UiPath®
Agentic automation

Agentic automation in life sciences

The future of automation is agentic and robotic.

Contents

The AI revolution continues: Agentic AI

In the spring of 2024, we released the Intelligent Document Processing in Life Sciences white paper that focused on digitizing paper based and manual processes across the life sciences industry. Although IDP continues to help organizations transform their manual processes, the emergence of agentic automation is now providing additional capabilities that will allow life sciences organizations to leverage an autonomous agent that is focused on achieving goals in a much more probabilistic manner.

Unlike IDP, agentic automation can manage and orchestrate complex end-to-end enterprise processes using agents, robots and humans. Specifically, agents can:

- Provide context-based information to users based on their queries. For life sciences organizations, they could determine eligibility to donate blood, plasma or tissue or participate in a clinical trial based on inclusion/exclusion criteria contained in a clinical protocol.
- Create and run workflows across enterprise applications, find the right application programming interfaces (APIs), plan the best ways to do things, and meet user requests. For example, an agent could reject an invoice sent to SAP after a robot reviewed the strategic sourcing agreement in Coupa stipulating that the organization is due an additional 10% discount when shipments are sent more than 10 days after purchase order release.
- Streamline end-to-end processes by assuming a virtual coworker role and executing actions autonomously. A quality agent, for example, could manage operational data and leverage robots that analyze a variety of historical manufacturing variables in quality assurance data to identify batches that may be trending negatively even though the metrics are still within tolerance.

Also, the matrix above highlights early, high value agent types that have been identified when discussing various agent strategies in life sciences. The agents can understand the specific steps of an end-to-end process while leveraging robots and humans to assemble analytic information and ensuring the process proceeds in a timely, accurate and transparent manner. This would include insertion of a human reviewer for instances where the data does not meet the defined validations, or the agent is unclear as how to proceed in the process.

Specifically, we have mapped out the quadrants of the matrix and have detailed the upper right and left sections below:

- **Upper right quadrant:** The "quick wins" section highlights agents that would focus on process areas and/or document types that are more common including sales, finance, or quality agents. These agents target more horizontal use cases that are common to many industries and, in most cases, require a significant amount of manual effort. The agentic orchestration will leverage a combination of robots, humans or other agents, leveraging both generative and specialized AI, that can deliver an "out of the box" model with a high level of accuracy and predictability. These scenarios represent excellent starting points in your agent journey as we combine agents, robots and humans to enable agentic orchestration across an end-to-end process capable of transforming the process and user experience.
- **Upper left quadrant:** The "strategic bets" section identifies process agents and document types that are more complex but also have higher value. While these agents are highly strategic, they may still be out of reach relative to the perceived risk of the organization and/or the maturity of the AI technology.

> **"AI agents will become the primary way we interact with computers in the future. They will be able to understand our needs and preferences and proactively help us with tasks and decision making."**
>
> Satya Nadella, CEO of Microsoft.Jun 14, 2024

