

How to Spot Where Agentic AI Works Best

The latest evolution of generative AI technology

CHARLES RINK, Senior Principal, Information Management & Analytics Technology, IQVIA,
WILLIAM O'REILLY, Director, Product Offering Development, IQVIA

Agentic AI is the latest evolution of generative AI technology to automate tasks that were not possible with older technologies. Large Language Models can now execute steps that previously derailed workflow automation platforms and robotic process automation software, unlocking the potential for end-to-end process management. Software providers are adding agentic AI capabilities to most enterprise software platforms with Microsoft Copilot and Google Gemini as the most visible examples.

Despite huge investments and a near-constant stream of announcements of breakthrough capabilities, real business impact with AI agents can still be a challenge. A recent study found that only a quarter of generative AI and agentic AI proofs of concept make it to production, and half of those struggle with adoption.¹ Business leaders are struggling to prioritize agentic AI efforts where they'll have success and impact.

In an environment saturated with articles praising the transformative potential of AI agents, we use this space to examine where they work well and where they don't, and we give practical insights and tips for successful programs.

Prioritizing agentic AI use cases

Creating the best chances of success requires a structured evaluation of potential use cases. Rating each project for impact and effort should be complemented by a risk assessment.

Business impact in life sciences requires one or more of revenue generation, cost savings, speed for decision-making, and/or meeting regulatory milestones. Because AI agents are an automation tool, cost saving is often the first place to look and the easiest to measure. Speed for decision-making and



regulatory milestones are a crucial value driver for pipeline products. And for commercial organizations, revenue generation almost always comes from improvements in customer engagement.

When evaluating costs, consider assessing some key factors:

- **Data preparation** — data mart creation, populating agent-ready metadata
- **AI product engineering** — interfaces, workflows, prompt sheets, MCP integrations
- **Change management** — communication, training, and (for the largest programs) role redefinition or restructuring
- **Scaling compute** — usage frequency, query length, model choice

Deploying new services can take time to build evidence, so include measures such as customer satisfaction or user adoption are often leading indicators of value.

Successful generative AI implementation in life sciences demands careful prioritisation and setting realistic expectations.

A shortcut to identify high impact use cases

Our experience indicates the most successful generative AI implementations have targeted three broad areas:



REPETITIVE MANUAL TASKS

Processes occurring with high frequency usually justify an automation investment. Process steps that had to be conducted by people, such as reviewing documents for compliance checks or tagging/categorising free text data fields, can now be automated or semi-automated using generative AI.



LOTS OF TEXT

Generative AI delivers transformative capabilities for text analysis compared to older, natural-language processing techniques. This is especially valuable for hard-to-mine text repositories. Many organizations have centralized numerical data, while leaving text data scattered. This presents an opportunity gap that generative AI can fill.



CODE GENERATION

Generative AI significantly accelerates coding tasks. This includes including coding for advanced analytics to accelerate data science. Coding copilots are now widely used, for example with > 85% adoption in a pilot of >500 IQVIA software developers.

The emerging killer app

People use reports and dashboards for different purposes, and some use cases are particularly well-suited to AI Agents.



DATA ANALYSIS

Agentic AI toolkits with reasoning agents can now analyze data on the fly with a high degree of accuracy. Implementation is still complex, requiring detailed metadata and extensive testing for correct data interpretation. When given access to large data repositories, the latest agents can write detailed reports that previously took weeks of analyst time.

- Many reports are designed for data browsing where users need to look up facts or extract certain pieces of data, and a prompt to an agent can be faster and easier than many filters and settings.
- Some dashboards get overloaded with metrics, facts and KPIs. AI Agents can summarize this information overload to surface the interesting insights. Customer360 dashboards for sales reps are a common example.
- User non-adoption, especially for senior or executive users, is a common problem for performance dashboards. AI Agents can summarize and push KPIs and personalized insights to these users through other channels such as Slack, Teams or email, also driving traffic back to the dashboard.

Common to these analytics and reporting use cases is a less expert or casual user. Expert users and very frequent users quickly learn how to navigate a system to get exactly what they need quickly, so chatbot interfaces only slow experts down.

Agentic AI toolkits with reasoning agents can now analyze data on the fly with a high degree of accuracy.

Understanding scaling of agentic AI

Proof of Concept (POC) projects demonstrate capabilities in a controlled environment and can often deliver an exciting initial output. Production systems, however, must perform consistently at larger scale, automation, reliability and under the demands of important business processes operated by a diverse set of users.

