

WHITE PAPER

AI in Healthcare: Core Principles, Real Impact, and What's Next

Table of Contents

Executive Summary 3

Inovalon’s AI Principles..... 7

Real-World Examples of AI Driving ROI and Better Outcomes 8

Foundational Pillars of AI in Healthcare..... 11

Regulatory and Ethical Considerations 15

Exploring the Future 17

Conclusion..... 20

Executive Summary

Artificial intelligence (AI) is poised to disrupt the healthcare industry like no other technology. While it won't replace human experts, AI is transforming how care is delivered – impacting the roles of doctors, nurses, pharmacists, and all other care providers, as well as the teams who support them, from payers and life sciences teams to policymakers.

AI is now permanently part of the healthcare conversation. In the last few years alone, “AI” went from a futuristic concept to table stakes; it's now a necessary cog in the care continuum, on par in importance with electronic health records (EHRs) and revenue cycle management (RCM) platforms, which AI is already optimizing. However, it's important to understand the mechanics of AI, the guardrails to keep in mind, and the delicate balance of innovating quickly while maintaining an appropriate level of caution, keeping humans in the loop and at the heart of healthcare. Of course, the patient's role amid AI's rising influence must be considered as well. As AI's presence in healthcare increases, patient perception ranges from enthusiasm to skepticism.

Meanwhile, industry professionals are in serious need of greater efficiency and better data analytics to inform key decisions. One-third of physicians are likely to leave their jobs in the next 5 years,¹ citing administrative burden and staff burnout as a major driver for this departure. AI can help; it holds the promise to greatly improve healthcare outcomes and economics throughout the patient journey – and Inovalon is delivering models to help teams across the care continuum capitalize on the potential of AI in a responsible, ethical way. Our AI capabilities are trained on the industry's largest primary source healthcare dataset, with experts in healthcare and technology working closely together to validate use cases, train models, and continuously refine the models.

We are committed to producing highly impactful AI tools, each one with a sharp focus on data quality and interoperability, privacy and security, ethical and accountable AI governance, and keeping a human in the loop across use cases.

Disclaimer: Our AI capabilities are intended to support human experts in their respective work. Each model must be used with a human-in-the-loop working in partnership with the AI; the model is not a replacement for a professional's experience.

1. “The physician shortage isn't going anywhere,” McKinsey & Company, Laura Medford-Davis, Rupal Malani, September 10, 2024, <https://www.mckinsey.com/industries/healthcare/our-insights/the-physician-shortage-isnt-going-anywhere>

Understanding AI and related technologies

As AI in healthcare continues to advance, our industry's understanding of this technology and related capabilities is varied. For clarity:

AI refers to the capability of technology to perform tasks that typically require human judgment — such as reasoning, recognizing patterns, or making decisions. It is the umbrella concept under which Machine Learning (ML) and Natural Language Processing (NLP) fall.


ML is a subset of AI focused on teaching systems to learn from data and improve performance over time without being explicitly programmed. In healthcare, ML powers predictions and automation, such as forecasting claims denials or identifying gaps in care.

NLP is a specialized branch of large language models that focuses on understanding and generating human language. It incorporates many ML techniques, and is uniquely suited to handle the text-rich and documentation-heavy environment of healthcare, such as summarizing physician notes or extracting diagnosis codes.

Additionally, generative AI and agentic AI are other types of AI gaining significant interest in healthcare. Generative AI can create new, original content including text, images, and code in response to a user's prompt. Its outputs are the result of the prompting guidance and the patterns the model has learned from the data it was trained on. Agentic AI is similar, but has more advanced reasoning abilities. Its purpose is to achieve specific goals with minimal human oversight, such as inferring decisions and interacting with other models or datasets with a low level of human prompting or involvement needed to complete a task (though that human presence must always remain).

How they are different:

- AI = The overarching field of “intelligent” technology.
- ML = The method AI uses to learn from large datasets and improve predictions.
- NLP = The application of AI/ML to human language, enabling systems to read, interpret, and generate text.



