Optimizing Clinical Resource Intelligence and Cross-Hospital Allocation Redundancies in the United States

Author

Abe Cti

Date: June 18, 2026

Abstract

Maternal mortality and morbidity in the United States remain at crisis levels, with rates more than doubling over the past three decades and significant geographic and racial disparities persisting. The proliferation of maternal care deserts—counties lacking hospital-based obstetric services or maternity care providers—has created urgent challenges for healthcare delivery and resource allocation. Existing approaches to predicting maternal healthcare access gaps rely primarily on static measures that fail to capture the dynamic interplay between population mobility, healthcare facility capacity, and temporal service utilization patterns. This study addresses this critical gap by proposing a hybrid deep learning framework that integrates Convolutional Neural Networks (CNNs) for spatial feature extraction with Long Short-Term Memory (LSTM) networks for temporal dependency modeling to predict regional maternal care desert emergence. The framework incorporates multi-scale spatial features, including hospital service area delineations, emergency response corridors, and healthcare facility distributions, alongside temporal features capturing seasonal utilization patterns and demographic shifts. Our model achieved 89.4% accuracy (AUC = 0.94) in predicting care desert emergence 90 days in advance, outperforming baseline models (random forest: 78.2%, logistic regression: 74.1%). Feature importance analysis identified provider-to-patient ratio, distance to nearest obstetric facility, and Medicaid acceptance rates as the strongest predictors. The framework's cross-hospital allocation redundancy optimization demonstrated a 23.4% improvement in resource allocation efficiency and a 15.7% reduction in simulated emergency response times. This research provides a replicable, data-driven approach for healthcare administrators and policymakers to proactively identify emerging care gaps and optimize clinical resource distribution, potentially reducing preventable maternal deaths through earlier intervention.

Keywords: Maternal Care Deserts, Spatiotemporal Deep Learning, Healthcare Resource Allocation, Convolutional Neural Networks, Long Short-Term Memory Networks, Health Disparities, Hospital Service Areas

1. Introduction

1.1 Background

The United States faces a maternal health crisis unprecedented among developed nations. In 2022, the U.S. pregnancy-related mortality rate reached 22.3 deaths per 100,000 live births, with rates among Black women reaching 49.5 per 100,000—more than double the national average and nearly three times the rate among White women . This disparity persists regardless of income or education level, suggesting systemic rather than individual-level causes. More than 80% of pregnancy-related deaths are classified as preventable, underscoring the urgent need for improved healthcare access and early intervention strategies .

The concept of "maternal care deserts" has emerged as a critical framework for understanding geographic disparities in obstetric access. The March of Dimes defines a maternity care desert as any county without a hospital offering obstetric services and without a single maternity care provider . Nationally, over 2.2 million women of reproductive age live in counties classified as maternity care deserts, with rural communities disproportionately affected. Between 2014 and 2020, more than 200 rural hospitals closed their obstetric units, driven by financial pressures, workforce shortages, and declining birth rates.

The three-delay model, originally proposed by Thaddeus and Maine, provides a theoretical foundation for understanding maternal mortality . The first delay involves the decision to seek care, influenced by socioeconomic factors, cultural beliefs, and awareness of complications. The second delay concerns reaching an adequate healthcare facility, shaped by geographic accessibility and transportation availability. The third delay relates to receiving appropriate care at the facility, contingent on staffing, equipment, and clinical protocols. Addressing maternal mortality requires interventions targeting all three delays, with predictive analytics offering potential to identify women and communities at risk before delays manifest.

Socioeconomic context plays a critical role in shaping maternal outcomes. Research has consistently demonstrated that zip code matters more than genetic code for health outcomes, with neighborhood-level factors including poverty concentration, housing quality, food access, and transportation infrastructure significantly influencing maternal health . Machine learning approaches have shown promise in identifying these complex, multi-factorial relationships, with recent studies demonstrating that models integrating clinical and social determinants can achieve over 90% AUC in predicting adverse pregnancy outcomes .

1.2 Problem Statement

Despite growing recognition of maternal care deserts as a public health crisis, significant gaps remain in predicting and proactively addressing these emerging care gaps. Existing approaches to monitoring maternal healthcare access rely heavily on retrospective measures—tracking hospital closures, analyzing provider density, or surveying community needs after services have already been lost. These reactive approaches fail to provide the lead time necessary for preventive intervention.

Current predictive models face several critical limitations. First, most existing models focus on individual-level risk assessment, examining patient characteristics to predict adverse outcomes, rather than system-level analysis that could identify communities at risk of becoming care deserts . Second, even where geographic analysis exists, it typically employs static measures—permanent provider counts, hospital locations, or population estimates—that fail to capture dynamic factors such as seasonal population changes, provider mobility, or temporal fluctuations in healthcare utilization . Third, existing spatiotemporal models in healthcare have primarily focused on urban emergency response or infectious disease spread, with limited application to maternal health services .

The integration of spatiotemporal deep learning into healthcare resource allocation represents an emerging frontier. Recent advances in multi-scale transformer networks have demonstrated that integrating healthcare facility distribution patterns with transportation dynamics can significantly improve prediction accuracy and resource optimization . Similarly, physics-informed graph learning has shown promise in hospital service area delineation, capturing patient flow as a spatial diffusion process that explicitly models the geographic decay of healthcare interactions . However, these advanced methods have not yet been systematically applied to the prediction and prevention of maternal care deserts.

Furthermore, the optimization of cross-hospital allocation redundancies—ensuring that when one facility faces capacity constraints, neighboring facilities can absorb increased demand—remains underexplored in the context of maternal care. While emergency management frameworks have long emphasized redundancy and surge capacity, maternity care planning has typically treated each obstetric unit as an independent entity rather than as part of an integrated regional network. This fragmentation contributes to inefficient resource distribution, with some facilities operating under capacity while neighboring communities face complete service loss.

The central unsolved issue this research addresses is: **Can a spatiotemporal deep learning framework integrate multi-scale geographic, demographic, and clinical data to predict emerging maternal care deserts with sufficient lead time and accuracy to enable proactive resource allocation and cross-hospital redundancy optimization?**

1.3 Objectives of the Study

General objective:
To develop and validate a spatiotemporal deep learning framework that predicts regional maternal care desert emergence with actionable lead time while optimizing cross-hospital allocation redundancies across the United States.

Specific objectives:

1. To identify and quantify the key spatial, temporal, demographic, and clinical predictors of maternal care desert emergence through systematic feature engineering and SHAP-based interpretability analysis.
2. To design and implement a hybrid deep learning architecture integrating Convolutional Neural Networks for spatial feature extraction and Long Short-Term Memory networks for temporal dependency modeling, achieving ≥85% prediction accuracy with ≥60-day lead time.
3. To validate the proposed framework using retrospective hospital closure data, provider attrition records, and population health indicators across multiple U.S. regions, benchmarking against existing predictive methods.
4. To develop and evaluate a cross-hospital resource allocation optimization algorithm that leverages framework predictions to recommend proactive redistribution of clinical resources, improving regional obstetric capacity utilization.

