

ARTIFICIAL INTELLIGENCE-BASED EARLY WARNING SYSTEMS FOR PREDICTING EMERGING INFECTIOUS DISEASE OUTBREAKS USING MULTISOURCE POPULATION HEALTH DATA

Ajmal Khan¹, Sumbal Shahtaj², Saadia³, Ubaid wazir⁴, Farhad Khattak^{*5}

^{1,2}Lady reading hospital, Peshawar, KPK, Pakistan Department of Pediatrics Medicine

³DHQ Hospital Kamber, Pakistan Department of Internal Medicine

⁴Pakistan Institute of Medical Sciences (PIMS), Islamabad, KPK, Pakistan Department of Surgery

^{*5}Gajajju Khan Medical College Department of internal medicine

¹drajmalkhan426@gmail.com, ²sumbulshahtaj25@gmail.com, ³sadiano2727@gmail.com

⁴ubaidwazirr@gmail.com, ^{*5}authorkhattakfarhad1@gmail.com

DOI: <https://doi.org/10.5281/zenodo.23310204>

Keywords

: artificial intelligence; early warning systems; infectious disease outbreaks; wastewater-based epidemiology; multisource data integration; graph neural networks; physics-informed neural networks; federated learning; health equity; epidemic forecasting

Article History

Received: 25 June 2026

Accepted: 16 September 2026

Published: 12 October 2026

Copyright @Author

Corresponding Author: *
Farhad Khattak

Abstract

The rapid global emergence of novel infectious diseases has exposed the limitations of traditional, reactive public health surveillance systems, which depend on clinical case reporting and consequently suffer from multi-week diagnostic delays. This review examines the paradigm shift toward Artificial Intelligence (AI)-based early warning systems that leverage multisource population health data to enable proactive, predictive epidemic intelligence. We synthesize current evidence on four primary data streams: wastewater-based epidemiology, pathogen genomics and metagenomic deconvolution, digital syndromic signals including Internet-of-Things data, and clinical electronic health records. The analytical architectures underpinning these systems spatiotemporal graph neural networks with dynamic attention mechanisms, epidemiology- and physics-informed neural networks, metapopulation attention networks, and reinforcement learning agents are critically evaluated for their forecasting capabilities. Empirical evidence indicates that wastewater surveillance can detect pathogen signals four to fourteen days before clinical surges, while epidemiology-informed models retain balanced accuracy of 81.53% and AUROC of 96.79% at seven-day horizons, substantially outperforming unconstrained deep learning baselines that underestimate outbreak peaks by 20–40%. However, significant operational bottlenecks persist, including a performance-translation gap, severe false-alert burdens with positive predictive values as low as 12%, algorithmic bias arising from the digital divide, and cross-jurisdictional data governance barriers. We conclude that realizing the full potential of AI-driven early warning requires embedding epidemiological constraints within model architectures, deploying privacy-preserving federated learning frameworks, implementing mandatory fairness auditing, and developing closed-loop prediction-intervention-evaluation systems that translate predictive intelligence into equitable, actionable public health responses.

1. INTRODUCTION

1.1 The Paradigm Shift in Epidemic Surveillance

Global health security relies fundamentally on the rapid detection, characterization, and containment of emerging infectious diseases. Historically, public health surveillance has operated on reactive, clinical-case-dependent paradigms. Traditional passive surveillance relies on healthcare providers manually reporting confirmed laboratory cases or symptomatic illness clusters to public health registries (Zeng et al., 2021). This creates an inherent diagnostic lag, as infected individuals must become symptomatic, seek medical care, undergo diagnostic evaluation, and have their test results aggregated through administrative channels before an outbreak is formally recognized. This multi-week delay severely impairs outbreak response capabilities, allowing novel pathogens or emerging variants to establish widespread community transmission before public health interventions such as ring vaccination, targeted isolation, or travel advisories can be deployed (Rudwan et al., 2026).

The integration of Artificial Intelligence (AI) into public health informatics marks a structural transition from reactive medical models to proactive, predictive, and precision-based population health management. Modern deep learning frameworks, high-throughput biosensing, and cloud computing allow public health systems to process high-velocity, heterogeneous data streams that far exceed human cognitive capacity (Saadatzaheh, 2024). By continuously ingesting non-traditional data including municipal wastewater viral loads, pathogen genomic

sequences, digital syndromic signals, atmospheric metrics, and clinical electronic health records AI-based early warning systems can detect subtle anomalies and project epidemic trajectories well before clinical presentation peaks (Kumar et al., 2027).

1.2 Architectural Framework of Modern AI Early Warning Systems

An AI-enabled early warning system operates across a multi-tier data ingestion and analytics pipeline. The base ingestion layer aggregates environmental, biological, digital, and clinical inputs (BESLEAGA, 2026). The analytical engine transforms these heterogeneous inputs into structured representations, utilizing spatial-temporal graph neural networks, natural language processing, and physics-informed neural networks. These models process the integrated data streams to execute three core functions: anomaly detection to identify abnormal signal spikes, time-series forecasting to project case trajectories and peak magnitudes, and spatial risk mapping to isolate geographic transmission hotspots (Baneres et al., 2021).

The top layer converts these predictive outputs into actionable public health intelligence, enabling risk stratification and automated decision support. The structural goal of this framework is to maximize lead time the operational window between an AI-generated alert and peak community transmission while minimizing false alarms that induce alert fatigue and misallocate scarce public health resources (Maheswaran et al., 2025).




