Deloitte.

The Life Sciences & Health Care AI Dossier

A selection of high-impact use cases

The Deloitte AI Institute™

About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect all the different dimensions of the robust, highly dynamic, and rapidly evolving Artificial Intelligence ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With™."

The Deloitte AI Institute aims to promote the dialogue and development of AI, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, the future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you are in: whether you are a board member or a C-Suite leader driving strategy for your organization—or a hands-on data scientist bringing an AI strategy to life—the Deloitte AI Institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet-ups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute

Foreword

Artificial intelligence (AI) continues to advance by leaps and bounds, delivering breathtaking capabilities once thought to be far off in the future. With a remarkable capacity to understand complex inputs and generate valuable outputs—and the rapidly emerging ability to execute real-world actions through intelligent agents and physical AI—AI is opening the door to innovations and new ways of working that were almost unthinkable just a few years ago.

As the AI landscape evolves, so does this compendium. Our latest edition features 130 of the most compelling use cases for AI across six major industries:

- **Consumer**
- **Energy, Resources & Industrials**
- **Financial Services**
- **Government & Public Services**
- **Life Sciences & Health Care**
- **Technology, Media & Telecommunications**

For each of these industries, we explore innovative uses for AI that can address enterprise challenges in new ways, expand and improve capabilities in every business function, and deliver advantages in efficiency, speed, scale, and capacity. To further provide context and clarity, each case specifies the primary business function it supports and whether agentic and/or physical AI is used. These labels are presented for informational purposes, helping you quickly grasp the intention and scope of each case.

Of course, every powerful tool presents potential risks, and AI is no exception. To help you better understand and manage the risks associated with AI, we use Deloitte's Trustworthy AI™ framework throughout this compendium to illuminate factors that contribute to trust and ethics in AI deployments, and to offer practical steps for strengthening governance and risk mitigation. The specific objective of our Trustworthy AI™ framework is to help organizations create AI systems that are (1) fair and impartial, (2) robust and reliable, (3) transparent and explainable, (4) safe and secure, (5) responsible and accountable, and (6) private.

Given AI's rapidly expanding scope and reach, this compendium offers just a glimpse of what the technology can do. Our goal is to convey what AI is currently capable of, and even more important, to inspire the next wave of AI-driven innovation. As AI technology continues to improve and organizations increasingly embrace it, we anticipate even more impressive and compelling use cases in the future—including those that have yet to be imagined.

We hope the use cases highlighted here will spark new ideas, provide a foundation for successful deployments, and set organizations on a path to harness the maximum value from this powerful new technology.

Nitin Mittal
Global Head of Generative AI
Deloitte Consulting LLP

Beena Ammanath
Global Deloitte AI Institute Leader
Deloitte Consulting LLP

Jim Rowan
US Head of AI
Deloitte Consulting LLP

The Life Sciences & Health Care AI Dossier

AI is accelerating transformation across life sciences and health care—sectors where precision, timeliness, and trust can mean the difference between life and death. From research labs to hospital wards, AI innovations (including LLMs, AI agents, and physical AI) are enabling faster discovery, more targeted treatments, and more efficient and effective delivery of care.

In life sciences, AI is streamlining drug discovery and development—helping to identify promising new molecules, optimizing clinical trial design, and identifying patient cohorts for precision therapies. AI is also being used to automate regulatory compliance and to optimize supply chain performance through improved supply and demand forecasting.

In health care, AI is improving diagnostic accuracy, personalizing treatment pathways, and easing administrative burdens. Increasingly, agentic AI and physical AI are automating complex workflows—supporting clinical decision-making, coding and billing, and even orchestrating and executing elements of care delivery with minimal human intervention.

With growing volumes of clinical, genomic, and real-world data, AI is helping organizations extract insights at a speed and scale that traditional methods cannot match. And as the technology matures, AI will not just enhance isolated functions—it will likely become integral to advancing scientific discovery, improving population health, and building more resilient, responsive health systems with ongoing human oversight.

That said, the use of AI in regulated, high-stakes environments demands rigorous validation, explainability, and alignment with ethical standards. Data privacy, equity in outcomes, and trust in automation are paramount. To succeed, companies will need to effectively integrate AI within existing care and research frameworks while navigating evolving regulatory requirements and public expectations.

AI is helping organizations extract insights at a speed and scale that traditional methods cannot match.

Note: *The tags below each use case indicate its primary business function and whether Agentic or Physical AI is used.*

Tags | Primary business function | Agentic AI | Physical AI

Multi-modal diagnosis and clinical decision support

Collaborating across data types to improve diagnostic accuracy

Agentic AI systems can act as a clinical co-pilot, with specialized agents analyzing different data sources—such as imaging, lab results, and patient records—and then combining insights to support more accurate and timely diagnoses.

ISSUE/OPPORTUNITY

Diagnosing complex conditions often requires synthesizing data from multiple modalities: radiology scans, lab tests, patient histories, genetic profiles, and physician notes. This process traditionally depends on coordination across multiple human specialists, each interpreting one slice of the evidence. Limited time, high patient volumes, and fragmented systems make it difficult to integrate all relevant data into a cohesive picture, which can delay treatment or lead to missed diagnoses.

As health systems face mounting pressure from clinician shortages and rising patient loads, the need for decision support that is both comprehensive and efficient is urgent. Multi-agent AI provides an opportunity to automate routine analysis across diverse data sources, highlight patterns that humans might overlook, and generate integrated recommendations that help physicians make faster and more informed decisions.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Specialized data analysis
Specialized AI agents can apply deep expertise in focused areas. For example, imaging agents can flag abnormalities on scans, NLP agents can parse physician notes and pathology reports, and other agents can interpret lab results and genetic markers.

Continuous learning from records
Agents can draw on the electronic health record as common context and shared memory, adapting recommendations over time based on a patient's history and treatment responses (in addition to population-level outcomes).

Contextual integration
A planning or coordination agent can link evidence across modalities—for example, connecting a lung nodule on a CT scan with smoking history and lab findings—before proposing next steps such as a biopsy or targeted treatment.

Workflow automation
By triaging scans, pre-annotating records, and uncovering key insights, agents can reduce the time physicians spend on routine analysis, allowing them to focus on patient interactions and issues requiring complex human judgment.

Multi-modal diagnosis and clinical decision support

MANAGING RISK AND PROMOTING TRUST

Safe and secure
Errors in diagnosis can be life-threatening. As such, agents should be deployed with strong safeguards, including rigorous validation of outputs and controlled simulation testing before integration into live workflows.

Transparent and explainable
Because clinical decisions carry major consequences, agents should provide transparent reasoning, including references to underlying data. Validator agents can be used to audit and cross-check other agents' conclusions.

Private
Medical data is highly sensitive, which means that agent systems must follow strict access controls, comply with HIPAA, GDPR, and similar regulations, and use privacy-preserving techniques to minimize unnecessary data exposure.

Responsible and accountable
Since liability in clinical settings is complex and risky, agent recommendations should initially be deployed under physician supervision, with final accountability for diagnosis and treatment decisions retained by human clinicians.

POTENTIAL BENEFITS

Faster, more accurate diagnoses
Automated integration of multimodal data helps generate insights earlier and reduces diagnostic errors, enabling quicker, more effective interventions.

Reduced clinician workload
By offloading burdensome, data-heavy tasks such as image triage and report synthesis, agents give physicians more time for patient care and complex decision-making.

Improved patient outcomes
Earlier detection and personalized treatment plans can lead to better survival rates, reduced complications, and higher overall quality of care.

Hyper-personalized health care

Delivering tailored, continuous care through multi-agent collaboration

Agentic AI systems can act as a 24/7 virtual care team, with specialized agents monitoring patient data, providing personalized coaching, and coordinating interventions to deliver hyper-personalized health care at scale.