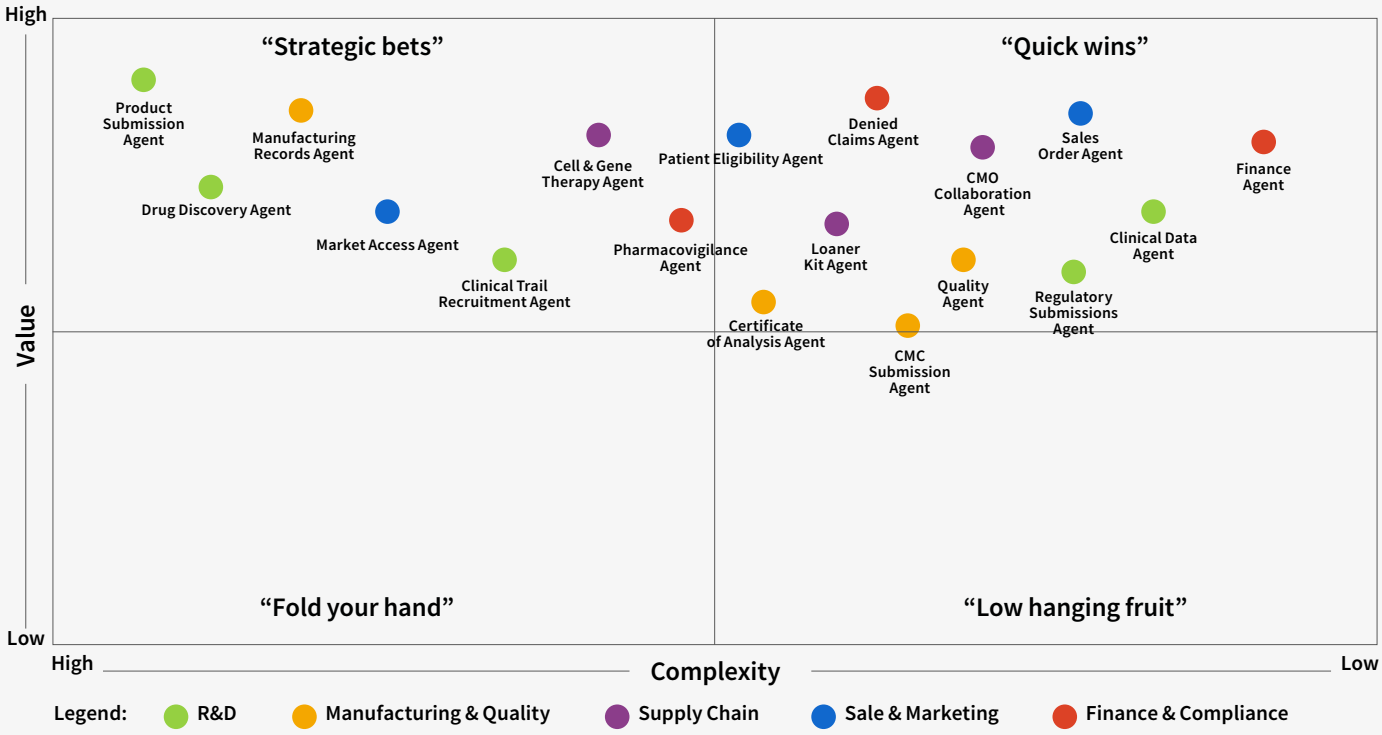

For example, a drug discovery agent may not be capable of completing an entire new drug application (NDA) or biologic license agreement (BLA) filing given this would entail thousands of pages of analysis and scientific documentation accumulated over the course of many years and numerous individual clinical trials. Instead, organizations may want to focus on smaller components of the overall NDA/BLA end-to-end process. Robots could leverage generative AI to create the narratives based on the clinical data associated with various spreadsheets and/or quality documentation who would then route the narratives to an agent who would review and, based on policy documents or predefined validations, route the document for approval or approve it themselves. An approved version of the document would be routed to your document management system, for example Veeva Vault, where the electronic Trial Master File (eTMF) would maintain the most current, approved version of the documents related to a clinical trial. An agent of this type could shave hours, days or even weeks off time to market, which has a real impact of accelerating license revenue by $2M - $4M per day for each day of time to market reduced.

AI continues to evolve across life sciences

What started in December 2022 with the release of generative AI has continued to evolve and iterate. A revolution has begun as we move from knowledge-based, generative-AI-powered tools, like chatbots that answer questions and generate content, to AI-enabled "agents" who act autonomously, focus on achievement of goals rather than tasks and leverage a probabilistic approach to solving complex problems. Agentic automation leverages humans, robots, and agents to orchestrate foundational models to execute complex end-to-end workflows across an increasingly digital world. In short, agents allow organizations to move from thought to action.

With agentic orchestration, organizations gain the unique capability to orchestrate UiPath agents, third-party agents, new or existing UiPath automations, and humans across end-to-end agentic workflows and systems. UiPath agents can leverage integrated process intelligence and agentic orchestration to identify valuable insights into key performance indicators, helping to uncover inefficiencies and drive continuous process improvements.

"For a long time, we've been working towards a universal AI agent that can be truly helpful in everyday life."

Demis Hassabis, Co-founder and CEO of DeepMind

Competitors from across the technology landscape have created a new language of terms, buzzwords, and concepts that have flooded the market as organizations look to create a competitive advantage for their products and services. Unfortunately for the industry, this has simply created a lot of noise and confusion.

Robots

- **Rules based**
- **Predictable / repeatable**
- **Determinstic**

Best for routine tasks that require **high reliability & efficiency**

Agents

- **Goals based**
- **Independent / autonomous**
- **Probabilistic**

Best for **ad hoc tasks** that require **high adaptability**

Historically, organizations have leveraged AI-powered automation platforms, including attended and unattended robots, that have automated and saved millions of labor hours for healthcare and life sciences companies across the globe.

Life sciences organizations have utilized enterprise-scale automation platforms that can develop, manage and maintain rules-based, highly predictable, and deterministic robots that are now extending that strategy to include agents that are simply an extension of that AI powered automation platform. Global life sciences companies count on automations to consistently and predictably execute their automations to ensure accurate data and decision-making that could impact patient safety or product quality.

Contrary to traditional automation, agents are goals-based, independent decision makers. Agents are more probabilistic than robots, ideal for ad hoc scenarios, and do not require rules-based criteria to consistently and effectively make decisions. Agents are highly adaptable in their decision making and are appropriate for use cases that tolerate some degree of uncertainty. Agents will learn from prior decisions and those learnings can be captured, tested and implemented, usually with the approval of a human reviewer, to improve the performance of the model.

Central to the success and execution of an agent strategy is the ability to leverage data from across the enterprise landscape. Global life sciences organizations have a complex and heterogeneous mix of enterprise systems. You may be familiar with the many acronyms for enterprise systems leveraged to develop, manufacture, and distribute lifesaving therapies. These include ERP (enterprise resource planning), CRM (customer relationship management), SCP (supply chain planning), LIMS (laboratory information management systems), MES (manufacturing execution systems) or QMS (quality management systems) whose organizations may have limited or no access to the process data across those systems that will minimize the decision making effectiveness of an agent. Leveraging a proven, highly extensible solution to orchestrate data across enterprise systems of record and systems of engagement will empower an agent and enterprise to operate at scale.

More importantly, agents will allow life sciences enterprises to transform regulated and non-regulated processes by allowing them to include a human in the loop for any scenario that is unclear and/or does not meet the data validation requirements defined for the agent to proceed. An agent can escalate the scenario to an appropriate human resource who can provide additional guidance and insight. The agent can recognize the steps taken and learn from the human in the loop decision making that could and would minimize or eliminate a similar occurrence in the future.

