Insight Series

Building the bridge—Generative AI and the future of clinical knowledge

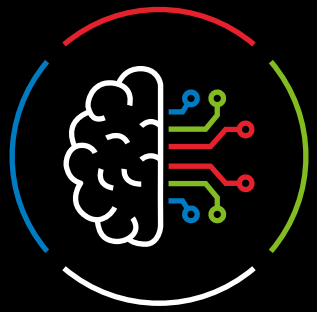

Special Edition: UpToDate Point of Care Report

Shelia Bond, MD, Director, Clinical Content Strategy, Wolters Kluwer Health

Wolters Kluwer

GenAI:
A new epistemic era

Medicine has always evolved alongside the ways we transmit and validate knowledge. From oral traditions passed between healers, to written texts that codified clinical knowledge, to digital systems that indexed and retrieved vast repositories of evidence—each shift has reshaped how clinicians learn, decide, and act.

We are at the threshold of another shift: from human-curated, clinically-anchored knowledge to machine-generated, probability-driven output. Generative AI (GenAI) doesn't just retrieve facts; it synthesizes, contextualizes, and translates across vast datasets—blurring once-clear lines between fact and fiction. This shift changes how knowledge is created, evaluated, and trusted, and it demands a new framework for clinical use.

This special edition of the UpToDate® Point of Care Report sets out to build a bridge between the rigor of traditional clinical knowledge and the promise of GenAI. We aim to engage deliberately and responsibly: preserving provenance, anchoring outputs to evidence, and keeping clinicians at the center of judgment.

This report also outlines our philosophy and the practical systems we've built to support this transition. Our goal is to make GenAI clinically dependable: transparent in its sources, auditable in its reasoning, and accountable to both evidence and expert clinical oversight—so innovation advances patient care without fragmenting truth.

"GenAI has the potential to be a powerful tool for supporting sustainability in healthcare organizations right now, as well as preparing them for a more efficient future."

Greg Samios, CEO,
Wolters Kluwer Health

Building the bridge: Our philosophy

Every major shift in how knowledge is created and shared brings both progress and risk. What makes this moment different is that we can act intentionally. In medicine, that means preserving what matters most: validated knowledge, sound reasoning, and critical thinking—the bedrock of safe, effective care.

UpToDate® Expert AI[1] was built with this principle at its core. It starts with expert-authored, peer-reviewed content from UpToDate—written by clinical experts who use their judgment to interpret the evidence and apply it to real-world scenarios. GenAI extends that foundation, reinforced by multifaceted validation and secured with built-in safeguards. The result is augmentation, not replacement: technology that supports clinical reasoning rather than shortcutting it.

Our approach rests on three core tenets:

Provenance—Clinical insights must be traceable to trusted sources, with transparent processes that show how knowledge is created, validated, and maintained.

Validation—No single measure is sufficient. Meaningful evaluation at the point of care requires multiple, complementary, yet orthogonal methods tailored to clinical use and context.

Guidance—We design to support clinical reasoning and judgment, presenting relevant clinical considerations and supporting clinicians as they navigate decisions, with safeguards that preserve autonomy and control.

These principles are designed to counter real risks: erosion of reliable knowledge, conflation of engagement with effectiveness, and loss of clinical skill through over-reliance on opaque systems.

Think of our approach as a bridge: a foundation of provenance; a substructure of validation for the point of care; and a superstructure that guides clinical reasoning with guardrails and feedback. This design connects proven, evidence-based knowledge and clinical expertise with transformative technology—advancing patient care while promoting reliability, safety, and trust.

1. UpToDate Expert AI is available for clinical decision support (CDS) purposes, including in the context of patient care, in the United States only. Authorized users outside of the United States may access UpToDate Expert AI solely for internal evaluation purposes and to provide feedback.

"For clinicians, UpToDate Expert AI harnesses the clinical wisdom of our thousands of expert contributors to think like a clinical colleague—answering complex, highly specific questions with nuance and clarity. We are providing a new level of support for better decisions by the care team and better care for patients, while meeting the needs of enterprises for transparency and governance."