Production systems have more extensive requirements than POCs such as systems integrations, data modeling, system monitoring, training, support, user administration and many other functions for a robust service. The agentic AI steps may be an essential component but only a small part of the work required. With these additional requirements, it's often better to integrate generative AI services into existing platforms, or to acquire pre-integrated agentic tools from software platform providers, versus deploying standalone tools.

Quality and consistency present formidable obstacles when scaling up a service from a POC or pilot. Creating a single, high-quality AI agent can be challenging, but creating more general-purpose agents that can scale across data sets and use cases is even more difficult.

Luckily, with huge investments from major technology companies, agentic AI capabilities are improving at a rapid pace, and there are more options to buy packaged solutions with agentic workflows tuned for stricter guardrails, deterministic steps, and additional ways to suppress LLM hallucinations. Developments such as Model Context Protocol (MCP) and Agent-2-Agent (A2A) are making it easier to connect modular components for more powerful and complete agentic solutions. Reasoning agents are better at explaining logic for the user to check for accuracy of understanding and interpretation, and often they provide code for data queries and calculations that allows further transparency and validation.

Change management for agentic AI

While Agentic AI tools are typically easy for end users to learn and use, adoption requires much more than training, and organization change management is an important workstream for any technology deployment.

The core principles of organization change management apply for agentic AI adoption, including the importance of leadership commitment, clear communication, and forums for stakeholder feedback. Larger changes require more planning and effort.

Successful adoption of AI

AGENTIC AI CREATES SOME ADDITIONAL CONSIDERATIONS FOR SUCCESSFUL ADOPTION

Organizational inertia is strong, so large efficiency gains need clear efficiency goals, and often a change in the mix and number of people deployed to a business function.

Ensure outputs undergo human review for consequential decisions or sensitive topics, such as compliance reporting. Over time, monitoring of quality metrics can allow more automation with lower rates of human review.

Non-expert user groups tend to adopt Agentic AI tools faster, as experts have a lower tolerance for any limitations of agentic AI solutions.²

Validation processes must be robust and ongoing. Unlike deterministic software with predictable error patterns, generative AI errors manifest unpredictably. When there is no provable validation step, programs must set thresholds for acceptable.

While Agentic AI tools are typically easy for end users to learn and use, adoption requires much more than training, and organization change management is an important workstream for any technology deployment.

Balancing innovation with reality

Successful generative AI implementation in life sciences demands careful prioritisation and setting realistic expectations. By targeting applications that combine high business impact with manageable cost and risk, organizations can avoid the pitfalls that derail so many agentic AI projects.

This article deliberately simplifies the topic to give rules-of-thumb and tips for business leaders to make more sense of the cacophony of options and products on offer. There will be important exceptions to the advice provided, and the pace of technology development will make some points obsolete. We hope that this nonetheless give some insight for program successes or delays and helps leaders invest in the solutions that drive real business value.

References

1. Source: BVP, BAIN and AWS survey of 400+ Healthcare execs. Available from: <https://www.bvp.com/atlas/the-healthcare-ai-adoption-index#Full-slide-deck>
2. Rink C, O'Reilly W. Where to Deploy Generative AI Solutions in Life Sciences. IQVIA White Paper. 2025 Mar 19. Available from: <https://www.iqvia.com/library/white-papers/where-to-deploy-generative-ai-solutions-in-life-sciences>

Authors



CHARLES RINK

Sr. Principal,
Information Management &
Analytics Technology, IQVIA

Charles Rink has over 20 years of experience in strategy consulting, technology, and analytics for the life sciences industry globally. Based in London, he brings extensive experience advising clients in commercial strategy, analytics automation, and omnichannel operations. He holds a BS degree in Molecular Biochemistry & Biophysics and Economics from Yale University.



WILLIAM O'REILLY

Director,
Product Offering Development,
IQVIA

Will leads the Generative AI and Machine Learning product capabilities across Orchestrated Analytics, including the IQVIA AI Assistant. With a diverse background as a clinical doctor, healthcare data analysis consultant, and data product manager, Will leverages his varied expertise to bring a unique perspective and create innovative solutions in Life Science technology products and services.



Embrace the future of healthcare with IQVIA's AI-driven solutions. Our expertise in AI can transform your commercial strategies and drive success. Discover more on our [website](#).