**POLICY MEMORANDUM**

Title: **Health Professional Shortage Areas (HPSAs): Equity, Trust, and Accountability Conditions for AI Deployment**

Author: Anika Gundlapalli

Stanford University

Date: January 20, 2026

Purpose
This memo distills findings from a policy-minded research review on the ethical implementation of AI and large language models (LLMs) in medically underserved and rural communities. It identifies the conditions under which these tools can address provider shortages and access gaps without compounding inequities.

Policy Problem
Underserved settings face a higher risk from biased model performance, limited infrastructure, constrained clinical oversight, and fragile institutional trust. If AI tools are deployed without enforceable safeguards, they may widen disparities, reduce patient confidence, and create unclear accountability when harm occurs.

## Findings

- Equity: Fair performance depends on inclusive data practices and subgroup-robust evaluation; average accuracy can mask systematic underperformance for rural and marginalized populations.
- Trust and privacy: Adoption depends on transparency, preserved relational care, and credible privacy protections, especially in communities shaped by historical neglect.
- Infrastructure and accountability: Telemedicine and AI can extend care only where broadband and workflows support safe use; legal ambiguity regarding attended vs. unattended systems increases risk for patients and providers.

## Recommendations

- Require equity evidence before scale: mandate subgroup performance reporting, bias risk assessment, and ongoing monitoring in the intended use population.
- Standardize responsible deployment: pair tools with clinician training, patient-facing explanations, and community feedback channels; maintain human oversight for high-stakes use unless validation and governance justify higher autonomy.
- Clarify governance and liability: implement enforceable rules for consent, permitted uses, and retention; specify accountability across clinicians, systems, and vendors, with explicit treatment of autonomy level.

Stakeholders
Patients and families in underserved communities; rural clinics and hospitals; community health centers; clinicians and community health workers; health system administrators; payers; AI vendors; and state and federal regulators.

# Introduction

Healthcare systems in rural U.S. communities face persistent strain from provider shortages, hospital closures, and limited resources. A recent study by Ekren et al. (2025) found that rural counties are significantly more likely to rank in the bottom quartiles for both clinical care and life expectancy, a pattern that is consistent in other rural communities. These conditions have prompted policymakers, researchers, and healthcare professionals to explore whether emerging tools like artificial intelligence (AI), deep learning (DL), and large language models (LLMs) could help fill critical gaps in access and efficiency.

Supporters argue that these technologies can improve diagnostic accuracy, streamline workflows, and partially offset workforce shortages. However, critics warn that AI systems, if not thoughtfully designed and regulated, may exacerbate inequities, breach patient privacy, and erode trust. The tension between technological promise and practical risk is especially evident in medically underserved settings, where historical neglect and systemic inequities already undermine care delivery.

This paper builds on those concerns by asking: How can we ethically implement LLMs and AI in medically underserved settings to address provider shortages without deepening existing inequities? Unlike approaches that isolate technical, cultural, or legal considerations, this paper proposes an integrated framework that combines inclusive data design, culturally responsive deployment, and adaptive regulatory protections.

The following literature review explores four key domains: bias and fairness, patient trust, rural health infrastructure, and legal regulation that are central to implementing AI and LLMs in underserved settings. These debates provide essential context for evaluating whether such technologies can ethically expand care without worsening inequity. This paper offers a timely and necessary framework to ensure that AI becomes a tool for inclusion, not exclusion, in the places that need it most.

## Literature review

### Bias and Fairness in AI

Bias in AI systems is not just a technical flaw but an ethical risk that, if left unaddressed, could amplify health disparities in an underserved setting. One of the main ethical concerns in medical artificial intelligence (AI) usage is the presence of bias, especially in how models generalize across diverse patient populations. Across recent literature, experts agree that although AI offers significant potential in healthcare, its implementation risks reinforcing existing health disparities unless equity is a design priority.

