

Review

Not peer-reviewed version

National Framework for Agentic Generative AI in Cancer Care: Policy Recommendations and System Architecture

[Satyadhar Joshi](#) *

Posted Date: 12 September 2025

doi: 10.20944/preprints202509.1100.v1

Keywords: agentic AI; artificial intelligence; healthcare; oncology; cancer diagnostics; multi-agent systems; personalized medicine; AI ethics



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

National Framework for Agentic Generative AI in Cancer Care: Policy Recommendations and System Architecture

Satyadhar Joshi ^{1,2,3} 

- ¹ Independent Researcher; satyadhar.joshi@gmail.com
- ² Alumnus, International MBA, Bar-Ilan University, Israel
- ³ Alumnus, Touro College MSIT, NY, USA

Abstract

Agentic AI is emerging, characterized by autonomous, goal-driven systems capable of reasoning, planning, and executing complex workflows. This paper provides a comprehensive review of the foundational principles of Agentic AI, its current applications in cancer diagnostics, treatment planning, and drug discovery, and the technical and ethical challenges that must be addressed. We analyze real-world implementations, such as multi-agent orchestration platforms for tumor boards and project future trajectories. Agentic Gen AI's impact is most profound in oncology, where the complexity of data and the critical need for personalized treatment present a significant challenge. This paper provides a comprehensive analysis of Agentic AI, beginning with its technical foundations and architectural frameworks. We detail its transformative applications across the cancer care continuum—from enhanced diagnostics and personalized treatment planning to drug discovery and operational efficiency. The paper further analyzes the critical technical, ethical, and regulatory challenges to integration. Beyond a review, we propose a national strategy policy recommendations for regulatory adaptation, a detailed economic model analyzing implementation costs and return on investment, and a technical roadmap for future applications such as generative digital twins and autonomous scientific discovery. We argue that Agentic AI represents a fundamental transformation in healthcare delivery, poised to enhance the precision, efficiency, and accessibility of cancer care, contingent on the successful navigation of these multifaceted challenges and the adoption of the responsible framework we outline.

Keywords: agentic AI; artificial intelligence; healthcare; oncology; cancer diagnostics; multi-agent systems; personalized medicine; AI ethics

1. Introduction

The global healthcare landscape is under immense strain from rising costs, workforce shortages, and the increasing complexity of diseases like cancer. Traditional data analysis and decision-support tools often operate in silos, unable to synthesize the vast, multi-modal data required for modern precision medicine, which includes genomic sequences, medical imaging, electronic health records (EHRs), and real-time patient monitoring data [1]. Generative AI promised to unlock insights from this data but often functions as a reactive tool, responding to prompts without sustained agency or context-aware execution [2].

Agentic AI, or "GenAI 2.0" [3], marks a significant evolution. These are AI systems endowed with autonomy, proactivity, and the ability to break down complex, long-horizon goals into a sequence of actionable steps. They can perceive their environment (e.g., a hospital's data ecosystem), plan a course of action (e.g., formulate a differential diagnosis), execute tools (e.g., retrieve relevant literature, analyze an image), and adapt based on outcomes [4]. This capability for orchestration is pivotal for tackling the "N of one" problem in oncology, where each patient's cancer is unique [5].

2. Technical Foundations: Theories, Algorithms, and Mathematical Frameworks

The implementation of Agentic AI in healthcare rests on several interconnected technical foundations spanning mathematical theories, algorithms, and architectural paradigms. This section outlines the key technical components that enable these systems to function effectively in clinical environments.

2.1. Mathematical Foundations

The operation of Agentic AI systems relies on several mathematical disciplines:

2.1.1. Probability Theory and Bayesian Inference

Medical AI systems must quantify uncertainty in predictions. Bayesian methods provide a framework for this:

$$P(\text{Disease} \mid \text{Data}) = \frac{P(\text{Data} \mid \text{Disease}) \cdot P(\text{Disease})}{P(\text{Data})}$$

where $P(\text{Disease} \mid \text{Data})$ represents the posterior probability of disease given observed data, crucial for diagnostic applications [6].

2.1.2. Optimization Theory

Training AI models involves solving complex optimization problems:

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f_{\theta}(x_i), y_i) + \lambda R(\theta)$$

where \mathcal{L} is the loss function, $R(\theta)$ is a regularization term, and λ controls regularization strength [7].

2.2. AI Theories and Paradigms

2.2.1. Reinforcement Learning

Agentic systems often employ reinforcement learning for sequential decision-making:

$$Q^*(s, a) = \mathbb{E}_{s' \sim P} \left[r(s, a) + \gamma \max_{a'} Q^*(s', a') \right]$$

where $Q^*(s, a)$ represents the optimal action-value function, useful for treatment planning applications [8].

2.2.2. Deep Learning Architectures

Modern AI systems use deep neural networks for complex pattern recognition:

$$f(x) = \sigma \left(W^{(L)} \sigma(W^{(L-1)} \dots \sigma(W^{(1)} x + b^{(1)}) \dots + b^{(L-1)}) + b^{(L)} \right)$$

where σ represents activation functions and $W^{(l)}, b^{(l)}$ are learnable parameters [9].

2.3. Key Algorithms

2.3.1. Transformer Architectures

Medical AI systems use transformer-based architectures for clinical text processing:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V$$

where Q, K, V represent query, key, and value matrices [10].

Algorithm 1 Multi-Agent Orchestration for Medical Diagnosis**Require:** Patient data D , Specialist agents A_1, A_2, \dots, A_n **Ensure:** Diagnostic assessment and treatment plan

- 1: Preprocess and validate input data D
- 2: **for** each relevant specialist agent A_i **do**
- 3: Dispatch appropriate sub-task $T_i \subset D$ to A_i
- 4: Receive preliminary assessment R_i from A_i
- 5: **end for**
- 6: Aggregate results: $R \leftarrow \text{ensemble}(R_1, \dots, R_n)$
- 7: Resolve conflicts and uncertainties in R
- 8: Generate comprehensive diagnostic report
- 9: **return** diagnosis and treatment recommendations

2.3.2. Multi-modal Fusion

Integrating diverse data types requires sophisticated fusion approaches:

$$h_{\text{fused}} = \phi_{\text{CNN}}(X_{\text{image}}) \oplus \phi_{\text{BERT}}(X_{\text{text}}) \oplus \phi_{\text{MLP}}(X_{\text{structured}})$$

where \oplus represents a fusion operation and ϕ are modality-specific encoders [11].