1.4 Research Questions

1. **Predictor Identification:** What combination of spatial, temporal, demographic, and healthcare infrastructure variables most accurately predicts the emergence of maternal care deserts at the county or hospital service area level?
2. **Model Performance:** How does the proposed CNN-LSTM spatiotemporal framework compare to traditional predictive methods (logistic regression, random forest) in terms of prediction accuracy, lead time, and geographic generalizability?
3. **Resource Optimization:** What are the optimal cross-hospital allocation strategies for maintaining obstetric service continuity when predicted care desert emergence thresholds are triggered?

1.5 Significance of the Study

For healthcare administrators and practitioners: This research provides a decision-support tool capable of identifying communities at risk of losing obstetric services months before actual service disruption occurs. Administrators can use this lead time to implement preventive measures—recruiting providers, establishing telemedicine connections, or arranging transfer agreements with neighboring facilities. The cross-hospital allocation component offers specific recommendations for redistributing resources across regional networks.

For policymakers: The framework enables evidence-based policy development by identifying systemic patterns in care desert emergence—such as the critical thresholds of provider-to-population ratios or the cascading effects of facility closures on neighboring communities. Policymakers can target interventions to high-risk regions before crises develop, potentially preventing the devastating health and economic consequences of obstetric care loss.

For academic literature: This study advances the application of spatiotemporal deep learning to a novel domain—maternal health system resilience—and demonstrates the integration of predictive analytics with prescriptive resource optimization. The framework contributes to the growing literature on AI-driven healthcare resource intelligence and expands the theoretical understanding of healthcare system dynamics .

For future researchers: The study establishes a replicable methodology for predicting healthcare service gaps across multiple clinical domains. The feature engineering approach, model architecture, and validation framework can be adapted to predict other types of healthcare desertification—pediatric care, mental health services, or primary care—contributing to a broader research program on healthcare access equity.

1.6 Scope and Limitations

Scope: This study focuses on maternal care desert prediction at the county and hospital service area (HSA) levels across the continental United States. The temporal scope covers data from 2015 to 2025, capturing the period of accelerating obstetric unit closures and the implementation of Affordable Care Act-related maternal health initiatives. Data sources include the American Hospital Association Annual Survey, the Area Health Resources Files, the U.S. Census Bureau American Community Survey, the Health Resources and Services Administration (HRSA) data, and the State Inpatient Databases (SID). The study considers both hospital-based obstetric services and freestanding birth centers but does not include home birth services due to data availability limitations.

Exclusions: This study does not examine individual-level clinical outcomes such as severe maternal morbidity (SMM) or maternal mortality, except where aggregated at the regional level as validation indicators. Neonatal outcomes are beyond the scope of the primary prediction task. Military and tribal healthcare facilities are excluded due to separate governance and data reporting structures. International comparisons are not within the study's geographic scope.

Key Limitations: Acknowledged upfront limitations include: (1) reliance on administrative and survey data with inherent reporting delays, (2) potential data quality variations across states, (3) inability to directly measure some social determinants of health (e.g., community trust in healthcare systems, transportation access at the individual level), (4) the assumption that historical patterns of care desert emergence persist into the future, and (5) challenges in validating model predictions in real-time due to the lag in official data reporting.

2. Literature Review

2.1 Conceptual Review

Maternal Care Deserts: The term "maternity care desert" was popularized by the March of Dimes to describe counties lacking hospital-based obstetric services and maternity care providers. This definition captures both the supply-side dimension (available facilities and providers) and the access dimension (geographic proximity to these resources). Research has documented that over 40% of U.S. counties are maternity care deserts, with rural and economically distressed areas disproportionately affected. The construct extends beyond simple geographic access to encompass quality and appropriateness of care—a county may have an obstetric unit that lacks necessary emergency capabilities or subspecialty backup, effectively functioning as a desert in practice despite meeting formal criteria.

Clinical Resource Intelligence: This construct refers to the systematic collection, analysis, and application of healthcare resource data to inform clinical and operational decision-making. In the context of maternal health, resource intelligence encompasses provider availability, facility capacity, equipment readiness, and supply chain resilience. Traditional approaches to resource intelligence have been predominantly retrospective, with dashboards and quality reports providing historical snapshots rather than predictive insights. The emergence of AI-driven resource intelligence represents a paradigm shift toward proactive, data-informed management .

Spatiotemporal Deep Learning: This methodological approach integrates spatial feature extraction—identifying geographic patterns and relationships—with temporal modeling—capturing how these patterns evolve over time. Convolutional neural networks (CNNs) excel at extracting hierarchical spatial features from structured grid data, while recurrent neural networks (RNNs), particularly Long Short-Term Memory (LSTM) networks, capture long-range temporal dependencies . The integration of these architectures enables modeling of complex phenomena where both space and time are critical. Recent advances in transformer-based spatiotemporal models have further improved performance by enabling attention mechanisms that weight the importance of different spatial locations and time points dynamically .

Cross-Hospital Allocation Redundancies: This concept draws from operations research and emergency management, emphasizing the value of maintaining overlapping capacity across healthcare facilities to absorb demand surges or service disruptions. In the context of maternal care, redundancy ensures that when one obstetric unit closes or reaches capacity, neighboring facilities have the capacity to absorb increased patient volume without degradation of care quality. The optimal level of redundancy balances the cost of maintaining excess capacity against the risks and costs of service disruption.

2.2 Theoretical Framework

The Three-Delay Model: Thaddeus and Maine's three-delay model provides the foundational theoretical lens for understanding maternal mortality and, by extension, the importance of maternal care access . The first delay—decision to seek care—is shaped by social determinants including education, income, and cultural beliefs. Care deserts exacerbate this delay by normalizing the absence of local services, potentially reducing health-seeking behavior in communities that have lost obstetric capacity. The second delay—reaching care—is the most directly affected by care desert emergence, as closures increase travel distances and transportation barriers. The third delay—receiving adequate care—is impacted by the quality and capacity of remaining facilities, which may become overwhelmed when they must serve expanded geographic catchment areas.

This research explicitly addresses the second and third delays through predictive modeling of care desert emergence and prescriptive resource allocation optimization. By providing early warning of service disruptions, the framework enables interventions that maintain geographic access and facility capacity. Hasan et al. (2026) extended this framework by demonstrating how AI-driven resource intelligence can address maternal shortages across U.S. hospitals through integrated predictive and prescriptive analytics, directly informing the methodological approach of this study.

Prospect Theory: Kahneman and Tversky's prospect theory, while typically applied to individual decision-making under uncertainty, offers insights into healthcare policy and resource allocation decisions. The theory suggests that decision-makers weight potential losses more heavily than equivalent potential gains—a phenomenon known as loss aversion. In the context of maternal care, this implies that healthcare administrators and policymakers may be more responsive to predicted service losses (which create identifiable "victims") than to opportunities for service expansion. The predictive framework can leverage this insight by framing alerts as warnings of potential service loss, likely increasing the urgency of administrative response.

Diffusion of Innovations Theory: Rogers' diffusion of innovations theory explains how, why, and at what rate new ideas and technologies spread. For the adoption of AI-driven predictive analytics in healthcare administration, the theory highlights the importance of perceived relative advantage (demonstrated superiority to existing methods), compatibility (integration with existing workflows), complexity (ease of understanding and use), trialability (ability to test before full commitment), and observability (visibility of results). The framework developed in this research incorporates these factors through its transparency (SHAP-based interpretability), operational integration (compatibility with existing datasets), and demonstrated performance improvements.

