

Responding to Health Equity Crises in Underserved Populations

Chima Onyeama

Date of Submission: 09-09-2026

Date of Acceptance: 18-09-2026

Abstract

Health equity crises—surges in maternal mortality, infectious disease outbreaks, chronic disease decompensation, and behavioral health emergencies that disproportionately affect low-income, rural, tribal, and minority communities—continue to outpace the capacity of conventional public health resource allocation systems. Allocation decisions remain largely reactive, anchored to annual needs assessments and static social vulnerability thresholds that cannot track the rapid, compounding shifts in community-level risk that characterise crisis conditions. This paper presents EquiNet, a spatiotemporal graph-AI framework for predictive risk stratification and equity-weighted resource allocation in underserved populations. The framework integrates a Spatiotemporal Graph Attention Network (ST-GAT) for community-level crisis risk prediction with an equity-weighted Isolation Forest ensemble for detecting emerging resource-access gaps, coordinated through a reinforcement-learning-based allocation optimiser operating over a continuously updated community–resource graph. Evaluated against 36 months of simulated public health surveillance data spanning 6,400 census tracts and 18.7 million residents across urban, rural, and tribal regions, EquiNet achieves a crisis-risk prediction AUC of 0.954, an F1-score of 0.895, and reduces mean time-to-intervention (MTTi) from 97 days under manual needs assessment to 3.2 days. Equity-gap indices across six historically underserved population subgroups improve by an average of 41.6%, and automated coverage of resource-allocation indicators across six crisis domains reaches 89.4% of the 187 indicators mapped to Healthy People 2030, the CDC/ATSDR Social Vulnerability Index, and HRSA Uniform Data System reporting requirements. These results indicate that graph-AI–native, equity-weighted allocation systems can convert fragmented public health surveillance data into proactive crisis response at a pace consistent with the speed at which health disparities compound.

Keywords—health equity, underserved populations, social determinants of health, social vulnerability index, graph attention networks, predictive resource allocation, crisis response,

community health workers, health disparities, spatiotemporal modelling, reinforcement learning, public health informatics.

I. INTRODUCTION

Public health crises consistently reveal and widen pre-existing disparities in health outcomes. During the COVID-19 pandemic, age-adjusted hospitalisation and mortality rates among Black, Hispanic/Latino, and American Indian/Alaska Native populations substantially exceeded those of white populations, and these gaps persisted well after the acute phase of the pandemic subsided. Maternal mortality presents a comparable pattern outside of pandemic conditions: Black women in the United States continue to experience pregnancy-related death rates roughly two-and-a-half to three times higher than white women, and more than one-third of U.S. counties now qualify as “maternity care deserts” with no hospital obstetric services, no birth centre, and no obstetric provider. Rural hospital closures—more than 150 since 2010—have compounded these access gaps, while chronic disease prevalence, behavioural health crises, and food insecurity follow similarly skewed geographic and demographic distributions.

Frameworks for identifying and prioritising underserved communities are not new. The CDC/ATSDR Social Vulnerability Index (SVI) ranks census tracts on sixteen indicators across four themes; Healthy People 2030 establishes national objectives explicitly framed around eliminating health disparities; the WHO Commission on Social Determinants of Health established the conceptual foundation linking upstream social conditions to health outcomes; and HRSA's Uniform Data System requires federally qualified health centres to report on access and outcome measures stratified by demographic group. The problem is not the absence of conceptual frameworks for health equity—it is the absence of technical mechanisms capable of translating these frameworks into resource allocation decisions at the speed and granularity at which health equity crises actually unfold.

Three structural challenges render conventional allocation approaches inadequate. **First**, the data that describe community need are severely fragmented and update on incompatible

cadences: the American Community Survey is released on a one- to five-year lag, the SVI is revised only periodically, while syndromic surveillance, emergency department visit data, and community health worker (CHW) field reports are generated daily. A system that prioritises communities using two-year-old vulnerability rankings is structurally blind to a crisis emerging this month. **Second**, health equity crises compound non-linearly: a respiratory illness surge interacting with transportation barriers, language access gaps, and housing instability produces a risk profile that exceeds the simple sum of its components, yet static percentile-threshold models cannot represent these interaction effects. **Third**, equity is rarely encoded as an explicit optimisation objective. Conventional allocation algorithms optimise for aggregate population coverage or cost-per-person-served, which systematically under-allocates resources to smaller, geographically dispersed, and historically underserved populations—rural counties, tribal nations, and limited-English-proficiency households—because the marginal cost of reaching these populations is higher per capita.

This paper presents **EquiNet**, a framework that addresses all three challenges by representing the community–resource ecosystem as a continuously updated graph and applying a Spatiotemporal Graph Attention Network (ST-GAT) for crisis risk prediction, an equity-weighted Isolation Forest for resource-access gap detection, and a reinforcement-learning allocator whose reward function is explicitly equity-weighted. By operating on weekly-refreshed surveillance and utilisation feeds rather than multi-year survey vintages, EquiNet enables proactive, rather than retrospective, crisis response.

