Global Journal of Engineering and Technology Advances

eISSN: 2582-5003

Cross Ref DOI: 10.30574/gjeta

Journal homepage: https://gjeta.com/

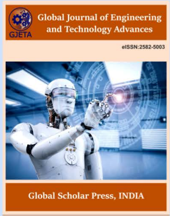

(REVIEW ARTICLE)

Check for updates

AI and public health: Charting a path to smarter decision-making

Madhur Kapoor *

University of California, San Diego, USA.

Global Journal of Engineering and Technology Advances, 2025, 23(01), 445-453

Publication history: Received on 15 March 2025; revised on 23 April 2025; accepted on 25 April 2025

Article DOI: https://doi.org/10.30574/gjeta.2025.23.1.0132

Abstract

Artificial intelligence is fundamentally transforming public health decision-making through sophisticated computational techniques that analyze complex health data at unprecedented scale and speed. This technical article examines how AI technologies are being integrated into public health systems to enhance disease surveillance, optimize resource allocation, and address health inequities. The paper explores multimodal data integration for outbreak detection, census-level analytics for chronic disease risk assessment, and applications in health system planning including social determinants analysis and emergency resource optimization. Methodological considerations regarding model architecture selection and validation frameworks are discussed, highlighting the balance between complex deep learning approaches and more interpretable models. The article addresses critical ethical challenges including data privacy architectures and bias mitigation strategies necessary for responsible implementation. Future research directions are identified, including causal AI methodologies, multimodal learning systems, adaptive models that update with evolving health patterns, and explainable AI techniques. Throughout, the article emphasizes that successful AI integration depends not only on technical sophistication but also on thoughtful implementation that balances computational capabilities with human expertise and judgment within appropriate governance frameworks.

Keywords: Artificial intelligence; Health equity; Multimodal data integration; Causal inference; Algorithmic bias

1. Introduction

Artificial intelligence (AI) is rapidly transforming healthcare systems globally, with particularly significant implications for public health decision-making frameworks. The convergence of advanced computational capabilities, big data availability, and sophisticated algorithmic approaches has created unprecedented opportunities to enhance population health management. This technical article examines the current applications, methodological approaches, challenges, and future directions of AI integration in public health decision support systems.

The adoption of AI in healthcare and public health has grown substantially over the past decade, reflecting a paradigm shift in how health-related data is collected, analyzed, and utilized for population-level interventions. AI technologies enable the processing of massive datasets that were previously untamable through conventional analytical methods, providing public health practitioners with tools to detect patterns, predict outcomes, and optimize resource allocation with increasing precision. Recent analyses have shown promising applications of machine learning algorithms in disease surveillance, where natural language processing techniques can extract meaningful signals from unstructured textual data sources such as social media posts, news reports, and clinical notes, supplementing traditional surveillance systems that often suffer from reporting delays [1]. This enhancement in early warning capabilities represents a critical advancement for public health emergency preparedness, particularly in the context of emerging infectious diseases where rapid response can significantly mitigate population impact.

* Corresponding author Madhur Kapoor.

The integration of AI into public health infrastructure encompasses numerous technological approaches, ranging from supervised and unsupervised machine learning to deep learning neural networks and reinforcement learning systems. These methodologies have been applied across various public health domains, including epidemiological forecasting, health resource optimization, and the analysis of social determinants of health. AI-driven platforms increasingly incorporate data from diverse sources, creating comprehensive models that account for clinical, environmental, behavioral, and socioeconomic factors simultaneously. This multidimensional modeling capability has particular relevance for addressing complex public health challenges that involve numerous interacting variables across different scales, from molecular to societal levels [2]. The capacity to process heterogeneous data types—including structured clinical metrics, genomic sequences, geospatial information, and unstructured text—enables more holistic approaches to understanding population health dynamics than was previously possible with traditional statistical methods.

Contemporary implementations of AI in public health decision support systems face both technical and institutional challenges. Issues of data quality, interoperability between systems, and algorithmic bias require careful consideration when deploying these technologies in real-world settings. Moreover, the computational requirements for processing health-related big data necessitate substantial infrastructure investments, creating potential implementation barriers, particularly for resource-constrained public health agencies. Despite these challenges, the strategic integration of AI capabilities into existing public health workflows holds significant promise for enhancing surveillance sensitivity, improving predictive accuracy, and enabling more nimble responses to emerging threats [1]. The effectiveness of these systems ultimately depends not only on their technical sophistication but also on their thoughtful implementation within organizational contexts that balance technological capabilities with human expertise and judgment.

