INTEGRATING HEALTH AI INTO SAFETY-NET AND RURAL HEALTH CARE SETTINGS: RESEARCH, INSIGHTS, AND RECOMMENDATIONS

STRATEGY PAPER

PREPARED FOR
Episcopal Health Foundation

PREPARED BY
IC^2 Institute
April 2025

CONTRIBUTORS

Matt Kammer-Kerwick, Ph.D., co-Principal Investigator
Gregory Pogue, Ph.D., co-Principal Investigator
Ishani Purohit, MPH
Emily Spandikow, MSc
Rayna De Jesus
James Lifton
S. Craig Watkins, Ph.D., Principal Investigator

Recommended Citation

Kammer-Kerwick, M., Pogue, G., Purohit, I., Spandikow, E., De Jesus, R., & Lifton, J., Watkins, S. C. (2025). *Integrating Health AI into safety-net and rural health care settings: Research, insights, and recommendations*. IC^2 Institute.

CONTENTS

TABLE OF FIGURES

EXECUTIVE SUMMARY

"I'm open to AI and I think it could be a game-changer in health care going forward, but as a resource, not a replacement for humans." –Nurse

Artificial intelligence holds great promise for transforming health care effectiveness and efficiency, particularly for populations with limited access to care, including people who receive safety-net services and in rural areas. Health artificial intelligence (Health AI) has the potential to make care more predictive, prevention-oriented, and personalized. When properly designed and deployed, Health AI could address persistent challenges faced by underserved populations, such as lack of access to care, high rates of chronic disease, poor health literacy, and social determinants of health.

Yet, Health AI must be ethically and thoughtfully integrated into safety-net service settings to keep human relationships and patient well-being central to care, and to ensure that these technologies do not overlook population-specific needs nor exacerbate inequities. There is a balance between leveraging rapid technological advances and taking the time to ensure such advances drive improved health outcomes for all.

With support from the Episcopal Health Foundation, the IC^2 Institute at The University of Texas at Austin developed a study and strategy paper identifying the most effective, actionable opportunities to introduce and integrate Health AI into care, especially in safety-net and rural services.

Our mixed-methods research included:

- An in-depth literature review, a survey of 229 Texas health care practitioners and in-depth interviews with 20 health care practitioners and researchers.
- Invited participants included health care providers and stakeholders – physicians, nurses, nurse practitioners, and thought leaders across a range of health care settings in Texas rural and underserved areas.

Key Findings Overview

Our analysis revealed important insights to inform best practices for Health AI development and deployment in safety-net health service settings. The full strategy paper details our findings and recommendations. Next is a high-level summary of findings and insights from our literature review and our quantitative and qualitative analysis of survey and interview findings.

KNOWLEDGE BASE

Artificial intelligence (AI) is generally recognized as the development of machines designed to perform many tasks long associated with human intelligence. Health AI includes the use of medical data (e.g., clinical, behavioral, genomic, and social determinants of health) to sharpen clinical insights and clinical decision-making.

Artificial intelligence, including machine learning and generative AI, is transforming health care by enhancing diagnostic accuracy, enabling personalized treatment, and improving operational efficiency — especially in complex, high-cost systems like the U.S. While Health AI holds great promise for early detection of disease, predictive medicine, and tailored care, it also introduces critical challenges related to bias, data privacy, and ethical responsibility, particularly for marginalized or underserved populations. Bridging the gap between visionary innovation and real-world application requires inclusive design, strong ethical frameworks, and regulatory oversight to ensure AI is trustworthy, equitable, and beneficial across all health care settings.

QUANTITATIVE RESULTS

Our survey findings and data centered on six themes related to health care providers' perceptions about AI, including its value and utility for themselves and their patients. The six themes are depicted in the figure below.

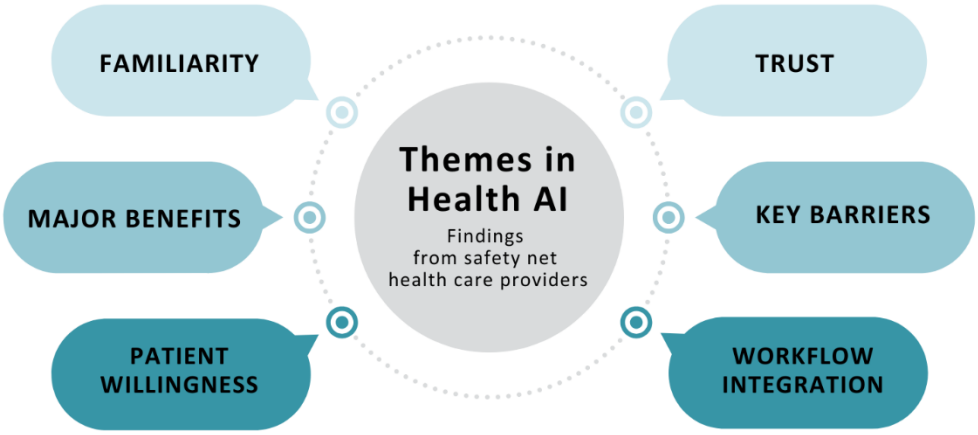

Throughout this study, we gauged respondents' attitudes about and disposition toward AI along several different dimensions. The figure below depicts respondents' familiarity with and their trust in AI, as well as their perception about their patients' responsiveness to AI and the degree to which they are confident that AI can be integrated into their current workflow. These and other items in the survey allow for a more in-depth understanding of what safety-net service providers think about AI and its prospects for enhancing health services. A significant portion of survey respondents, 73%, were either somewhat or very familiar with AI, while over one in four,

27%, reported that they were either somewhat or completely unfamiliar with AI. Whereas 48% of respondents reported that they somewhat or completely trust AI, one in five, 20%, reported that they either somewhat or completely distrust AI. As we note throughout the full strategy paper, a low level of trust in AI is a major barrier to AI adoption in safety-net health care settings. Just under half of respondents, 45%, believed that their patients would be responsive to AI tools being used in their care compared to three in ten, 30%, who believed their patients would be somewhat or very non-responsive to AI. And when it comes to integrating AI technology into health care, more than half of our respondents, 57%, were neutral or not very confident in their organization's ability to integrate AI into their workflow. See the full strategy paper for more details. The charts below illustrate the complexities of adopting and deploying AI in safety-net settings and suggest that, while safety-net providers are familiar with AI and even display modest levels of trust, there are notable barriers that must be addressed to fully unlock the potential of AI in safety-net health care environments.

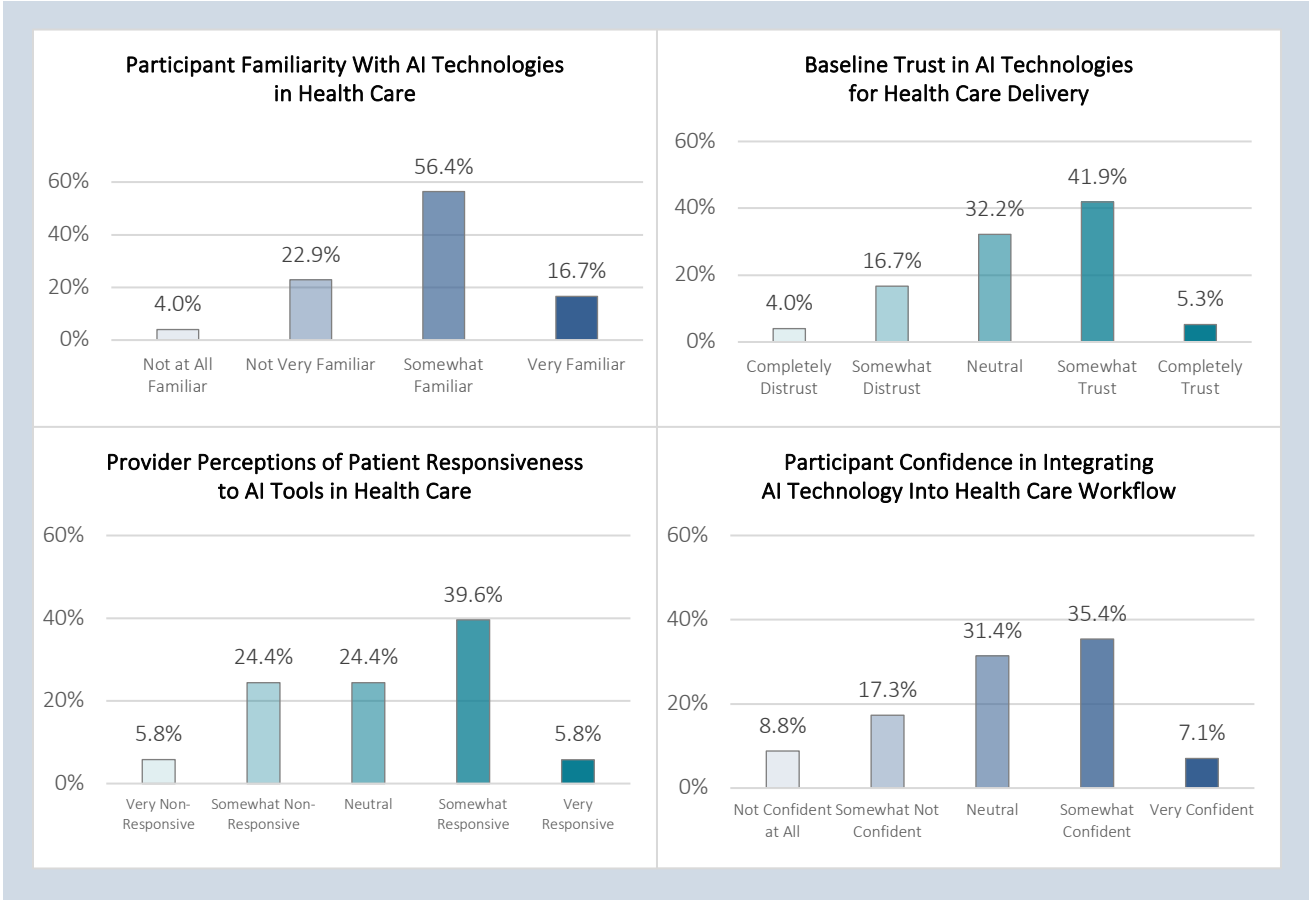

A high degree of trust was associated with survey participants who were more familiar with AI technologies in health care and perceived themselves to have a better understanding of AI and machine learning in general. Providers with low trust cited a need to build AI experience and exposure, and they had concerns around data bias and privacy/security. Degree of trust was not correlated with participants' role in their organization or their geographic location.

When asked to consider what they perceived to be the top benefit and top barrier to AI in safety-net service settings, participants gave the following leading three answers per category.

Top Benefit of Health AI		Top Barrier to Health AI Integration	
Streamlined administrative tasks	38.8%	Concerns about data privacy/security	31.9%
Enhanced patient outcomes	24.2%	Insufficient staff training/knowledge	19.5%
Improved diagnostic accuracy	23.3%	Lack of funding for AI implementation	17.3%

Specific benefits. Respondents also highlighted several areas where AI could significantly improve care for populations in safety-net services or rural areas, including remote patient monitoring, diagnostics, and enhancing health literacy. As Health AI augments health care practitioners' ability to enhance health literacy, it can empower patients and support better health outcomes.

Patient challenges. The health care providers surveyed noted challenges or barriers to patient adoption and responsiveness to Health AI technologies in rural and safety-net service settings, including technology access, technological literacy, and cultural differences such as language. To overcome these challenges, effective design and culturally tailored approaches are essential.

Provider challenges. In addition to the lack of training and preparation to use Health AI, highlighted in the figure above, health care providers in the study cited two additional barriers to their own adoption of AI technologies: inadequate technical infrastructure in underserved areas/service settings and resistance to change among both providers and patients alike.

Cautious optimism. Health care providers in our study were cautiously optimistic about the value and utility of four realistic but fictional Health AI solutions/tools presented in the survey, and cited their reasons for caution and optimism. The tools presented could increase provider efficiency, develop patient-specific therapy plans, increase precision in risk assessment, and improve patient and care team communications and engagement.

QUALITATIVE INSIGHTS

Through in-depth interviews and open-ended survey questions, health care providers, researchers, and other thought leaders shared their perspectives on Health AI, including the reasoning behind their views. They identified several factors that influence provider or patient trust in Health AI. These factors are highlighted in the figure at the top of the next page.

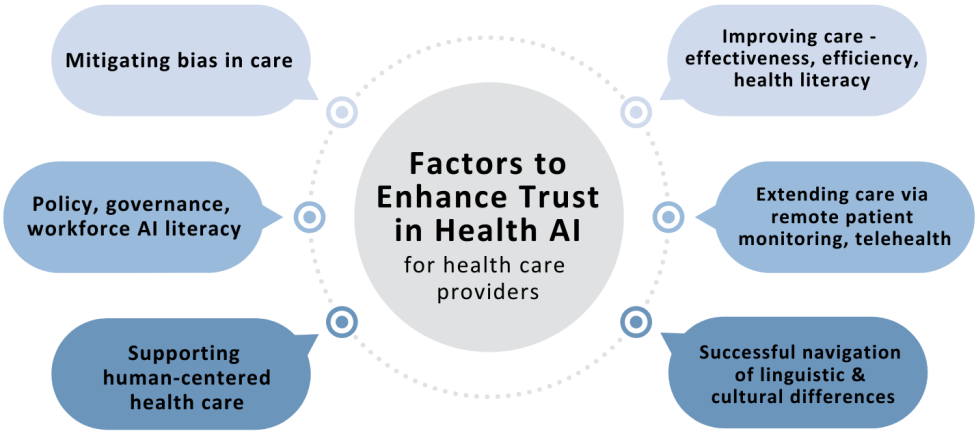

A common theme across these qualitative insights is the importance of keeping human-centered care at the core, even as AI-based solutions are adopted. In this study, we define AI-enabled, human-centered care as: *Systems that augment the abilities of human health care practitioners to deliver care that is always deeply empathetic and responsive to the unique needs, clinical and social, of a patient.*

4 Strategic Recommendations:

Successful AI Integration in Safety-Net and Rural Health Care Settings

Considering our research findings, we offer four strategy recommendations on how stakeholders and leaders in health policy, philanthropy/funding, and services can actively improve care for rural communities and others in need of safety-net health care services through a holistic and strategic approach to Health AI. The main report also includes action steps.

1	**Educate: Build AI Literacy and Workforce Capacity** Collaborate with local organizations to host an AI health care governance conference; develop statewide educational programs for health care professionals, focusing on responsible AI use.
2	**Participate: Engage Rural and Safety-Net Services Populations** Actively involve safety-net and rural service populations in AI adoption by conducting region-specific research and piloting AI solutions.
3	**Evaluate: Establish Strong Evaluation Frameworks** Create key performance indicators (KPIs) to evaluate AI tools for consideration and integration, and to assess their impacts.
4	**Innovate: Develop Locally Driven AI Solutions** Support the development of AI solutions tailored to Texas' unique health care challenges through hackathons and innovation incubators.

Conclusion:

Toward an Effective, Efficient, Inclusive Health AI

While intelligent machines may be an inevitable part of our future, the role they play in society is not yet determined. By combining educational, participatory, evaluative, and innovative strategies, Health AI can be responsibly introduced in safety-net and rural health care settings to improve care delivery, boost operational efficiency, and build trust in technology. Certain structural barriers must be addressed, including strengthening related technology infrastructure and access, improving digital literacy through culturally and linguistically appropriate training, and addressing resistance to change in communities. Our strategic recommendations offer a framework for deploying AI in ways that center equity, accountability, and collaboration – ensuring that underserved populations benefit from Health AI advances. They also include guidance for policy, governance, and evaluation of Health AI tools. The long-term success of Health AI will depend on thoughtful implementation in environments that have been historically overlooked, paving the way for more equitable, high-quality care across all communities.

The full strategy paper that follows offers additional data points, details on our findings, and insights from our analysis to drive strategic investments in the future of Health AI.

PART ONE: CONTEXT AND RESEARCH APPROACH

With the arrival of generative artificial intelligence (AI) systems like ChatGPT, the presence of AI technologies in our lives is more pronounced than ever. AI currently shapes our media consumption (think Netflix), our transit and navigation experiences (think Google Maps), and how we access information (think social media platforms). Even more, AI is active in many high-stakes contexts, such as criminal justice, education, employment, and health care. This raises important questions about the risks and opportunities of deploying AI in critical, high-stakes contexts, including life-and-death situations. In this strategy paper, our focus is on the promises and perils of AI in health care (Health AI).

Broadly speaking, Health AI refers to the application of artificial intelligence techniques and technologies to health care and medical contexts. This includes, for example, creating intelligent systems that can analyze, interpret, and process medical data to assist health care professionals in making better decisions. Some of the widely cited promises of Health AI include the delivery of care that is more predictive, prevention oriented, and personalized. While the use of Health AI is expanding, its benefits certainly do not reach all patients and medical settings equally. One of the primary challenges will be to design Health AI in ways that cut across the spectrum of providers and patients to ensure *that broader segments of the population can access and benefit from the promise of precision health.*

Focus, Methodology, and Framework of the Current Analysis

This strategy paper addresses key questions and concerns about Health AI. For example, what risks are associated with the application of machine learning techniques in health services delivery? We address Health AI risks such as bias, inadequate transparency and explainability, and the threat of substituting machine intelligence for human intelligence. Along the way, the strategy paper will address **three core questions**:

1. What are the risks and opportunities associated with developing Health AI for underserved populations?
2. What are the barriers to building a more robust and diverse AI workforce among health care providers who work in safety-net service settings?
3. How can the perspectives and expertise of health care providers of safety-net services inform the design and deployment of Health AI in ways that transform their capacity to deliver impactful care to underserved populations?

Our analysis, conducted by a team of researchers from the IC^2 Institute, has a specific focus on Health AI for **health care providers in safety-net settings** who serve rural populations, with a secondary focus on urban center populations. This strategic analysis is generously supported by the Episcopal Health Foundation, a visionary team that prioritizes these same focus areas and populations for services and investments. To complement the findings and insights in the full

strategy paper, we offer examples and a vision for targeted AI efforts related to maternal health, diabetes, and food insecurity, specific Foundation priority areas (see Part Two Spotlight box).

