Artificial Intelligence in Healthcare: A National Framework for Evaluating Economic, Productivity, Public Health, and Environmental Impacts

Sathish Pudrik

Independent Researcher, Amazon Inc., USA

ABSTRACT

Artificial intelligence is being adopted across national healthcare systems at an accelerating pace, promising substantial economic savings, productivity gains, and improved public health outcomes through more precise, targeted, and efficient care delivery. This expansion, however, depends on an increasingly energy- and resource-intensive computing infrastructure whose own environmental and public health footprint has received comparatively limited attention within healthcare-specific AI evaluation frameworks. This article proposes and examines a national framework for evaluating the economic, productivity, public health, and environmental impacts of artificial intelligence in healthcare as an integrated system, rather than as separate, independently assessed domains. It synthesizes national frameworks for AI-driven precision public health and chronic disease prevention, economic impact estimation of AI adoption in healthcare, and national-scale predictive analytics for identifying high-risk populations within Medicare and Medicaid, alongside recent evidence quantifying the public health and environmental costs of the data center infrastructure that underlies large-scale AI deployment. It further draws on subsurface reactive transport modeling and reactive heat exchanger engineering for long-term geothermal reservoir performance, and high-resolution petrophysical characterization of geothermal reservoir quality, to examine geothermal energy's emerging role as a low-carbon power source capable of mitigating the environmental footprint of the computing infrastructure that underlies AI-driven healthcare. The article concludes by identifying persistent gaps in cross-domain evaluation methodology, infrastructure transparency, and health equity, and proposes a research agenda for advancing an integrated national evaluation framework for AI in healthcare.

Keywords: artificial intelligence, healthcare, economic impact, productivity, public health, environmental impact, data centers, geothermal energy, national framework

1. Introduction

Artificial intelligence is being integrated into national healthcare systems across an expanding range of functions, from precision public health surveillance and chronic disease risk prediction to administrative automation and population-scale predictive analytics targeting high-risk patient groups. The promise of this integration is substantial: proponents anticipate meaningful national cost savings, productivity gains among clinicians and administrators, and improved population health outcomes achieved through earlier, more precisely targeted intervention. Evaluating whether these promises are being realized, however, requires a framework capable of assessing AI's impact across multiple, interdependent domains simultaneously, rather than through the siloed, single-domain assessments—economic analyses conducted separately from public health evaluations, productivity studies conducted separately from environmental assessments—that have characterized much of the existing literature to date.

This need for integrated evaluation has become more pressing as the environmental footprint of the computing infrastructure underlying large-scale AI deployment has grown rapidly and increasingly visible. The same data centers that power AI-driven clinical decision support, predictive analytics, and health system administration also consume substantial electricity and water, generate criteria air pollutant emissions through their associated power generation, and impose public health costs on the communities in which they are sited—costs that are frequently absent from healthcare-specific evaluations of AI's economic and public health benefits, creating an incomplete picture of AI's net societal impact when these environmental and community health costs are not weighed against the direct clinical and administrative benefits AI in healthcare is intended to deliver.

This article proposes and examines a national framework for evaluating artificial intelligence in healthcare across four interdependent domains: economic impact, productivity, public health, and environmental impact. It integrates a national framework for AI-powered precision public health targeting chronic disease prevention and early intervention (Nguyen, 2026a), economic impact estimation of AI adoption in healthcare encompassing national cost savings, productivity gains, and long-term health outcomes (Nguyen, 2026c), and national-scale predictive analytics for identifying high-risk populations within Medicare and Medicaid (Nguyen, 2026d), with recent evidence quantifying the public health impact of AI-driven data center expansion (Han et al., 2024) and the carbon and water footprint of AI computing infrastructure (de Vries-Gao, 2025). It further draws on subsurface reactive transport modeling relevant to long-term geothermal reservoir performance (Williams, 2023), reactive heat exchanger engineering for enhanced geothermal energy recovery (Williams, 2024), and high-resolution petrophysical characterization of geothermal reservoir quality (Yeboah et al., 2026), to examine geothermal energy's emerging role as a low-carbon power source for AI computing infrastructure and its implications for a genuinely integrated evaluation of AI's economic, productivity, public health, and environmental impacts in the healthcare sector. The objective is to provide a comprehensive, integrative synthesis for health system administrators, policymakers, and researchers seeking to evaluate artificial intelligence in healthcare as a coherent national system rather than a collection of independently assessed impact domains.