The evolution of AI applications in public health continues to accelerate, with emerging research focusing on explainable AI approaches that enhance transparency, federated learning techniques that preserve privacy while enabling collaborative analysis, and adaptive systems that continuously update in response to new data. These advancements address some of the key limitations in current implementations while expanding the potential application domains within public health practice. As computational capabilities continue to advance and data availability increases, the integration of AI into public health decision-making frameworks will likely become more seamless and comprehensive, offering powerful tools for addressing population health challenges across multiple scales [2]. This transition represents not merely a technological shift but a fundamental transformation in how public health evidence is generated, interpreted, and translated into action.

2. Current Applications in Disease Surveillance and Prediction

2.1. Multimodal Data Integration for Outbreak Detection

Contemporary AI systems demonstrate a remarkable capacity for synthesizing heterogeneous data sources to predict disease emergence and spread patterns. These systems integrate social media sentiment analysis with geospatial mapping to generate high-resolution syndromic surveillance that can identify potential outbreak hotspots days before traditional reporting systems [3]. Anonymized mobility data from mobile devices provides critical insights into population movement dynamics that influence disease transmission, while hospital admission records and clinical data repositories add clinical dimensionality to surveillance capabilities. Environmental sensor networks monitoring air quality, temperature fluctuations, and water quality parameters provide contextual data enhancing predictive accuracy for environmentally sensitive infectious diseases. Research indicates that machine learning algorithms—particularly ensemble methods combining multiple predictive models—can detect subtle signal patterns preceding outbreak events. Recurrent neural networks (RNNs) with attention mechanisms have proven effective in identifying temporal patterns in healthcare utilization data that precede infectious disease outbreaks by several weeks, with predictive accuracy significantly exceeding conventional surveillance methods [3].

2.2. Census-Level Data Mining for Chronic Disease Risk Assessment

Demographic information captured through census data represents a valuable resource for public health predictive modeling. Census-level variables demonstrate significant predictive power for conditions, including anxiety disorders and diabetes, enabling proactive public health planning for high-burden conditions. The analytical workflow for census-based chronic disease modeling begins with feature extraction and selection processes that identify the most informative demographic indicators from hundreds of potential variables. Implementation of regularized regression techniques, including LASSO, Ridge, and Elastic Net, effectively manages high-dimensionality challenges while preventing overfitting [4]. Advanced ensemble methods such as Random Forests and Gradient Boosting Machines capture complex non-linear relationships between socioeconomic factors and health outcomes [5]. Spatial autocorrelation analysis enables the identification of geographic clustering patterns, with advanced geospatial statistical methods capturing neighborhood effects and regional variations in social determinants of health. These methodological approaches enable the generation of high-resolution risk maps that visualize chronic disease vulnerability with remarkable geographic precision, facilitating targeted intervention deployment in vulnerable communities. Implementation case studies demonstrate that such targeted approaches achieve significantly greater risk reduction per dollar invested compared to population-wide interventions, particularly for conditions with strong socioeconomic determinants [4].