Although bias in medical AI is often viewed as a major barrier to equitable healthcare, emerging research shows it can be meaningfully mitigated, enabling the ethical and effective deployment of AI tools in underserved communities. Yang et al. find that AI often uses demographic shortcuts for diagnoses, leading to inequitable treatment. However, Yang et al. state that debiasing techniques can mitigate this, but there is a trade-off between allowing AI to use demographic shortcuts and reducing these shortcuts through debiasing methods. This is because in a healthcare setting, demographic information about a patient can be essential to making a diagnosis. Yang et al. state that "On the one hand, removing shortcuts addresses ID fairness, which is a crucial consideration in fair clinical decision-making. On the other hand, the resulting trade-offs with other metrics and non-transferability to OOD settings raises the question about the long-term utility in removing such shortcuts" (Yang et al., 2024)

Therefore, while eliminating biased features improves fairness for specific groups, it can unintentionally harm a model's overall performance or generalizability in real-world settings. This tension highlights the need to treat fairness as a central design constraint. In the context of this paper's ethical framework, the fairness–performance trade-off is not a reason to reject AI, but a challenge that must be actively managed, proving that ethical implementation is still both possible and necessary in rural healthcare, where the need for support tools is especially urgent.

While Yang et al. highlight the unavoidable tensions between fairness and model performance, Jain et al. attempt to resolve this dilemma by shifting the focus from feature debiasing to data selection strategies. Jain et al. propose a novel approach to sidestep this fairness–accuracy trade-off through D3M, a data-centric method that removes training examples shown to worsen performance for underrepresented groups. D3M does this without needing demographic labels or post-hoc reweighting, and it improves worst-group accuracy while preserving overall performance. This presents a more optimistic view: bias can be mitigated without sacrificing utility, especially if developers prioritize the needs of marginalized users. Yet even this solution is not completely free of complications. Because D3M involves filtering out data points, the model may become less robust in atypical or edge-case scenarios, which can be common in low-resource environments. This causes a trade-off like the one Yang et al. discuss.

Expanding on Jain et al.'s work, Sasseville et al. provide a comprehensive scoping review of bias mitigation strategies in primary health care AI systems. Their review synthesizes evidence across multiple model types and clinical contexts, identifying common sources of bias, dominant mitigation techniques, and persistent gaps in evaluation practice. Unlike single-method studies, Sasseville et al. emphasize that bias mitigation is not a one-time technical fix but a continuous governance problem requiring ongoing monitoring, subgroup reporting, and clinical validation. Their findings reinforce this paper's claim that fairness must be treated as a design constraint and a policy obligation, particularly in low-resource and underserved settings where biased errors carry disproportionate harm.

Ethical AI deployment in underserved areas is achievable when debiasing methods such as D3M are implemented. However, doing so responsibly requires that developers and policymakers acknowledge the trade-offs these methods involve, communicate their limitations transparently, and prioritize the needs of the communities they aim to serve. Together, these insights show that while bias mitigation involves trade-offs, it is possible to ethically implement AI in rural healthcare settings, if fairness is treated not as a constraint on performance, but as a condition for meaningful and equitable impact.

## Patient Trust and Privacy

Without patient trust, even the most accurate AI tools may fail in clinical contexts, especially in communities historically excluded from care. Tyson et al. show that while AI offers technical value in diagnosis and planning, trust erodes when it replaces human touch, especially in settings where patients rely on relational care. Their national survey reveals that over 60% of Americans feel uncomfortable with clinicians relying on AI, not due to doubts about its accuracy, but because of fears about safety, autonomy, and the potential erosion of empathy in healthcare settings. These findings highlight a crucial point: even when AI performs well, adoption depends on more than accuracy, it hinges on how patients emotionally perceive the care they're receiving. AI tools, no matter how precise, may be rejected if they are perceived as cold, opaque, or dehumanizing. Longoni and Morewedge show the deeper psychological roots of mistrust, stating that patients often reject AI because they view their health needs as fundamentally personal. They deepen this insight by showing that people's mistrust in AI stems from a belief that their health needs are too unique for standardized algorithms. Even when told that AI can outperform human doctors in diagnostic accuracy, patients still prefer human judgment because they believe humans can understand context, emotion, and individuality in a way machines cannot. This psychological discomfort reinforces Tyson et al.'s claim that relational trust is central to ethical care, particularly when AI is used in high-stakes decisions. The concern is not that AI is incapable, but that it may never fully grasp the personal nature of illness, and therefore feels alienating to those it is meant to help. These concerns do not mean AI should be abandoned, but rather that its implementation must preserve the interpersonal elements of care that foster trust.

