




WHEN AI SETS THE TERMS OF CARE: MAKING HIDDEN CHOICES VISIBLE IN HEALTH SYSTEMS

Hospitals rely on AI to guide care, but hidden configuration choices shape how these tools operate and whom they prioritize. New research shows AI embeds trade-offs into care delivery and calls for greater transparency and oversight to align these decisions with clinical judgment and organizational values.

FEATURING

CIL Affiliate Shefali Patil, and CIL Core Faculty
Christopher Myers and Tinglong Dai





Summary

Hospitals are rapidly adopting AI to support decisions about patient care, from predicting deterioration to prioritizing treatment. But this research reveals a critical leadership blind spot: before these tools ever reach clinicians, hospitals are making decisions about how they are configured—and those choices quietly shape the care every patient receives. These configurations embed trade-offs between competing priorities such as caution and cost, speed and clinical judgment. The key insight is a shift in perspective. AI implementation is not simply a technical rollout. It is a moment where organizational values are set, often out of view, and then scaled across the system.

The Leadership Challenge

Across health care, AI is becoming embedded in everyday clinical decisions. Systems are used to predict which patients are getting sicker, determine who should be seen first in the emergency department, and identify warning signs in medical imaging. These tools promise better outcomes and more efficient operations.

But long before clinicians encounter these systems, critical decisions have already been made. During procurement and implementation, organizations determine how the system will operate. Should it cast a wide net, identifying as many potential cases as possible while generating frequent false alarms? Or should it be more selective, reducing noise but risking missed patients? Should it prioritize speed and throughput, or preserve space for clinicians to pause, question, and intervene?

These are not minor technical settings. They are decisions about what kind of care the organization is prepared to deliver. Yet they are typically made by IT teams, vendors, or administrators during the purchasing process, often without meaningful input from the clinicians who will rely on these systems in real time.

Once these choices are locked into the software, they do not remain localized. They are applied to every patient, every day. And because many AI systems continue to update over time, the rules governing care can shift without clinicians at the bedside ever being aware. What appears to be a neutral tool is, in practice, Over time, this can lead to alert fatigue, where caregivers begin to tune out warnings altogether. A system designed to improve safety can unintentionally make it harder to act on critical signals.

Getting to the Source of the Problem

CIL Affiliate and Associate Professor at the Texas McCombs School of Business, Shefali V. Patil, and her co-authors—Professors Christopher Myers and Tinglong Dai of the Johns Hopkins Carey Business School—examine a critical but often overlooked question: how do the hidden configuration decisions of AI systems influence clinical care, and who is responsible for making them?

The prevailing assumption is that AI tools function as objective supports for clinical judgment. Leaders focus on adoption, integration, and performance metrics. But this research challenges that assumption by showing that the most consequential decisions occur before the tool is ever used—at the point of selection and setup.

At that stage, organizations are not simply choosing a tool. They are deciding how that tool will behave in practice. Configuration decisions determine how sensitive an alert system is, how aggressively patients are flagged, and how much autonomy the system has to act. Each of these reflects a trade-off between competing priorities. A system can be tuned to maximize early detection, minimize false alarms, reduce cost, or increase efficiency. It cannot do all of these equally well.

The issue is not that these trade-offs exist. It is that they are rarely surfaced as explicit leadership decisions. Instead, they are embedded in vendor defaults and implementation settings, creating a gap between how systems are designed and how care is experienced on the front lines.

New Findings with Implications to Leadership and Organizational Performance

The research makes clear that AI systems operationalize trade-offs that directly shape patient outcomes and clinician experience—and once set, those trade-offs are applied consistently across the organization.

When a hospital increases the sensitivity of an AI alert system, it catches more patients who are truly at risk. At the same time, clinicians are flooded with false alarms. Over time, this can lead to alert fatigue, where caregivers begin to tune out warnings altogether. A system designed to improve safety can unintentionally make it harder to act on critical signals.



In another scenario, a hospital under financial pressure may configure its AI to flag fewer cases. This reduces unnecessary interventions and controls costs. But it also means some patients—often the most vulnerable—are not identified early enough. A decision made during implementation in the name of efficiency quietly reshapes who receives timely care.

In systems designed for speed, AI may take on more autonomous roles, routing patients or prioritizing cases. Throughput improves, but clinicians have less opportunity to intervene when something does not align with their judgment. Efficiency increases, but the space for professional discretion narrows.

These examples illustrate a central insight. What appears to be a technical configuration becomes a consistent pattern of care.

The research also highlights a significant consequence for clinicians. When outcomes are uncertain, responsibility becomes ambiguous. If clinicians follow the AI and something goes wrong, they may be blamed for relying too heavily on it. If they override the AI and something goes wrong, they may be blamed for not trusting it. The authors describe this dynamic as “accountability ping-pong.” Over time, it erodes confidence in clinical judgment and places clinicians in an untenable position within a system whose rules they did not help define.


To address these challenges, the authors propose a coordinated governance framework that connects external transparency with internal decision-making. First, AI developers should provide standardized, plain-language summaries—Model Cards—that clearly explain how the system works and what trade-offs are built into it. These disclosures make design choices visible.

Second, hospitals should establish multidisciplinary review processes that include clinicians, nurses, administrators, and ethicists. These groups would use that transparency to deliberate the trade offs before deployment, determine which priorities align with the organization’s values, and communicate those decisions clearly to staff.

Together, these elements form a system of governance. External transparency enables internal deliberation, and internal deliberation ensures that the values embedded in AI systems are chosen intentionally rather than inherited implicitly.

Interpretation: What This Means for Leaders

This research reframes AI implementation as a leadership responsibility centered on how systems are configured. The most important leadership moment is not



when clinicians use the tool. It is when the organization decides how the tool will operate.

For leaders, this means shifting attention to where these decisions are made. Procurement and implementation are not administrative steps. They are the moments where organizational values are translated into operational rules that will guide care at scale.

This also requires expanding who is involved in those decisions. When clinicians are excluded, systems are designed without the insight of those who must navigate their consequences. Bringing multidisciplinary voices into the process is not simply inclusive practice. It is essential to aligning system behavior with clinical reality.

The proposed governance framework clarifies accountability by making trade-offs visible and shared. When decisions are explicit and documented, organizations can stand behind them, and clinicians are no longer left navigating unseen constraints.

Transparency, in this sense, is not just about information. It is about restoring the connection between judgment and responsibility in an increasingly automated environment.

Conclusion

AI is often introduced as a way to improve decision-making in complex environments. This research shows that it also relocates where decisions are made—into system design and configuration.

The leadership task is to engage that shift directly. When organizations make these choices visible and subject to collective deliberation, they retain the ability to shape how care is delivered and how clinical judgment is exercised.

The question is not whether AI will influence patient care. It already does. The question is whether those influences will remain hidden in the background, or whether leaders will bring them into the open and lead them with intention.



Access the full research paper here: Patil, S. V., Myers, C. G., & Dai, T. (2026). Protecting clinical value judgment in the age of AI. *Npj Digital Medicine*, 9(1).

<https://doi.org/10.1038/s41746-026-02561-1>