Figure 1: Conceptual Framework and Integrated Ecosystem of Artificial Intelligence Applications Across Biomedical and Public Health Domains

2. Multisource Population Health Data Integration

2.1 Environmental and Wastewater-Based Epidemiology

Wastewater-based epidemiology has emerged as a non-invasive, cost-effective tool for population-level health monitoring. Individuals infected with enteric, respiratory, or systemic pathogens shed viral particles, bacterial cells, or genetic material in urine and feces, regardless of whether they develop clinical symptoms. Consequently, municipal sewage systems act as aggregated biological sensor arrays, capturing shedding signals from both symptomatic and asymptomatic cohorts across entire sewersheds (Saingam et al., 2024).

Historically applied to track wild poliovirus and enteric enteroviruses, wastewater surveillance expanded significantly during the COVID-19 pandemic to track SARS-CoV-2 dynamics and

variant substitutions. Recent advances pair wastewater sampling with automated AI analytics to identify low-abundance pathogen signals weeks before clinical escalation (Therrien et al., 2024). For example, machine learning algorithms applied to municipal wastewater sequencing data have demonstrated the capacity to detect distinct genomic signatures for influenza, respiratory syncytial virus, mpox, measles, *Candida auris*, and novel SARS-CoV-2 sublineages with high confidence using as few as two to five wastewater samples. This capability effectively uncouples disease surveillance from healthcare-seeking behavior and clinical testing access (Oh & Wijaya, 2025).

Commercial and open-access platforms, such as Biobot Analytics and SmartHealth-Track, operationalize wastewater surveillance by combining quantitative reverse transcription PCR

viral loads, digital PCR, and next-generation sequencing data with regression models and isolation forest algorithms. These models normalize raw genetic concentrations against human fecal cross-assembly phage markers and wastewater flow rates, using anomaly detection to flag statistically significant deviations from baseline viral shedding (Shao et al., 2026).

2.2 Pathogen Genomics and Metagenomic Deconvolution

While RT-qPCR confirms the presence of targeted, known pathogens in wastewater or clinical samples, unbiased metagenomic next-generation sequencing enables the discovery of novel or mutated biological agents. However, sequencing pooled environmental samples presents a significant computational challenge: a single wastewater sample contains a complex mixture of degraded genetic material from thousands of host organisms, commensal microbiota, and multiple overlapping viral lineages (Stalder et al., 2019).

AI algorithms solve this deconvolution challenge by analyzing deep sequencing read depths and single nucleotide variant frequencies. Machine learning architectures trained on global genomic repositories can reconstruct individual viral strain haplotypes from mixed reads, estimating the relative prevalence of emerging variants within a community (Brito & Alm, 2016). When paired with predictive algorithms, genomic wastewater surveillance can identify fitness-enhancing mutations such as amino acid substitutions in viral surface proteins that confer immune evasion or heightened receptor affinity prior to their widespread clinical proliferation (O'Donald et al., 2026).

2.3 Digital Health, Syndromic Signals, and Internet-of-Things

Digital syndromic surveillance captures real-time health-seeking behaviors that precede formal

clinical diagnosis. Search engine queries regarding specific symptoms, social media posts, over-the-counter pharmaceutical sales, and workplace absenteeism records provide high-velocity proxies for community disease activity (Lee et al., 2023).

Natural Language Processing models scan unformatted public texts, regional news outlets, and informal health worker reports to detect emerging disease clusters. India's "Health Sentinel" tool demonstrates the efficiency of this approach: by leveraging NLP for automated media scanning, the system reduced manual monitoring workloads by 98% while identifying potential disease outbreaks with 96% efficiency, generating thousands of real-time public health alerts. Similarly, the integration of Internet-of-Things streams including continuous physiological data from consumer wearables and smart thermometers provides early population-level signals of systemic infection prior to formal medical evaluation (Hassanalieragh et al., 2015).

2.4 Clinical Electronic Health Records and Health System Intelligence

Within institutional healthcare settings, Electronic Health Record systems offer deep, patient-level clinical data. Real-time algorithms embedded in clinical workflows analyze triage notes, emergency department chief complaints, vital sign fluctuations, and routine laboratory orders to monitor nosocomial transmission and unexpected surges in acute respiratory or gastrointestinal illness (Ye et al., 2024). Advanced clinical decision support algorithms, including reinforcement learning agents, utilize these continuous clinical time-series to stratify patient risk, optimize therapeutic resource allocation, and predict regional intensive care unit capacity limits during large-scale public health emergencies (Combi & Pozzi, 2019).

