INTELLEGENS REPORT

The future of agentic AI in R&D

Strategic summary, June 2025

© 2025 Intellegens Limited
intellegens.com | info@intellegens.com
Intellegens, The Studio, Chesterton Mill, Cambridge, CB4 3NP, UK

intellegens

Executive Summary

This short report summarizes insights presented by Intellegens at an invitation-only online seminar in April 2025 alongside feedback from attendees at the session, who were senior R&D personnel in sectors including chemicals, materials, foods and beverages, and life sciences. Agentic AI – autonomous, tool-using AI systems with Large Language Models (LLMs) as decision engines – is emerging as a transformative force for R&D. Rather than isolated models, agentic systems coordinate and automate workflows across diverse software and data ecosystems. The seminar discussions and survey results reported here show a high degree of industry optimism around the future implementation of this technology but also highlight critical challenges around orchestration of proprietary tools, system integration, IP protection, and data governance. Companies that harness agentic AI effectively will unlock new levels of R&D productivity, agility, and innovation.

Background

The case for agentic AI

Agentic AI is emerging at a pivotal moment for research and development. As data volumes grow, tools diversify, and pressure mounts for faster innovation cycles, traditional software workflows are no longer sufficient. Scientists and engineers need systems that not only generate insight, but actively orchestrate decision-making across fragmented tools, platforms, and datasets. This is where agentic AI offers a step change. Unlike traditional AI models that focus on prediction or classification in isolation, agentic systems coordinate tasks, select tools, and act on behalf of users. They do this by combining reasoning engines — typically Large Language Models (LLMs) — with orchestration frameworks that integrate across diverse software environments. For R&D, this means accelerating everything from experimental design to data analysis to knowledge discovery — not by replacing human scientists, but by enhancing and automating the workflows that support them.

Intellegens and agentic AI

Intellegens is well positioned to help realize this potential based on core competencies in:

- **Machine Learning and AI:** Intellegens is a proven developer, implementor, and integrator of AI technologies, with best practices in machine learning established through years of successful application in scientific and industrial domains

© 2025 Intellegens Ltd.

- **R&D Tooling:** The Intellegens focus has been on delivering AI within a modular, production-ready suite of decision-support tools built specifically for discovery and development in chemicals, materials, life sciences, and manufacturing

- **Data Security:** Intellegens has deep expertise in secure system design and governance, with ISO 27001 certification.

The focus on tools is particularly pertinent. Agentic AI systems require more than intelligence; they require access to the right tools. Intellegens already has many of the building blocks that will be needed — a modular stack for prediction, optimization, design of experiments, analytics, and data integration.

At the heart of the agentic system lies the LLM — the "brain" of the agent, responsible for interpreting user intent, planning a solution, and orchestrating actions. Initially, our agents will leverage the Alchemite™ tool stack — orchestrating across predictive modeling, explainable AI, adaptive experiment design, and workflow integration — all within a tightly governed, enterprise-grade environment.

Looking ahead, these agents will operate across broader R&D ecosystems. They won't just use Intellegens-built tools — they'll need to interact with and orchestrate external systems. Examples might include Microsoft Teams, Google Cloud Platform (GCP), SQL databases, SAP systems, and bespoke in-house infrastructure. In this future landsape, the result will be a collaborative, cross-platform system that uses the best available resources to solve scientific challenges more effectively and efficiently than ever before.

Strategic seminar, April 2025

On 3 April 2025, Intellegens hosted a strategic seminar for some key customer organizations. We shared insights into the current status of agentic AI technology and ideas for rolling this technology out alongside our Alchemite™ solution. We also gathered input from attendees through the use of live polls. Attendees included senior R&D personnel from top chemicals companies, several materials producers, one of the world's leading multi-national food and beverage corporations, and biotechnology innovators. Their feedback thus constitutes a good sample of industry sentiment relating to agentic AI across a range of strategic sectors.