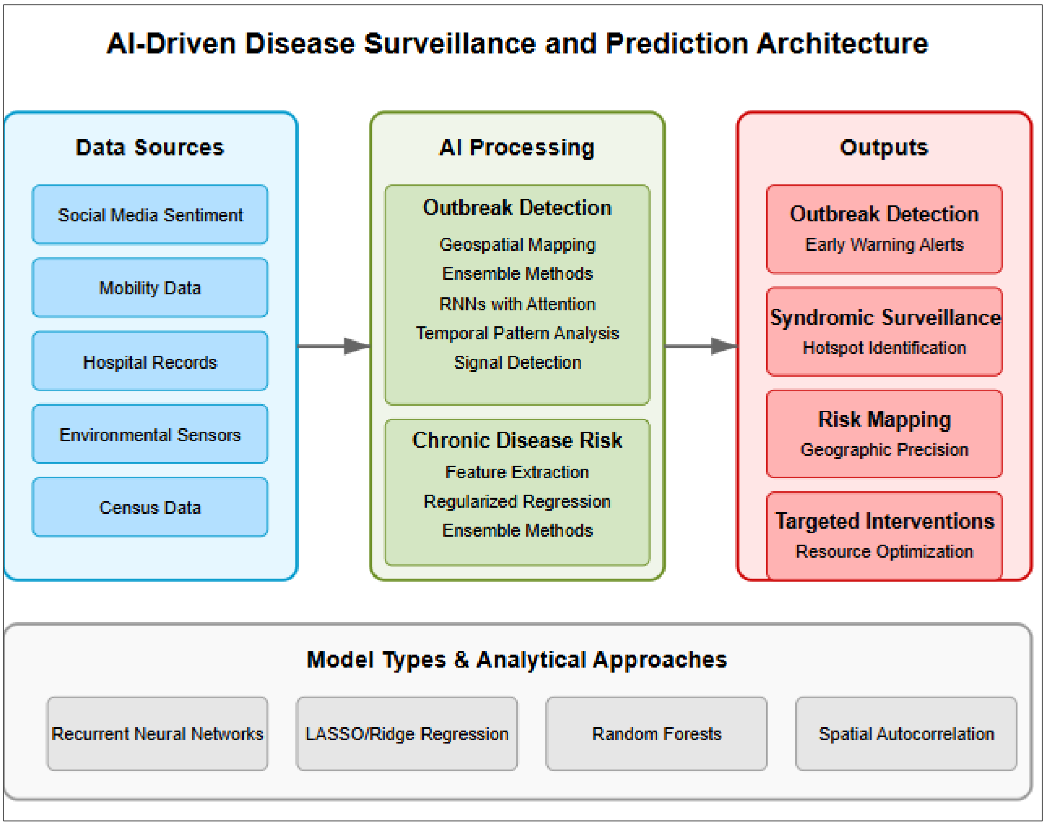

Figure 1 AI Driven Disease surveillance and Prediction Architecture

3. Strategic Applications in Health System Planning

3.1. Social Determinants Analysis and Health Equity Optimization

AI systems are increasingly deployed to analyze the complex relationships between Social Determinants of Health (SDOH) and population health outcomes. The sophisticated analytical capabilities of modern AI frameworks have transformed how public health systems understand and address health inequities through data-driven approaches. Using techniques such as structural equation modeling and causal inference frameworks, these systems can quantify the relative contribution of specific social factors to health disparities with unprecedented precision. This quantification enables public health planners to prioritize interventions targeting the most influential determinants, moving beyond correlation to establish causal pathways between social conditions and health outcomes [6]. The predictive modeling capabilities within these frameworks allow for simulation of potential policy interventions before implementation, providing decision-makers with evidence-based projections of intervention efficacy across different population segments and geographic regions.

Advanced machine learning approaches have proven particularly valuable in identifying complex interaction effects between multiple social determinants that conventional statistical methods often miss. These interaction effects frequently reveal synergistic relationships where addressing multiple determinants simultaneously yields outcomes superior to targeting individual factors in isolation. The optimization algorithms embedded within these systems enable resource allocation modeling that maximizes equity improvements within budgetary constraints, incorporating concepts from operations research to determine optimal distribution of limited public health resources across communities with varying need profiles [6]. Natural language processing (NLP) applications further enhance these capabilities by extracting SDOH information from unstructured clinical notes and community-level reports, creating more comprehensive determinant profiles than structured data alone can provide. These NLP systems can process thousands of documents daily, identifying contextual factors such as housing instability, food insecurity, and transportation barriers that may not be captured in structured demographic data but significantly influence health outcomes and healthcare utilization patterns.