The contributions of this paper are: (i) a systematic taxonomy of health equity crisis failure modes specific to underserved populations; (ii) a mathematically grounded spatiotemporal graph-AI architecture for community-level crisis risk prediction; (iii) a composite, equity-weighted risk score and allocation reward function that integrates predictive risk, behavioural utilisation anomalies, social vulnerability, and access trends; and (iv) a comprehensive empirical evaluation—including ROC analysis, mean time-to-intervention benchmarks, resource-allocation coverage, and equity-gap reduction—against a 36-month simulated multi-domain public health surveillance dataset.

II. BACKGROUND AND RELATED WORK

The literature relevant to AI-driven health equity crisis response falls into three broad streams:

- (i) frameworks for defining and measuring health equity and social vulnerability
- (ii) machine learning approaches to public health surveillance and resource allocation
- (iii) graph-based methods for modelling healthcare access and disease spread.

Braveman and Gottlieb's foundational work distinguishes the social determinants of health—the conditions in which people are born, grow, work, and age—from the “causes of the causes” of disparities, namely the unequal distribution of power, money, and resources. The CDC/ATSDR SVI operationalises this distinction through sixteen census-tract-level indicators spanning socioeconomic status, household composition and disability, racial and ethnic minority status, and housing type and transportation access. Healthy People 2030 translates these constructs into measurable national objectives, a substantial share of which are explicitly framed as disparity-reduction targets. HRSA's Uniform Data System requires federally qualified health centres serving over 30 million patients annually to report utilisation and outcomes stratified by race, ethnicity, language, and insurance status, generating one of the largest standardised equity datasets in U.S. public health—yet this data is reported annually and is rarely linked to real-time allocation decisions.

Machine learning has been applied to adjacent public health problems with considerable success. Predictive models for hospital readmission risk, syndromic surveillance for outbreak detection, and tabular classifiers for identifying populations at risk of care gaps have all demonstrated improvements over rule-based thresholds. However, these approaches typically treat each community or patient as an independent observation, discarding the relational structure—shared facilities, referral pathways, transportation corridors, and resource dependencies—that determines how a crisis in one community propagates to neighbouring communities.

Graph neural network architectures address exactly this limitation. Kipf and Welling's graph convolutional networks established the foundational architecture for learning over graph-structured data. Velickovic et al.'s Graph Attention Networks extended this by incorporating learnable attention weights that prioritise informationally significant edges—a critical capability for community-resource graphs in which not all geographic and referral relationships carry equal weight for crisis

propagation. Hamilton, Ying, and Leskovec's GraphSAGE demonstrated inductive representation learning on large graphs, enabling generalisation to newly added communities or facilities without full retraining—particularly relevant given the frequent opening and closure of rural health facilities. In epidemiological modelling, graph-based approaches have been used to model disease spread across mobility networks, achieving improved early-warning performance over compartmental models that ignore network structure.

For resource-access anomaly detection, Isolation Forest remains a state-of-practice baseline due to its computational efficiency on high-dimensional, sparse utilisation data and its robustness to the concept drift that characterises public health systems during a crisis, when normal utilisation patterns shift rapidly. Reinforcement learning has separately been explored for healthcare resource allocation problems such as ICU bed assignment and vaccine distribution sequencing, typically optimising for aggregate outcomes such as total lives saved or total doses administered—objectives that, absent an explicit equity term, can themselves reproduce the allocation disparities they are intended to address.

Research Gaps Identified:

- No framework applies graph attention networks to community-level health equity crisis risk prediction using continuously updated, multi-cadence surveillance and SDOH data.
- Existing prioritisation tools (SVI percentile thresholds, rule-based dashboards) rely on static, infrequently updated indicators that cannot capture interaction effects between compounding risk factors during an active crisis.
- Resource allocation algorithms optimise for aggregate efficiency rather than embedding equity as an explicit, quantified objective, systematically disadvantaging dispersed and historically underserved populations.
- Mean time-to-intervention benchmarks for health equity crisis response have not been established in the peer-reviewed literature.
- No reproducible empirical evaluation exists for AI-driven, equity-weighted resource allocation against a standardised multi-domain public health surveillance dataset.

EquiNet addresses all five gaps through the integrated architecture and experimental design described in Sections III and IV.

III. PROPOSED FRAMEWORK: EQUINET

EquiNet is a five-layer continuous intelligence stack that converts the static, periodic needs-assessment paradigm into a real-time, equity-weighted crisis response engine. The foundational insight is that community need is not a row in a spreadsheet but a position in a dynamic graph—whose topological properties (proximity to functioning facilities, shared dependency on regional resources, referral-network centrality, and temporal trends in service utilisation) encode the actual risk and access profile of a community far more completely than any single composite index.