To date, little research or experimentation in Health AI has focused on safety-net care providers. We believe this to be a missed opportunity, one that perpetuates a core concern about AI: the degree to which AI design and deployment exclude the perspectives and participation of diverse ideas and expertise. Our focus on safety-net service providers is crucial for many reasons. For one, these providers are responsible for delivering health services and care to populations more likely to be underinsured or uninsured and underserved. Second, we believe the study of safety-net service providers and Health AI offers the chance to consider a range of social, environmental, and behavioral factors that impact the health outcomes of underserved populations, thus ultimately enriching the kinds of Health AI systems developed for all populations.

A Mixed-Methods Approach

For this study and strategy analysis, our research team took a mixed methods approach with:

- An extensive **review of literature** relevant to the state of AI in health care, focusing on concepts related to AI, health care, equity, and marginalized communities. A full version of the literature review is provided in a separate companion document to the strategy paper.
- A **survey** with results from 229 Texas health care practitioners in a range of safety-net health care settings and roles, focused on current awareness and general perceptions of AI as well as reactions to specific vignette-based Health AI deployment scenarios.
- A set of in-depth **interviews** with 20 health care practitioners, researchers, and thought leaders to further explore issues and themes from the literature review and survey.

See Appendix A for further details on our methodology.

Anchored in the findings from our literature review (see separate document), survey, and interviews with health care professionals working in safety-net services and settings, the strategy paper offers an in-depth look at some of the challenges and opportunities AI presents in the health care space. In the pages that follow, you'll find:

- A primer on AI, Health AI, and AI use in health care.
- Ethical considerations for Health AI, particularly for safety-net services and settings.
- Findings from the current survey and interviews.
- A set of 4 strategic recommendations for Health AI integration in safety-net and rural health care settings and services.

PART TWO: FOUNDATIONS OF HEALTH AI

Primer on Artificial Intelligence and Health AI

We begin with some basic definitions and descriptions of AI and Health AI. Researchers began speaking about "computer intelligence" in 1947, exploring the possibility that machines could imitate humans based on the data provided and the programming conditions. While the scope and definition of artificial intelligence continue to evolve at a steady pace, **AI is generally recognized as the development of machines designed to perform many tasks long associated with human intelligence**. More specifically, today's AI systems are able to execute tasks like visual perception, language translation, speech recognition, and decision making. Due in part to significant advances in computing power, AI can perform tasks in a manner that far surpasses human ability. This capacity makes the technology a powerful tool across a number of domains. Whereas humans contend with biological and cognitive constraints that limit both the amount of information we can process and the speed at which we can process it, computers do not have the same limits. That gives computers a noteworthy advantage over humans.

As AI systems continue to evolve, the technology steadily creeps into many different spheres of our daily lives. For example, if you use a smartphone, you are guided by an intelligent system that can perform navigation or translate language. If you subscribe to Netflix or shop on Amazon, AI systems routinely make recommendations that may influence your viewing or consumer behavior. As AI systems become more capable, they will continue to become more deeply implicated in the lives of individuals, organizations, and society. While intelligent machines may be an inevitable part of our future, the roles that these systems play in society is not inevitable.

Types of AI

Machine Learning. There are different subsets of AI. One common example of AI is machine learning (ML), which involves computers programmed to make predictions or decisions based on particular forms of data and training procedures. In supervised ML, humans select which features of the data computers should focus on in order to facilitate a prediction or decision-making task. For example, ML systems can predict the onset of depression based on features like mood patterns, social interactions, and/or social media posts. In this instance, humans can identify the most relevant features that precede an episode of depression and train a supervised AI model to identify those patterns to make a prediction. By contrast, unsupervised ML models learn to make predictions or detect patterns in data without any form of human guidance like data annotation or labeling of which features a model should identify while performing a task, predictive or otherwise.

Deep Learning. Deep learning is another branch of AI. More specifically, deep learning systems use artificial neural networks (information networks modeled after neurons) to find distinct

patterns in large datasets with very little human instruction. Deep learning models are engineered to simulate the human brain, using neural networks or interconnected neurons to process and identify very specific features in raw data. An example of a deep learning systems would be an image recognition tool. After supplying an AI model with a large dataset of dog images, a deep learning model would be able to discern the features such as tail, teeth, and legs that identify a dog in a photo.

Generative AI. The arrival and popularity of ChatGPT has introduced the world to yet another subset of AI, generative artificial intelligence. Generative AI is a branch of AI that uses neural networks to generate an output — a piece of text or an image — similar to a large set of data inputs that a system has been trained to simulate. Unlike deep learning systems, generative AI systems do not just recognize patterns in a specific set of human artifacts or datasets; they also produce outputs that realistically simulate the artifacts in any given dataset.

Today's AI systems are designed for specific tasks or a narrow set of related tasks. This is widely referred to as narrow or "weak AI." These systems thus operate under predefined parameters and are unable to perform tasks outside of their scope. Take for example, a postpartum risk prediction AI system designed to detect early signs of depression among new mothers. That system would analyze very specific forms of data (e.g., sleep patterns, mood and emotional changes, social support, heart rate) to perform a very specific task: predicting whether or not a new mother is vulnerable to postpartum depression. The term *narrow* is used because this system cannot be used in any other clinical context — like cancer detection or diabetes — or any other domain — like finance or criminal justice.

Artificial General Intelligence. Humans can learn to do a variety of things such as create recipes, navigate neighborhood traffic, write a legal brief, or respond empathetically to someone in crisis. Many of the big AI companies are pursuing another type of AI that is capable of learning like humans. Artificial general intelligence (AGI) or "strong AI" is often described as a form of AI that can simulate, for example, the multimodal intelligence of humans, including visual recognition, language translation, and decision making. What makes AGI "strong" is its capacity to learn, reason, adapt, and solve novel problems in diverse contexts without manual intervention. AGI is still a theoretical concept and does not currently exist.

The Promise, Peril, and Practicalities of Health AI[1]

Our focus is Health AI, which uses medical data (e.g., clinical, behavioral, genomic, and social determinants of health) to sharpen clinical insights and clinical decision making. This involves, among other things, deploying AI to improve health care delivery, patient outcomes, and/or

[1] For more extensive explorations into the risks and opportunities associated with the development and deployment of Health AI, see National Academy of Medicine, Whicher, D., Ahmed, M., Israni, S. T., & Matheny, M. (Eds.). (2023). *Artificial intelligence in health care: The hope, the hype, the promise, the peril*. The Learning Health System Series. National Academies Press (US).

operational proficiency. As the U.S. healthcare system becomes more complex, inefficient, costly, and uneven in terms of the quality of care delivered to patients, many are turning to AI to address these challenges. For example, researchers explored the use of AI to significantly improve the diagnoses and detection of diseases like breast cancer via medical images. A reliable AI diagnostic tool could reduce the number of errors and operational inefficiencies to create a more fluid and effective workflow.

One great promise of Health AI is its potential to refine predictive medicine. Because the vast majority of chronic diseases in the U.S. are preventable (think diabetes or heart disease), there is widespread consensus in the need to develop and refine models of health care that emphasize disease prevention via early detection tools and protocols. This will not only lead to improved health outcomes, it will also significantly reduce the high economic costs associated with managing chronic diseases.

AI systems' ability to analyze massive amounts of data, identify patterns, and make inferences about the likelihood of the onset of a disease suggest that these systems, when deployed properly, can give practitioners more capacity to practice preventative medicine. For example, a remote patient monitoring system might detect that a client who has previously presented with mild depression has consistently reported low mood scores and has rarely left home in the last two weeks. An AI system might alert a behavioral health professional of this pattern, predict that the patient may be approaching depression or crisis and, indicate that the patient may require immediate attention. Such an approach could help avert a clinical emergency.

A notable challenge regarding the predictive attributes of Health AI is the need to develop use cases and data streams that draw from diverse population sets. A predictive model in the chronic disease space trained on data from a relatively healthy population will be less reliable when used, for example, to assess safety-net patient populations.

Another great promise of Health AI is its potential to personalize health care. AI's ability to analyze large and varied data streams — genomic and clinical data, social determinants of health — suggests that it is poised to individualize aspects of medicine. Imagine scenarios in which a Health AI system could help practitioners tailor clinical treatment based on some combination of an individual's genetic, health, and behavioral profile. AI might be leveraged to develop a targeted therapeutic plan for a patient's unique set of social, environmental, and clinical circumstances, thus avoiding current one-size-fits-most approaches to care delivery.

However, the ability to capture and analyze personal data poses serious questions related to data privacy and security. As AI systems continue to evolve, measures that ensure safe and ethical uses of patient data will be crucial. Moreover, patients should have the right to know how their data is being used and how an AI model made a specific decision about them.

Researchers are also beginning to experiment with the use of generative AI models embedded in medical products. One example is the use of virtual assistants or chatbots that can deliver medical advice, medication reminders, psychoeducation, or symptom checks. Unlike humans, this chatbot is available 24/7/365 to serve as a resource for information, emotional support, and adherence to instructions from a health care professional or regimen. However, the challenges associated with deploying AI chatbots are substantial. For example, generative AI models have a tendency to hallucinate or make things up, deliver content that may be inaccurate or irrelevant, and fail to exhibit empathy in highly sensitive uses cases.

Ethical Considerations for AI in Safety-Net Service Settings

There are key ethical considerations at the foundation of our analysis and strategy for Health AI. For much of its existence, AI has primarily been treated as a technical problem to be solved. This includes focusing on building systems that function technically and computationally, and produce relevant outputs. However, the steady adoption and deployment of AI in high-stakes domains like education, employment, criminal justice, finance, and health care means that **AI is not only a technical problem to solve, there are social and ethical considerations to address as well**. The use of machine intelligence to impact decisions regarding who receives a loan, who is hired, or who is most likely to receive an accurate clinical diagnosis raises significant ethical questions. Indeed, a whole field of knowledge related to what has variously been referred to as Ethical AI, Responsible AI, Fair AI, and AI Safety has emerged in academia, industry, and community and public sector domains. Generally speaking, this field of inquiry raises questions that focus attention on the many ways AI, if not properly developed, deployed, and monitored, can cause societal harm.

The ethical challenges posed by AI are especially resonant in health care, a sector where individuals health, well-being, and lives are literally at stake. If an AI system makes an error in predicting a movie for you to watch, it may be annoying but not likely hazardous to your health. If an AI system makes an error in predicting the onset of a chronic health condition, the results could be life threatening.

Some of the most cited ethical AI principles include:

- Transparency
- Explainability
- Accountability
- Bias mitigation

Transparency, Explainability, and Accountability

In practice, a more transparent AI system might subject its algorithmic operations to third-party evaluation, versus maintaining a closed or “black box” posture that keeps algorithmic decision-making processes secretive. A more explainable AI system would establish the conditions for a

patient to receive an informed explanation of how an AI-assisted clinical decision was made. Researchers are also demanding that there be clear accountability for algorithms that discriminate or cause harm to the patient the AI was intended to serve. Who should be responsible for a clinical AI error that causes some discernible harm? Would it be the clinical team that used the AI system, the health care organization that purchased the AI system, or the team that developed the system? Currently, there are no clear legal or policy guardrails for how these and other critical matters related to Health AI would be resolved.

Addressing Bias With Ethical AI[2]

As we've noted, AI systems can perform cognitive tasks that far surpass human cognition, thus making them powerful and hyper-efficient tools. But this also means that AI systems have the capacity to automate and scale various forms of inequality. Research has consistently found bias, often implicit, in AI systems. Take, for example, the creation of a clinical machine learning algorithm that is designed to predict the likelihood of patients developing diabetes based on their medical records and lifestyle data. These algorithms often rely, in part, on human expertise to function. Humans, for example, select and label the data that trains an AI system to perform a narrow task. In this case, human annotators may be used to help label the data during the data preprocessing and training stage. Human annotators may rely on pre-existing clinical guidelines that are based on historical research, which may have excluded certain underrepresented populations. Let's say annotators decide to use body mass index (BMI) thresholds in the diabetes-related algorithm. These thresholds, determined by clinical studies, overlook cultural differences in body composition and health outcomes. The use of BMI would, in effect, integrate historical bias into a machine learning algorithm being developed to predict the onset of diabetes. Though unintended, the bias could lead to systematic disparities in the functioning of this predictive AI model, thus producing real harm to patients whose BMI may not fall within the thresholds established in biased clinical studies.

Researchers from the University of California at Berkeley and The University of Chicago examined a risk-based algorithm that was used by a hospital to determine which patients needed access to a high-risk care management program and its available services. The researchers found that, even though Black patients suffered poorer health outcomes such as being diagnosed with more chronic diseases, they were much less likely to be recommended for this health risk management program than White patients. The algorithm was engineered to use cost spent for medical services as a proxy for identifying who needed health care. The researchers concluded that this technical decision was implicitly biased against Black patients. This particular AI model was, in effect, trained to recommend people who could pay for health care rather than those who suffered from poorer health outcomes. To guard against the implicit

[2] For an example of racial bias in health-oriented algorithms see, Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science (New York, N.Y.)*, *366*(6464), 447–453. https://doi.org/10.1126/science.aax2342

bias in this system, researchers recommended training the algorithm to determine risk based on other variables, including the number of chronic conditions that needed treatment in a year.

The importance of representing people with different demographic characteristics in training data and testing AI operations goes beyond just technical accuracy. The need for representation reflects the need for AI to act according to the highest expectations of human ethics.

Key components of ethical AI include:

- Protecting data providers' privacy and safety.
- Using systems that are explainable, accountable, and transparent in action.
- Focusing on outcomes that are robust, fair, non-discriminatory, ultimately exhibiting beneficence, promoting well-being, and doing no harm.[3]

AI researchers and developers have devised other techniques for mitigating bias in AI, such as developing representative datasets and building fairness and equity metrics into algorithmic procedures (i.e., predictive parity, demographic parity), along with adequate monitoring and auditing procedures to identify outputs that may lead to bias or disparate impacts. In addition, there is growing consensus that the design, development, and deployment of AI should include a diversity of voices and expertise. This principle is sometimes referred to as **inclusive design**. Take for example, a Health AI system that is used to analyze the management of postpartum depression among mothers in rural communities. An inclusive design approach would ensure that the voices, expertise, and perspectives of rural stakeholders — health care practitioners, patients, and community advocates — were fully incorporated into the system's design and deployment, thus strengthening its ability to deliver solutions relevant to mothers and practitioners in rural settings.

Practicing these and other ethical principles will not only create AI solutions that are fine-tuned for safety-net services environments. Doing so will also strengthen trust in these systems among rural and safety-net services populations, thus addressing one of the major barriers to the spread and adoption of AI systems in the range of diverse health care settings. *If we attend to the ethics of AI, the focus can shift to alleviating structural barriers in health awareness, access, and delivery of quality care.*[4]

[3] Bertoncini, A. L. C., & Serafim, M. C. (2023). Ethical content in artificial intelligence systems: A demand explained in three critical points. *Frontiers in Psychology, 14*, 1074787. https://doi.org/10.3389/fpsyg.2023.1074787

[4] Patel, M. I., Lopez, A. M., Blackstock, W., Reeder-Hayes, K., Moushey, E. A., Phillips, J., & Tap, W. (2020). Cancer disparities and health equity: A policy statement from the American Society of Clinical Oncology. *Journal of Clinical Oncology, 38*(29), 3439–3448. https://doi.org/10.1200/JCO.20.00642

Spotlight on AI's Potential to Impact Three Health Priority Areas

The Episcopal Health Foundation has identified three priority areas for deeper consideration, investment, and action: maternal health, diabetes care, and food insecurity. The Foundation and others invested in improving outcomes in these three areas can consider the following targeted examples and vision for how AI can be a resource for mitigating the prevalence of disease, poor health outcomes, and systemic disparities.

Maternal Health

AI is already used in maternal health to improve prenatal care through risk assessments for conditions such as gestational diabetes and preeclampsia, and to enhance diagnostic accuracy with AI-powered ultrasounds. It also helps predict potential labor complications, enabling earlier intervention. AI could revolutionize maternal care by offering real-time, personalized care plans and continuous monitoring via wearable devices, optimizing both prenatal and postpartum care. This type of research is currently ongoing by a team of social and computational scientists at UT Austin. However, challenges remain, including limited access to advanced technologies in rural areas, provider skepticism about AI's reliability, and the risk of biased outcomes if AI models are not trained on diverse, representative datasets.

Diabetes Care

In diabetes care, AI currently supports glucose monitoring, insulin dosage adjustments, and blood sugar trend forecasting, helping patients manage their condition with greater precision. In the future, AI could predict complications such as neuropathy or retinopathy and provide continuous health monitoring by integrating with wearable data collection devices, offering a more personalized approach to diabetes management. However, the effectiveness of these tools depends on the accuracy of data from devices, and device malfunctions or inaccuracies could lead to incorrect recommendations. Additionally, continuous monitoring may cause patient fatigue, and some individuals might struggle with fully engaging with AI tools or understanding how to use the tools effectively.

Food Insecurity

AI is currently optimizing food distribution networks by predicting shortages and helping to connect volunteers with food banks, improving the logistics of food assistance. Looking ahead, AI could offer personalized food assistance plans tailored to individuals' nutritional needs, preferences, and health conditions. It could also predict food insecurity patterns, enabling proactive interventions before shortages become critical. However, challenges include data reliability from individuals experiencing food insecurity, which can impact the accuracy of AI models and the complexity of scaling personalized solutions. There is also the potential for community resistance if AI-driven solutions are perceived as impersonal or disconnected from local realities.

PART THREE: ANALYSIS OF QUANTITATIVE FINDINGS

Participant Roles and Settings

The health care providers and stakeholders in our survey sample represent different segments of the safety-net health care workforce. Figure 1 illustrates the distribution of survey participants' organization types.

