Editorial

ISSN 2639-8478

Cancer Science & Research

The Agentic Shift in Oncology: How Multi-Agent AI Systems Will Reshape Cancer Care by 2028

Enrique Díaz Cantón, MD, MSc, MSc[1]* and Matías Díaz Cantón[2]

[1]*Professor of Oncology and Artificial Intelligence in Medicine, Instituto Universitario CEMIC, Director, Postgraduate Program in Artificial Intelligence and Medicine, Academia Nacional de Medicina, Buenos Aires, Argentina.*

[2]*AI Engineer, Instituto Tecnológico de Buenos Aires (ITBA), Buenos Aires, Argentina.*

***Correspondence:**
Enrique Díaz Cantón, MD, MSc, MSc, Professor of Oncology and Artificial Intelligence in Medicine, Instituto Universitario CEMIC, Director, Postgraduate Program in Artificial Intelligence and Medicine, Academia Nacional de Medicina, Buenos Aires, Argentina.

Received: 10 Jan 2026; **Accepted:** 23 Feb 2026; **Published:** 03 Mar 2026

Citation: Enrique Díaz Cantón, Matías Díaz Cantón. The Agentic Shift in Oncology: How Multi-Agent AI Systems Will Reshape Cancer Care by 2028. Cancer Sci Res. 2026; 9(1): 1-5.

ABSTRACT

Medicine is undergoing a paradigm shift from artificial intelligence (AI) as a passive analytical tool to AI as an active, semi-autonomous teammate. This transition is fueled by the convergence of frontier large language models (LLMs) with open-source agentic frameworks and standardized protocols such as the Model Context Protocol (MCP), which enable systems capable of complex reasoning, multi-step task execution, and real-world interaction with clinical environments. As of early 2026, multi-agent AI systems are demonstrating the capacity to transform oncology practice: augmenting multidisciplinary tumor boards with decision quality comparable to expert panels, compressing drug discovery timelines from years to months, and more than doubling clinical trial matching recall. The evidence is now substantial: an autonomous agent leveraging GPT-4 with multimodal precision oncology tools achieved 91% accuracy in clinical decision-making compared to 30.3% for the base model alone (Nature Cancer, 2025), while the MDAgents framework achieved best performance in seven of ten medical benchmarks at NeurIPS 2024. In this editorial, we examine the current evidence base and project the trajectory of agentic AI in oncology through 2028. We hypothesize that by the end of this decade, the integration of AI teammates—orchestrating complex care pathways, designing novel therapeutics, and simulating their effects in patient-specific digital twins—will begin to yield measurable improvements in patient outcomes, potentially establishing a pathway toward a 70% 10-year disease-free survival in advanced triple-negative breast cancer. The implications for clinical practice, regulatory frameworks, and the identity of the oncologist are profound.

Keywords

Agentic AI, Multi-Agent Systems, Precision Oncology, Digital Twin, Triple-Negative Breast Cancer, Large Language Models, Clinical Decision Support, Drug Discovery, Tumor Board, Model Context Protocol, Hallucination Reduction, Regulatory Frameworks.

From Tools to Teammates: A Conceptual Revolution

The history of AI in oncology has been one of incremental, tool-based progress: algorithms that classify radiological images, predict genomic signatures, or stratify patients by risk. These systems, however powerful, share a fundamental limitation—they are passive. They wait for a query, process it, and return a result. The clinician must provide the input, interpret the output, and decide what to do next. As Zou and Topol articulated in their landmark *Lancet* commentary, the next era demands a conceptual leap: from AI as a sophisticated calculator to AI as a healthcare teammate [1]. An AI agent, as defined by the comprehensive review by Truhn and colleagues in *Nature Reviews Cancer*, is a "(semi-) autonomous system capable of sensing, learning and acting upon its environment" [2]. Unlike a tool, an agent can take initiative: maintaining long-term memory of a patient's complex history, proactively monitoring incoming data streams, autonomously navigating databases such as electronic health records (EHRs), and orchestrating specialized sub-tools to execute multi-step workflows [1,2].