Key pillars of AI in healthcare

- **Data quality and interoperability:** AI is only as effective as the data that powers it.
- **Privacy and security:** The data informing AI models needs to be securely shared, connected, and deidentified - all in compliance with industry regulations such as HIPAA and HITRUST.
- **Transparent, human-in-the-loop decisions:** AI should always remain as a supporting tool, not a deciding factor. Humans retain the final say.
- **Ethical and accountable AI governance:** Appropriate guardrails are needed to uphold AI accuracy, consistency, and performance.
- **Reliability and safety at scale:** For every AI tool, the developers who built it should check accuracy, further train the model, and make course corrections as needed on an ongoing basis.

*To explore these concepts further, please see page 11.

The promise of AI in healthcare

Whatever mixture of AI technologies is utilized, the potential for positive impact – to truly transform healthcare so it works better for all people – is exciting and necessary.

The healthcare system is at a tipping point: clinician shortages, escalating administrative costs, regulatory mandates, and demand for value-based care make it imperative to adopt AI that drives efficiency, accuracy, and trust. Regulatory compliance and reimbursement models are constantly changing, causing health plans, providers, and pharmacies to adjust to the latest expectations for reporting, prior authorization, value-based care, and beyond. Meanwhile, rising costs are deterring some patients from seeking care ([36% of adults](#)²) and challenging providers to keep their doors open. [One study](#) found 4 in 10 hospitals had negative operating margins, and of those with positive margins, some are alarmingly slim – 22% of hospitals had a positive margin of less than 5%.³

2. "Americans' Challenges with Health Care Costs," KFF, Grace Sparks, Lunna Lopes, Alex Montero, Marley Presiado, Liz Hamel, July 11, 2025, <https://www.kff.org/health-costs/americans-challenges-with-health-care-costs/>

3. "What are the Implications of the 2025 Budget Reconciliation Bill for Hospitals?," KFF, Zachary Levinson, Tricia Neuman, Jamie Godwin, June 12, 2025, <https://www.kff.org/medicaid/what-are-the-implications-of-the-2025-budget-reconciliation-bill-for-hospitals/>

Additionally, clinicians are leaving practice at a growing rate. With one-third of physicians likely to leave their jobs in the next five years,¹ patients will be left with limited options, creating access issues and potential delays in care. A major driver for this large departure is administrative burden, which AI can help reduce, potentially supporting long-term retention of clinical professionals. This is one of the biggest opportunities for AI in healthcare: doctors can get more time back in their day, and patients can gain real access to their care team. The medical encounter can become one in which the physician is able to focus fully on their patient, instead of on a computer screen.

1/3 of physicians

are likely to leave their jobs in the next five years, creating access issues for patients. AI can help. Inovalon is addressing these pressures with models powered by real-world, de-identified healthcare data.

That's just one example of AI's potential influence on the care continuum, and there's much more opportunity beyond completing a single task more quickly or understanding a large amount of data efficiently. While these alone are valuable, it's the sum of these small wins happening throughout healthcare that will achieve meaningful impact – improving efficiency, allowing for more engaging, thoughtful, productive clinical encounters, and significantly enhancing the patient experience.



For payers, using AI means faster, more accurate medical reviews and prior authorization decisions, saving significant time and labor costs.



For specialty and infusion pharmacies, it means automating prior authorization submissions and dispensing workflows to improve access to therapy and reduce delays in care.



For providers, it means less time spent on administrative tasks, allowing more time focused on patient care.



For life sciences and research teams, it means finding patients for clinical trials faster and accelerating evidence generation by reducing manual data curation and enabling richer analytics at scale.

Inovalon's AI Principles

AI holds the promise to greatly improve healthcare outcomes and economics throughout the care journey, and Inovalon is exploring that potential with responsible, ethically built models. We are committed to a patient-first mission, strongly believing that AI must ultimately support human decision-makers in every use.


This mission keeps healthcare experts at the helm of every healthcare function. AI is an augmentation, not a replacement; its purpose is to support critical decisions with transparent analyses that a model can explain, that can be audited, and that are compliant with industry standards. Our models deliver evidence-backed outputs that can cite the exact data point that informed a particular finding (from a chart, claim, or ruleset, for example), enabling rapid validation by the human expert using an AI solution.

At Inovalon, we approach model development holistically. Each model is the result of a unique combination of extensive data, advanced technology, and experts with deep subject matter knowledge working to bring all the components of a model together in a focused, thoughtful manner.