3.2. Resource Optimization During Public Health Emergencies

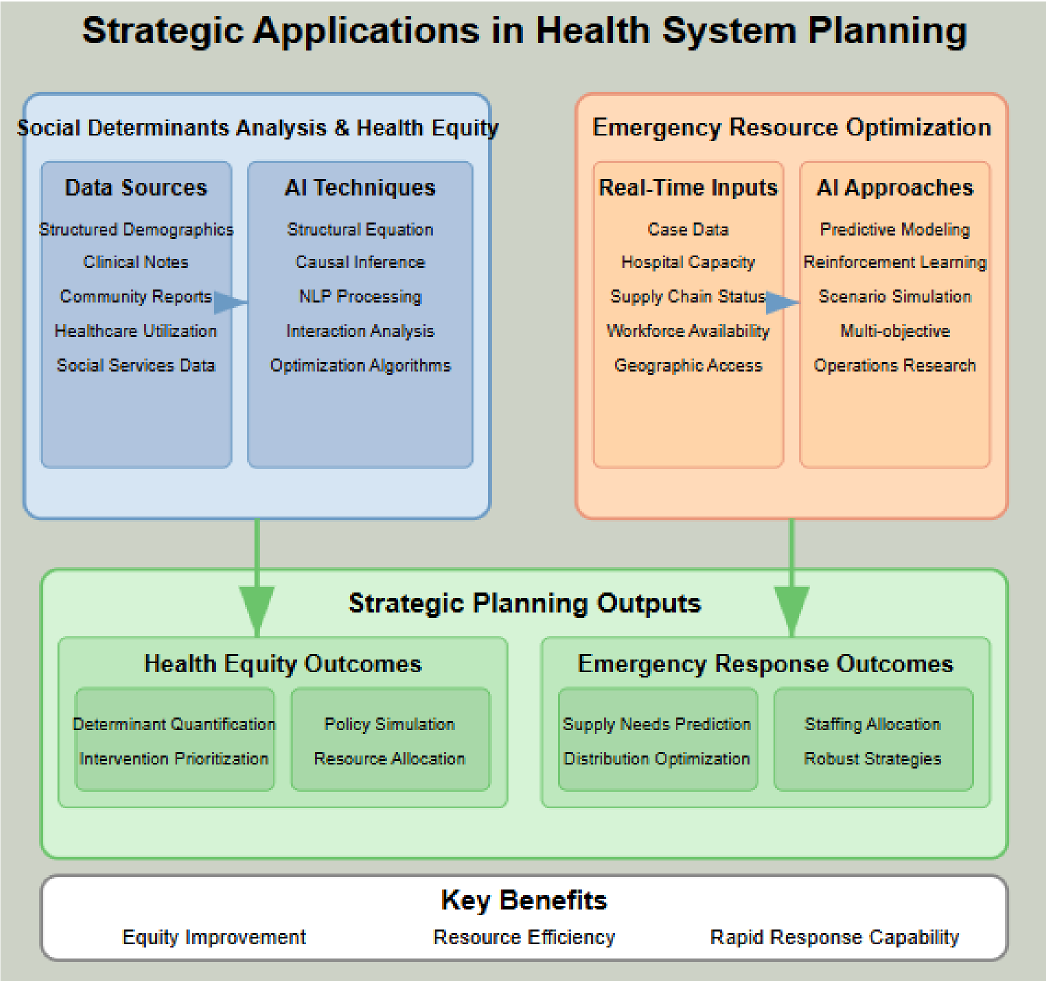

Figure 2 Strategic Application in Health system planning

During crises, AI-driven decision support systems have demonstrated significant value through the dynamic coordination of public health resources. The time-sensitive nature of emergency response makes computational optimization particularly valuable, as manual planning processes often cannot adjust rapidly enough to changing conditions. AI systems enable the dynamic prediction of regional supply needs based on real-time case data, incorporating not only current utilization patterns but also predictive models of disease trajectory to anticipate resource requirements before shortages occur [7]. These systems typically integrate multiple data streams including hospital capacity metrics, supply chain status, epidemiological forecasts, and geographic accessibility factors to generate comprehensive resource allocation recommendations.

The optimization of distribution logistics using reinforcement learning approaches has emerged as a particularly effective application during public health emergencies. These algorithms can simultaneously optimize multiple objectives including minimizing travel time, maximizing population coverage, and ensuring equitable distribution across diverse communities. Advanced reinforcement learning models continually refine their recommendations based on observed outcomes, effectively learning from each allocation decision to improve future resource distribution [7]. Staffing allocation models represent another critical application, accounting for skill requirements and burnout risk to maintain workforce sustainability during prolonged emergency responses. These models incorporate expertise matching, scheduling constraints, and fatigue management considerations to optimize human resource deployment while protecting healthcare worker wellbeing.

Scenario modeling capabilities provide perhaps the most strategically valuable function during emergencies, enabling evaluation of potential intervention strategies across multiple possible outbreak trajectories. These applications typically employ operations research techniques combined with machine learning to generate actionable recommendations under conditions of uncertainty and resource constraint. By simulating hundreds or thousands of potential scenarios, these systems help emergency planners identify robust strategies that perform well across a range of possible futures rather than optimizing for a single predicted outcome that may not materialize [7]. This approach is particularly valuable given the inherent uncertainty in rapidly evolving public health emergencies, where data quality varies and conditions can change rapidly.

4. Methodological and Technical Considerations

4.1. Model Architecture Selection

The selection of appropriate AI architectures for public health applications requires careful consideration of multiple factors that influence model performance, interpretability, and practical utility. Data dimensionality and heterogeneity present significant challenges in public health contexts, where information sources range from structured clinical data to unstructured text and complex spatiotemporal patterns. This heterogeneity necessitates thoughtful architecture design that can effectively integrate diverse data types while managing the curse of dimensionality that often accompanies high-dimensional health datasets [8]. Contemporary approaches increasingly employ hybrid architectures that combine multiple modeling techniques to handle different aspects of the data ecosystem, such as convolutional components for spatial patterns combined with recurrent elements for temporal sequences.