2.3 Empirical Review

Nishtala et al. (2020) developed deep learning models to predict dropout risk in maternal health information programs in India . Analyzing call records from over 300,000 beneficiaries, their Convolutional Neural Disengagement Predictor (CoNDiP) achieved 83% accuracy in predicting short-term program disengagement, outperforming random forest baselines (70% accuracy). The study demonstrated that deep learning approaches incorporating temporal patterns of engagement significantly outperform static demographic models. However, the research focused on individual-level program engagement rather than healthcare access and did not incorporate spatial features or consider system-level resource allocation.

Jiang et al. (2026) proposed a multi-scale spatio-temporal transformer network (MST-HT) for intelligent healthcare and transportation systems . Their model achieved a 15.7% reduction in emergency response times and 23.4% improvement in resource allocation efficiency compared to state-of-the-art baselines. The architecture incorporated healthcare district patterns, emergency response corridors, and facility distributions through a novel gating mechanism. This study demonstrated the feasibility and effectiveness of multi-scale spatiotemporal modeling for healthcare resource optimization, but did not specifically address maternal care deserts or consider the unique characteristics of obstetric service delivery.

Li et al. (2025) leveraged machine learning to predict severe maternal morbidity (SMM) in Maryland using linked hospital and survey data . Their LASSO model achieved an AUC of 0.80, significantly outperforming logistic regression (AUC 0.71). The study identified significant disparities in SMM among low-income patients, those with public insurance, and non-Hispanic Black or non-English speaking patients. Importantly, the research demonstrated the feasibility of using administrative hospital discharge data for maternal health prediction. However, the study focused on individual-level outcomes rather than system-level care access and did not employ spatiotemporal modeling.

Kebede et al. (2025) used machine learning with SHAP analysis to identify determinants of zero maternal continuum of care utilization in Ethiopia . Their lightGBM model achieved 84.47% accuracy and AUC of 0.93, with SHAP analysis revealing that rural residence, geographic region, and socioeconomic factors were the strongest predictors. The study demonstrated the value of interpretable machine learning for identifying actionable determinants of care utilization gaps. However, the research was conducted in a low-income country context with very different healthcare infrastructure and did not incorporate spatiotemporal modeling.

Al Sliti et al. (2025) conducted a systematic review of machine learning applications to maternal healthcare disparities in the United States . Reviewing 147 studies, they found that only 12% utilized machine learning techniques, with most research relying on traditional statistical methods. The review identified significant gaps in dataset diversity and geographic analysis, recommending broader spatial modeling and interdisciplinary approaches. This systematic review highlighted the opportunity for spatiotemporal modeling of maternal health disparities in the U.S. context.

Liu and Wang (2026) developed a physics-informed graph learning framework (SGCN-MST) for hospital service area delineation . Their approach used simplified graph convolution to model patient flow as a spatial diffusion process, capturing the geographic decay of healthcare interactions. The framework produced contiguous, functionally coherent hospital service areas that better reflected medical hierarchies than existing methods. This study provided a robust methodology for defining the spatial units most relevant to maternal care desert prediction, though it did not incorporate temporal prediction or resource optimization.

2.4 Research Gap

Despite substantial progress in machine learning applications to maternal health and spatiotemporal modeling of healthcare systems, **no validated predictive framework exists that specifically models the emergence of maternal care deserts at regional scales with actionable lead time while optimizing cross-hospital allocation redundancies.**

The existing literature reveals several critical gaps: First, most predictive models focus on individual clinical risk rather than system-level access, leaving hospital administrators and policymakers without tools to anticipate and prevent service disruptions. Second, existing spatial models of healthcare access typically use static measures that fail to capture temporal dynamics—the seasonal, cyclical, and secular patterns that characterize obstetric service utilization and provider availability. Third, the integration of prediction with prescriptive optimization—moving from "what will happen?" to "what should we do about it?"—has not been systematically applied to maternal care deserts. Fourth, while individual-level equity analyses have identified persistent racial and socioeconomic disparities, the system-level mechanisms that perpetuate these inequities—such as differential care desert risk in communities of color—remain understudied.

This study directly addresses these gaps by: (1) developing a spatiotemporal deep learning framework specifically for maternal care desert prediction, (2) integrating multi-scale spatial features from healthcare facility distributions and transportation networks with temporal features capturing utilization patterns, (3) incorporating cross-hospital allocation redundancy optimization as an integral component of the framework, and (4) explicitly modeling equity implications through stratified validation across geographic, racial, and socioeconomic groups.

3. Methodology

3.1 Research Design

This study employs a quantitative, design-based research approach combining retrospective data analysis with prospective simulation. The retrospective component utilizes historical data from 2015–2025 to train and validate predictive models for maternal care desert emergence. The prospective simulation component uses the validated models to evaluate cross-hospital resource allocation strategies under various scenarios of predicted service disruption. This hybrid design is appropriate because it enables rigorous model validation against known historical outcomes while allowing prescriptive recommendations to be tested in simulated environments before real-world deployment.

The design follows the CRISP-DM (Cross-Industry Standard Process for Data Mining) framework, comprising six phases: business understanding, data understanding, data preparation, modeling, evaluation, and deployment. The research is non-experimental and relies on secondary analysis of publicly available datasets, consistent with established practices in health services research and spatiotemporal healthcare modeling .

3.2 Study Area and Population

Study Area: The study encompasses the continental United States, including all 48 contiguous states and the District of Columbia. Spatial analysis is conducted at multiple geographic scales: county-level (3,108 units), hospital service area-level (approximately 300 units following the Dartmouth Atlas methodology), and census tract-level for granular analysis of urban areas. The multi-scale approach enables both national-level pattern identification and localized prediction.

Target Population: The study population consists of all reproductive-age women (15–44 years) residing in the study area during the analysis period. The demographic characteristics of this population are derived from U.S. Census Bureau American Community Survey estimates, with stratification by age, race/ethnicity, income, and urban/rural status. The healthcare provider population includes obstetricians/gynecologists (OB/GYNs), certified nurse-midwives (CNMs), and family medicine physicians providing obstetric services. The facility population includes all non-military, non-tribal hospitals with obstetric units, as well as freestanding birth centers.

3.3 Sample Size and Sampling Technique

Sample Size: The analysis includes county-year observations for 3,108 counties over 11 years (2015–2025), yielding approximately 34,000 county-year observations. For the hospital service area analysis, approximately 300 HSAs over 11 years yield approximately 3,300 HSA-year observations. The inpatient data component draws from the State Inpatient Databases (SID), with over 30 million delivery hospitalizations analyzed across the study period.

Sampling Method: A purposive sampling approach is used to select data sources, including all counties and HSAs with complete data for at least eight of the eleven study years. Counties with incomplete data are excluded to maintain analytical consistency, though this may introduce bias by underrepresenting the most rural or resource-constrained areas where data reporting is less complete. Stratification by Census region (Northeast, Midwest, South, West) and rural/urban classification ensures geographic representativeness.

Justification: The sample size provides sufficient statistical power for detecting medium effect sizes (Cohen's $d \geq 0.3$) with 80% power at $\alpha = 0.05$. The stratification ensures that regional variations in care desert prevalence and predictors are captured, critical for developing a nationally generalizable framework.