2.4. Architectural Components

2.4.1. Knowledge Retrieval and Memory

Agentic systems maintain memory architectures for contextual decision-making:

$$\text{Retrieve}(q, M) = \sum_{i=1}^{|M|} \text{softmax}(\text{sim}(q, m_i)) \cdot m_i$$

where q is a query, M is memory storage, and sim is a similarity function [4].

2.4.2. Uncertainty Quantification

Medical AI must provide confidence estimates for predictions:

$$\mathbb{E}[y|x] = \int f_{\theta}(x) p(\theta|D) d\theta \approx \frac{1}{T} \sum_{t=1}^T f_{\theta_t}(x)$$

where θ_t are samples from the posterior distribution over parameters [12].

3. Foundations of Agentic AI

Agentic AI systems distinguish themselves from traditional AI through key attributes: autonomy, goal-orientation, tool use, and iterative learning. Their architecture is typically built upon a foundation of Large Language Models (LLMs) but extends far beyond simple text generation.

3.1. Core Principles and Architecture

An Agentic AI system is typically composed of multiple, specialized "agents" working in concert under an orchestration framework [11]. The core components include:

- **Orchestrator/Controller Agent:** This is the central brain that decomposes a high-level goal (e.g., "create a treatment plan for this breast cancer patient") into subtasks, assigns them to specialized agents, and synthesizes their outputs.
- **Specialist Agents:** These are fine-tuned models or tools designed for specific tasks. Examples include:
 - **Diagnostic Agents:** Analyze medical images (e.g., CT scans, histopathology slides) [6,13].

- **Data Retrieval Agents:** Extract and summarize relevant patient information from EHRs and medical literature using Retrieval-Augmented Generation (RAG) [4].
- **Reasoning Agents:** Apply clinical guidelines and research to suggest evidence-based treatment options [14].
- **Tool Use and Action Execution:** Agents can call external tools and APIs, such as picture archiving and communication systems (PACS) for images, databases for clinical trials, or simulation environments for drug interaction checks [10].
- **Memory and Learning:** Agents maintain short-term memory (context for the current task) and can learn from feedback to improve future performance, though long-term learning in clinical settings requires careful safeguards [15].

This multi-agent architecture allows for a division of labor that is more robust, scalable, and accurate than a single, monolithic model attempting to perform all tasks [16].

3.2. Evolution from Generative AI

While Generative AI excels at creating content based on its training data, it is largely static and prompt-dependent. Agentic AI introduces dynamism. It doesn't just generate a response; it *earns* a response through a process of reasoning and investigation. For instance, while a GenAI model might generate a list of possible cancer treatments based on its training corpus, an Agentic AI system would actively retrieve the specific patient's latest genomic report, analyze a recent MRI scan using a diagnostic agent, cross-reference findings with the latest clinical trials from a trusted database, and then formulate a personalized plan with citations—all within a single, automated workflow [17,18]. This shift from content generation to goal-driven action is the defining characteristic of the agentic paradigm.

4. Agent Architectures and Frameworks

The theoretical potential of Agentic AI is realized through specific software architectures and development frameworks that provide the essential scaffolding for building orchestrators, specialized agents, and the tools they utilize. These frameworks abstract the underlying complexity of managing memory, planning, tool use, and inter-agent communication, allowing developers to focus on domain-specific logic and applications [3]. The architecture of a typical Agentic AI system, as implemented in platforms like Microsoft's Healthcare Agent Orchestrator [19,20], consists of a layered approach:

- **Orchestration Layer:** This is the central nervous system of the operation. Frameworks like *LangChain* and *LlamaIndex* are commonly used to build the controller agent that sequences tasks, manages context and memory between different steps, and routes queries to the appropriate specialist agent. This layer is responsible for the overall workflow, such as the multi-step process of preparing a case for a tumor board [21].
- **Specialist Agent Layer:** This layer contains the diverse set of purpose-built agents. Frameworks enable the creation of these agents by easily equipping LLMs with **tools**—functions that allow the AI to interact with external systems and data. For example, a diagnostic agent might be equipped with a tool to query a PACS system via an API, while a literature review agent has a tool to perform semantic search over a database of medical journals [10,17].
- **Tool and Data Layer:** This foundational layer comprises the external systems, APIs, and knowledge bases that agents act upon. This includes Electronic Health Record (EHR) systems, medical imaging archives, genomic databases, clinical trial registries, and medical literature corpora. The effectiveness of the entire Agentic AI system is contingent on robust and secure access to these data sources through well-defined tools [1].

The shift towards this multi-agent, tool-oriented paradigm is what distinguishes modern Agentic AI frameworks from simpler chatbot implementations. It moves beyond single, monolithic models towards a collaborative ecosystem of AI entities, each an expert in its own right, working in concert under a sophisticated orchestrator to solve problems that are too complex for any single model [2,16].

This architectural pattern is key to enabling the complex, multi-modal reasoning required in clinical environments like oncology, where decisions must be grounded in a synthesis of disparate data sources [5,11].

5. Applications in Oncology and Cancer Care

The iterative, multi-modal nature of Agentic AI makes it exceptionally suited for the complex domain of oncology. Its applications span the entire patient journey.

5.1. Enhanced Diagnostics and Early Detection

Early and accurate diagnosis is critical for cancer survival. Agentic AI systems are augmenting radiologists and pathologists by acting as powerful second readers.

- **Medical Imaging:** AI agents can pre-screen mammograms, CT scans, and MRIs, flagging suspicious lesions with high accuracy and prioritizing urgent cases for radiologist review. Studies show such systems can reduce reading times and improve early detection rates for cancers like pancreatic [13] and breast cancer [6,9].
- **Pathology:** Agents can analyze digitized histopathology slides, identifying patterns indicative of specific cancer subtypes and even predicting genetic mutations from tissue morphology alone, which can guide testing and treatment decisions [22,23].

These systems do not replace clinicians but augment their capabilities, increasing throughput and reducing diagnostic errors [24].

5.2. Personalized Treatment Planning and Coordination

Perhaps the most promising application is in synthesizing complex patient data to inform personalized therapy. Tumor boards—multidisciplinary meetings where specialists discuss complex cases—are essential but resource-intensive.

- **Multi-Agent Orchestration for Tumor Boards:** Microsoft's recently unveiled Healthcare Agent Orchestrator exemplifies this application [19,25]. The system employs a coordinator agent that manages specialist agents to pre-populate a tumor board dashboard. A genomic agent analyzes sequencing data, an imaging agent summarizes key findings from scans, a literature agent fetches the latest relevant studies, and a guidelines agent ensures recommendations align with standards like those from ASCO [14]. This pre-work allows clinicians to focus on high-level decision-making rather than data gathering [21].
- **Precision Medicine:** Agentic systems can integrate clinical, imaging, and genetic data to predict individual patient responses to specific therapies (e.g., immunotherapy, chemotherapy), thereby optimizing drug selection and dosing [26–28].