This evolution is powered by the synergy between frontier LLMs and robust agentic frameworks. LangGraph, the dominant framework for stateful graph-based multi-agent workflows, enables persistent state management, human-in-the-loop capabilities, and hierarchical orchestration patterns critical for clinical deployment [3]. CrewAI offers role-based agent design that mirrors the structure of a tumor board, where a coordinating oncologist delegates to pathologists, radiologists, and surgeons [3]. The Model Context Protocol (MCP), introduced by Anthropic in November 2024, has become the de facto standard for connecting AI agents to external tools and data sources—often described as "USB-C for AI"—using a client-server architecture with JSON-RPC 2.0 messages that expose tools, resources, and prompt templates [4]. For oncology, healthcare-specific MCP implementations such as BioMCP (connecting to PubMed, ClinicalTrials.gov, and genomic variant databases) and Health Record MCP (connecting to EHRs via SMART on FHIR) are already enabling AI agents to interact directly with clinical data infrastructure [4,5].

As of February 2026, the evidence is clear: the era of the AI teammate has begun, and oncology stands at the vanguard.

The Oncologist's New Teammate: Evidence as of February 2026

The burden of complexity

Every year, 20 million people are diagnosed with cancer globally, each presenting a unique constellation of histological subtypes, genomic alterations, comorbidities, and treatment histories [6]. Multidisciplinary tumor boards (MDTs) are the gold standard for personalized care planning, yet clinicians spend 1.5–2.5 hours per patient reviewing imaging, pathology, clinical notes, and genomic data in preparation [6]. MDT teams often have less than two minutes of discussion per case, with 7% of cases postponed [7]. The result: less than 1% of cancer patients globally access truly personalized treatment planning [6]. Administrative burden consumes twice the time physicians spend in direct patient interaction, contributing to a $4.6 billion annual physician turnover cost in the United States [8].

Multi-agent systems outperform single models

The evidence that multi-agent architectures surpass single-LLM approaches is now substantial. The MDAgents framework (Kim et al., NeurIPS 2024) introduced the first adaptive multi-agent framework mirroring real-world medical decision-making, dynamically routing simple cases to individual agents while escalating complex cases to multidisciplinary teams. MDAgents achieved best performance in seven of ten medical benchmarks, with group collaboration yielding an average accuracy improvement of 11.8% [9]. The most striking oncology result comes from Ferber et al. (*Nature Cancer*, 2025): an autonomous agent leveraging GPT-4 with multimodal precision oncology tools—including vision transformers for MSI/KRAS/BRAF detection and MedSAM for radiology—achieved 91.0% correct clinical conclusions compared to 30.3% for GPT-4 alone, a near-tripling of accuracy [10]. A Multi-Agent Conversation (MAC) framework for rare disease diagnosis deployed four doctor agents plus one supervisor, outperforming single models, Chain-of-Thought, and Self-Refine methods across 302 cases [11]. A systematic review of 20 studies found that all agent systems outperformed baseline LLMs, with single-agent tool-calling studies showing a median 53 percentage point improvement [12].

Institutional deployments

The Microsoft Healthcare Agent Orchestrator, piloted at Stanford Health Care, Johns Hopkins, Providence Genomics, and Mass General Brigham, deploys a coordinated team of specialized AI agents within Microsoft Teams [6]. Its Patient History Agent transforms three hours of manual record review into minutes; its Pathology Agent (Paige AI's Alba) delivers real-time digital pathology insights; and its Clinical Trials Agent achieves more than double the recall of the Criteria2Query baseline [6]. The Oxford-led TrustedMDT has received NHS Health Research Authority approval (25/HRA/5004) for a pilot study, representing a rigorous clinically validated approach [7]. EvoMDT, published in *npj Digital Medicine* in January 2026, demonstrated decision quality comparable to human MDTs while reducing response times by 30–40% across multiple cancer types [13]. For oncology-specific applications, OncoPainBot (*npj Digital Medicine*, 2026) deploys four specialized agents incorporating RAG with NCCN, CSCO, and WHO guidelines [14], while a Virtual AML Panel uses five agents for WHO classification and ELN prognostication [15].

