



AI in Action: Turning Health Data into Decisions

A Not-for-Attribution
Roundtable Summary
(Chatham House Rules)

IN PARTNERSHIP WITH



Executive Summary

Federal healthcare organizations are entering a decisive phase in their adoption of artificial intelligence. What began as exploration and isolated pilots is rapidly evolving into early operational deployment across agencies responsible for veterans' care, military readiness, and public health. The conversation is no longer about whether AI has value. That question has been answered. The real challenge is how to implement it responsibly, effectively, and at scale.

The roundtable revealed a shared recognition that artificial intelligence is not constrained by technical capability. The tools are accessible, increasingly intuitive, and capable of delivering immediate impact. Instead, the constraints lie within data, workforce readiness, trust, and institutional alignment. Agencies are grappling with how to integrate AI into environments that are complex, regulated, and mission-critical.

At the same time, there is a growing sense of urgency. The potential benefits are too significant to ignore. AI offers a path to reduce administrative burden, improve decision-making, and unlock insights from vast amounts of data that have historically been underutilized. But realizing that potential requires a disciplined approach—one that prioritizes data quality, builds trust through transparency, and equips the workforce to engage with these tools effectively.

This paper reflects a federal healthcare ecosystem in transition. The experimentation phase is ending. The execution phase has begun.

The Shift from Curiosity to Capability

For several years, artificial intelligence in government has been characterized by curiosity. Agencies explored use cases, conducted pilots, and tested the boundaries of what the technology could do. That phase is now giving way to something more consequential. Leaders are beginning to see AI not as a novelty, but as an operational capability that can be embedded into everyday workflows.

This shift is being driven by a combination of factors. The technology itself has matured significantly, particularly in the area of generative AI. Tools that once required specialized expertise can now be used with minimal training. At the same time, external pressures—ranging from workforce shortages to increasing mission complexity—are forcing agencies to find new ways to operate more efficiently.

What is emerging is a pragmatic approach to adoption. Rather than pursuing ambitious, high-risk applications, agencies are focusing on areas where AI can deliver immediate and measurable value. This often means starting with internal processes and administrative functions, where the stakes are lower and the benefits are easier to demonstrate.

This transition marks an important inflection point. AI is no longer something that agencies are experimenting with on the margins. It is becoming part of how they operate.



Data as the Defining Constraint

Despite rapid advances in AI technology, participants consistently emphasized that the primary constraint is not the models themselves, but the data that feeds them. Federal healthcare systems generate enormous volumes of data, yet much of it is fragmented, inconsistent, or difficult to access.

This reality creates a paradox. On one hand, agencies possess a wealth of information that could be used to improve outcomes and drive efficiency. On the other hand, the lack of cohesion within that data limits its usefulness. Artificial intelligence, rather than solving this problem outright, often exposes it more clearly.

Participants described how AI systems can surface inconsistencies in policy documents, identify duplicative records, and highlight gaps in data coverage. While this can be valuable, it also underscores the need for a more deliberate approach to data management. Without a strong foundation, even the most advanced AI tools will produce results that are difficult to trust.

The concept of a “single source of truth” emerged as a central objective. Achieving this requires not only technical solutions, but also organizational alignment. Data must be standardized, governed, and maintained in a way that supports consistent use across the enterprise.

In this context, the focus of AI adoption is shifting. The challenge is no longer building models. It is building the data infrastructure that allows those models to function effectively.

Building Trust Through Transparency

Trust is a critical factor in the adoption of artificial intelligence, particularly in healthcare environments where decisions can have significant consequences. Participants made clear that trust cannot be assumed. It must be earned.

One of the most effective ways to build trust is through transparency. Users need to understand how AI systems arrive at their conclusions. They need visibility into the data sources, the processing steps, and the logic behind the outputs. Without this, even accurate results may be viewed with skepticism.

This requirement has important implications for system design. Transparency must be built into AI systems from the beginning. It cannot be added as an afterthought. This includes not only technical explainability, but also clear communication about the limitations of the technology.

The need for transparency is particularly acute in high-stakes environments. In clinical settings, for example, practitioners must be able to validate AI-generated insights against their own expertise. In policy contexts, decision-makers must be confident that recommendations are grounded in reliable data.

As AI becomes more integrated into federal healthcare operations, trust will become a key differentiator. Systems that are transparent and understandable will be adopted more readily than those that are not.

A Risk-Based Approach to Adoption

Given the complexity and sensitivity of healthcare environments, participants emphasized the importance of a risk-based approach to AI adoption. Not all use cases carry the same level of risk, and agencies must be deliberate in how they prioritize their efforts.

Low-risk use cases provide an opportunity to build experience and demonstrate value without exposing the organization to significant downside. These include tasks such as drafting documents, summarizing information, and supporting internal workflows. In these scenarios, AI serves as an assistant rather than a decision-maker, and human oversight remains central.

High-risk use cases, such as clinical diagnosis or treatment recommendations, require a much higher level of confidence. These applications involve complex ethical, legal, and operational considerations that cannot be addressed through technology alone. As a result, most agencies are approaching these areas with caution.

This phased approach allows organizations to build momentum while managing risk. It also creates a foundation for more advanced applications in the future. By starting with low-risk use cases, agencies can develop the capabilities and confidence needed to tackle more complex challenges over time.