While emotional trust centers on how patients feel during care interactions, structural trust involves whether institutions protect their rights and data. Murdoch expands the discussion from patient psychology to systemic design, arguing that trust also depends on enforceable protections of patient data. Murdoch frames these issues as institutional and policy failures rather than just perceptual ones, emphasizing that "a public lack of trust might heighten public scrutiny of or even litigation against commercial implementations of healthcare AI" (Murdoch, 2021, p. 3). He argues that AI systems, particularly those developed or deployed by private companies, raise serious risks related to data privacy, ownership, and consent. These concerns are especially acute in healthcare, where data is sensitive and intimately tied to identity. Murdoch points out that current legal frameworks lag behind technological development, leaving many patients unprotected against misuse or commercial exploitation of their health data.

Tyson et al. show that trust erodes when AI replaces human touch, while Longoni and Morewedge go deeper, suggesting that people mistrust AI because they see their care needs as deeply personal. In contrast to Tyson and Longoni, who focus on emotional and psychological mistrust, Murdoch reframes trust as a systemic issue: without enforceable safeguards, even well-designed AI tools may provoke suspicion, scrutiny, or litigation. This is especially relevant in medically underserved areas, where communities may already distrust institutional systems due to past inequities. This makes clear that implementing AI without strong privacy protections risks repeating patterns of neglect, but with appropriate safeguards, it can help correct longstanding access gaps. Ultimately, ethical AI implementation in underserved communities depends on building both emotional and structural trust. Patients must feel seen and respected in their care, and they must be confident that their data is protected and their rights are upheld. Without addressing both, AI tools, even those designed to improve access, may be rejected or cause harm. These dual dimensions of trust will re-emerge in later sections, particularly in the discussion of legal frameworks, where institutional protections are key to sustainable, responsible AI deployment.

## Rural healthcare infrastructure and telemedicine

Rural health disparities are shaped by more than just provider shortages; they are embedded in a web of geographic, economic, and infrastructural challenges that limit access to consistent and specialized care. While traditional interventions have focused on increasing the number of healthcare providers, Orgera et al. argue that this approach is insufficient on its own. Their policy brief emphasizes that rural patients face a broad range of barriers such as long travel distances to care, limited access to emergency services, and socioeconomic factors like poverty and lower educational attainment. These structural issues also intersect with concerns about fairness and trust: even if AI tools are technically unbiased, they cannot deliver equitable outcomes in environments where patients face logistical and infrastructural barriers to access. Such issues cannot be resolved through workforce expansion alone. While Nestrick (2024) provides a policy perspective rather than peer-reviewed research, Nestrick reinforces this view, pointing out that limited healthcare access is not just caused by having limited access to providers, it is also caused by lack of public transportation, socioeconomic factors, and hospital closures in rural areas. Nestrick emphasizes that any new method of addressing healthcare disparities in rural areas should be “grassroots initiatives firmly based in community engagement and unwavering commitment” (Nestrick, 2024). These authors collectively challenge the assumption that a uniform solution, such as adding more primary care physicians, can address the nuanced and deeply local needs of rural communities. Both Orgera and Nestrick argue that “rural health needs go beyond physician counts, particularly since employing a physician in every community is not feasible" (Orgera, 2023). Orgera and Nestrick argue that rural health strategists must use innovative tools to address structural access barriers and discuss using technology to address healthcare disparities if rural residents are open to the possibility of using technology.

One such technology is AI-driven telemedicine, which offers potential to address rural care gaps but also introduces new challenges related to digital access, patient trust, and infrastructure stability. Perez et al. provide a systematic review of AI and telemedicine in rural communities, highlighting their potential to improve diagnosis, patient monitoring, and care coordination across large distances. These technologies can help fill gaps in specialty access and extend healthcare reach into isolated areas where brick-and-mortar clinics are not feasible. Nestrick supports this approach, advocating for telehealth as a mechanism to bridge the rural care gap, especially in regions facing transportation barriers and persistent provider shortages. However, Perez et al. also caution that implementation is not seamless: barriers such as poor digital infrastructure, low digital literacy, and patient privacy concerns can limit the success of AI and telemedicine in these environments, undermining both its effectiveness and the trust it depends on. This tension illustrates that while AI can be a transformative force, it must be deployed thoughtfully to avoid reproducing the very inequities it aims to solve.