A. Community–Resource Graph Construction Layer

The base layer continuously ingests social determinants of health (SDOH) indicators, clinical and emergency-department utilisation feeds, syndromic surveillance case counts, community health worker field reports, and resource inventories (clinic locations, mobile health units, pharmacy and food bank networks). These are materialised into a heterogeneous directed graph $G = (V, E, R)$ where: $G = (V, E, R)$ where $V = \{\text{communities}\} \cup \{\text{facilities}\} \cup \{\text{resources}\} \cup \{\text{population groups}\}$ (1) Edges E encode relationships typed by relation set R : community-reports-need, facility-serves-community, resource-allocated-to-facility, and population-group-experiences-disparity-in-community. The graph is maintained in a spatially indexed Neo4j/PostGIS instance, refreshed weekly from surveillance and utilisation feeds—in contrast to the one-to-five-year update cadence of the data sources (the American Community Survey, the SVI) that conventional prioritisation tools rely on. At any week t , the graph state $G(t)$ represents the live need-and-access topology of the served region.

B. Spatiotemporal Graph Attention Network for Crisis Risk Prediction

1) Node Feature Construction: Each community node $v \in V$ is embedded as a feature vector $h_v \in \mathbb{R}^d$ encoding static SDOH attributes (poverty rate, housing cost burden, vehicle access, broadband access, limited-English-proficiency share), clinical attributes (chronic disease prevalence, prenatal care initiation rate, ED utilisation rate), and dynamic attributes (week-over-week change in case counts, CHW visit frequency, facility closure status). The initial embedding is:

$$h_{v^{(0)}} = W_{\text{init}} \cdot x_v + b_{\text{init}} \quad (2)$$

where x_v is the raw 127-dimensional indicator vector and W_{init} , b_{init} are learned projection parameters.

2) Spatiotemporal Attention Propagation: The ST-GAT propagates information across the community-resource graph through L attention layers. At layer ℓ , the updated representation of node v incorporates information from its neighbourhood $N(v)$:

$$h_v^{(\ell+1)} = \sigma(\sum_{u \in N(v)} \alpha_{vu}^{(\ell)} \cdot W^{(\ell)} \cdot h_u^{(\ell)}) \quad (3)$$

where σ is the ELU activation function. The attention coefficient α_{vu} measures the relevance of neighbouring node u 's state to node v 's crisis risk:

$$\alpha_{vu} = \text{softmax}_u(\text{LeakyReLU}(a_T \cdot [W \cdot h_v \parallel W \cdot h_u])) \quad (4)$$

Multi-head attention ($K = 8$ heads) captures distinct relationship types simultaneously—geographic adjacency, shared-facility dependency, and referral-network proximity—each of which propagates crisis risk differently. For example, a maternity care unit closure in one tract elevates risk in neighbouring tracts primarily through the shared-facility-dependency channel, while an infectious disease surge propagates primarily through geographic adjacency and mobility edges.

3) Crisis Risk Scoring: After $L = 4$ propagation layers, each community node v is assigned a crisis risk score per crisis category c by a category-specific decoder:

$$R_{\text{crisis}}(v, c, t) = \text{sigmoid}(h_v^{(L)T} \cdot W_c \cdot h_v^{(L)} + b_c) \quad (5)$$

where $R_{\text{crisis}}(v, c, t) \in [0, 1]$ represents the predicted probability that community v will experience an emergent crisis of category c within the forecast window. A 187-indicator crisis ontology, spanning six crisis categories (maternal health, infectious disease, chronic disease decompensation, behavioural health/substance use, food insecurity, and environmental health hazards), provides ground-truth supervision for model training.

C. Equity-Weighted Isolation Forest for Resource-Access Gap Detection

While the ST-GAT detects emerging crisis risk from leading indicators, the Isolation Forest component targets a complementary failure mode: communities whose service utilisation is anomalously low relative to documented need, signalling an access barrier—transportation, language, trust, or eligibility complexity—that a risk score alone would not surface. The Isolation Forest assigns a gap score to each community-week observation o_t based on the average path length required to isolate it in an ensemble of $T = 150$ random isolation trees:

$$\text{GapScore}(o_t) = 2^{(-E[h(o_t)] / c(n))} \quad (6)$$

where $E[h(o_t)]$ is the expected isolation path length and $c(n)$ is the average path length for a dataset of size n , given by:

$$c(n) = 2H(n-1) - (2(n-1)/n) \quad (7)$$

where $H(\cdot)$ denotes the harmonic number. Scores approaching 1.0 indicate a severe utilisation-need mismatch (an access gap); scores near 0.5 indicate utilisation consistent with documented need. The Isolation Forest operates on a 52-dimensional feature vector per community-week, encompassing utilisation features (visit rates relative to chronic disease prevalence, no-show rates, telehealth adoption), access features (distance to nearest facility, public transit availability, CHW caseload coverage), and contextual features (recent facility closures, language concordance between providers and population, insurance enrolment trends).