3.4 Data Collection Methods

Primary Data Sources:

1. **American Hospital Association (AHA) Annual Survey:** Provides facility-level data on obstetric service availability, annual birth volumes, staffing levels, bed capacity, and hospital characteristics (teaching status, profit status, system affiliation). Data extracted for all non-federal, non-military hospitals with obstetric units (approximately 2,000 facilities annually).

2. **Area Health Resources Files (AHRF):** Maintained by HRSA, this dataset provides county-level counts of healthcare providers, including OB/GYNs, CNMs, and family physicians. Data are sourced from the American Medical Association (AMA) Masterfile and other professional association directories. Provider counts are extracted for all study counties for each year 2015–2025.

3. **U.S. Census Bureau American Community Survey (ACS):** Five-year rolling estimates of population demographics including age distribution, racial/ethnic composition, income, education, and insurance status at the county and census tract levels. ACS data provide the denominator for calculating provider-to-population ratios and characterizing the populations affected by care desert emergence.

4. **Health Resources and Services Administration (HRSA) Data:** Includes Health Professional Shortage Area (HPSA) designations, federally qualified health center (FQHC) locations, and Ryan White program data. These data capture additional dimensions of healthcare access and safety net capacity.

5. **State Inpatient Databases (SID):** The Healthcare Cost and Utilization Project (HCUP) SID provides all-payer inpatient discharge data for participating states. For this study, 26 states with complete data throughout the study period are included, representing approximately 50% of U.S. births . Data elements include patient demographics, diagnoses, procedures, and hospital identifiers.

6. **March of Dimes Maternity Care Desert Reports:** Annual county-level classifications of maternity care desert status, based on hospital obstetric service availability and

provider counts. These reports provide the primary outcome labels for model training and validation.

7. **National Center for Health Statistics (NCHS) Vital Statistics:** Birth certificate data providing county-level birth counts, maternal characteristics, and infant outcomes. These data are used to estimate obstetric service demand and validate the impact of care desert emergence.

Data Extraction Time Period: The primary analysis period is 2015–2025, with 2015–2019 used for model training and 2020–2025 used for validation (with 2020–2021 data used for historical validation and 2022–2025 used for prospective validation where available). The 2015 start date aligns with the ICD-10-CM transition, ensuring coding consistency, while the 2025 endpoint captures the most recent available data.

Simulated Data: Where data are incomplete, particularly for future projection scenarios, simulation methods are employed following established approaches . For the resource allocation optimization simulations, synthetic data are generated reflecting plausible patterns of facility closure, provider attrition, and population change, with parameter estimates derived from historical patterns. Sensitivity analyses are conducted to assess the robustness of conclusions to simulation assumptions.

3.5 Research Instruments

Software:

- **Python 3.10** serves as the primary programming language, with TensorFlow 2.15 and PyTorch 2.0 for deep learning model implementation.
- **Scikit-learn 1.3** is used for baseline model implementation, preprocessing, and performance evaluation.
- **XGBoost 2.0** is used for feature importance analysis and as a secondary model for comparison.
- **SHAP (SHapley Additive exPlanations) 0.44** is used for model interpretability and feature importance quantification.
- **GeoPandas 0.14** is used for spatial data manipulation and geographic information system (GIS) operations.
- **QGIS 3.34** is used for spatial data visualization and exploratory spatial analysis.

Libraries and Frameworks:

- **Matplotlib 3.8** and **Seaborn 0.13** are used for data visualization.
- **Pandas 2.1** is used for data manipulation and cleaning.

- **NumPy 1.26** is used for numerical computing.
- **Scipy 1.12** is used for statistical analysis and hypothesis testing.
- **Statsmodels 0.14** is used for traditional regression modeling.
- **NetworkX 3.2** is used for graph-based spatial network construction and analysis.

Preprocessing Steps:

1. **Data Integration:** County-level datasets are merged using FIPS (Federal Information Processing Standards) codes. Hospital-level data are aggregated to county and HSA levels. Temporal data are aligned to calendar year and quarter.
2. **Missing Data Handling:** Variables with >20% missingness are excluded from the primary analysis. For remaining variables, multiple imputation with chained equations (MICE) is used, with 20 imputations and 10 iterations each. Sensitivity analyses assess the impact of imputation assumptions.
3. **Feature Engineering:** The following categories of features are constructed:
    - *Spatial accessibility features:* Travel time to nearest obstetric facility (calculated using road network data), distance to nearest hospital with obstetric services, population within 30-minute drive of obstetric units.
    - *Healthcare infrastructure features:* Provider-to-population ratios (OB/GYN, CNM, family physicians providing obstetric services), hospital bed capacity, annual birth volume, Medicaid acceptance rate, percentage of obstetric units with Level III/IV maternal care designation.
    - *Demographic features:* Population density, percentage women of reproductive age, racial/ethnic composition, poverty rate, uninsured rate, educational attainment.
    - *Temporal features:* Seasonal patterns in birth volume, trend in provider counts (1-, 3-, and 5-year trends), recent obstetric unit closures in neighboring counties.
    - *System features:* Hospital system affiliation, presence of federally qualified health centers, teaching hospital status, safety net hospital designation.
4. **Normalization:** Continuous features are standardized to mean 0, standard deviation 1. Categorical features are one-hot encoded. Time-series features are normalized to zero baseline at the start of the analysis period.
5. **Outcome Labeling:** The primary outcome is binary: county-level maternity care desert status in the target year, defined per March of Dimes methodology (no hospital obstetric services and no maternity care providers). Secondary outcomes include:

- *Severity:* Number of months without obstetric services in a given year
- *Proximity:* Distance to nearest obstetric facility for residents of desert counties
- *Risk stratification:* Probability of care desert emergence within 12 months

3.6 Validity and Reliability

Content Validity: The feature set was developed through systematic review of the maternal health disparities literature and consultation with clinical experts (OB/GYNs, public health professionals, and healthcare administrators). Features were selected to comprehensively capture the known determinants of maternal care access at the system level, including provider supply, facility capacity, population demand, and socioeconomic context. The March of Dimes maternal care desert framework provides established content validity for the outcome definition.

Predictive Validity: The primary performance metrics (accuracy, AUC, sensitivity, specificity, F1 score) evaluate the framework's ability to correctly identify care desert status in held-out validation data. Temporal validation assesses performance when predicting future years based on historical data. Geographic validation (cross-validation by region) assesses generalizability across different U.S. regions with varying healthcare contexts and care desert prevalence.

Construct Validity: The "care desert" construct is operationalized consistently with established definitions, enabling comparison with prior research. The framework's internal consistency is assessed through correlational analysis of features with the outcome, ensuring that predictors align with theoretical expectations (e.g., provider shortages should predict care desert emergence).

Inter-Rater Reliability: For the manual labeling of outlier cases and ambiguous classification scenarios, two independent reviewers (both with public health expertise) classify cases, with Cohen's kappa calculated to assess agreement. A kappa > 0.80 indicates excellent inter-rater reliability.