This report provides a brief summary of the meeting, enabling readers to get an overview of agentic AI, its potential and some of the barriers for implementing it, and current industry thinking on the topic.

Introducing agentic AI

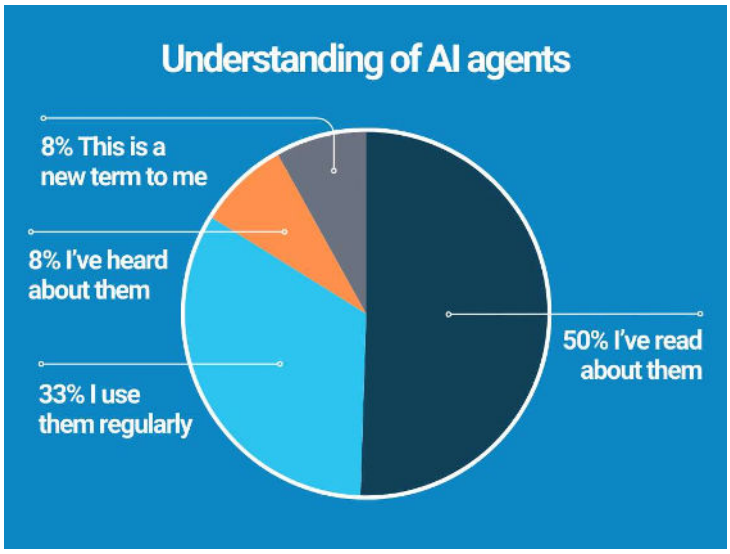

What is Agentic AI?

Agentic AI refers to the next evolution of intelligent systems: software-based agents that not only process information but actively solve tasks by orchestrating the right tools and workflows across complex digital environments. When questioned about their understanding of AI agents (left), seminar attendees gave a good demonstration of current market awareness. Agentic AI is already being applied in R&D organizations, but the great majority – even among this relatively 'AI aware' group of R&D professionals – are in the research phase, recognising the potential of agentic systems but not yet operationalizing them.

The Intellegens team gave a deeper perspective. These software-based systems are designed to perceive information, make decisions, and propose actions in digital environments. In R&D contexts, agents typically embody 'personas' (for example, a research scientist) and can plan strategies, dynamically select tools, and execute multi-step workflows to achieve scientific goals. At the heart of an agent lies a Large Language Model (LLM), which acts as the reasoning engine, planning how best to solve a user's query. However, true capability emerges only when the LLM is combined with an orchestration framework that gives the agent access to a full ecosystem of specialized tools. A complete agentic system integrates:

- **LLMs** – Core reasoning and multimodal understanding
- **Embeddings and vectors** – Long-term memory and context
- **RAG** – Access to live/external data
- **LangChain / orchestration** – Dynamic tool selection and task routing
- **Memory and context** – Personalized, adaptive interactions
- **APIs and plugins** – Integration into lab systems, databases, enterprise tools
- **Cloud and scale** – Scalability, uptime, and resilience
- **Token optimization** – Efficient resource management for complex workflows

Rather than simply answering queries, an agentic system analyzes the user's goal (for example, "What experiments should I do next to find a lower cost material?"), strategizes the best approach by recalling relevant past actions and planning appropriate next steps, selects and applies the most suitable tools — whether in analytics, simulation, design of experiments, or knowledge discovery — and collates the outputs into a final actionable proposal (such as "Test the following 10 candidate formulations suggested in the attached experimental plan.").

