



Generative AI and the changing dynamics of clinical consultations

David Fraile Navarro and colleagues consider how generative AI can be safely integrated into clinical consultations to ensure patient centred decision making and accountability

David Fraile Navarro,¹ Marcus Lewis,² Charlotte Blease,^{3,4} Rupal Shah,⁵ Sara Riggare,³ Sylvie Delacroix,⁶ Richard Lehman⁷

Generative artificial intelligence (AI) has moved rapidly from experiment to everyday use in clinical encounters. Powered by large language models (LLMs), generative AI systems generate novel content in response to prompts rather than following scripted rules. Clinicians and patients commonly access generative AI through conversational interfaces such as ChatGPT that respond in natural language, while specialised applications such as AI powered medical scribes are being increasingly adopted for clinical documentation. Generative AI's fluent but only partially verifiable reasoning changes how clinical knowledge is formed, communicated, and tested within the doctor-patient relationship, shifting consultations from exchanging facts to co-producing explanations.

The consultation is therefore becoming a three way conversation in which explanations are shaped by doctor, patient, and AI, creating what we call triadic care. As part of this shift it is important to record use of AI so that healthcare systems, researchers, clinicians, and patients can learn and improve practice. Openness also clarifies how AI influenced clinical decision making and supports patient understanding and trust. When AI use is not explicitly discussed, it becomes harder to assess its effect on clinical judgment, patient autonomy, and the therapeutic relationship.

The challenge for clinicians is to integrate AI in ways that preserve the benefits while managing potential risks. This article, part of a BMJ series on use of generative AI, focuses on consultation based care, especially primary care, where generative AI most directly affects communication and decision making. Other articles in the series will consider the patient experience and the competencies that clinicians need to use AI transparently and effectively within the clinical encounter.^{1,2}

AI has entered the consultation room

Clinicians report widespread use of AI across diagnostic support, documentation, patient communication, and administrative tasks.³ In 2024, a survey of 1183 US physicians found that two thirds used AI tools in practice, up from 38% the previous year.³ In the UK, general practitioners report using AI for checking drug interactions, diagnostic suggestions, and drafting letters.⁴ Early studies suggest generative AI can broaden differential diagnoses and enhance clinical reasoning,⁵ while ambient scribes reduce documentation burden and improve the perceived quality of notes.^{6,7}

Patients and carers are also using AI. In Australia, one in 10 adults reports seeking health advice from ChatGPT.⁸ People describe using it to search for explanations, assemble care plans, or seek second opinions. In one case, a child with tethered cord syndrome received a diagnosis only after his mother consulted ChatGPT; in another, a carer used AI to complete a care plan while awaiting specialist input.^{9,10} Yet similar systems have misclassified neurological symptoms, delaying stroke treatment.¹¹

Generative AI is not just another information tool. Unlike a web search, which yields links, chatbots deliver synthesised reasoning in natural language. This can make medical thinking feel co-produced, extending cognition beyond the individual into interaction with an artificial partner.¹² Patients can now attempt differential diagnoses once reserved for specialists, and clinicians can rapidly test hypotheses about rare diseases.¹³

Generative AI also changes how medical knowledge can be verified. Traditional sources, papers, and guidelines provide citations that can be checked, whereas AI provides confident answers without complete and verifiable reasoning.¹⁴ Importantly, only 19% of users crosscheck chatbot outputs compared with 50% for search engines,¹⁵ yet users trust AI authored responses as much as those from doctors, even when inaccurate.¹⁶ This changes something fundamental about medical information: when neither clinician nor patient can trace how advice was reached, it becomes something to interpret rather than verify.

Working without AI may come to feel like working without the internet: possible, but increasingly impractical.¹⁷ This will shift clinical expertise from producing answers to interpreting them with patients, translating AI suggestions, testing them against clinical context, and weighing them alongside patient values and lived experience. Generalists who already integrate uncertainty and conflicting information may find this transition natural, but the shift is substantial: from knowing answers to interpreting AI generated suggestions, from reasoning alone to helping patients weigh generative explanations against their own circumstances and values.

