

## **Evaluating the National Value of AI in Healthcare: Economic Savings, Productivity Enhancement, Long-Term Health Outcomes, and Environmental Impact**

**Halil Kose**

Independent Researcher, Netherland

### **Abstract**

Artificial intelligence is increasingly positioned as a transformative force for national healthcare systems, promising economic savings through more efficient resource allocation, productivity enhancement through automation of administrative and clinical documentation tasks, improved long-term health outcomes through precision prevention and early intervention, and, less visibly, carrying its own environmental footprint through the computing infrastructure that underlies its deployment. Evaluating AI's true national value requires weighing these four dimensions together rather than assessing any one in isolation, and requires holding genuine benefits alongside documented risks, including instances in which AI-driven administrative tools have themselves generated new burdens rather than reducing existing ones. This article proposes and examines a national evaluation framework spanning economic savings, productivity enhancement, long-term health outcomes, and environmental impact, synthesizing national frameworks for AI-powered precision public health and AI-driven prior authorization transformation, empirical evidence on ambient AI scribe technology's impact on clinical documentation burden, critical evidence on AI-assisted prior authorization denial practices in Medicare Advantage, and recent quantification of the public health and environmental costs of AI computing infrastructure. It further draws on high-resolution petrophysical characterization of geothermal reservoir quality and machine learning-enhanced geomechanical characterization of wellbore stability to examine the subsurface science underlying geothermal energy's emerging role in mitigating the environmental footprint of AI infrastructure. The article concludes by identifying persistent gaps in cross-domain evaluation methodology, regulatory oversight, and infrastructure sustainability, and proposes a research agenda for advancing a genuinely integrated national evaluation of AI's value in healthcare.

**Keywords:** artificial intelligence, healthcare, national value, economic savings, productivity, health outcomes, environmental impact, prior authorization, geothermal energy

## 1. Introduction

The national value proposition for artificial intelligence in healthcare rests on a combination of anticipated economic savings, productivity gains among clinicians and administrators, and improved long-term population health outcomes achieved through more precise, earlier, and better-targeted intervention. These anticipated benefits have driven substantial national investment in AI-powered clinical decision support, administrative automation, and population health analytics across health systems and payers alike. Realizing genuine national value from this investment, however, requires an evaluation framework capable of assessing these benefits honestly against their costs and risks, including administrative and regulatory risks that have emerged as AI tools have moved from pilot demonstration into large-scale operational deployment, and environmental costs associated with the computing infrastructure that AI-driven healthcare applications depend upon.

Nowhere is this need for honest, balanced evaluation more apparent than in the domain of AI-driven administrative automation, where the same technology promoted as a solution to administrative burden has, in specific documented instances, become a source of new burden and harm when deployed without adequate transparency, clinical oversight, or regulatory safeguard. Prior authorization, the utilization management process through which health insurers require advance approval before covering certain services, exemplifies this duality: AI-driven prior authorization tools hold genuine promise for reducing the substantial administrative time burden that manual prior authorization processes impose on clinicians and patients, while recent evidence documents instances in which AI-assisted prior authorization review has instead been associated with increased claim denials, reduced transparency, and mounting legal and regulatory scrutiny.

This article proposes and examines a national framework for evaluating artificial intelligence's value in healthcare across four domains—economic savings, productivity enhancement, long-term health outcomes, and environmental impact—integrating a national framework for AI-powered precision public health targeting chronic disease prevention and early intervention (Nguyen, 2026a), a national framework for transforming prior authorization through artificial intelligence to reduce administrative burden and improve patient access (Nguyen, 2026b), and

economic impact estimation of AI adoption in healthcare (Nguyen, 2026c), with empirical evidence on ambient AI scribe technology's demonstrated impact on clinical documentation burden and efficiency (Duggan et al., 2025) and critical evidence documenting the risks that AI-assisted prior authorization review has introduced into the Medicare Advantage program (Raza et al., 2026). It further draws on recent quantification of 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), alongside high-resolution petrophysical characterization of geothermal reservoir quality and machine learning-enhanced geomechanical characterization of wellbore stability, to examine geothermal energy's emerging role in mitigating the environmental footprint of the computing infrastructure that underlies national-scale AI deployment in healthcare. The objective is to provide a comprehensive, honestly balanced synthesis for health system administrators, policymakers, and researchers seeking to evaluate the true national value of artificial intelligence in healthcare.