ISSUE/OPPORTUNITY

Managing chronic conditions such as diabetes, hypertension, or heart disease requires continuous monitoring, timely adjustments, and patient engagement outside of clinical visits. Human care managers and clinicians often have limited bandwidth to track individual patients in real time, leaving care gaps where small issues can escalate into serious and costly hospitalizations or complications. Patients, meanwhile, often struggle to follow generalized care plans that don't adapt to their lifestyles, personal preferences, or daily health fluctuations. The result is higher readmission rates, preventable emergency visits, and reduced patient satisfaction and health outcomes. Multi-agent AI offers a way to make health care more proactive and personalized, addressing issues early and potentially scaling up to provide individualized support to large populations

Tags: **Customer Experience** | **Agentic AI**

HOW AI CAN HELP

Continuous patient monitoring
A monitoring agent can analyze real-time data from wearable devices, glucometers, or other sensors, detecting subtle patterns such as recurring glucose spikes or irregular heart rhythms.

Personalized coaching
A coaching agent can engage patients directly through apps or chatbots, offering tailored advice, timely reminders, and personal encouragement that reflects their specific health needs and behaviors.

Evidence-based insights
A knowledge agent can stay updated with medical guidelines and population-level data, interpreting broader patient patterns and suggesting personalized, evidence-based adjustments to care.

Automated, coordinated care
Agents can schedule follow-ups, suggest medication adjustments (with clinician approval), offer healthy recipes, and even auto-order groceries—all orchestrated through a shared patient profile.

Hyper-personalized health care

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	Biased training data could lead to less effective guidance for certain demographic groups. To minimize problems and avoid reinforcing existing disparities, AI agents should be trained and validated on diverse patient datasets.
Private	Because agentic AI systems continuously collect and analyze sensitive patient information, they must provide strict privacy protections and fully comply with data regulations such as HIPAA and GDPR. This includes obtaining clear patient consent.
Transparent and explainable	Patients and clinicians could initially be skeptical of following AI-driven guidance. To help build trust, AI agents should provide understandable explanations for their recommendations, including specific references to the data or guidelines that informed them.
Responsible and accountable	Clinical decisions can have major safety implications. As such, AI agent outputs should be reviewed by care teams where appropriate, with final accountability for treatments and adjustments retained by licensed clinicians.

POTENTIAL BENEFITS

Improved patient outcomes at lower cost

Timely, personalized interventions can help improve patient adherence to treatments, reduce complications, prevent costly emergency visits, and enable earlier corrections when managing chronic diseases. More broadly, agentic AI systems can boost operational efficiency and clinician productivity.

Scalable personalized care

AI agents can enable clinicians to provide individualized care to thousands of patients simultaneously, increasing reach and helping to address workforce shortages.

Higher patient engagement and satisfaction

AI can help patients feel supported between clinical visits, improving adherence to treatment plans and fostering trust in their care providers.

Automated customer service

Improving health plan member interactions through multi-agent systems

Agentic AI systems can handle a wide variety of customer service tasks—such as benefit inquiries, claims status updates, and enrollment support—seamlessly escalating complex cases to human representatives.

ISSUE/OPPORTUNITY

Health plans face large volumes of inquiries from members and prospective members across multiple channels, ranging from routine coverage questions to complex claims issues. Traditional call centers and support teams often struggle with long wait times, inconsistent service quality, and high staffing costs, all of which can frustrate members and erode customer satisfaction.

In today's competitive marketplace, a poor service experience can increase churn, while inefficient processes can drive up operating costs through the roof. Health plans need a way to deliver fast, accurate, and personalized support without continuously scaling human staffing levels, particularly during high-volume periods such as open enrollment.

Tags: **Customer Service** | **Agentic AI**

HOW AI CAN HELP

Specialized AI agents
Individual agents can handle specific tasks such as answering benefit questions, processing claims inquiries, and assisting with new member onboarding, providing responses that are timely, accurate, and consistent.

Seamless collaboration
AI agents can share context across tasks, improving their service quality and allowing for smooth escalation to human representatives when needed, without requiring members to repeat their information.

Dynamic prioritization
Requests can be triaged based on urgency, member status, or complexity, ensuring that high-priority cases receive timely attention.

Proactive engagement
AI agents can initiate reminders or follow-ups, such as prompting prospective members with enrollment information or reminding existing members of required documentation.

Automated customer service

MANAGING RISK AND PROMOTING TRUST

Private	Health plan support requires handling sensitive personal and medical information. Agentic AI systems must be designed with strict data protection protocols and full compliance with regulations like HIPAA.
Robust and reliable	Since members rely on accurate information to make critical financial and health decisions, AI-generated responses should be validated against authoritative data sources and continuously monitored for accuracy.
Responsible and accountable	Health-related interactions often require a strong dose of empathy and humanity. AI agents should be implemented as a complement to human staff, with clear escalation paths to ensure that members always have access to human support when needed.

POTENTIAL BENEFITS

Improved member experience

Faster, more accurate, and more consistent responses across channels help improve the overall customer experience, build member trust and satisfaction—and reduce churn.

Increased operational efficiency

Automating routine tasks frees up human staff to focus on complex or high-priority cases, boosting effectiveness while at the same time improving productivity and efficiency.

Enhanced scalability

Agentic AI can help health plans maintain consistent levels of service and responsiveness in times of fluctuating inquiry volumes (e.g., during open enrollment) without proportionally increasing staffing.

Smarter clinical trials

Improving trial design, recruitment, and monitoring with multi-agent systems

Agentic AI systems can optimize clinical trials by designing protocols, selecting sites, identifying participants, and monitoring progress in real time, helping trials run faster, safer, and more efficiently.

ISSUE/OPPORTUNITY

Clinical trials are one of the most expensive and resource-intensive stages of drug and therapy development, often slowed by challenges such as inefficient site selection, poor patient recruitment from narrow patient pools that don't fully reflect the broader population, and frequent protocol amendments. These inefficiencies can lead to delays, higher costs, and a higher risk of trial failure—ultimately impacting the speed-to-market and efficacy of new treatments.

In the face of growing therapeutic pipelines and increasing pressure to reduce time-to-market, life sciences organizations need tools that can improve trial design and execution while maintaining compliance with strict regulatory standards.

Tags: R&D/Product Development | Agentic AI

HOW AI CAN HELP

Optimized site and protocol design

AI agents can analyze hospital capabilities, historical trial data, and patient access to recommend optimal trial sites, appropriate sample sizes, and robust protocols that balance scientific rigor with operational feasibility.

Data integration and harmonization

By combining data from electronic health records, registries, and wearable devices, agents can create a holistic, real-time view of trial operations, ensuring decisions are made with the most current and complete information.

Adaptive trial management

Agents can continuously monitor trial progress and outcomes, dynamically adjusting resource allocation or protocol details to keep trials on track and compliant.

Patient stratification and recruitment

Specialized agents can identify and stratify eligible patient populations, improving recruitment efficiency while reducing delays caused by insufficient enrollment.

Smarter clinical trials

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	AI models should be trained and validated using demographically broad datasets to avoid systemic bias in recruitment strategies. Representation audits can help ensure no population segments are inadvertently favored or excluded.
Private	Since clinical trials depend on sensitive patient data, agents must be designed with strong privacy protections, data governance controls, and full compliance with regulations such as HIPAA and GDPR.
Robust and reliable	Trial outcomes directly impact patient safety and therapeutic approval. AI-generated recommendations need to be validated against clinical expertise and continuously refined through real-world feedback.
Responsible and accountable	Over-reliance on automation can lead humans to become complacent and miss important clinical nuances. AI agents should be treated as decision-support tools with physicians and trial managers maintaining final accountability.

POTENTIAL BENEFITS

Lower trial costs and improved trial success rates

Improved trial design and adaptive management reduce delays and inefficiencies, driving down the high costs associated with clinical development. Also, real-time recruitment monitoring increases the likelihood of trial completion.

More accurate trial results

By including a broader and more demographically representative participant pool, clinical trials are more likely to produce results that reflect the real-world effectiveness of a drug across different populations.

Accelerated time-to-market

Smarter site selection, faster recruitment, and fewer protocol amendments shorten trial timelines, enabling faster delivery of new therapies to the marketplace.

Better commercial outcomes

Drugs that are tested and proven effective across a wide range of demographic groups are more likely to gain broad clinical adoption, which increases their commercial potential.

Dynamic inventory management

Using coordinated AI agents to optimize supply chains and inventory

Agentic AI systems can transform supply chain and inventory management through specialized agents that help optimize inventory levels and supply chain performance while reducing costs.