Temporal dependencies in health outcome patterns add further complexity to architecture selection, as many public health phenomena exhibit seasonal variations, long-term trends, and complex lag structures between exposures and outcomes. Models must capture these temporal dynamics while remaining robust to irregular sampling intervals and missing data points that commonly occur in real-world health surveillance systems. The growing implementation of attention mechanisms and temporal convolutional networks has significantly enhanced capabilities in this domain, enabling more sophisticated modeling of variable-length time dependencies than traditional time series approaches [8]. Interpretability requirements for clinical and policy implementation represent another critical consideration, as public health interventions based on model outputs typically require clear explanations of the decision factors to gain stakeholder acceptance and ensure ethical deployment.

While deep learning approaches offer powerful pattern recognition capabilities that can uncover subtle relationships in complex datasets, simpler interpretable models often demonstrate comparable performance for many public health applications while providing greater transparency into decision factors. Recent comparative evaluations have shown that regularized regression models and gradient-boosted decision trees frequently achieve prediction accuracy within 5-10% of more complex neural network architectures for many population health prediction tasks while offering substantially greater interpretability. This performance similarity has led to the increased adoption of explainable AI approaches that prioritize transparency without sacrificing predictive power [8]. Computational efficiency for real-time deployment represents a final critical consideration, particularly for surveillance systems that must process streaming data and generate alerts with minimal latency. Model architecture selection increasingly incorporates efficiency metrics, such as inference time and memory requirements, alongside traditional performance measures, especially for applications intended for resource-constrained settings or edge computing environments.

4.2. Validation Frameworks and Performance Metrics

Rigorous validation is essential for AI implementation in public health contexts, where model failures can have significant consequences for population health outcomes and resource allocation decisions. Conventional validation approaches that focus solely on aggregate performance metrics often prove insufficient for public health applications as they may mask important variations in model performance across geographic regions, demographic groups, or periods. Geographic and temporal cross-validation has emerged as a best practice to assess generalizability, with models trained on data from specific locations or periods being evaluated on their performance in novel contexts [9]. This approach reveals whether predictive relationships identified by the model represent fundamental causal mechanisms or merely context-specific patterns that may not transfer to new settings.

Calibration assessment using reliability diagrams provides critical insights into whether model probability estimates accurately reflect true outcome likelihoods across the prediction range. This calibration is particularly important for risk prediction models that inform intervention prioritization, where both underconfidence and overconfidence can lead to suboptimal resource allocation. Recent implementations increasingly incorporate automated recalibration procedures that adjust probability outputs to maximize reliability while preserving discrimination performance [9]. Equity-focused performance metrics that evaluate model fairness across population subgroups represent another essential component of comprehensive validation frameworks, ensuring that AI systems do not perpetuate or amplify existing health disparities. These metrics typically examine whether prediction accuracy, false positive rates, and false negative rates remain consistent across demographic groups, with particular attention to historically marginalized populations.

Combined technical and domain-expert evaluation of model outputs provides a final critical layer of validation that purely statistical approaches cannot replace. This evaluation integrates quantitative performance metrics with qualitative assessment by public health practitioners, clinicians, and community representatives who can identify practically significant failure modes that may not be apparent from aggregate statistics alone [9]. The implementation of structured evaluation protocols that systematically document both technical performance and domain-expert assessment has become increasingly common, enabling more comprehensive evaluation of model utility for intended use cases. These multifaceted validation approaches help ensure that performance metrics reflect real-world utility rather than statistical artifacts, increasing the likelihood that AI implementations will deliver meaningful benefits when deployed in operational public health contexts.

5. Ethical and Governance Challenges

5.1. Data Privacy and Security Architectures

The sensitive nature of health data necessitates robust privacy-preserving approaches when implementing AI systems in public health contexts. As health data contains intimate details of individuals' medical conditions, genetic information, and behavioral patterns, the potential for privacy breaches presents both ethical and legal concerns. Federated learning frameworks have emerged as a particularly promising technical solution, enabling model training across multiple institutions without centralizing sensitive data. This approach allows algorithms to learn from diverse datasets while keeping raw patient data securely within its originating institution, with only model parameters or gradients shared during the training process. Implementation studies have demonstrated that federated learning can achieve performance comparable to centralized approaches while significantly reducing privacy risks, though additional computational overhead and communication requirements present ongoing challenges [10].