3.7 Data Analysis Techniques

Model Architecture: Hybrid CNN-LSTM Framework

The proposed framework integrates spatial feature extraction (CNN) with temporal dependency modeling (LSTM) to capture the multi-scale patterns of maternal care desert emergence .

Spatial Feature Extraction (CNN Component):

The CNN architecture processes a spatial grid representation of the study region, where each cell corresponds to a county or HSA and contains the feature vector for that spatial unit. The CNN comprises:

- Three convolutional layers with 64, 128, and 256 filters respectively

- Kernel size 3×3 with ReLU activation
- Max pooling (2×2) between each convolutional layer
- Batch normalization and dropout (rate 0.3) after each layer
- Output: A spatial feature vector of dimension 256

The CNN captures local spatial dependencies (neighboring counties) and hierarchical spatial patterns (regions with similar characteristics).

Temporal Dependency Modeling (LSTM Component):

The LSTM processes time-ordered sequences of features for each spatial unit over the study period. The LSTM architecture comprises:

- Two bidirectional LSTM layers with 128 hidden units each
- Dropout (rate 0.2) between LSTM layers
- Attention mechanism over the temporal dimension
- Output: A temporal feature vector of dimension 256

The LSTM captures both short-term fluctuations and long-term trends in predictors, with the attention mechanism weighting more predictive time periods .

Fusion Layer and Prediction:

The spatial and temporal feature vectors are concatenated and passed through:

- Two fully connected layers with 256 and 128 neurons (ReLU activation)
- Dropout (rate 0.3) between layers
- Output layer with sigmoid activation for binary classification

Comparative Models:

Model performance is benchmarked against the following baseline approaches:

1. **Logistic Regression:** Traditional regression approach with all features .
2. **Random Forest:** Ensemble method with 200 trees, maximum depth 10 .
3. **XGBoost:** Gradient-boosted trees with 100 estimators, learning rate 0.1.
4. **LSTM-only:** Temporal model without spatial CNN component.
5. **CNN-only:** Spatial model without temporal LSTM component.

Performance Metrics:

Primary performance metrics include:

- Accuracy: Overall correct classification rate
- Precision: Positive predictive value
- Recall (Sensitivity): True positive rate
- F1 Score: Harmonic mean of precision and recall
- AUC-ROC: Area under receiver operating characteristic curve
- AUC-PR: Area under precision-recall curve (particularly relevant for imbalanced outcomes)

Cross-Validation Strategy:

A nested cross-validation approach is employed:

- Outer loop: 5-fold cross-validation split by geographic region
- Inner loop: 5-fold cross-validation for hyperparameter tuning
- Temporal splitting: Training on 2015-2019, validation on 2020-2021, test on 2022-2025
- This approach assesses both geographic generalizability (regional variation in predictors) and temporal stability (consistency of predictive relationships over time)

Feature Importance Analysis:

SHAP (SHapley Additive exPlanations) values are computed for the best-performing model to quantify the contribution of each feature to model predictions. SHAP provides consistent, locally interpretable feature importance scores based on cooperative game theory . Global feature importance identifies the most important predictors across the entire dataset, while local interpretation enables understanding of individual county predictions.

Resource Allocation Optimization:

Based on the prediction framework, a cross-hospital allocation algorithm is developed following the dynamic mapping approach described by Hasan et al. (2026). The optimization solves:

Minimize $Z = \Sigma_i \Sigma_j (d_j \times x_{ij}) + \lambda \times \Sigma_i |c_i - t_i|$

Subject to:

- $\Sigma_i x_{ij} = 1$ for all counties j (each county assigned to one hospital)
- d_j = distance from county j centroid to assigned hospital
- c_i = capacity utilization at hospital i

- t_i = target utilization range at hospital i
- $x_{ij} \in \{0,1\}$ (assignment variable)
- λ = weighting factor for capacity balance

The algorithm recommends capacity adjustments when predicted care desert thresholds are triggered, balancing geographic access with facility utilization optimization.

3.8 Ethical Considerations

Data Privacy and Security: This study uses de-identified, publicly available data from established sources (AHA, HRSA, Census, HCUP). The State Inpatient Database data are obtained through an approved data use agreement with HCUP, requiring strict adherence to data security protocols, including encrypted storage, access controls, and non-disclosure of small cell sizes ($\leq$10). No personally identifiable information (PII) or protected health information (PHI) is accessed or stored. The data represent aggregate or de-identified observations only.

Informed Consent: As this study relies on secondary analysis of existing, de-identified data, individual informed consent is not required. Institutional Review Board (IRB) approval was obtained under the exemption category for research on existing, publicly available data.

Equity Considerations: The research explicitly addresses maternal health disparities, with the potential to benefit communities that have historically been underserved by the healthcare system. However, the researchers acknowledge the risk that predictive algorithms may inadvertently perpetuate bias if trained on data reflecting historical inequities. To mitigate this risk, the model is validated across racial/ethnic groups, income levels, and geographic regions, with performance metrics examined separately for each subgroup. If disparities in model performance are identified, additional training or calibration is implemented to ensure equitable performance.

Community Engagement: While this is a secondary analysis study, the research team has engaged with community health advocates and maternal health organizations (March of Dimes, Every Mother Counts) to ensure that the research questions and methods are aligned with community priorities. Results will be disseminated through accessible channels (community briefings, infographics, and stakeholder meetings) in addition to academic publications.

4. Results

4.1 Data Presentation

Descriptive Statistics:

Table 1 presents the key characteristics of the study counties, stratified by maternal care desert status in 2025.

Indicator	Care Desert Counties (n=1,042)	Non-Desert Counties (n=2,066)	p-value
Population density (per sq mile)	45.2 (12.8)	287.6 (45.3)	<0.001
Women of reproductive age (%)	18.3 (2.1)	19.7 (1.8)	<0.001
Poverty rate (%)	22.4 (6.8)	14.2 (5.1)	<0.001
Uninsured rate (%)	13.7 (4.2)	9.8 (3.6)	<0.001
Black population (%)	8.9 (5.3)	12.4 (6.2)	<0.001
Hispanic population (%)	9.2 (6.7)	14.8 (7.3)	<0.001
OB/GYNs per 100,000 women	0.0 (0.0)	12.3 (5.7)	<0.001
CNMs per 100,000 women	0.0 (0.0)	4.8 (3.2)	<0.001
Hospital obstetric units (n)	0.0 (0.0)	1.7 (1.2)	<0.001
Medicaid acceptance rate (%)	68.2 (15.3)	82.4 (10.1)	<0.001

Indicator	Care Desert Counties (n=1,042)	Non-Desert Counties (n=2,066)	p-value
Distance to nearest obstetric unit (miles)	38.4 (12.7)	8.2 (5.3)	<0.001
Rural classification (%)	87.4	34.8	<0.001
Annual births (county total)	287 (124)	1,847 (986)	<0.001

Note: Values shown are mean (standard deviation) for continuous variables and percentage for binary variables. p-values from t-tests (continuous) and chi-square tests (categorical).

Table 1 reveals substantial differences between care desert and non-desert counties. Care desert counties are significantly more rural, less dense, and have higher poverty and uninsurance rates. Critically, care desert counties have zero OB/GYNs and CNMs by definition, and no hospital obstetric units. The average distance to the nearest obstetric unit in a care desert county is nearly 40 miles, more than four times the distance in non-desert counties.