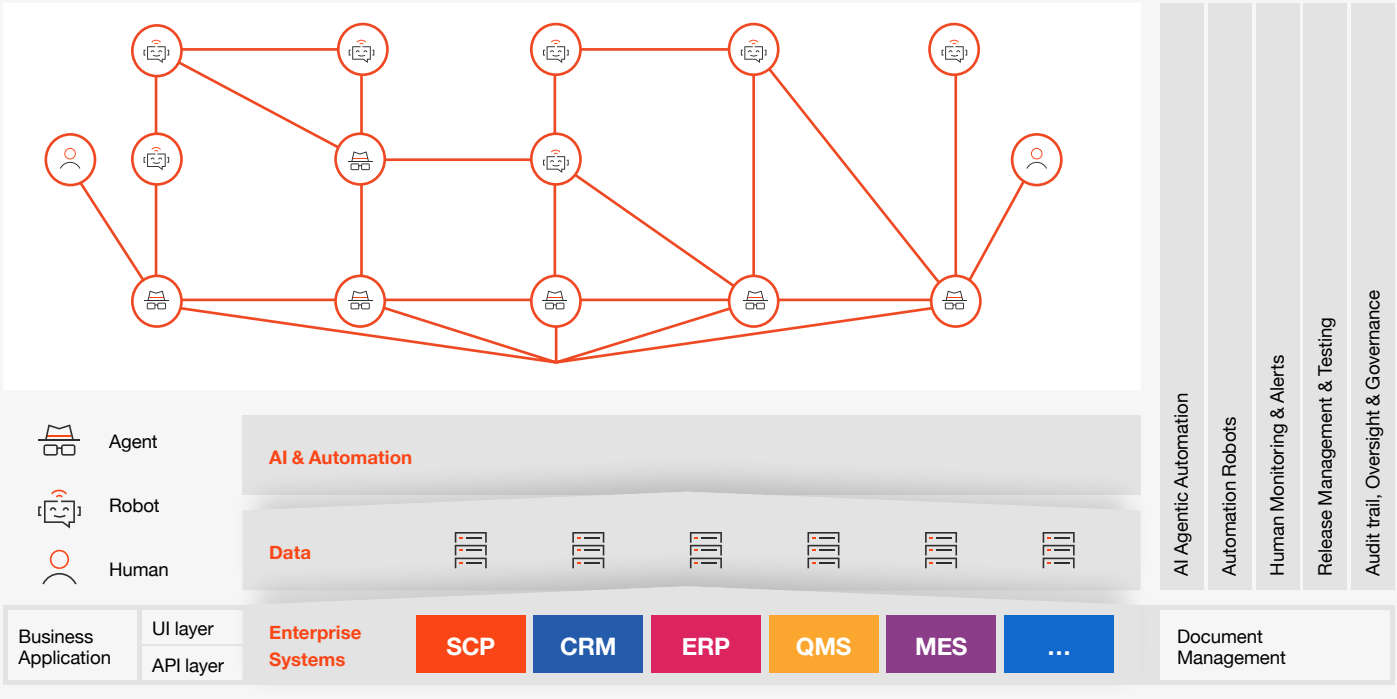

Quality assurance, quality control, and AI

As organizations look to innovate with AI, they must also ensure compliance with industry regulations, such as Good Manufacturing Practices (GMP), 21 CFR Part 11 or 21 CFR Part 211. These regulations have been around for decades and are not expected to change to accommodate AI or any other technology. Consequently, any opportunity to leverage AI must be done in a manner that is compliant and in alignment with traditional quality goals and objectives of global regulators.

Regardless of the use case, it is important to ensure that any AI-powered agentic automation strategy meets, at a minimum, the following directives:

- **Governed, secure, and transparent data**
  Any data related to a model, robot or agent will require that the data (training, tuning and testing) is governed, secure & transparent. This would include maintaining copies of the model, version control and leverage change management best practices.
- **Robot or agent actions are attributable and recorded in an audit trail**
  This would include log files with full transparency to ensure visibility into decision making and data validations, where appropriate.
- **Human in the loop**
  Any agentic use case should provide for an escalation to human approval when necessary.
- **Real-world performance monitoring**
  Agent models will require monitoring to ensure the accuracy of the model while also looking for opportunities to update models with higher levels of accuracy.
- **Testing is not optional**
  Any update of an AI model requires full testing to ensure the organization fully understands the agent's decision-making process. Self-learning and/or real time updates of a model without sufficient testing will be a violation of life sciences quality standards.

"Agentic systems refer to digital systems that can independently interact in a dynamic world. While versions of these software systems have existed for years, the natural-language capabilities of gen AI unveil new possibilities, enabling systems that can plan their actions, use online tools to complete those tasks, collaborate with other agents and people, and learn to improve their performance."

McKinsey, July 2024

Agentic use cases for life sciences

Finance agent

AI-powered automation has been widely deployed in finance. But agents could allow life sciences organizations to go further in improving their productivity, reducing their cycle times, and eliminating costs. The use of a probabilistic agent in finance could have a tremendous impact on an individual process area.