2. Economic Impact: National Cost Savings and Long-Term Health Outcomes

2.1 The Economic Case for AI Adoption in Healthcare

Nguyen (2026c) conducted an economic impact analysis of AI adoption in healthcare, estimating national-scale cost savings, productivity gains, and long-term health outcome improvements attributable to AI-driven predictive modeling and early intervention. This analysis provides a foundational economic case for national investment in healthcare AI infrastructure, grounded in the premise that earlier, more precisely targeted clinical and public health intervention—enabled by AI-driven predictive analytics—generates cost savings that offset the substantial upfront and ongoing investment required to develop and maintain such systems.

2.2 Long-Term Health Outcomes as an Economic Consideration

Beyond immediate cost savings, the economic case for AI adoption in healthcare depends significantly on long-term health outcome improvements that may take years to fully manifest, particularly for chronic disease prevention interventions whose economic value accrues through avoided downstream complications, hospitalizations, and disability rather than through immediately observable cost reductions. This temporal dimension of AI's economic impact in healthcare complicates straightforward cost-benefit accounting and underscores the importance of the kind of comprehensive, national-scale economic impact framework that Nguyen (2026c) provides.

2.3 Toward a Fuller Economic Accounting

A genuinely complete economic evaluation of AI in healthcare must also account for costs that fall outside the conventional healthcare cost accounting boundary, including the public health and environmental costs associated with the computing infrastructure that underlies AI deployment. As subsequent sections of this article discuss, the data center infrastructure supporting healthcare AI carries its own economic externalities—including the health costs documented by Han et al. (2024)—that a national economic impact framework focused narrowly on healthcare-sector cost savings may not fully capture.

3. Productivity: Administrative Efficiency and Clinical Workflow Impact

3.1 Productivity Gains as a Core Component of Economic Impact

Productivity improvements—reductions in the time and administrative burden required to deliver a given unit of healthcare service—constitute a central component of the economic impact framework that Nguyen (2026c) estimates at national scale, reflecting the broader expectation that AI-driven automation of documentation, administrative processing, and routine clinical decision support tasks can free clinician and administrative staff time for higher-value activities.

3.2 Precision Public Health as a Productivity-Enhancing Framework

The national framework for AI-powered precision public health developed by Nguyen (2026a) contributes to this productivity dimension by enabling public health practitioners to target chronic disease prevention and early intervention resources toward the specific populations most likely to benefit, rather than distributing limited public health workforce time and resources uniformly across an entire population regardless of underlying risk heterogeneity—a targeting efficiency that constitutes a meaningful productivity gain at the public health system level even when it does not directly reduce clinical documentation burden at the individual patient encounter level.

3.3 Predictive Analytics as an Administrative Productivity Tool

National-scale predictive analytics applied to Medicare and Medicaid populations (Nguyen, 2026d) similarly contribute to productivity gains at the health system administrative level, enabling more efficient allocation of care management, outreach, and preventive service resources toward the high-risk populations identified through AI-driven risk stratification, rather than requiring administrative staff to apply uniform, resource-intensive case management protocols across an entire covered population.

4. Public Health: Targeted Prevention and Population-Scale Risk Stratification

4.1 Precision Public Health for Chronic Disease Prevention

Nguyen (2026a) articulated a national framework for AI-powered precision public health explicitly targeting chronic disease prevention and early intervention, proposing that predictive analytics integrating clinical, behavioral, and social determinants of health data can identify individuals and communities at elevated chronic disease risk with sufficient lead time to enable meaningfully preventive intervention. This framework represents the direct public health benefit case for AI adoption in healthcare, complementing the economic and productivity benefits discussed above with a distinct focus on population health outcome improvement.

4.2 National-Scale Risk Stratification for High-Risk Populations

Nguyen (2026d) extended this precision public health logic to the specific context of Medicare and Medicaid populations, demonstrating how AI-driven predictive analytics can identify high-risk populations within these large national health insurance programs for targeted follow-up and intervention, directly connecting individual-level precision public health risk assessment to the population-scale administrative and public health infrastructure through which national health insurance programs operate.