Table 1: Taxonomy and Operational Dynamics of Multisource Population Health Data Streams

Data Stream	Primary Sources	Key Analytical AI Techniques	Predictive Lead Time	Major Operational Strengths	Inherent Limitations & Biases
Wastewater Epidemiology	Municipal sewage, influent collection points	Anomaly detection, regression models, viral load normalization	4 to 14 days prior to clinical surge	Unbiased population coverage; includes asymptomatic cases	Flow-rate dilution; degradation of viral RNA in sewage pipes
Pathogen Genomics	Metagenomic sequencing, targeted amplicon NGS	Deep mixture deconvolution, haplotype reconstruction, SNV analysis	Concurrent to early lineage emergence	Identifies novel variants and antigenic drift mutations	High cost; complex bioinformatic pipelines; batch latency
Digital Syndromic Data	Search logs, social media NLP, OTC pharmacy sales, IoT wearables	Long Short-Term Memory, Isolation Forests, Transformer NLP	2 to 7 days prior to clinical visits	High velocity; captures early pre-diagnostic behavior	High background noise; digital divide bias; public panic distortion
Clinical Health System Data	EHR time-series, ED chief complaints, laboratory orders	Time-series classifiers, Reinforcement Learning (MDP), Survival Models	Real-time to 48 hours for clinical deterioration	High diagnostic specificity and granular physiological data	Healthcare-seeking bias; exclusion of uninsured cohorts; inter-EHR interoperability gaps
Environmental & Spatial Data	Land-surface temp, absolute humidity, land-use maps, mobility feeds	Spatiotemporal Graph Neural Networks, GATv2	Seasonal to 4-week horizon forecasting	Captures environmental transmission drivers and spatial decay	Spatial coarseness; complex non-linear confounding factors

3. Advanced Algorithmic Architectures for Outbreak Forecasting

3.1 Spatial-Temporal Graph Neural Networks and Attention Mechanisms

Infectious disease transmission inherently manifests across both spatial and temporal dimensions. Pathogens spread through physical space driven by human mobility, transportation networks, and localized environmental conditions, while disease progression follows time-dependent epidemiological trajectories.

Traditional sequence models such as Recurrent Neural Networks or standard Transformers process temporal dependencies effectively but treat geographic locations as independent, isolated channels, thereby ignoring spatial interactions (Zhu et al., 2024).

Spatiotemporal Graph Neural Networks (ST-GNNs) overcome this limitation by modeling geographic regions, wards, or municipalities as nodes V within a dynamic graph $G = (V, E, W)$, where edges E represent connectivity and weights

W capture spatial interaction strength over time t . Advanced architectures incorporate dynamic graph attention mechanisms (such as GATv2) to update edge weights continuously based on time-varying environmental and human factors (Wang et al., 2022).

A prominent example is meteorology-aware dynamic graph construction, which addresses the reality that geographic distance alone cannot capture pathogen transmission efficiency. Specific meteorological precursors such as sustained temperature drops, absolute humidity falling below 7 g/m^3 , and abrupt barometric pressure shifts exceeding 5 hPa frequently precede viral transmission spikes (Luo et al., 2025). To capture this, models dynamically construct topological edge weights $w(i, j, t)$ by combining Gaussian-kernel meteorological similarity $S_{\text{met}}(i, j, t)$ with exponentially decaying geographic distance $A_{\text{geo}}(i, j)$:

$$S_{\text{met}}(i, j, t) = \exp\left(-\frac{\|x_{i,t} - x_{j,t}\|^2}{2\sigma^2}\right)$$

$$A_{\text{geo}}(i, j) = \exp(-\lambda \cdot \text{dist}(i, j)) \cdot \mathbb{1}[\text{dist}(i, j) < \delta]$$

$$w(i, j, t) = \text{Sigmoid}(W_w \cdot [S_{\text{met}}(i, j, t); A_{\text{geo}}(i, j); h_{i,t}])$$

Here, $x_{i,t}$ represents the meteorological feature vector for region i at time t , $\text{dist}(i, j)$ measures great-circle geographic distance, $h_{i,t}$ denotes hidden spatial embeddings, σ is the bandwidth scale parameter, λ is the spatial decay scale, and δ is the adjacency threshold mask. This dynamically weighted topological graph feeds into spatial graph convolutions, allowing attention mechanisms to prioritize transmission pathways between climatologically aligned and geographically connected zones (Pham et al., 2026).

3.2 Physics-Informed and Epidemiology-Informed Neural Networks

Despite their expressive capacity, unconstrained deep learning models face significant failure modes when applied to epidemic forecasting. Standard mean-squared-error loss functions push pure data-driven models toward the temporal mean, as non-outbreak baseline periods

outnumber epidemic peak periods (Nelson et al., 2025). Consequently, black-box deep learning architectures routinely underestimate outbreak peak magnitudes by 20% to 40% and introduce peak timing errors exceeding two weeks. Furthermore, when exposed to out-of-distribution data such as structural shifts caused by non-pharmaceutical interventions or novel pathogen variants purely data-driven models often yield unfeasible predictions that violate basic mass-balance constraints (Vindas et al., 2026).

To resolve these vulnerabilities, Epidemiology-Informed Graph Neural Networks (EIGNNs) and Physics-Informed Neural Networks (PINNs) explicitly integrate classic compartmental differential equations into the deep learning architecture (Treibert & Ehrhardt, 2022). Rather than relying solely on empirical data matching, EIGNNs embed deterministic mathematical formulations such as the Susceptible-Exposed-Infectious-Recovered continuous-time differential system directly into the loss function or layer transitions:

$$\begin{aligned} \frac{dS}{dt} &= -\beta \frac{S \cdot I}{N}, & \frac{dE}{dt} &= \beta \frac{S \cdot I}{N} - \sigma E, & \frac{dI}{dt} &= \sigma E - \gamma I, & \frac{dR}{dt} &= \gamma I \end{aligned}$$

In an EIGNN framework, the total optimization objective balances empirical prediction error with epidemiological physics constraints:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}}(\hat{y}, y) + \mu \mathcal{L}_{\text{epi}}\left(\frac{d\hat{y}}{dt}, f_{\text{SEIR}}(\hat{y}; \hat{\beta}, \hat{\sigma}, \hat{\gamma})\right)$$

Here, $\mathcal{L}_{\text{task}}$ measures cross-entropy or prediction error against observed incidence rates y , while $\mathcal{L}_{\text{epi}}$ penalizes predictions that deviate from the governing differential equations f_{SEIR} parameterised by neural-estimated transmission ($\hat{\beta}$), incubation ($\hat{\sigma}$), and recovery ($\hat{\gamma}$) rates, scaled by regularization weight μ (Torku et al., 2021).

