Generative Health Consulting, LLC

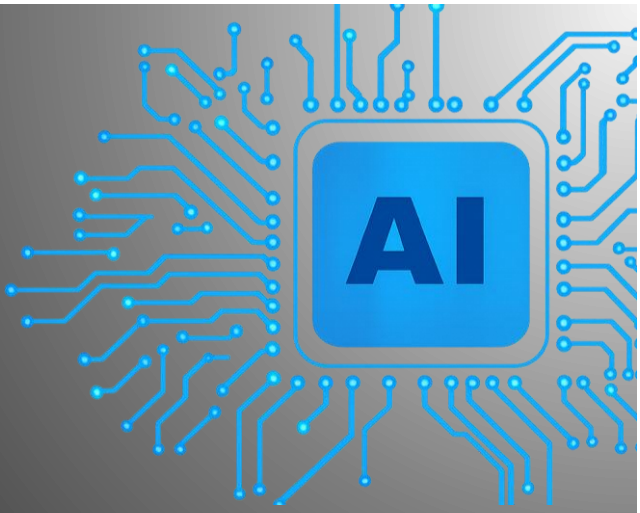

Assess | Enable | Transform

www.genhealthconsult.ai

GHC Insights

August 4th, 2025

RAG vs Generative AI vs Agentic AI: A Guide to Best Practices in Healthcare

The pharmaceutical industry stands at a critical juncture where artificial intelligence technologies are fundamentally reshaping drug discovery, regulatory compliance, and patient care. As pharmaceutical, biotech, biopharma, and life science companies navigate this transformation, three distinct AI approaches have emerged as game-changers: Retrieval-Augmented Generation (RAG), Generative AI, and Agentic AI. Understanding how these technologies differ and when to deploy each is crucial for regulated healthcare companies seeking to assess, enable, and transform their AI capabilities.

The AI Landscape in Pharmaceutical Companies

The pharmaceutical industry's adoption of AI technologies has accelerated dramatically, with the global market for AI in drug discovery estimated at $1.5 billion in 2023 and projected to reach over $15 billion by 2032, growing at a CAGR of almost 30 percent[1]. This explosive growth reflects the industry's recognition that AI can address longstanding challenges in drug development, regulatory compliance, and operational efficiency.

Pharmaceutical companies face unique pressures that make AI adoption particularly attractive. Mountains of clinical data, ever-changing regulations, and the need to innovate while maintaining strict compliance create an environment where traditional approaches often fall short[2]. The industry spends up to 25% of

a medium or large pharmaceutical manufacturing site's operational budget on compliance efforts alone[3], highlighting the need for more efficient solutions.

Understanding the Three AI Paradigms

Retrieval-Augmented Generation (RAG): The Precision Tool

RAG represents a sophisticated approach that combines the power of large language models with the ability to retrieve information from external knowledge sources[4]. In pharmaceutical applications, RAG systems excel at providing accurate, up-to-date information by grounding responses in current regulatory guidelines, clinical trial data, and scientific literature.

Aspect	RAG	Generative AI	Agentic AI
Core Function	Info retrieval + generation	Creates new content	Autonomous task execution
Autonomy Level	Semi-auto needs queries	Reactive needs prompts	Highly auto independent
Pharma Uses	Compliance data retrieval	Protocol content gen	Screening optimization
Knowledge Src	External DB + LLM	Pre-trained model	Multiple systems
Complexity	Medium vector DBs	Low existing models	High integration
Key Advantage	Current data accuracy	Creative generation	End-to-end completion

Comparison of RAG, Generative AI, and Agentic AI in Pharmaceutical Applications

The key advantage of RAG in pharmaceutical settings lies in its ability to address the critical limitation of traditional AI models: their knowledge cutoff date. While standard LLMs can only reference information from their training data, RAG systems can access the latest FDA guidelines, recent clinical trial results, and updated safety protocols[5]. This capability is essential in an industry where regulations change frequently and accuracy is paramount.

Generative AI: The Creative Engine

Generative AI focuses on creating new content based on training data and user prompts[2]. In pharmaceutical contexts, these systems draft clinical trial protocols, generate regulatory submissions, and create patient education materials. However, generative AI operates reactively—it requires human input for each task and cannot act independently[6].

The pharmaceutical industry has embraced generative AI for tasks that benefit from creative output while maintaining human oversight. Companies use it to generate marketing materials, draft initial research

proposals, and create synthetic data for testing scenarios[7]. However, its reactive nature means it serves primarily as a sophisticated assistant rather than an autonomous system.

Agentic AI: The Autonomous Executor

Agentic AI represents the most advanced paradigm, capable of autonomous decision-making and task execution with minimal human intervention[2]. These systems don't just generate content or retrieve information, they actively monitor processes, make decisions, and execute complex workflows independently.