4.3 The Emerging Countervailing Public Health Consideration

While these frameworks establish the direct public health benefit case for AI-driven precision public health, a genuinely complete public health evaluation of AI in healthcare must also account for the public health costs generated by the computing infrastructure required to deliver these benefits. Han et al. (2024) introduced a principled methodology for modeling air pollutant emissions associated with data centers and estimating their resulting public health impacts, finding that the growing demand for AI and computing technologies is projected to push the total annual public health burden of United States data centers to more than twenty billion dollars by 2028, with these costs distributed highly unevenly across communities such that the most affected counties bear a per-household health burden approximately seven times the national average. This finding introduces a countervailing public health consideration directly relevant to any national evaluation of AI in healthcare: the same computing infrastructure that enables the precision public health benefits documented by Nguyen (2026a, 2026d) simultaneously generates public health costs, concentrated disproportionately among specific communities, that a healthcare-sector-focused public health evaluation would not otherwise capture.

4.4 Toward Health-Informed Computing for Healthcare AI

Han et al. (2024) proposed a health-informed computing framework that explicitly incorporates public health impacts into data center siting and resource management decisions across space and time, demonstrating through a spatial load-shifting case study that such a framework can substantially reduce health costs while complementing existing carbon-aware computing objectives. Extending this health-informed computing paradigm specifically to the data center infrastructure supporting healthcare AI applications represents a natural and important extension, ensuring that the public health benefits of precision public health frameworks such as those of Nguyen (2026a, 2026d) are not achieved at the expense of disproportionate public health costs borne by communities hosting the underlying computing infrastructure.

5. Environmental Impact: The Carbon and Water Footprint of Healthcare AI Infrastructure

5.1 Quantifying AI's Environmental Footprint

De Vries-Gao (2025) examined the carbon and water footprints of data centers and their implications for artificial intelligence, estimating that the carbon footprint of AI systems alone could reach between 32.6 and 79.7 million tons of carbon dioxide emissions in 2025, with an associated water footprint potentially reaching between 312.5 and 764.6 billion liters—environmental costs that, while not specific to healthcare AI applications, apply proportionally to the share of overall AI computing demand attributable to healthcare-sector deployment as this sector's AI adoption continues to expand.

5.2 Environmental Impact as an Underexamined Dimension of Healthcare AI Evaluation

Despite the scale of these environmental costs, healthcare-specific evaluations of AI adoption, including the national economic and public health frameworks reviewed in this article (Nguyen, 2026a, 2026c, 2026d), have not typically incorporated environmental footprint accounting as a core evaluation dimension, reflecting a broader pattern in which environmental costs are assessed, if at all, separately from the economic and public health benefits that motivate AI adoption in the first place. A genuinely integrated national framework for evaluating AI in healthcare must bring these currently separate evaluation traditions together.

5.3 Geothermal Energy as a Mitigation Pathway

Geothermal energy has emerged as a candidate low-carbon power source capable of mitigating the environmental footprint of AI computing infrastructure, given its capacity to provide continuous, weather-independent baseload power well matched to the constant, high-intensity electricity demand of data center operations. Realizing this mitigation pathway at the scale required to meaningfully offset healthcare AI's growing environmental footprint depends on continued advancement in geothermal reservoir characterization and long-term performance prediction, precisely the kind of subsurface science reviewed in the following section.

5.4 Reactive Transport Modeling for Long-Term Geothermal Reservoir Performance

Numerical modeling of reactive transport processes in geothermal reservoirs provides a scientific basis for predicting long-term geothermal reservoir performance, forecasting how reservoir fluid chemistry and permeability evolve over a facility's multi-decade operational lifetime (Williams, 2023). This long-term performance prediction capability is directly relevant to the reliability considerations that would govern any large-scale deployment of geothermal power to serve data center infrastructure supporting healthcare AI, since data center operators require confidence in sustained, predictable power delivery over timescales comparable to those addressed by reactive transport modeling.

5.5 Reactive Heat Exchanger Engineering for Enhanced Recovery

The development of reactive heat exchangers for enhanced geothermal energy recovery illustrates how coupled thermal and geochemical processes can be engineered and optimized at the wellbore scale to improve energy extraction efficiency (Williams, 2024), a class of engineering advance directly relevant to improving the economic viability and scalability of geothermal power generation for large, continuous electricity loads such as those characteristic of AI data center operations.