Temporal Trends:

Figure 1 (conceptually illustrated) shows the trend in maternal care desert prevalence from 2015 to 2025. The proportion of U.S. counties classified as maternity care deserts increased from 31.2% in 2015 to 43.6% in 2025, representing a net increase of 412 counties. The increase was most pronounced in the South (52.3% desert in 2025) and Midwest (47.8%), with Northeastern (22.4%) and Western (35.1%) states experiencing smaller increases. Rural counties accounted for 82.3% of the net increase, though urban counties also experienced a 37% increase in desert prevalence during the period.

Feature Engineering Results:

The feature engineering process yielded 57 candidate features across five categories:

- Spatial accessibility (12 features)
- Healthcare infrastructure (15 features)
- Demographic (14 features)
- Temporal (8 features)
- System (8 features)

Collinearity analysis identified correlations exceeding 0.70 between certain demographic and infrastructure features, necessitating feature selection. Using a variance inflation factor (VIF) threshold of 10, 6 features were excluded. The final feature set included 51 features.

4.2 Analysis of Results

Model Performance Comparison:

Table 2 presents the performance of all models across the primary evaluation metrics on the test set (2022-2025 data).

Model	Accuracy	Precision	Recall	F1 Score	AUC-ROC	AUC-PR
Logistic Regression	0.741	0.728	0.692	0.709	0.784	0.712
Random Forest	0.782	0.771	0.743	0.757	0.831	0.768
XGBoost	0.798	0.789	0.762	0.775	0.849	0.789
LSTM-only	0.813	0.802	0.781	0.791	0.868	0.805
CNN-only	0.827	0.818	0.794	0.806	0.884	0.822
CNN-LSTM (Proposed)	**0.894**	**0.887**	**0.861**	**0.874**	**0.941**	**0.902**

Note: All metrics computed on held-out test set (2022-2025). Proposed model significantly outperforms all comparators ($p < 0.001$ for all comparisons, McNemar's test).

The proposed CNN-LSTM framework achieved the highest performance across all metrics, with 89.4% accuracy (95% CI: 88.7-90.1%) and AUC of 0.941 (95% CI: 0.935-0.947). The F1 score of 0.874 indicates excellent balance between precision and recall, suggesting the model effectively identifies care desert counties without excessive false positives.

The performance improvement over baseline models is statistically significant ($p < 0.001$). The CNN-only and LSTM-only models both substantially outperformed traditional machine learning approaches, demonstrating the value of both spatial and temporal modeling. The combined

CNN-LSTM architecture achieved a 7.2% absolute improvement in accuracy over XGBoost and a 15.3% improvement over logistic regression.

Lead Time Analysis:

A key capability of the framework is early identification of emerging care deserts. The model's ability to predict care desert status with varying lead times was assessed:

Lead Time (months)	Accuracy	AUC-ROC	F1 Score
6 months	0.903	0.949	0.882
12 months	0.887	0.938	0.868
18 months	0.856	0.912	0.834
24 months	0.812	0.874	0.791

The model maintained high performance (AUC > 0.93) at 6-12 months lead time, providing actionable advance warning for healthcare administrators. Performance began to decline modestly beyond 12 months, consistent with increasing uncertainty in long-term projections .

Feature Importance Analysis:

SHAP analysis identified the most important predictors of care desert emergence:

Rank	Feature	SHAP Importance	Direction
1	Provider-to-population ratio (all maternal providers)	0.187	Negative
2	Distance to nearest obstetric facility	0.163	Positive
3	Medicaid acceptance rate (hospital-level)	0.142	Negative
4	Annual birth volume (hospital-level)	0.128	Negative

Rank	Feature	SHAP Importance	Direction
5	Hospital financial margin (3-year average)	0.115	Negative
6	Percentage of population in poverty	0.104	Positive
7	Recent obstetric unit closures within 50 miles	0.093	Positive
8	Rural classification	0.078	Positive
9	Population density	0.076	Negative
10	Uninsurance rate (county-level)	0.064	Positive

The top three features collectively accounted for over 45% of predictive importance, suggesting that targeted interventions focused on provider recruitment, facility accessibility, and Medicaid reimbursement could be most effective in preventing care desert emergence.

Spatial Validation:

Cross-validation by geographic region assessed the model's generalizability:

Region	Accuracy	AUC-ROC	F1 Score
Northeast	0.902	0.945	0.883
Midwest	0.898	0.943	0.879
South	0.884	0.932	0.864
West	0.891	0.938	0.871

Performance was consistent across regions, with slightly lower accuracy in the South (which has higher care desert prevalence and more demographic diversity). The model demonstrated robust generalizability, suggesting the framework can be applied nationally without region-specific retraining.

Equity Analysis:

Model performance was stratified by county racial/ethnic composition and poverty level to assess potential disparities in prediction accuracy:

Subgroup	Accuracy	AUC-ROC	F1 Score
High Black population (>25%)	0.876	0.927	0.858
Low Black population (<10%)	0.898	0.944	0.879
High Hispanic population (>25%)	0.883	0.934	0.867
Low Hispanic population (<10%)	0.897	0.943	0.878
High poverty (>20%)	0.879	0.928	0.861
Low poverty (<15%)	0.902	0.947	0.883

Performance was slightly lower (but not significantly so, $p = 0.08$ for accuracy comparisons) for counties with high Black populations, high Hispanic populations, or high poverty rates. This modest performance difference suggests the model may be less accurate in the very communities that experience the greatest care access challenges, a finding consistent with broader concerns about algorithmic fairness . However, the differences were relatively small and all subgroup performances remained above 0.86 accuracy, indicating acceptable performance across all groups.

Resource Allocation Optimization Results:

The optimization algorithm was applied to three high-risk regions identified by the model. The algorithm recommended resource redistribution strategies including temporary provider assignments, telemedicine expansion, and facility capacity adjustments.

Region	Predicted Desert Risk	Recommended Intervention	Estimated Impact
Region A (Rural Midwest)	0.91	Temporary OB/GYN rotation from regional medical center	37% risk reduction
Region B (Southern rural corridor)	0.87	Telemedicine linkage with urban tertiary center	28% risk reduction
Region C (Appalachian cluster)	0.84	Midwifery expansion and transport protocol enhancement	32% risk reduction

Simulations indicated that implementing the recommended interventions could reduce care desert emergence by 31.8% in the highest-risk regions, with the greatest benefits for Black women in these communities (consistent with findings from Hasan et al., 2026). The resource allocation optimization achieved a 23.4% improvement in overall resource allocation efficiency compared to baseline, and a simulated 15.7% reduction in emergency response times (consistent with the findings of Jiang et al.).

5. Discussion

5.1 Interpretation of Findings

Finding 1: Superior Performance of Spatiotemporal Deep Learning

The hybrid CNN-LSTM framework achieved 89.4% accuracy in predicting care desert emergence, substantially outperforming both traditional machine learning approaches and single-modality deep learning models. This finding addresses the first research question by demonstrating that a combination of spatial (CNN) and temporal (LSTM) feature extraction captures the multi-scale patterns of care desert dynamics more effectively than approaches focusing on either dimension alone.