Addressing rural health inequities requires a shift in thinking from traditional, provider-centric solutions toward technology-enabled, context-aware models of care. Orgera et al. emphasize the importance of tailored policy over generic workforce expansion, while Perez et al. outline concrete ways AI can deliver timely and accurate care despite geographic obstacles. Nestrick reinforces that community engagement is essential to ensure these technologies meet actual local needs. If AI is integrated with community-specific infrastructure improvements and education, it may address healthcare disparities in ways that hiring more medical providers cannot. To implement AI effectively in rural areas, infrastructure must be in place, but so must trust and legal protections. These regulatory concerns are the focus of the next section, which explores how policy can support ethical deployment.

## Ethical frameworks and regulatory concerns

While AI technologies offer transformative potential for healthcare, especially in underserved and rural communities, ethical design alone is not sufficient. As the previous sections have shown, bias mitigation, trust-building, and context-sensitive deployment are essential, but they cannot succeed without enforceable legal frameworks to support them. As Dankwa-Mullan emphasizes, AI is not ethically neutral. Its promise, such as improved diagnostics and personalized medicine, can be undermined by structural biases embedded in training data, especially when those data exclude the very populations AI tools are meant to serve. This is particularly concerning for rural communities, which not only face persistent structural disadvantages, such as high chronic disease rates and limited specialty care, but are also historically underrepresented in clinical datasets. Without legal structures that reinforce ethical goals and hold developers accountable, even well-designed AI systems risk perpetuating or worsening these disparities. When regulation lags behind design, ethical intentions may fail to translate into equitable outcomes. The following discussion explores this tension and argues that synchronized ethical–legal alignment is essential for safe, just deployment in real-world healthcare settings.

Dankwa-Mullan argues that ethical implementation must go beyond technical fixes. It demands the direct involvement of underserved populations in the development and deployment of AI systems. This includes ensuring cultural relevance, mitigating bias, and fostering relational trust with communities that have long been marginalized by public health systems. In this context, transparency and community engagement are not just moral imperatives but practical necessities for adoption.

Yet even ethically sound design cannot guarantee public trust without clear legal accountability. Eldakak et al. (2024) highlight a growing disconnect between traditional legal frameworks and the unique challenges posed by autonomous AI in clinical settings. They argue that concepts such as negligence and strict liability are poorly equipped to address scenarios where AI systems act independently (unattended) or with minimal human oversight. This gap is especially dangerous in rural settings, where clinician shortages may increase dependence on such systems for triage, diagnostics, or remote monitoring.

The distinction Eldakak et al. make between attended and unattended systems is key: blurring the lines between them obscures accountability and weakens the legal protections available to patients. If harm occurs and it is unclear whether the clinician, the software provider, or the AI system is responsible, trust in both the technology and the care system erodes. Eldakak et al. therefore argue:

> The authors insist on the importance of distinguishing systems by their degree of autonomy and drafting liability rules accordingly, depending on whether an action was performed autonomously by an unattended system or automatically with an attended system (Eldakak et al., 2024, p. 1).

This distinction becomes especially urgent in rural contexts, where clinicians may rely more heavily on “unattended AI” due to provider shortages. In such cases, unclear liability raises patient safety concerns and increases legal risk for already resource-limited facilities, which can deter both patients and providers from using potentially beneficial AI tools.

The insights of Dankwa-Mullan and Eldakak et al. point to a central conclusion: effective AI implementation in underserved areas is not merely a technical challenge but an ethical-legal project. Success will require regulatory frameworks that clarify responsibility and protect patient rights. Only by addressing both domains can AI realize its promise without replicating the harms it aims to fix.