Peter Bonis, MD,
Chief Medical Officer,
Wolters Kluwer Health

The foundation: Why provenance matters

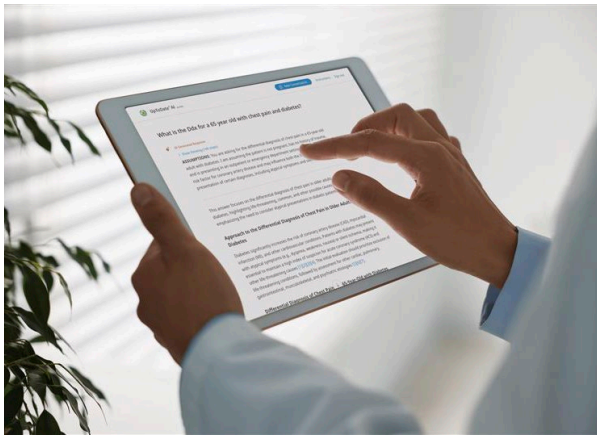

Provenance means knowing where knowledge comes from, how it is created, and who stands behind it. In clinical decision-making, that traceability is essential. Amid information overload and evolving standards, provenance helps clinicians separate what is merely available from what is current, reliable, and clinically relevant so judgments are grounded in trustworthy knowledge.

Origin—Where the knowledge comes from
AI-generated content must be grounded in purpose-built, point-of-care resources—not an undifferentiated ingestion of the entire medical literature. Clinical decision-making requires discerning which evidence matters, when, and why; not simply a summarization. UpToDate Expert AI is grounded only in the expert-authored, peer-reviewed, continuously updated content from the UpToDate team. Every day, new clinical research is evaluated for clinical relevance, context, methodological rigor, readiness to inform or change practice, and incorporated into UpToDate's corpus.

Derivation— How answers are produced
Trustworthy output depends on a trustworthy process. The system is designed to mirror clinical reasoning: embedding logic into how content is indexed, how prompts are constructed, and when probabilistic synthesis is appropriate versus when deterministic retrieval is required. Transparency is built in: Clinicians can see how an answer was formed and trace sources instantly—with one-click access to the exact UpToDate citations—so synthesis never obscures evidence.

Accountability— Who stands behind it
GenAI-supported solutions reflect the judgment and integrity of their makers. UpToDate Expert AI is shaped by practicing physicians and pharmacists with backgrounds in specialty care, hospital administration, patient safety, medical education, publishing, and content technology. It draws on Wolters Kluwer's global network of 7,600+ contributors supported by internal physician editors trained in evidence-based methodology. Visible authorship and editorial stewardship reinforce responsibility and trust—avoiding anonymity that may erode confidence.

Why this matters

Provenance is not a technical preference; it is a clinical requirement. By anchoring GenAI in accountable, expert-authored, and traceable knowledge—and by making derivation transparent—clinicians can question, verify, and apply information with confidence.

