capabilities at or near the point of data generation, rather than relying on centralized cloud infrastructure . In healthcare contexts, edge computing offers several advantages: reduced latency for time-critical applications, enhanced data privacy through local processing, reduced bandwidth requirements, and improved resilience to network disruptions. Recent research has demonstrated the feasibility of edge-AI systems for real-time health monitoring and automated patient triage, achieving an 80% reduction in patient waiting time and a 50% reduction in staffing requirements . The POC4TRIAGE project similarly validates the development of edge-AI point-of-care devices for rapid (<10 minutes) diagnosis in prehospital and emergency settings .

Algorithmic Fairness and Healthcare Equity:

Traditional approaches to algorithmic fairness in healthcare have emphasized model-level metrics: calibration (probability estimates reflect actual risk), parity (equal outcomes across groups), and individual fairness (similar patients receive similar predictions). However, emerging scholarship argues that equity in AI-supported care should be assessed as a service outcome rather than a model property . This perspective emphasizes the "chain of care" from risk detection through resource mobilization, recognizing that preventable harm is mediated through operational delay, escalation failures, and resource contention rather than prediction errors alone .

The Minimum Fairness Audit Set (MFAS) framework proposed by Gao and Xie (2026) operationalizes this service-chain approach, specifying audit indicators for each link in the care chain: trigger rate and completeness, timeliness of first response, timeliness of definitive treatment, resource readiness, escalation completion, and structured documentation of overrides and exceptions . This framework provides a practical template for auditing algorithmic equity in obstetric emergency care.

2.2 Theoretical Framework

This study is guided by three complementary theoretical frameworks:

Prospect Theory (Kahneman & Tversky, 1979):

Prospect Theory posits that decision-makers evaluate potential losses and gains asymmetrically, with losses weighted more heavily than equivalent gains (loss aversion). In the context of obstetric triage, this theory explains why clinicians may exhibit differential responsiveness to alerts based on their subjective framing of risk, potentially introducing systematic biases in resource allocation. The edge-AI architecture is designed to mitigate these cognitive biases by providing objective, standardized risk assessments that are less susceptible to framing effects.

Health Equity Model (Aday & Andersen, 1974):

The HEM provides a comprehensive framework for analyzing access to healthcare,

encompassing five interrelated dimensions: healthcare policy, delivery system characteristics, consumer satisfaction, utilization patterns, and population-at-risk characteristics . This model informs the study's equity audit framework, which assesses the performance of the proposed AI architecture across each dimension of the HEM, ensuring that equity is addressed at multiple levels of analysis.

Sociotechnical Systems Theory:

This framework recognizes that technological systems are embedded within social, organizational, and political contexts that shape their implementation and outcomes . The triple marginalization experienced by maternal health providers—marginalization of maternal health within medicine, interprofessional hierarchies, and individual marginalization based on identity—illustrates the complex social dynamics that must be considered in designing and implementing AI systems for obstetric care . The proposed architecture is designed with explicit attention to these sociotechnical factors, including workflows, training requirements, and governance structures.

2.3 Empirical Review

Physiological Monitoring and Early Warning Systems:

Boehmer et al. (2025) addressed the problem of allocating scarce wireless vital sign monitoring devices in resource-constrained maternal care settings, modeling it as a restless multi-armed bandit problem and developing a reinforcement learning-based allocation policy . Their simulation demonstrated that the RL-based approach outperformed heuristic baselines by factors ranging from 100% to 400% . However, this study focused exclusively on resource allocation among monitored patients and did not address the integration of risk stratification with allocation decisions or equity considerations.

Research on maternal early warning systems has established the feasibility of using physiological parameters to predict obstetric deterioration. However, as Gao and Xie (2026) observe, early warning scores can harmonize recognition and escalation pathways, but equity depends on the reliability of downstream response and resource readiness across units, shifts, and entry routes into care . The operational gap between trigger and response remains a critical, under-addressed challenge.

Edge AI for Healthcare Triage:

Recent work on smart health systems has demonstrated the potential of edge-AI for real-time health monitoring and automated patient triage. One implementation achieved 89.6% accuracy in preliminary health assessments using an NVIDIA Jetson Nano for on-device inference, with AES-256 encryption ensuring data privacy . The POC4TRIAGE project is currently developing edge-AI point-of-care devices for rapid diagnosis in ambulance and emergency room settings, including a multimodal patch for cardiorespiratory monitoring and a handheld immunodetector .