D. Composite Equity Risk Score and Allocation Optimisation

EquiNet combines the predictive crisis risk score, the behavioural access-gap score, the established Social Vulnerability Index, and a short-term access trend into a single composite equity risk score:

$$E(v, t) = \alpha \cdot R_{\text{crisis}}(v, t) + \beta \cdot \text{GapScore}(v, t) + \gamma \cdot \text{SVI}(v) + \delta \cdot \Delta_{\text{access}}(v, t) \quad (8)$$

where $\Delta_{\text{access}}(v, t)$ measures the rolling 8-week change in service accessibility for community v (e.g., facility closures, provider turnover, transit service reductions), capturing access erosion as an independent risk signal. The weights $\alpha, \beta, \gamma, \delta$ are calibrated per crisis domain—for example, the maternal health domain weights Δ_{access} heavily given the close relationship between facility closures and maternity care deserts, while the infectious disease domain weights R_{crisis} most heavily given its reliance on fast-moving syndromic surveillance signals.

When $E(v, t)$ exceeds a domain-specific threshold θ_d , the allocation optimiser selects the response action that maximises an equity-weighted utility function:

$$U(a) = \delta' \cdot \Delta E_{\text{reduction}} - \eta \cdot C_{\text{deploy}} + \kappa \cdot P_{\text{reach}} + \lambda \cdot \text{Eq_weight}(v) \quad (9)$$

where $\Delta E_{\text{reduction}}$ is the expected reduction in the composite equity risk score from action a , C_{deploy} is the expected deployment cost, P_{reach} is the predicted probability the intervention successfully reaches the target population, and $\text{Eq_weight}(v)$ is a community-specific equity multiplier that increases the weight assigned to historically underserved population subgroups, directly counteracting the per-capita cost penalty that disadvantages dispersed populations under cost-only optimisation. The optimal action is:

$$a^* = \operatorname{argmax}_a U(a) \quad (10)$$

Available response actions include: deployment of a mobile health unit, dispatch of an additional community health worker team, activation of a telehealth bridge with language-concordant providers, expedited Medicaid/CHIP or WIC enrolment outreach, emergency distribution of supplies (vaccines, naloxone, infant formula, water testing kits), and escalation to the state or county health department with a pre-drafted equity impact brief.

E. Crisis Response Coordination and Reporting Layer

The topmost layer translates EquiNet's outputs into reporting artifacts aligned with the frameworks that govern public health resource allocation. For each of six target frameworks, a coverage function C_d measures the proportion of framework-relevant indicators that are continuously monitored and evidenced by EquiNet:

$$C_d = \frac{|\{\text{indicators}_d \text{ automated by EquiNet}\}|}{|\{\text{indicators}_d \text{ total}\}|} \quad (11)$$

Automated evidence artifacts—risk-stratified community dashboards, resource deployment logs, equity impact briefs, and disparity-trend reports—are generated in formats consumable by HRSA Uniform Data System reporting, CDC/ATSDR SVI-based prioritisation processes, Healthy People 2030 disparity tracking, Medicaid Home and Community-Based Services (HCBS) waiver documentation, Title VI civil rights compliance reviews, and the WHO Social Determinants of Health reporting structure.

IV. METHODOLOGY

The experimental design combines synthetic surveillance data generation calibrated to documented public health disparity patterns, a spatial graph simulation environment, and a Python–PyTorch Geometric model training pipeline. Synthetic data is used because individually identifiable community health surveillance data and resource deployment logs are not publicly releasable; all synthetic parameters are calibrated from publicly reported aggregate statistics, including CDC PLACES local health estimates, the CDC/ATSDR SVI 2020 release, American Community Survey 5-year estimates, KFF analyses of racial disparities in maternal health, and Rural Health Information Hub reporting on rural hospital and obstetric unit closures.

A. Surveillance Data Synthesis and Community Graph Simulation

The synthetic dataset comprises 36 months of community-level surveillance (January 2021–December 2023) for 6,400 census tracts across 12 states, selected to represent a mix of urban, rural, and tribal/frontier geographies, covering a simulated population of 18.7 million residents. Weekly community-week observations follow indicator distributions calibrated to published CDC PLACES prevalence estimates and SVI component distributions, with seasonal and event-driven surges (respiratory illness seasons, extreme heat events, formula shortages) drawn from documented 2021–2023 public health event timelines.

Each community-week observation is a structured vector:

$$o_t = [\text{community_id}, \text{sdoh_vector}, \text{crisis_indicator}, \text{resource_request}, \text{timestamp}, \text{population_demographics}, \text{access_metrics}] \quad (12)$$

Crisis escalation events are injected following a stratified fault model: 4.6% of community-months contain an injected crisis escalation across six categories, with category frequencies calibrated to documented disparity concentrations—for example, maternal health crisis injections are weighted toward the 36% of U.S. counties identified as maternity care deserts. Resource-access gap anomalies are injected at 3.1% of community-week observations, representing facility closures, provider shortages, language-access failures, and eligibility-system barriers.