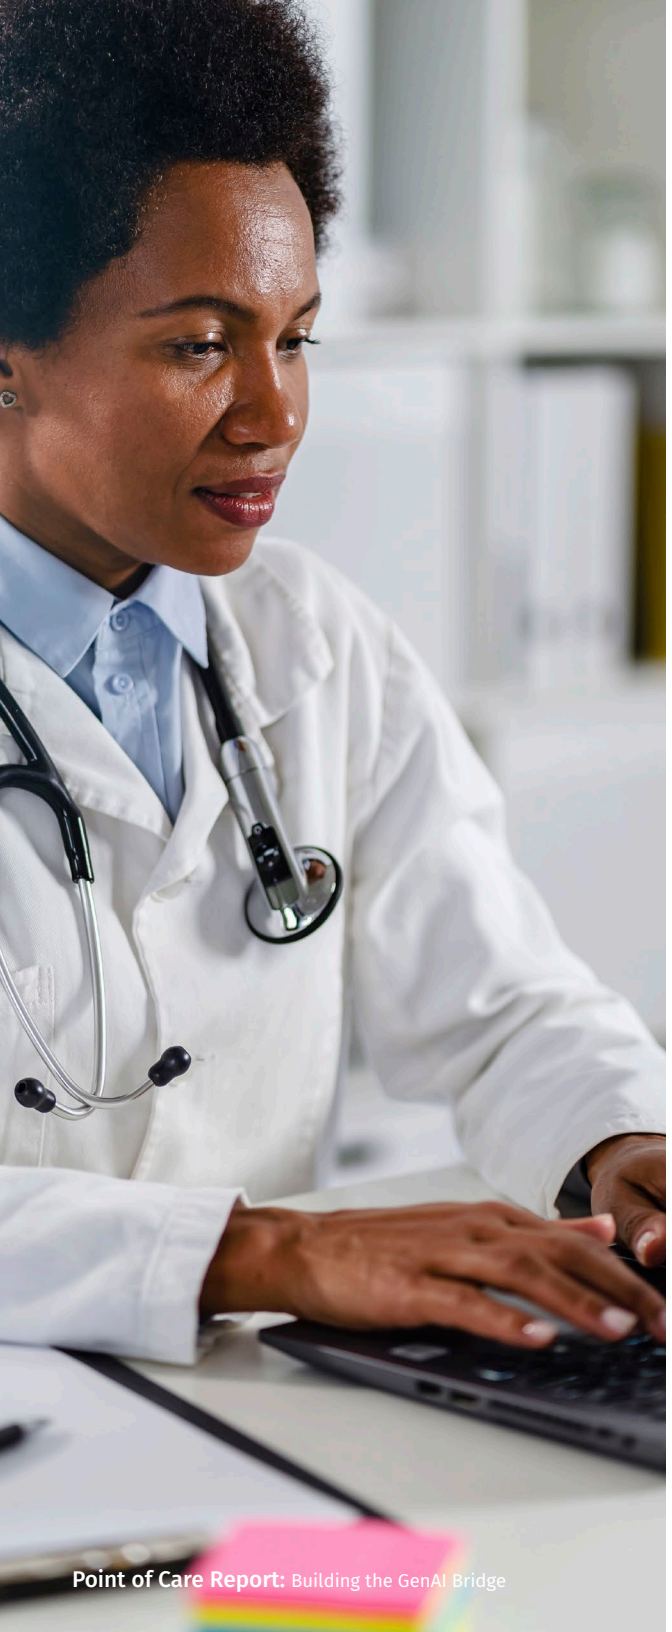

Substructure: Validation for the point-of-care

In the era of generative technology, every output is a calculation—a probability, not a fixed truth. The same question can yield different answers depending on context or phrasing. For clinical use, validation must focus on reliability, relevance, and actionability at the point of care. Here, we review some commonly used benchmarks in the industry and elements of our approach to validation.

Prominent external benchmarks (examples) — useful, but not sufficient on their own.

- **USMLE-style exams**. Provide a comparative signal of clinical knowledge and scenario interpretation. Limitations may include potential training-data contamination, curated question sets, multiple-choice constraints, and weak alignment to real-time clinical needs (e.g., they don't assess source grounding, transparency, or appropriate abstention).

- **Clinical case challenges.** CPCs, like those published in The New England Journal of Medicine or provided by NEJM are closer to clinical reality with multi-step reasoning under uncertainty. They are potentially still limited by possible presence in training corporate, reliance on an individual expert's approach, emphasis on rare/teaching cases, and scoring variability.

Together, these yield directional signals—not proof of safety or effectiveness at the bedside—so they must be complemented by validation designed for clinical care.

Internal, point-of-care validation—multi-layered and orthogonal

- **Expert review and red teaming**. Internal physician editors and our global contributor network evaluate reliability, relevance, and clarity, and stress-test edge biases and other failure modes that matter in practice.