In pharmaceutical applications, agentic AI systems can autonomously screen thousands of compounds, adjust clinical trial parameters in real-time, and optimize manufacturing processes without waiting for human input[6]. This level of autonomy makes agentic AI particularly valuable for repetitive, data-intensive tasks that require consistent execution and rapid adaptation to changing conditions.

How Pharmaceutical Companies View These Technologies

Current Adoption Patterns

Pharmaceutical companies are taking a measured approach to AI adoption, with different technologies finding their niches based on specific use cases and regulatory requirements. According to recent industry surveys, 44% of compliance officers in pharmaceutical companies are already leveraging AI capabilities, though implementations remain limited to simple use cases targeting operational efficiency[8].

The industry's perspective on these technologies reflects a careful balance between innovation and risk management:

RAG Systems are viewed as the most trustworthy option for regulatory compliance and clinical decision support. Companies appreciate RAG's ability to provide transparent, citation-backed responses that can withstand regulatory scrutiny[9]. Major pharmaceutical companies are implementing RAG for regulatory document search, clinical trial data retrieval, and drug interaction checking.

Generative AI is seen as a valuable tool for content creation and ideation, but with clear boundaries. Companies use it for drafting documents, generating hypotheses, and creating training materials, always with human review and validation[10]. The technology is viewed as augmenting human capabilities rather than replacing them.

Agentic AI is regarded with both excitement and caution. While companies recognize its potential to transform operations, concerns about autonomy in regulated environments mean adoption is currently

focused on well-defined, lower-risk applications like laboratory automation and supply chain optimization[11].

Industry Challenges and Considerations

Pharmaceutical companies face unique challenges when implementing these AI technologies:

1. **Regulatory Compliance**: Every AI-driven decision must meet FDA and EMA requirements, creating additional complexity for implementation[12]
2. **Data Security**: Protecting proprietary research and patient data requires robust security measures[9]
3. **Validation Requirements**: AI systems must undergo rigorous testing and validation before deployment[8]
4. **Integration Complexity**: Legacy systems and siloed data create technical barriers[2]

Best Practices for RAG Implementation

Architecture and Technical Considerations

RAG Implementation Workflow in Pharmaceutical Companies

Successful RAG implementation in pharmaceutical settings requires careful attention to architecture and design. Based on industry experience and research, here are the key best practices:

1. Chunking Strategies

The most effective approach for pharmaceutical documents involves using larger chunk sizes than typical implementations. Research shows that using 10,000 character chunks with 2,000 character overlap

provides optimal results for complex regulatory documents[3]. This approach ensures that context is preserved while maintaining efficient retrieval.

2. Embedding Models

Domain-specific embedding models significantly outperform general-purpose models. Pharmaceutical companies should use specialized models like:

- SciBERT for scientific literature[13]
- BioGPT for biomedical text[13]
- LLM-Embedder for capturing complex semantic relationships[3]

3. Retrieval Enhancement

The QA-RAG approach, which uses both queries and hypothetical answers from fine-tuned models, demonstrates superior performance in pharmaceutical applications. This dual-track retrieval method achieves context precision scores of 0.717 compared to 0.556 for conventional RAG[3].

4. Document Processing

- Use OCR technology specifically designed for scientific texts (e.g., Nougat transformer)[3]
- Implement semantic indexing rather than keyword-based organization[14]
- Maintain comprehensive metadata for audit trails and compliance[15]

Implementation Framework

Phase 1: Data Preparation

- Clean and structure existing documentation
- Implement version control for regulatory documents
- Create comprehensive metadata schemas

Phase 2: System Development

- Deploy vector databases optimized for pharmaceutical data
- Implement reranking algorithms for improved precision
- Integrate with existing systems through APIs

Phase 3: Validation and Testing

- Conduct thorough accuracy assessments using domain experts
- Implement continuous monitoring and feedback loops
- Ensure compliance with regulatory requirements

Phase 4: Deployment and Scaling

- Start with pilot programs in low-risk areas
- Gradually expand based on performance metrics
- Maintain human-in-the-loop validation

Best Practices for Generative AI and Agentic AI

Generative AI Best Practices

1. Content Generation Guidelines

- Always maintain human review processes for generated content
- Implement guardrails to prevent hallucinations and ensure accuracy
- Use fine-tuned models trained on pharmaceutical-specific data

2. Use Case Selection

- Focus on applications where creativity adds value (marketing materials, hypothesis generation)
- Avoid using for critical safety or regulatory decisions without human oversight
- Implement version control and approval workflows

3. Quality Assurance

- Establish clear evaluation metrics for generated content
- Regular audits of output quality and compliance
- Continuous training updates based on feedback

Agentic AI Best Practices

1. Autonomous System Design

- Define clear boundaries for autonomous decision-making
- Implement fail-safe mechanisms and human override capabilities
- Ensure complete audit trails for all autonomous actions

2. Implementation Strategy

- Start with well-defined, repetitive tasks
- Gradually increase autonomy based on performance metrics
- Maintain regulatory compliance throughout