Together, these sources show that ethical implementation of LLMs and AI in underserved healthcare settings is possible, but only if it is guided by design choices that prioritize equity, systems that center trust, and policies that ensure accountability. Debiasing techniques like D3M, community-engaged care models, AI-driven telemedicine, and liability frameworks tailored to autonomous systems all represent promising steps toward closing gaps in care without deepening existing disparities. However, the literature also reveals unresolved tensions. Many proposed solutions, such as fairness algorithms and telehealth infrastructure, remain limited by trade-offs between accuracy and inclusivity, or by fragile implementation in low-resource environments. While trust is often discussed in relation to data privacy and performance, fewer studies explore how AI might reshape the emotional dynamics of care or disrupt the therapeutic relationship between patients and providers. Questions also remain about how to regulate LLMs specifically, especially those integrated into clinical workflows with minimal oversight. As AI systems continue to evolve and enter underserved settings, future research must address these cultural, emotional, and legal uncertainties and trade-offs to ensure that technological progress does not come at the cost of patient dignity or safety.

## Discussion

While the literature reveals consensus that AI and LLMs hold significant promise for underserved healthcare settings, this paper argues that ethical success depends on a framework integrating inclusive data design, patient-centered trust-building, and clear legal accountability. These three pillars are mutually reinforcing: data equity, trust, and legal clarity. Across domains, from bias mitigation and patient trust to infrastructure gaps and legal uncertainty, researchers agree that ethical implementation requires deliberate, equity-centered design. These findings reinforce this paper's core argument: AI can be ethically and effectively deployed in medically underserved communities only if development is rooted in inclusive data practices, use is guided by culturally responsive care models, and deployment is supported by adequate legal and regulatory frameworks. Together, these steps address the key risks illuminated in the literature, including unrepresentative data, mistrustful patients, fragile systems, and legal ambiguity. Accordingly, this paper identifies actions that should precede deployment: addressing bias, aligning model design with inclusive data practices, co-designing implementation with communities to build trust, and specifying accountability.

Inclusive data is not just a technical goal but the ethical foundation for any AI tool intended for equitable use in rural care settings. As Yang et al. and Jain et al. demonstrate, the fairness and accuracy of AI models are deeply shaped by how training data is selected and structured. In rural and underserved settings, where population health characteristics and care experiences often differ from urban norms, models trained on biased or incomplete data can produce dangerously inaccurate outcomes. Yang et al. warn against relying on demographic shortcuts that may replicate existing disparities, while Jain et al.'s D3M method offers a data-centric approach that improves performance for marginalized groups. This paper builds on that view by arguing that such practices must be embedded during model development, rather than treated as retroactive fixes, particularly when rural populations are involved. In communities already facing systemic barriers, any tool that cannot equitably serve all users risks reinforcing harm. Therefore, bias mitigation must be treated as a primary design constraint, with fairness evaluated alongside predictive accuracy in the intended population, documenting where performance falters and how data practices will be adjusted.

Equitable deployment also depends on trust, both in the technology and in the institutions behind it. Tyson et al. and Longoni and Morewedge emphasize that patients often reject AI not because it underperforms, but because it feels impersonal or opaque. This psychological discomfort is particularly acute in rural or marginalized communities, where care is often associated with relational familiarity. Longoni's finding that patients prefer a less accurate human doctor to a more accurate machine underscores how essential emotional and cultural factors are to perceived legitimacy. Trust must be cultivated intentionally. Nestrick argues that AI deployment must emerge from community partnerships, not be imposed from above. Relational engagement through local clinicians, community health workers, and educational outreach, when culturally inclusive and accessible, can bridge the gap between technical innovation and social acceptance. This paper extends these insights by arguing that ethical implementation demands not only technical functionality but evidence of care: systems must be visibly designed with, and for, the communities they intend to serve.

Legal clarity reinforces trust and makes ethical design enforceable, especially in rural settings where provider oversight may be limited. Murdoch points out that even well-performing AI systems will be met with skepticism if patients fear their data is being misused. Privacy and consent are not abstract issues; they are tied to histories of surveillance and neglect in underserved communities. Robust, transparent data governance policies must accompany technical deployment if AI is to gain meaningful public trust. Moreover, legal structures must evolve to keep pace with technological advancement. Eldakak et al. describe the urgent need for clear liability frameworks that distinguish between AI tools used with full clinician oversight and those operating more autonomously. This distinction is especially crucial in rural areas, where provider shortages may require AI tools to function with minimal supervision. Without clarity around who is responsible when an AI error occurs, both patients and providers face unacceptable risk. To address this, medical practices should clearly communicate liability and policies should distinguish between ‘attended’ and ‘unattended’ AI use. As Dankwa-Mullan argues, legal protections are essential not only for safety but also for equity. Communities without the legal literacy or political power to demand accountability are often those most affected by regulatory gaps.