ISSUE/OPPORTUNITY

Managing inventory in life sciences and pharmaceuticals is particularly challenging due to complex supply chains, strict regulatory requirements, and unpredictable demand for critical products. Traditional inventory management often relies on fixed rules and periodic actions, which can lead to costly overstocks, stockouts that disrupt patient care, and inefficiencies in production and procurement.

In addition to regulatory and operational complexity, companies must also navigate factors such as geographic variations in demand, perishability of products, and external disruptions like transport delays or weather events. These challenges can result in shortages with serious health implications, or excess inventory that drives up costs and waste.To remain competitive and compliant, organizations need inventory management systems that can respond dynamically to real-time changes in demand, supplier performance, and production capacity—minimizing waste and ensuring that essential therapies and products are always available.

Tags: **Procurement/Sourcing & Supply Chain** | **Agentic AI**

HOW AI CAN HELP

Demand forecasting agents
AI agents can analyze historical sales data, market trends, and external signals to help predict demand fluctuations and proactively adjust inventory levels and plans. AI can also integrate local geographical characteristics, disease prevalence, and socioeconomic factors to generate highly accurate, micro-market-specific demand forecasts.

Supplier and procurement agents
Agents can monitor supplier performance, lead times, and procurement activities, ensuring purchasing decisions align with both short-term needs and longer-term forecasts.

Production and scheduling agents
AI can optimize production schedules, identifying bottlenecks and aligning output with available resources and expected demand. By leveraging cross-silo data and advanced analytics, AI can identify patterns and trends that traditional methods might miss, helping to mitigate product shortages.

Distribution and logistics agents
Agents can monitor distribution flows, detect inefficiencies, and propose adjustments in real time to help ensure inventory is positioned where it is most needed.

Coordinated orchestration
An orchestration layer confirms that actions across procurement, production, and logistics remain aligned, enabling holistic, end-to-end inventory optimization.

Dynamic inventory management

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Poor forecasts or system errors can trigger significant supply chain disruptions. AI agents should be validated against human planners and real-world performance, with continuous monitoring to ensure accuracy and stability.

Transparent and explainable

Inventory decisions can have a major impact on costs, compliance, and patient access. As such, agent-driven decisions should be supported by clear rationales that supply chain professionals can understand and verify.

Responsible and accountable

Automated systems can overlook subtle operational factors that a human expert might consider crucial. Conversely, over-reliance on automation can make human experts complacent. Human oversight of AI is essential, with supply chain professionals empowered to override or adjust agent recommendations as needed.

POTENTIAL BENEFITS

Reduced costs and improved financial performance

Dynamic, data-driven inventory planning powered by agentic AI can lower inventory holding costs, reduce excess stock, and minimize spoilage of perishable products.

Improved service levels

Real-time monitoring and coordinated decision-making reduces the risk of stockouts, ensuring patients and providers receive products when needed.

Greater supply chain agility

Organizations can respond more quickly to disruptions or demand spikes, maintaining resilience in the face of changing market or regulatory conditions.

End-to-end autonomous drug discovery and development

Accelerating drug discovery through collaborative AI agents

Agentic AI systems can drive the drug discovery and development process—from molecular design to clinical trials—using specialized agents that work together to accelerate innovation and reduce costs.

ISSUE/OPPORTUNITY

Drug discovery and development is one of the most expensive and time-consuming challenges in life sciences, often costing billions of dollars and spanning many years or even decades. Researchers must manually generate candidate molecules, predict their properties and efficacy, assess toxicity, and evaluate manufacturability (often across disconnected systems and siloed teams). This fragmented approach slows innovation, drives up costs, and increases the risk of late-stage failures.

With growing demand for faster therapeutic innovation and mounting R&D pressures, life sciences organizations need ways to both increase the throughput of candidate discovery and improve confidence in early-stage selections. Multi-agent AI offers an opportunity to streamline molecule design, integrate diverse data sources, and automate iterative refinement of promising compounds.

Tags: R&D/Product Development | Agentic AI

HOW AI CAN HELP

Agent-driven design of novel compounds
Specialized AI agents can generate novel chemical structures, simulate compound interactions, and predict pharmacological properties, providing a broad and high-quality candidate pool. AI can also propose modifications to known molecules in early-stage drug development, rapidly evaluating their therapeutic potential and feasibility.

Multi-modal integration
AI systems can incorporate data from different domains—including chemical, biological, and safety data—creating a holistic view of each compound's potential and reducing the risk of pursuing weak candidates.

Collaborative refinement
Toxicity, manufacturability, and regulatory-focused agents can evaluate compounds in parallel, sharing data with molecule generation and simulation agents to iteratively improve designs. An orchestration agent can oversee and manage the specialized agents, integrate outputs, and prioritize workflows, facilitating seamless progression from initial ideation to candidate selection.

Clinical trials and regulatory submissions
Agents can design preclinical studies and trial protocols based on predicted outcomes. Also, agents can prepare submission-ready documentation, track compliance requirements across regions, and adapt protocols to evolving regulatory standards.

End-to-end autonomous drug discovery and development

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Early design errors can lead to costly downstream failures. AI outputs should be validated against human benchmarks and continuously monitored.
Transparent and explainable	Regulatory approval requires meeting strict standards for safety, efficacy, and traceability. Agent reasoning and decision paths should support explainable outputs, including transparency into the data and assumptions used.
Responsible and accountable	Drug discovery is a complex challenge with life-or-death implications. AI agents should be deployed as co-pilots to human scientists, with humans retaining accountability for final go/no-go decisions.

POTENTIAL BENEFITS

Accelerated time-to-market

Automated generation and evaluation of compounds can significantly shorten early-stage development timelines—from years to months—enabling rapid responses to emerging diseases and other unmet needs.

Lower R&D costs

By rapidly identifying and prioritizing candidates that meet key property constraints, AI can reduce development costs and improve the quality of compounds advancing to later stages.

Higher likelihood of success

Comprehensive, multi-modal analysis and real-time adaptation improve the chances that selected candidates will advance successfully through clinical trials. The results could include reduce late-stage attrition, improved ROI, and timely delivery of safer, more effective therapies.

A co-writer for appeals

Denial appeal letters

AI can be used to draft denial appeal letters, drawing from patient records and medical policies and guidelines in a faster, more cost-effective way.

ISSUE/OPPORTUNITY

When a medical insurance claim is denied, hospital billing staff face a costly and lengthy process of reviewing patient records and medical policies to create an appeal letter. For US hospitals, appeals-related administrative costs are measured in billions of dollars. Part of the challenge is the amount of time required for staff to compile an appeal. While more than 60% of denied claims are recoverable, vague reasons for denial and limited hospital billing resources result in only 0.2% of in-network claims being appealed, with millions of dollars written off as uncollectible loss each year.[1]

Tags: **Operations**

HOW AI CAN HELP

Retrieving policies and guidelines
An AI-based retrieval model can reach across large volumes of medical policies and member plans to identify the necessary information for a claims appeal.

Extracting patient data
Using extractive algorithms, the organization can rapidly consult unstructured medical notes, medications, lab results, and other electronic health records.

Writing the appeal
With the necessary information gathered by AI, an LLM can be used to generate an appeal letter.

A co-writer for appeals

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

When consulting highly detailed guidelines, policies, and records to appeal a claim denied for vague reasons, AI models may misinterpret the denial or the records, leading to an unsuccessful appeal. Ultimately, a human needs to be accountable for validating appeal letters.

Private

By drawing from electronic health records, the model is consuming health information whose protection is subject to laws and regulations. Organizations must ensure that the data ingested and information outputted aligns with data protection and patient privacy expectations.

POTENTIAL BENEFITS

Reclaim revenue

Automating the denial appeal process can supplement hospital billing resources, leading to more denial appeals filed and potentially more revenue recovered.

Efficiency improvement

Compared to traditional manual methods, AI technologies can enhance the speed and efficiency of appeals activities, such as substantiating claims and drafting appeals. Their potential to streamline processes for both simple and complex cases can make the legal workflow more timely and cost-effective.

Faster admin for payers, providers, and patients

Accelerated prior authorization

Using AI to analyze medical policies, guidelines, and provider-submitted information about underlying issues, patient needs, and medical history, an organization can automate a Prior Authorization submission (Provider) or generate a Prior Authorization approval or denial (Payer).