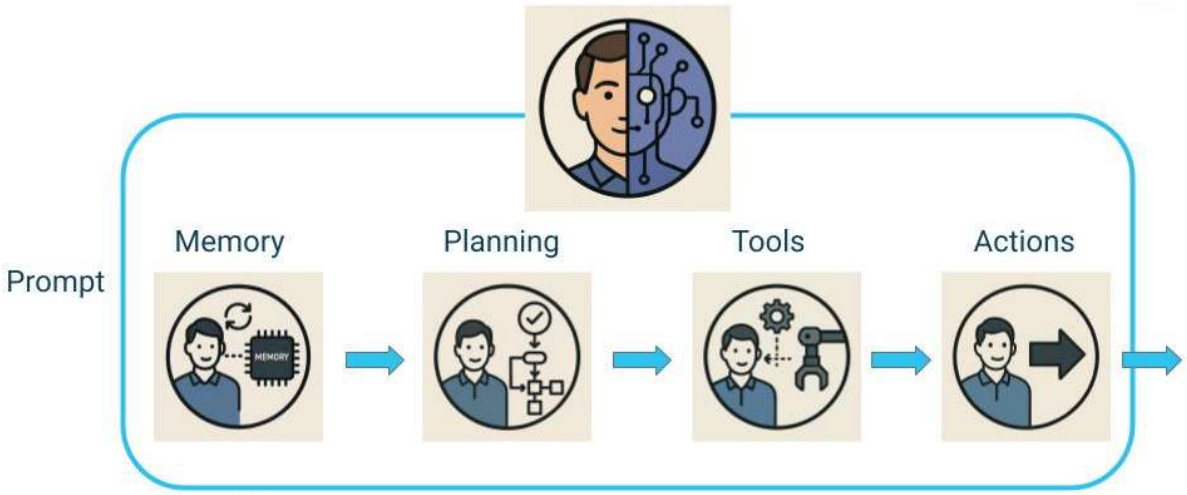

For particularly complex challenges, Multi-Agent Systems (MAS) come into play. Here, multiple specialized agents — each focused on different aspects of the problem (e.g., hypothesis generation, evaluation, optimization) — collaborate and coordinate to reach higher-quality outcomes than any single agent could achieve alone. An early example of MAS success is Google's AI Co-Scientist (2025), which used a team of agents to propose novel hypotheses, validate experimental strategies, and uncover scientific breakthroughs — demonstrating that agentic AI is capable of real, autonomous research, not just assisting human scientists.

Adoption – potential impact and barriers

Seminar attendees were almost universally positive about the potential of agent-based AI to transform R&D decision-making. In live polling, 87% thought the technology would be transformational in how R&D decisions are made, while a further 7% saw it driving impact in specific niche areas; the remainder were still evaluating its full implications.

87% of attendees believed Agentic AI will result in a **transformational shift** in how R&D decisions are made.

intellegens

Although the technology is developing at an extraordinary pace only three years after the democratization of LLMs through platforms like ChatGPT, the vision of fully agent-powered R&D is not yet fully realized. Delivering on this promise will require agentic systems tailored to the real-world workflows of scientists, not generic AI deployments. A crucial success factor will be the ability of agents to dynamically access, select, and orchestrate the right set of tools, since the true value of any agent is determined by the quality, breadth, and integration of the tools it can leverage.

It is the **tools** that the agent is able to call on that really determine its value

intellegens

When asked about barriers to adoption, seminar participants highlighted several critical considerations. Interestingly, trust in AI outputs was rated as a relatively minor concern — likely reflecting the

scientific community's growing familiarity with AI systems like Alchemite™, where AI is used to guide and optimize experimental work rather than replace it. Built-in validation through experimental confirmation remains the norm. By far the greatest concerns related to business risks around intellectual property (IP) protection, data security, and governance. As agentic systems evolve to integrate across multiple internal and external tools, ensuring robust security frameworks, maintaining data sovereignty, and managing IP ownership will be fundamental to successful adoption. Any serious rollout of agentic AI in R&D must prioritize these practical realities from the outset.