In the example and graphic below we have identified three types of agents that will be leveraged to support procure to pay agent. Here's a brief introduction:

- **Process agents:** High-level agents that manage processes by coordinating task/skill agents and workflows.
- **Skill agents:** Specialized agents that perform a single, specific task without relying on other agents.
- **Task agent:** Agents that manage specific tasks by orchestrating single/multiple skill-level agents, robots and workflows.

Use case 1: Procure to pay

Let's use the procure to pay function to highlight how a finance agent could create a significant transformation of the process.

In the graphic below you can see how the process agent manages the end-to-end process that may include assistance from other agents or robots. The skill agent is focused on a specific activity, like purchase requisition or invoice processing, which may include the use of task agents, robots or workflows that provide the insight to drive an accurate response. The process agent is aware of the end-to-end process and uses the other agents and robots to acquire insights to make the necessary decision or escalation.

One of the benefits of an agent is its ability to leverage existing automation to drive a more compelling process. Robots, agents, and humans support the end-to-end process and can reuse automations across the enterprise for any other use case that may require the insights.

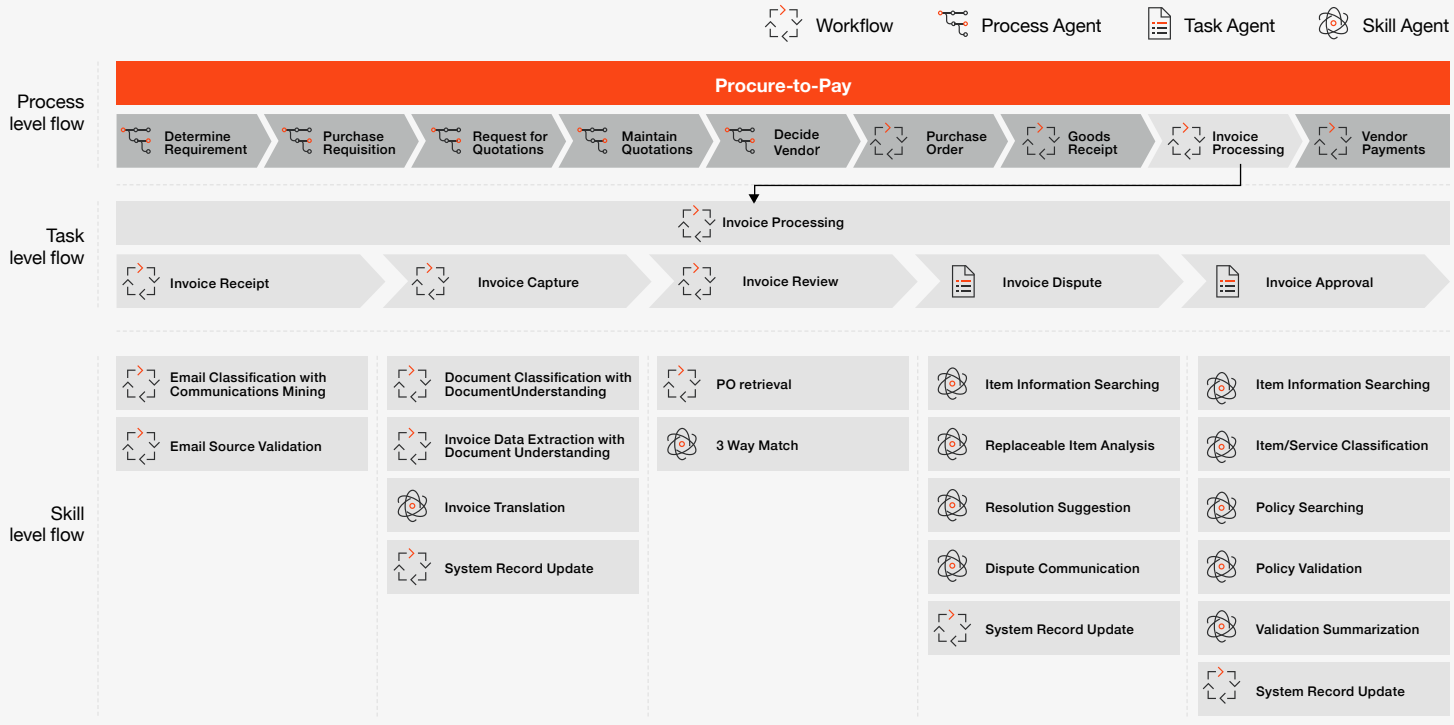

For example:

- A skill agent reacts to the updated status of a batch that uses a contract manufacturing organization (CMO) for the manufacture of a cell and gene therapy product.
- The finance skill agent initiates the request for pickup by a 3PL cold chain logistics provider after notification from the CMO that the batch has been approved and is available for pick up. The finance robot creates the purchase requisition for the services provided by the CMO based on the notification of completion.
- The finance skill agent secures the 3PL logistics reservation via a confirmation number and validates the pickup time with the operational team who secures the reservation and slotting time for the pickup and delivery.
- The robot and the agent work in tandem to coordinate the activity and ensure operational alignment across the enterprise.
- There is no requirement for human involvement, although the agent can escalate to a person in the case of cancelation or operational anomaly.

Scenario 2:

Automation can also enhance or extend processes to include a more comprehensive and robust view of an operation. Organizations may have an existing two-way match process that a robot enables across multiple ERP solutions.