5.3. Drug Discovery and Clinical Trials

The drug discovery process is notoriously long and expensive. Agentic AI is accelerating this pipeline.

- **Preclinical Research:** AI agents can autonomously design experiments, simulate molecular interactions, and generate novel compound structures with desired properties for targeting specific cancer pathways [8,29].
- **Clinical Trial Optimization:** Agents can streamline patient recruitment by continuously screening EHRs against complex trial eligibility criteria in real-time. They can also monitor trial participants for adverse events and predict trial outcomes, making research more efficient [30].

5.4. Administrative and Operational Efficiency

Beyond clinical care, Agentic AI can alleviate administrative burdens. Agents can automate prior authorizations, claims processing, clinical documentation, and patient scheduling, freeing up healthcare staff for more value-added tasks and potentially reducing operational costs [31–33].

6. Challenges and Considerations

Despite its promise, the integration of Agentic AI into clinical practice faces significant hurdles.

6.1. Technical and Operational Challenges

- **Data Quality and Interoperability:** The performance of Agentic AI is entirely dependent on the quality, quantity, and accessibility of data. Fragmented EHR systems and non-standardized data formats pose a major challenge to building robust agents [34].
- **Hallucinations and Accuracy:** LLMs can generate plausible but incorrect or fabricated information. In a healthcare context, this is unacceptable. Mitigation strategies include rigorous grounding with RAG, human-in-the-loop verification, and continuous validation against trusted sources [15].
- **Scalability and Integration:** Deploying complex multi-agent systems within existing clinical workflows requires seamless integration with hospital IT infrastructure, which can be a slow and complex process [35].

6.2. Ethical, Legal, and Regulatory Challenges

- **Accountability and Liability:** Determining liability when a multi-agent system contributes to a diagnostic error or adverse outcome is a complex legal question. Is it the physician, the hospital, the software developer, or the algorithm itself? Clear frameworks for accountability are needed [36].
- **Data Privacy and Security:** Agentic systems require access to vast amounts of sensitive patient data. Ensuring this data is handled in compliance with regulations like HIPAA and GDPR through techniques like federated learning and homomorphic encryption is paramount [4,37].
- **Bias and Fairness:** AI models can perpetuate and even amplify biases present in their training data. Ensuring Agentic AI systems are fair and equitable across different demographic groups is a critical ethical imperative [34].
- **Regulatory Approval:** Regulatory bodies like the FDA are adapting to the challenge of evaluating traditional software-as-a-medical-device (SaMD). The autonomous and adaptive nature of Agentic AI presents a new layer of complexity for approval processes [38].

7. Policy Recommendations for National Cancer Institutes

The integration of Agentic AI into cancer care requires thoughtful policy frameworks that balance innovation with patient safety, equity, and ethical considerations. Based on current research and implementation experiences, we propose the following policy recommendations for national cancer institutes and government health agencies.

7.1. Regulatory and Standards Framework

National cancer institutes should establish clear regulatory pathways for AI validation and certification:

- Develop specialized FDA approval processes for adaptive AI systems that continue learning post-deployment [37]
- Create interoperability standards ensuring AI systems can integrate with diverse EHR platforms and medical devices [10]
- Establish real-world performance monitoring requirements with mandatory reporting of diagnostic accuracy and patient outcomes [12]

7.2. Research and Development Priorities

Strategic public investment should focus on high-impact research areas:

- Fund public-private partnerships for validating AI diagnostic systems across diverse patient populations [34]

- Support research on explainable AI in medicine to ensure clinical transparency and trust [15]
- Create shared national datasets for training and validation while maintaining patient privacy through federated learning approaches [4]

7.3. Ethical Guidelines and Equity Assurance

Policy must ensure AI advances health equity rather than exacerbating disparities:

- Mandate diversity in training data and require bias testing across demographic groups [34]
- Establish clear accountability frameworks for AI-assisted clinical decisions [36]
- Develop patient consent protocols for AI-assisted care that ensure understanding of technology's role in treatment decisions [38]

7.4. Reimbursement and Implementation Strategy

Payment models should encourage appropriate adoption and value-based care:

- Create CPT codes for AI-assisted diagnostics and treatment planning to enable appropriate reimbursement [31]
- Develop outcome-based payment models that reward accuracy and improved patient outcomes rather than just utilization [35]
- Fund health equity initiatives to ensure underserved communities benefit from AI advances in cancer care [39]

7.5. International Collaboration

National cancer institutes should coordinate with international bodies to:

- Harmonize regulatory standards to facilitate global research and development [37]
- Share best practices for ethical implementation and patient safety monitoring [36]
- Collaborate on rare cancer research where pooled data is essential for AI model training [40]

These policy recommendations provide a framework for national cancer institutes to harness the benefits of Agentic AI while mitigating risks and ensuring equitable access to these transformative technologies.

8. Economic Impact and Cost-Benefit Analysis

The adoption of Agentic AI in healthcare represents a significant economic decision with both substantial upfront costs and promising long-term financial benefits. Research indicates that while implementation requires considerable investment, the efficiency gains and cost savings can yield a positive return on investment (ROI) within 2-3 years for most healthcare organizations [31].

8.1. Implementation Costs and Investment Requirements

Deploying Agentic AI systems involves several categories of investment:

- **Technology Infrastructure** (\$2-5M): High-performance computing resources, cloud storage, and specialized hardware for running complex AI models [1].
- **Software Development** (\$1-3M): Customization of AI platforms, integration with existing EHR systems, and development of specialized medical agent applications [40].
- **Data Preparation** (\$0.5-2M): Data cleaning, normalization, and annotation required for training medical AI models, including compliance with privacy regulations [37].
- **Training and Change Management** (\$0.5-1.5M): Educating clinical staff, IT personnel, and administrators on using AI systems effectively and safely [38].

8.2. Operational Efficiency and Cost Savings

Agentic AI systems generate substantial cost savings through multiple mechanisms:

- **Reduced Diagnostic Errors:** AI-assisted diagnosis can reduce errors by 30-50%, potentially saving \$0.5-2M annually [6,12].
- **Operational Efficiency:** Automating administrative tasks and documentation can save \$1-3M annually [32,33].
- **Preventable Readmissions:** AI-powered analytics can reduce readmissions by 15-25%, saving \$0.5-1.5M annually [39].