By enforcing structural epidemiological constraints via automatic differentiation, physics-informed models maintain strong predictive stability across long forecast horizons. Empirical evaluations on hospital transmission datasets demonstrate that epidemiology-informed models

retain a balanced accuracy of 81.53% and an AUROC of 96.79% at a 7-day forecast horizon, substantially outperforming non-constrained baseline deep learning models (Rifanti et al., 2026).

3.3 Metapopulation Attention Networks and Heterogeneity Modeling

In large-scale regional or international forecasting, population heterogeneity across age brackets, contact rates, and socio-environmental vulnerability introduces complex transmission dynamics. Metapopulation Spatio-Temporal Attention Networks (MPSTAN) handle these variations by utilizing separate, adaptive neural parameter generators (Mao et al., 2024). Rather than assuming static, uniform mixing, these parameter generators process distinct embedding spaces for intra-patch (within-community) and inter-patch (between-community) disease dynamics. This setup enables the neural network to infer spatial transmission rates dynamically without requiring real-time, privacy-restricted population mobility tracking matrices (Cao et al., 2023).

3.4 Reinforcement Learning for Sequential Public Health Interventions

As early warning systems evolve, public health analytics is shifting from passive risk scoring toward sequential, dynamic decision support. Reinforcement Learning (RL) agents model public health response as a Markov Decision Process (MDP) defined by a tuple (S, A, P, R, γ) , where states S represent population disease states, actions A represent countermeasure interventions, P defines state transition dynamics, R represents the reward function, and γ is the discount factor (Weltz et al., 2022).

Similar to clinical RL implementations like the "AI Clinician" which extracts optimal treatment policies for sepsis management from high-dimensional intensive care time-series public health RL agents evaluate thousands of simulated epidemic trajectories. By identifying intervention policies that maximize long-term survival rewards, these models provide decision-makers with optimized countermeasure strategies tailored to evolving outbreak scenarios (Kwak et al., 2021).

Table 2: Comparative Analysis of Core AI Model Architectures in Infectious Disease Early Warning

Model Architecture	Input Data Modalities	Spatial-Temporal Mechanics	Peak Prediction Reliability	Computational & Mathematical Complexity	Primary Failure Modes
Standard Sequence Models (LSTM, GRU)	Univariate/Multivariate time-series case counts	Temporal dependencies via recurrent gating mechanisms	Poor; systematically underestimates peaks by 20-40%	Low to Moderate computational footprint	Ignores spatial connectivity; susceptible to temporal phase lag
Spatiotemporal Graph Neural Networks (ST-GNN / GATv2)	Spatial matrices, mobility feeds, climate inputs	Graph convolutions coupled with temporal attention layers	Moderate to High; captures spatial expansion patterns	High; requires matrix operations over dynamic graph topologies	Vulnerable to distribution shifts and unconstrained out-of-bounds errors
Epidemiology-Informed	Multisource population health	Physics-informed	High; maintains	Very High; requires	High sensitivity to

GNNs (EIGNN / PIGNN)	data, compartment ODEs	loss regularization enforcing differential equations	accuracy across 7- to 28-day horizons	continuous automatic differentiation and ODE solver integration	initial parameter bounds (β, γ)
Metapopulation Attention Networks (MPSTAN)	Regional multi-patch case series, demographic profiles	Dual neural parameter generators for intra- and inter-patch dynamics	High; robust across heterogeneous regional trends	High; multi-stream parameter generator optimization	Parameter identifiability challenges under sparse regional testing
Reinforcement Learning (MDP Agents)	Longitudinal state vectors, intervention histories, capacity levels	Policy iteration optimization over simulated action-state trajectories	N/A (Focused on optimal intervention execution)	Extremely High; demands off-policy evaluation and vast simulation spaces	Reward function misspecification; off-policy distribution shifts

4. Operational Bottlenecks, Clinical Translation, and System Implementation

4.1 The Performance-Translation Gap and False-Alert Pathology

A primary challenge in deploying AI early warning systems is the gap between offline predictive validation and real-world operational performance. Machine learning models frequently achieve high area under the receiver operating characteristic curve scores (0.85 – 0.95) during retrospective evaluations on curated historical datasets. However, when deployed in prospective, live operational environments, many algorithms experience significant performance drops (Mercer, 2026).

This translation gap is illustrated by external prospective validations of commercial health system early warning models. A notable evaluation of Epic's proprietary sepsis early warning model revealed an operational sensitivity of only 33% and a positive predictive value (PPV) of 12%, meaning 88% of generated alerts were false positives. This low PPV creates a severe false-alert

burden, leading to alert fatigue among clinical staff and public health officers. When clinical and operational teams are inundated with false alarms, they naturally adapt by overriding, delaying, or ignoring system outputs (Otero et al., 2026).