Figure 1. Organization Types in Sample

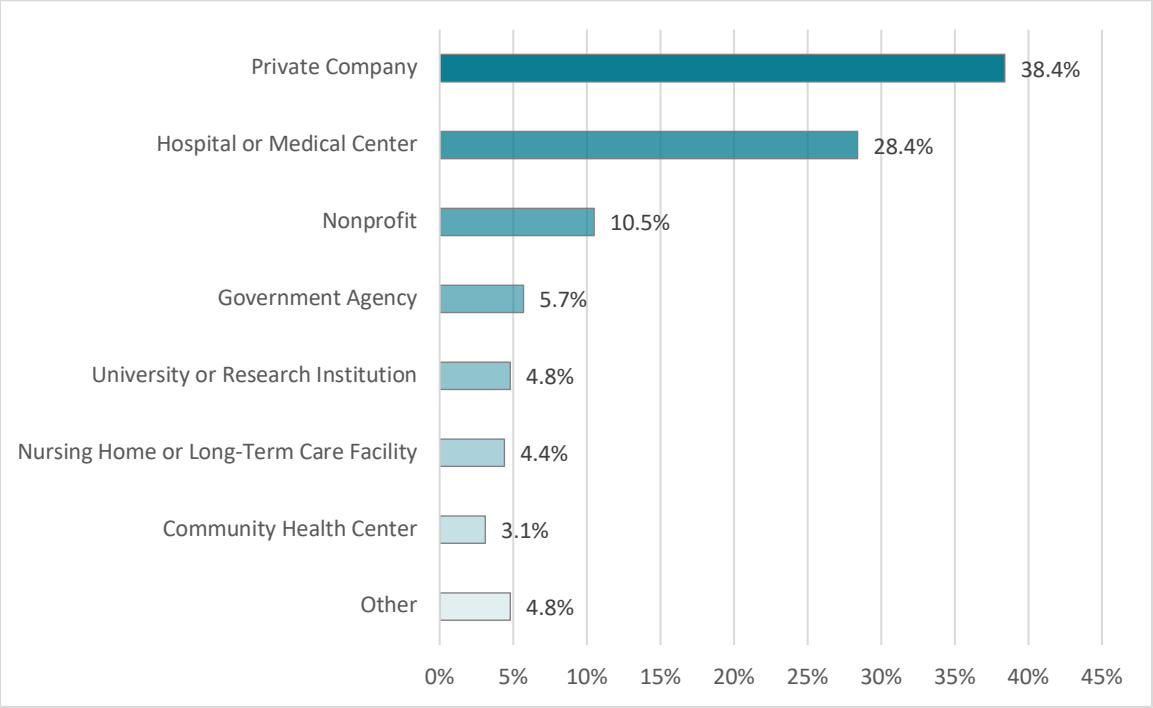

Figure 2 depicts the distribution of roles represented in the survey. The largest proportion of respondents identify as nurses (25.6%), closely followed by physicians (24.7%), indicating that direct clinical care providers made up the majority of the sample. Administrative staff (14.1%) and nurse practitioners or physician assistants (13.2%) also represent significant portions, while smaller groups included medical technicians, health care IT professionals, and allied health professionals (each 2.6%), as well as researchers/scientists (1.3%) and graduate student health care practitioners (0.9%). Medical students comprised a very small proportion at 0.4%. The "other clinical staff" category accounted for 11.9%, suggesting a variety of additional roles not explicitly listed.

Figure 2. Roles/Professional Titles Represented in Sample

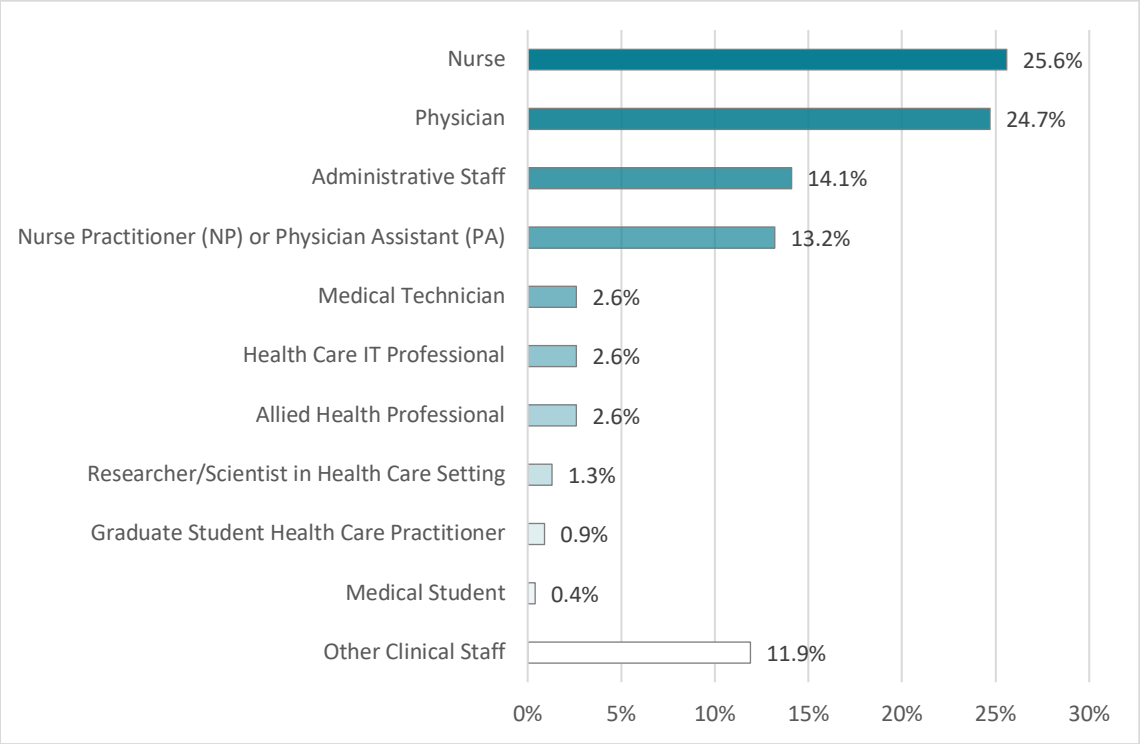

Figure 3 illustrates the patient segments served by respondents in the survey. The majority of respondents indicated serving all demographics (51.5%) and the elderly (52.4%). Other prominent segments of the population served include individuals with disabilities (40.6%), patients with chronic conditions (41.0%), and children (37.6%). Additionally, those serving low-income individuals or families (37.1%), patients with diabetes (36.7%) patients with mental health conditions (31.4%), veterans (27.1%), pregnant women and new mothers (24.9%) suggest a focus on populations requiring tailored support. Notably, patient segments served by respondents in the survey also include underserved urban (17.9%) and rural communities (24.9%), reflecting efforts to bridge health care access gaps. Food insecurity (18.3%) also emerges as a critical area of service, tying health care needs to broader social determinants.

Figure 3. Patient Segments Served

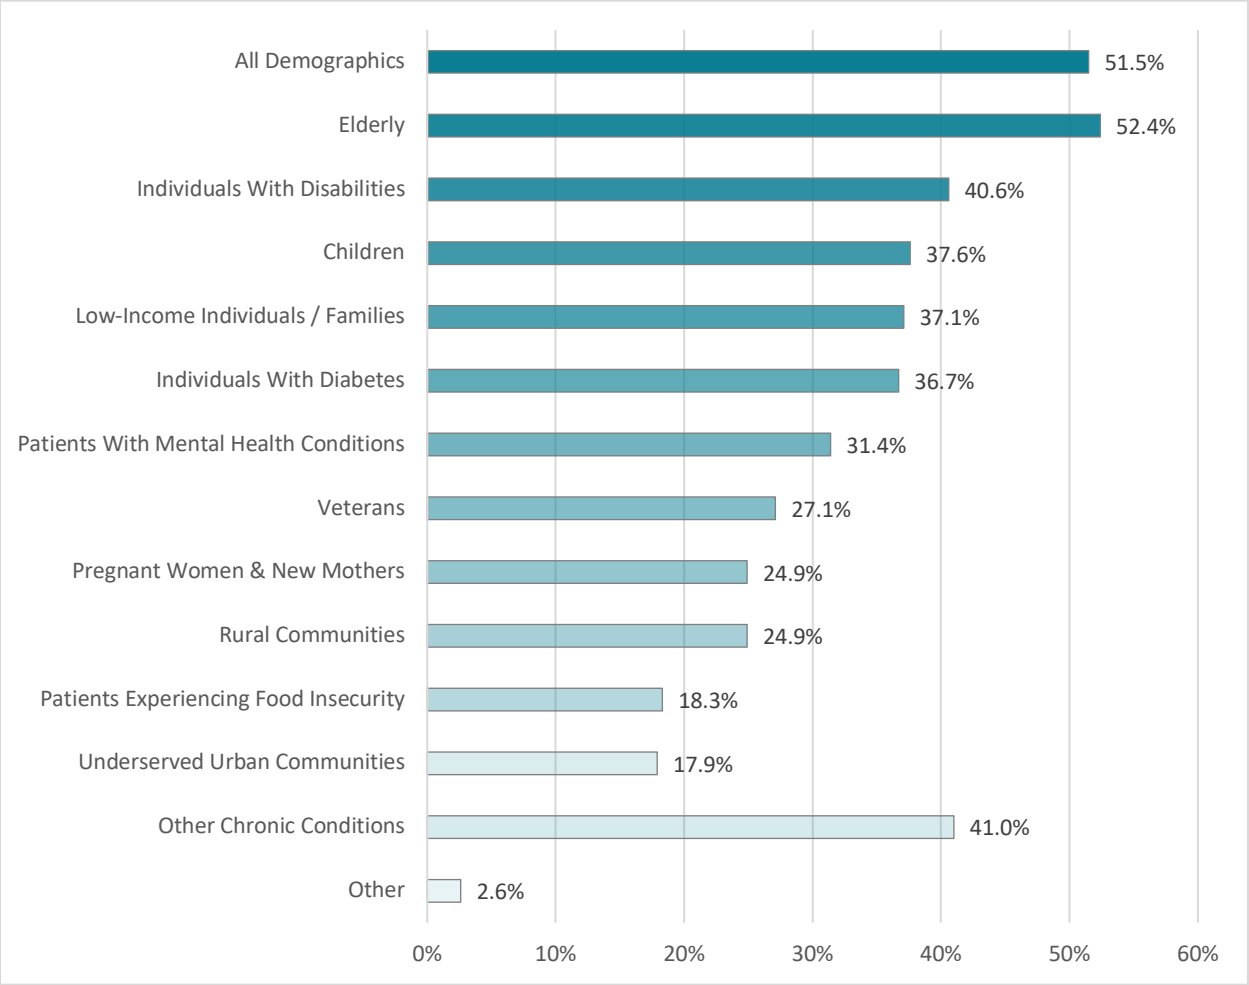

Note: Totals exceed 100% as survey respondents could select multiple patient segments served.

Familiarity and Trust in AI

The survey assessed baseline awareness and trust in Health AI. Figure 4 indicates that 56.4% of respondents are somewhat familiar with AI technologies, suggesting moderate awareness. A smaller group (16.7%) reported being very familiar, while the remaining participants were largely or fully unfamiliar with AI technologies.

Figure 4. Participant Familiarity With AI Technologies in Health Care

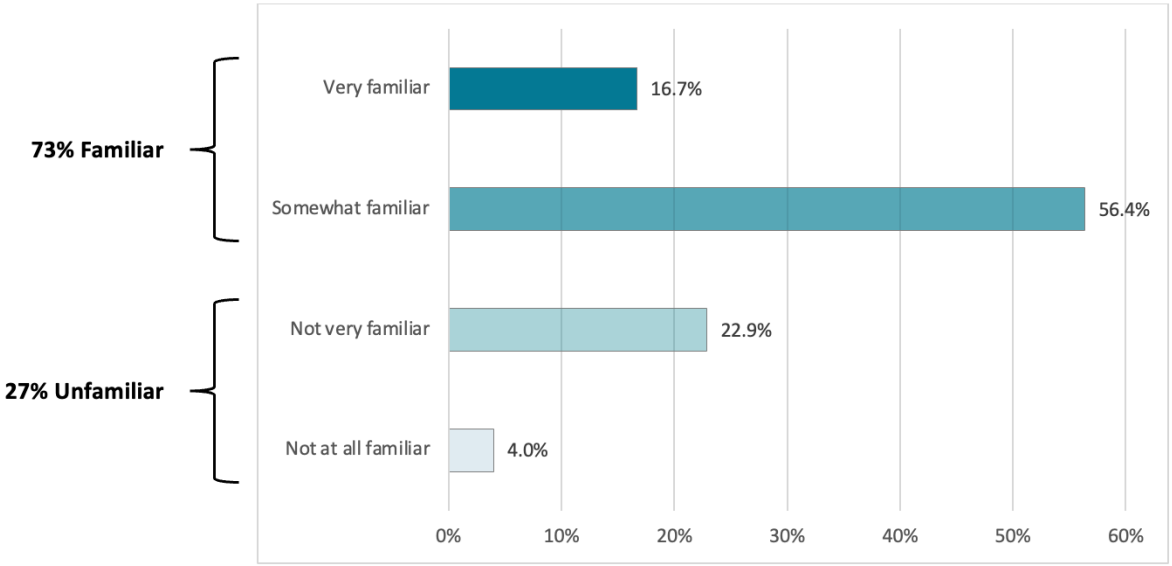

Figure 5 examines current levels of trust in AI technologies for health care delivery. While 41.9% expressed some trust, 32.2% were neutral, and 16.7% showed some distrust. Only 4.0% indicated complete distrust, while 5.3% expressed complete trust. These findings suggest cautious optimism tempered by hesitancy, consistent with earlier research.

Figure 5. Baseline Trust in AI Technologies for Health Care Delivery

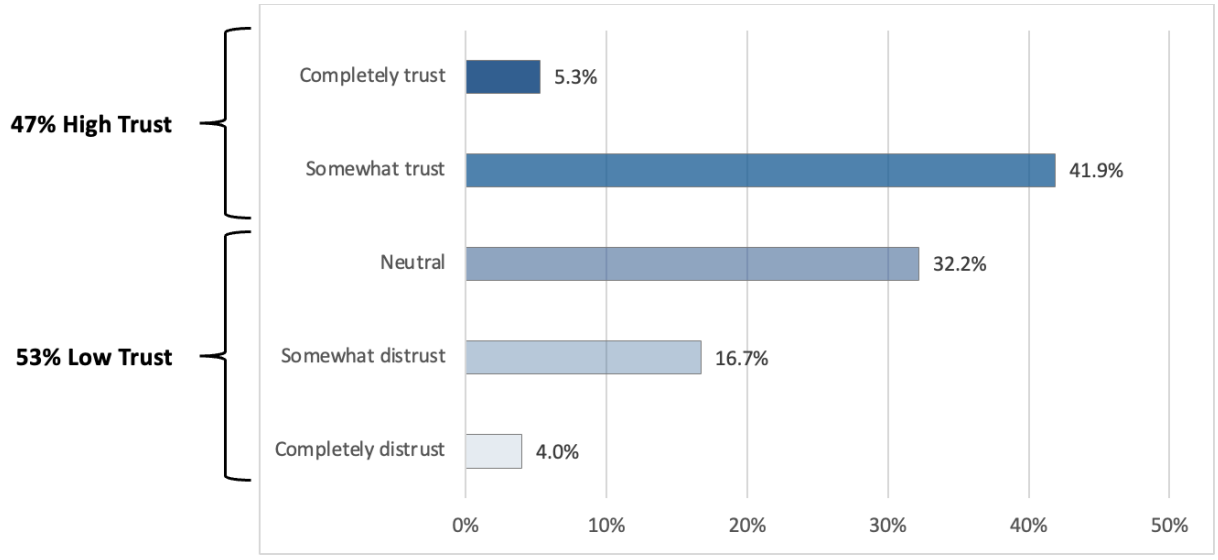

We have used trust in several additional analyses presented next by stratifying survey responses by those who expressed high baseline trust (somewhat trust or completely trust) versus low

trust (neutral, somewhat distrust, or completely distrust). Using this stratification, 47.2% of our participants were in the high trust group, and 52.8% were in the low trust group.

Those with greater baseline familiarity in AI had higher levels of current trust in AI technologies. In fact, when the data from Figures 4 and 5 are combined in Figure 6, we can see that 71.1% of those who perceived themselves to be somewhat or very familiar with the use of AI in health care had high trust in AI to support health care delivery. By comparison, only 42.9% of those currently "not at all" or "not very" familiar with the use of AI in health care had high trust in AI to support health care delivery.

Figure 6. Relationship Between Familiarity and Trust in AI

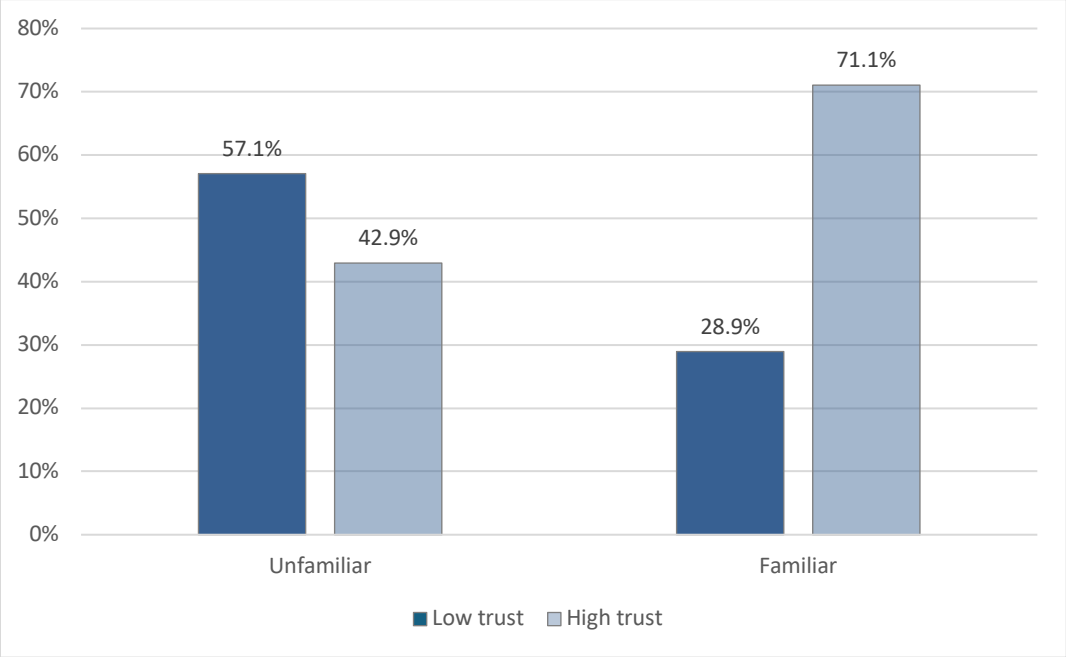

Figure 7 presents the themes that emerged from an analysis of open-ended survey data. It shows the frequency of the various themes influencing participants' degree of trust in AI technologies for health care delivery among low trust and high trust groups. Building trust was the most frequently mentioned factor by both groups, when asked what they perceive impacts their current level of trust. (Note the high trust group included a majority of respondents stating they "somewhat trust" AI. Only 5.3% said that they completely trust AI.)