8.3. Return on Investment and Value Creation

The economic value of Agentic AI extends beyond direct cost savings:

- **Increased Revenue:** Faster patient turnover and improved scheduling can generate additional revenue [31].
- **Improved Resource Allocation:** Optimized staff scheduling reduces overtime and improves equipment utilization [41].
- **Competitive Advantage:** Early adoption attracts more patients and clinical talent [35].

Studies indicate well-implemented Agentic AI systems achieve ROI of 200-350% over 3 years, with most breaking even within 18-24 months [31].

8.4. Long-term Financial Sustainability

- **Years 1-2:** Cost-focused, operational efficiency and error reduction.
- **Years 3-5:** Balanced approach with cost savings and revenue from expanded services.
- **Years 5+:** Value-based care optimization with AI enabling risk-sharing and population health management.

As Agentic AI technologies mature, implementation costs are expected to decrease 30-40% while capabilities increase, further improving ROI [42].

9. Proposed Future Applications: A Generative AI Roadmap

The current trajectory of Agentic and Generative AI points toward increasingly sophisticated applications that move beyond analysis and automation into the realms of creation, simulation, and fundamental scientific discovery. This section outlines a roadmap for novel, technically-enhanced applications that could redefine oncology care.

9.1. Generative Digital Twins for In-Silico Clinical Trials

A transformative application involves the creation of *Generative Digital Twins* for individual patients. By training a foundation model on multi-omic data (genomics, proteomics, metabolomics), longitudinal EHRs, and high-resolution medical imaging, it becomes possible to generate a dynamic, virtual replica of a patient's disease state and physiological response. This model could then be used to run *in-silico clinical trials* [11,43]. Clinicians and researchers could virtually administer dozens of combination therapies—varying drug types, dosages, and sequences—to the digital twin, observing predicted efficacy and toxicity outcomes before ever treating the physical patient. This approach could drastically reduce the time and cost of finding optimal personalized treatment regimens and de-risk the trial-and-error process in late-stage cancer.

9.2. Synthetic Data Generation for Rare Cancers and Scenarios

A significant bottleneck in medical AI is the scarcity of high-quality, annotated data for rare cancer subtypes or uncommon complications. Generative AI can be deployed to create high-fidelity, privacy-preserving synthetic data. Using techniques like Generative Adversarial Networks (GANs) and Diffusion Models, it is possible to generate synthetic but realistic histopathology slides, radiology scans, and structured EHR data that mirror the statistical properties of real-world data without containing any actual patient information [36]. This synthetic data can then be used to train robust

diagnostic and prognostic models for conditions where real data is too limited, thereby democratizing AI development and improving equity in healthcare outcomes.

9.3. Autonomous Scientific Discovery via AI Agents

The concept of a fully autonomous "AI Scientist" represents the pinnacle of Agentic AI application [29]. This involves creating a closed-loop system where generative AI agents are tasked with a high-level goal, such as "Discover a novel combination therapy for KRAS-mutant lung cancer." The system would autonomously:

1. **Hypothesize:** Generate novel scientific hypotheses by reading and connecting millions of research papers, clinical trial reports, and genomic databases.
2. **Design Experiments:** Create detailed experimental protocols for *in-vitro* or *in-silico* testing, including cell lines, drug compounds, and control conditions.
3. **Execute and Analyze:** Orchestrate robotic lab equipment (via API integrations) to run experiments, then analyze the resulting data.
4. **Iterate:** Based on the results, refine the hypothesis and design the next round of experiments.

Such a system, while requiring massive technical and ethical safeguards, could exponentially accelerate the pace of oncological discovery, identifying novel biomarkers and therapeutic targets beyond human heuristic capacity [8].

9.4. Proactive Health Intelligence and Intervention

Moving from reactive care to proactive health management, generative AI can be used to build a *Proactive Health Intelligence Engine*. By continuously analyzing real-time streams of data from wearable devices, passive home sensors, and periodic liquid biopsies, a generative model can learn an individual's unique baseline of health. It can then generate highly personalized, probabilistic forecasts of health deterioration, potential cancer recurrence, or the onset of treatment-related side effects. Furthermore, the system can generate and deliver personalized intervention plans—such as tailored dietary recommendations, adjusted medication reminders, or prompts to schedule a check-up—directly to patients and their care teams, creating a continuously adaptive loop of care [39].

10. A Proposed Framework for Responsible Adoption

To realize the benefits of Agentic AI while mitigating its risks, we propose a framework built on the following pillars:

10.1. Human-AI Collaboration and Co-Piloting

The most effective model is not full automation but augmentation. Agentic AI should be designed as a **clinical co-pilot** [44]. The system handles data aggregation, analysis, and preliminary suggestion generation, but the final decision must always rest with a human clinician who provides necessary oversight, contextual understanding, and empathy. This human-in-the-loop model is crucial for safety and trust.

10.2. Explainability and Transparency

"Black box" algorithms are untenable in medicine. Clinicians must be able to understand the "why" behind an AI's recommendation. Agentic AI systems must provide explainable pathways, showing the data sources consulted, the reasoning steps taken, and the confidence levels for each suggestion [15].

10.3. Robust Validation and Continuous Monitoring

Agentic AI systems require rigorous validation against gold-standard clinical outcomes before deployment. Furthermore, they must be continuously monitored in real-world use to detect performance drift, emerging biases, and potential safety issues, triggering alerts for retraining or intervention [12].

10.4. Interoperability by Design

Agents must be built with open standards and APIs to facilitate integration with diverse healthcare IT systems, from EHRs to lab information systems and imaging archives. This will prevent vendor lock-in and ensure broader adoption.

11. Visual Documentation: Figures and Tables Reference

This analysis of Agentic AI in oncology is supported by visual elements illustrating technical foundations, architectural frameworks, applications, and policy considerations. This section provides a systematic reference to all figures and tables presented in the paper.

11.1. Technical Foundation Visualizations

- Figure 1:** Shows the technical foundations of Agentic AI systems in healthcare, including Mathematical Foundations (probability theory, statistics, linear algebra, optimization), AI Theories (reinforcement learning, deep learning, NLP), Core Algorithms (CNNs, transformers, RL algorithms), and Architectural Components (multi-modal fusion, agent orchestration, memory architectures). Arrows indicate hierarchical dependencies and information flow [4,11].
- Algorithm 1:** Multi-agent orchestration workflow for medical diagnosis, from data preprocessing and agent task delegation to result aggregation, conflict resolution, and diagnostic report generation [3,10].