Medicine has undergone similar transitions before, such as when laboratory tests and diagnostic radiology displaced clinical observation as the arbiter of truth. But generative AI introduces a different challenge. The quality of triadic care depends on whether its reasoning can be examined and tested: what did the AI consider, why, and with what

¹ Centre for Health Informatics, Australian Institute of Health Innovation, Macquarie University, Sydney, Australia

² Abbey Medical Centre, London, UK

³ Patient representative, Department of Women's and Children's Health, Participatory eHealth and Health Data Research Group, Uppsala University, Uppsala, Sweden

⁴ Digital Psychiatry, Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, Massachusetts, USA

⁵ Professional Development Team, NHS England Workforce Training and Education, London, UK

⁶ Centre for Data Futures, Dickson Poon Law School, King's College London, London, UK

⁷ King's College London, London, UK

Correspondence to: D Fraile Navarro david.frailenavarro@mq.edu.au

Cite this as: *BMJ* 2025;391:e085325 <http://doi.org/10.1136/bmj-2025-085325>

uncertainty? The question is not whether clinicians can adapt, but whether healthcare systems will provide the infrastructure to make this adaptation visible, safe, and improvable. Without open examination of AI responses, doctors and patients may form different understandings of a condition or its management, making it harder to build trust and reach shared decisions.

Trust and accountability in triadic care

AI integration in consultations promises better decisions, less onerous documentation, and potentially stronger relationships. But how AI use becomes visible matters. Policies remain uneven and documentation seldom notes AI's role, making learning and oversight difficult.³ Some patients may hesitate to mention that they have used AI, just as they feared being dismissed for consulting "Dr Google."¹⁸ A brief, non-judgmental prompt, "Have you used AI to look into this? Shall we review it together?" can reveal patient use without stigma.

Patients also say they want to know when doctors have used AI, though preferences vary by age, sex, and income.^{19 20} In addition, trust will affect how patients view AI's contribution. Patients who already trust their clinician or health system are more likely to expect AI to help,²¹ while those with lower trust may view AI either as an alternative authority or with added scepticism. Transparency must go beyond explaining what the tool does, why it is being used, and how outputs are monitored. An open discussion of its results, along with clear explanation of its purpose and oversight, is required to ensure AI enhances agency rather than erodes confidence.²²

Evidence from a US qualitative study using cardiovascular AI scenarios supports the value of transparency. Participants reported higher trust in both the clinician and the decision when use was acknowledged and outputs were reviewed together.²³ This aligns with triadic care, where brief acknowledgment is paired with joint review and lightweight documentation. In ambient scribe workflows, explicit consent and clinician review of generated notes have been crucial for sustaining trust.⁶

Yet transparency can backfire if overused. In a 2023 survey of 1455 members of the patient advisory members of Duke Health, a health system provided by Duke University in the US, AI drafted portal messages were rated more empathetic than human ones, but satisfaction fell once participants learnt the messages came from AI.²⁴ To avoid consent fatigue, proportionate approaches are needed, **prioritising consent for tools that pose meaningful risks** (such as AI systems recommending treatment changes or predicting disease progression) or offer genuine opportunities for patient action (such as AI generated care summaries patients can review and edit before they become part of the medical record).²²

Accountability is closely tied to observability. Primary care decisions are iterative and uncertain, and generative AI can blur where responsibility lies. Trials show that even when AI models outperform physicians, the accuracy of clinicians with and without access to them is similar,⁵ underscoring that judgment remains human and that knowing who used what, for what, with what effect matters for learning and governance. Commercial systems require greater clinical insight because of their market driven development and limited visibility into training data, updates, and data reuse policies.²⁵ Without this visibility, product changes can outpace safety processes and undermine trust. Commercial influence may also indirectly shape care—for example, generative AI systems trained on data from highly specialised clinical settings could risk nudging clinicians and patients towards defensive medicine.

Finally, design and workflow determine whether transparency becomes collaboration. Experience with electronic health records shows that poor human-computer interaction creates burden rather than value.²⁶ Design principles for human-AI interaction,²⁷ showing why a system made a recommendation, displaying how certain it is, allowing corrections, and providing secure sharing, make it possible for clinicians and patients to examine AI outputs together. These features do not resolve deeper questions about judgment or trust, but they are the practical foundation for integrating AI safely and transparently into care.

Integrating triadic care into the clinical encounter

Triadic care does not require wholesale redesign of healthcare systems. It can be embedded through consistent transparency practices, light touch policies, and practical tools that reinforce trust, accountability, and shared understanding.