Hallucination Reduction Through Multi-Agent Consensus

Hallucination remains the critical safety concern for clinical AI deployment. Multi-agent consensus mechanisms now offer quantitative solutions. The Iterative Consensus Ensemble (ICE) method demonstrated the most compelling results: three LLMs iteratively critique each other until consensus, achieving accuracy gains from 60.2% to 74.03% at final consensus—a 23% improvement—with the medical subset reaching 81.17%. Most items settled in 2–3 rounds with no fine-tuning required [16]. At ICML 2024, Smit et al. benchmarked multi-agent debate strategies on medical datasets, finding that every evaluated multi-agent system outperformed single-agent methods, with Multi-Persona achieving approximately 15% improvement on USMLE [17].

Retrieval-augmented generation (RAG) in clinical oncology has become a critical complementary strategy. A study of RAG chatbots using curated Cancer Information Service sources achieved 0% hallucination rate for GPT-4 [18]. RadOncRAG evaluated 15 LLMs and found that RAG improved non-reasoning models significantly (Gemini-2.0 +6.7%, GPT-4o +5.7%), though reasoning models already achieving 91.6% showed no additional benefit [19]. Ke et al. showed GPT-4 with RAG achieved 96.4% accuracy versus 86.6% for humans ($p=0.016$) with zero hallucinations across 14 clinical scenarios [20].

Important caveats have emerged. Recent theoretical work formalized multi-agent debate as Bayesian posterior belief updates, proving that debate induces a martingale over beliefs, meaning simple majority voting accounts for most observed gains and performance is bounded by the strongest individual agent's accuracy [21]. These findings suggest that heterogeneous agents with structured interventions—rather than homogeneous debate alone—are essential for genuine deliberative benefit.

Clinical Trial Matching: The Most Deployment-Ready Application

Clinical trial matching represents the most mature application of agentic AI in oncology. TrialGPT (Jin et al., *Nature Communications*, 2024) achieved 87.3% matching accuracy and reduced screening time by 42.6% across 183 patients with over 75,000 trial annotations [22]. A multi-agent platform with specialized OncoAgents and an oncology-specific knowledge graph (Kurnaz et al., ASCO 2025) prospectively evaluated 3,804 patients, achieving 82% trial matching accuracy versus 47% for zero-shot GPT-4o [23]. Cleveland Clinic's AI analyzed 840,523 patients against 74 trials, identifying over 350,000 preliminary matches and enrolling 189 patients over six months [24]. Memorial Sloan Kettering deployed Triomics' AI platform across its full trial portfolio in 2025 [24], and Mount Sinai launched an AI-powered clinical trial-matching platform identifying opportunities "earlier and more consistently" [25].

The Two-Year Horizon: From Augmentation to Transformation (2026–2028)

Gartner predicts that by 2028, 15% of day-to-day work decisions will be made autonomously by agentic AI—up from 0% in 2024—and that 70% of healthcare stakeholders will have adopted AI [26]. The Boston Consulting Group projects that agentic AI will compress drug development timelines from years to months [27]. Based on these trajectories, we project three areas of fundamental transformation.

The digital twin as standard of care

MD Anderson Cancer Center is already using digital twin models to simulate tumor responses to therapies in silico, with published results in predicting bone metastasis treatment response [28]. The UT Austin Oden Institute and the National Cancer Institute are developing computational frameworks for "theranostics digital twins" [29]. By 2028, we anticipate that agentic AI will make this a clinical standard: an AI teammate creating, maintaining, and continuously updating a patient's digital twin from real-time multi-omic data, simulating treatment regimens, and presenting ranked personalized recommendations with confidence scores.