ISSUE/OPPORTUNITY

The Prior Authorization process is manual and labor-intensive for both health care payers and providers. The process requires the input of coders who understand the intent of a payer's Prior Authorization policies, as well as the need for medically necessary care management plans. The time required to analyze medical records and policies to make determinations on Prior Authorization submission, approval, or denial can lead to a long administrative process between the payer and provider, which can negatively impact patient satisfaction and the customer experience.

Tags: **Operations**

HOW AI CAN HELP

Supporting the provider
For providers, AI can help prepare a Prior Authorization submission by analyzing submission requirements and guidelines and cross-referencing with a patient's medical records to ensure necessary requirements are met. AI can then aid in submission to the payer and continually learn which best practices tend to lead to Prior Authorization approvals.

Supporting the payer
For payers, AI can help reduce the time required to make a Prior Authorization decision, impacting the patient experience. It also helps mitigate fraud by determining if there are anomalies in a provider's coding practices and supports compliance by analyzing submitted Prior Authorization requests and records against the payer's policies and procedures.

More efficient operations
For both payers and providers, using AI can reduce work burdens and streamlines the ability to handle Prior Authorizations, which can reduce costs while improving patient experiences.

Faster admin for payers, providers, and patients

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Prior Authorization requires the provider and payer to communicate sensitive patient data, such as protected health information (PHI) and personally identifiable information (PII), etc., which means this data is exposed to the model. Risks include unauthorized third-party access, as well as AI systems inadvertently revealing sensitive information during the generation process, thus compromising patient data confidentiality.

Fair and impartial

The process for submitting and responding to a Prior Authorization involves a standard set of rules and the patient's medical history, which introduces the potential for bias in AI models. This bias might arise from the historical data used to train the model (e.g., disparities in health care treatment or outcomes), which could lead the AI model to inadvertently perpetuate and even amplify such biases by making its own biased decisions or recommendations. The use of standardized authorization rules and patient-specific medical history, alongside continuous monitoring and careful evaluation, helps mitigate this risk and promotes fairer outcomes.

Robust and reliable

While the process for submitting and responding to a Prior Authorization revolves around a standard set of rules and the patient's medical history, there is a risk the model will misinterpret nuanced medical conditions of underrepresented populations that were not in the training dataset, and thus falsely deny the need.

POTENTIAL BENEFITS

Speed and efficiency

With AI, providers and payers may require less time to understand policies, research patient medical records for compliance, and generate, approve, or deny a Prior Authorization request.

Continuous learning

An AI feedback loop refers to the cyclical process whereby the AI model's output is presented to users or evaluators for feedback, which is then used to iteratively update and refine the model. This enhances the consistency and quality of outputs, enables providers to gain a deeper understanding of payer policies, streamlines decision-making processes, and helps payers optimize their procedures.

Improved patient experience

As the Prior Authorization process becomes more efficient, patients can receive the care management they need without needless waiting for administrative processes to conclude. This can increase patient satisfaction by improving the administrative and patient experience.

Simplifying claims submission

Medical coding

AI can be used to create codes for a claims department to categorize incoming claims and billing for medical services and procedures, which can improve the accuracy, efficiency, and speed of the claims process.

ISSUE/OPPORTUNITY

The claims submission process in the medical industry can be laborious and error-prone, involving the manual categorization of a large volume of incoming claims with complex medical codes. This time-consuming task leads to backlogs, delays, and potential payment issues for health care providers.

Tags: **Operations**

HOW AI CAN HELP

Transformed claims processing
Using AI to help categorize incoming claims and analyze and assign accurate codes can improve the overall accuracy, efficiency, and speed of claims processing. This can result in faster reimbursements for providers and a streamlined experience for both the claims department and patients.

Reduced labor burden
By leveraging AI, human effort in the claims submission process can be redirected to higher value tasks, which could result in administrative cost savings for the payer.

Simplifying claims submission

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

A Large Language Model (LLM) used in medical billing may be susceptible to bias from skewed training data, incorrect labels, and under-represented cases, potentially leading to incorrect claim categorization. To mitigate these issues, careful data collection, diverse model testing, and continuous monitoring and adjustment are vital for ensuring fair and accurate performance.

Private

To assess coding accuracy, the LLM compares the billed codes with the patient medical history, which exposes the patient's data to the underlying model and creates potential privacy risks that need to be mitigated.

Robust and reliable

Medical coding is highly regulated with strict penalties for over/under coding. The accuracy and reliability of LLM outputs in this regard is essential, as mistakes could have serious consequences. Reliability may be challenging in part because patient medical history may contain multiple modalities (e.g., text, images, and video).

POTENTIAL BENEFITS

Accuracy to limit revenue loss

Leveraging an LLM can help reduce the risk of coding errors. This can increase billing accuracy and decrease revenue loss due to errors.

Time efficiency

Automating the review of medical records can save valuable time for health care practitioners, enabling them to focus on more meaningful work.

A physician's message manager

Provider in-basket management

An LLM can be used to process messages in a health care provider's in-basket, accelerating responses while enabling physicians to focus on patient-facing care.

ISSUE/OPPORTUNITY

The amount of time required for primary care providers (PCPs) to accomplish both administrative and patient care responsibilities can exceed what is possible in a day. In some cases, upwards of two-thirds of their time is spent on administrative, non-patient facing work.

The 21st Century Cures Act encourages electronic medical records (EMR) in-basket usage, which led to a dramatic increase in in-basket messages during the COVID-19 pandemic. The result is a significant burden on PCPs, which is contributing to physician burnout.

Tags **Customer Service**

HOW AI CAN HELP

Triaging the in-basket
AI can be used to review routine messages (e.g., Rx refills, scheduling) and delegate simpler tasks to automation.

Message assistant
PCPs can leverage AI to summarize complex clinical messages for review and use the model to draft replies for provider input and response. AI models consult prior in-basket replies and EHR data when creating drafts.

Insights at scale
By using AI-enabled in-basket message systems at scale, organizations can identify issues related to patient negativity in their messages. The insights into complaints, expressions of dissatisfaction, frustration, confusion, or concern about care can inform interventions that may improve the patient experience.

A physician's message manager

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Use of AI for in-basket systems involves collecting, processing, and storing large amounts of sensitive patient data, such as medical history, diagnoses, and treatment plans. This data is subject to strict privacy laws, and any unauthorized third-party access could result in legal and financial consequences for health care providers.

Responsible and accountable

If messages are composed or summarized with inaccurate information, it could lead the PCP to erroneous decision-making or poor patient engagement, which can have significant consequences for patient health, trust in the health care provider, and the reputation of the organization.

POTENTIAL BENEFITS

Physician support

By using an AI-enabled in-basket system, the PCP's time-consuming administrative tasks are reduced, permitting more patient-facing work and mitigating one cause of physician burnout.

Timely responses

A more efficient process for working through in-basket messages can lead to faster responses to patient needs, contributing not only to a better patient experience but potentially also better health outcomes.

Patient sentiment

By identifying and tracking signals of negativity at scale, health care providers can gain insights into common pain points in the patient experience. This could help them proactively address these issues, whether by adjusting their practices, improving communication, or implementing other interventions to enhance patient satisfaction.

Democratizing AI model creation

Knowledge domain model development

Generative AI can be used to improve existing AI models by removing user interface (UI) hurdles through reinforcement learning (RL) without the need for technical staff.

ISSUE/OPPORTUNITY

Developing novel AI models for life sciences and health care organizations continues to demand a high degree of technical proficiency to perform data exploration, feature engineering, model training, and evaluation. Frequently, the steps involved in model training lack a user-friendly interface, posing accessibility challenges for health care professionals and domain experts who may not possess extensive technical backgrounds. Simultaneously, the quality and relevance of model outputs hinges significantly on domain expertise and practical experience. Overcoming this divide between technical acumen and domain knowledge remains an obstacle in harnessing the complete capabilities of AI within the field of life sciences and health care.

Tags | R&D/Product Development

HOW AI CAN HELP

Empowering professionals
With its capacity for learning from and adapting to iterative feedback, generative AI can act as an enabler for professionals across various sectors. It offers the opportunity to continually refine existing domain-specific AI models by adding new training data. This iterative enhancement increases the model's accuracy, utility, and relevance to the user's specific professional needs. In this way, generative AI can empower professionals by providing them with tailored, precision AI tools that evolve with their work.