Health systems can support transparency with standardised templates that clinicians adapt to individual consultations. Brief, plain language scripts normalise AI openness for common scenarios such as ambient scribes, decision support tools, and patient generated AI content. Specific communication approaches clinicians need to implement these practices effectively are dealt with elsewhere in this series.²

Simple documentation structures can make AI use auditable without creating burdensome paperwork. Adding an "AI involvement" field in the electronic health record, with structured options for tool name, purpose (eg, interaction check, differential diagnosis, note drafting), and clinician response (accepted/modified/rejected with reason) would require minimal effort. Recording a brief rationale for rejecting an AI contribution makes patterns visible for safety learning and monitoring equity, as the performance of models can vary across different patient populations.^{28 29}

Institutional policies can reinforce these habits. Organisations can include AI related consent options in notification templates, maintain registries of approved tools linked to regulatory clearance, and adapt case reviews to ask what AI has added. Incident reporting systems that flag AI related near misses create a record that services can learn from.²⁰ Since patients already use AI to research symptoms and prepare questions, simple prompts such as, "If you've used any online tools or AI, feel free to show me and we'll review it together," normalise patient disclosure and create opportunities for joint review.

Technology design is central. Vendors and health IT teams should provide systems that let clinicians and patients see how AI arrived at a suggestion, what inputs were used, and how confident the system is. Useful features include explanations of recommendations, uncertainty displays, editable outputs with version history, and audit logs. Secure channels (for example, patient portals) can support sharing when needed. These elements should be evaluated based on whether they improve decision quality and patient understanding.²⁷

Regulation and governance should keep accountability proportionate, **clearly defining who is accountable for what while avoiding excessive regulatory burden**. Clinicians will retain responsibility for final clinical decisions. Organisations should supply safe, validated tools and set out the conditions for their use, procurement standards, staff training, performance monitoring, and incident reporting (box 1).

Box 1: Minimum transparency standards for generative AI tools used in care

- *Purpose and validation*—State the intended uses, clinical contexts evaluated, and headline performance with typical error rates
- *Limits*—Describe known limitations and situations where performance falls off
- *Data and equity*—Summarise the demographic and clinical data used for development and testing and report performance across key patient groups
- *Updates and versioning*—Note the update schedule, what changed, and how changes are monitored
- *Governance and data use*—Explain data handling, auditability, and routes for incident reporting and review

Vendors should be obliged to disclose purpose, limits, and update schedules of their systems. Minimum disclosure standards should apply whenever AI materially informs advice or documentation, aligned with existing informed consent requirements.²² Regulatory frameworks must also include technology companies directly, mandating transparency in model capabilities and limitations.

Given the rapid evolution of generative AI, the immediate task is to make its use sufficiently transparent to understand its effects on clinical reasoning, patient autonomy, and the doctor-patient relationship. Research priorities include describing real world use by patients and clinicians; documenting prevalence and drivers of non-disclosure^{20 21}; testing the safety and effectiveness of disclosing AI involvement and jointly reviewing AI outputs^{5 15 27}; developing minimal standards for documentation and audit⁵; designing interfaces that support inspectable reasoning and collaboration²⁷; and equity and language barriers.^{28 29} Beyond documenting use of AI, research must examine how it transforms the therapeutic relationship itself, from one based on the doctor's expertise to one centred on helping patients interpret what AI tells them.

AI is already reshaping consultations in ways we are only beginning to understand. Adoption of triadic care makes AI's role visible and enables inspection of its reasoning, collaborative testing, and documenting of outcomes. This not only meets the immediate needs for transparency and accountability but prepares medicine for a deeper transformation as AI diagnostic capabilities approach or exceed human performance. Transparency alone cannot resolve the challenge of relying on reasoning neither party can fully verify, but without it we cannot examine how AI affects clinical judgment, patient autonomy, and therapeutic relationships, or develop frameworks for medicine's evolution from knowledge holder to interpreter of AI generated advice. With observable practices and institutional support, medicine can grapple with questions it hasn't fully articulated: what does clinical expertise mean when knowledge is abundant, but verification is scarce?