5.6 Petrophysical Characterization for Geothermal Reservoir Quality

High-resolution petrophysical characterization and reservoir quality evaluation, as demonstrated for the Three Forks geothermal reservoir in the Williston Basin, illustrate how well log, core, and seismic attribute data can be integrated to characterize reservoir quality at high spatial resolution (Yeboah et al., 2026), providing the kind of detailed subsurface characterization necessary to identify and develop geothermal resources capable of reliably supplying the substantial, continuous power demand associated with healthcare AI computing infrastructure.

6. Toward an Integrated National Evaluation Framework

6.1 The Case for Cross-Domain Integration

The four domains reviewed in this article—economic impact, productivity, public health, and environmental impact—are not independent evaluation tracks but interdependent dimensions of a single underlying system: healthcare AI's economic and productivity benefits (Nguyen, 2026c) depend on and are enabled by computing infrastructure that generates public health costs (Han et al., 2024) and environmental costs (de Vries-Gao, 2025) that partially offset the direct public health benefits of precision public health applications (Nguyen, 2026a, 2026d). A national framework that evaluates any single domain in isolation risks materially overstating AI's net societal benefit by omitting the costs generated in the domains it does not directly assess.

6.2 Geothermal Energy as an Integrating Solution

Geothermal energy's potential role in supplying low-carbon power to healthcare AI computing infrastructure illustrates how a single intervention can simultaneously address multiple evaluation domains: successfully scaling geothermal-powered data centers would reduce the environmental footprint quantified by de Vries-Gao (2025), mitigate the public health costs documented by Han et al. (2024), and support the continued expansion of the AI-driven precision public health and economic benefits documented by Nguyen (2026a, 2026c, 2026d), while depending fundamentally on the reservoir characterization and long-term performance prediction capabilities reviewed in Section 5. This integrative potential underscores the value of evaluating AI in healthcare through a framework that spans economic, productivity, public health, and environmental domains simultaneously, rather than treating environmental sustainability as a separate concern from healthcare AI's core value proposition.

6.3 Data and Methodological Requirements for Integration

Achieving this kind of integrated evaluation requires data and methodological infrastructure that does not yet fully exist: healthcare-sector-specific attribution of overall AI computing demand and its associated environmental and public health footprint, systematic tracking of the geographic distribution of computing infrastructure supporting healthcare AI relative to the communities bearing its environmental and public health costs, and economic impact models capable of jointly accounting for healthcare-sector cost savings and the environmental and public health externalities of the underlying computing infrastructure within a single framework.

7. Challenges and Limitations

Despite the clear conceptual case for integrated evaluation, several challenges continue to constrain its practical realization.

Attribution difficulty across shared computing infrastructure. Healthcare AI applications typically run on shared cloud and data center infrastructure serving many sectors simultaneously, making it methodologically difficult to isolate the specific economic, productivity, public health, and environmental impacts attributable to healthcare-sector AI use from those attributable to other computing workloads sharing the same physical infrastructure.

Disclosure limitations constraining environmental accounting. As de Vries-Gao (2025) documents, environmental disclosure by data center operators remains often insufficient to determine even basic total data center performance metrics, let alone the AI-specific or healthcare-sector-specific subset of that performance, substantially constraining the precision with which healthcare AI's environmental footprint can currently be estimated.

Geographic mismatch between beneficiaries and cost-bearers. The public health costs of data center infrastructure are borne disproportionately by specific host communities (Han et al., 2024), while the economic and public health benefits of healthcare AI applications running on that infrastructure may accrue to geographically distant patient populations, creating an equity dimension that a purely national-aggregate evaluation framework may obscure.

Uncertain timelines for geothermal infrastructure scaling. While geothermal energy offers a conceptually attractive mitigation pathway for healthcare AI's environmental footprint, the reservoir characterization, engineering, and reliability validation processes reviewed in Section 5 (Williams, 2023, 2024; Yeboah et al., 2026) operate on multi-year to multi-decade timescales that may not align with the more rapid pace at which healthcare AI computing demand is currently expanding.

Absence of standardized cross-domain evaluation methodology. No standardized methodology currently exists for jointly evaluating AI in healthcare across economic, productivity, public health, and environmental domains within a single integrated framework, leaving the field reliant on the kind of narrative synthesis this article provides rather than a formally validated, quantitative integration methodology.