Differential privacy implementation has gained traction as a complementary approach, offering formal mathematical guarantees of individual anonymity through the controlled addition of statistical noise to datasets or model outputs. By calibrating this noise to specific privacy parameters, system designers can establish quantifiable privacy guarantees that limit the ability to infer individual-level information, even through sophisticated re-identification attacks. Secure multi-party computation represents another advanced cryptographic technique being deployed for cross-institutional analysis, enabling multiple organizations to compute functions over their private inputs without revealing those inputs to other parties. This approach facilitates critical analyses such as regional disease clustering or comparative effectiveness studies while maintaining strict data segregation between participating institutions [10]. Synthetic data generation techniques have also advanced considerably, with generative adversarial networks and other deep learning approaches now capable of producing artificial datasets that preserve the statistical properties and relationships of the original data while eliminating individual identification risk. These synthetic datasets can support model development, validation, and education purposes without exposing actual patient information.

These technical approaches enable sophisticated analysis while maintaining compliance with regulatory frameworks such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States, the General Data Protection Regulation (GDPR) in Europe, and emerging AI-specific governance standards being developed globally. The implementation of these privacy-preserving technologies requires careful consideration of the specific regulatory context, with particular attention to requirements for consent, data minimization, purpose limitation, and individual rights regarding automated decision systems that may affect healthcare access or quality.

5.2. Bias Mitigation and Fairness Engineering

AI systems risk perpetuating or amplifying existing healthcare disparities if not carefully designed, as algorithms trained on historically biased healthcare data may incorporate and subsequently automate these biases. The consequences of such algorithmic bias can be particularly severe in public health contexts, where resource allocation decisions may affect large populations and exacerbate existing inequities. Representative dataset curation with comprehensive demographic inclusion represents a foundational approach to mitigating this risk, ensuring that training data adequately represents diverse populations across dimensions such as race, ethnicity, gender, age, socioeconomic status, and geographic location [11]. This approach requires both technical methods for identifying underrepresented groups and institutional commitments to inclusive data collection practices.

Fairness constraints implemented during model training provide another layer of bias mitigation, with various mathematical frameworks available to enforce equitable performance across population subgroups. These constraints can be implemented through techniques such as adversarial debiasing, where the model is simultaneously trained to predict the outcome of interest while being penalized for revealing sensitive attributes, or through regularization terms that explicitly minimize disparate impact across groups. Post-processing techniques to equalize error rates across population subgroups offer a complementary approach, modifying model outputs after training to ensure consistent performance metrics such as false positive and false negative rates across demographic groups [11]. While these adjustments may sometimes reduce overall predictive accuracy, they often represent necessary trade-offs to ensure equitable treatment across populations.