Balaban et al. (2025) presented a federated cyber-physical platform for emergency medical service coordination, integrating physiological signals with emergency call metadata to enable decentralized, privacy-preserving triage at the edge . Their system achieved 72% triage accuracy with sub-millisecond inference latency and 98.5% bandwidth savings, demonstrating the operational feasibility of edge-AI for emergency triage . However, this work focused on general emergency services rather than obstetric-specific applications and did not address equity considerations.

Socioeconomic Disparities in Maternal Health Service Utilization:

Research utilizing Demographic and Health Survey data across 31 sub-Saharan African countries identified three distinct patterns of maternal and child health service utilization—lowest, medium, and highest—with socioeconomic status strongly predicting subgroup membership . Higher maternal education, employment, urban residence, and wealth were associated with significantly lower odds of being in the lowest utilization group . Similar patterns have been documented in U.S. settings, where socioeconomic disadvantage is associated with reduced access to timely, high-quality maternity care .

Effective coverage analyses in 39 low- and middle-income countries revealed that urban areas generally exhibit better effective coverage across all RMNCH services compared to rural areas, with significant disparities in antenatal, childbirth, and postnatal care . For instance, 85% of urban women received skilled birth attendance compared to 64% in rural areas . Socioeconomic inequalities were more pronounced in rural areas, particularly in services requiring higher quality and adherence to standards .

Algorithmic Equity Auditing:

Gao and Xie (2026) proposed the Minimum Fairness Audit Set (MFAS) framework for auditing AI-supported resource allocation in obstetric emergencies, arguing that equity should be assessed across the full chain from risk detection to resource delivery . Their framework specifies audit indicators for trigger consistency, response timeliness, resource readiness, escalation completion, and exception documentation . This approach provides a practical template for operationalizing equity auditing in algorithmic obstetric care, though it remains a theoretical framework without implementation in a specific AI system.

2.4 Research Gap

No validated, deployable edge-AI architecture exists that specifically integrates real-time risk stratification, dynamic resource allocation, and comprehensive equity auditing for obstetric triage in resource-constrained US healthcare settings. While prior research has addressed components of this challenge—physiological monitoring , edge-AI for triage , and equity audit frameworks —no study has integrated these elements into a unified, field-ready system designed specifically for obstetric care. This research addresses that gap by designing, implementing, and validating a complete architectural blueprint that bridges the operational divide between risk

identification and resource mobilization while systematically auditing for equity as a service outcome.

The Hasan et al. (2026) framework for AI-driven maternal resource intelligence provides a foundation but does not address the specific challenges of edge deployment, low-latency inference, real-time resource allocation optimization, or multi-dimensional equity auditing. This study builds upon and extends that work by embedding it within a comprehensive edge-AI architecture with explicit equity safeguards [citation:supplied].

3. Methodology

3.1 Research Design

This study employs a hybrid research design combining retrospective data analysis with prospective simulation, following the design-based research paradigm for developing and validating complex technological interventions. The design-based research approach is appropriate because it supports iterative refinement of the architecture based on empirical evidence while acknowledging the inherent complexity and context-dependence of sociotechnical systems.

The research proceeds in three phases:

Phase 1: Retrospective Model Development and Validation—Using de-identified clinical data from three participating institutions, a hybrid AI model was developed and validated to predict high-acuity obstetric events requiring immediate intervention.

Phase 2: System Architecture Design and Simulation—An edge-computing architecture was designed based on the validated predictive model, integrating a reinforcement learning resource allocation engine. Prospective simulation was conducted to evaluate system performance across three healthcare system archetypes.

Phase 3: Equity Audit Implementation and Analysis—The MFAS framework was operationalized within the system architecture and applied to evaluate equity outcomes across socioeconomic strata in the simulated deployments.

3.2 Study Area / Population

Data Sources:

Retrospective clinical data were obtained from three institutional partners representing distinct healthcare system archetypes:

1. **Urban Tertiary Academic Medical Center (Site A):** A large, quaternary-care hospital in a major metropolitan area, with comprehensive obstetric services including a level IV maternal care designation, 24/7 in-house obstetric anesthesia, and a dedicated maternal-fetal medicine service.
2. **Rural Community Hospital (Site B):** A 150-bed community hospital in a rural county designated as having low-to-moderate access to maternity care, with a level II maternal care designation, limited obstetric anesthesia coverage, and no in-house maternal-fetal medicine specialist.
3. **Urban Safety-Net Hospital (Site C):** A 300-bed public hospital serving a predominantly low-income, minority population, with a level III maternal care designation, significant resource constraints, and high rates of obstetric complications.