A robot can extend the 2-way match process to a 3-way or 4-way match, if required. Inclusion of purchase order, invoice, receiving, and quality data can be a complicated scenario if they originate across multiple enterprise systems. A robot can aggregate the required data that is compiled, validated, and, when necessary, escalated to a human for subsequent review. Robots will assist an agent in gathering, compiling, and analyzing data while the agent manages a higher-level process including procure to pay.

Scenario 3:

A life sciences company has an agreement with a direct materials supplier that all orders must ship within 10 days of purchase order placement. Both the supplier and the life sciences organization may miss applying the additional 10% discount because it is contingent upon operational performance as the basis for including the discount on the invoice.

An agent, however, can review the full terms and conditions of any agreement at the point of invoice receipt and flag an invoice where an additional 10% discount was not applied based on the actual ship data.

While it may not be realistic or even possible for finance personnel to find these random errors, the result is unnecessary payments to vendors. Agents can provide a more robust and detailed review of, for example, the process documentation at a contractual level to ensure suppliers are compliant with the terms of their agreements.

> **"AI agents will become an integral part of our daily lives, helping us with everything from scheduling appointments to managing our finances. They will make our lives more convenient and efficient."**
>
> Andrew Ng, Co-founder of Google Brain and Coursera

Patient eligibility agent

Eligibility is a nuanced and elusive topic that can take on many dimensions and perspectives in life sciences. For the purposes of our discussion, we are going to focus on the health, genomic, and/or biomarker aspect of eligibility. There are a variety of processes and treatments that require a patient to meet a set of criteria that would make them eligible for therapy, donation, or trial.

Use case 1:

Many patients would like to donate human blood, plasma or tissue that could be used for a downstream therapy or process. Blood or plasma donation is dependent on a variety of factors such as the health of a patient or current or previous medications. For example, if a patient is currently on a blood thinner or was undergoing chemotherapy, they would not be eligible to donate blood or plasma.

Typically, an organization maintains an extensive amount of documentation that defines who is and isn't eligible to donate blood or tissue. Digesting the vast amount of standard operating procedure (SOP) documentation on this topic is a time-consuming task that will limit an organization's ability to recruit, train, and deploy resources across the network of donation clinics in a specific market.

A digital agent, however, can create a quarriable interface that can accelerate knowledge transfer and transform a user's experience. Users can ask questions using natural language and, in response, the agent will review the hundreds or thousands of pages of documentation and return a definitive response to the user. It will not only answer the question but also provide the citation of the information

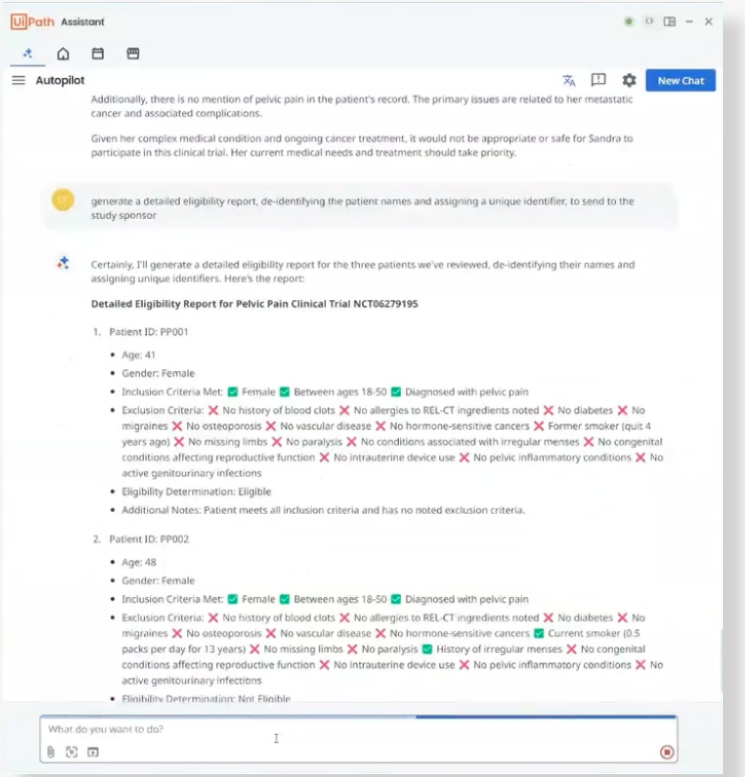

used to make that determination including the actual location in the reference documentation where they were able to make the determination.

In the example below, a user is leveraging AI to determine the eligibility of a patient to donate blood, plasma, or tissue. The AI ingests the organizational documentation, typically standard operating procedures (SOP) that defines the eligibility of a patient relative to that specific donation. The AI agent can analyze the clinical history of the patient, either in person or via an electronic medical record, and query that information with the SOP data that defines eligibility. The agent returns a simple, concise, and accurate assessment of the medical history of the patient that could be used as an initial screening of patients and/or routed to a human for follow up directly with the patient on specific medical questions.

Note that the response is cited and includes links back to the specific section of the medical record including the precise location in the document where the information was found.

Additionally, because the agent is reviewing a defined set of proprietary data specific to this life sciences organization, it eliminates the possibility of hallucinations inherent when scanning the internet. Examples of proprietary data sources an organization might use are SOPs, clinical trial protocols, or product information leaflets. Leveraging an organization's own documentation ensures a consistent and accurate response to operational, product, or technical queries to determine eligibility or enrollment in a variety of functions including trials or donation scenarios.