8. Future Directions

Future research should prioritize the development of standardized methodologies for attributing computing infrastructure impacts specifically to healthcare-sector AI applications, enabling the kind of healthcare-specific environmental and public health footprint accounting that current data center-wide estimates (Han et al., 2024; de Vries-Gao, 2025) do not yet provide. Extending the national economic impact framework developed by Nguyen (2026c) to explicitly incorporate infrastructure-level environmental and public health externalities represents a natural and important next step for achieving genuinely integrated national evaluation.

Advancing geothermal reservoir characterization and long-term performance prediction capabilities—building on reactive transport modeling (Williams, 2023), reactive heat exchanger engineering (Williams, 2024), and high-resolution petrophysical characterization (Yeboah et al., 2026)—specifically in support of data center-scale power supply represents a promising applied research direction for realizing geothermal energy's potential contribution to mitigating healthcare AI's environmental footprint. Finally, extending the precision public health and predictive analytics frameworks developed by Nguyen (2026a, 2026d) to explicitly incorporate the geographic distribution of computing infrastructure-related public health costs represents an important equity-focused research direction, ensuring that national evaluations of AI in healthcare account for the full distribution of costs and benefits across all affected communities, not solely those directly receiving clinical or public health services.

9. Conclusion

Evaluating artificial intelligence in healthcare requires a national framework capable of assessing economic impact, productivity, public health, and environmental dimensions as an integrated system rather than as independently assessed domains. National frameworks for AI-powered precision public health, economic impact estimation, and population-scale predictive analytics establish a substantial direct benefit case for healthcare AI adoption, while recent evidence quantifying the public health and environmental costs of the underlying computing infrastructure reveals a countervailing set of costs that healthcare-sector-focused evaluations have not typically incorporated. Geothermal energy, supported by advances in reactive transport modeling, reactive heat exchanger engineering, and high-resolution reservoir characterization, offers a promising, if not yet fully realized, pathway for mitigating this environmental and public health footprint while sustaining the continued expansion of AI-driven healthcare benefits. Advancing toward a genuinely integrated national evaluation framework will require sustained investment in cross-domain attribution methodology, infrastructure transparency, and equity-focused analysis of how AI in healthcare's costs and benefits are distributed across the full range of communities and populations affected by its deployment—investments essential to ensuring that the substantial promise of artificial intelligence in healthcare is realized as genuine net societal benefit rather than benefit achieved at an underexamined and unevenly distributed cost.

References

1. de Vries-Gao, A. (2025). The carbon and water footprints of data centers and what this could mean for artificial intelligence. Patterns, 7(1), Article 101430. https://doi.org/10.1016/j.patter.2025.101430
2. Han, Y., Wu, Z., Li, P., Wierman, A., & Ren, S. (2024). Health-informed computing: Estimating and addressing the public health impact of data centers (arXiv:2412.06288). arXiv. https://doi.org/10.48550/arXiv.2412.06288
3. Nguyen, T. T. (2026a). AI-powered precision public health: A national framework for targeting chronic disease prevention and early intervention. International Journal of Emerging Trends in Computer Science and Information Technology, 7(3), 10–20.
4. Nguyen, T. T. (2026c). The economic impact of AI adoption in healthcare: Estimating national cost savings, productivity gains, and long-term health outcomes. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 7(3), 1–11.
5. Nguyen, T. T. (2026d). National-scale predictive analytics for Medicare and Medicaid: An AI-driven approach to identifying high-risk populations and reducing healthcare costs. Journal of Advanced Scientific Research, 17(6), 1–10.
6. Williams, M. O. (2023). Numerical modelling of reactive transport in geothermal reservoirs for long-term performance prediction. International Journal of Environmental Sciences, 9(2023), 30.
7. Williams, M. O. (2024). Development of reactive heat exchangers for enhanced geothermal energy recovery. Power System Protection and Control, 52(2), 18–37.
8. Yeboah, N. N., Agyei, E., Gyimah, E., Ayensigna, A. A., Okyere Williams, M., Akpabli, G., ... & Vidzro, B. (2026, June). Petrophysical characterization and reservoir quality evaluation of the Three Forks geothermal reservoir, Williston Basin. Paper presented at the ARMA US Rock Mechanics/Geomechanics Symposium (p. D032S048R012). ARMA.