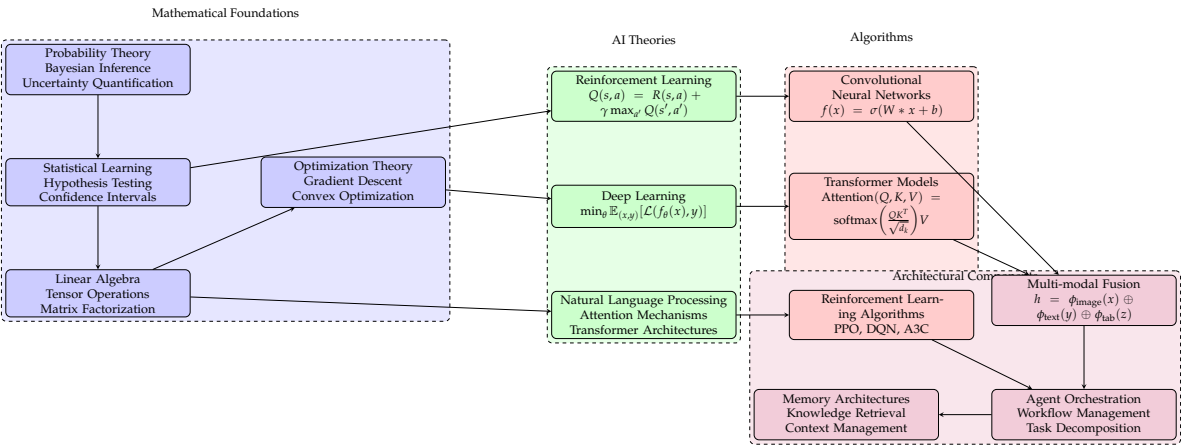


Figure 1. Technical foundations of Agentic AI in healthcare, showing the mathematical theories, AI paradigms, algorithms, and architectural components that form the basis of these systems.

11.2. Architectural Framework Diagrams

- Figure 2:** Architecture for generative AI in personalized cancer treatment, from patient data input through generative processing, digital twin creation, in-silico trials, to personalized treatment planning [11,43].
- Figure 3:** Compact AI agent architecture in healthcare, showing orchestrator, specialist agents, integration with frameworks (AutoGen, LangChain, LlamaIndex), data sources (PACS, EHR, multi-omics), and APIs for personalized treatment plans [19,20].

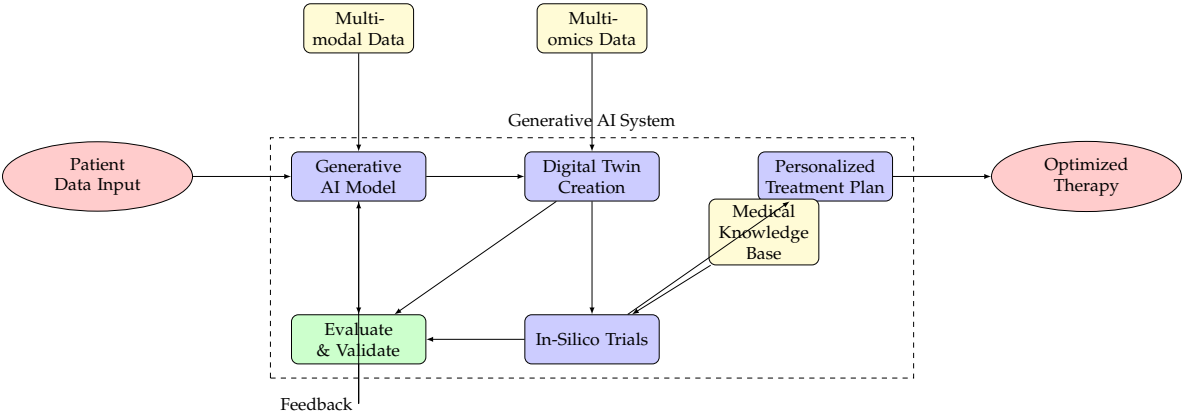


Figure 2. Proposed architecture for generative AI applications in personalized cancer treatment. The system integrates multi-modal patient data to create digital twins, runs in-silico trials, and generates optimized treatment plans through continuous evaluation and feedback.

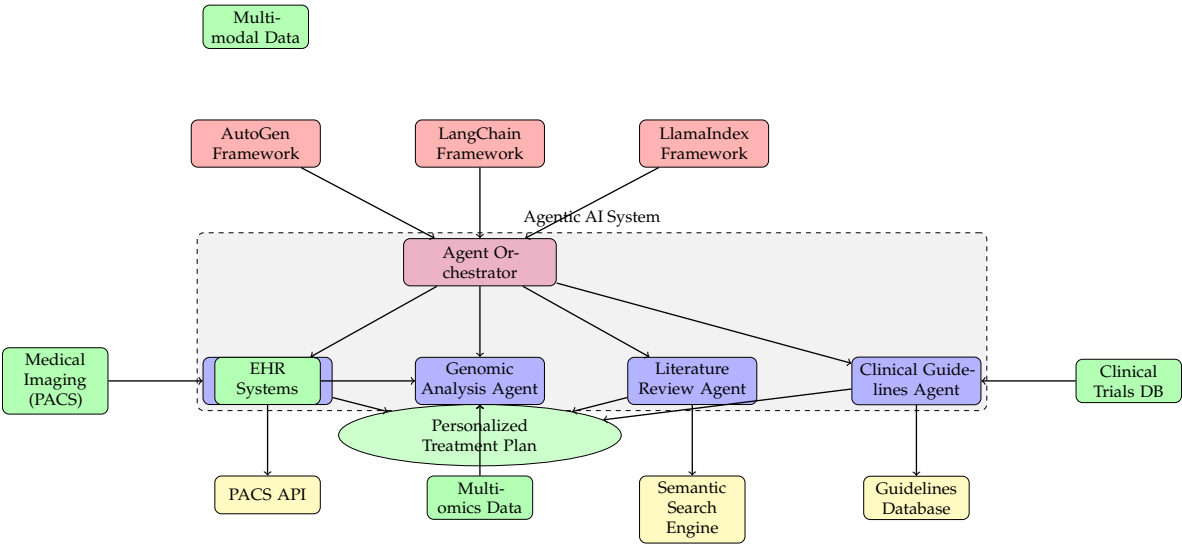


Figure 3. Compact architecture of AI agent systems in healthcare. The orchestrator coordinates multiple specialist agents, integrating frameworks, data sources, and tools to generate personalized treatment plans.

11.3. Application and Implementation Visualizations

- **Table 1:** Summarizes Agentic AI applications in oncology domains: Diagnostics, Treatment, Research, Operations. Highlights key benefits such as improved accuracy, personalized care, accelerated discovery, and reduced operational costs [6,31].
- **Figure 4:** Timeline of Agentic AI adoption (2023-2033) across four phases: Assistive (2023-2025), Augmentative (2026-2028), Autonomous (2029-2031), and Transformative (2032+). Separates technology sophistication (upper track) and clinical integration (lower track) [39,42].