Drug discovery at machine speed

AI-native biotechnology firms report Phase I safety success rates between 80% and 90% for AI-designed molecules, nearly doubling the historical average of approximately 50% [30]. Iambic Therapeutics' IAM1363, the first AI-discovered cancer drug, entered clinical trials in October 2025 [31]. Multi-agent systems are accelerating this: the Bio AI Agent integrates target discovery, toxicity prediction, and clinical protocol design for autonomous CAR-T therapy development [32], while PharmaMar and Globant announced a multi-agent AI collaboration for cancer drug discovery in February 2026 [33].

The closed-loop paradigm: a hypothesis for advanced TNBC

The convergence of intelligent MDT augmentation, AI-accelerated drug discovery, and patient-specific digital twins enables hypotheses once unthinkable. Consider advanced triple-negative breast cancer (TNBC), where metastatic disease carries a median overall survival of 12–18 months and a 5-year survival of approximately 11–15% [34]. While KEYNOTE-522 demonstrated 5-year event-free survival of 81.2% in early-stage disease [35] and ASCENT-04 showed a 35% reduction in progression risk with sacituzumab govitecan plus pembrolizumab in first-line metastatic PD-L1-positive TNBC (HR 0.65, $p<0.001$) [36], a cure for advanced disease remains elusive.

We propose that by the mid-2030s, a 70% 10-year disease-free survival in advanced TNBC is achievable through a closed-loop, AI-driven ecosystem: (1) early interception via continuous multi-omic monitoring detecting metastatic relapse months before imaging; (2) personalized therapeutic design by drug-discovery sub-agents tailoring AI-optimized ADCs or neoantigen vaccines to the tumor's mutational landscape; (3) digital twin simulation optimizing dosing, scheduling, and combination strategies; and (4) proactive resistance management through real-time ctDNA monitoring with autonomous second-line strategy design before clinical relapse. This represents a model where the oncologist anticipates and preempts disease progression, supported by an AI teammate with access to the entirety of the world's oncological knowledge.

The Regulatory Gap

As of December 2025, the FDA has authorized over 1,300 AI-enabled medical devices, with 258 in 2025 alone—the most in any single year—yet none incorporates LLMs [37]. The vast majority (97%) entered via the 510(k) pathway, with radiology dominating at approximately 76% of authorizations [37]. The EU AI Act, classifying most healthcare AI as "high-risk" with strict requirements effective August 2026, and WHO guidance flagging hallucinations, data bias, and automation bias, represent evolving but incomplete frameworks [38].

Critically, no regulatory framework exists specifically for multi-agent AI systems. Freyer, Jayabalan, Kather, and Gilbert (*Nature Medicine*, 2025) argued that frameworks must evolve "beyond static device paradigms" to incorporate adaptive oversight for autonomous agents, proposing a spectrum of AI autonomy with corresponding regulatory intensity, regulatory sandboxes, and continuous real-world performance monitoring [39]. Key unresolved challenges include accountability gaps when multiple models interact, emergent behaviors not predictable from individual components, and the absence of validated methodologies for testing multi-agent clinical workflows [39,40]. Furthermore, 88% of US health systems are using AI internally but 80% lack formal governance [41], and less than 2% of FDA-cleared AI/ML devices are supported by randomized controlled trials [37].

Challenges and the Road Ahead

We acknowledge substantial barriers. Gartner predicts that over 40% of agentic AI projects will be canceled by end of 2027 due to governance failures [26]. Security vulnerabilities in open-source frameworks underscore the need for robust cybersecurity in clinical deployments [42]. The fundamental question of liability when an AI agent makes a clinical recommendation leading to harm remains unresolved. The 70% 10-year DFS benchmark for advanced TNBC is deliberately aspirational, requiring resolution of deep biological challenges including tumor heterogeneity, immune evasion, and the blood-brain barrier.

Yet the history of oncology teaches that audacious goals have a way of becoming reality when the right tools emerge. The 5-year survival for childhood acute lymphoblastic leukemia was less than 10% in the 1960s; today it exceeds 90%. The JAMA Summit Report (Angus, Khera et al., 2025)—featuring over 50 expert authors—has called for real-world testing, responsible data sharing, and aligned incentives [43]. *Nature Medicine*'s editorial warned that medical LLM disclaimers dropped from 26% in 2022 to just 1% in 2025, and that 58% of US physicians worry about AI over-reliance [44]. The path forward demands rigorous validation, transparent governance, and a regulatory framework fit for autonomous, adaptive systems.