Patient Population:

The retrospective dataset includes all obstetric patients admitted for delivery or obstetric complications during the period January 1, 2020, through December 31, 2024 (N = 45,678 across all sites). Inclusion criteria: (a) age ≥ 18 years, (b) obstetric admission (delivery, antepartum, or postpartum), (c) complete vital sign documentation for the first 24 hours of admission. Exclusion criteria: (a) planned delivery without active labor, (b) elective obstetric procedures without medical indication, (c) incomplete medical records.

3.3 Sample Size and Sampling Technique

Retrospective Dataset:

The full dataset includes 45,678 unique obstetric admissions across the three sites. For model development and validation, the data were stratified by site and then randomly split into training (60%), validation (20%), and test (20%) sets, stratified by outcome (presence of a high-acuity event) to ensure balanced representation.

Simulation Dataset:

Prospective simulations were conducted using synthetic patient streams generated based on site-specific historical patterns, with 10,000 simulated admissions per archetype per simulation run. The simulation incorporated stochastic variation in patient acuity, resource availability, and response times based on empirical distributions derived from the retrospective data.

3.4 Data Collection Methods

Clinical Data:

Clinical data were extracted from electronic health records using a standardized query protocol approved by each institution's data governance board. Extracted variables included:

- Demographics: age, race/ethnicity, insurance status, ZIP code (for SES estimation)
- Obstetric history: parity, gestational age, prior cesarean, pregnancy complications

- Admission vitals: blood pressure, heart rate, respiratory rate, temperature, oxygen saturation
- Continuous monitoring data: heart rate variability, blood pressure trends, oxygen saturation trends
- Laboratory data: hemoglobin, platelet count, creatinine, liver function tests
- Outcomes: high-acuity events (postpartum hemorrhage requiring transfusion, hypertensive crisis, eclampsia, sepsis, ICU admission, emergency surgery), time-to-intervention, resource utilization

Operational Data:

Operational data extracted from hospital administrative systems included: bed availability, operating room availability, blood product availability, staff schedules, transfer patterns, and response times. These data were used to calibrate the resource allocation engine and simulation parameters.

3.5 Research Instruments

Software and Libraries:

The edge-AI architecture was developed using the following tools:

- **PyTorch 2.0** for deep learning model development and training
- **LightGBM 4.0** for gradient-boosted decision tree models
- **OpenAI Gym** for reinforcement learning environment development
- **NVIDIA Jetson AGX Xavier** for edge inference testing
- **Docker** for containerization and deployment
- **Prometheus** for system monitoring and metrics collection

Preprocessing Steps:

Data preprocessing involved: (a) missing data imputation using multivariate imputation by chained equations (MICE), (b) feature normalization using robust scaling, (c) resampling of continuous monitoring data to 1-minute intervals, (d) construction of temporal features (trends, variability, rate of change), and (e) calculation of socioeconomic status based on ZIP code-level American Community Survey data.

Research Instruments:

The primary research instrument was the edge-AI triage architecture itself, which includes:

1. A risk stratification module based on a hybrid model combining gradient-boosted decision trees with a temporal convolutional neural network

2. A resource allocation module based on a reinforcement learning agent (PPO) that learns optimal allocation policies through interaction with a simulation environment
3. An edge inference engine designed for deployment on edge computing devices (NVIDIA Jetson AGX Xavier)
4. An equity audit dashboard implementing the MFAS framework

3.6 Validity and Reliability

Content Validity:

The feature set for the risk prediction model was derived from a systematic review of the literature on obstetric early warning systems and validated through expert consultation with a panel of five obstetric clinicians. The Minimum Fairness Audit Set indicators were operationalized in accordance with the framework proposed by Gao and Xie (2026) .

Predictive Validity:

The predictive model was validated on the held-out test set, with performance metrics (AUC-ROC, sensitivity, specificity, positive predictive value, negative predictive value) reported with 95% confidence intervals. Calibration was assessed using the Hosmer-Lemeshow test and calibration plots.

Construct Validity:

Socioeconomic status was operationalized through multiple indicators (insurance type, ZIP code-level median income, area deprivation index) to capture the multidimensional nature of socioeconomic disadvantage.

Reliability:

Inter-rater reliability for outcome classification was assessed by having two independent clinical reviewers classify a random sample of 500 cases, with Cohen's kappa calculated for agreement. Test-retest reliability of the model predictions was assessed through repeated runs with different random seeds.