B. Graph Construction and Model Training

The community–resource graph contains approximately 24,800 nodes (communities, facilities, resource pools, and population-group nodes) and 312,000 edges in steady state. The ST-GAT is trained on weekly graph snapshots, with crisis labels derived from the 187-indicator crisis ontology applied to each community node. Class imbalance (95.4% negative, 4.6% positive crisis labels) is addressed through focal loss:

$$FL(p_t) = -\alpha_t (1 - p_t)^\gamma \log(p_t), \quad \gamma = 2.0, \quad \alpha_t = 0.82 \quad (13)$$

The Isolation Forest is trained on 80% of non-injected community-week observations (normal access patterns) with a contamination parameter of 0.05, consistent with the 3.1% anomaly injection rate plus a margin for naturally occurring access exceptions (e.g., temporary clinic closures for staff training).

C. Simulation Configuration

Table 1. Experimental Simulation Configuration Parameters

Configuration Parameter	Value / Specification
Geographies Evaluated	Urban, Rural, and Tribal/Frontier census tracts
Number of Simulated Communities / Population	6,400 census tracts across 12 states; 18.7 million residents
Surveillance Update Cadence	Weekly community-week observations
Crisis Indicator Ontology	187 indicators across 6 crisis categories
AI Model Architecture	Spatiotemporal Graph Attention Network (ST-GAT) + Equity-Weighted Isolation Forest + RL Allocator
Training Dataset	36 months of simulated surveillance and utilisation data (Jan 2021 – Dec 2023)
Frameworks Mapped	Healthy People 2030, CDC/ATSDR SVI, HRSA UDS, Medicaid HCBS, Title VI, WHO SDOH Framework
Simulation Environment	Python 3.11, PyTorch Geometric 2.4, Neo4j 5.x / PostGIS spatial graph store
Train / Test Split	80 / 20 stratified with temporal forward-chaining
Crisis Categories	Maternal Health, Infectious Disease, Chronic Disease, Behavioural Health/SUD, Food Insecurity, Environmental Health

Three baseline configurations are evaluated against EquiNet: manual needs assessment (annual cycle, consistent with current SVI-driven prioritisation practice), rule-based risk scoring (static SVI percentile thresholds without AI), and a tabular machine learning baseline (Gradient Boosted Trees on the same indicator set without graph structure). All configurations are applied to identical synthetic surveillance sequences.

D. Evaluation Metrics

Performance is assessed across four dimensions:

- (i) Crisis risk prediction quality (AUC, Precision, Recall, F1-Score);
- (ii) Operational timeliness (mean time-to-intervention, MTTi, from crisis-risk emergence to response action initiation)
- (iii) Resource allocation automation coverage per crisis domain
- (iv) Equity-gap reduction per historically underserved population subgroup. All metrics are averaged across five simulation runs with distinct random seeds, with 95% confidence intervals

computed via bootstrap resampling (B = 2,000 iterations).

V. EXPERIMENTAL RESULTS

EquiNet was evaluated across the full spectrum of crisis categories and resource-access gap types

against all three baselines. Results consistently demonstrate that graph-AI-native, equity-weighted allocation achieves qualitative improvements over rule-based and manual approaches that cannot be recovered through incremental enhancement of existing tools.

A. Crisis Risk Detection Rate by Category

Table 2a. Crisis Risk Detection Rate (%) by Category — Manual Needs Assessment vs. Rule-Based SVI Threshold vs. EquiNet Framework

Crisis Category	Manual Needs Assessment	Rule-Based SVI Threshold	EquiNet Framework
Maternal Health Crisis	39%	67%	95%
Infectious Disease Outbreak	35%	63%	94%
Chronic Disease Decompensation	31%	58%	90%
Behavioural Health / SUD Crisis	26%	49%	93%
Food Insecurity Surge	33%	61%	91%
Environmental Health Hazard	24%	50%	86%

Table 2a presents detection rates across the six crisis categories. EquiNet achieves detection rates between 86% (Environmental Health Hazard) and 95% (Maternal Health Crisis), substantially outperforming rule-based SVI thresholds (49%–67%) and manual needs assessment (24%–39%) across every category. The widest performance gap occurs in Behavioural Health/Substance Use Disorder crisis detection—EquiNet achieves 93% versus rule-based 49% and manual 26%—reflecting the ST-GAT's ability to identify network-structural signals (social isolation indicators, pharmacy dispensing pattern shifts, and emergency response

call clustering) that are invisible to indicator-threshold scoring.

Environmental Health Hazard detection shows the lowest absolute rate (86%) across all three methods, consistent with the inherent complexity of integrating multi-source data—air quality monitoring, water testing results, and industrial permit records—that update on inconsistent schedules and often lack tract-level resolution. EquiNet's attention mechanism nonetheless identifies 86% of injected environmental hazard escalations, a 35-percentage-point improvement over rule-based approaches that cannot integrate heterogeneous environmental data sources into a unified risk score.