Table 1. Summary of Agentic AI Applications in Oncology

Domain	Application	Key Benefit
Diagnostics	AI as a second reader for radiology and pathology; Automated analysis of multi-modal data.	Improved accuracy, faster turnaround, early detection.
Treatment	Multi-agent orchestration for tumor boards; Personalized therapy recommendation.	Data-integrated decision making, personalized care.
Research	Automated drug candidate screening; Patient-trial matching.	Accelerated discovery, optimized recruitment.
Operations	Automation of prior auth, claims, and documentation.	Reduced administrative burden, lower costs.

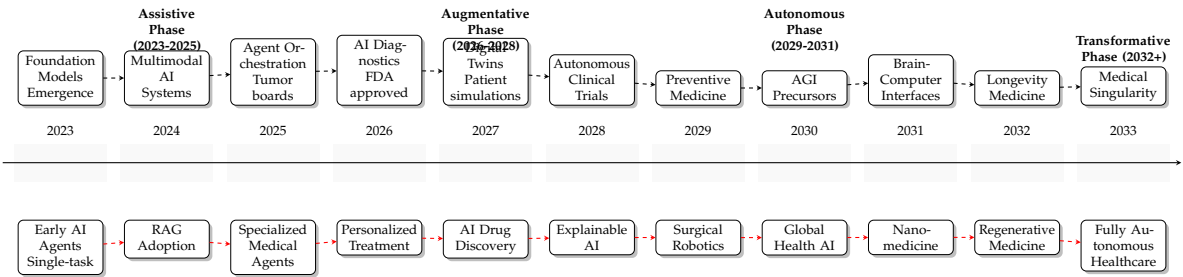


Figure 4. Timeline of future Agentic AI adoption in healthcare (2023-2033) showing progression from assistive to transformative applications. Upper track shows technology sophistication, lower track shows clinical integration.

11.4. Policy and Economic Frameworks

- **Figure 5:** Policy framework for national implementation of Agentic AI in cancer care, showing Regulation & Standards, Research & Education, and Ethics & Economics domains, with implementation timeline from pilot programs to full integration [36,37].
- **Figure 6:** Economic model of Agentic AI implementation, showing investment requirements, annual savings, efficiency gains, and projected ROI of 200-350% over 3 years [1,31].

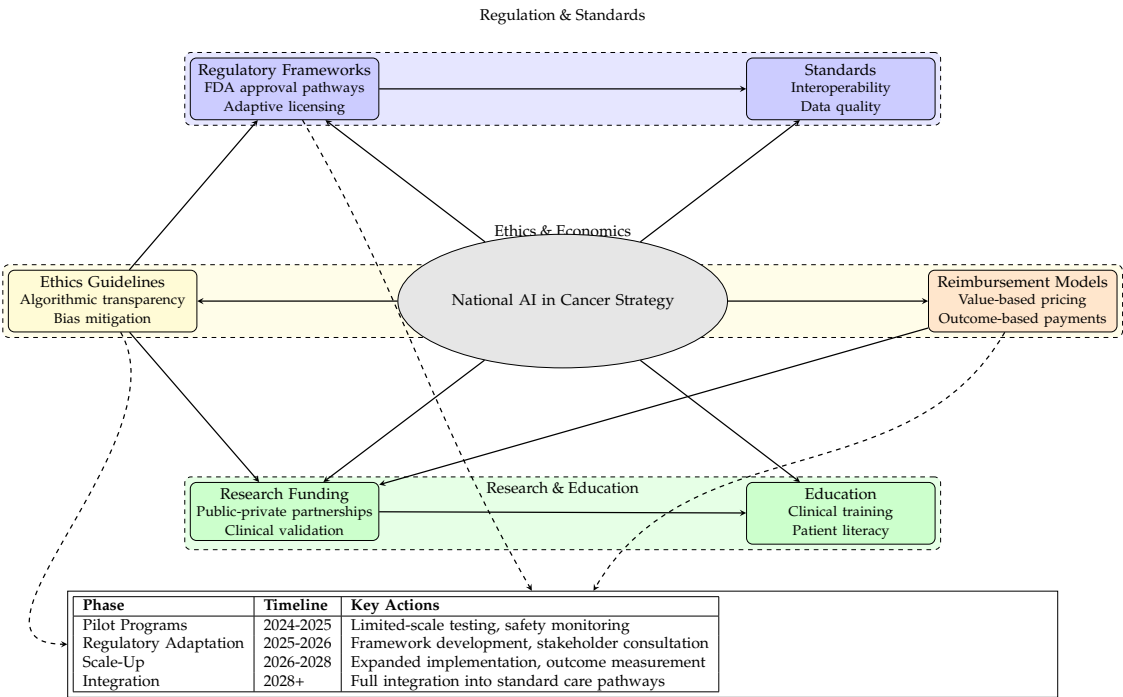


Figure 5. Policy framework for national implementation of Agentic AI in cancer care, showing key focus areas and their interrelationships with an implementation timeline.