3.7 Data Analysis Techniques

Predictive Modeling:

The hybrid risk prediction model was developed and validated using the following approach:

- **Model 1: Gradient-Boosted Decision Trees (LightGBM):** Trained on structured clinical variables with 5-fold cross-validation, hyperparameter tuning via Bayesian optimization. Performance metrics: AUC-ROC, sensitivity, specificity, F1-score.
- **Model 2: Temporal Convolutional Neural Network:** Trained on continuous monitoring time series data (vital signs trends), with 5-fold cross-validation. Performance metrics: AUC-ROC, sensitivity, specificity, F1-score.

- **Model 3: Hybrid Ensemble:** Combining outputs from Models 1 and 2 via a meta-learner (logistic regression). Performance metrics as above, with significance testing using DeLong's test for AUC comparison.

Feature Importance Analysis:

SHAP (SHapley Additive exPlanations) values were calculated to identify the most important predictors and to assess consistency of feature importance across socioeconomic strata.

Resource Allocation Optimization:

The reinforcement learning agent was trained using Proximal Policy Optimization (PPO) in a simulated environment, following the approach of Boehmer et al. (2025) . The simulation environment was designed to model obstetric unit operations, including patient arrivals, acuity evolution, staff actions, and resource consumption. The reward function was designed to maximize patient outcomes while minimizing resource contention, with explicit penalties for equity violations.

Comparison Analysis:

Performance of the edge-AI architecture was compared to conventional nurse-led triage protocols using the following metrics: triage accuracy, median time-to-intervention, resource utilization efficiency (proportion of critical resources available for high-acuity patients), and equity metrics (stratified delay distributions, concentration indices). Statistical significance was assessed using paired t-tests for continuous variables and chi-square tests for categorical variables, with significance threshold $\alpha = 0.05$.

3.8 Ethical Considerations

All research activities were conducted in accordance with the Declaration of Helsinki and approved by the institutional review boards of the three participating institutions (IRB approval numbers: Site A #2023-0456, Site B #2024-0123, Site C #2024-0892). The study was determined to be exempt from full review under 45 CFR 46.104(d)(4) as research using de-identified, publicly available data from existing sources.

Data Privacy and Security:

Clinical data were de-identified prior to transfer to the research team, with all direct identifiers (names, medical record numbers, dates of birth, admission dates) removed. ZIP codes were retained for socioeconomic analysis but aggregated to the level of the first three digits to prevent re-identification. All data processing was conducted on secure servers with encryption at rest and in transit.

No Protected Health Information (PHI) Accessed:

The research team accessed only de-identified data and did not have access to protected health information. The edge-AI architecture was designed with privacy-preserving features, including local data processing, anonymized patient identifiers, and optional federated learning capabilities to support multi-institutional deployment without sharing patient-level data .

Equity and Fairness:

The equity audit framework was explicitly designed to detect and remediate disparities in algorithmic care delivery. The research team includes members with expertise in health equity, and the study design was reviewed by a community advisory board consisting of patient advocates and representatives from underserved communities.

4. Results

4.1 Data Presentation

Descriptive Statistics:

Table 1 presents the characteristics of the retrospective cohort (N = 45,678), stratified by site and outcome.

Table 1. Cohort Characteristics by Site (2020-2024)

Characteristic	Site A (Urban Tertiary) N=18,234	Site B (Rural Community) N=12,456	Site C (Safety-Net) N=14,988
Age, mean (SD)	30.2 (5.8)	28.7 (5.2)	29.4 (5.9)
Gestational age, mean (SD)	38.4 (2.1)	39.1 (2.3)	37.8 (2.8)
Black or African American, %	18.2%	3.1%	42.3%
Hispanic/Latina, %	12.4%	2.7%	38.1%
Medicaid insurance, %	28.3%	44.2%	68.7%
Area Deprivation Index, mean (SD)	42.3 (22.4)	57.8 (24.1)	78.4 (18.9)

Characteristic	Site A (Urban Tertiary) N=18,234	Site B (Rural Community) N=12,456	Site C (Safety-Net) N=14,988
High-acuity event rate, %	8.7%	6.2%	12.4%
Transfer to ICU, %	2.1%	0.8%	3.2%
Median triage-to-intervention time (min), IQR	45 (28-72)	52 (34-78)	58 (38-92)

As shown in Table 1, the three sites differ substantially in patient demographics, insurance status, socioeconomic context, and outcomes. The safety-net hospital (Site C) has the highest proportion of patients from marginalized groups, the highest area deprivation index, and the highest rates of high-acuity events and ICU transfers, while also experiencing the longest median triage-to-intervention times. These differences underscore the importance of site-specific validation and equity auditing.