Conclusion

The transition from AI tools to AI teammates is not a distant vision; it is unfolding now—in the tumor boards of Stanford, the research labs of Oxford, and the drug discovery pipelines of AI-native biotechs. Three developments converge to make this transformation viable: the technical infrastructure—with LangGraph, CrewAI, MCP, and cloud-native platforms—is production-ready; the evidence base demonstrates multi-agent systems reliably outperform single models by substantial margins; and domain-specific implementations show this is no longer theoretical but already being built for oncology.

The most compelling insight from this evidence is not that AI can match clinical guidelines—that was established by 2024—but that multi-agent architectures with RAG, consensus mechanisms, and safety-check agents can approach the collaborative intelligence of a multidisciplinary tumor board. The gap between this promise and clinical reality lies in regulation, governance, and validation. The question is no longer whether agentic AI will transform oncology, but whether we, as a profession, will be ready to embrace it.

Disclosure of AI Use

The authors used Claude (Anthropic) as an AI writing assistant during the preparation of this editorial. Specifically, AI tools were employed for literature search synthesis, structural organization of evidence, and drafting assistance. All content was critically reviewed, verified, and edited by the authors, who take full responsibility for the accuracy and integrity of the work. All references were individually verified by the authors against their original sources. The final manuscript reflects the authors' independent analysis and scientific judgment.

References

1. Zou J, Topol EJ. The rise of agentic AI teammates in medicine. Lancet. 2025; 405: 457.
2. Truhn D, Azizi S, Zou J, et al. Artificial intelligence agents in cancer research and oncology. Nat Rev Cancer. 2026.
3. Lang Chain. LangGraph Multi-agent orchestration framework. GitHub. 2025. https://github.com/langchain-ai/langgraph.
4. Anthropic. Model Context Protocol MCP. Anthropic. 2024. https://modelcontextprotocol.io
5. Masayoshi Y. EHR-MCP evaluating GPT-4.1 connected to hospital EHR via MCP tools. arXiv. 2025.
6. Lungren M. Developing next-generation cancer care management with multi-agent orchestration. Microsoft Industry Blogs. 2025. https://www.microsoft.com/en-us/industry/blog/healthcare/2025/05/19/developing-next-generation-cancer-care-management-with-multi-agent-orchestration/
7. University of Oxford. Oxford-built multi-agent assistant for cancer care to be piloted in collaboration with Microsoft. University of Oxford News. 2025. https://www.ox.ac.uk/news/2025-12-16-oxford-built-multi-agent-assistant-cancer-care-be-piloted-collaboration-microsoft
8. Afkhami A. How AI agents and tech will transform health care in 2026. Boston Consulting Group. 2026. https://www.bcg.com/publications/2026/how-ai-agents-will-transform-health-care
9. Kim YB, Park C, Jeong H, et al. MDAgents an adaptive collaboration of LLMs for medical decision-making. NeurIPS. arXiv. 2024.
10. Ferber D, El Nahhas OSM, Wölflein G, et al. Development and validation of an autonomous artificial intelligence agent for clinical decision-making in oncology. Nat Cancer. 2025; 6: 1337-1349.
11. Chen M. Multi-agent conversation framework for rare disease diagnosis. npj Digit Med. 2025.
12. Gorenshtein A, Omar M, Glicksberg BS, et al. AI agents in clinical medicine a systematic review. medRxiv. 2025.
13. Huang T. EvoMDT a self-evolving multi-agent system for structured clinical decision-making in multi-cancer. npj Digit Med. 2026.
14. OncoPainBot four-agent system for cancer pain management. npj Digit Med. 2026.
15. Virtual AML Panel multi-agent system for WHO classification and ELN prognostication. Blood/ScienceDirect. 2025.
16. Omar M. Iterative Consensus Ensemble ICE method for LLM consensus. Comput Biol Med. 2025.
17. Smit A. Multi-agent debate strategies on medical datasets. Proceedings of Machine Learning Research ICML. 2024; 235: 45883-45905.