Key messages

- Generative AI is already used by both clinicians and patients in consultations, but its role often goes undocumented
- Triadic care—making AI use visible and its reasoning inspectable—enables healthcare systems to study and govern AI's effects on clinical decisions
- Clinical expertise is shifting from knowing answers to helping patients interpret AI generated information in context
- Simple infrastructure (documentation standards, transparent technology, proportionate accountability) can make this shift observable and safe

- Research must examine how AI transforms the doctor-patient relationship and develop frameworks for this evolution

Contributors and sources: The authors have emerged from our diverse expertise and perspectives on the intersection of artificial intelligence with clinical practice. DFN, the lead author, is a GP with clinical experience in Spain and the UK and currently in an academic role in Australia researching health informatics and applying generative AI technologies to healthcare problems. ML brings frontline clinical perspective as a practising GP and GP training programme director in London. CB contributes expertise in digital health and ethics from her work at both Uppsala University and Harvard Medical School. RS offers insights as a GP with professional development expertise and interest in philosophical hermeneutics. SR brings the crucial patient-researcher perspective, having been involved from the early planning stages. SD contributes philosophical and legal expertise in AI ethics. RL offers perspective from his extensive work on evidence based medicine, uncertainty in general practice, and shared decision making. All authors contributed to conceptualising, drafting, and revising this article. DFN is the guarantor.

Public and patient involvement: As a coauthor and patient-researcher, SR has been involved from the earliest discussions and planning of this series, contributing to the conceptualisation of triadic care and ensuring that patient considerations remain central to our analysis. Her insights have strengthened our understanding of how patients are already using AI tools to navigate healthcare and have shaped our recommendations for future developments.

AI use: Generative AI (mostly Claude 4 Opus, but also ChatGPT) was used as an initial soundboard to expand on the original series ideas and article structure; to craft a cohesive narrative avoiding textual duplication and providing clarity and conciseness, and to help craft and refine use cases and examples.

Competing interests: We have read and understood BMJ policy on declaration of interests and declare that SR is a member of the BMJ patient panel.

Provenance and peer review: Commissioned; externally peer reviewed.

We thank Iain Marshall, King's College London, for contributing to the initial discussions on this article and series.

- 1 Blease C, Lewis M, Riggare S, Fraile Navarro D, Lehman R. Generative AI and patient agency in the clinical encounter. *BMJ* 2025;391:e085323.
- 2 Clinical competencies for using generative AI in patient care. *BMJ* 2025;391:e085324.
- 3 American Medical Association. 2025 2 in 3 physicians are using health AI—up 78% from 2023. 2025. <https://www.ama-assn.org/practice-management/digital-health/2-3-physicians-are-using-health-ai-78-2023>
- 4 Blease CR, Locher C, Gaab J, Hägglund M, Mandl KD. Generative artificial intelligence in primary care: an online survey of UK general practitioners. *BMJ Health Care Inform* 2024;31:e101102. doi: 10.1136/bmjhci-2024-101102 pmid: 39288998
- 5 Goh E, Gallo R, Hom J, et al. Large language model influence on diagnostic reasoning: a randomized clinical trial. *JAMA Netw Open* 2024;7:e2440969. doi: 10.1001/jamanetworkopen.2024.40969 pmid: 39466245
- 6 Shah SJ, Devon-Sand A, Ma SP, et al. Ambient artificial intelligence scribes: physician burnout and perspectives on usability and documentation burden. *J Am Med Assoc* 2025;328:80. doi: 10.1093/jama/ocae295 pmid: 39657021
- 7 Tierney AA, Gayre G, Hoberman B, Mattern B, Balleca M, Wilson Hannay SB, et al. Ambient artificial intelligence scribes: learnings after 1 year and over 2.5 million uses. *NEJM Catal* 2025;6:CAT.25.0040doi: 10.1056/CAT.25.0040.
- 8 Ayre J, Cvejic E, McCaffery KJ. Use of ChatGPT to obtain health information in Australia, 2024: insights from a nationally representative survey. *Med J Aust* 2025;222:-2. doi: 10.5694/mja2.52598 pmid: 39901778
- 9 deBronkart D. Forced to wait while Dad suffers, Hugo Campos takes action: #PatientsUseAI. Patients Use AI. 5 Aug 2024. <https://patientsuseai.substack.com/p/forced-to-wait-while-dad-suffers>
- 10 Partners in diagnosis: ChatGPT, a mother's intuition, and a doctor's expertise with Courtney Hofmann and Dr. Holly Gilmer. *NEJM AI Grand Rounds* 28 Aug 2025. <https://ai-podcast.nejm.org/e/partners-in-diagnosis-chatgpt-a-mother-s-intuition-and-a-doctor-s-expertise-with-courtney-hofman-and-dr-holly-gilmer/>
- 11 Saenger JA, Hunger J, Boss A, Richter J. Delayed diagnosis of a transient ischemic attack caused by ChatGPT. *Wien Klin Wochenschr* 2024;136:-8. doi: 10.1007/s00508-024-02329-1 pmid: 38305909
- 12 Clark A, Chalmers D. The extended mind. *Analysis* 1998;58:-19doi: 10.1093/analysis/58.1.7.
- 13 Madrid Health Service: a pioneer in applying generative artificial intelligence to improve diagnosis for patients with rare diseases. Microsoft Pensa, 2023. <https://news.microsoft.com/es-es/2023/09/15/madrid-health-service-a-pioneer-in-applying-generative-artificial-intelligence-to-improve-diagnosis-for-patients-with-rare-diseases/>
- 14 Chen Y, Benton J, Radhakrishnan A, Uesato J, Denison C, Schulman J, et al. Reasoning models don't always say what they think. *ArXiv* 2025 [Preprint.] <https://arxiv.org/abs/2505.05410>
- 15 Yun HS, Bickmore T. Online health information-seeking in the era of large language models: cross-sectional web-based survey study. *J Med Internet Res* 2025;27:e68560. doi: 10.2196/68560 pmid: 40163112
- 16 Shekar S, Pataranutopom P, Sarabu C, Cecchi GA, Maes P. People overtrust ai-generated medical advice despite low accuracy. *NEJM AI* 2025;2doi: 10.1056/Aloa2300015.