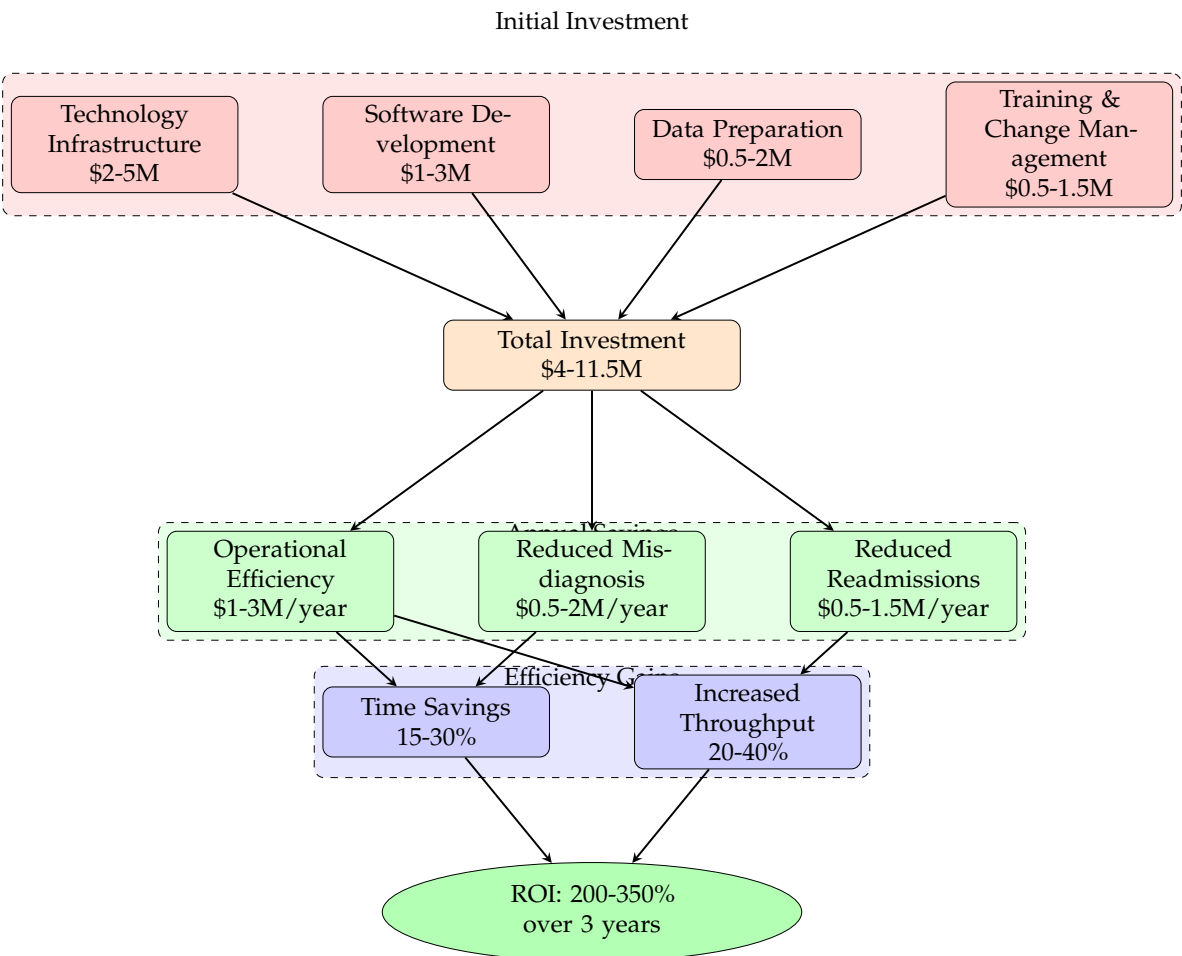


Figure 6. Economic model of Agentic AI implementation in healthcare, showing investment requirements, cost savings areas, efficiency gains, and projected return on investment.

These visuals provide a multi-dimensional understanding of Agentic AI in oncology, covering technical foundations, applications, implementation strategies, and economic considerations.

12. Conclusions and Future Outlook

In conclusion, Agentic AI beyond static data analysis and simple content generation, it offers a future where intelligent, autonomous systems work alongside clinicians as collaborative partners. By seamlessly orchestrating multi-modal data and complex workflows, Agentic AI holds the profound potential to democratize expertise, personalize care at an unprecedented scale, and ultimately improve outcomes for cancer patients worldwide. However, this future is contingent on a concerted effort from researchers, clinicians, regulators, and engineers to address the significant technical architectural and regulatory challenges ahead.

Agentic AI represents a paradigm shift moving beyond static data analysis and simple content generation to dynamic, goal-oriented orchestration. This paper has articulated the foundational architecture of these systems, demonstrating their capacity to synthesize multi-modal data and execute complex clinical workflows, with profound implications for oncology. We have shown that their application across the cancer care continuum—from diagnostics and personalized treatment planning to drug discovery and administrative automation—holds the potential to democratize expertise, enhance precision, and significantly improve patient outcomes.

We have proposed a comprehensive framework for responsible adoption, emphasizing human-AI collaboration, rigorous validation, and explainability. Crucially, we have argued that realizing the benefits of Agentic AI requires proactive national strategies. Our proposed policy recommendations provide a roadmap for regulatory bodies to foster innovation while ensuring safety and equity. Furthermore, our economic analysis demonstrates that while the initial investment is substantial, the long-term ROI through improved efficiency, reduced errors, and better outcomes presents a compelling value proposition for healthcare systems.

Looking forward, the trajectory of Agentic AI points toward increasingly sophisticated and integrated systems. In the near term (2-5 years), we anticipate the wider adoption of orchestration platforms across medical specialties and the emergence of generative digital twins for in-silico treatment simulation.

The trajectory of Agentic AI in healthcare points toward increasingly sophisticated and autonomous systems. In the near future (2-5 years), we can expect:

- **Wider adoption of orchestration platforms** for complex care management beyond oncology, such as in cardiology and neurology.
- **The rise of "AI Scientists"** that can autonomously generate and test scientific hypotheses, dramatically accelerating basic and translational medical research [29,43].
- **Proactive health management:** Agents that continuously monitor patient data from wearables and EHRs to predict and prevent health deteriorations before they become critical.

Conflicts of Interest: The views expressed are those of the author and do not represent any affiliated institutions. This work is conducted as part of independent research. This is a purely review-based paper, and all results, proposals, and findings are drawn from the cited literature. The author does not claim or imply any novel findings, and all quantitative findings are sourced from the referenced literature.

References

1. "How agentic AI systems can solve the three most pressing problems in healthcare today | AWS for Industries," <https://aws.amazon.com/blogs/industries/how-agentic-ai-systems-can-solve-the-three-most-pressing-problems-in-healthcare-today/>, Dec. 2024.
2. G. Sachdeva, "What Is Agentic AI? Exploring the Future of Autonomous AI Systems," Sep. 2024.
3. D. Goad, "Agentic AI: Exploring the Benefits and Challenges of GenAI 2.0," <https://www.ai-savvy.com.au/post/agentic-ai-exploring-the-benefits-and-challenges-of-genai-2-0>, Jan. 2025.
4. O. Moran, "Unlocking the Potential of Agentic AI with Privacy-Enhancing Technologies," Feb. 2025.