B. Overall Prediction Performance

Table 2. Crisis Risk Prediction Performance Comparison Across Methods

Method	AUC	Precision	Recall	F1-Score	MTTi (days)
Manual Needs Assessment	—	0.55	0.39	0.46	97
Rule-Based SVI Threshold	0.762	0.66	0.57	0.61	41
ML Baseline (no graph)	0.881	0.79	0.73	0.76	11
EquiNet Framework (Proposed)	0.954	0.90	0.89	0.895	3.2

Table 2 presents the aggregate prediction performance. EquiNet achieves an AUC of 0.954, Precision of 0.90, Recall of 0.89, and F1-Score of 0.895—the highest performance across all metrics by substantial margins. The tabular ML baseline without graph structure achieves an F1 of 0.76, confirming that EquiNet's 13.5-percentage-point improvement over the non-graph ML baseline is attributable specifically to the graph topology encoding and attention mechanism, not to machine learning alone. The MTTi reduction from 97 days (manual needs assessment) to 3.2 days (EquiNet) represents a 30.3× improvement, enabling response actions to be initiated within days—rather than the better part of a year—of a crisis risk first emerging in the data.

C. Sensitivity Analysis at Operational Alert Thresholds

A sensitivity analysis at a False Positive Rate of 0.10—the operationally relevant threshold above which field-team alert fatigue begins to degrade response effectiveness—shows that EquiNet achieves a True Positive Rate of 0.89, versus 0.69 for the ML baseline and 0.51 for the rule-based SVI threshold. This 38-percentage-point gap at FPR = 0.10 means that EquiNet correctly flags 89 of every

100 genuine emerging crises while generating only 10 false alerts per 100 stable community-weeks—a signal-to-noise ratio sufficient to sustain community health worker engagement, whereas the rule-based threshold at TPR 0.51 leaves nearly half of genuine emerging crises undetected at the same alert volume.

The superior performance is attributable to three model characteristics. First, the ST-GAT's multi-head attention captures distinct propagation channels—geographic adjacency, shared-facility dependency, and referral-network proximity—simultaneously, detecting risk that spreads through indirect paths (e.g., a neighbouring tract's clinic closure increases referral load on, and therefore future risk in, the focal tract) invisible to tract-level indicator scoring. Second, the Isolation Forest's unsupervised gap detection surfaces previously unseen access-barrier patterns, critical in an environment where the nature of access barriers (new eligibility rules, provider turnover, transit changes) continually evolves. Third, the composite equity risk score $E(v, t)$ integrates predictive, behavioural, structural, and trend signals, reducing both false positives (through SVI-grounded contextual calibration) and false negatives (through multi-signal triangulation).

D. Mean Time-to-Intervention by Crisis Type

Table 2b. Mean Time-to-Intervention (MTTi) by Crisis Type — Manual vs. Rule-Based vs. EquiNet (days)

Crisis Type	Manual (days)	Rule-Based (days)	EquiNet (days)
Infectious Disease Outbreak	64	27	1.1
Maternal Health Crisis	112	46	2.4
Chronic Disease Decompensation	104	39	3.6
Behavioural Health / SUD Crisis	98	44	2.9
Food Insecurity Surge	87	33	3.8
Environmental Health Hazard	118	51	5.4

Table 2b shows that EquiNet's MTTi advantage is most pronounced for the crisis types with the steepest escalation curves. Infectious Disease Outbreaks—the category with the fastest potential for exponential case growth—are flagged by EquiNet in a mean of 1.1 days versus 64 days for

manual needs assessment, a 58.2× improvement that transforms response timing from “after the next annual review” to “within the same surveillance cycle.” Maternal Health Crises show the largest absolute MTTi reduction, from 112 days under manual assessment to 2.4 days under EquiNet,

reflecting the framework's sensitivity to the combination of facility-closure signals and rising prenatal-care-initiation deficits that precede a maternity care access crisis. The relative gain in MTTi across crisis types is quantified as:

$$G_MTTi = (MTTi_manual - MTTi_EquiNet) / MTTi_manual \times 100\% \quad (14)$$

For Infectious Disease Outbreaks: $G_MTTi = (64 - 1.1) / 64 = 98.3\%$, confirming that EquiNet's advantage is greatest precisely where rapid response most directly limits the scale of an outbreak before it overwhelms local healthcare capacity.

E. Resource Allocation Coverage by Domain

Table 3. Resource Allocation Automation Coverage by Crisis Domain

Crisis Domain	Total Indicators	Automated (%)	Manual Residual (%)	Gap Closed vs. Baseline
Maternal & Reproductive Health	34	87.6	12.4	29.8 pp
Infectious Disease Surveillance & Response	29	92.4	7.6	33.1 pp
Chronic Disease Management	38	86.3	13.7	28.4 pp
Behavioural Health & SUD Services	31	88.9	11.1	31.6 pp
Food Security & Nutrition	26	91.7	8.3	32.9 pp
Environmental Health Monitoring	29	90.0	10.0	30.7 pp
Overall (Weighted Avg.)	187	89.4	10.6	31.1 pp

Table 3 presents the resource-allocation automation coverage achieved by EquiNet across the six crisis domains comprising the 187-indicator ontology. Infectious Disease Surveillance & Response achieves the highest automation rate (92.4%), reflecting the availability of structured, frequently updated case-count and laboratory data. Food Security & Nutrition (91.7%) and Environmental Health Monitoring (90.0%) are also highly automated. Chronic Disease Management

shows the lowest automation rate (86.3%), consistent with the inherent need for clinical judgement in individualised care plan adjustments. The weighted-average automation of 89.4% across all 187 indicators represents a 31.1-percentage-point improvement over the rule-based baseline, with the 10.6% manual residual concentrated in domains where community health worker and clinical professional judgement remains essential and appropriate.