- 17 Flinn R. Dr ChatGPT will see you now. *Wired* 10 Jul 2025. <https://www.wired.com/story/dr-chat-gpt-will-see-you-now-artificial-intelligence-llms-openai-health-diagnoses/>
- 18 Cuteanu A, Seguin M, Ziebland S, et al. Qualitative study: patients' enduring concerns about discussing internet use in general practice consultations. *BMJ Open* 2021;11:e047508. doi: 10.1136/bmjopen-2020-047508 pmid: 33910948
- 19 Park HJ. Patient perspectives on informed consent for medical AI: A web-based experiment. *Digit Health* 2024;10:20552076241247938. doi: 10.1177/20552076241247938 pmid: 38698829
- 20 Fehr J, Citro B, Malpani R, Lippert C, Madai VI. A trustworthy AI reality-check: the lack of transparency of artificial intelligence products in healthcare. *Front Digit Health* 2024;6:1267290. doi: 10.3389/fdgth.2024.1267290 pmid: 38455991
- 21 Nong P, Ji M. Expectations of healthcare AI and the role of trust: understanding patient views on how AI will impact cost, access, and patient-provider relationships. *J Am Med Inform Assoc* 2025;32:-9. doi: 10.1093/jamia/ocaf031 pmid: 40036944
- 22 Mello MM, Char D, Xu SH. Ethical obligations to inform patients about use of AI tools. *JAMA* 2025;334:-70. doi: 10.1001/jama.2025.11417 pmid: 40690211
- 23 Stroud AM, Minter SA, Zhu X, Ridgeway JL, Miller JE, Barry BA. Patient information needs for transparent and trustworthy cardiovascular artificial intelligence: a qualitative study. *PLOS Digit Health* 2025;4:e0000826. doi: 10.1371/journal.pdig.0000826 pmid: 40258073
- 24 Cavalier JS, Goldstein BA, Ravitsky V, et al. Ethics in patient preferences for artificial intelligence-drafted responses to electronic messages. *JAMA Netw Open* 2025;8:e250449. doi: 10.1001/jamanetworkopen.2025.0449 pmid: 40067301
- 25 Snider D. *Economic incentives and mass-market training: how profit motives led to inevitable hallucinations in large language models*. Social Science Research Network, 2025. <https://papers.ssrn.com/abstract=5181079>
- 26 Fry E, Schulte F. Death by a thousand clicks: where electronic health records went wrong. *Fortune* 2019. <https://fortune.com/longform/medical-records/>
- 27 Amershi S, Weld D, Vorvoreanu M, et al. Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* 2019:1-13.
- 28 Tang A, Tung N, Nguyen HQ, et al. Health information for all: do large language models bridge or widen the digital divide? *BMJ* 2024;387:e080208. doi: 10.1136/bmj-2024-080208 pmid: 39393817
- 29 Omar M, Soffer S, Agbareia R, et al. Sociodemographic biases in medical decision making by large language models. *Nat Med* 2025;31:-81. doi: 10.1038/s41591-025-03626-6 pmid: 40195448