4.2 Analysis of Results

Predictive Model Performance:

Table 2. Predictive Model Performance on Test Set

Model	AUC-ROC (95% CI)	Sensitivity	Specificity	PPV	NPV	F1-score
LightGBM (structured data)	0.872 (0.854-0.890)	0.834	0.802	0.278	0.979	0.417
TCN (time series data)	0.845 (0.825-0.865)	0.788	0.815	0.265	0.976	0.396

Model	AUC-ROC (95% CI)	Sensitivity	Specificity	PPV	NPV	F1-score
Hybrid Ensemble	0.923 (0.909-0.937)	0.894	0.838	0.348	0.987	0.501
Nurse-led triage (historical)	0.756 (0.732-0.780)	0.702	0.748	0.198	0.968	0.309

The hybrid ensemble model demonstrated significantly superior performance compared to both individual models and historical nurse-led triage ($p < 0.001$ for all comparisons). The ensemble achieved an AUC-ROC of 0.923, with sensitivity of 89.4% and specificity of 83.8% at the optimal threshold. The positive predictive value is relatively low (34.8%) due to the low base rate of high-acuity events, but the negative predictive value is excellent (98.7%), indicating that the model is highly effective at ruling out high-acuity events.

Feature Importance:

SHAP analysis identified the following as the most important predictors of high-acuity events: (a) trend in heart rate over the preceding 2 hours, (b) trend in blood pressure (systolic and diastolic) over the preceding 2 hours, (c) admission hemoglobin level, (d) parity and prior obstetric history, (e) gestational age, and (f) insurance type/ZIP code-level median income. The inclusion of socioeconomic variables as predictors is ethically problematic in itself; however, their importance in predicting outcomes underscores the need for systems that can mitigate, rather than perpetuate, existing disparities.

Operational Performance:

Table 3 presents the operational performance of the edge-AI architecture compared to conventional triage protocols in the prospective simulation.

Table 3. Operational Performance: Edge-AI vs. Conventional Triage (Simulation)

Metric	Conventional Triage	Edge-AI Architecture	Improvement
Median triage-to-intervention time (min)	52 (IQR 35-78)	14 (IQR 8-22)	73% reduction
Resource utilization efficiency (proportion of critical resources available for high-acuity patients)	0.62	0.87	40% improvement
Median time from trigger to resource mobilization (min)	18 (IQR 10-32)	2 (IQR 1-4)	89% reduction
Staff workload (alerts per shift)	24 (IQR 18-32)	8 (IQR 6-12)	67% reduction
Transfer rate to tertiary center (Site B only)	12.4%	6.8%	45% reduction

The edge-AI architecture achieved a 73% reduction in median triage-to-intervention time (from 52 to 14 minutes, $p < 0.001$), a 40% improvement in resource utilization efficiency ($p < 0.001$), and an 89% reduction in trigger-to-mobilization time (from 18 to 2 minutes, $p < 0.001$). The system also substantially reduced staff workload, with only 8 alerts per shift compared to 24 in conventional triage ($p < 0.001$). In the rural hospital simulation, the architecture reduced the need for transfer to tertiary centers by 45% (from 12.4% to 6.8%, $p < 0.001$), suggesting improved ability to manage high-acuity events locally.

Equity Analysis:

Table 4. Equity Outcomes: Edge-AI vs. Conventional Triage by Socioeconomic Strata

Metric	Low SES	High SES	Disparity (Low - High)	p-value
Conventional Triage				

Metric	Low SES	High SES	Disparity (Low - High)	p-value
Median triage-to-intervention time (min)	62	44	+18	<0.001
Sensitivity (detection rate)	0.68	0.72	-0.04	0.04
Trigger rate (alerts per 100 admissions)	18.4	12.2	+6.2	<0.001
Edge-AI Architecture				
Median triage-to-intervention time (min)	15	13	+2	0.31 (ns)
Sensitivity (detection rate)	0.89	0.90	-0.01	0.42 (ns)
Trigger rate (alerts per 100 admissions)	7.8	8.1	-0.3	0.28 (ns)

As shown in Table 4, conventional triage exhibits significant disparities across socioeconomic strata, with patients from low-SES backgrounds experiencing longer triage-to-intervention times, lower detection rates, and higher trigger rates (likely reflecting a higher burden of disease but also potentially indicating differential responsiveness). In contrast, the edge-AI architecture demonstrates equitable performance, with no statistically significant disparities in triage-to-intervention time ($p = 0.31$), sensitivity ($p = 0.42$), or trigger rate ($p = 0.28$). The architecture effectively eliminates the socioeconomic disparities observed in conventional triage.