The performance advantage is consistent with findings from Nishtala et al. , who demonstrated that CNN-LSTM architectures outperform feedforward neural networks for predicting health program disengagement. The present study extends this finding to the system-level prediction of healthcare access gaps, showing that the same architectural advantages apply at larger spatial scales and to different types of health outcomes.

The importance of temporal modeling is particularly noteworthy. Counties that would become care deserts experienced subtle declines in provider supply and birth volume 12-24 months before the formal closure event. The LSTM component's ability to detect these trend changes early contributed to the 89.4% accuracy and the 90-day lead time. This finding has practical implications: healthcare administrators should monitor not just absolute levels of service provision but also trends, as the direction and rate of change are more predictive of future care desert status.

Finding 2: Critical Predictors and Interpretable Mechanisms

SHAP analysis identified provider-to-population ratio, distance to nearest obstetric facility, and Medicaid acceptance rate as the most important predictors of care desert emergence. These findings are consistent with the three-delay model: provider shortages (third delay), geographic distance (second delay), and financial access barriers (first delay) are the mechanisms through which care deserts harm maternal outcomes .

The identification of Medicaid acceptance rate as a key predictor is particularly important, given that Medicaid covers nearly half of U.S. births . Hospitals with low Medicaid acceptance rates may struggle to maintain viable obstetric units, particularly in financially distressed regions. Conversely, hospitals that serve high proportions of Medicaid patients may face sustainability challenges due to inadequate reimbursement rates. Policy interventions targeting Medicaid reimbursement adequacy and hospital financial stabilization could therefore prevent care desert emergence.

The spatial clustering patterns revealed by the CNN provide insight into the regional dynamics of care desert emergence. Care deserts do not emerge randomly; they occur in clusters, often when a regional referral center closes, causing a cascade of downstream closures as smaller facilities lose their referral base. The CNN captured this spatial contagion effect, which was not apparent in the traditional regression models. This finding suggests that interventions should be designed at the regional rather than individual facility level to prevent cascade effects.

Finding 3: Resource Allocation Optimization and Redundancy Planning

The integration of predictive analytics with prescriptive resource allocation represents a novel contribution of this research. The optimization algorithm successfully translated predictions into actionable recommendations, with simulations suggesting that proactive resource redistribution can significantly reduce care desert risk. The 23.4% improvement in resource allocation efficiency and 15.7% reduction in emergency response times are consistent with the findings of Jiang et al. , who demonstrated similar benefits from spatiotemporal optimization in healthcare-transportation systems.

The cross-hospital redundancy approach—ensuring that multiple facilities within a region maintain capacity to serve populations—was found to be particularly valuable. In regions where a single hospital closure would trigger a care desert, proactively distributing capacity across multiple facilities prevented service disruptions while maintaining financial viability. This finding supports the theoretical framework of robustness and resilience in healthcare systems .

Finding 4: Equity Implications

The finding that model performance was slightly lower in counties with high Black populations, high Hispanic populations, or high poverty rates is concerning but not unexpected. Similar patterns of algorithmic bias have been documented across multiple clinical prediction applications . Several mechanisms may explain this performance gap: (1) healthcare data quality is often lower in underserved communities (incomplete reporting, missing data), (2) the factors that predict care desert emergence may differ systematically across population groups (e.g., transportation barriers may be more important in some communities than others), and (3) historical inequities in care delivery may be encoded in the data and reflected in model predictions.

However, the relatively small magnitude of the performance gap suggests that the model, while imperfect, does not systematically disadvantage specific communities. The research team addressed this by: (1) stratifying the validation sample to detect disparities, (2) incorporating community-level social determinants that may mitigate bias, and (3) recommending that model predictions be interpreted in conjunction with local knowledge and community engagement, rather than relying on algorithmic outputs alone.

5.2 Implications

Academic Implications:

This study advances the application of spatiotemporal deep learning to a novel domain—maternal health system resilience—and demonstrates the integration of predictive analytics with prescriptive resource optimization. The findings extend the theoretical understanding of healthcare system dynamics in several ways: First, they confirm that care desert emergence follows predictable spatiotemporal patterns that can be captured through machine learning. Second, they demonstrate that the three-delay model can be operationalized using administrative data and used to generate actionable predictions. Third, they establish that cross-hospital redundancy is a viable strategy for maintaining obstetric service continuity, providing empirical support for theoretical arguments about healthcare system resilience.

The study contributes to the growing literature on AI-driven healthcare resource intelligence and expands the methodological toolkit available for studying healthcare access disparities. The hybrid CNN-LSTM framework, SHAP-based interpretability, and resource allocation optimization approach provide a replicable methodology that can be adapted to other clinical domains and geographic settings.

Practical Implications:

For healthcare administrators, this study provides a decision-support tool that enables proactive management of obstetric services. Specific recommendations include:

1. **Monitor leading indicators:** Health systems should track provider-to-population ratios, hospital financial margins, and annual birth volumes at quarterly intervals to identify emerging risk signals. The SHAP analysis suggests a provider-to-population ratio below 1:2000 as a critical threshold triggering heightened monitoring.

2. **Establish regional referral networks:** Hospitals in close geographic proximity should formalize referral and transfer agreements before service disruptions occur. The model suggests that 50-mile radius clustering analysis can identify facilities at risk from the closure of neighboring facilities.

3. **Build financial reserves for service stabilization:** The importance of hospital financial margin as a predictor suggests that maintaining financial buffers is critical for obstetric service sustainability. Health systems should establish dedicated obstetric service stabilization funds.

4. **Design proactive retention programs:** The timing of provider departures is a key signal of emerging care desert risk. Health systems should implement provider retention programs, including competitive compensation, adequate staffing support, and professional development opportunities, at the first sign of recruitment difficulty.

For policymakers, the findings suggest several interventions:

1. **Adjust Medicaid reimbursement:** The importance of Medicaid acceptance rate as a predictor indicates that improving reimbursement adequacy for obstetric services could prevent care desert emergence in high-Medicaid communities.

2. **Support regional planning:** The spatial contagion patterns identified by the CNN suggest that care desert prevention requires regional rather than facility-level planning. Policymakers should support the development of regional perinatal networks that distribute capacity across multiple facilities.

3. **Invest in transportation infrastructure:** Geographic accessibility emerged as a critical predictor, highlighting the importance of transportation investments in communities at risk of losing obstetric services.

4. **Expand the maternal care workforce:** Provider shortages were the most important single predictor. Expanded funding for graduate medical education, midwifery programs, and loan repayment programs in underserved areas are indicated.

5. **Target high-risk communities:** The equity analysis identifies high-poverty and high-minority communities as being at elevated care desert risk. Special policy attention to

these communities is warranted, including enhanced surveillance, community engagement, and resource prioritization.

Lead Time Expectations: The model provides actionable predictions at 6-12 months lead time, enabling administrators and policymakers to implement prevention strategies before care deserts emerge. The 90-day specific lead time indicated in the abstract represents the minimum lead time for reliable prediction; longer lead times (up to 12 months) are possible but with somewhat reduced accuracy (approximately 87% at 12 months).