- **Semi- and fully-automated evaluation.** Proprietary, UpToDate-authored clinical rubrics define high-quality responses across medicine, pharmacy, and patient. Automated harnesses track variability, source grounding and evidence use, appropriateness of abstention, and alignment with intended interaction design.

- **Continuous monitoring.** Not just human, but clinical expert-in-the-loop feedback captures clinician ratings and flags; findings drive rapid iteration and updates under clinical governance.

Why this matters

Validation is the safeguard that turns possibility into reliability. In clinical care, it's not enough for GenAI to sound plausible—it must be provably reliable and contextually relevant. That's why point-of-care validation must be multifaceted: combining expert judgment, automated rigor, and continuous feedback, and ultimately demonstrating improvement in real-word outcomes to meet the standards of medicine.

Validation isn't a final step—it's a continuous process of refinement and accountability. It protects against innovation outpacing responsible development, and promotes the delivery of insights at the bedside that are worthy of the decisions they inform.

"There is a transformation in how younger generations want to consume and use content, with a practical need to save as many seconds as possible while in workflow. We're aiming to deliver a simulated conversation experience in clinical decision support, aligning with evolving clinician needs."

Yaw Fellin, SVP and General Manager,
Clinical Decision Support and Provider Solutions,
Wolters Kluwer Health

Superstructure: Guidance, guardrails, and feedback

Even with a strong foundation of expert-authored content and rigorous validation, GenAI requires more to be clinically dependable. Its potential to support decision-making must be matched by systems that constrain risk, guide use, and continuously learn from practice. That's why we've built a superstructure that includes embedded guardrails, cues that encourage deeper critical thinking, and expert-reviewed feedback—so that interactions reinforce sound clinical reasoning and appropriate care.

Embedded guardrails
Safety was prioritized from the outset. Hard exclusions support the prevention of generation in high-risk domains, and logic determines when generative synthesis is appropriate versus when deterministic retrieval is preferable. Extensive targeted controls address areas where precision is critical, such as drug information, adverse effects, and bias. These guardrails are foundational and expand as the platform evolves.

Critical thinking cues
The interface is designed to support informed, clinical use. Visual cues and structural elements help clinicians engage thoughtfully, while transparent sourcing allows instant verification by the clinician. Every response shows how it was formed and where it came from—reinforcing that generative output is a starting point for reasoning, not a final answer.

Clinical reasoning and decision support
Generative responses are enriched with structured logic derived from the UpToDate editorial process. Assumptions, suggested next steps, and decision branch points reflect the same clinical reasoning clinicians rely on in UpToDate. These aren't generic features or technologic tricks—they are purpose-built to anticipate user need at the point of care.

Expert-in-the-loop feedback
User feedback on responses, queries, and system behavior is reviewed by clinicians. This expert-in-the-loop model promotes a system that evolves with real-world use, not just data. What makes our approach unique is how tightly feedback connects to content development. When users flag issues or submit questions, it informs not just the model, but the underlying content itself. That content is reviewed and, as necessary, refined and expanded by experts, creating a dynamic loop where clinician input drives meaningful improvement.

Why this matters

Guardrails, expert guidance, open feedback, and transparent continuous improvement are not enhancements—they are essential. Without them, even well-validated systems can mislead, confuse, or erode clinical judgment. But with them, GenAI becomes a trustworthy partner: one that respects complexity, reinforces reasoning, and improves continuously under expert stewardship. This superstructure supports innovation that doesn't just reach the bedside—it arrives responsibly and ready to support care and safety.

Beyond the bridge: Sustaining clinical knowledge in a generative era

As GenAI becomes embedded in clinical workflows, the work of maintaining a stable, trustworthy knowledge base becomes more urgent—not less. Without it, we risk systems that appear capable but lack the rigor that clinical care demands. That means evolving content stores to support generative synthesis without compromising reliability and developing real-world validation methods that reflect how these tools perform in practice—not just in benchmarks.