Streamlining health care model development
Generative AI can help simplify model development in the complex and highly regulated health care industry. By focusing on intuitive user interface designs and automated processes, generative AI minimizes UI obstacles, making it more accessible for professionals to refine and improve their existing models.

Improving alignment
Generative AI leverages reinforcement learning (RL) techniques, a type of machine learning where an AI system learns to make decisions by trial and error, to validate and improve its own outputs. This process assists in mitigating prevalent AI challenges, including hallucinations or confabulations, ambiguity, and colloquialism misuse. As a result, it bolsters AI's reliability and furnishes professionals with more precise models and predictions, thus aligning AI capabilities more closely with user requirements.

Democratizing AI model creation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Hallucinations or confabulation could lead to the execution of incorrect procedures or use of suboptimal reagents and equipment, causing inaccurate experiments and inefficient use of resources. Particularly in medical or pharmaceutical labs, inaccurate information could even lead to compliance or regulatory issues.

Transparent and explainable

The generative AI system incorporates tools that offer transparency into data engineering pipelines, including data preparation stages. This inherent transparency facilitates an understanding of the AI model's functioning within the organization, fostering trust in the accuracy and reliability of the AI system's outputs. It is a crucial component of the AI use case, demonstrating the system's accountability and promoting its acceptance across the organization.

POTENTIAL BENEFITS

Enhance institutional knowledge access

AI can help reduce institutional knowledge loss due to employee exits and enable on-demand access to domain-specific knowledge across the organization.

Increase development throughput

Domain area experts can drive more self-sufficient model experimentation and development by utilizing natural language model outputs and synthesizing insights about optimal procedures, reagents, equipment, and techniques into a comprehensive and accessible format.

Cost management

This democratic approach to AI model development empowers employees to take part in model experimentation, reducing costs associated with machine learning operations and technical specialists.

Optimizing lab procedures

Experimental design

AI can be used to create procedural templates and recommendations on best practices (e.g., reagents, equipment, techniques).

ISSUE/OPPORTUNITY

Laboratory personnel, including researchers, technicians, and managers, often face challenges in maintaining up-to-date procedural templates and maintaining the consistent application of best practices, especially as scientific knowledge evolves rapidly. These challenges can lead to inefficiencies, errors, and inconsistency in experiments or analyses. Additionally, without a central source of curated recommendations, time and resources may be wasted sourcing and comparing various reagents, equipment, and techniques. These pain points present an opportunity for AI to streamline and enhance laboratory processes.

Tags | **R&D/Product Development**

HOW AI CAN HELP

Generation of novel processes
Leveraging historical data and scientific principles, an AI model could suggest novel experimental designs, more efficient processes, or alternate uses of reagents and equipment, stimulating innovation in laboratory procedures.

Data analysis and interpretation
AI uses a large language model (LLM) to analyze data from lab protocols, equipment specifications, previous experimental designs, reagent usage, and techniques, providing a holistic understanding of laboratory procedures and principles.

Optimizing lab procedures

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

The integration of multimodal text and images of complex structures and processes in experimental design presents complexity. This can heighten the risk of unworkable, unfeasible, or inefficient designs, as interpreting and accurately representing this diverse and intricate data can be challenging. These challenges could potentially lead to errors in the design and execution of experiments, resulting in failed or less reliable outcomes and unnecessary time and resource expenditure.

Responsible and accountable

In the event of erroneous design recommendations, accountability may be an issue. Determining who bears the responsibility for incorrect designs and their potential consequences is important. The roles of human oversight and system validation need to be clearly defined.

Transparent and explainable

With the application of AI in experimental design, there may be challenges related to explainability. If scientific or academic papers are to be published based on the results, authors need to be able to adequately explain the methodology behind the AI-recommended designs, which can be inherently complex due to the black-box nature of some AI models.

POTENTIAL BENEFITS

Efficiency

LLMs can reduce the time and effort needed for experimental design by streamlining and accelerating data analysis and procedure consolidation, and by providing best practice recommendations.

Lower cost

With less time required for experimental design, organizations can reduce the overall operational costs of experiments while also increasing throughput.

Revealing the rules

Automated regulatory compliance

AI can be used to support compliance by processing large amounts of regulatory documents from multiple geographies.

ISSUE/OPPORTUNITY

Compliance with ever-changing regulations in every geography is a costly, time-consuming process for pharmaceutical companies. Even with significant investment in legal help, regulatory compliance can be hard to achieve.

Tags **Compliance & Risk**

HOW AI CAN HELP

Text processing
AI can be used to extract regulations for a specific purpose from thousands of pages of regulatory texts, enabling compliance.

Transforming the legal support ecosystem
As AI handles the laborious, detail-oriented process of regulatory text processing, it can also lead to a commensurate decrease in the need for third-party legal and compliance support.

Mitigating financial risk
By employing AI in regulatory compliance, the potential financial risk associated with non-compliance can be reduced.

Revealing the rules

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable	AI models may produce outputs that are hard to interpret, making it difficult to validate them and explain the reasoning to regulatory authorities.
Private	While regulatory authorities may vary, data privacy around personal health information remains a priority, and data that is not anonymized first may leak and become inappropriately disclosed.
Robust and reliable	An AI model trained to extract compliance factors from regulatory documents may be susceptible to outputting information that looks accurate but is a hallucination, making human validation an important element for mitigating risks around reliability.

POTENTIAL BENEFITS

Cost reduction

Using AI to process regulatory documents reduces the need for humans to perform time-consuming tasks, potentially lowering the cost of compliance.

Fuel for growth

When regulatory compliance becomes tractable across geographies because of AI processing capabilities, it helps the organization confidently expand business operations globally.

20/20 impurity detection

AI-driven visual inspection for particulate matter in IV fluids

Computer vision powered by AI can be used to detect particulate contamination in IV bags, reducing product waste and improving patient safety in life sciences manufacturing.

ISSUE/OPPORTUNITY

Pharmaceutical manufacturers, particularly those producing IV fluids and life-saving therapies, face a persistent and costly challenge: detecting particulate matter in sterile products. Despite sterile manufacturing environments, small particles—such as plastic, dust, or other foreign materials—can still enter IV bags, posing serious health risks to patients.

This issue is not new—dating back to the 1940s—and despite ongoing improvements, a scalable, reliable, and cost-efficient solution has remained elusive. Historically, detection has relied on manual inspection, often using contingent labor, leading to inconsistent results, high labor costs, and significant product waste. Also, every incident of contamination risks brand reputation, regulatory scrutiny, and potential product recalls.

Tags | **Manufacturing & Quality**

HOW AI CAN HELP

Automated, real-time inspection
By combining AI vision capabilities with generative AI models trained on synthetic and real-world data, manufacturers can automate the inspection process at scale. High-resolution imaging and computer vision detect anomalies in fluid packaging with greater precision and consistency than the human eye.

Continuous learning
AI enhances the system by learning from historical defect data, adapting to new defect types, and identifying potential causes through pattern recognition across vast datasets. The AI system not only flags potential contamination in real time but also enables traceability—helping identify root causes by analyzing patterns across manufacturing lines, geographies, or specific production lots. This insight enables proactive correction and long-term process improvements.

20/20 impurity detection

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

AI models should undergo rigorous testing across multiple manufacturing lines and environments to help ensure high accuracy and minimal false positives/negatives. Redundancy checks, human-in-the-loop validation, and performance monitoring help ensure reliable operation even under variable lighting or packaging conditions.

Transparent and explainable

AI-based contamination detection can provide clear, trackable results, allowing manufacturers to understand why a product passes or fails inspection. Detailed imaging and reporting help ensure accountability and regulatory compliance while enhancing confidence in quality control.

POTENTIAL BENEFITS

Improved patient safety and confidence

Consistent detection of contaminants before they enter the supply chain increases patient safety, and fewer quality incidents enhance brand trust with hospitals, regulators, and patients.

Operational efficiency and scalability

Replacing manual inspection with AI reduces reliance on contingent labor and speeds up quality control processes. Also, AI-based solutions can be deployed at scale across multiple products, manufacturing lines, and facilities worldwide.

Improved ROI and reduced costs/waste

AI helps drive measurable ROI through improved efficiency, reduced labor costs, and avoidance of costly recalls. Also, fewer discarded bags due to false positives or late-stage detection leads to significant material savings.