## **2. Economic Savings: The National Cost-Benefit Case**

### **2.1 Estimating National-Scale Economic Impact**

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 the foundational economic case underlying much of the current national investment in healthcare AI infrastructure, premised on the expectation that earlier, more precisely targeted intervention generates savings that outweigh the substantial costs of developing, deploying, and maintaining AI systems at scale.

### **2.2 Administrative Cost Savings Through Prior Authorization Transformation**

A significant component of AI's potential national economic value lies in administrative cost reduction, particularly within the prior authorization process, which imposes substantial direct costs on both payers and providers. Nguyen (2026b) developed a national framework for transforming prior authorization through artificial intelligence, explicitly targeting the reduction of administrative burden and improvement of patient access as core economic and operational objectives, reflecting the broader recognition that prior authorization administration

represents one of the largest sources of avoidable administrative cost within the United States healthcare system.

### 2.3 The Necessity of Weighing Costs Against Realized Risks

A genuinely complete economic evaluation must weigh these anticipated savings against the costs generated when AI-driven administrative tools fail to perform as intended or are deployed without adequate safeguards. As Section 4 of this article discusses, documented instances of AI-assisted prior authorization tools contributing to increased claim denials and subsequent legal and regulatory response (Raza et al., 2026) generate their own economic costs—litigation expense, regulatory compliance burden, provider network attrition—that a purely benefit-focused economic impact estimate would not capture.

## 3. Productivity Enhancement: Evidence From Clinical and Administrative Deployment

### 3.1 Precision Public Health as a Public Health Workforce Productivity Tool

The national framework for AI-powered precision public health developed by Nguyen (2026a) contributes to productivity enhancement at the public health system level by enabling 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 uniformly across an entire population regardless of underlying risk heterogeneity.

### 3.2 Empirical Evidence From Ambient AI Scribe Deployment

Beyond conceptual frameworks, empirical evidence from real-world clinical deployment provides direct evidence of AI-driven productivity enhancement in healthcare. Duggan et al. (2025) conducted a quality improvement study examining clinician experiences with ambient artificial intelligence scribe technology, an AI tool that passively records patient-clinician conversations and automatically generates draft clinical documentation, finding meaningful reductions in time spent on note-writing and overall visit-related documentation burden among participating clinicians; this study, alongside a growing body of similar real-world evidence on ambient AI scribe technology, represents some of the most directly measurable evidence currently available of AI-driven productivity gains in frontline clinical practice, complementing

the more conceptual, framework-level productivity case articulated in national policy documents.

### 3.3 Administrative Productivity Through Prior Authorization Automation

Nguyen (2026b) similarly frames AI-driven prior authorization transformation explicitly as a productivity intervention, aiming to reduce the substantial staff time that both payer and provider organizations currently devote to prior authorization submission, review, and appeal processes. When implemented with appropriate transparency and clinical oversight, this class of administrative automation holds genuine potential to redirect staff time toward higher-value activities, mirroring the documentation-time savings demonstrated for ambient AI scribes in the clinical domain.

### 3.4 Productivity Gains Are Not Unconditional

The productivity case for AI-driven administrative automation, however, is not unconditional. As the following section discusses in greater detail, AI-assisted prior authorization systems deployed without adequate transparency and oversight have in specific instances generated additional administrative burden—through increased appeals, denial navigation, and regulatory response—rather than reducing it, underscoring that productivity gains from healthcare AI are contingent on responsible implementation rather than an automatic consequence of AI adoption itself.