The operational risk of alert-induced workflow disruption was highlighted in the NAVOY prospective randomized controlled trial evaluating a sepsis early warning system. The trial found that prospective deployment of the alert model resulted in numerically higher intensive care unit mortality in the intervention group compared to usual care (13.2% vs. 11.6%) (Bayat et al., 2026). This outcome demonstrated that early warnings do not confer clinical benefit if they disrupt workflows without offering clear, actionable, and protocolized intervention pathways. Applying a formal four-level clinical actionability scoring framework reveals that 61.5% of published medical and epidemiological AI systems remain at Level 0 (no demonstrated actionability), with only 9.2% reaching Level 3 (proven prospective operational impact) (Reichstein et al., 2025).




Figure 2: Performance Evaluation and Empirical Outbreak Surveillance Metrics of the SmartHealth-Track AI Platform

4.2 Algorithmic Bias, Health Equity, and the Digital Divide

AI-based population health monitoring systems risk exacerbating existing health inequities if algorithmic biases and digital coverage disparities are unaddressed. Predictive models trained predominantly on data from affluent, high-income healthcare jurisdictions perform poorly when applied to low-resource or under-monitored communities (Osonuga et al., 2025).

This disparity stems fundamentally from the digital divide. Digital syndromic feeds (search logs, social media NLP, smart wearable metrics) disproportionately sample younger, urban, and higher socioeconomic demographics with ready access to digital technology. Conversely, rural populations, older age brackets, marginalized

groups, and low-and-middle-income countries are systematically underrepresented in digital data streams (Koehle et al., 2022).

Furthermore, model accuracy often degrades non-linearly across age groups and hardware configurations. For example, studies on deep learning diagnostic models demonstrate that algorithmic predictive performance drops significantly in older demographics (e.g., accuracy declining from 84.5% in patients under 40 years to 66.2% in individuals aged 75 and older). When public health early warning systems allocate testing resources, mobile clinics, or medical stockpiles based on biased digital signals, they risk directing resources toward digitally visible cohorts while under-serving vulnerable populations (Chin et al., 2023).

4.3 Data Governance, Privacy-Preserving Federated Learning, and Interoperability

Integrating multisource health data across jurisdictional and institutional boundaries introduces regulatory, legal, and data governance hurdles. Aggregating individual-level clinical records, high-resolution geolocation tracking, and consumer digital footprints triggers significant data privacy concerns under frameworks such as HIPAA and GDPR (Hasan & Kudapa, 2021).

To overcome privacy barriers without sacrificing predictive performance, modern early warning architectures are increasingly adopting Federated Learning. In a federated framework, a centralized global early warning model distributes its base parameters to local nodes (such as regional health

authorities, individual hospital systems, or municipal wastewater treatment labs). Each node trains the model locally on its private data streams, sending only anonymized parameter gradient updates back to the central server for aggregation (Li et al., 2025).

This decentralized approach preserves data sovereignty and security while enabling the collective intelligence of multi-institutional datasets. However, implementing federated learning requires robust data standardization protocols (such as HL7 FHIR for clinical records and standardized schemas for pathogen genomics) to ensure interoperability across heterogeneous hardware and software infrastructures (Loftus et al., 2022).

Table 3: Critical Implementation Bottlenecks and Strategic Engineering Solutions

Category	Primary Implementation Bottleneck	Underlying Cause	Technical & Structural Solution
Operational & Clinical	Severe False-Alert Burden & Alert Fatigue	High false-positive rates (e.g., 12% PPV in clinical CDS alerts)	Implementing human-in-the-loop validation & multi-stage alarm thresholds
Statistical & Algorithmic	Out-of-Distribution Peak Underestimation	Unconstrained MSE loss functions favoring baseline temporal means	Embedding physics/epidemiology constraints (EIGNNs / PINNs)
Equity & Social	Digital Divide & Demographic Bias	Under-sampling of low-income, elderly, and offline populations	Mandatory fairness auditing & hybrid WBE/clinical data weighting
Governance & Privacy	Cross-Jurisdictional Data Sharing Blockades	Regulatory restrictions (HIPAA/GDPR) & data sovereignty concerns	Deploying Privacy-Preserving Federated Learning frameworks
Technical Integration	Data Heterogeneity & Interoperability Gaps	Unstructured text, varying genomic depth, siloed EHR schemas	Harmonizing data via standardized HL7 FHIR & automated bioinformatic pipelines

5. Emerging Paradigms and Strategic Directions

5.1 Foundation Models and Cross-Modal Generative AI

The field of epidemic intelligence is moving toward multi-modal public health foundation models. Inspired by large language models, these foundation architectures undergo self-supervised pre-training on vast repositories of biological,

meteorological, demographic, and spatial-temporal time-series data (Awais et al., 2025).

By learning unified generalized representations of population health dynamics, cross-modal foundation models can be fine-tuned for downstream public health tasks such as predicting variant-specific transmission trajectories, forecasting regional hospital bed demand, or

deconvolving multi-pathogen wastewater signals with minimal task-specific labeled data. This flexibility is particularly valuable during novel pathogen scenarios, where an emerging disease presents unknown biological transmission parameters that render historical statistical models obsolete (Chen et al., 2024).