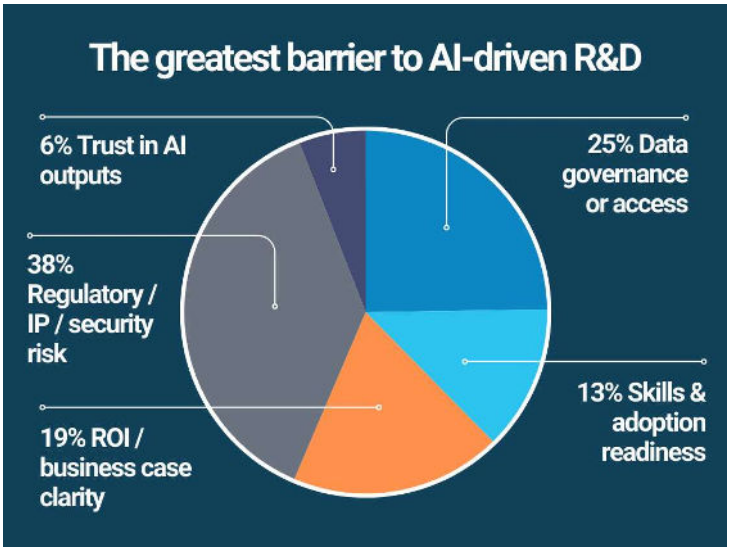

A model for agentic AI adoption

Intellegens has identified a clear strategic model for helping R&D organizations adopt agentic AI technologies effectively, building on both our technical capabilities and insights gathered during the strategic seminar. Successful adoption will hinge on three key principles:

- Competitive advantage for user organizations will come from building **proprietary agents** that orchestrate the company's internal research tools, data assets, and workflows — not simply from relying on general-purpose AI. Solutions must leverage proprietary knowledge alongside commercial and open-source resources.
- **Standardized, secure integration is critical**. Proprietary agents should be constructed using proven, extensible technologies and embedded into existing R&D platforms, ensuring that scientific teams benefit from trusted security frameworks and familiar workflows.
- **Phased, targeted rollout** will accelerate adoption. Early projects should focus on high-value, high-impact tasks, using persona-based designs to support specific R&D roles where agentic systems can deliver measurable gains.

To begin building agentic capabilities, organizations should map their existing business hierarchies to understand how R&D roles and personas interact in practice. It is critical to mirror the real structure of scientific teams and workflows. First, organizations should define key personas — for example, experimental scientists, data analysts, project managers, and domain specialists — and analyze the responsibilities and interactions associated with each

role. Each persona should then be broken down into key tasks: repeatable decisions, actions, or analyses that agents can be designed to support. Using the agentic AI framework:

- Individual task-level agents are built to solve specific problems.
- Persona-level orchestration agents manage the planning and prioritization of those tasks for each role.
- Multi-agent systems allow different personas — and their associated agents — to collaborate dynamically toward shared R&D goals.

For example, a manager agent might set a high-level research objective. Experimentalist and analyst agents would then coordinate autonomously, planning experiments, analyzing results, and refining approaches — mirroring how human teams operate, but with greater speed, scalability, and precision.

Intellegens' Alchemite™ Suite already provides a strong foundation for this phased approach. Each Alchemite™ app targets a core R&D persona:

- **Managers** reviewing and directing experimental programs (Alchemite™ Viewer)
- **Scientists** testing hypotheses / exploring relationships in data (Alchemite™ Explorer)
- **Experimentalists** designing optimal experiments, fast (Alchemite™ Designer)
- **Subject matter experts** conducting advanced ML analysis (Alchemite™ Innovator)
- **Data scientists** integrating ML into lab systems and workflows (Alchemite™ Architect)

Today, these apps allow users to apply structured-data machine learning to accelerate discovery. Tomorrow, agentic systems will allow R&D teams to intelligently orchestrate a broader range of tools — including simulation, modelling, optimization, and knowledge extraction — dynamically and autonomously across workflows.