The largest distinction between low and high trust groups lies in the relative frequency with which factors are mentioned to explain their level of trust in AI. Among low trust respondents, over half (50.5%) cited building trust as a core issue, followed closely by data bias and accuracy (33.0%) and maintaining data security, privacy, and anonymity (12.8%). These participants expressed deep concern over the fairness, reliability, and safety of AI systems.

In contrast, high trust respondents were more likely to mention the practical benefits of AI: although 28.4% mentioned building trust, 23.5% mentioned efficient and effective care, followed by diagnostic capabilities (11.8%), and automating documentation (10.8%) as reasons for their confidence in AI. Additionally, high trust respondents were far more likely to mention the accuracy of data (11.8%) compared to their low trust counterparts (1.8%).

Figure 7. Issues Impacting Trust in AI in Health Care, Stratified by Providers' Trust in AI

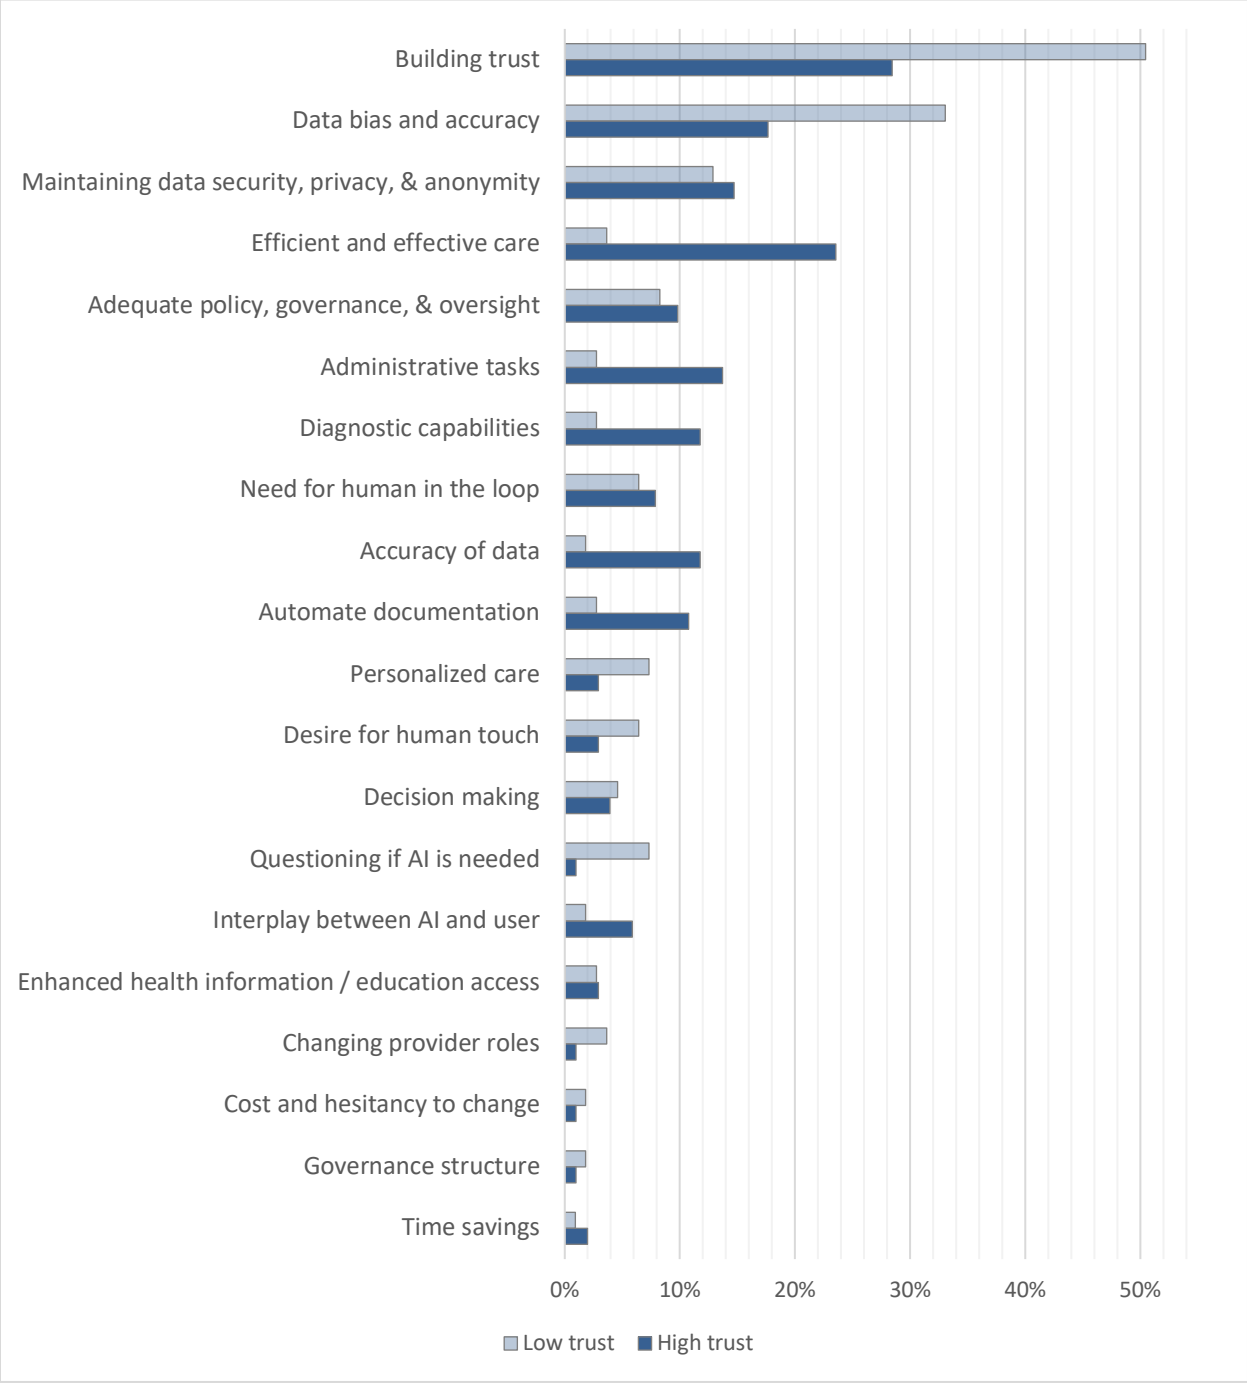

Note: Totals exceed 100% as survey respondents could select multiple issues/factors.

For this survey, we used the following definitions of AI and related disciplines:

1) Artificial Intelligence (AI): In health care, AI refers to the use of advanced algorithms and software to simulate human intelligence for tasks like diagnosing diseases, personalizing treatment plans, and automating administrative processes.
2) Machine Learning (ML): Machine learning in health care is a subset of AI that uses statistical models and algorithms to predict patient outcomes, identify trends, and assist in clinical decision making without needing explicit programming for each task.
3) Data Science (DS): Data science in health care involves techniques like data mining, predictive modeling, and statistical analysis to uncover patterns and trends that can lead to better health outcomes.

As depicted in Figure 8, the majority of respondents reported a moderate or strong grasp of these topics (combined total of 63.4%). A smaller proportion, 6.6%, said they understand these topics as extremely well, suggesting a subset of respondents with a high level of confidence in their knowledge. Conversely, a notable percentage of 29.9% expressed less confidence, which includes 5.7% who admitted they do not understand AI and related concepts well at all.

Figure 8. How Well Participants Understand AI and ML, at a Basic Level

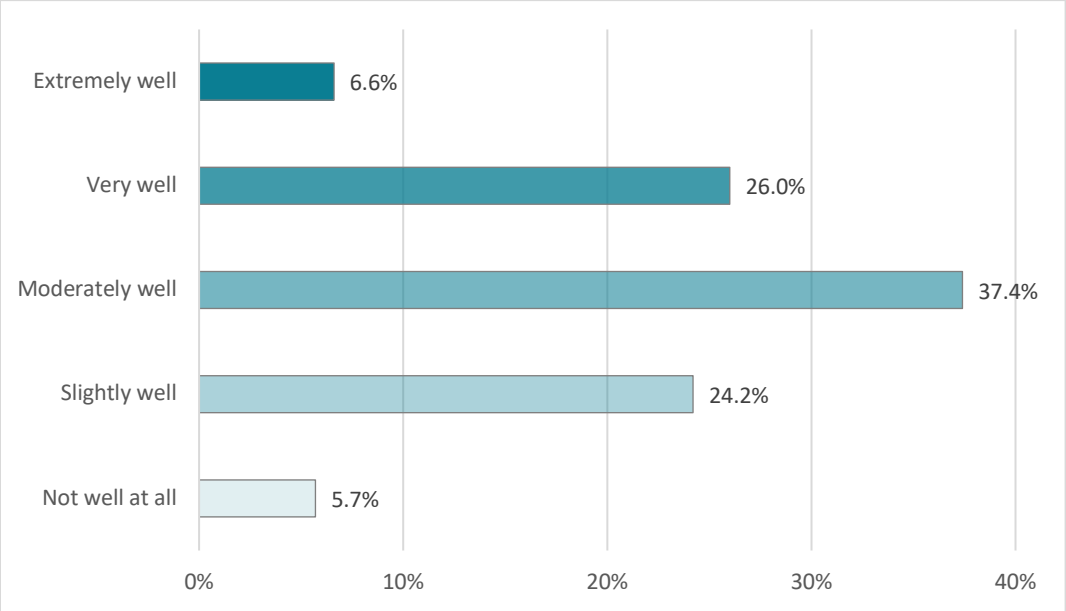

Note: Participants were given a basic definition of AI and ML to consider.

Benefits and Barriers for Patients and Providers

Figure 9 highlights the most significant benefits that respondents believe AI could bring to health care delivery. The majority, at 38.8%, identified streamlined administrative tasks as the primary advantage, suggesting that reducing the burden of paperwork and bureaucratic processes would

be a welcomed priority when integrating AI into health care. Following this, enhanced patient outcomes (24.2%) and improved diagnostic accuracy (23.3%) were recognized as critical benefits, reflecting the potential of AI to directly improve the quality of care and the precision of medical decisions. A smaller proportion, 11.0%, emphasized the importance of personalized treatment plans, while 2.6% of respondents noted other benefits. *These findings underscore the perception that AI's greatest impact lies in its ability to improve operational efficiency and patient care outcomes, with a focus on addressing administrative inefficiencies and enhancing clinical accuracy.*

Figure 9. Most Significant Benefit Perceived From AI to Health Care Delivery

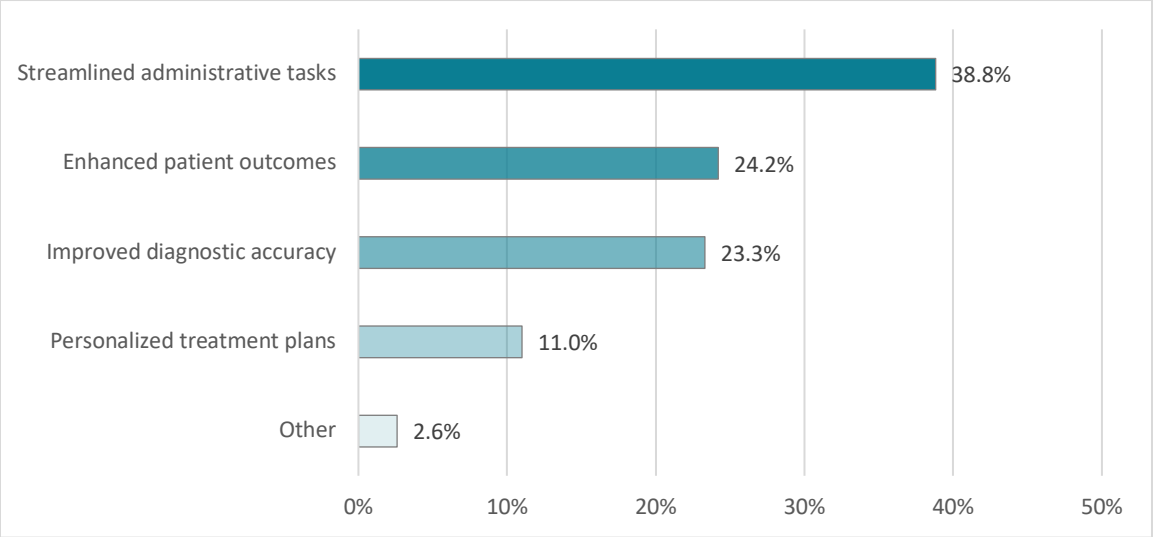

Figure 10 identifies the primary barriers to adopting AI technologies in safety-net health care settings. The most significant obstacle, cited by 31.9% of respondents, is concerns about data privacy and security, reflecting widespread apprehension about the protection of sensitive patient information. The second most frequently noted barrier, at 19.5%, is insufficient training or knowledge about AI, emphasizing the need for education and workforce development to support AI implementation. Lack of funding for AI implementation follows at 17.3%, highlighting financial constraints as a critical hurdle, particularly in settings with limited resources. Additionally, limited technological infrastructure (14.6%) and resistance to change among staff or leadership (13.7%) underscore challenges related to organizational readiness and cultural acceptance of new technologies. A small percentage of respondents (3.1%) cited other barriers.

Figure 10. Primary Barrier Perceived to Prevent Adoption of AI in Safety-Net Health Care Settings

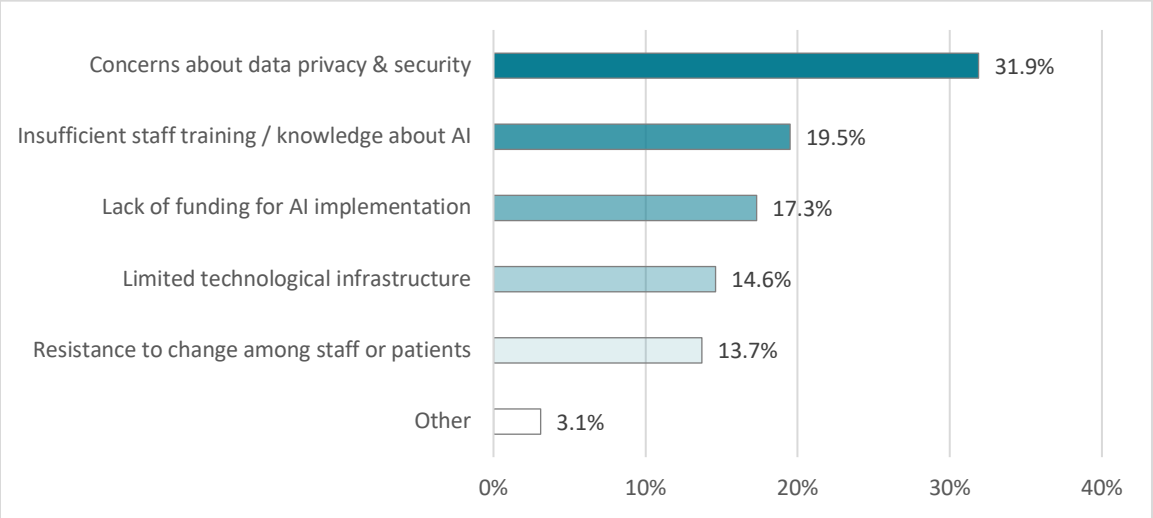

Figure 11 explores how providers perceive the responsiveness of patient towards AI-based tools to support better health outcomes (we did not survey patients directly). Most respondents expressed varying levels of optimism, with 39.6% indicating that they believe that their patients would be somewhat responsive to AI tools. Further, 24.4% believe that their patients would be neutral, suggesting uncertainty about how patients might react. Only 5.8% of respondents felt that their patients would be very responsive. On the other hand, a notable proportion expressed some doubt about patient responsiveness (a combined 30.2%), indicating some skepticism about patients' willingness or ability to engage with AI technologies for their health care.