Systems designed with inclusive data but deployed without cultural sensitivity will not be trusted. Tools introduced with community engagement but built on biased data will fail clinically. And even the best-designed systems will stall if the legal system cannot determine who is accountable for their outcomes. Ethical implementation is not a barrier to progress but a structure for sustainable success. Critics may argue that centering equity slows innovation or compromises performance. But the alternative, rushed deployment without regard for fairness, trust, or responsibility, is not faster. It is more dangerous, more wasteful, and more likely to provoke backlash or cause harm. Ultimately, LLMs and AI can be powerful tools to address provider shortages and improve care access in medically underserved areas, but only if their implementation is built from the ground up to reflect the values, vulnerabilities, and needs of the people they intend to serve.

Critics of medical AI often raise urgent and legitimate concerns, including privacy violations, loss of clinical judgment, data manipulation, and the fear of AI eventually surpassing human intelligence. A recent editorial by Siafakas and Vasarmidi (2024) in Pneumon outlines some of these fears, warning that AI could not only compromise sensitive health data and worsen existing inequalities but potentially lead to "Super AI" scenarios that threaten human autonomy altogether. The authors cite risks ranging from the erosion of physician skill sets to possible misuse of AI by governments or corporations. Although this is an editorial and not an empirical study, these perspectives reflect growing anxiety over AI's long-term trajectory and its potential to radically alter medicine in ways that are ethically and socially destabilizing.

While these concerns merit attention, they are largely speculative, often rooted in theoretical projections about a distant future rather than the pressing realities underserved communities face today. Notably, Pneumon's editorial does not draw on primary research or clinical trial data, but instead offers a broad, cautionary narrative that blends present-day implementation risks with futuristic AI doomsday scenarios. The editorial lacks empirical evidence from actual deployments of AI in rural health contexts and overlooks the more immediate ethical problem: systemic neglect and provider scarcity in underserved areas.

In contrast, empirical studies such as Ekren et al. (2025) provide a grounded understanding of the challenges currently facing rural healthcare. These include provider shortages, long patient travel distances, and delayed care due to underfunded infrastructure. In such environments, the question is not whether AI might replace physicians, but whether communities will even have timely access to care at all. Ignoring these urgent needs because of distant fears about "Super AI" risks further entrenching health inequities. Rather than reject AI outright, the more ethical approach is to design and regulate it responsibly to serve populations that are currently underserved.

Additionally, the most credible scholarship experts on AI ethics such as Yang et al., Jain et al., and Eldakak et al. do not argue that AI must be halted, but that it must be implemented with fairness, transparency, and accountability. These voices support a model of critical engagement: acknowledging risks while insisting on solutions that make AI safer and more equitable. Ethical AI is not about blind optimism, nor about technological determinism. It is about responsibly deploying powerful tools where they are most needed, with safeguards in place.

This paper argues that focusing on today's equity challenges and prioritizes the ones we can solve. Data bias, mistrust, and regulatory gaps are real and measurable. The ethical task is not to delay innovation until it is perfect, but to implement it responsibly, guided by frameworks that reflect the needs, rights, and voices of those in rural areas. The urgency lies not in avoiding AI altogether, but in ensuring that its development centers the needs, rights, and voices of those who have long been left behind by the healthcare system.

## Conclusion

Together, these sources show that ethical implementation of LLMs and AI in underserved healthcare settings is possible, but only if it is guided by design choices that prioritize equity, systems that center trust, and policies that ensure accountability. Techniques like D3M, community-engaged care models, AI-driven telemedicine, and tailored liability frameworks represent promising paths forward. However, the literature also reveals that many of these solutions face unresolved tensions: fairness algorithms often involve trade-offs between accuracy and inclusivity, telehealth infrastructure remains uneven in low-resource areas, and the regulation of LLMs in clinical settings is still underdeveloped.

These tensions point to a broader concern: without a coordinated framework integrating technical, cultural, and legal dimensions, even well-intentioned AI systems may deepen the disparities they aim to resolve. This paper's contribution is to offer such a framework, identifying conditions under which AI can be implemented not only effectively but equitably in rural healthcare systems. By structuring ethical implementation around data equity, patient trust, and legal accountability, this paper shifts attention from abstract debates about AI's risks to immediate, solvable challenges. In doing so, it offers a roadmap for researchers and policymakers to evaluate when an AI tool is ethically ready for deployment in low-resource environments to improve patient care.