Reducing Administrative Burden as a Strategic Entry Point

One of the clearest areas of opportunity identified in the discussion is the reduction of administrative burden. Across healthcare systems, administrative tasks consume a significant portion of time and resources. These tasks are often repetitive, time-consuming, and not directly tied to mission outcomes.

Artificial intelligence offers a way to streamline these processes. By automating documentation, summarizing information, and assisting with reporting, AI can significantly reduce the workload associated with administrative functions. This has immediate benefits. It frees up time for clinicians, improves efficiency, and allows organizations to operate more effectively.

In many cases, the impact is dramatic. Tasks that once required days can be completed in a matter of hours. This shift not

only improves productivity, but also enhances job satisfaction by allowing staff to focus on more meaningful work.

Importantly, these use cases are relatively low risk. They do not require AI to make critical decisions, and they can be easily validated by human users. This makes them an ideal starting point for broader adoption.

The Challenge of Fragmentation and Interoperability

While the potential of AI is significant, it is constrained by the fragmented nature of healthcare data. Information is often stored across multiple systems, each with its own standards and formats. This fragmentation limits the ability of AI to provide comprehensive insights.

Participants described a landscape in which data is abundant but disconnected. Patients may have records spread across multiple

providers, systems, and locations. Providers, in turn, must navigate a complex web of data sources to access the information they need.

Efforts to improve interoperability are underway, but progress has been uneven. Existing frameworks have made it easier for providers to share information, but they have not fully addressed the needs of patients or the challenges of integrating new data sources.

This fragmentation represents a structural barrier to AI adoption. Until it is addressed, the full potential of AI will remain out of reach. Agencies must continue to invest in interoperability and data integration as part of their broader AI strategy.

Workforce Readiness and the Human Element

Technology alone is not sufficient to drive adoption. The effectiveness of AI depends on the people who use it. Participants highlighted workforce readiness as a critical factor, noting that even the best tools can be undermined by a lack of understanding or confidence.

There is a growing recognition that AI requires a new set of skills. Users must be able to interact with these systems effectively, interpret their outputs, and apply them in context. This is not a trivial challenge, particularly in large organizations with diverse workforces.

At the same time, there is a cultural dimension to consider. Many employees are uncertain about AI, unsure of how it will affect their roles or whether it can be trusted. Addressing these concerns is essential to building adoption.

The most effective approaches to workforce development are those that emphasize practical experience. Rather than relying solely on formal training, organizations are finding success with hands-on learning, mentorship, and real-time demonstrations. These methods allow users to build confidence and develop skills in a way that is directly relevant to their work.

Overcoming Cultural Resistance

Cultural resistance is a natural response to change, particularly when that change involves new and unfamiliar technologies. In the context of AI, this resistance often manifests as skepticism, fear, or hesitation.

Participants noted that one of the most effective ways to address this resistance is through exposure. When users see how AI can be applied to their work, and when they experience its benefits firsthand, their perceptions begin to shift. What was once viewed as a threat becomes a tool.

Communication also plays a critical role. Leaders must be clear about what AI can and cannot do, and they must set realistic expectations. Overpromising can lead to disappointment and reinforce skepticism, while honest and transparent communication builds credibility.

Ultimately, cultural change is a gradual process. It cannot be forced, but it can be guided. By creating an environment that encourages experimentation and supports learning, organizations can foster a culture that is more open to innovation.



Public Health Applications and Opportunities

In the public health domain, AI presents a unique set of opportunities and challenges. Public health systems rely on large-scale data to monitor trends, identify risks, and inform policy decisions. AI has the potential to enhance these capabilities by automating data analysis and improving the accuracy of insights.

Participants discussed how AI can be used to streamline surveillance, reduce duplication in reporting, and support more timely decision-making. These applications are particularly important in situations where rapid response is critical.

However, public health systems also face significant challenges. Data is often collected from a wide range of sources, each with its own limitations. Ensuring consistency and reliability is a complex task. There are also concerns about bias and equity, which must be carefully managed.

As a result, AI adoption in public health is likely to proceed with both urgency and caution. The potential benefits are substantial, but they must be balanced against the need for accuracy and fairness.

Toward an AI-Enabled Federal Healthcare Enterprise

Looking ahead, participants described a future in which AI is fully integrated into federal healthcare operations. In this vision, AI is not a separate tool, but a core component of how work gets done. It supports decision-making, streamlines workflows, and enables more effective use of data.

Achieving this vision will require sustained effort. Agencies must continue to invest in data infrastructure, develop their workforce, and refine their governance frameworks. They must also remain adaptable, as the technology continues to evolve.

The path forward is not without challenges, but the direction is clear. Artificial intelligence is becoming an essential part of federal healthcare, and its role will continue to grow.

Conclusion

Federal healthcare agencies are at a pivotal moment in their adoption of artificial intelligence. The technology is ready, and the

opportunities are significant. The challenge now is to move from experimentation to execution in a way that is thoughtful, disciplined, and aligned with mission needs.

Success will depend on more than technology. It will require strong data foundations, a capable and confident workforce, and a commitment to transparency and trust. Agencies that can bring these elements together will be well positioned to realize the full potential of AI.

The transition is already underway. The next phase will determine how far and how fast it progresses.