Figure 11. Provider Perceptions of Patient Responsiveness to AI Tools in Health Care

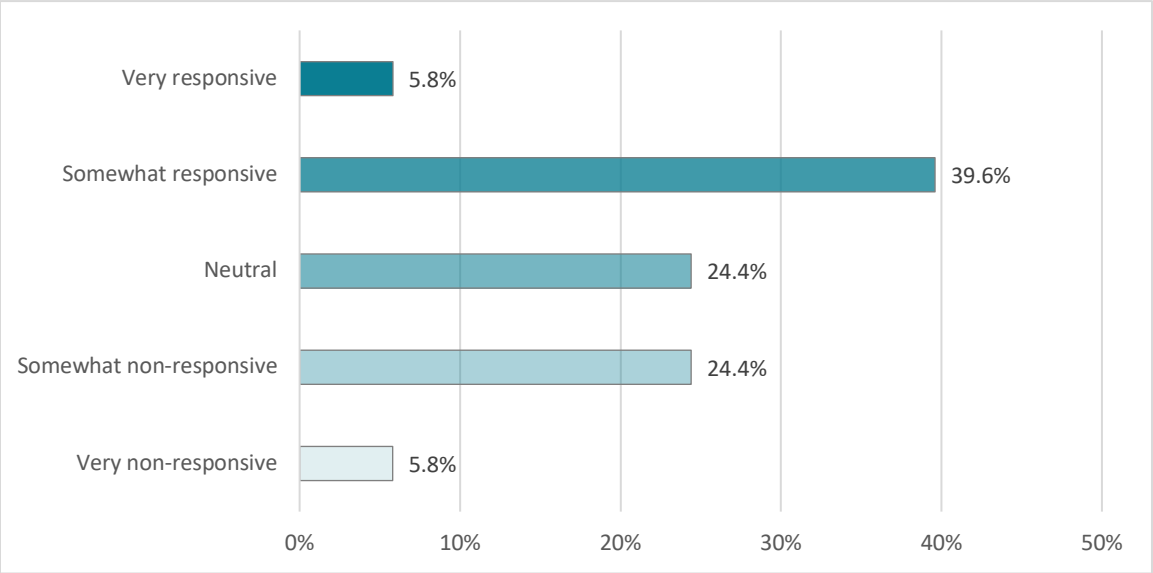

Figure 12 highlights differences between low trust and high trust respondents regarding the factors influencing their perceptions of patient responsiveness to AI technologies for better health outcomes. Both low and high trust respondents believe that "building trust" in Health AI is the most important factor in whether or not their patients will respond favorably to the use of AI to support their health care. Similarly, both groups identified "unequal access to technology/ infrastructure" as the second most significant factor. The two groups varied in their third cited factor: For low trust respondents, the "interplay between AI and user" (or "cost and hesitancy to change) was the third leading factor, whereas "efficient and effective care" emerged as the third highest factor among high trust respondents.

Figure 12. Issues Impacting Perceived Patient Responsiveness to AI in Health Care, Stratified by Providers' Trust in AI

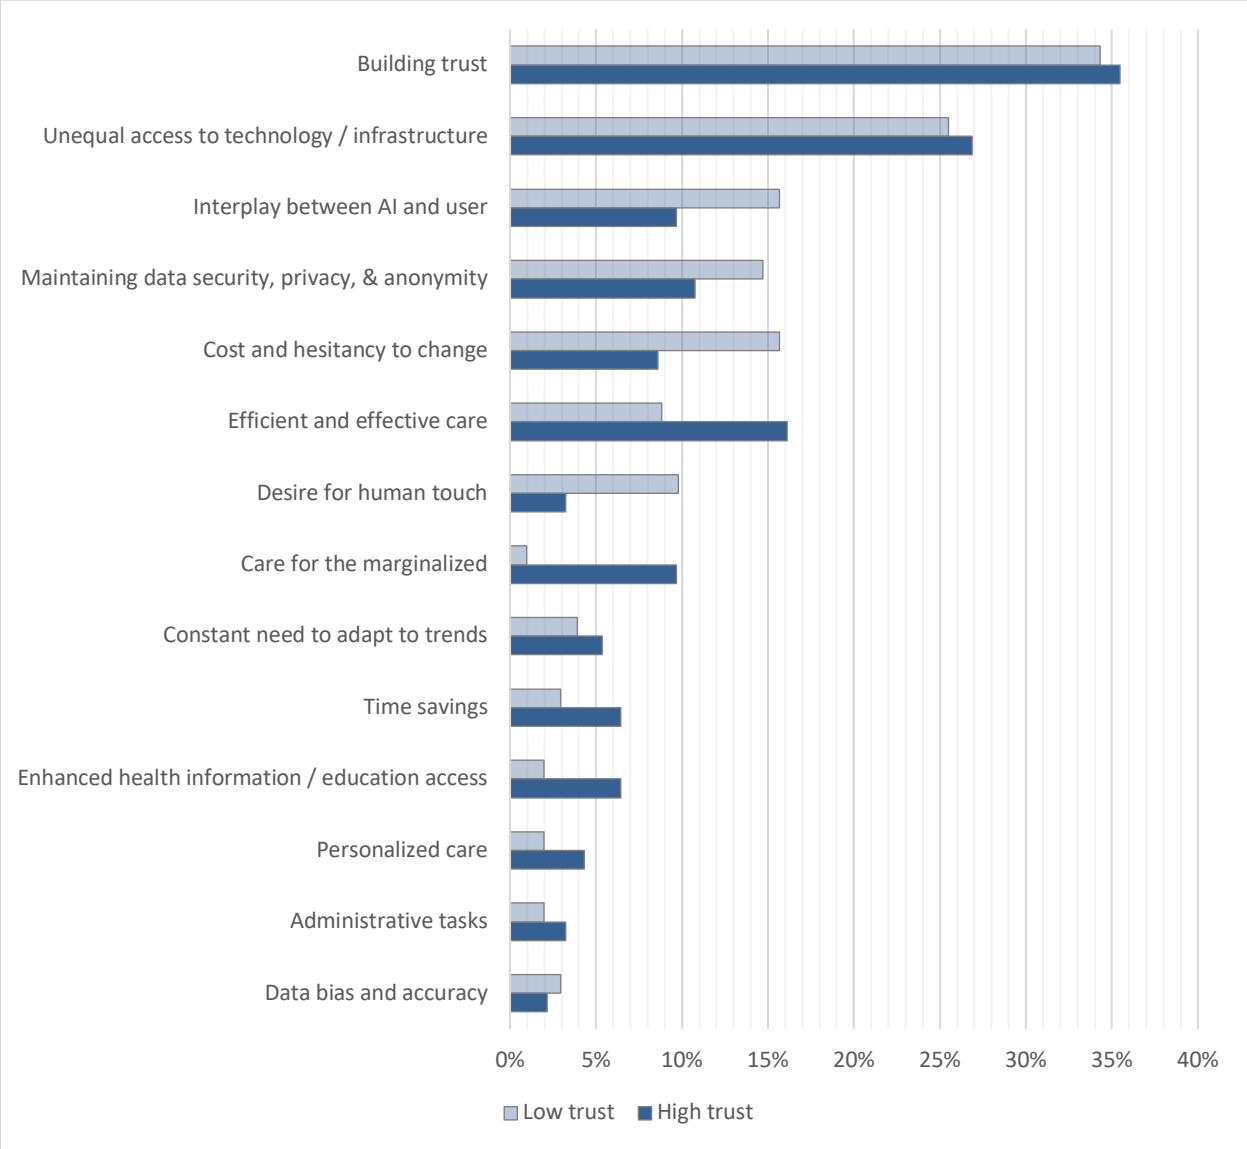

Figure 13 illustrates respondents' confidence levels in integrating AI technologies into current health care workflows. Approximately one third, or 35.4%, described themselves as somewhat confident, indicating moderate confidence in adopting AI tools. Meanwhile, 48.7% of respondents were either neutral or somewhat not confident, suggesting a mix of uncertainty or hesitation in implementing these technologies. Only 7.1% reported being very confident about integrating AI technologies into their current health care workflows.

Figure 13. Participant Confidence in Integrating AI Technology Into Health Care Workflow

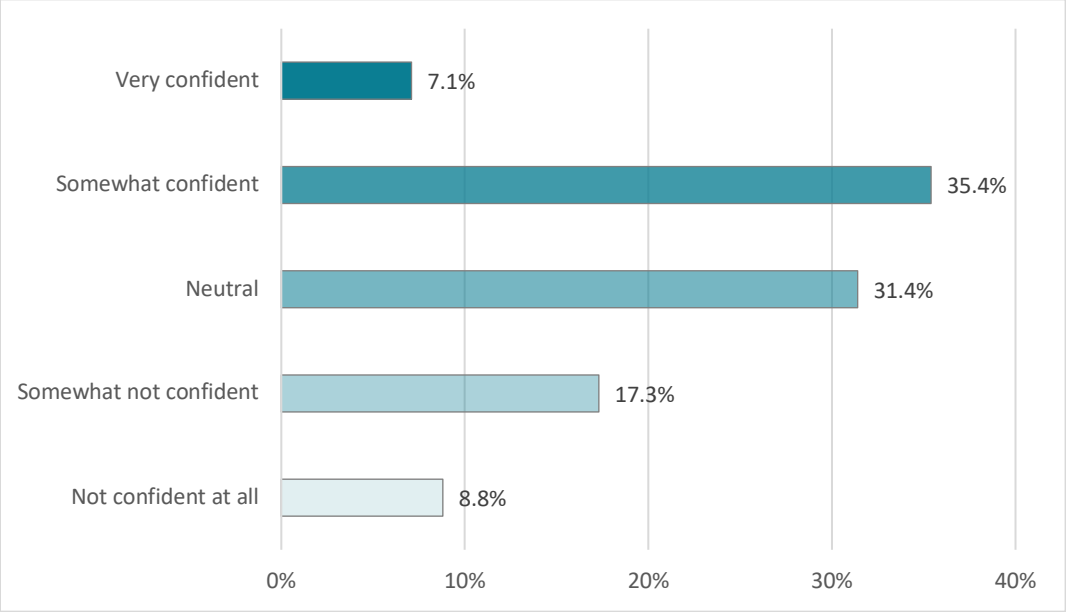

Figure 14 depicts how low trust and high trust respondents perceive the potential benefits of AI for rural and safety-net services populations. Among all respondents, “efficient and effective health care” emerged as the leading Health AI benefit. For low trust respondents, “decision making” (38%) was the second leading benefit, followed by “administrative tasks.” For high trust respondents, “administrative tasks” (32%) was the second leading benefit, and “personalized care” was third. Telehealth/telemedicine was significantly more emphasized by high trust respondents compared to low trust respondents, highlighting that individuals who trust AI technologies may recognize the potential of AI-powered telehealth solutions to bridge gaps in service delivery and enhance care availability for remote or resource-limited populations.

Figure 14. Potential Perceived Benefits of Health AI for Safety-Net and Rural Populations, Stratified by Providers’ Trust in AI

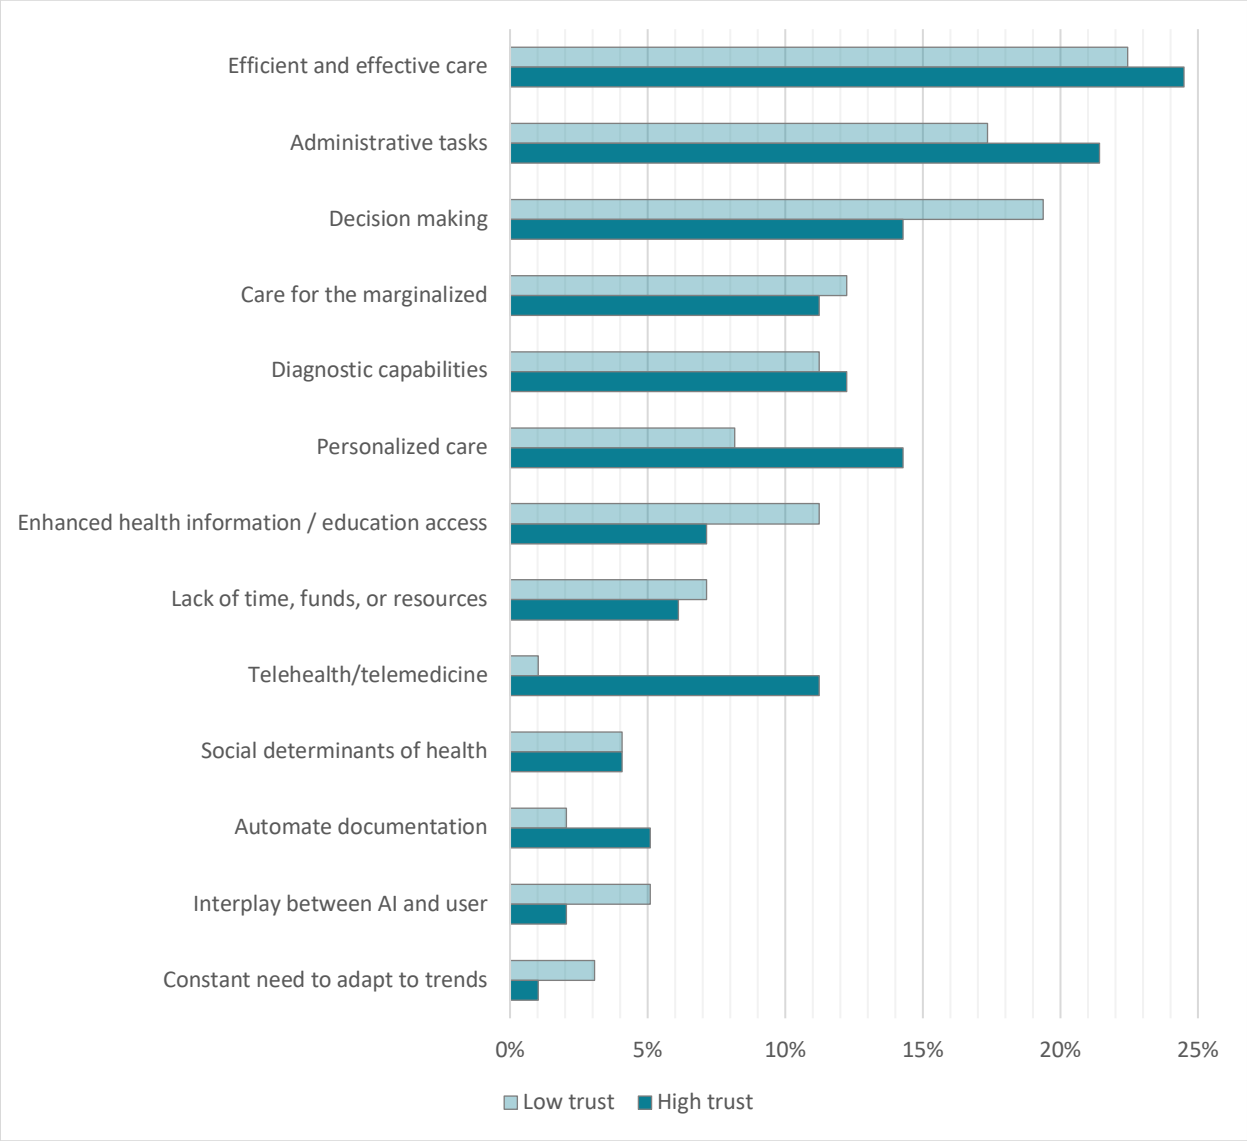

Figure 15 highlights additional factors that respondents believe are relevant to their organization's adoption and deployment of Health AI. Cost and hesitancy to change were significantly more emphasized by low trust respondents, reflecting their concerns about financial barriers and organizational resistance to AI integration. On the other hand, high trust respondents were much more likely than their low trust counterparts to emphasize adequate policy, governance, and oversight as barriers to AI adoption among health care organizations. This group was also more likely to emphasize issues related to workforce and training as barriers. These findings illustrate that addressing financial and cultural resistance is key for low trust groups, while high trust groups prioritize infrastructure equity as a critical enabler of AI adoption.

Figure 15. Considerations for Organization Adoption of Health AI, Stratified by Providers' Trust in AI

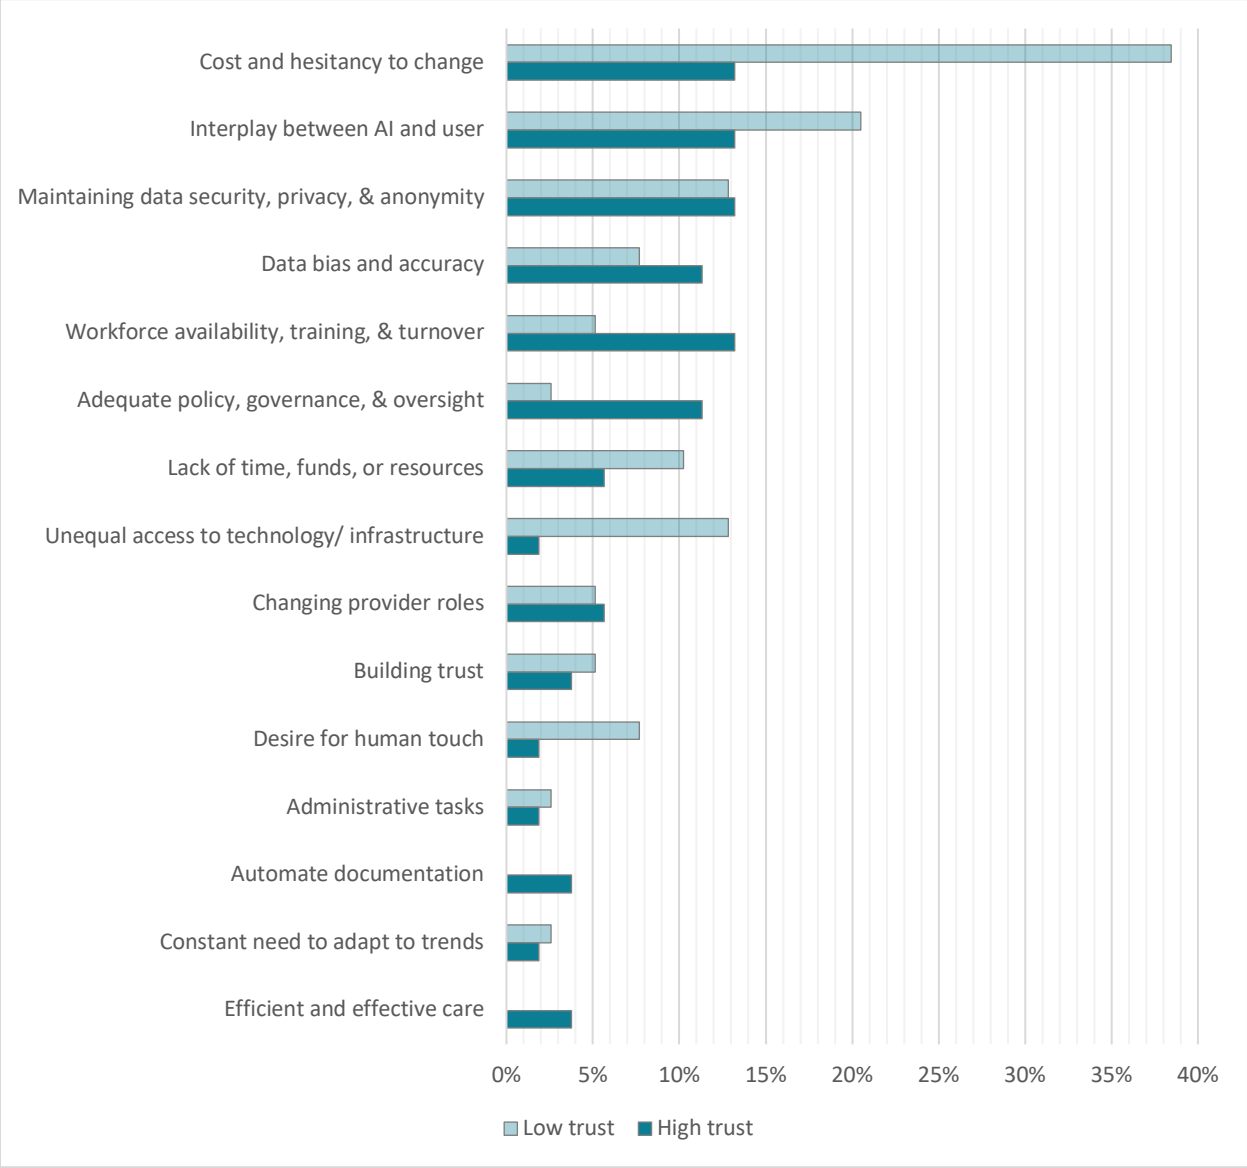

Reaction to Specific Health AI Concepts/Tools

The survey included four vignettes about four realistic but fictional AI tools designed to assess appeal and concern for specific Health AI concepts and technologies. See Appendix A for the full vignettes, which can be summarized as follows:

- DocuScribeAI is an AI tool designed to assist in transcribing and organizing clinical notes during patient interactions.
- TherapiaAI analyzes real-time data from patient histories, genetic information, and current health conditions. It suggests customized therapeutic plans based on clinical guidelines and patient-specific factors.
- HealthRiskAI integrates patient-specific social determinants of health and provides real-time risk assessments.
- CommCare analyzes communication preferences and engagement patterns, offering insights to improve patient-provider interactions.