It also means preparing clinicians to use these tools wisely. AI literacy must become part of medical education, reinforcing the habits of critical thinking that define good medicine. As generative systems reshape how knowledge is created and interpreted, we must confront deeper questions: What does it mean for something to be reliable when it is synthesized? Who is accountable for probabilistic knowledge? How do we preserve nuance and uncertainty in systems designed to produce answers?

These are not questions for one company or one generation. They require open, ongoing dialogue across the medical community. We invite that conversation—not to defend the past, but to shape a future where technology strengthens clinical judgment and where innovation remains anchored in trust, transparency, and care.

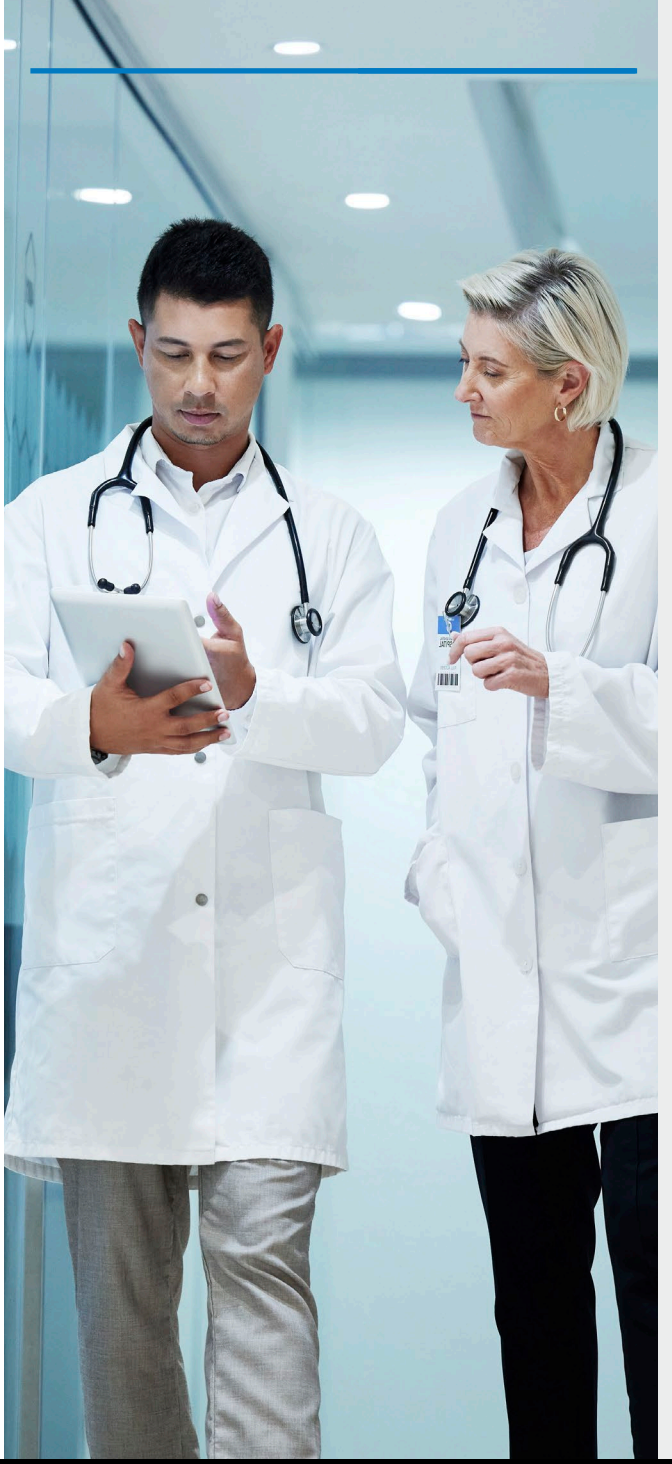

Contact our team →

Learn about about UpToDate Expert AI, our generative solution built on the trusted legacy of UpToDate clinical decision support.

Explore more issues →

The UpToDate Point of Care Report series has insights that can help address your organizational challenges.

Wolters Kluwer

 12.25 | 001371