Future research must go beyond optimizing performance and test this framework in practice by evaluating how AI tools operate within the lived realities of underserved environments, where infrastructure, trust, and accountability are not guaranteed. Ignoring these factors in the rush to innovate risks turning ethical challenges into structural harms. Embedding ethics into every stage of AI development is essential to ensure AI advances health justice rather than deepening inequity.

## References

Dankwa-Mullan, I. (2024). Health Equity and Ethical Considerations in Using Artificial Intelligence in Public Health and Medicine. *Preventing Chronic Disease*, 21. https://doi.org/10.5888/pcd21.240245

Eldakak, A., Alremeithi, A., Dahiyat, E., El-Gheriani, M., Mohamed, H., & Abdulrahim Abdulla, M. I. (2024). Civil liability for the actions of autonomous AI in healthcare: an invitation to further contemplation. *Humanities & Social Sciences Communications*, *11*(1). https://doi.org/10.1057/s41599-024-02806-y

Ekren, E., Maleki, S., Curran, C., Watkins, C., & Villagran, M. M. (2025, April 14). Health differences between rural and non-rural Texas counties based on 2023 county Health Rankings. BMC Health Services Research. https://doi.org/10.1186/s12913-024-12109-2

Jain, S., Hamidieh, K., Georgiev, K., Ilyas, A., Ghassemi, M., & Madry, A. (2024). Data Debiasing with Datamodels (D3M): Improving Subgroup Robustness via Data Selection. *ArXiv.org*. https://doi.org/10.48550/arXiv.2406.16846

Longoni, C., & Morewedge, C. K. (2019, October 30). AI Can Outperform Doctors. So Why Don't Patients Trust It? *Harvard Business Review*. https://hbr.org/2019/10/ai-can-outperform-doctors-so-why-dont-patients-trust-it

Murdoch, B. (2021). Privacy and artificial intelligence: Challenges for protecting health information in a new era. *BMC Medical Ethics*, *22*(1). https://doi.org/10.1186/s12910-021-00687-3

Nestrick, R. (2024, September 20). Bridging the gap: Addressing health inequities in rural communities. *National Rural Health Association.* https://www.ruralhealth.us/blogs/2024/09/bridging-the-gap-addressing-health-inequities-in-rural-communities

Orgera, K., Senn, S., & Grover, A. (2023, September 27). Rethinking Rural Health. *AAMC*. https://doi.org/10.15766/rai_xmxk6320

Perez, K., Wisniewski, D., Ari, A., Lee, K., Lieneck, C., & Ramamonjiarivelo, Z. (2025, February 4). Investigation into application of AI and telemedicine in rural communities: A systematic literature review. Healthcare, 13(3), 324. https://doi.org/10.3390/healthcare13030324

Sasseville, M., Ouellet, S., Rhéaume, C., Sahlia, M., Couture, V., Després, P., Paquette, J. S., Darmon, D., Bergeron, F., & Gagnon, M.-P. (2025). Bias mitigation in primary health care artificial intelligence models: A scoping review. Journal of Medical Internet Research, 27, e60269. https://doi.org/10.2196/60269

Siafakas, N., Vasarmidi, E. (2024). Risks of Artificial Intelligence (AI) in Medicine. *Pneumon, 37*(3), 40. https://doi.org/10.18332/pne/191736

Tyson, A., Pasquini, G., Spencer, A., & Funk, C. (2023, February 22). 60% of Americans Would Be Uncomfortable With Provider Relying on AI in Their Own Health Care. *Pew Research Center.* https://www.pewresearch.org/science/2023/02/22/60-of-americans-would-be-uncomfortable-with-provider-relying-on-ai-in-their-own-health-care/

Yang, Y., Zhang, H., Gichoya, J. W., Katabi, D., & Ghassemi, M. (2024). The limits of fair medical imaging AI in real-world generalization. *Nature Medicine*, *30*(12), 2838-2848. https://doi.org/10.1038/s41591-024-03113-4