Appeal and concern were measured on 5-point scales, as shown next.

Appeal

1) No Appeal
2) A Little Appeal
3) Some Appeal
4) Highly Appealing
5) Extremely Appealing

Concern

1) No Concern
2) A Little Concerning
3) Some Concern
4) Highly Concerning
5) Extremely Concerning

Figure 16 shows the mean appeal and concern rating for each vignette, including the 95% confidence interval (CI). *Across all AI concepts/tools, the average appeal is greater than the average level of concern.* The appeal of CommCare is the lowest and is significantly lower than the most appealing concept, DocuScribeAI.

Figure 16. Average Appeal and Concern Rating for Each Vignette

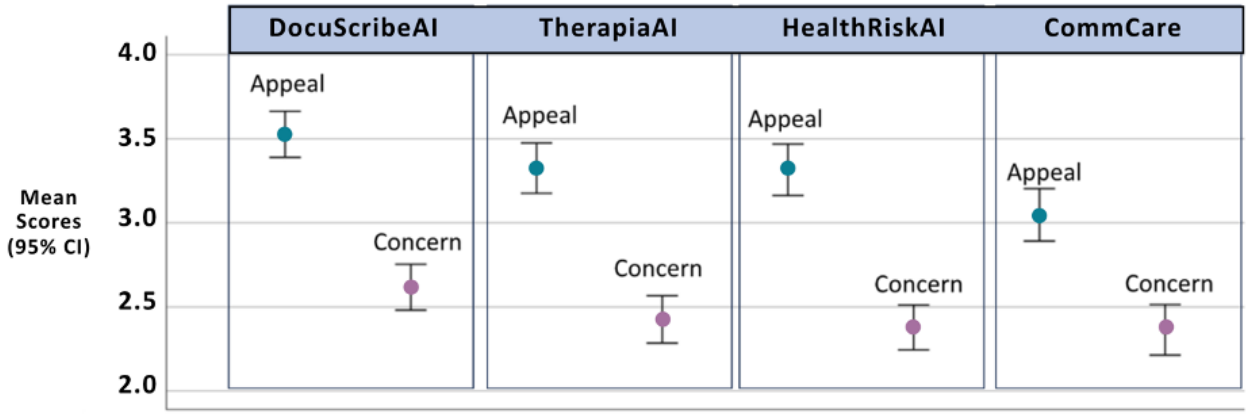

Participants were also asked: “In a sentence or two, tell us a little more about the appeal you perceive [for each fictional AI tool] and any concerns you have.” Responses were analyzed thematically and coded. Figures 17 through 20 show these responses for each vignette, with the results stratified by participants as follows:

- Those labeled **Promise** in the figures gave an appeal rating greater than or equal to their concern rating.
- Those labeled **Peril** in the figures gave a concern rating greater than their appeal rating.

Figure 17 summarizes reactions to DocuScribeAI based on this promise versus peril contrast. These data also highlight that, while DocuScribeAI is widely recognized for its ability to automate documentation and save time, skepticism remains about its accuracy and security and there is a need to keep humans in the loop. The tool's potential to streamline administrative tasks and improve efficiency is promising, yet concerns about data reliability and privacy must be addressed to build trust among skeptical stakeholders. Balancing automation with human involvement and ensuring a governance network is in place are critical to maximizing the perceived benefits while mitigating risks. Participants who more strongly recognized the promise of DocuScribeAI were more likely to mention the automation of documentation (47.4% promise, 21.4% peril). Conversely, participants who recognized the peril of DocuScribeAI were more likely to: have concerns about the accuracy of data (46.4% peril, 33.3% promise); recognize that humans are needed in the loop (46.4% peril, 26.9% promise); question if AI is needed (21.4% peril, 9.4% promise); and comment on the need to focus on the interplay between language, culture, and physicality (14.3% peril, 4.1% promise).

Figure 17. DocuScribeAI: Reasons for Appeal and Concern, Stratified by Participants With Higher Promise vs. Higher Peril

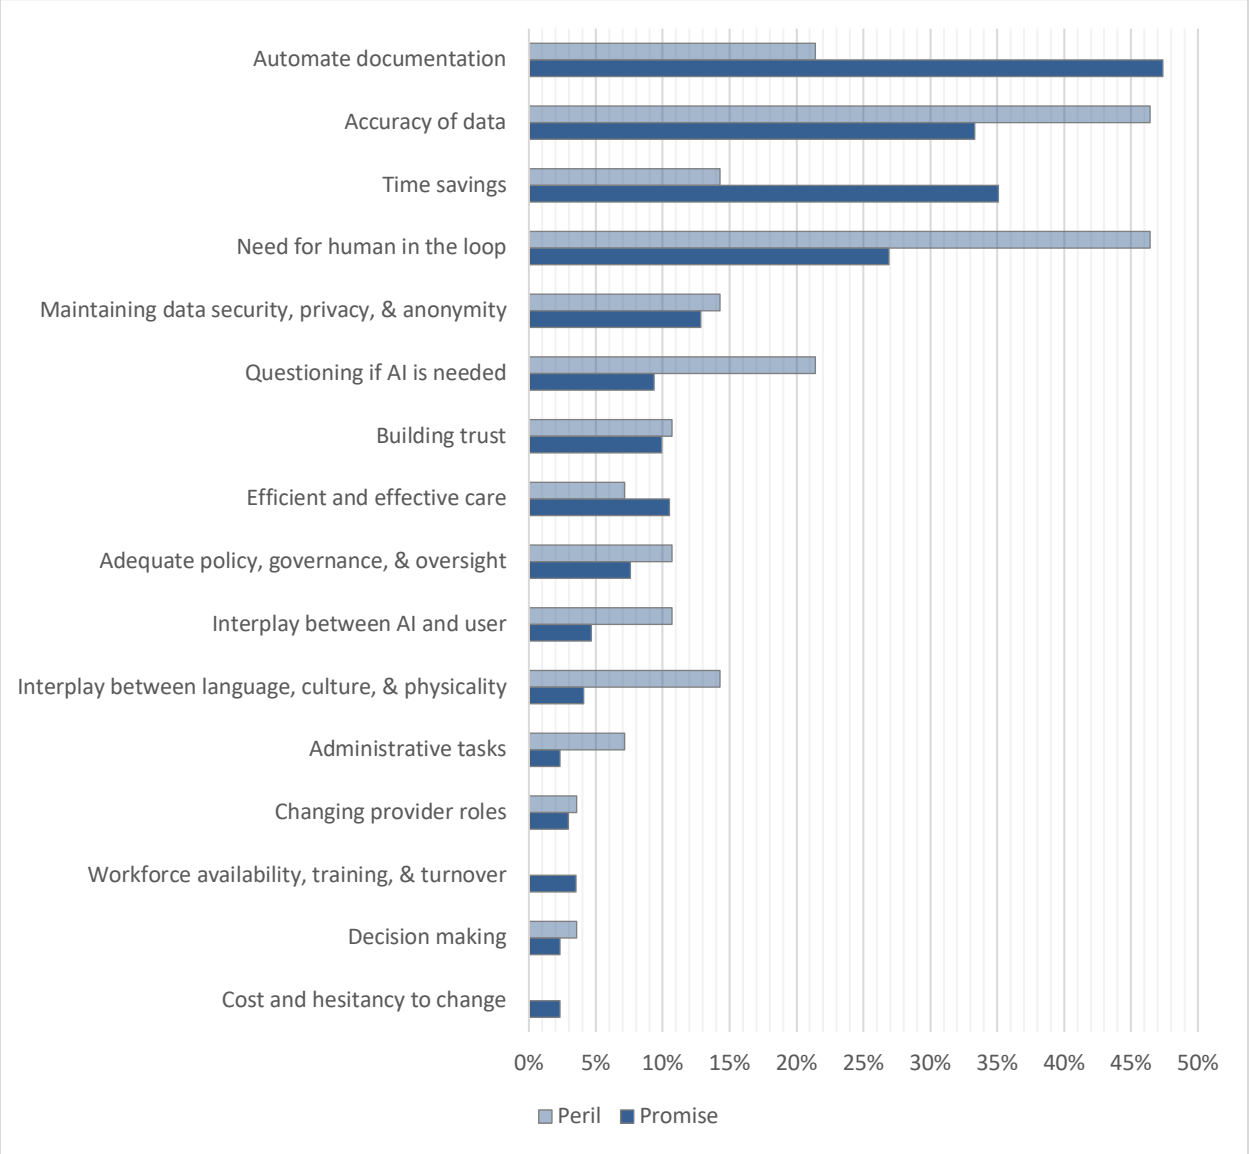

Figure 18 highlights perceptions of TherapiaAI, a tool designed to enhance treatment personalization, which diverged between those who saw promise with the tool and those who saw peril. Both groups highlighted the need for human oversight (20.8% peril, 18.1% promise) and the potential for personalized care (20.8% peril, 18.1% promise), but their priorities differed. Peril respondents focused more on data bias and accuracy (25.0%), data privacy (16.7%), trust (16.7%), and the desire for human touch (16.7%), reflecting caution toward AI-driven care. In contrast, those who viewed promise in TherapiaAI as a tool prioritized benefits like efficient and effective care (18.1%), enhanced health information and education access (14.1%), and time savings (8.7%), indicating greater confidence in the tool's ability to improve both patient and provider experiences. These differences suggest that perceptions of risk versus benefit shape how individuals evaluate the role of AI in health care innovation.

Figure 18. TherapiaAI: Reasons for Appeal and Concern, Stratified by Participants With Higher Promise vs. Higher Peril

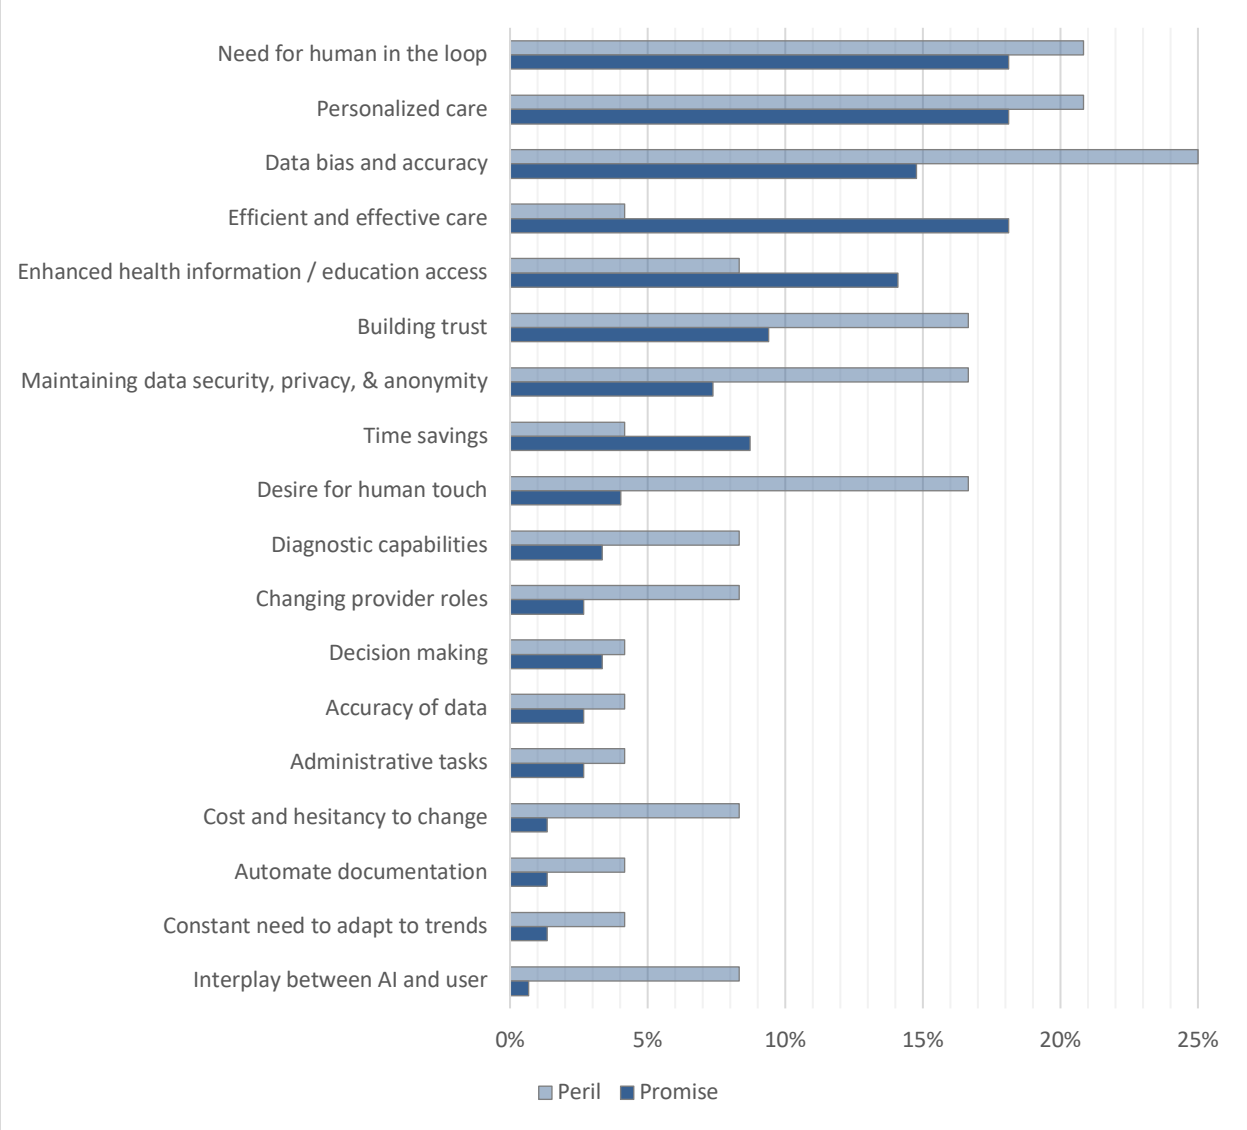

Figure 19 highlights perceptions of HealthRiskAI, a tool that integrates patient-specific social determinants of health (SDoH) and provides real-time risk assessments. Respondents who saw the tool with peril emphasized concerns around data accuracy (31.8%), privacy and security (22.7%), and the need for human oversight (22.7%), reflecting caution about automation and trust in AI. In contrast, those who viewed the tool with promise highlighted potential benefits such as efficient and effective care (13.1%), enhanced decision-making (10.9%), and time savings (5.1%), indicating optimism about AI's ability to support clinical practice. Both groups noted the importance of data accuracy and human involvement, but their different priorities reveal how perceptions of risk versus benefit can shape provider evaluations of AI in health care.