Also, prompts are leveraged to further simplify the interaction with the user as the AI can create standard and repeated queries that help standardize the interaction with the AI agent. The consistency of the query improves the quality of the response as the system continues to learn from every interaction.

Agents can also embed a complex end-to-end business process to prioritize and/or escalate individual scenarios that may occur during their analysis. Complex end-to-end processes typically require data from a variety of enterprise systems and an agent's ability to orchestrate data across the enterprise is critical to enabling an optimal insight and/or solution.

Use case 2:

In a clinical trial, administrators will spend days, weeks or even months scanning medical records, reviewing registries, or evaluating requests to find patients who can participate in each trial. This is an extremely important function, as without the appropriate patients, trial and potential lifesaving treatment could be delayed or stopped entirely.

It is important to note that 11% of all clinical trials fail to recruit a single patient, and 20% of all cancer trials fail due to insufficient patient recruitment.

Using the power of AI agents, organizations, who have appropriate access to patient data, can take a more holistic and strategic approach to recruiting patients. For example, leveraging AI to review large cohorts of medical records, registries, and/or documentation to identify high probability candidates for trials based on their medical history can dramatically accelerate time to market for new therapies.

For example, an electronic medical record (EMR) is compared to the clinical trial protocol that defines the parameters of the clinical trial and allows the AI to scan an individual or group of patient records to determine potential eligibility. Patients could be accepted directly into a trial if they met specific criteria, or they could be flagged as a high potential candidate. A trial manager could then follow up with the patient to validate their information and determine their interest and/or ability to participate in the trial.

A patient recruitment agent could dramatically accelerate the recruitment timeline and allow organizations to accelerate their trial process and, consequently, accelerate time to market for life saving therapies.

"AI agents will transform the way we interact with technology, making it more natural and intuitive. They will enable us to have more meaningful and productive interactions with computers."

Fei-Fei Li, Professor of Computer Science at Stanford University

Cell and gene therapy agent

The development of cell and gene therapies (CGT) represents both exciting opportunities and significant challenges when it comes to patient engagement. CGT holds immense potential for revolutionizing how we treat diseases. Their importance stems from their unique ability to address diseases at their core, by leveraging a patient's own cells or genes to fight and defeat the disease.

Use case 1:

The continued development of personalized medicines offers the potential for highly effective treatments tailored to an individual's unique genetic makeup. The result is fewer side effects and dramatically improved health outcomes.

CGT has shown remarkable success in treating certain conditions and offering hope for patients with previously limited treatment options. Many of these therapies aim to provide durable or even curative effects, potentially eliminating the need for ongoing treatments. Patients are also better able to participate in their own care journey, better understand their condition and treatment options, and make better-informed decisions.

Conversely, CGT are often complex treatments that require specialized knowledge. The care protocols are unique to the therapy and include a new ecosystem of partners to enable the therapy. Therapies can be expensive, require a diagnostic, and may not be readily available or accessible to all patients. CGT involves lengthy and sometimes burdensome treatment processes. It may require patients to travel to treatment centers, stay in that area for the duration of the initial treatment, and return for any additional treatments.

The creation of a patient engagement agent could significantly impact the processing of CGT patients by reducing cycle times, transforming the patient experience, and improving the outcome for the patient.

The image below highlights an example process from patient engagement through treatment delivery for a patient receiving a CGT. A patient engagement process agent improves the end-to-end CGT process by leveraging other agents, combined with humans and robots, who focus on ensuring the patient successfully schedules, attends, and receives the expected therapy while also providing a support network to assist throughout the therapy.

For example, an agent is empowered to book the time for the treatment at the closest available center while also arranging travel for the patient and any guardian or family member. Other agents can communicate with the patient to confirm appointment times, clarify requirements or special requests, and distribute therapy communications to patient care circles while ensuring appropriate privacy requirements. The agent could also provide non-treatment information to help transform the experience with various tips on how to deal with potential side effects of the treatment while also explaining the next steps in the therapy.

Once a patient is approved for the therapy, an agent can schedule the logistics pickup and delivery of the patient and/or therapy while another agent could notify the contract manufacturing organization with the required information and then communicate the expected delivery date and timelines. Timing is critically important with CGT as the effectiveness of the therapy for many CGT will drop significantly after certain time horizons. For example,