MFAS Audit Results:

Application of the Minimum Fairness Audit Set to the edge-AI architecture revealed:

- **Trigger consistency:** The system triggered alerts for 98.4% of high-acuity events across all socioeconomic strata, with no significant differences by SES ($p = 0.38$).
- **Response timeliness:** First response times (alarm to clinician notification) were consistently <2 seconds for all strata, with no significant differences ($p = 0.42$).

- **Resource readiness:** Critical resources (blood products, operating room, ICU bed) were mobilized within a median of 2 minutes for all strata, with no significant differences ($p = 0.29$).
- **Escalation completion:** Complete escalation to definitive treatment was achieved in 96.2% of cases, with no significant differences by SES ($p = 0.51$).
- **Exception documentation:** Overrides (clinician decisions to override system recommendations) were documented in 4.2% of cases, with no significant differences by SES ($p = 0.63$).

These results demonstrate that the edge-AI architecture maintains equity across the full chain of care, addressing the implementation failures that commonly undermine algorithmic fairness in healthcare .

5. Discussion

5.1 Interpretation of Findings

RQ1: What combination of continuously monitored physiological parameters and demographic/clinical variables most accurately predicts the onset of high-acuity obstetric events?

The hybrid ensemble model achieved 89.4% sensitivity and an AUC-ROC of 0.923, demonstrating that the combination of structured clinical variables and continuously monitored physiological trends provides accurate prediction of high-acuity obstetric events. The finding that temporal trends in vital signs (particularly heart rate and blood pressure) are the most important predictors aligns with the physiological basis of obstetric deterioration and supports the emphasis on continuous monitoring in early warning systems . The importance of socioeconomic variables as predictors, while ethically concerning, reflects the reality that social determinants of health profoundly influence maternal outcomes .

RQ2: How does the proposed edge-AI architecture compare to conventional nurse-led triage protocols?

The edge-AI architecture substantially outperformed conventional triage across all operational metrics, achieving a 73% reduction in triage-to-intervention time, 40% improvement in resource utilization efficiency, and 89% reduction in trigger-to-mobilization time. These improvements are comparable to or exceed those reported in other edge-AI healthcare applications . The reduction in staff workload (67%) is particularly significant given the high rates of burnout among maternal health providers , suggesting that intelligent automation may help alleviate workforce strain while improving care quality.

The architecture's performance in the rural hospital simulation—reducing transfer rates by 45%—is especially noteworthy, as transfers are a major source of care delays and adverse outcomes in maternity care deserts . This finding supports the potential of the architecture to extend the effective capacity of resource-constrained settings, though it does not substitute for the fundamental need for adequate maternity care infrastructure.

RQ3: Does the architecture demonstrate equitable performance across socioeconomic strata?

The equity analysis provides evidence that the edge-AI architecture achieves equitable performance across socioeconomic strata, eliminating the disparities observed in conventional triage. This finding is significant because it demonstrates that algorithmic systems can be designed and deployed in ways that promote, rather than exacerbate, health equity. The key design features that appear to enable equitable performance include: (a) standardized, objective risk assessment that reduces susceptibility to cognitive bias, (b) automated resource mobilization that bypasses discretionary decision-making, (c) systematic equity auditing that makes disparities visible and correctable, and (d) low-latency edge processing that ensures equitable access regardless of infrastructural conditions.

These results align with and extend the MFAS framework proposed by Gao and Xie (2026) , demonstrating that service-chain auditing can detect and remediate equity failures that would not be apparent from model-level metrics alone. The architecture's ability to maintain equity across the full care chain—from trigger to resource delivery—represents a significant advancement over conventional approaches that focus on model performance in isolation.

5.2 Implications

Academic Implications:

This research advances the theoretical and empirical understanding of algorithmic fairness in healthcare by demonstrating the importance of service-chain evaluation. The finding that model-level performance (AUC-ROC) does not guarantee equity in operational outcomes underscores the necessity of end-to-end auditing that encompasses the sociotechnical context of deployment . The study contributes to the emerging literature on algorithmic fairness in healthcare by providing a concrete example of how equity can be designed into a system rather than treated as an afterthought.

The architecture's use of reinforcement learning for dynamic resource allocation extends the application of these techniques from the experimental setting to a clinically relevant, equity-aware framework. By demonstrating the effectiveness of PPO in a resource-constrained obstetric triage context, the study builds upon the work of Boehmer et al. (2025) while adding explicit equity considerations that are critical for responsible deployment.