18. RAG chatbot with Cancer Information Service sources achieving 0% hallucination rate. PubMed Central. 2025.
19. Thaker DA. RadOncRAG: evaluation of 15 LLMs with retrieval-augmented generation. JCO Clin Cancer Inform. 2025.
20. Ke Y. GPT-4 with RAG achieving 96.4% accuracy in clinical oncology scenarios. Npj Digit Med. 2025.
21. Wu Y. Multi-agent debate as Bayesian posterior belief updates. arXiv. 2025.
22. Jin Q. TrialGPT matching patients to clinical trials. Nat Commun. 2024.
23. Kurnaz C. Multi-agent platform with OncoAgents for clinical trial matching. J Clin Oncol. 2025; 43: 1554.
24. Cleveland Clinic AI for clinical trial matching Memorial Sloan Kettering deployment of Triomics platform. 2025.
25. Mount Sinai Health System. Mount Sinai launches AI-powered clinical trial-matching platform. Mount Sinai Newsroom. 2026.
26. Gartner Inc. Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027. Gartner Newsroom. 2025.
27. Afkhami A. Future of digital health 2026. Boston Consulting Group Report. 2026.
28. Anderson Cancer Center. Digital twin model shows promise for predicting bone metastasis treatment response. MD Anderson Newsroom. 2025.
29. National Cancer Institute. Theranostics digital twins: a path to personalized medicine. NCI Events. 2026.
30. Artificial intelligence designed drugs hit 90% Phase I success rate in trials. 2 Minute Medicine. 2026.
31. Iambic Therapeutics' AI-discovered cancer drug shows promise in early-stage trial. Reuters. 2025.
32. Chen Z. Bio AI Agent a multi-agent artificial intelligence system for autonomous CAR-T cell therapy development. arXiv. 2025.
33. PharmaMar, Globant. Multi-agent AI collaboration for cancer drug discovery. Press release. 2026.
34. Waks AG, Winer EP. Breast cancer treatment a review. JAMA. 2019; 321: 288-300.
35. Schmid P, Javier Cortes, Rebecca Dent, et al. Overall survival with pembrolizumab in early-stage triple-negative breast cancer. N Engl J Med. 2024; 391: 1981-1991.
36. Bardia A. Sacituzumab govitecan plus pembrolizumab in untreated advanced triple-negative breast cancer. N Engl J Med. 2025; 393: 1529-1540.
37. Singh RP. Taxonomy of FDA-authorized AI-enabled medical devices. npj Digit Med. 2025.
38. World Health Organization. Ethics and governance of artificial intelligence for health WHO guidance. Geneva WHO. 2024.
39. Freyer O, Jayabalan N, Kather JN, et al. Regulatory frameworks for autonomous AI agents in medicine. Nat Med. 2025.
40. Gilbert S. Regulatory sandboxes for controlled testing of agentic AI. npj Digit Med. 2026.
41. Trilliant Health. What differentiates 2026 AI predictions. Trilliant Health Analysis. 2026.
42. OpenClaw ClawBot in healthcare security risks and vulnerabilities. Ventus.ai Blog. 2026.
43. Angus DC, Khera R. JAMA Summit Report on AI in medicine. JAMA. 2025.
44. For trustworthy AI, keep the human in the loop editorial. Nat Med. 2025.
45. Qiu P, Lam C, Li Z, et al. LLM-based agentic systems in medicine. Nat Mach Intell. 2024; 6: 1418-1420.
46. Lee Y, Ferber D, Rood JE, et al. How AI agents will change cancer research and oncology. Nat Cancer. 2024; 5: 1765-1767.
47. Rajpurkar P, Topol EJ. Clinical certification pathway for generalist medical AI. Lancet. 2025; 405: 20.
48. Wu E. Value alignment in agentic AI for medicine. NEJM AI.