Quality and traceability

Root cause analysis helps address upstream issues in the manufacturing line, improving overall process quality.

Transforming pharmaceutical research & development

Physical AI smart labs for drug discovery

Physical AI-powered "smart labs" autonomously execute the entire design-make-test-analyze (DMTA) cycle in drug discovery. These systems design molecules using AI-driven computational models ("in silico"), generate synthesis plans, and directly execute them through integrated robotic lab platforms.

These systems orchestrate execution, physically synthesize molecules, transfer samples via automated handling systems, and run assays using connected lab instruments, analyze results, and initiate the next iteration—with minimal human intervention beyond oversight and exception handling.

ISSUE/OPPORTUNITY

Pharmaceutical discovery is constrained by slow, manual DMTA cycles that take 8–12 weeks per iteration, limiting the number of molecules that can be tested and delaying time-to-market for new therapies. Traditional discovery requires manual setup of experiments, physical sample movement, and fragmented instrument workflows—creating bottlenecks and increasing error risk in repetitive tasks. Capital-constrained biotech companies face existential time pressure but lack the workforce to scale discovery operations. Physical AI and robotics offer an opportunity to automate up to 70–80% of standardized processes, accelerating cycle times and enabling faster kill decisions on failing candidates while progressing winners. Autonomous systems can work continuously, handle greater experimental loads, and improve portfolio NPV by compressing discovery timelines.

Tags: R&D/Product Development | Physical AI

HOW AI CAN HELP

AI-driven molecule design & synthesis

AI algorithms design novel molecules in silico, predict synthetic routes, and generate execution plans including reagent plating, liquid handling sequences, and robotic orchestration.

AI-optimized experimental design

Machine-learning models predict which experiments will yield the most information, reducing unnecessary data generation by up to 50% and maximizing throughput from existing lab capacity.

Autonomous sample transport

AMRs transport samples between synthesis, purification, and testing stations. Automated assay platforms run efficacy and safety screens, upload results directly to AI systems, and close the loop for next-cycle decision-making.

Real-time data analysis & adaptive iteration

AI monitors experimental outcomes in real time, evaluates performance against objectives, identifies failures early, and autonomously designs the next DMTA cycle—compressing iteration time from weeks to days.

Robotic synthesis & purification

Robotic arms and automated synthesis platforms execute AI-generated protocols autonomously, performing liquid handling, reaction setup, purification, and quality control (QC) without human intervention, operating 24/7 when facilities permit.

Transforming pharmaceutical research & development

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Silent error propagation is the most distinctive reliability risk here. Errors in robotic synthesis, sample handling, or AI iteration decisions can compound across multiple cycles before detection, advancing flawed candidates and discarding viable ones. This corrupts the scientific foundation on which regulatory submissions rest. Reliability must be validated across the full automated workflow, not just individual steps.

Responsible and accountable

Regulatory authorities expect documented human accountability for scientific decisions. A system operating with minimal human intervention must have an exceptionally robust audit trail to satisfy that expectation. Organizations must clearly document the boundary between AI-generated recommendations and human-validated judgments at every stage, making full traceability of experimental decisions a deployment prerequisite, not a backward-looking governance addition.

Transparent and explainable

Regulatory accountability requires that qualified scientists can assess and stand behind the AI's scientific reasoning—not just validate its outputs. When AI makes decisions to drop candidates or advances molecules based on predicted properties, the reasoning must be traceable to a standard that satisfies both internal scientific review and external regulatory examination.

POTENTIAL BENEFITS

Radical cycle time compression
Autonomous smart labs reduce DMTA cycle times from 8–12 weeks to under 2 weeks, enabling faster progression of viable drug candidates, earlier termination of failures, and significant improvement in portfolio Net Present Value (NPV).

Scalable innovation infrastructure
Pharma companies can scale discovery operations by adding parallel robotic workstations, creating a path to continuous, factory-style drug discovery without workforce constraints.

Increased discovery capacity
By automating 70–80% of standardized workflows, labs can handle larger portfolios with the same physical infrastructure and workforce, focusing human effort on the 20% of artisanal, non-standardizable tasks.

Capital efficiency for biotech
Small biotech companies can deploy smart labs to replace large discovery teams, reducing headcount requirements while achieving higher throughput—critical for one-product companies racing against cash burn.

Enhanced reliability & reproducibility
Robotic systems eliminate human error in repetitive tasks, ensure consistent execution of protocols, and capture tacit operator knowledge that would otherwise be lost, improving data quality and experimental reproducibility.

Automated drug dispensing systems

AI-powered robotic systems for accurate and scalable medication management

AI-powered robotic pharmacy systems automate medication dispensing using physical AI and robotics. Prescription dispensing robots handle counting, packaging, and labeling, while prescription filling robots accurately fill vials or blister packs. Automated storage and retrieval systems organize inventory and enable rapid access. Combined with AI algorithms and computer vision, these systems ensure accurate verification, dosage calculation, and low-error, continuous pharmacy operations with minimal manual intervention.

ISSUE/OPPORTUNITY

Manual pharmacy workflows rely on repetitive tasks such as counting, filling, and inventory handling, which increase errors and slow down operations. These processes are difficult to scale while maintaining accuracy and consistency, leading to delays and higher workload for pharmacists. Physical AI-driven robotic systems offer an opportunity to automate dispensing, filling, and storage, reducing errors, improving efficiency, and enabling scalable, high-throughput pharmacy operations.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

AI-driven prescription processing

AI validates prescriptions, ensures accurate data verification, and reduces manual review effort to speed up processing.

Computer vision verification

Ensure correct medication identification and labeling while detecting discrepancies to prevent dispensing errors.

Robotic dispensing systems

Automate counting, packaging, and labeling of medications while ensuring precision and consistency in repetitive tasks.

Automated storage & retrieval

Manage inventory and enable fast medication access while optimizing stock organization and reducing retrieval time.

Real-time safety checks

Detect drug interactions and patient risks while supporting informed decision-making for safer dispensing.

Robotic filling systems

Fill vials and blister packs with high precision, reducing variability and improving dosage accuracy.

Automated drug dispensing systems

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Medication dispensing errors cause direct patient harm—and a system that introduces its own error modes in place of human ones does not improve patient safety; it substitutes one risk for another. Reliability must be validated continuously across the full range of medications and dosage forms handled, not just the products well-represented during design and testing.

Safe and secure

A compromised dispensing system that alters dosage instructions or suppresses drug interaction checks is not an IT problem—it is literally a matter of life and death. Security must therefore be treated as a patient safety issue, not an IT governance issue, with protections reflecting the direct clinical consequences of system compromise or manipulation of prescription data and dispensing logic.

Responsible and accountable

Pharmacist accountability for final dispensing decisions is a regulatory and professional requirement that AI assistance cannot displace. The governance design—including which checks are AI-assisted and which require human sign-off—must be explicit, enforced, and documented. This is the dimension that ultimately determines whether the system is deployable in a regulated healthcare environment.

POTENTIAL BENEFITS

Reduced errors & improved safety
Minimize manual errors in dispensing and filling while enhancing patient safety and reducing adverse events.

Higher workforce efficiency
Reduce repetitive workload for pharmacy staff while enabling greater focus on patient care activities.

Scalable 24/7 operations
Support continuous, high-volume dispensing while enabling pharmacies to handle increased demand efficiently.

Faster processing time
Accelerate prescription handling and delivery while reducing patient wait times and improving service levels.

Consistent & accurate operations
Ensure standardized and reliable processes while improving quality control and reproducibility.

Optimized inventory management
Improve stock organization and retrieval speed while reducing shortages and enhancing inventory visibility.

Surgical robotics for microsurgery

AI-assisted precision surgery

High-precision surgical robots/robotic hands support microsurgical procedures under direct human supervision. These systems feature extremely low latency, high accuracy, and tight integration between AI perception and physical actuation during live operations.

ISSUE/OPPORTUNITY

Microsurgery demands precision beyond normal human motor capability, while safety requirements leave no tolerance for error. Manual procedures limit consistency and accessibility. Delicate operations such as nerve repair, vascular reconstruction, or ophthalmic surgery require movements measured in fractions of millimeters, performed while the surgeon's hands naturally experience tremor and fatigue. Surgical outcomes vary based on individual surgeon dexterity, experience level, and physical condition, creating inconsistency in patient results for identical procedures.