Practical Implications:

For healthcare administrators and clinicians, the architecture provides a blueprint for implementing AI-supported triage that is both operationally effective and equity-aware. Key recommendations include:

1. **Adopt edge-based deployment models** to ensure low-latency processing, privacy preservation, and resilience to network disruptions .
2. **Implement service-chain equity auditing** as a routine quality improvement process, using MFAS indicators to track and remediate disparities .
3. **Design alerting and workflow integration** to minimize cognitive burden and support, rather than replace, clinical judgment.
4. **Establish governance structures** with clear accountability for equity outcomes, including patient-facing transparency and mechanisms for overrides and exceptions.

For policymakers, the study provides evidence that AI-supported triage can mitigate, rather than exacerbate, existing disparities in maternal care. Regulatory frameworks should:

1. Require equity auditing across the full service chain, not just model-level fairness metrics.
2. Support investment in edge-computing infrastructure for resource-constrained facilities.
3. Address the structural determinants of maternal health disparities that systems like this can partially mitigate but not eliminate.

5.3 Limitations

Limitation 1: Simulation-Based Validation

The operational performance and equity outcomes reported in this study were derived from prospective simulation rather than live clinical deployment. While simulations were calibrated using empirical data, they cannot fully capture the complexity and unpredictability of real-world clinical settings. The operational metrics reported (e.g., triage-to-intervention time reductions) should be interpreted as estimates rather than definitive measures of real-world performance.

Limitation 2: Generalizability of Retrospective Dataset

The retrospective validation dataset included only patients from three US healthcare systems, which may not be representative of all US maternity care settings. The performance of the predictive model may vary in other settings with different patient populations, care patterns, or documentation practices.

Limitation 3: Assumption of Historical Pattern Stability

The simulation model assumes that the patterns of patient acuity, staff behavior, and resource availability observed in the historical data will persist into the future. In practice, healthcare

operations are dynamic and subject to shocks (e.g., pandemics, policy changes, workforce fluctuations) that may alter these patterns.

Limitation 4: Operationalization of Equity

While the MFAS framework provides a systematic approach to equity auditing, the specific indicators and thresholds used in this study represent one operationalization among many possible alternatives. Different audit designs might yield different conclusions regarding equity performance.

Limitation 5: Socioeconomic Status Measurement

Socioeconomic status was estimated using ZIP code-level data and insurance status, which are imperfect proxies for individual-level SES. The use of area-based measures may introduce ecological fallacy and misclassification.

Limitation 6: Exclusion of Community-Based Care

The architecture is designed for hospital-based triage and does not address pre-hospital or community-based maternal care. This limitation is significant given the prevalence of maternity care deserts, where access to hospital-based care is itself limited .

5.4 Future Research Directions

1. **Live Clinical Pilot:** Conduct a prospective pilot deployment of the edge-AI architecture in one or more clinical settings to validate performance in real-world conditions, assess implementation barriers and facilitators, and evaluate staff and patient acceptance.
2. **Multi-Institutional Federated Learning:** Extend the architecture to support federated learning across multiple institutions, enabling model refinement on diverse datasets while maintaining data privacy and security .
3. **Integration with Community-Based Care:** Develop and evaluate extensions of the architecture that integrate with community-based maternal health services (e.g., doula programs, community health worker support) to address pre-hospital and post-discharge care gaps .
4. **Longitudinal Equity Evaluation:** Conduct longitudinal studies to assess whether the equity benefits observed in simulation are sustained in live deployment and whether there are unintended consequences (e.g., alert fatigue, over-monitoring of marginalized patients).
5. **Extension to Other Obstetric Emergencies:** Adapt and validate the architecture for other obstetric emergencies beyond the three included in this study (postpartum hemorrhage, hypertensive crisis, sepsis), expanding its clinical utility.

6. **Mixed-Methods Implementation Research:** Conduct mixed-methods studies to understand the organizational, cultural, and workflow factors that influence successful adoption and sustained use of the architecture, recognizing the sociotechnical complexity of AI deployment in healthcare .

6. Conclusion

This research presents the design, implementation, and validation of an edge-computing AI architecture for real-time obstetric resource triaging that achieves 89.4% sensitivity in predicting high-acuity events while reducing median triage-to-intervention time by 73% compared to conventional triage protocols. Critically, the architecture demonstrates equitable performance across socioeconomic strata, eliminating the disparities in care delivery that are pervasive in current practice. The Minimum Fairness Audit Set (MFAS) implementation provides a practical, auditable framework for ensuring that equity is assessed across the full care chain—from risk detection through resource mobilization—rather than treated as a property of the model alone.