Figure 19. HealthRiskAI: Reasons for Appeal and Concern, Stratified by Participants With Higher Promise vs. Higher Peril

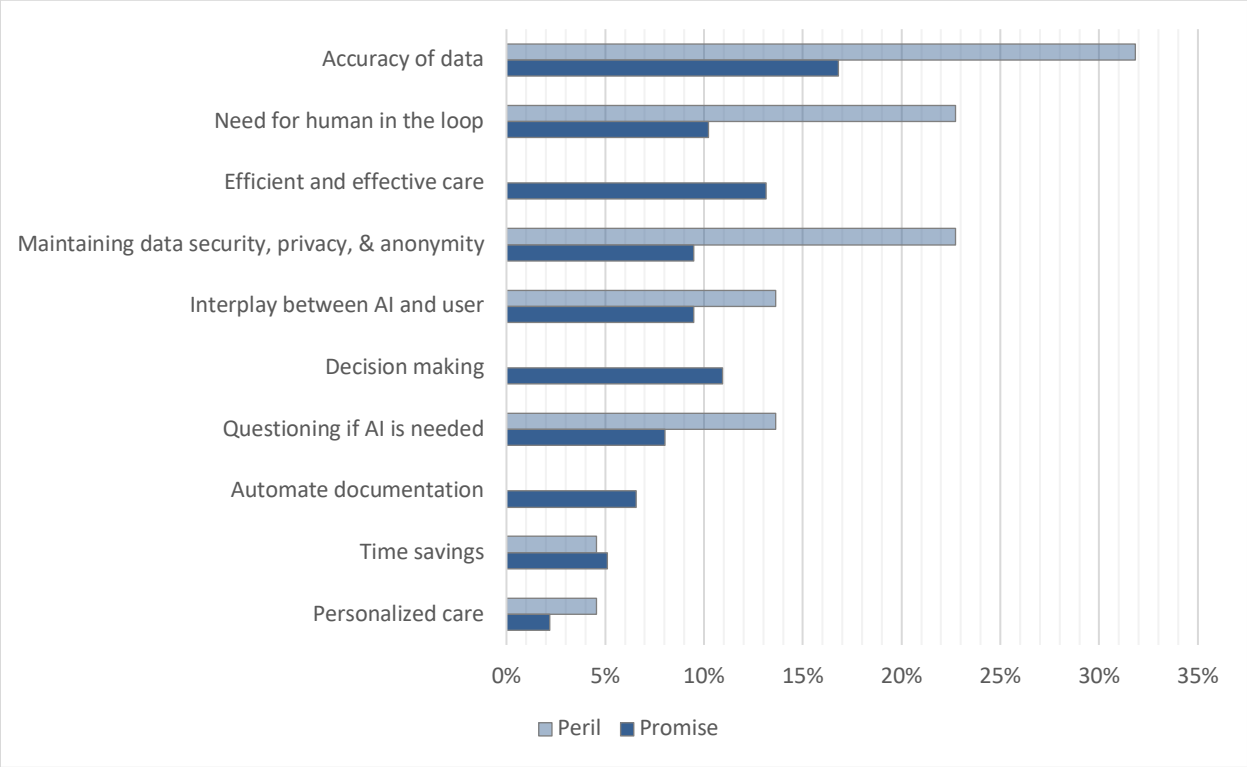

Figure 20 explores perceptions of CommCare, an AI tool designed to evaluate patient communication patterns to enhance engagement and satisfaction. Efficient and effective care was mentioned approximately three times more frequently by respondents who saw promise in AI (23.8% promise, 8.3% peril). Conversely, respondents more concerned with the perils of AI questioned if AI was needed approximately four times more frequently (37.5% peril, 9.2% promise).

Figure 20. CommCare: Reasons for Appeal and Concern, Stratified by Participants With Higher Promise vs. Higher Peril

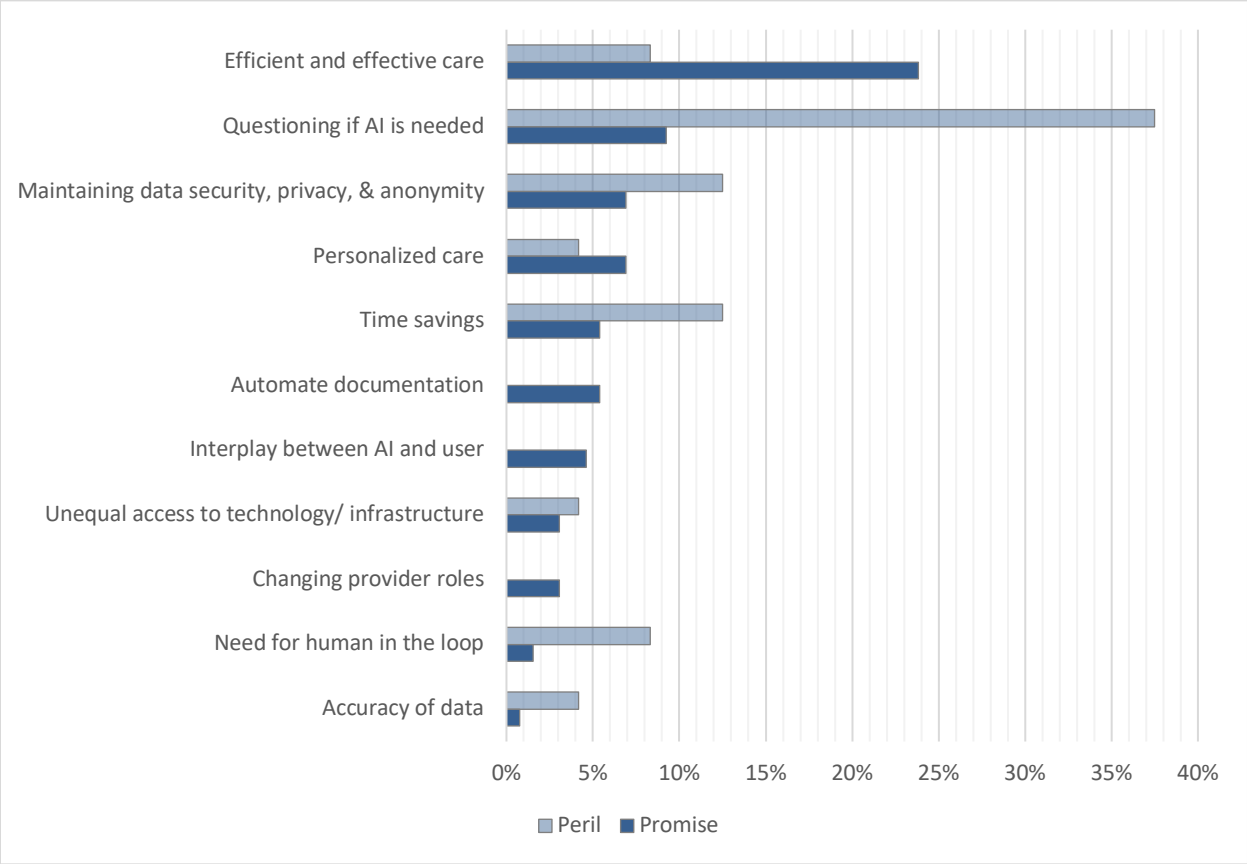

PART FOUR: DEEPER QUALITATIVE INSIGHTS

In this section of the strategy paper, we turn our attention to the open-ended questions and responses collected as part of the survey and during our thought-leader interviews. Whereas the survey/quantitative data in Part Three may tell us *what* safety-net health care practitioners and stakeholders feel about a specific aspect of Health AI, interview/qualitative data can help explain the reasons *why* they have such feelings and perceptions.

Our analysis of interview transcripts provided further evidence for the importance of trust in AI to be able to introduce these tools into real-world clinical environments. Indeed, the need to build trust was one of the most common themes in interviews, suggesting that future implementation of Health AI must enhance confidence in AI's ability to improve health care delivery and patient outcomes. Our interviewees identified several **factors associated with enhancing health care practitioners' trust in Health AI**, including:

- Addressing and reducing bias in care.
- Building an AI Literate Healthcare Workforce
- Supporting human-centered health care.
- Augmenting and improving care efficiency and effectiveness.
- Extending care via remote patient monitoring, telehealth.
- Navigating linguistic and other cultural differences successfully.

Mitigating Bias in Care

As we've noted, one of the most significant challenges in Health AI is the threat of bias, a theme echoed frequently among study participants. This suggests that health care stakeholders are aware of how bias can undermine the performance of Health AI systems. When introducing Health AI into a health care setting, providers sought assurance that the data used to train those AI systems was inclusive and representative of the population receiving their care. Further, outcomes from AI are not innately trusted; AI must demonstrate accuracy in algorithm function and output data.

> *"If we take existing data, which we know is going to be biased for all sorts of different reasons, and we train a new AI system on that [data], then we're going to have to expect that that AI system is going to structuralize those biases in that algorithm, and is going to produce biased decision making or prioritization. This worries me a tremendous amount, the fact that those biases are just going to make additional inequity in our decision making and drive the disparities even further." – Researcher*

Interestingly, some interviewees were quick to point out that humans, though well-intentioned, are also not completely accurate, trustworthy, or unbiased. Some AI advocates note that AI can be trained to mitigate the known biases of health care practitioners and, as a result, play a role in

building more equitable health care models. AI, for example, could assist in the reduction of human biases in diagnosis and treatment plans. This view of AI as a mitigator of bias was reflected in our interviews.

> *"I think that AI can overcome some of the implicit biases that humans have. So, can AI craft a message, you know, incorporating race and ethnicity or age into their educational thinking. Can it? Yeah, I'm sure it can. Can humans do it? No, humans cannot do it. …We, for the most part, repeat what we know and how we've grown up. We have our upbringing, [our] background shows off in different ways, whether we consciously try to modify our answers or not, many times it just comes through. But I think the AI may be able to be less biased than humans." – Senior Leader in Health Care*

Building an AI-Literate Healthcare Workforce

A recurring theme in interviews was the critical need to invest in the AI literacy of the health care workforce. As Health AI tools are introduced into safety-net and rural health care settings, frontline staff – clinicians, administrators, and support teams – must be equipped with the skills to understand, question, and responsibly use these technologies. This includes training on how AI systems work, how to verify outputs, and how to identify potential risks.

> *"There needs to be some kind of training that happens on how to utilize [AI], how to think about some of these safety risks or issues that we've talked about. And so I think that's where the training happens, at the institution. I mean, like the clinical setting level [or] the hospital." – Senior Leader for Regulatory Compliance*

Interviewees emphasized that such training must go beyond technical instruction. It should foster a culture of transparency, shared learning, and accountability – empowering staff to engage with AI systems thoughtfully, rather than passively accepting outputs.

> *"...putting in policies and procedures and accountable governance in place...and not have it punitive, but educational. Transparency is everything. So I think, what is that additional level of transparency and accountability? ...whoever is implementing an AI tool is accountable for what the outcomes are and how it's used." – Senior Operations Leader*

While institutional governance remains important, it must be paired with hands-on, locally embedded training that prepares workers to navigate the ethical and operational dimensions of AI use. Investing in AI literacy is essential for building trust, ensuring safety, and enabling equitable adoption of Health AI tools in under-resourced community settings.

Creating Opportunities for Human-Centered Care

In both our survey and interview findings, health care practitioners were clear about the one condition that Health AI must and should satisfy before widespread trust is possible: Applications

of AI in health care must enable workers to deliver more human-centered forms of care. We define AI-enabled, human-centered care as: *Systems that augment the abilities of human health care practitioners to deliver care that is always deeply empathetic and responsive to the unique needs, clinical and social, of a patient or client.*

> *"Our staff mentioned a client shaking their leg [or] fidgeting with their hands. In a space like this, you could read my face a little if I rolled my eyes, or the smack of my lips and so forth. But what kind of AI can read that language? In addition to the fact that clinicians are highly trained in voice inflection, and I would think that AI [or] the people who are creating [AI would] understand that voice inflection is really important, language is really important." – Senior Leader in Health Care*

> *"…will [AI] be able to be sensitive to the emotional reality of the individual? Will you be asking a question without [knowing that a patient] is in a vulnerable moment? So, I think that those emotional vulnerabilities may be lost in a world where the questions are standardized, but there's no concrete data. [Can AI] measure emotional state? Maybe AI can ask a couple of the PHQ-2 questions and see if, indeed, there is an emotional component to this patient and modify [its actions]?" – Senior Leader in Health Care*

Quotes like these underscore the fact that, no matter how sophisticated machine intelligence becomes, interviewees are not ready to trust AI to replace human perception, responsiveness, and emotional intelligence. Indeed, these are elements that make human intelligence unique from machine intelligence.

Improving Efficiency, Effectiveness, and Patient Education

Increasing the efficiency and effectiveness of care is critical to practitioners. There is widespread consensus among our study participants that the future we create with Health AI should be one that augments rather than automates some of the more delicate aspects of health care. *While most practitioners do not see a future in which AI completely replaces trained human experts, they can imagine a future in which AI systems will deepen human expertise* and, as a result, significantly improve the quality of care. In fact, demonstrating that AI can improve the quality of care by augmenting the diagnostic process was viewed as strong motivator of trust.

> *"When a client arrives at an appointment and when they're sitting in front of a clinician, that clinician [can] read a prepared summary of everything that that AI has captured in [a set of] assessments. And they can be ready with a prepared preliminary list of possible diagnoses they might be looking for. And a preliminary list of possible problem areas that they might need to focus on in questioning that client. I think that that would allow our clinicians at intake to have a much more focused discussion with the client, rather than having to go through two hours of question-and-answer assessments. That is tedious for clinicians and for clients. And [after] that two-and-a-half-hour assessment process, clients walk out of there going, 'Okay, but I didn't get any treatment.'" – Senior Health Care Leader*

The next quote underscores how Health AI can improve care by improving communication between a health care practitioner and their patient throughout the entire cycle of care.

> *"Keeping artificial intelligence will enable us to communicate much better with the people we serve if used appropriately at the right place. [AI] can enable a bond between providers and patients. So, if programmed correctly and put in for the right purpose, it can absolutely help an organization — from pre-procedure to when they're in the bed to when they go home to when they pay. I think it will help patient compliance." – Senior Operations Leader*

There is growing consensus that one of the most effective ways to improve health care is to empower patients to be engaged in their own health and well-being. Among other things, this includes enhancing their health literacy, which can help improve their understanding, for example, of a chronic condition and how best to manage that condition. As a result, health care practitioners understand the value of psychoeducation, that is, the need to help enhance patient knowledge about the impact of medication, environment, and behavior in their health outcomes. To the degree that Health AI can augment health care practitioners' abilities to enhance health literacy, Health AI will drive greater levels of trust among them.

> *"But there's also a lot of educational elements to what we do with clients. You know, basic kinds of education about certain life skills, or health and aspects of health and well-being that they can learn. I'm hopeful that AI can fill that void, the void being the discrepancy between the providers that we need and the number that we actually have in order to be able to provide this service to people in our communities. I'm hopeful that the AI can help fill some of those spaces to make our work more efficient." – Senior Health Care Leader*

Extending Care: Remote Patient Monitoring

Participants saw AI with the potential to extend care access through remote and personal device engagement too. When treating people, communities, and sites for health care that lack resources, AI was seen as one way to improve care.

> *"So, if we're consciously trying to close the gap in access to health care, [AI will] help us accelerate that closure. If we don't try to do it, [AI] will accelerate …the health disparities that we see. …I want everybody to know that …health disparities and what we see today is not just about race, it's not just about income levels. I live in a rural area and it's about real, true access to health care. And it's like every single one of us has a health care access issue today, and we can use AI to [address] those access issues. Or will [Health AI] become something that only a few people have access to? And everybody else gets left behind." – Senior Leader for Data Strategy and Innovation*

Health AI was viewed as an important opportunity to improve the delivery of telehealth and strengthen remote patient monitoring. This aspect of Health AI — the ability to deliver care to

and engage patients who live in health service deserts — has the potential to significantly enhance the value of Health AI in safety-net care settings and, consequently, improve trust in the technology. However, as noted in the next quote, remote patient monitoring platforms must be designed in ways that sustain rather than diminish patient engagement.

> *"We're moving to an app to see if an app would be more accessible for the population. We had this dashboard where people had to log in and log out and click on links. It was very cumbersome, so it just didn't work for the community. We also tried to do self-monitoring, [integration] with Fitbits and dietary self-monitoring. It was too much, did not work."*
> *– Senior Leader for Health and Equity Initiatives*

Navigating Cultural Diversity in Care

One of the advantages of centering this strategy paper around the views, experiences, and expertise of safety-net health care professionals is the opportunity to learn more about the unique dimensions of lived experience in these settings. Texas has a strikingly diverse population, which poses significant challenges to the delivery of health care. Population diversity introduces cultural differences and linguistic diversity. As a result, engaging patients receiving safety-net services with Health AI demands fitting AI systems and interface to users of different languages, cultures, technology proficiency levels, and/or physical capabilities. Designing AI to be sensitive to cultural and linguistic diversity was a common refrain in our interviews and is eloquently expressed in the next quote related to patient literacy in Health AI tools.

> *"Do they [patients] know what they're actually asking [AI] for? And then can they reliably interpret or accept what AI gives back? If AI says something stupid, is that patient going to know that's not a good [output]? Or is the patient going to be able to express [themselves] in a way to AI that it can generate something reasonable? And can [AI] do it in ways that are linguistically, culturally sensitive? Can it understand what a native English speaker is trying to communicate versus someone who's a non-native English speaker? …I'm not convinced at this point that it's there yet, but you know, hopefully it will be."*
> *-Senior Executive and Medical Leader*

Integrated Analysis and Discussion

Insights from our in-depth interviews support many of the literature's findings while emphasizing practical applications of Health AI. The literature highlighted the significant value of AI in diagnostic capabilities, telehealth, and the need for adequate data representation.[5] It also recognized opportunities for automating administrative tasks and stressed the importance of

[5] Iacobelli, F., Yang, A., Tom, L., Leung, I. S., Crissman, J., Salgado, R., & Simon, M. (2023). Predicting social determinants of health in patient navigation: Case study. *JMIR Formative Research, 7*, e42683. https://doi.org/10.2196/42683

regulatory oversight and keeping up with emerging AI trends.[6] However, there was less focus in the literature on the necessity of maintaining human oversight in AI processes and less focus on the potential of AI to improve health information accessibility and patient education.[7]

A comparative analysis between the literature[8] and the primary perspectives of health care professionals revealed how both acknowledge AI's potential to enhance decision making and medical treatment, provide personalized care tailored to individual patients, facilitate telehealth and remote monitoring, and reduce biases in health care data and algorithms. *While the literature tends to concentrate on data-driven challenges and visionary, long-term solutions that have yet to be fully realized, the health care professionals and thought leaders we interviewed focus on immediate, practical applications that integrate AI to support healthy patient outcomes.*

For example, long-term visions for Health AI often include the development of "digital twin" technology. This concept involves creating virtual replicas of patients using integrated data from genomics, imaging, environmental factors, and electronic health records to simulate and predict outcomes of treatments in a digital environment before implementing them in real life.[9] Such systems, while not yet widely realized, hold the potential to revolutionize personalized medicine by enabling clinicians to optimize treatment plans, reduce risks, and anticipate complications with unprecedented precision.