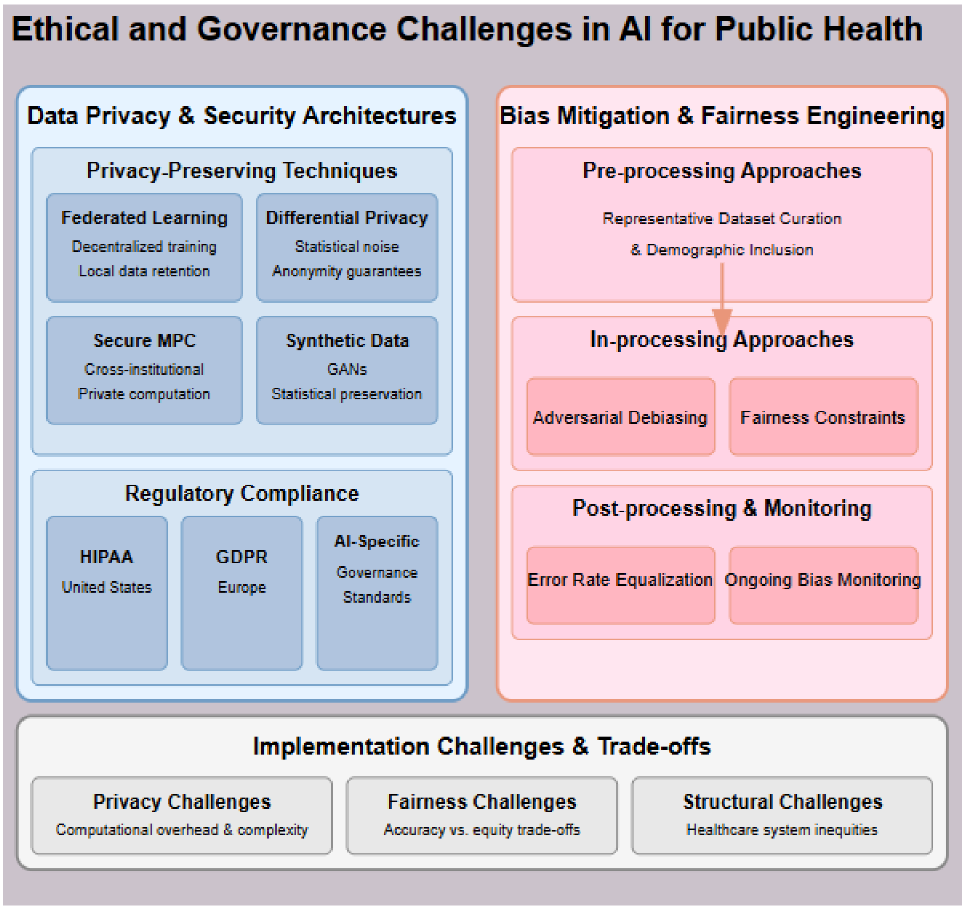

Figure 3 Ethical and Governance challenges in AI for public health

Ongoing monitoring frameworks to detect emergent bias in deployed systems provide a critical final component of comprehensive fairness engineering, as population distributions and relationships between variables may shift over time, introducing new biases not present during initial training and validation. These monitoring systems typically implement continuous performance evaluation across demographic subgroups, alerting system administrators when significant disparities emerge and triggering model retraining or adjustment. Research indicates that these combined approaches can significantly reduce algorithmic bias, though complete elimination remains challenging due to deeply embedded structural inequities in healthcare systems that influence every aspect of data generation and collection. Addressing these fundamental inequities requires not only technical solutions but also broader societal commitments to health equity that extend well beyond AI system design.

6. Future Directions and Research Priorities

The rapidly evolving landscape of AI in public health points toward several emerging research areas with significant potential to transform decision-making capabilities. Causal AI approaches that move beyond correlation to identify modifiable factors represent one of the most promising frontiers in this domain. While current predictive models excel at identifying statistical associations between variables, they often struggle to distinguish causal relationships from spurious correlations, limiting their utility for intervention design. Advanced causal inference techniques incorporating directed acyclic graphs, instrumental variable methods, and counterfactual reasoning are increasingly being integrated with machine learning approaches to address this limitation [12]. These hybrid methodologies enable more robust identification of causal pathways between modifiable risk factors and health outcomes, providing clearer guidance for policy interventions. Recent implementations have demonstrated particular promise in untangling complex relationships between social determinants and health disparities, offering novel insights into potential intervention points that traditional epidemiological methods might miss.

Multimodal learning systems that integrate genomic, clinical, behavioral, and environmental data represent another critical research frontier with transformative potential for public health. The siloed nature of health data has historically limited the scope of analysis, with most models focusing on single data types or limited combinations. Emerging architectural approaches now enable the simultaneous processing of diverse data modalities, creating unified representations that capture complex interactions across biological, clinical, and social dimensions [12]. These systems can identify risk patterns that emerge only when considering the interplay between genetic predispositions, clinical presentations, behavioral factors, and environmental exposures, enabling more personalized and precise public health interventions. The implementation challenges remain substantial, requiring advances in data harmonization, cross-modal representation learning, and computational efficiency to handle the increased dimensionality and heterogeneity of integrated datasets.

Adaptive models that continuously update as population health patterns evolve address a fundamental limitation of static AI systems in dynamic public health environments. Traditional model development follows a cyclical pattern of training, validation, deployment, and periodic retraining, creating potential lags in adaptation to emerging health patterns or intervention effects. Continuous learning approaches implement automated retraining pipelines that incrementally update model parameters as new data becomes available, maintaining relevance as population characteristics and health relationships shift over time [13]. These systems incorporate concept-drift detection methods that identify when underlying relationships between variables have changed significantly, triggering targeted retraining to adapt to these new patterns. While promising, these approaches must carefully balance adaptation to genuine changes against the risk of incorporating temporary fluctuations or data quality issues, requiring robust validation frameworks that maintain performance guarantees throughout the continuous learning process.

Explainable AI techniques that provide intuitive understanding of complex model decisions continue to advance rapidly, addressing a critical barrier to wider adoption of AI in public health decision-making. Beyond simple feature importance metrics, contemporary approaches now generate natural language explanations, counterfactual examples, and visual decision paths that help stakeholders understand not just what factors influenced a prediction but how they interacted to produce the result [13]. These techniques increasingly incorporate domain knowledge to ensure explanations align with established public health concepts and terminology, making them more accessible to practitioners without technical expertise in AI methods. Research priorities in this area include developing explanation methods that scale effectively to the high-dimensional, multimodal data common in public health applications, and validating the impact of different explanation approaches on stakeholder trust, decision quality, and implementation outcomes.