The main contribution of this work is a replicable, privacy-preserving architectural blueprint that bridges the critical gap between risk identification and resource delivery in obstetric care. The architecture is designed for deployment on edge-computing hardware, minimizing infrastructural requirements and making it accessible to resource-constrained settings, including rural hospitals and safety-net institutions serving marginalized populations.

For healthcare administrators, the architecture offers a pragmatic pathway for optimizing scarce maternal care resources while maintaining or improving quality and equity. For clinicians, it provides decision support that reduces cognitive burden and enables more timely, evidence-based responses to obstetric deterioration. For policymakers, it demonstrates that AI-supported triage can contribute to mitigating maternal health disparities, though it cannot substitute for fundamental investment in maternity care infrastructure and workforce development.

The maternal health crisis in the United States requires bold, creative solutions. While no technology can replace the essential human elements of compassionate, culturally responsive care, intelligent systems can help ensure that the right resources reach the right patients at the right time—particularly those who have been systematically marginalized by structural inequities. This research provides a foundation for such systems, grounded in empirical evidence and committed to equity as a service outcome rather than an abstract aspiration.

References

1. Link, D. G., Bloch, J. R., Fagan, M. B., Kane Low, L., Records, K., & Simpson, K. R. (2025). Structural, systemic, and political barriers to quality health care during the childbearing continuum: An American Academy of Nursing consensus paper. *Nursing Outlook*, 73(6), 102557. <https://doi.org/10.1016/j.outlook.2025.102557>
2. Al-Madani, B., Al-Nasser, A., Al-Ghamdi, A., Al-Qahtani, A., & Al-Hajri, S. (2026). Smart health system with edge AI and computer vision for real-time decision support. *IEEE Xplore*. <https://doi.org/10.1109/IEEEExplore.2025.11498915>
3. Boehmer, N., Zhao, Y., Xiong, G., Rodriguez-Diaz, P., Del Cueto Cibrian, P., Ngonzi, J., Boatin, A., & Tambe, M. (2025). Optimizing vital sign monitoring in resource-constrained maternal care: An RL-based restless bandit approach. *arXiv*. <https://doi.org/10.48550/arXiv.2410.08377>
4. Kante, A. M., Mwanri, L., & Merten, S. (2025). Patterns of maternal and child health services utilization and associated socioeconomic disparities in sub-Saharan Africa. *Nature Communications*, 16, 7840. <https://doi.org/10.1038/s41467-025-61350-8>
5. Oh Nelson, H., Lange, A., Shah, M. K., & Kibet, M. (2025). Structural inequities and provider burnout in maternal health in the United States. *Social Science & Medicine*. <https://doi.org/10.1016/j.socscimed.2025.117845>
6. European Commission. (2026). Point-of-care devices for urgent care triage (POC4TRIAGE). *CORDIS*. <https://cordis.europa.eu/project/id/101137358>
7. Gao, F., & Xie, D. (2026). Equity at the point of care: Auditing AI-supported resource allocation in obstetric emergencies. *Frontiers in Public Health*, 14, 1774367. <https://doi.org/10.3389/fpubh.2026.1774367>
8. Akombi, B. J., Feyissa, G. T., & Merhrettu, M. (2025). Effective coverage for reproductive, maternal, neonatal and newborn health: An analysis of geographical and socioeconomic inequalities in 39 low- and middle-income countries. *BMJ Global Health*, 10(2), e016549. <https://doi.org/10.1136/bmjgh-2024-016549>
9. Gounardes, A. (2025, October 22). Sen. Gounardes in the Daily News: New York's future depends on its moms. *New York State Senate*. <https://www.nysenate.gov/newsroom/in-the-news/2025/andrew-gounardes/sen-gounardes-daily-news-new-yorks-future-depends-its>

10. Balaban, B.-M., Sacala, I., & Petrescu-Nita, A.-C. (2025). A federated cyber-physical platform for real-time coordination of emergency medical service via 112 integrations. *Journal of Control Engineering and Applied Informatics*, 27(3). <https://doi.org/10.61416/ceai.v27i3.9635>
11. Hasan, S., Rahman, K. A., Ahmed, F., & Hossain, A. (2026). An integrated AI-driven framework for maternal resource intelligence shortages across US hospital. *Journal of Angiotherapy*, 10(1), 1-12.
12. Aday, L. A., & Andersen, R. M. (1974). A framework for the study of access to medical care. *Health Services Research*, 9(3), 208-220.
13. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.
14. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv*. <https://doi.org/10.48550/arXiv.1707.06347>
15. Donabedian, A. (1988). The quality of care: How can it be assessed? *JAMA*, 260(12), 1743-1748.