## 4. Long-Term Health Outcomes: Benefits and Documented Risks

### 4.1 Precision Public Health for Long-Term Chronic Disease Outcomes

Nguyen (2026a) articulated a national framework explicitly targeting chronic disease prevention and early intervention through AI-driven precision public health, 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 genuinely preventive intervention, generating long-term health outcome benefits that Nguyen (2026c) further characterizes in economic terms as a component of AI's broader national value proposition.

### 4.2 The Prior Authorization Case Study in AI-Related Health Outcome Risk

A rigorous, honest evaluation of AI's contribution to long-term health outcomes must also account for documented instances in which AI deployment has been associated with adverse, rather than beneficial, health outcome risk. Raza et al. (2026) examined the use of artificial intelligence in prior authorization review within Medicare Advantage, documenting an increase in reported instances of AI-assisted, blanket denials of coverage in recent years, resulting in insurer criticism, class action litigation, congressional investigation, and provider network attrition; the authors identify provider burnout and patient burden from navigating prior authorization requests and appeals as direct consequences of this pattern, representing a health outcome risk pathway—delayed or denied access to medically necessary care—directly counter to the health outcome benefits that AI-driven healthcare tools are intended to deliver.

#### 4.3 Transparency and Oversight as Preconditions for Realizing Health Outcome Benefits

Raza et al. (2026) explicitly call for improved transparency in how AI tools approve or deny coverage claims as a necessary condition for ensuring a fair healthcare system, alongside attention to provider burnout and patient burden arising from prior authorization navigation. This finding directly qualifies the long-term health outcome case for AI-driven prior authorization transformation articulated by Nguyen (2026b): realizing genuine health outcome benefit from AI-driven administrative automation depends on the transparency, clinical oversight, and regulatory safeguards that Raza et al. (2026) identify as currently insufficient in at least some real-world deployments, rather than following automatically from AI adoption itself.

#### 4.4 A Balanced Long-Term Health Outcome Assessment

Taken together, the precision public health framework (Nguyen, 2026a) and the documented prior authorization risk evidence (Raza et al., 2026) illustrate that AI's long-term health outcome impact in healthcare is neither uniformly positive nor uniformly negative, but contingent on the specific application, implementation safeguards, and regulatory context in which a given AI system operates—a nuance that a national evaluation framework must explicitly incorporate rather than assuming health outcome benefit as an automatic consequence of AI deployment.

### 5. Environmental Impact: The Infrastructure Cost of National-Scale AI Deployment

### 5.1 Quantifying the Public Health Cost of AI Computing Infrastructure

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 host communities such that the most affected counties bear a per-household health burden approximately seven times the national average—costs that apply proportionally to the share of national computing demand attributable to healthcare-sector AI applications as this sector's AI adoption continues to expand.

### 5.2 Quantifying the Carbon and Water 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, underscoring that the environmental cost of national-scale AI deployment, including healthcare-sector applications, extends well beyond electricity consumption alone.

### 5.3 Environmental Impact as the Underexamined Fourth Domain

Despite the scale of these environmental and public health costs, national evaluations of AI's value in healthcare, including the economic and productivity frameworks reviewed in this article (Nguyen, 2026a, 2026b, 2026c), have not typically incorporated environmental footprint accounting as a core evaluation dimension. A genuinely complete national framework for evaluating AI's healthcare value must treat environmental impact as a fourth, coequal evaluation domain alongside economic savings, productivity enhancement, and long-term health outcomes, rather than as an externality assessed separately, if at all.

### 5.4 Geothermal Energy as a Mitigation Pathway Requiring Rigorous Subsurface Science

Geothermal energy has emerged as a candidate low-carbon power source for AI computing infrastructure, given its capacity to provide the continuous, weather-independent baseload power that data center operations require. Realizing this mitigation pathway depends on

rigorous subsurface characterization capable of establishing both the resource quality and the operational safety of candidate geothermal reservoirs. 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, providing the detailed subsurface characterization necessary to identify geothermal resources capable of reliably supplying data center-scale power demand.