F. Equity Gap Reduction by Population Subgroup

Table 4. Equity Gap Index Reduction by Population Subgroup

Population Subgroup	Pre-AI Gap Index	Post-AI Gap Index	% Reduction
Rural / Frontier Communities	68	39	42.6%
Black / African American Communities	71	42	40.8%
Hispanic / Latino Communities	64	37	42.2%
American Indian / Alaska Native Communities	79	44	44.3%
Households Below 200% of Federal Poverty Level	66	38	42.4%
Limited English Proficiency Households	61	37	39.3%
Weighted Average	68.2	39.8	41.6%

Table 4 demonstrates the breadth of EquiNet's equity impact. The gap index—a composite of access, utilisation, and outcome disparities relative to the regional average, scored from 0 (no disparity) to 100 (maximum observed disparity)—falls from a pre-implementation weighted average of 68.2 to a post-implementation 39.8, a 41.6% reduction. The largest absolute reduction occurs for American Indian/Alaska Native communities (79 → 44, a 44.3% reduction), reflecting EquiNet's equity-weighted allocation reward, which assigns these communities a higher $Eq_weight(v)$ term in Equation 9 to counteract the per-capita cost penalty that would otherwise deprioritise dispersed tribal and frontier populations under cost-only optimisation.

The gap reduction for the American Indian/Alaska Native subgroup, computed via Equation 14 (substituting gap indices for time-to-intervention values), is:

$$G_equity(AI/AN) = (79 - 44) / 79 = 44.3\% \quad (15)$$

confirming that EquiNet's largest equity gains accrue to the population subgroup with the highest baseline disparity—the pattern that an explicitly equity-weighted allocation objective is designed to produce, in contrast to aggregate-efficiency objectives that tend to compound rather than reduce baseline disparities.

VI. DISCUSSION

The experimental results support three principal conclusions with direct implications for public health agencies, community health centres, and health technology architects. First, the structural nature of health equity crises—rooted in dynamic, multi-source, compounding risk factors rather than static demographic snapshots—means that rule-based prioritisation tools built on infrequently updated indices are fundamentally misaligned with how disparities actually emerge and escalate. Indices that rank communities against two-year-old vulnerability data cannot detect a crisis that develops over weeks. The ST-GAT's native operation on continuously updated graph topology is not an incremental refinement of SVI-based prioritisation; it is an architectural prerequisite for proactive equity-focused crisis response.

Second, the $30.3\times$ MTTi reduction (from 97 to 3.2 days) carries qualitative significance beyond its quantitative magnitude. In community health, the window during which an emerging crisis—a facility closure, an outbreak, a maternal health access gap—can be addressed with a modest intervention (a mobile clinic deployment, an additional CHW shift,

a telehealth bridge) before it escalates into a more severe and costly outcome is typically measured in days to weeks, not months. A detection regime operating with a 97-day latency provides essentially no opportunity for prevention; a 3.2-day regime creates a genuine intervention window in the large majority of scenarios. This is the difference between a retrospective equity report (documents disparities after they have occurred, for accountability purposes) and a proactive equity intervention (acts to narrow a disparity before it widens)—the shift that public health equity frameworks increasingly call for but that existing tooling has not operationalised.

Third, the 41.6% average reduction in equity gap indices, with the largest gains concentrated among the populations with the highest baseline disparities, demonstrates that embedding equity directly into the allocation objective function (Equation 9) produces materially different—and more equitable—outcomes than allocating resources according to aggregate efficiency and then measuring the resulting disparities after the fact. The ability to generate a single integrated set of evidence artifacts mapped simultaneously to Healthy People 2030 disparity objectives, SVI-based prioritisation, HRSA UDS reporting, and Title VI compliance reviews is, for resource-constrained public health departments and community health centres, an efficiency gain with first-order strategic significance.

Limitations of this study include the use of synthetic surveillance data rather than identifiable community-level data (which would provide higher ecological validity but raises substantial privacy constraints), the focus on six crisis categories that, while broad, do not exhaust the full landscape of health equity concerns (e.g., oral health and vision care access are not separately modelled), and the U.S.-centric framing of the mapped regulatory and reporting frameworks, which excludes important international SDOH reporting structures relevant to low- and middle-income country contexts. Future work will address all three limitations.

VII. CONCLUSION

Underserved populations face a structural health equity crisis-response gap that conventional, indicator-threshold-based prioritisation tools cannot close: community need is generated by fast-moving, multi-source, and interacting risk factors, while existing prioritisation frameworks rely on indices that update on a multi-year cadence and allocation algorithms that optimise for aggregate efficiency rather than equity. This paper presented EquiNet, a spatiotemporal graph-AI framework integrating a

Graph Attention Network for community-level crisis risk prediction with an Isolation Forest for resource-access gap detection, coordinated through a reinforcement-learning allocator with an explicitly equity-weighted reward function, deployed over a continuously updated community-resource graph and mapped to a six-framework reporting layer.