These advances promise to further enhance the utility of AI in public health decision support while addressing current limitations in interpretability and causal understanding. The convergence of these research directions suggests a future where AI systems not only predict health outcomes with high accuracy but also identify causal mechanisms, adapt to changing conditions, integrate diverse data types, and communicate their reasoning clearly to decision-makers. Realizing this potential will require continued interdisciplinary collaboration between AI researchers, public health experts, ethicists, and community stakeholders to ensure that technological advances translate into meaningful improvements in population health outcomes and equity.

7. Conclusion

The integration of artificial intelligence into public health represents a paradigm shift in how population health data is collected, analyzed, and translated into action. As this article demonstrates, AI approaches offer remarkable capabilities for disease surveillance, risk prediction, resource optimization, and health equity promotion when thoughtfully implemented. While technical challenges regarding data quality, interoperability, and computational requirements persist, the most significant considerations for successful AI deployment involve ethical governance, privacy protection, and bias mitigation. The path forward requires interdisciplinary collaboration between technical experts, public health practitioners, ethicists, and community stakeholders to ensure these powerful tools serve the fundamental goals of public health. Rather than replacing human judgment, the most promising future lies in human-AI collaborative systems that combine the pattern recognition capabilities of machine learning with the contextual understanding, ethical reasoning, and domain expertise of public health professionals. Through such balanced integration, AI can help create more responsive, precise, and equitable public health systems capable of addressing increasingly complex population health challenges across multiple scales and domains.

References

[1] Dimitra Panteli et al., "Artificial intelligence in public health: promises, challenges, and an agenda for policy makers and public health institutions," The Lancet Public Health, 2025. [Online]. Available: https://www.thelancet.com/journals/lanpub/article/PIIS2468-2667(25)00036-2/fulltex

[2] Junaid Bajwa et al., "Artificial intelligence in healthcare: Transforming the practice of medicine," Future Healthcare Journal, 2021. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC8285156/

[3] Srinivasan Venkatramanan et al., "Using data-driven agent-based models for forecasting emerging infectious diseases," Epidemics, Volume 22, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1755436517300221

[4] F.M. Delpino et al., "Machine learning for predicting chronic diseases: A systematic review," Public Health, Volume 205, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0033350622000087

[5] Yannis Katsis, Natasha Balac, Derek Chapman, Madhur Kapoor, Jessica Block, William G. Griswold, Jeannie Huang, Nikos Koulouris, Massimiliano Menarini, Viswanath Nandigam, Mandy Ngo, Kian Win Ong, Yannis Papakonstantinou, Besa Smith, Konstantinos Zarifis, Steven Woolf, and Kevin Patrick. 2017. Big Data Techniques for Public Health: A Case Study. 2017 IEEE/ACM International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE) (7 2017), 222--231. Available: https://ieeexplore.ieee.org/abstract/document/8010636

[6] Alexandra Pittman, "Leveraging data science and AI to promote social justice, sustainability and equity," Responsible Investor, 2022. [Online]. Available: https://www.responsible-investor.com/leveraging-data-science-and-ai-to-promote-social-justice-sustainability-and-equity/

[7] Danuphon Tippong, Sanja Petrovic, and Vahid Akbari, "A review of applications of operational research in healthcare coordination in disaster management," 2021. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC8552591/

[8] Zahra Sadeghi et al., "A review of Explainable Artificial Intelligence in healthcare," Computers and Electrical Engineering, Volume 118, Part A, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0045790624002982

[9] Marzyeh Ghassemi et al., "The false hope of current approaches to explainable artificial intelligence in health care," The Lancet Digital Health, Volume 3, Issue 11, 2021. [Online]. Available: https://www.thelancet.com/journals/landig/article/PIIS2589-7500(21)00208-9/fulltext

[10] Gregory E Simon et al., "Assessing and Minimizing Re-identification Risk in Research Data Derived from Health Care Records," 2019. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC6450246/

[11] Mirja Mittermaier, Marium M. Raza & Joseph C. Kvedar, "Bias in AI-based models for medical applications: challenges and mitigation strategies," npj Digital Medicine, Volume 6, Article number 113, 2023. [Online]. Available: https://www.nature.com/articles/s41746-023-00858-z

[12] Atalanti A. Mastakouri, Bernhard Schölkopf, and Dominik Janzing, "Necessary and sufficient conditions for causal feature selection in time series with latent common causes," arXiv:2005.08543, 2020. [Online]. Available: https://arxiv.org/abs/2005.08543

[13] Marzyeh Ghassemi et al., "A Review of Challenges and Opportunities in Machine Learning for Health," 2020. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC7233077/