5.2 Closed-Loop "Prediction-Intervention-Evaluation" Systems

To bridge the gap between predictive intelligence and public health response, future early warning platforms are aiming to integrate prediction, intervention, and evaluation into unified closed-loop decision support engines (Chang et al., 2025). In this integrated workflow, continuous multisource data ingestion feeds into spatial-temporal AI engines that identify emerging outbreak risks. These risk profiles trigger causal AI models that evaluate potential countermeasure strategies such as localized rapid-testing deployment, targeted public advisories, or pharmaceutical distribution and project their comparative impacts (Chen et al., 2026).

Once public health authorities initiate a chosen intervention, the platform continuously monitors incoming population health data streams to evaluate real-world countermeasure effectiveness, dynamically updating its underlying causal models. This closed-loop approach ensures that AI early warning insights directly inform, optimize, and refine public health responses over time (Wu, 2025).

Conclusion

This review demonstrates that AI-based early warning systems represent a fundamental structural transition in epidemic surveillance moving from reactive, clinical-case-dependent paradigms toward proactive, predictive, and precision-based population health management. By integrating heterogeneous data streams including wastewater viral loads, pathogen genomic sequences, digital syndromic signals, environmental metrics, and clinical electronic health records, modern analytical architectures can detect subtle anomalies and project epidemic

trajectories well before clinical presentation peaks. The algorithmic evidence synthesized here reveals a clear hierarchy of model performance. Unconstrained sequence models systematically underestimate outbreak peaks by 20–40% and introduce timing errors exceeding two weeks, while spatiotemporal graph neural networks capture spatial expansion patterns but remain vulnerable to distribution shifts. Epidemiology-informed and physics-informed neural networks, which embed compartmental differential equations directly into their loss functions, achieve superior predictive stability across extended forecast horizons, retaining balanced accuracy of 81.53% and AUROC of 96.79% at seven days. Metapopulation attention networks further address population heterogeneity, while reinforcement learning agents offer sequential decision support for intervention optimization. Nevertheless, substantial barriers separate predictive promise from operational impact. The performance-translation gap remains stark: models achieving AUROC scores of 0.85–0.95 in retrospective validation frequently collapse in prospective deployment, with commercial sepsis early warning systems demonstrating sensitivity of only 33% and positive predictive values of 12%. Such false-alert burdens induce alert fatigue, workflow disruption, and, as the NAVOY trial demonstrated, potential patient harm. Concurrently, algorithmic bias rooted in the digital divide threatens to exacerbate health inequities by systematically underrepresenting rural, elderly, and low-income populations. Cross-jurisdictional data sharing remains constrained by regulatory frameworks including HIPAA and GDPR. Addressing these challenges requires a multi-pronged strategy. First, epidemiological and physical constraints must be embedded within model architectures to ensure biologically plausible predictions under out-of-distribution conditions. Second, privacy-preserving federated learning frameworks, coupled with standardized interoperability protocols such as HL7 FHIR, can enable multi-institutional intelligence without compromising data sovereignty. Third, mandatory fairness auditing and hybrid data weighting

strategies are essential to ensure equitable resource allocation. Fourth, human-in-the-loop validation and multi-stage alarm thresholds can mitigate false-alert pathology while preserving sensitivity. Looking forward, the emergence of cross-modal public health foundation models pre-trained on vast repositories of biological, meteorological, demographic, and spatial-temporal data offers unprecedented flexibility for novel pathogen scenarios where historical statistical models become obsolete. The ultimate vision is a closed-loop prediction-intervention-evaluation system in which continuous multisource data ingestion feeds spatial-temporal AI engines that identify emerging risks, trigger causal models to evaluate countermeasure strategies, and dynamically update based on real-world intervention effectiveness. Realizing this vision will require sustained transdisciplinary collaboration among epidemiologists, computer scientists, clinicians, ethicists, and policymakers. The technical capacity for early, accurate, and equitable outbreak prediction now exists; the imperative is to build the governance, infrastructure, and translational pathways that convert algorithmic insight into protection of global health security.