Complex microsurgical techniques remain accessible only to specialists at major medical centers, limiting patient access based on geography and surgeon availability.

The opportunity is to use physical AI to enhance surgical precision while keeping clinicians fully in control. Leveraging robotic hands enables direct translation of surgeon intent into controlled end effector motions at the point of contact, reducing tremor at the source and enabling precise force control.

Tags: **Field Services** | **Physical AI**

HOW AI CAN HELP

High-resolution perception

AI interprets visual and sensor data during surgery, enhancing the surgeon's view with magnification, filtering, and highlighting of critical anatomical structures that guide precise intervention.

Latency-constrained execution

Decisions and actions occur in real time with minimal delay between surgeon input and robotic response, maintaining the natural feel of direct tissue interaction critical for surgical judgment.

Precision motion control

Robotic hand executes movements at sub-millimeter accuracy, filtering out hand tremor and scaling surgeon hand movements to finer instrument movements for delicate tissue manipulation.

Safety-critical integration

AI operates within strict procedural constraints including movement boundaries, force limits, and anatomical safety zones that prevent inadvertent damage to surrounding structures.

Human-controlled operation

Surgeons retain authority over all actions, with robotic systems translating their commands into precise physical movements rather than making autonomous surgical decisions.

Surgical robotics for microsurgery

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

A latency spike or motion control failure during nerve repair or vascular reconstruction is not a recoverable error—it's a matter of life and death. Validation must reflect actual surgical conditions, not laboratory benchmarks, and must demonstrate performance across the full range of procedures, tissue types, and edge cases the system might encounter in live use.

Safe and secure

The safety envelope—movement boundaries, force limits, and anatomical safety zones—is the primary mechanism protecting patients during live procedures. Any failure that causes the system to act outside that envelope, whether from software error, sensor malfunction, or external compromise, has direct and potentially irreversible consequences. Fail-safe design is not a feature; it is the fundamental condition for clinical deployment.

Responsible and accountable

Surgeon accountability for patient outcomes is both a clinical and a legal requirement that the system's design must actively support. Logs of all commands, system states, and force feedback throughout each procedure are the evidentiary record on which that accountability rests. Without a complete and reliable audit trail, the governance principle ("surgeon retains authority") applies in theory but cannot be demonstrated in practice.

POTENTIAL BENEFITS

Improved surgical outcomes

The robotic hand enables finer tissue handling and more accurate suture placement, reducing trauma and complication rates and improving healing times for microsurgical cases.

Consistency

End effector level actuation and sensing from the robotic hand deliver uniform precision regardless of surgeon fatigue or tremor, reducing variability in outcomes.

Patient benefit

Recovery outcomes improve through less tissue trauma, reduced scarring, shorter hospital stays, and lower complication rates that translate into better long-term function.

Expanded treatment capability

More procedures become feasible as robotic precision enables operations that exceed manual capabilities, bringing advanced microsurgical techniques to cases previously considered too complex.

Service robots for hospital operations

Autonomous execution of routine logistics

Robots perform repetitive service tasks in hospitals, such as transporting medications, meals, linens, and supplies. These systems navigate human environments and interact with staff and patients. Robots physically move items through corridors, elevators, and clinical areas while making real time route and task decisions.

ISSUE/OPPORTUNITY

Hospitals face labor shortages and high turnover in repetitive service roles. Manual task execution consumes skilled staff time. Nurses and porters spend significant portions of shifts transporting medications, lab samples, linens, and supplies between nursing stations, pharmacies, laboratories, and patient rooms.

The opportunity is to deploy AMRs to automate routine point to point logistics, freeing clinical staff for direct patient care. On robot physical AI enables resilient local decision making (obstacle avoidance, dynamic re routing, queued handoffs) that preserves service continuity during peak demand or partial network outages.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

Indoor navigation

Robots move safely in human spaces, navigating hallways, elevators, and doorways while avoiding obstacles like patients, visitors, equipment, and staff moving through dynamic hospital environments.

Human interaction handling

Robots operate around people, yielding right-of-way, communicating arrival through alerts or displays, and responding appropriately when staff or patients need to access items or clear pathways.

Task execution

Deliveries are completed autonomously as robots transport items between designated locations, confirm delivery through digital handoffs, and return for the next assignment without human intervention.

Operational reliability

Tasks are executed consistently with predictable delivery times, proper handling of temperature-sensitive medications or fragile items, and automated documentation of completed deliveries.

Fleet coordination

Multiple robots operate together, sharing elevator access, coordinating routes to avoid congestion, and balancing workload distribution to maintain service levels during peak demand periods.

Service robots for hospital operations

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Hospital service robots transporting medications, lab samples, and temperature-sensitive supplies must perform reliably across a live clinical environment that includes crowded corridors, elevator queues, and potential network outages. Delivery failures or unexpected stops can disrupt care workflows, delay medication administration, and create congestion in spaces where patient safety depends on clear access.

Safe and secure

Autonomous robots in hospitals must respond safely to unpredictable human behavior, including patients with impaired mobility, children, and individuals in distress. Physical collisions in clinical spaces could cause injury or disrupt care. Safety boundaries must be validated across the full diversity of people and conditions encountered in live hospital environments, not just controlled testing environments.

Responsible and accountable

When a robot fails to deliver medication on time, delivers to the wrong location, or causes a care disruption, accountability must be clearly allocated between the AI developer, manufacturer, integrator, and hospital operator. Delivery logs and operational records must be sufficient to support investigations and demonstrate compliance with hospital governance and patient care requirements.

POTENTIAL BENEFITS

Labor relief
Staff focus on higher-value work; for example, nurses can spend more time on patient assessment and care rather than transport and other routine service tasks.

Service consistency
Tasks are completed reliably with uniform timing and quality regardless of shift staffing levels, employee experience, or competing demands on human workers' attention.

Scalability
Operations expand without staffing increases as additional robots handle growing delivery volumes during facility expansions and workload peaks without proportional labor costs.

Operational efficiency
Routine tasks are automated, eliminating delays from staff unavailability, reducing delivery errors from miscommunication, and maintaining consistent service during peak periods.

Eldercare and memory care support systems

Addressing critical care gaps through embodied AI

Humanoid robots deployed in memory care and eldercare settings provide companionship, medication management, and emotional support for individuals with Alzheimer's, dementia, and age-related care needs.

ISSUE/OPPORTUNITY

Aging populations are pushing care needs beyond available caregiving capacity, leaving facilities and families struggling with staffing shortages, medication adherence, and meaningful engagement for patients with cognitive decline. Overburdened caregivers—both professional and familial—cannot provide continuous supervision or companionship, while isolation and medication errors accelerate health risks.

Humanoid AI assistants offer a path forward, they provide consistent, always available support for routine monitoring, medication reminders, and engagement—while keeping caregivers firmly in control of all clinical and judgment intensive decisions. By extending care continuity without replacing human empathy or accountability, Physical AI helps stabilize care quality, reduce caregiver burden, and improve patient well being at scale.

Tags: **Customer Experience** | **Physical AI**

HOW AI CAN HELP

Emotional connection & companionship
Humanoid form factor enables elderly individuals to personify and connect with the AI assistant, leveraging the innate human tendency to bond with physical objects. The robot engages in conversations on topics the patient raises, providing stimulation and reducing isolation.

Medication compliance & health monitoring
AI monitors medication schedules and reminds patients to take prescribed medications at appropriate times. System tracks compliance and can alert caregivers to missed doses or concerning patterns.

Governed intervention
Escalation to caregivers when thresholds are crossed. No autonomous clinical decision making.

Consistent, reliable presence
Unlike human caregivers managing multiple patients or working shifts, the AI assistant provides continuous availability. Patients benefit from predictable, patient interaction without fatigue or irritability.

Eldercare and memory care support systems

MANAGING RISK AND PROMOTING TRUST

Private

Robots monitoring elderly patients with cognitive decline collect highly sensitive data—medication adherence, behavioral patterns, and indicators of deterioration—about individuals who may lack capacity to consent. Strict governance is essential, with meaningful consent processes involving family members or legal guardians and clear limits on data access, retention, and secondary use.