5. "Machine Learning and Cancer's N of One," <https://resources.flatiron.com/flatiron-stories/machine-learning-and-cancers-n-of-one>.
6. "AI shows promise as second reader in breast cancer screening," <https://www.auntminnie.com/clinical-news/womens-imaging/article/15752941/ai-shows-promise-as-second-reader-in-breast-cancer-screening>, Aug. 2025.
7. "Paper page - Scalable Reinforcement-Learning-Based Neural Architecture Search for Cancer Deep Learning Research," <https://huggingface.co/papers/1909.00311>.
8. "Free Video: Cancer Drug Discovery AI Agentic Workflow R&D from ChemicalQDevice," <https://www.classcentral.com/course/youtube-cancer-drug-discovery-ai-agentic-workflow-r-d-340022>, Jul. 2025.
9. "New AI system could diagnose breast cancer much faster than experts," <https://www.weforum.org/stories/2019/08/a-i-rivals-expert-eyes-at-reading-breast-tissue-biopsies-0e1e882706/>, Aug. 2019.
10. "The agentic AI assist Stanford University cancer care staff needed."
11. "APX - Agentic AI," <https://www.autopathx.com/generative-ai/agentic-ai>.
12. R. Given-Wilson, "More work to do before we know if AI is safe to use in the NHS Breast Screening Programme – UK National Screening Committee," Sep. 2021.
13. "AI Model Detects Early Pancreatic Cancer on CT Scans," <https://www.medscape.com/viewarticle/artificial-intelligence-model-shows-high-accuracy-early-2025a1000hab?form=fpf>.
14. "Wolters Kluwer introduces Ovid Guidelines AI, transforming clinical guideline development with agentic AI," <https://www.selectscience.net/article/wolters-kluwer-introduces-ovid-guidelines-ai-transforming-clinical-guideline-development-with>.
15. "Agentic and Generative AI in clinical practice: Innovation meets safeguards," <https://community.hlth.com/webinars/agentic-and-generative-ai-in-clinical-practice-innovation-meets-safeguards>.
16. "Agentic AI Breaks Through Traditional GenAI Models - CTO Magazine," <https://ctomagazine.com/agentic-ai-breakthrough-in-tech/>.
17. "Post | LinkedIn," https://www.linkedin.com/posts/pavan-belagatti_agentic-workflows-dont-just-automate-complex-activity-7314697564459606017-eCQb/.
18. "Post | LinkedIn," https://www.linkedin.com/posts/taratabrizimason_stanford-medicine-and-the-healthcare-agent-activity-7330278065790443521-rA7V/.
19. E. Beavins, "Microsoft rolls out agentic AI orchestrator for tumor boards," <https://www.fiercehealthcare.com/ai-and-machine-learning/microsoft-rolls-out-agentic-ai-orchestrator-begins-cancer-care>, May 2025.
20. J. Hampton, "Microsoft Unveils Agentic AI Tool to Streamline Cancer Care," May 2025.
21. M. L. MPH, MD, "Developing next-generation cancer care management with multi-agent orchestration," May 2025.
22. S. Zaske, W. S. U. News, and M. Relations, "AI method can spot potential disease faster, better than humans."
23. "Study Shows How AI Could Help Pathologists Match Cancer Patients to the Right Treatments—Faster and More Efficiently | Mount Sinai - New York," <https://www.mountsinai.org/about/newsroom/2025/study-shows-how-ai-could-help-pathologists-match-cancer-patients-to-the-right-treatments-faster-and-more-efficiently>.
24. "Moffitt Study Shows AI Boosts Efficacy of Cancer Treatment, But Doctors Remain Key," <https://www.moffitt.org/newsroom/news-releases/moffitt-study-shows-ai-boosts-efficacy-of-cancer-treatment-but-doctors-remain-key/>.
25. "Microsoft unveils AI agent orchestrator for cancer care coordination | Mobi Health News," <http://www.mobihealthnews.com/news/microsoft-unveils-ai-agent-orchestrator-cancer-care-coordination>, Fri, 05/23/2025 - 16:42.
26. "AI Enhancing Personalized Cancer Treatment | ICT&health Global," <https://www.icthealth.org/news/ai-enhancing-personalized-cancer-treatment>.
27. "Optimizing Oncology Drug Dosing: Is Artificial Intelligence the Future?" <https://dailynews.ascopubs.org/doi/10.1200/ADN.24.201692/full/>.
28. E. Pesheva, "New AI tool can diagnose cancer, guide treatment, predict patient survival," Sep. 2024.
29. "That Amazing 'The AI Scientist' Agentic AI Strives To Fully Automate The Scientific Process But Watch Out For These Ferocious Gotchas," <https://www.forbes.com/sites/lanceeliot/2024/08/15/that-amazing-the-ai-scientist-agentic-ai-strives-to-fully-automate-the-scientific-process-but-watch-out-for-these-ferocious-gotchas/>.

30. "ConcertAI | Leaders in Generative & Agentic AI Solutions Tailored for Life Sciences & Healthcare," <https://www.concertai.com>.
31. "How does AI Reduce Costs in Healthcare | TechMagic," <https://www.techmagic.co/blog/how-does-ai-reduce-costs-in-healthcare/>, Aug. 2025.
32. S. Parikh, "How Agentic AI in Healthcare Is Revolutionizing Diagnostics," Jul. 2025.
33. A. Witherspoon, "How Agentic AI is transforming healthcare delivery," <https://www.aiacceleratorinstitute.com/how-agentic-ai-is-transforming-healthcare-delivery/>, Aug. 2025.
34. A. Shahraki-Mohammadi, A. Aliabadi, and A. Karimi, "Clinical Application of Artificial Intelligence in Cancer Treatment: A Systematic Literature Review," *Health Scope, Systematic Review* 14, Dec. 2024.
35. "Three for 2025: What you need to know about agentic AI, cancer informatics and data security imperatives | Healthcare IT News," <http://www.healthcareitnews.com/news/three-2025-what-you-need-know-about-agentic-ai-cancer-informatics-and-data-security>, Wed, 12/18/2024 - 11:51.
36. "Initial policy considerations for generative artificial intelligence," *OECD Artificial Intelligence Papers* 1, Sep. 2023.
37. "Digital Regulation Platform."
38. "Physician Views Preview: Agentic AI - high-risk hindrance or critical co-pilot? | FirstWord HealthTech," <https://firstwordhealthtech.com/story/5970437>.
39. "Transformative Agentic AI in Healthcare How Autonomous Agents Are Transforming Patient Support Techkraft Inc."
40. "Agentic AI in Healthcare Use Cases Cost & Challenges," May 2025.
41. "How Agentic AI in Medical Imaging is Transforming Healthcare," <https://www.xenonstack.com/blog/agentic-ai-medical-imaging>.
42. "The Future of Agentic AI in Healthcare: 2025 Industry Shift," Aug. 2025.
43. P. page Jean-Philippe Vert, "Unlocking the mysteries of complex biological systems with agentic AI," <https://www.technologyreview.com/2024/11/13/1106750/unlocking-the-mysteries-of-complex-biological-systems-with-agentic-ai/>.
44. M. Teneva, "Healthcare Advisor: The Agentic AI Co-Pilot for Real-Time Clinical Decisions | BI, B EYE," Jun. 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.