## 5.5 Geomechanical Safety as a Prerequisite for Infrastructure-Scale Deployment

Beyond resource characterization, safely scaling geothermal energy production to serve data center-scale power demand requires proactive geomechanical risk management. Machine learning-enhanced geomechanical characterization and wellbore stability assessment, integrating stress field estimation, rock strength prediction, and operational parameter analysis into a proactive risk-management workflow, has been demonstrated in offshore energy reservoir settings, offering a directly transferable methodology for ensuring that geothermal energy expansion intended to mitigate AI's environmental footprint does not introduce new geomechanical or induced seismicity risks of its own—a consideration essential to a genuinely responsible, rather than merely convenient, environmental mitigation strategy.

## 6. Toward an Integrated National Evaluation Framework

### 6.1 The Case for Honest, Cross-Domain Integration

The evidence reviewed across this article's four domains reveals a more complex picture than a purely benefit-focused national narrative would suggest: AI-driven precision public health (Nguyen, 2026a) and prior authorization transformation (Nguyen, 2026b) hold genuine promise for economic savings, productivity enhancement, and improved long-term health outcomes, empirically supported in specific applications such as ambient AI scribe deployment (Duggan et al., 2025), while documented risks in AI-assisted prior authorization review (Raza et al., 2026) and the environmental and public health costs of the underlying computing infrastructure (Han et al., 2024; de Vries-Gao, 2025) represent genuine, currently underweighted costs within most national evaluation narratives.

### 6.2 Geothermal Energy as a Case Study in Integrated Mitigation

Geothermal energy's potential contribution to reducing AI computing infrastructure's

environmental footprint illustrates how responsible infrastructure development can address multiple evaluation domains simultaneously, provided it is grounded in the kind of rigorous reservoir characterization and geomechanical risk management that ensures the mitigation solution does not introduce new categories of risk in the process of addressing another.

### 6.3 Transparency and Oversight as Cross-Cutting Requirements

A recurring theme across the economic, productivity, health outcome, and environmental domains reviewed in this article is the centrality of transparency and oversight to realizing genuine national value from AI in healthcare. Raza et al. (2026) explicitly identify transparency in AI-driven coverage determinations as necessary to a fair healthcare system, a principle that extends naturally to environmental disclosure by data center operators (de Vries-Gao, 2025) and to honest, methodologically rigorous economic impact estimation (Nguyen, 2026c) that does not omit realized risks and costs from its accounting.

## 7. Challenges and Limitations

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

Asymmetric evidence maturity across domains. Empirical evidence on AI's productivity impact, exemplified by ambient AI scribe studies (Duggan et al., 2025), and on AI-related prior authorization risk (Raza et al., 2026), has matured considerably in specific applications, while national-scale economic and public health impact estimation (Nguyen, 2026a, 2026c) remains comparatively more projective and framework-level, complicating efforts to integrate these evidentiary types within a single quantitative national assessment.

Attribution difficulty for shared computing infrastructure. As in other domains of AI deployment, healthcare AI applications typically run on shared cloud and data center infrastructure serving multiple sectors simultaneously, making it methodologically difficult to isolate the environmental and public health costs specifically attributable to healthcare-sector AI use (Han et al., 2024; de Vries-Gao, 2025).

Regulatory lag relative to deployment pace. The documented risks in AI-assisted prior authorization review (Raza et al., 2026) illustrate a broader pattern in which regulatory and oversight frameworks have lagged behind the pace of AI deployment in high-stakes healthcare

administrative functions, a gap that must be closed for the productivity and economic benefits of administrative AI automation to be realized without generating offsetting harms.

Uncertain timelines for geothermal infrastructure scaling. While geothermal energy offers a conceptually attractive environmental mitigation pathway, the reservoir characterization and geomechanical risk validation processes reviewed in Section 5 operate on timescales that may not align with the more rapid pace at which healthcare AI computing demand is currently expanding.