3. Risk Management

- Conduct thorough risk assessments before deployment
- Implement continuous monitoring systems
- Establish clear escalation procedures

Comparative Analysis and Strategic Recommendations

When to Use Each Technology

RAG is optimal for:

- Regulatory compliance queries requiring current information
- Clinical trial data retrieval and analysis
- Drug interaction checks and safety assessments
- Any application requiring transparent, citation-backed responses

Generative AI excels at:

- Creating first drafts of protocols and reports
- Generating patient education materials
- Developing marketing content
- Hypothesis generation for research

Agentic AI is best for:

- High-throughput compound screening
- Automated laboratory workflows
- Supply chain optimization
- Real-time clinical trial monitoring

Integration Strategies

Successful pharmaceutical companies are adopting a hybrid approach that leverages the strengths of each technology:

1. **Complementary Deployment**: Use RAG for information retrieval, Generative AI for content creation, and Agentic AI for process automation
2. **Layered Architecture**: Implement RAG as the knowledge foundation, with Generative AI for user interaction and Agentic AI for workflow orchestration
3. **Phased Implementation**: Start with RAG for compliance, add Generative AI for efficiency, then introduce Agentic AI for transformation

Future Outlook and Recommendations

The pharmaceutical industry's AI journey is just beginning. McKinsey estimates that generative AI alone could add $60-110 billion annually to the pharmaceutical and medical-product industries[7]. However, realizing this potential requires thoughtful implementation that balances innovation with regulatory compliance.

Key Recommendations for Pharmaceutical Companies:

1. **Start with RAG**: Its transparency and accuracy make it ideal for building trust in AI systems
2. **Experiment with Generative AI**: Use it for non-critical applications to build expertise
3. **Prepare for Agentic AI**: Develop governance frameworks and technical infrastructure
4. **Invest in Training**: Build internal capabilities through comprehensive education programs
5. **Collaborate with Regulators**: Work proactively with FDA and EMA to shape AI guidelines

Critical Success Factors:

- **Data Quality**: Clean, well-structured data is essential for all AI implementations
- **Human Oversight**: Maintain appropriate human involvement, especially for critical decisions
- **Continuous Improvement**: Implement feedback loops and regular model updates
- **Regulatory Alignment**: Ensure all AI systems meet current and anticipated regulations
- **Cultural Change**: Foster an AI-positive culture while maintaining scientific rigor

Conclusion

The pharmaceutical industry's adoption of RAG, Generative AI, and Agentic AI represents a fundamental shift in how companies approach drug discovery, regulatory compliance, and patient care. Each technology offers unique advantages, and successful implementation requires understanding their distinct capabilities and limitations.

RAG provides the accuracy and transparency essential for regulatory compliance, making it the foundation for AI adoption in pharmaceutical settings. Generative AI enhances creativity and efficiency in content creation, while Agentic AI promises to revolutionize operational workflows through autonomous execution.

As the industry continues to evolve, companies that thoughtfully implement these technologies—starting with robust RAG systems, experimenting with Generative AI, and preparing for Agentic AI—will be best positioned to deliver innovative therapies while maintaining the highest standards of safety and compliance. The key is not choosing one technology over another, but rather understanding how to orchestrate all three to create a comprehensive AI strategy that transforms pharmaceutical operations while upholding the industry's commitment to patient safety and regulatory excellence.

1. https://www.biostrand.ai/retrieval-augmented-generation-in-drug-discovery
2. https://www.salesforce.com/healthcare-life-sciences/life-sciences-artificial-intelligence/agentic-ai-in-pharma/
3. https://arxiv.org/html/2402.01717v1
4. https://a21.ai/how-new-generative-ai-tech-rag-can-increase-pharma-revenue/
5. https://www.maargasystems.com/2024/09/09/rag-in-non-hospital-healthcare-settings/
6. https://www.salesforce.com/healthcare-life-sciences/life-sciences-artificial-intelligence/ai-agents-in-pharma/
7. https://www.edvantis.com/blog/generative-ai-in-pharma/
8. https://www.iqvia.com/locations/emea/blogs/2025/02/how-ai-is-shaking-up-compliance-in-the-life-sciences-industry
9. https://talbotwest.com/services/retrieval-augmented-generation/rag-in-the-pharmaceutical-industry
10. https://medcitynews.com/2025/07/from-hype-to-impact-how-pharma-can-use-agentic-ai-to-drive-efficiency-and-trust/
11. https://www.consultancy.eu/news/11963/agentic-ai-transforming-the-pharma-lifecyle-from-rd-through-to-commercialization
12. https://www.cognizant.com/uk/en/insights/blog/articles/medtech-gets-ready-for-an-agentic-ai-future
13. https://intuitionlabs.ai/articles/rag-performance-pharmaceutical-documents
14. https://www.leewayhertz.com/advanced-rag/

15. https://community.databricks.com/t5/technical-blog/the-ultimate-guide-to-chunking-strategies-for-rag-applications/ba-p/113089