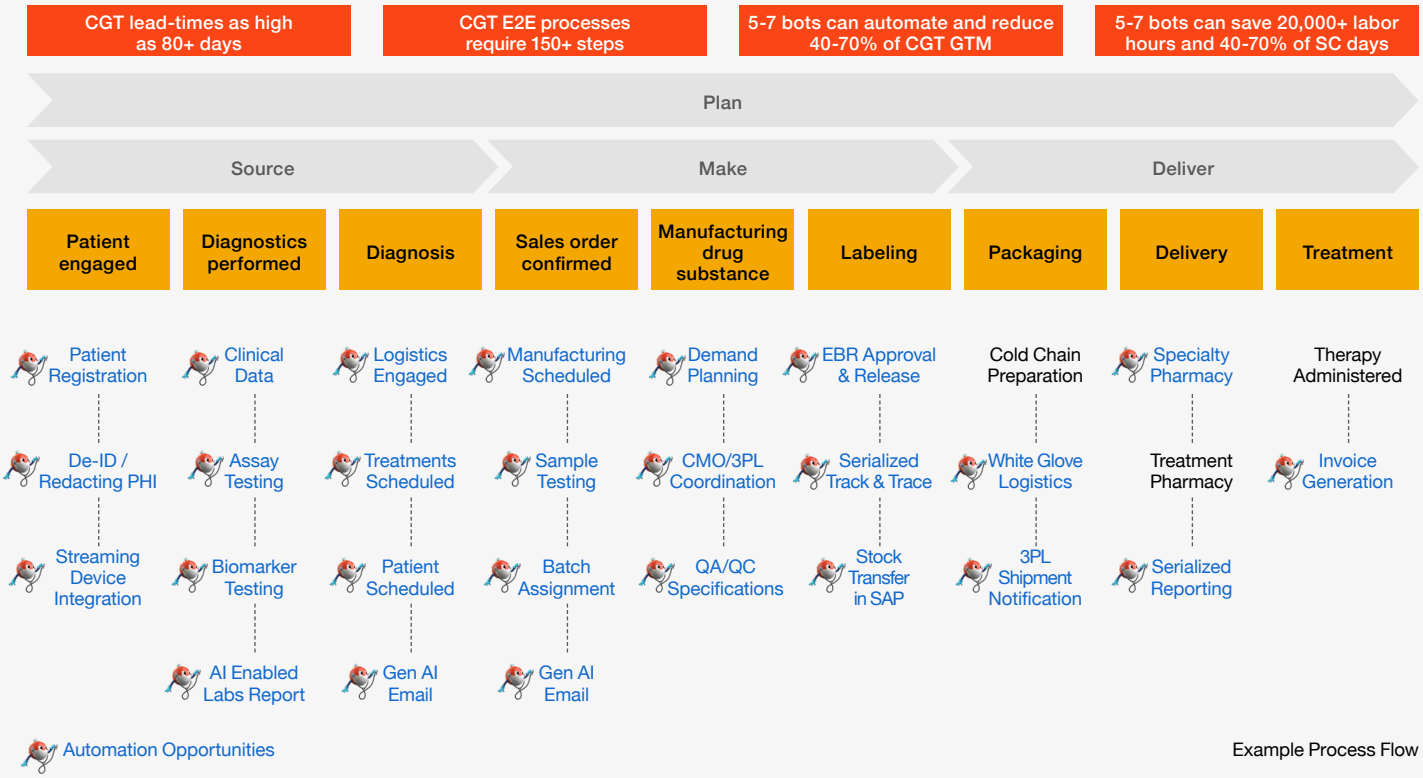

the efficacy of CAR-T therapies drop dramatically after approximately 30 days. Slotting the patient and/or cell or tissue specimen into the production schedule and then aligning specimen pickup times and dates eliminates unnecessary delays in the process while reducing the cycle time from cell/tissue collection to administering of the therapy to the patient.

Use case 2:

From a reuse perspective, the earlier discussed finance agent can also be integrated into the CGT process. Once the cold chain pickup and delivery are completed and manufacturing has been confirmed, the finance agent will be able to evaluate and process the respective invoices for the services rendered. The finance agent can confirm the signed proof of delivery and execute a two- or three-way match or another financial controls to validate the approval of the payment for the services. The agent would be able to contact the vendor regarding any operational or contractual issues identified or escalate the issue to an appropriate human for follow-up.

CGT and/or patient engagement agents are also able to communicate and coordinate with patients to ensure they are aware of all upcoming appointments, fully understand the therapy regimen and instructions, and have confirmed follow-up appointments with healthcare practitioners. As mentioned earlier, existing agents, combined with robots and humans can leverage a natural language quarriable interface with product information leaflets as the proprietary data set with precise and accurate information about a drug, device, or therapy. The interface could also be available directly to patients who can query the repository in natural language that accelerates knowledge transfer on specific data elements, for example symptoms, side effects or drug interactions, to ensure the safety and efficacy of the therapy for the patient.

> **"Agents are not only going to change how everyone interacts with computers. They're also going to upend the software industry, bringing about the biggest revolution in computing since we went from typing commands to tapping on icons."**
>
> **Bill Gates**, Co-founder of Microsoft

More to consider: Drug discovery agent

In the white paper "Intelligent automation in life sciences R&D," we highlighted the many challenges life sciences organizations have in reducing time to market when developing a new drug, device or therapy and how intelligent document management can provide a game changing capability to digitize paper based or manual processes and transform the process.

For many life sciences organizations, research, clinical, or scientific data may reside in a variety of locations, including company research sites, clinical research organizations, clinical trial sites, cloud-based data repositories, electronic lab notebooks, or research institutions. Data management, consolidation, visibility, validation, and analysis can be very cumbersome, difficult, and time-consuming.

Ensuring the access, accuracy and integrity of the data is also one of the most important tasks within a research project. AI-enabled automation can simplify, improve, and accelerate this process in a variety of ways that could provide a significant reduction in employee and researcher cycle times.

Agents can manage this and other complex end-to-end processes while also interacting directly with the research. Using analytical strategies, agents can identify optimized approaches—a process that would take an individual researcher days or weeks to complete.