REFERENCES

- Awais, M., Naseer, M., Khan, S., Anwer, R. M., Cholakal, H., Shah, M., ... & Khan, F. S. (2025). Foundation models defining a new era in vision: a survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(4), 2245-2264.
- Baneres, D., Guerrero-Roldán, A. E., Rodríguez-González, M. E., & Karadeniz, A. (2021). A predictive analytics infrastructure to support a trustworthy early warning system. *Applied Sciences*, 11(13), 5781.
- Bayat, A., Pavakopethan, A., Kashmar, N., Yousuf, I., Saab, A., Salem Sokhn, E., ... & El Morr, C. (2026). AI-based epidemic early warning using social media and search engine signals: A scoping review. *Health Informatics Journal*, 32(3), 14604582261479106.
- BESLEAGA, A. N. (2026). Live Mobile Edge Sensors Swarm System: Decentralized AI-Driven Early Warning Architecture for Disaster Response and Climate Monitoring.
- Brito, I. L., & Alm, E. J. (2016). Tracking strains in the microbiome: insights from metagenomics and models. *Frontiers in Microbiology*, 7, 712.
- Cao, Q., Jiang, R., Yang, C., Fan, Z., Song, X., & Shibasaki, R. (2023). Metapopulation graph neural networks: Deep metapopulation epidemic modeling with human mobility. *arXiv preprint arXiv:2306.14857*.
- Chang, Y. H., Chen, F. C., & Lee, C. I. (2025). Developing an early warning system with personalized interventions to enhance academic outcomes for at-risk students in Taiwanese higher education. *Education Sciences*, 15(10), 1321.
- Chen, Y., Qin, X., & Chen, S. (2026). From graph models to intelligent decision-making: a review of spatio-temporal graph neural networks for regional disease risk prediction and etiology mining. *Frontiers in Public Health*, 14, 1830634.
- Chen, Z., Xu, L., Zheng, H., Chen, L., Tolba, A., Zhao, L., ... & Feng, H. (2024). Evolution and Prospects of Foundation Models: From Large Language Models to Large Multimodal Models. *Comput. Mater. Continua*, 80(2), 1753-1808.
- Chin, M. H., Afsar-Manesh, N., Bierman, A. S., Chang, C., Colón-Rodríguez, C. J., Dullabh, P., ... & Ohno-Machado, L. (2023). Guiding principles to address the impact of algorithm bias on racial and ethnic disparities in health and health care. *JAMA network open*, 6(12), e2345050.
- Combi, C., & Pozzi, G. (2019). Clinical information systems and artificial intelligence: recent research trends. *Yearbook of medical informatics*, 28(01), 083-094.

- Hasan, M. T., & Kudapa, S. P. (2021). Data privacy-aware machine learning and federated learning: A framework for data security. *American Journal of Interdisciplinary Studies*, 2(03), 01-34.
- Hassanalieragh, M., Page, A., Soyata, T., Sharma, G., Aktas, M., Mateos, G., ... & Andreescu, S. (2015, June). Health monitoring and management using Internet-of-Things (IoT) sensing with cloud-based processing: Opportunities and challenges. In *2015 IEEE international conference on services computing* (pp. 285-292). IEEE.
- Koehle, H., Kronk, C., & Lee, Y. J. (2022). Digital health equity: addressing power, usability, and trust to strengthen health systems. *Yearbook of medical informatics*, 31(01), 020-032.
- Kumar, R., Mitra, D., Alloun, W., Sierra, B. G., & Sağlam, K. (2027). AI in Epidemiology, Pandemic Prediction, and Disease Modeling. In *Geographic Information Systems for Medical Data Analysis* (pp. 393-422). IGI Global Scientific Publishing.
- Kwak, G. H., Ling, L., & Hui, P. (2021). Deep reinforcement learning approaches for global public health strategies for COVID-19 pandemic. *PloS one*, 16(5), e0251550.
- Lee, H. Y., Lee, K. H., Lee, K. H., Erdenbayar, U., Hwang, S., Lee, E. Y., ... & Youk, H. (2023). Internet of medical things-based real-time digital health service for precision medicine: Empirical studies using MEDBIZ platform. *Digital health*, 9, 20552076221149659.
- Li, N., Lewin, A., Ning, S., Waito, M., Zeller, M. P., Tinmouth, A., ... & Canadian Transfusion Trials Group. (2025). Privacy-preserving federated data access and federated learning: Improved data sharing and AI model development in transfusion medicine. *Transfusion*, 65(1), 22-28.
- Loftus, T. J., Ruppert, M. M., Shickel, B., Ozrazgat-Baslanti, T., Balch, J. A., Efron, P. A., ... & Bihorac, A. (2022). Federated learning for preserving data privacy in collaborative healthcare research. *Digital Health*, 8, 20552076221134455.
- Luo, J., Wang, X., Fan, X., He, Y., Du, X., Chen, Y. Q., & Zhao, Y. (2025). A novel graph neural network based approach for influenza-like illness nowcasting: exploring the interplay of temporal, geographical, and functional spatial features. *BMC Public Health*, 25(1), 408.
- Maheswaran, N., Logeswari, G., Bose, S., Gokulraj, G., Vijayaraj, G., & Prabhu, D. (2025, November). AIOPS-Enabled IoT Framework for Intelligent Early Warning and Disaster Management Systems. In *2025 International Conference on Intelligent Computing, Information and Control Systems (ICOIICS)* (pp. 885-891). IEEE.
- Mao, J., Han, Y., & Wang, B. (2024). MPSTAN: Metapopulation-based spatio-temporal attention network for epidemic forecasting. *Entropy*, 26(4), 278.
- Mercer, D. R. (2026). AI-Driven Risk Identification and Early-Warning Systems: A Cross-Domain Systematic Survey of Methods, Applications, and Trustworthiness. *Annals of Applied Sciences*, 7(1).
- Nelson, S. P., Raja, R., Eswaran, P., Alzabut, J., & Rajchakit, G. (2025). Modeling the dynamics of COVID-19 in Japan: Employing data-driven deep learning approach. *International Journal of Machine Learning and Cybernetics*, 16(10), 7045-7058.
- O'Donald, S. N., Patel, F., Keen, P., Hanson, L. A., Cunningham, F., Lawrence, M. L., & Tekedar, H. C. (2026). Hi-C Metagenome Deconvolution of Double-Crested Cormorant (*Nannopterum auritum*) Fecal Samples Demonstrates Feasibility of Linking Microbial Genomes, AMR Genes, and Mobile Elements in Avian Microbiomes. *Microorganisms*, 14(6), 1198.