Safe and secure

Humanoid robots in close proximity to elderly patients with cognitive decline and impaired mobility must behave predictably under all conditions—including when patients become confused, agitated, or interact physically in unexpected ways. Safety boundaries must be validated specifically for memory care populations, whose responses cannot be assumed to mirror those of other user groups.

Responsible and accountable

Caregivers must retain full clinical accountability. In practice, this requires reliable escalation protocols that alert human caregivers when thresholds are crossed, and governance frameworks ensuring families and care facilities understand precisely what the robot does autonomously versus what triggers human involvement.

Transparent and explainable

Families and care facilities deploying robots with cognitively impaired patients must understand what the system does, what data it collects, and how it escalates concerns—since patients themselves may lack capacity to assess these things. Transparency obligations fall on deploying organizations, not just on individual users, given the vulnerability of the population being served.

POTENTIAL BENEFITS

Addresses eldercare labor shortages
Augments overstretched caregiving workforce by handling routine monitoring, companionship, and medication management tasks. Enables human caregivers to focus on tasks requiring judgment and empathy.

Reduces caregiver burden
Family members and professional caregivers gain peace of mind from continuous patient monitoring. System handles routine reminders and engagement, reducing stress on human caregivers.

Improved medication adherence
Timely reminders and compliance tracking reduce missed doses and escalation risk, with alerts routed to caregivers when needed.

Improves patient outcomes
Consistent monitoring and medication compliance support reduce adverse health events. Emotional engagement and mental stimulation may slow cognitive decline for dementia patients.

Cleanroom manufacturing automation

Precision operations in sterile environments

Robotic systems perform material handling, assembly, inspection, and localized environmental monitoring in pharmaceutical and biologics cleanroom spaces, reducing human presence in sterile zones while maintaining strict contamination control and validated quality standards.

ISSUE/OPPORTUNITY

Pharmaceutical and biologics manufacturing requires sterile production environments where contamination control is critical and quality failures carry massive costs. Cleanroom operations are labor-intensive, requiring workers to don extensive protective equipment, follow rigorous gowning procedures, and work in uncomfortable controlled environments with restricted airflow and temperature regulation. Human presence introduces contamination risk despite protective protocols, as people generate particulates through movement, breathing, and skin shedding—even when fully covered.

Quality failures in sterile manufacturing trigger costly batch rejections that can reach millions of dollars, regulatory investigations that threaten facility operating licenses, and potential patient safety issues that create liability exposure.

Current physical AI systems lack cleanroom certifications and ISO collaborative robotics standards for sterile environments, with updated standards expected in Q3 2028.

These automated cleanroom applications—especially post-certification—can deliver substantial value in high-margin sterile production settings where quality and consistency are paramount and where reducing human presence directly reduces contamination risk while maintaining production capacity.

Tags: **Manufacturing & Quality** | **Physical AI**

HOW AI CAN HELP

Certified cleanroom robotics

Purpose-designed systems that meet cleanroom classification standards perform material handling, assembly, and inspection tasks in sterile zones without introducing particulate or biological contamination, using specialized materials and designs validated for controlled environments.

AI-driven quality monitoring

Advanced vision systems with closed-loop process control detect anomalies in real-time and automatically adjust parameters to maintain sterile processing conditions and product quality, identifying deviations before they cause batch failures.

Collaborative sterile operations

Physical AI in cleanroom manufacturing is optimized for minimal human presence to reduce contamination risk. Over time, evolving safety and cleanroom standards may enable limited, governed human robot collaboration for exception handling and complex tasks—but this remains secondary to automation first sterile operations.

Cleanroom manufacturing automation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Robotic systems in pharmaceutical cleanrooms must perform reliably within validated sterile environments where any deviation—a dropped component, a missed quality check, or an uncontrolled movement—can contaminate a batch and trigger regulatory investigation. Reliability standards must meet pharmaceutical validation requirements, not just general industrial robotics benchmarks.

Responsible and accountable

Pharmaceutical manufacturing is one of the most heavily regulated production environments in the world. The audit trail requirement is not just good governance—it is a regulatory obligation. When AI contributes to a quality failure, the evidentiary record must be sufficient to satisfy FDA or equivalent scrutiny.

Transparent and explainable

Quality assurance teams and regulatory inspectors reviewing AI-driven inspection decisions need to understand what anomalies were detected, how severity was assessed, and what actions were taken—both to exercise oversight and demonstrate compliance. Opaque AI quality decisions that cannot be explained to a regulatory inspector are not compatible with pharmaceutical manufacturing standards.

POTENTIAL BENEFITS

Contamination risk reduction

Minimizing human presence in sterile zones reduces particulate generation and biological contamination while maintaining consistent environmental controls and product quality.

Regulatory compliance support

Automated documentation, process traceability, and consistent adherence to validated procedures strengthen an operation's regulatory compliance posture in heavily scrutinized pharmaceutical manufacturing.

Enhanced product consistency

Robotic systems eliminate human variability in repetitive tasks, delivering more consistent results in high-value sterile production where quality failures trigger costly batch rejections and regulatory scrutiny.

Contacts

Beena Ammanath
Global Deloitte AI Institute Leader, Deloitte AI Institute United States, Lead
Deloitte Consulting LLP
bammanath@deloitte.com

Audrey Ancion
Deloitte AI Institute Canada, Lead
Deloitte Canada
knuttall@deloitte.com

Naser Bakhshi
Deloitte AI Institute Netherlands, Lead
Deloitte Netherlands
nbakhshi@deloitte.nl

Dr. Bjoern Bringmann
Deloitte AI Institute Germany, Lead
Deloitte Germany
bbringmann@deloitte.com

Martin Cabrera
Deloitte AI Institute Chile, Lead
Deloitte
mcabreraa@deloitte.com

Jefferson Denti
Deloitte AI Institute Brazil, Lead
Deloitte Brazil
jdenti@deloitte.com

Tiago Durao
Deloitte AI Institute Portugal, Lead
Deloitte Portugal
tdurao@deloitte.com

Richard Eudes
Deloitte AI Institute France, Lead
Deloitte France
reudes@deloitte.fr

Roman Fan
Deloitte AI Institute China, Lead
Deloitte China
rfan@deloitte.com

Tomas Meca Figueres
Deloitte AI Institute Spain, Lead
Deloitte Spain
tomecafigueras@deloitte.es

Alfredo Maria Garibaldi
Deloitte AI Institute Italy, Lead
Deloitte Italy
agaribaldi@deloitte.it

Nicolas Griedlich
Deloitte AI Institute Luxembourg, Lead
Deloitte Luxembourg
ngriedlich@deloitte.lu

Geert Hallemeesch
Deloitte AI Institute Belgium, Lead
Deloitte Belguim
ghallemeesch@deloitte.com

Jan Hejtmanek
Deloitte AI Institute Central Europe, Lead
Deloitte Central Europe
jhejtmanek@deloitte.com

Prashanth Kaddi
Deloitte AI Institute India, Lead
Deloitte India
pkaddi@deloitte.com

Sultanbek Khunkaev
Deloitte AI Institute Middle East, Lead
Deloitte United Arab Emirates
sukhunkaev@deloitte.com

Tomotake Kozu
Deloitte AI Institute Japan, Lead
Deloitte Japan
tomotake.kozu@tohmatsu.co.jp

Carlos Labanda
Deloitte AI Institute South-Latin America, Lead
Deloitte Colombia
clabanda@deloitte.com

Chris Lewin
Deloitte AI Institute Asia Pacific Lead
Deloitte Singapore
chrislewin@deloitte.com

Wessel Oosthuizen
Deloitte AI Institute Africa, Lead
Deloitte Africa
woosthuizen@deloitte.com

Sulabh Soral
Deloitte AI Institute United Kingdom, Lead
Deloitte United Kingdom
ssoral@deloitte.com

Dr. Elea Wurth
Deloitte AI Institute Australia, Lead
Deloitte Australia
ewurth@deloitte.com.au

www.deloitte.com/us/AIInstitute

Deloitte.

Endnotes

1. Karen Pollitz, Justin Lo, Rayna Wallace, Salam Mengistu, "Claims Denials and Appeals in ACA Marketplace Plans in 2021," KFF, February 2023.