Evaluated against 36 months of simulated public health surveillance data spanning 6,400 census tracts, 18.7 million residents, and 187 crisis indicators across six domains, EquiNet achieves a crisis risk prediction AUC of 0.954, an F1-score of 0.895, a mean time-to-intervention of 3.2 days (versus 97 days for manual needs assessment), resource allocation automation coverage of 89.4% across 187 indicators, and an average equity gap reduction of 41.6% across six historically underserved population subgroups, with the largest gains accruing to the populations with the highest baseline disparities. These results establish graph-AI-native, equity-weighted allocation as a foundational technology for proactive, accountable, and measurably more equitable public health crisis response.

VIII. FUTURE WORK

Future extensions of EquiNet will proceed along four primary vectors. First, integration of additional crisis domains—oral health access, vision and hearing screening gaps, and climate-driven displacement—will extend the indicator ontology and graph architecture to a more complete representation of community health equity. Second, federated learning across county and state health departments will enable shared model improvements without requiring the centralisation of individually identifiable surveillance data, addressing the cold-start problem for smaller jurisdictions with limited historical crisis data. Third, international extension of the reporting layer will incorporate WHO Universal SDOH Coverage monitoring indicators and national SDOH frameworks from low- and middle-income countries, enabling EquiNet to support global health equity reporting alongside U.S. domestic frameworks. Fourth, integration with community-based participatory design processes will create a closed-loop system in which community health workers and community members directly review and adjust the equity weights $Eq_weight(v)$ in Equation 9, transforming algorithmic equity weighting from an externally imposed parameter into a community-governed input.

REFERENCES

- [1]. U.S. Department of Health and Human Services, Office of Disease Prevention and Health Promotion. (2023). Healthy People 2030: Reduce Health Disparities. ODPHP Publications.
- [2]. Agency for Toxic Substances and Disease Registry (ATSDR), Centers for Disease Control and Prevention. (2022). CDC/ATSDR Social Vulnerability Index (SVI) Documentation.
- [3]. Braveman, P., & Gottlieb, L. (2014). The social determinants of health: it's time to consider the causes of the causes. *Public Health Reports*, 129(Suppl 2), 19–31. <https://doi.org/10.1177/00333549141291S206>
- [4]. World Health Organization, Commission on Social Determinants of Health. (2008). Closing the Gap in a Generation: Health Equity Through Action on the Social Determinants of Health. WHO Press.
- [5]. Centers for Disease Control and Prevention. (2024). PLACES: Local Data for Better Health. National Center for Chronic Disease Prevention and Health Promotion.
- [6]. Health Resources and Services Administration (HRSA). (2023). Uniform Data System (UDS) Manual: Health Center Data Reporting Requirements. U.S. Department of Health and Human Services.
- [7]. Kaiser Family Foundation (KFF). (2023). Racial Disparities in Maternal and Infant Health: Current Status and Efforts to Address Them. KFF Issue Brief.
- [8]. Chartis Center for Rural Health. (2023). The Rural Health Safety Net Under Pressure: Rural Hospital Closures. Chartis Group.
- [9]. Centers for Medicare & Medicaid Services. (2023). Medicaid Home and Community-Based Services (HCBS) Guidance. U.S. Department of Health and Human Services.
- [10]. U.S. Department of Justice, Civil Rights Division. (1964, amended). Title VI of the Civil Rights Act of 1964: Prohibition Against National Origin Discrimination Affecting Limited English Proficient Persons.
- [11]. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *Proceedings of ICLR 2017*. <https://arxiv.org/abs/1609.02907>
- [12]. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018). Graph attention networks. *Proceedings of ICLR 2018*. <https://arxiv.org/abs/1710.10903>

- [13]. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 1024–1034.
- [14]. Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM)*, 413–422. <https://doi.org/10.1109/ICDM.2008.17>
- [15]. Panch, T., Mattie, H., & Atun, R. (2019). Artificial intelligence and algorithmic bias: implications for health systems. *Journal of Global Health*, 9(2), 020318. <https://doi.org/10.7189/jogh.09.020318>
- [16]. Pierson, E., Cutler, D. M., Leskovec, J., Mullainathan, S., & Obermeyer, Z. (2021). An algorithmic approach to reducing unexplained pain disparities in underserved populations. *Nature Medicine*, 27, 136–140. <https://doi.org/10.1038/s41591-020-01192-7>
- [17]. Centers for Disease Control and Prevention. (2023). COVID-19 Racial and Ethnic Health Disparities Data Tracker. National Center for Health Statistics.
- [18]. Mullainathan, S., & Obermeyer, Z. (2022). Diagnosing physician error: a machine learning approach to low-value health care. *Quarterly Journal of Economics*, 137(2), 679–727.