- Oh, S., & Wijaya, J. (2025). Predictive surveillance and diagnosis of COVID-19: An integrative machine learning and wastewater multi-omics approach. *Water Research*, 124981.
- Oh, S., & Wijaya, J. (2025). Predictive surveillance and diagnosis of COVID-19: An integrative machine learning and wastewater multi-omics approach. *Water Research*, 124981.
- Osonuga, A., Osonuga, A. A., Fidelis, S. C., Osonuga, G. C., Juckes, J., & Olawade, D. B. (2025). Bridging the digital divide: artificial intelligence as a catalyst for health equity in primary care settings. *International Journal of Medical Informatics*, 106051.
- Otero, N., Pelivan, I., Fitoz, B., Luterbacher, J., Kuglitsch, M. M., & Ma, J. (2026). Advancing early warning systems: AI applications, challenges and emerging insights. *Environmental Research Letters*, 21(8), 083005.
- Pham, D. T., Van Nguyen, D., Tran, T. T., & Nguyen, V. A. (2026). Casenet: a novel HIV forecasting model using attention-based spatio-temporal graph network leveraging city-wide infection correlations. *Neural Computing and Applications*, 38(5), 119.
- Reichstein, M., Benson, V., Blunk, J., Camps-Valls, G., Creutzig, F., Fearnley, C. J., ... & Weldemariam, K. (2025). Early warning of complex climate risk with integrated artificial intelligence. *Nature Communications*, 16(1), 2564.
- Rifanti, U. M., Susyanto, N., & Boonklurb, R. (2026). SUC-PINNs: A physics-informed neural networks approach to inverse problems in epidemic models with partial observability. *International Journal of Modeling, Simulation, and Scientific Computing*, 2650043.
- Rudwan, E., Ali, H. E. M., Ibrahim, M., Mohamedelnour, H., Almandalawi, Y. S., Hassan, M. K. M., ... & Rudwan, A. (2026). Artificial intelligence-driven approaches to disease surveillance and outbreak prediction. *Open Access Library Journal*, 13(2), 1-16.
- Saadatzadeh, R. (2024). The future of public health: integrating artificial intelligence in disease surveillance and prevention. *Journal of Foresight and Health Governance*, 1(2), 1-15.
- Saingam, P., Jain, T., Woicik, A., Li, B., Candry, P., Redcorn, R., ... & Winkler, M. K. (2024). Integrating socio-economic vulnerability factors improves neighborhood-scale wastewater-based epidemiology for public health applications. *Water Research*, 254, 121415.
- Shao, X. T., Li, H. Y., Guan, J. B., Ma, X. D., & Wang, D. G. (2026). A novel fever prevalence indicator for early warning of acute respiratory infections using wastewater-based epidemiology and machine learning. *Environmental Research*, 125279.
- Stalder, T., Press, M. O., Sullivan, S., Liachko, I., & Top, E. M. (2019). Linking the resistome and plasmidome to the microbiome. *The ISME journal*, 13(10), 2437-2446.
- Therrien, J. D., Thomson, M., Sion, E. S., Lee, I., Maere, T., Nicolaï, N., ... & Vanrolleghem, P. A. (2024). A comprehensive, open-source data model for wastewater-based epidemiology. *Water Science & Technology*, 89(1), 1-19.
- Torku, T. K., Khaliq, A. Q., & Furati, K. M. (2021). Deep-data-driven neural networks for COVID-19 vaccine efficacy. *Epidemiologia*, 2(4), 564-586.
- Treibert, S., & Ehrhardt, M. (2022). A Physics-Informed Neural Network to Model COVID-19 Infection and Hospitalization Scenarios.

Vindas, Y., Bornet, A., Abbas, M., Geissbühler, D., Rodrigues-Jr, J. F., & Teodoro, D. (2026). Epidemiology-Informed Graph Neural Networks for Predicting and Interpreting Transmissible Hospital-Acquired Infections: A Retrospective Cohort and Simulation Study. *medRxiv*, 2026-05.

Wang, L., Adiga, A., Chen, J., Sadilek, A., Venkatramanan, S., & Marathe, M. (2022, June). Causalgnn: Causal-based graph neural networks for spatio-temporal epidemic forecasting. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 36, No. 11, pp. 12191-12199).

Weltz, J., Volfovsky, A., & Laber, E. B. (2022). Reinforcement learning methods in public health. *Clinical therapeutics*, 44(1), 139-154.

Wu, Y. (2025, November). AI-Enabled Consumer Behavior Prediction and Marketing Strategy Experiments in Digital Transformation. In *Proceedings of the 2025 2nd International Conference on Economic Data Analytics and Artificial Intelligence* (pp. 310-315).

Ye, J., Woods, D., Jordan, N., & Starren, J. (2024). The role of artificial intelligence for the application of integrating electronic health records and patient-generated data in clinical decision support. *AMIA Summits on translational science proceedings*, 2024, 459.

Zeng, D., Cao, Z., & Neill, D. B. (2021). Artificial intelligence-enabled public health surveillance from local detection to global epidemic monitoring and control. In *Artificial intelligence in medicine* (pp. 437-453). Academic Press.

Zhu, X., Zhang, Y., Ying, H., Chi, H., Sun, G., & Zeng, L. (2024). Modeling epidemic dynamics using Graph Attention based Spatial Temporal networks. *Plos one*, 19(7), e0307159.