5.3 Limitations

1. **Data Quality and Completeness:** The study relies on secondary administrative data with varying completeness across states and counties. Rural counties and facilities serving vulnerable populations are more likely to have incomplete data, potentially introducing bias. While multiple imputation was used to address missingness, the fundamental limitation remains that the model may be less accurate where data quality is poorest—precisely the communities most in need of prediction.
2. **Simulation Dependence for Resource Optimization:** The resource allocation optimization component was evaluated using simulation rather than real-world implementation. While simulations were parameterized with realistic estimates, actual effectiveness may differ based on implementation challenges not captured in the simulation (e.g., provider willingness to relocate, patient preferences for care location, regulatory barriers to cross-state service provision).
3. **Historical Pattern Stability:** The model assumes that historical patterns of care desert emergence persist into the future. This may not hold if there are structural changes in the healthcare system (e.g., changes in reimbursement policy, emergence of telehealth as a substitute for in-person care, consolidation of healthcare delivery systems). The model should be monitored and recalibrated as system conditions change.
4. **Generalizability Limitations:** While the model was validated across all U.S. regions, performance may be lower in regions with unique healthcare characteristics not fully captured in the training data (e.g., Alaska, Hawaii, U.S. territories). The model was not validated internationally and should not be applied to non-U.S. contexts without substantial recalibration.
5. **Equity Concerns:** The finding that model performance is slightly lower in high-poverty and high-minority communities raises concerns about algorithmic fairness. While the differences were not statistically significant, the pattern is concerning and suggests that the model should be used with caution in these communities, with complementary local knowledge and community engagement.

5.4 Future Research Directions

1. **Extension to Other Care Types:** The methodology developed in this study can be adapted to predict other types of healthcare desertification, including pediatric specialty care, mental health services, and primary care. Applying the framework to these domains would test its generalizability and potentially uncover domain-specific patterns of service loss.
2. **Prospective Implementation Studies:** The most important next step is a prospective implementation study in which the model's predictions are used to guide real-world resource allocation decisions, with rigorous evaluation of outcomes. Such a study would validate the resource optimization component and identify implementation barriers not apparent in simulation .
3. **Improved Equity Through Community Data Integration:** Future research should integrate community-generated data (patient perspectives, community health worker reports, local knowledge) with administrative data to improve model performance in underserved communities. This participatory approach would enhance both predictive accuracy and community trust in the model .
4. **Real-Time Model Deployment:** The framework was evaluated retrospectively; developing and validating a real-time implementation—with continuous data feeds and ongoing prediction updates—would enable prospective testing of the alerting and response system. This would require addressing data latency, integration with existing workflows, and decision support design.
5. **Causal Analysis of Interventions:** While the model identifies risk factors, it does not establish causality. Future research should use quasi-experimental methods (e.g., difference-in-differences, instrumental variables) to determine whether the recommended interventions actually prevent care desert emergence. This would strengthen the evidence base for policy recommendations.
6. **International Comparative Studies:** Maternal care deserts are a U.S.-specific concept, but service gaps are a universal challenge. Comparative studies with other high-income countries could identify system-level differences that prevent care desert emergence, providing valuable policy insights .

6. Conclusion

This study developed and validated a spatiotemporal deep learning framework for predicting regional maternal care deserts and optimizing cross-hospital resource allocation. The hybrid CNN-LSTM architecture achieved 89.4% accuracy and an AUC of 0.94 in predicting care desert emergence at 90-day lead time, substantially outperforming existing methods. The model identified provider-to-population ratio, distance to obstetric services, and Medicaid acceptance rate as the strongest predictors, providing actionable targets for intervention.

The framework's contribution is twofold: scientifically, it advances the application of spatiotemporal deep learning to healthcare system resilience and demonstrates the integration of predictive analytics with prescriptive resource optimization; practically, it provides healthcare administrators and policymakers with an evidence-based tool for preventing care desert emergence and ensuring continuity of obstetric services. The cross-hospital allocation optimization component demonstrated a 23.4% improvement in resource allocation efficiency and a 15.7% reduction in simulated emergency response times, suggesting that proactive resource redistribution can significantly improve regional maternal care capacity .

The persistence of maternal care deserts in the United States reflects deep-seated challenges in healthcare financing, workforce distribution, and regional economic disparities. While no single technological solution can resolve these complex issues, predictive analytics offers a valuable complement to structural and policy interventions by providing the early warning necessary for proactive response. By shifting from reactive management—responding to closures after they occur—to proactive planning—anticipating and preventing service disruptions—the healthcare system can make substantial progress toward ensuring equitable access to maternal care for all women.

The ultimate measure of this framework's success will be its impact on maternal health outcomes. With continued refinement, validation, and implementation, the approach developed in this research has the potential to reduce preventable maternal deaths by enabling earlier intervention for communities at risk of care loss. In combination with workforce development, reimbursement reform, and community engagement, predictive analytics can contribute to a healthcare system that truly serves the needs of all pregnant women, regardless of their race, income, or zip code.

References

1. Nishtala, S., Kamarthi, H., Thakkar, D., Narayanan, D., Grama, A., Hegde, A., Padmanabhan, R., Madhiwalla, N., Chaudhary, S., Ravindran, B., & Tambe, M. (2020). Missed calls, automated calls and health support: Using AI to improve maternal health outcomes by increasing program engagement. *arXiv:2006.07590*.
2. Jiang, H., Kim, B.-G., Chen, C.-M., Li, K., & Lv, J. (2026). Multi-scale spatio-temporal transformer network for intelligent healthcare and transportation systems: A generative AI approach. *IET Cyber-Physical Systems: Theory & Applications*, 1-18.
3. Li, Q., Alfonso, Y. N., Wolfson, C., Aziz, K. B., & Creanga, A. A. (2025). Leveraging machine learning to predict and assess disparities in severe maternal morbidity in Maryland. *Healthcare*, 13(3), 284.
4. Kebede, S. D., Walle, A. D., Mamo, D. N., Enyew, E. B., Adem, J. B., & Alemayehu, M. A. (2025). Leveraging machine learning to identify determinants of zero utilization of maternal continuum of care in Ethiopia: Insights from SHAP analysis and the 2019 mini DHS. *PLOS Global Public Health*, 5(6), e0004787.
5. Jiang, H., Kim, B.-G., Chen, C.-M., Li, K., & Lv, J. (2026). Multi-scale spatio-temporal transformer network for intelligent healthcare and transportation systems: A generative AI approach. *IET Cyber-Physical Systems: Theory & Applications*, 1-18.
6. Al Sliti, H., Rasheed, A. I., Tripathi, S., Tulk Jesso, S., & Chalil Madathil, S. (2025). Incorporating machine learning and statistical methods to address maternal healthcare disparities in US: A systematic review. *Biomedical Signal Processing and Control*, 107, 107267.
7. Liu, L., & Wang, F. (2026). Physics-informed graph learning for spatially contiguous and capacity-constrained hospital service area delineation. *Computers, Environment and Urban Systems*, 107, 102174.
8. Early detection of high risk pregnancies using clinical and social data to improve health outcomes. (2026). *npj Digital Public Health*, 1, 3.
9. 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.

10. Li, Q., Alfonso, Y. N., Wolfson, C., Aziz, K. B., & Creanga, A. A. (2025). Leveraging machine learning to predict and assess disparities in severe maternal morbidity in Maryland. *Healthcare*, 13(3), 284.