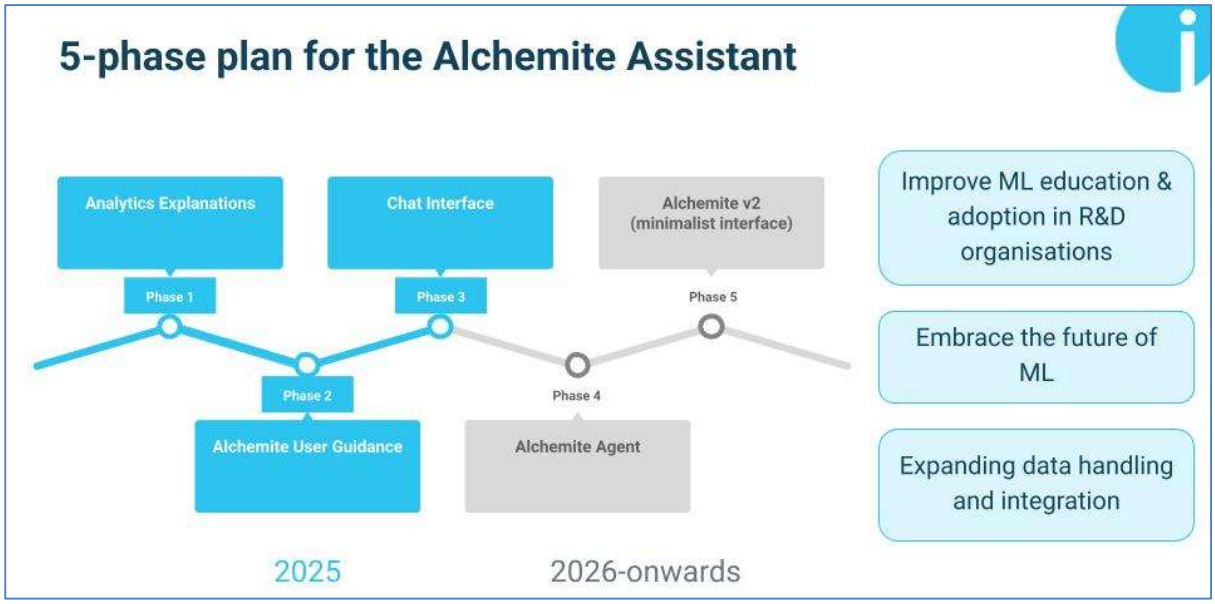

Intellegens has developed a 5-phase roadmap (above) to guide this evolution. Starting with enhanced user interaction through analytics explanation and guided decision support, agentic

capabilities will expand towards full autonomous orchestration, culminating in interconnected agent ecosystems operating across internal and external platforms.

This outline plan was shared with attendees at the seminar and interest in collaborating with Intellegens as it is further developed was established.

Conclusions

Agentic AI represents a critical inflection point for the future of R&D. Industry leaders recognize its transformational potential. However, realizing this potential will depend on more than deploying general-purpose AI systems. It will require the development of secure, intelligent agents that can orchestrate a company's proprietary data, tools, and workflows, embedded directly into the day-to-day operations of scientists.

The key priorities are clear:

- Security and governance must be foundational, not optional.
- Tool integration and orchestration must be the core of agentic capabilities, not just language generation.
- Phased, persona-driven rollouts must guide adoption, ensuring measurable value from the outset.

Intellegens is advancing a clear roadmap to pilot and expand agentic AI within real R&D environments using our proven Alchemite™ suite, secure infrastructures, and scientific expertise as a launchpad. We invite organizations that are ready to move beyond theory — to build the next generation of interconnected, intelligent R&D systems — to collaborate with as strategic partners in this transformation.

About Intellegens

Our vision is that machine learning will drive innovation and deliver value wherever data is used in R&D. We aim for best-in-class easy-to-use machine learning software for data analysis in chemicals, materials, life science, and manufacturing. Our Alchemite™ technology originated at the University of Cambridge and development is on-going at Intellegens, in close collaboration with our growing community of customer organizations. These represent sectors including additive manufacturing, aerospace, alloys, batteries, biopharmaceuticals, ceramics, chemical processes, composites, consumer products, cosmetics, drug discovery, energy, food and beverage, formulated products, paints, plastics, and printing technology.

www.intellegens.com | info@intellegens.com