Absence of standardized integrated evaluation methodology. No standardized methodology currently exists for jointly evaluating AI's economic, productivity, health outcome, and environmental impact in healthcare within a single national framework, leaving the field reliant on narrative synthesis of the kind this article provides rather than a formally validated, quantitative integration methodology.

## 8. Future Directions

Future research should prioritize the extension of rigorous, empirically grounded evaluation—of the kind exemplified by ambient AI scribe productivity studies (Duggan et al., 2025) and prior authorization risk documentation (Raza et al., 2026)—across a broader range of AI healthcare applications, reducing the current reliance on projective, framework-level estimation for domains where empirical, real-world deployment evidence remains comparatively limited. Extending the national economic impact framework developed by Nguyen (2026c) to explicitly incorporate both the documented risks of AI-driven administrative tools (Raza et al., 2026) and the environmental and public health externalities of AI computing infrastructure (Han et al., 2024; de Vries-Gao, 2025) represents a critical next step toward genuinely integrated, honest national evaluation.

Advancing geothermal reservoir characterization and geomechanical risk management specifically in support of data center-scale power supply, building on high-resolution petrophysical characterization and machine learning-enhanced geomechanical characterization, represents a promising applied research direction for realizing geothermal energy's potential environmental mitigation contribution responsibly. Finally, developing standardized transparency and oversight requirements for AI-driven healthcare administrative tools, directly responsive to the risks documented by Raza et al. (2026), represents an urgent

regulatory research priority necessary to ensure that the productivity and economic benefits anticipated from AI-driven administrative automation are realized in practice rather than offset by the kind of administrative and health outcome harms already observed in specific deployments.

## 9. Conclusion

Evaluating the true national value of artificial intelligence in healthcare requires an honest, integrated framework spanning economic savings, productivity enhancement, long-term health outcomes, and environmental impact, one that weighs documented benefits against documented risks rather than presenting AI adoption as an unqualified national good. National frameworks for AI-powered precision public health and prior authorization transformation, together with empirical evidence on ambient AI scribe productivity gains, establish a genuine and, in specific applications, empirically demonstrated benefit case for AI in healthcare. At the same time, documented evidence of AI-assisted prior authorization tools contributing to increased coverage denials, provider burnout, and regulatory scrutiny, alongside the substantial public health and environmental costs of the computing infrastructure underlying national-scale AI deployment, reveal a set of genuine, currently underweighted costs that an honest national evaluation cannot omit. Geothermal energy, grounded in rigorous reservoir characterization and geomechanical risk management, offers a promising pathway for mitigating AI's environmental footprint, provided its development is pursued with the same rigor and transparency that responsible AI deployment in healthcare more broadly requires. Advancing toward a genuinely integrated national evaluation framework will require sustained investment in empirical, real-world evidence generation, standardized transparency and oversight requirements, and honest accounting of environmental and public health externalities—investments essential to ensuring that artificial intelligence's substantial promise in healthcare translates into genuine, responsibly realized national value.

## 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. Duggan, M. J., Gervase, J., Schoenbaum, A., Johnson, K. B., et al. (2025). Clinician experiences with ambient scribe technology to assist with documentation burden and efficiency. *JAMA Network Open*, 8(2), Article e2460637. <https://doi.org/10.1001/jamanetworkopen.2024.60637>
3. 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>
4. 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.
5. Nguyen, T. T. (2026b). Transforming prior authorization through artificial intelligence: A national framework for reducing administrative burden and improving patient access. *International Journal of AI, BigData, Computational and Management Studies*, 7(3), 17–27.
6. 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.
7. Raza, S., Gerke, S., Silcox, C., Hendricks-Sturup, R., & Shachar, C. (2026). Medicare advantage becoming a disadvantage with use of artificial intelligence in prior authorization review. *npj Digital Medicine*, 9, Article 208. <https://doi.org/10.1038/s41746-026-02387-x>