The practitioners and researchers in our in-depth interviews and survey emphasize how AI can help reduce administrative burdens, allowing more time for patient care, enhance patient communication, and aggregate data to provide a comprehensive view of patient conditions. They also stress the importance of ensuring AI systems are adaptable to the language, cultural, and physical needs of patients. Bridging the researcher and practitioner perspectives requires a strategy that balances the potential of advanced AI capabilities with the need for immediate and clinically relevant solutions that support health care providers and foster trust among patients.

---

[6] Solomonides, A. E., Koski, E., Atabaki, S. M., Weinberg, S., McGreevey, J. D., Kannry, J. L., Petersen, C., & Lehmann, C. U. (2021). Defining AMIA's artificial intelligence principles. *Journal of the American Medical Informatics Association, 29*(4), 585–591. https://doi.org/10.1093/jamia/ocac006

[7] Dave, T., Athaluri, S. A., & Singh, S. (2023). ChatGPT in medicine: An overview of its applications, advantages, limitations, future prospects, and ethical considerations. *Frontiers in Artificial Intelligence, 6*, 1169595. https://doi.org/10.3389/frai.2023.1169595

[8] Our in-depth literature review is included as a separate document, a companion to the strategy paper.

[9] Bajwa, J., Munir, U., Nori, A., & Williams, B. (2021). Artificial intelligence in healthcare: Transforming the practice of medicine. *Future Healthcare Journal, 8*(2), e188–e194. https://doi.org/10.7861/fhj.2021-0095

PART FIVE: STRATEGIC RECOMMENDATIONS

for Successful AI Integration in Safety-Net and Rural Health Care Settings

Considering our research findings, we offer these strategy recommendations on how stakeholders and leaders in health policy, philanthropy/funding, and services can actively improve care for communities in rural areas and others in need of safety-net health care services through a holistic and strategic approach to Health AI. [10]

1) Educate: Build AI Literacy and Workforce Capacity

Collaborate with local organizations to host an AI health care governance conference and develop statewide educational programs for health care professionals, focusing on responsible AI use.

Rationale	*Actions*
Ensuring health care workers, especially in underserved areas, have the knowledge to integrate AI responsibly will build trust with vulnerable populations.	○ Host an AI governance conference with stakeholders across Texas, focusing on AI policy, ethics, and local health care needs. ○ Take learnings from the conference on the road to key institutions across Texas, holding practitioner and patient focus groups to receive feedback to mold a "governance handbook" for institutional reference. ○ Develop educational programs combining remote and hands-on learning for health care providers to foster trust-building and AI applications in patient care.

[10] *A Note:* It is well-established that Health AI has benefits in automating administrative tasks, thus saving time for clinicians and other health care providers so that they can focus on other aspects of patient engagement and care. We support the introduction of Health AI tools into safety-net and rural service settings to increase operational efficiency. Our findings and recommendations go further than administrative efficiency, however, and point out strategic steps to ethically integrate Health AI technologies into the full range of health care settings to benefit providers and patients both, improving health outcomes across the board.

2) Participate: Engage Rural and Safety-Net Services Populations

Actively involve rural and safety-net services populations in AI adoption by conducting region-specific research and piloting AI solutions.

Rationale	*Actions*
Involving these populations early in AI development ensures the technology is effective, relevant, and trustworthy for the communities it is meant to serve.	○ Fund research to improve the collection of data from rural and underserved populations, ensuring AI models are representative and equitable. ○ Deploy AI applications such as remote monitoring and mental health services in rural areas, and measure their impact on patient engagement and trust. ○ Conduct follow-up studies with patients, families, and caregivers to refine AI interventions and ensure they meet local needs. Learnings from these interventions would update the user manual and governance handbook referenced in Recommendations 1 and 3.

3) Evaluate: Establish Strong Evaluation Frameworks

Create key performance indicators (KPIs) to evaluate AI tools for consideration, integration, and assessment of their impacts.

Rationale	*Actions*
A standardized evaluation framework will aid institutions in assessing AI technologies, ensuring they align with patient care goals and demonstrate tangible benefits.	○ Pilot AI technologies in one institution to assess performance and refine KPIs for evaluating AI tools in active patient-provider settings. ○ From learnings, develop a “user manual” for AI consideration and piloting, to provide institutions with guidance as they select and assess AI tools. This user manual will complement the governance handbook described in Recommendation 1.

4) Innovate: Develop Locally Driven AI Solutions

Support the development of AI solutions tailored to Texas' unique health care challenges through hackathons and innovation incubators.

Rationale	*Actions*
Texas' diverse health care needs require locally driven AI solutions. Hackathons and incubators will help create impactful tools that address these challenges.	○ Host regional hackathons focused on developing AI solutions for rural and underserved populations. ○ Partner with universities and health care organizations to establish a Health AI incubator that supports the prototyping and commercialization of solutions that address local challenges. ○ Determine the fit of the "governance handbook" and "user manual" to diverse and innovative Health AI solutions and update them to accommodate cutting-edge technologies.

CONCLUSION

By combining educational, participatory, evaluative, and innovative strategies, we can support the introduction of Health AI in safety-net and rural health care settings. AI can be introduced and integrated in ways that improve care delivery and operational efficiency, while also building trust in the technology. Implementing the strategic recommendations described in this paper can ensure that AI is deployed responsibly, collaboratively, ethically, and effectively, with a focus on meeting the needs of underserved populations while fostering a balance between technology and the human touch that is crucial for health care success. Further, our recommendations produce deliverables that will be iteratively updated, which offers tangible guidance for policy and governance as well as for the selection, testing, and evaluation of Health AI in safety-net settings. The future success and impact of Health AI will depend, in part, on how well these innovations are designed and deployed in environments that have been overlooked and underserved, thus ensuring higher quality care for all populations.

APPENDIX

Appendix A: Methodology and Framework of the Current Analysis

This study utilized a mixed methods approach that included an extensive review of literature relevant to the state of AI in health care, a health care practitioner survey, and a set of in-depth interviews with health care providers and researchers. The survey focused on the current awareness, general perceptions of AI, and reactions to specific vignette-based deployment scenarios for Health AI among a range of health care providers in Texas. The interviews probed issues revealed in the literature review and survey more deeply.

Literature Review Framework

The Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework has been significant in our efforts to systematically explore and evaluate AI's potential in the healthcare system.[11] Adhering to the PRISMA guidelines allowed us to synthesize the current landscape more comprehensively and transparently, assessing AI technologies' opportunities and challenges when being implemented within the health care sector.

This approach allowed us to evaluate and interpret existing primary studies that explore the promises and perils of Health AI. Alongside the PRISMA method, the two independent reviewers were also using an eligibility criterion to assist with the selection process of the articles.

Following the PRISMA guidelines, we created an elaborate search strategy to identify the growing literature on the influence of promises and perils of AI in health care. We conducted a literature search in the following two databases: PubMed and Web of Science. The main concepts that we focused on during our searches were: artificial intelligence, health care, equity and marginalized communities. Table A1 outlines the keywords and synonyms used for the literature search, ensuring comprehensive coverage of relevant topics. Table A2 presents the inclusion and exclusion criteria applied to select articles, ensuring relevance and consistency in the review process.

[11] Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., … Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372*. https://doi.org/10.1136/bmj.n71

Table A1. Keywords and Synonyms

Artificial Intelligence	"ML OR AI" OR "Electronic health" OR "big data" OR "data mining" OR "decision support" OR "Data"
Healthcare/Health Care	"Health" OR "treatment" OR "patient experience" OR "provider experience" OR "clinical outcome" OR "rural health" OR "health information technology" OR "health informatics"
Equity	"Ethics OR equity" OR "accessibility" OR "health equity" OR "health disparity" OR "Social Determinants of Health"
Marginalized Communities	"Safety-net OR marginalized communities" OR "rural communities" OR "rural" OR "underserved population"

Table A2. Eligibility Criteria

Inclusion	Exclusion
Mentions AI/DL/ML or some other component of Artificial Intelligence	Does not mention AI/DL/ML or other components of Artificial Intelligence
Highlights health care or a component of health	Did not focus on health care or a component of health
Relates to patient or provider perspective	Does not relate to patient or provider perspective
Written in English	Any papers that focused on clinical trials, meta-analyses, or randomized controlled trials

Survey Methodology

A survey was used to investigate how experiences, perceptions, and roles influence the adoption and understanding of Health AI technologies. Specifically, the research sought to expand knowledge on how health care professionals perceive AI in health care settings, with a particular focus on rural and other populations in need of safety-net services.

The survey was piloted among the research team's health care provider network in August 2024 and fielded through a health care organization research panel, Dynata, from September 5-24, 2024. We received 229 completed surveys from Texas health care providers. The sample represents a cross section of organization types and health care provider roles.

Although the survey was largely quantitative in nature, there were opportunities for respondents to respond with qualitative content to four general open-response questions and four vignette-associated open-response questions, each with ~200 responses. Each of the four vignettes in the survey presented a different example of a Health AI tool (presented in the next section) and offered a single question of open-ended responses. We coded and bundled survey open-ended responses for analysis using the same approach used in interview analysis, described next.

Interview Methodology

A series of in-depth interviews were used to probe deeper into the experiences and perceptions of health care practitioners and researchers. The interview sample included 18 independent interviews, 14 with health practitioners and 4 with research leaders. Participants were recruited using the health care network of the research team and referrals beyond that network to other health care providers and researchers.

Interview transcripts were coded using 66 individual codes reflecting expressed sentiments. Codes addressing related sentiments were combined into common themes reflecting promise or peril in using Health AI. As noted above, the eight open-response questions in the survey were analyzed using the same codes as interviews and common themes were bundled.

Vignettes Presented in Survey

The following is verbatim content from the survey related to four realistic but fictional vignettes:

In the next section of our survey, we will present several scenarios for how AI might be deployed in healthcare settings. Please read each description and think about how that technology fits the patients that you and your practice serve. Think about the appeal of the technology as described and the level of concern, if any, that you perceive. We recognize that these technologies might be more or less appropriate for different types of patients. For this exercise, we ask that you think about all of the patients you serve.

V1. DocuScribeAI: *You are a healthcare professional working in a high-demand clinic serving a diverse population, including many underinsured patients. Managing multiple patient visits is complex, and clinical documentation consumes a significant amount of your time. Recently, you were introduced to* ***DocuScribeAI****, an AI tool designed to assist in transcribing and organizing clinical notes during patient interactions. DocuScribeAI listens to conversations, transcribes them in real-time, and categorizes key information using natural language processing. It integrates seamlessly with the clinic's electronic health records system.*

V2. TherapiaAI: *You are a healthcare professional in a clinic that serves a diverse and often underinsured population. Managing varied patient needs is challenging, and you discover* ***TherapiaAI****, an AI tool designed to enhance treatment personalization. As you interact with patients, TherapiaAI analyzes real-time data from patient histories, genetic information, and current health conditions. It suggests customized therapeutic plans based on clinical guidelines and patient-specific factors. This tool integrates with your clinic's electronic health records system, allowing for immediate updates and informed decisions.*

V3. HealthRiskAI: *You are a healthcare provider at a community clinic that serves a diverse safety-net population. Your clinic frequently faces the challenge of allocating limited resources to patients with the highest needs. Recently, you have become aware of a new tool called* ***HealthRiskAI****, which integrates patient-specific social determinants of health (SDoH) and provides real-time risk assessments. HealthRiskAI has the potential to enhance resource allocation by identifying patients who require immediate attention and those who could benefit from proactive outreach. For example, it can flag*

patients at risk of developing complications due to non-adherence to their treatment plan for follow-up. The tool integrates with the clinic's existing electronic health records, enabling immediate updates and informed decision making. By using HealthRiskAI, clinics can better prioritize patient care, ensuring timely and appropriate interventions for those with the highest needs, and ultimately improving health outcomes.

V4. CommCare: *You are a healthcare provider at a high-volume mental health clinic, managing a demanding schedule with diverse patient needs. Despite efforts to connect with each patient, you've observed frequent appointment no-shows and disengagement during sessions. Post-visit feedback often indicates patient dissatisfaction. You learn about* ***CommCare****, an AI tool designed to evaluate patient communication patterns to enhance engagement and satisfaction. CommCare analyzes communication preferences and engagement patterns, offering insights to improve patient-provider interactions.*

Appendix B: Glossary

AI (Artificial Intelligence)	The simulation of human intelligence processes by machines, particularly computer systems. In health care, AI is used to analyze complex medical data, predict patient outcomes, support clinical decision making, and automate administrative tasks.
AI-Driven Diagnosis	The use of AI algorithms to analyze medical data (e.g., imaging, lab results, patient history) to aid in diagnosing diseases. AI-driven diagnosis is increasingly being used in fields like radiology, pathology, and dermatology.
AI Governance	The framework of policies, principles, and processes that oversees the ethical development and use of AI in health care. AI governance aims to ensure that AI tools are implemented in ways that promote fairness, transparency, accountability, and patient safety.
AI Model Training	The process of feeding data into AI algorithms to help them learn and make predictions. In health care, model training typically involves large datasets – including patient records, medical imaging, and clinical trials – to develop AI tools that assist in diagnosis or treatment.
Automated Workflow	The use of AI and automation tools to streamline administrative and operational tasks in health care settings, such as scheduling, billing, and patient intake, reducing administrative burdens and improving efficiency.
Bias in AI	The presence of systematic errors in AI algorithms that lead to unfair, prejudiced, or discriminatory outcomes. In health care, bias can arise from biased training data or flawed model development, potentially leading to health disparities.
Clinical Decision Support Systems (CDSS)	AI-powered tools designed to assist health care providers in making clinical decisions. These tools/systems provide evidence-based recommendations, alerts, and guidelines to improve patient outcomes and reduce errors.
Clinical Trial Optimization	AI applications used to streamline the design, recruitment, and management of clinical trials. AI can improve patient recruitment, data analysis, and trial outcomes, ultimately speeding up the development of new treatments.
Clinical Validation	The process of ensuring that an AI tool or model performs reliably and accurately in real-world clinical settings. Clinical validation is crucial to ensuring that AI applications in health care are safe, effective, and beneficial for patients.
Data Governance	The policies, procedures, and standards that ensure the quality, security, and privacy of data. In Health AI, data governance ensures that data used to train AI models are accurate, representative, and ethically sourced.

Data Integrity	The accuracy, consistency, and reliability of data used in AI systems. In health care, ensuring data integrity is vital for creating accurate models and avoiding the risks of incorrect or incomplete information influencing clinical decisions.
Deep Learning	A type of machine learning that uses neural networks with many layers (hence "deep") to model complex patterns in large datasets. It's commonly used in image recognition, such as analyzing medical imaging data for disease detection.
Electronic Health Records (EHRs)	Digital versions of patients' paper charts, containing medical history, diagnoses, treatments, medications, test results, and other critical health data. AI is often used to analyze and mine EHRs for insights that can improve patient care.
Ethical AI	The development and use of AI in ways that are consistent with ethical principles, such as fairness, transparency, accountability, and respect for privacy. Ethical AI in health care ensures that technologies do not exacerbate existing health inequalities and are aligned with patient rights and societal values.
Explainability (or Interpretability)	The degree to which AI models can be understood and explained by humans. In health care, explainable AI is critical for clinicians to trust and effectively use AI tools in practice, as they need to understand the rationale behind AI recommendations.
Federated Learning	A machine learning technique that allows AI models to be trained on decentralized data (e.g., patient data from various hospitals) without the data ever leaving the local site. This helps to preserve privacy while enabling the sharing of insights across institutions.
Health AI Adoption	The process by which health care organizations, providers, and patients integrate AI technologies into daily practice. This includes the readiness of health care institutions, training of staff, and fostering of trust among patients and providers in AI-driven tools.
Health AI Ecosystem	The interconnected set of technologies, organizations, regulations, and stakeholders involved in the development, implementation, and use of AI in health care. This includes AI developers, health care providers, policymakers, and patients.
Health Equity	The principle that all individuals should have fair access to health care resources and opportunities for good health, regardless of their socioeconomic status, race, or geographic location. Health AI must be designed and implemented with a focus on reducing disparities in care.
Health Informatics	The field that involves the use of technology and data to improve the quality and safety of health care. It includes the collection, analysis, and management of health data to support clinical decision making, policy development, and patient care.

Interoperability	The ability of different health care systems and technologies to work together and exchange information seamlessly. AI can help enhance interoperability by enabling more efficient data sharing and integration across various health care platforms.
Machine Learning (ML)	A subset of AI that involves training algorithms to identify patterns in data and make decisions without being explicitly programmed. In health care, ML is often used for diagnostics, predictive analytics, and personalized treatment plans.
Natural Language Processing (NLP)	A branch of AI that focuses on the interaction between computers and human language. NLP is used in health care to analyze and interpret unstructured text from medical records, patient notes, and clinical literature.
Patient-Centered Care	An approach to health care that respects and responds to individual patient preferences, needs, and values. AI in health care should be designed to support, rather than replace, patient-centered care, ensuring that patients are actively involved in decision making.
Personalized Medicine	Medical treatment tailored to individual patients based on their genetic, environmental, and lifestyle factors. AI can help analyze vast amounts of data to create personalized treatment plans, improving the precision of care.
Predictive Analytics	The use of AI and statistical techniques to analyze historical data and predict future outcomes. In health care, predictive analytics can help identify patients at risk for diseases, predict hospital readmissions, or anticipate disease progression.
Privacy and Security	In Health AI, privacy refers to the protection of patient data from unauthorized access, while security involves safeguarding systems from cyber threats. Ensuring the privacy and security of health data is critical to building trust in AI applications.
Regulatory Compliance	The adherence to laws, regulations, and standards governing health care and data privacy. In the context of Health AI, this includes complying with regulations like HIPAA (Health Insurance Portability and Accountability Act) in the U.S., which governs patient privacy and the use of health data.
Telemedicine	The remote delivery of health care services using technology, such as video consultations or remote monitoring. AI is often integrated into telemedicine platforms to provide clinical decision support, monitor patients, or analyze health data.