Agents can analyze vast amounts of data while leveraging standard analytical or scientific methodologies defined in organizational source data documents. Leveraging defined scientific methodologies, agents can analyze the vast amounts of data and identify high potential candidates to pursue that, ideally, should significantly reduce the analysis time for a researcher and time to market for the therapy.

Further, agents can leverage robots that can orchestrate the data from across enterprise platforms or repositories while also understanding the end-to-end process around the data. The agentic orchestration must be both extensible and secure to capture data from across a multitude of enterprise systems while also understanding the process steps necessary to accomplish the analysis.

Process steps would be defined in organizational source data repositories that would allow the agents to be more proactive in their approach including but not limited to making suggestions or recommendations on next steps for the various participants.

Lastly, the agent can quickly and efficiently analyze the data and deliver the results to a human for review. Or, if data validation criteria are met, it can deliver the data in a touchless manner directly to the next step in the process. In either case, the agent will accelerate the product development process in ways not previously possible and significantly reduce the time to market for the development of new drugs and devices.

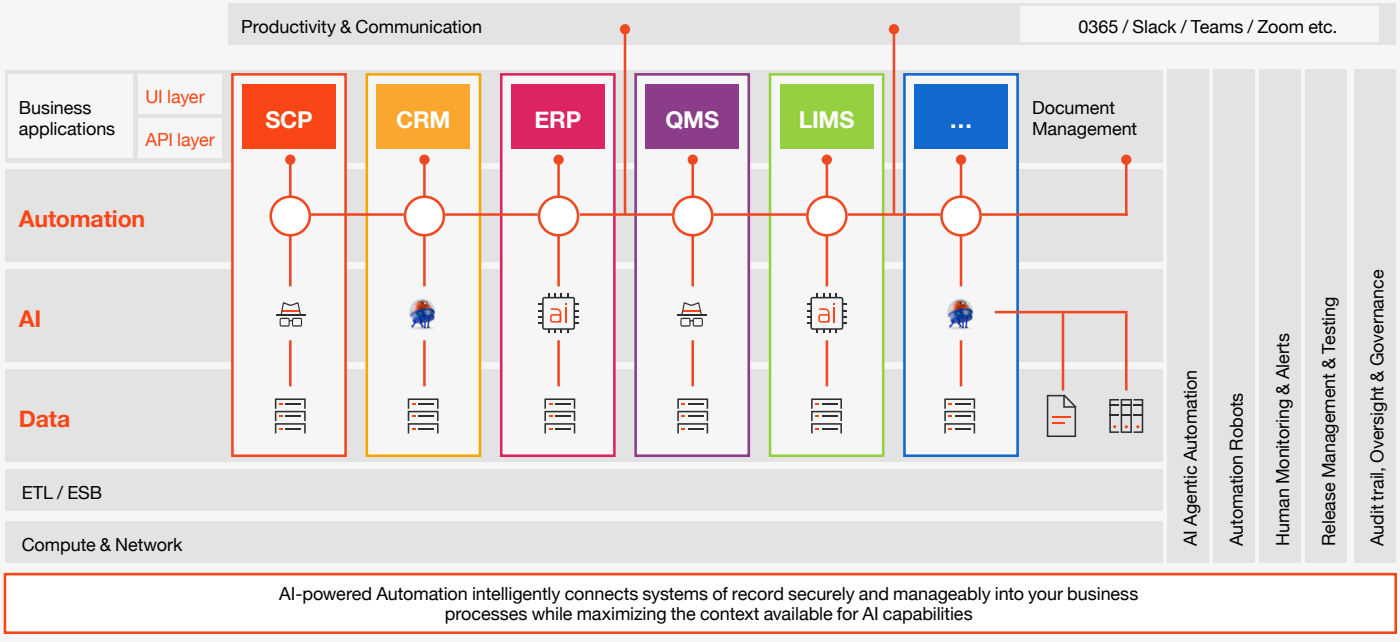

Steps to getting started

Determining where an organization should get started on their agentic automation journey can be a difficult decision given all the noise that exists across the marketplace. We thought it would be helpful to include a few simple criteria to help provide you with some guidance on developing your agentic automation use cases. Included below are suggestions for the criteria to determine your initial use cases:

- **Start simple:** Look for less complex end-to-end use cases that will allow you to leverage the technology while also getting more familiar with the capabilities and use cases.
- **High level of manual effort:** Use cases that have high levels of manual processes are excellent candidates to transform the process and/or provide significant payback for the elimination of the manual efforts. This can also improve quality by reducing the number of errors.
- **Deterministic use cases:** Although agentic AI can support a more probabilistic approach to AI, you may want to leverage a more deterministic, rules-based use cases as you get comfortable with the power of the agentic AI. This may minimize the power of the agents but will also allow you to get familiar with the technology.
- **Non-regulated processes:** From a risk mitigation perspective, targeting non-regulated processes as you begin your journey into the world of agentic AI will help avoid regulatory issues around, for example, patient safety or product quality.
- **Leverage existing automation:** Agents will need the assistance of robots, humans and other agents so look to leverage your existing automation investments to help improve the autonomy and intelligence of the agents.
- **Human in the loop:** Agents should have an escalation option for instances where the agent may not understand how to proceed. Allowing human oversight will ensure appropriate decision making while also improving the performance of the model.

Ready to learn more about agentic automation for life sciences?

Contact us →