Our expertise reaches across healthcare, with seasoned leaders in quality measurement, risk adjustment, clinical practice and research, and technology and AI – and we bring all their experience together when solutioning how to build and apply AI for healthcare needs. Additionally, partnerships with world-leading technology companies like AWS, Google, and Snowflake help ensure our solutions and capabilities keep up with the rapid pulse of change in the world of AI.



Inovalon's AI models deliver explainable source citations behind every output.



Real-World Examples of AI Driving ROI and Better Outcomes

Inovalon has released and is currently building a range of AI models to help organizations across healthcare achieve better care outcomes and economics.

Disclaimer: Our AI capabilities are intended to support human experts in their respective work. Each model must be used with a human-in-the-loop working in partnership with the AI; the model is not a replacement for a professional's experience.

Clinical semantics analytics for health plans

Clinical semantics models use LLMs to understand clinical text, map it to standardized vocabularies, and generate computable logic. The result is faster, more reliable extraction of valuable clinical insights pulled from unstandardized, narrative healthcare data.

For health plans, these tools are helping quality measurement teams dig through massive amounts of patient charts to improve accuracy in quality reporting and member care outcomes. Patient charts are full of structured and unstructured data, and for plans managing thousands of members, understanding each member's specific needs often feels like looking for a needle in a haystack.

Clinical semantics analytics are changing that. This means that out of 1,000 charts that may have been individually reviewed in the past, AI is now helping coders focus their time on the 100 charts, for example, that actually need their attention.

With AI taking the first pass to read through chart data, coders and clinical staff can “read” thousands of charts in a matter of seconds and focus their time acting as the verification expert to the model's findings. This solves the challenge that nurses and doctors all document differently, and sometimes in a way that's ambiguous, as well as the traditional amount of time and labor it takes to review each chart. AI can understand their notes, make assessments on the notes, and share findings with a human expert to either agree, reject, or dig deeper into that particular chart.

The AI-supported review helps meet quality measurement goals and improves risk score documentation. Similarly, for every member confirmed to have diabetes, for example, AI can assist the plan in identifying the appropriate interventions and ensuring it receives appropriate funding to help manage member conditions, and ultimately, achieve better outcomes.

Furthermore, every analysis is saved and can be reproduced; humans can later check the work of this model and the expert who verified its findings. This maintains an audit trail as the model evolves and helps ensure the tool continues to learn with human experts still involved at every iteration.

Quality and risk intelligence for health plans

As health plans manage thousands (or millions) of members – and even more charts across all members – applying AI to scan medical records drives faster, more accurate decision-making for quality measurement, risk adjustment and audit readiness, and ultimately, member care.

Now available! NLP for HCC coding

Manual chart reviews for Hierarchical Condition Category (HCC) coding accuracy are slow, costly, and often inconsistent. In traditional retrospective record reviews, a team of medical coders reviews each member chart one by one, amounting to hours of labor to help a health plan capture risk-adjustable conditions more accurately.

Using [NLP for medical record review](#) turns a tedious process into a more efficient review of the records that need the most attention. AI can determine potential risk adjustment diagnoses in a chart, identifying charts that may require the reviewer's attention and excluding those that may not.

Coming soon: AI for accurate quality measure documentation

The National Committee for Quality Assurance's (NCQA) Healthcare Effectiveness Data and Information Set (HEDIS[®]) measures health plans' performance on member care for more than 90%⁵ of U.S. health plans. HEDIS includes 90 quality measures across six categories, and measure reporting requirements can change each year as new measures or new reporting specifications are introduced.

Quality reporting influences a plan's Star Ratings, their performance incentive amounts, and member perceptions. Despite the importance of accurate reporting, health plans struggle to keep up with the demands of each season, during which plans invest a significant amount of time and labor costs having nurses perform manual chart reviews.

Fortunately, generative AI can transform the chart review process to ease the burden on nurse reviewers, decrease the time reviews take, and improve accuracy – all to support effective quality measurement while lowering costs for the health plan. AI-enhanced solutions can evaluate

5. "Healthcare Effectiveness Data and Information Set (HEDIS)," U.S. Department of Health and Human Services, Office of Disease Prevention and Health Promotion, Accessed September 9, 2025 <https://odphp.health.gov/healthypeople/objectives-and-data/data-sources-and-methods/data-sources/healthcare-effectiveness-data-and-information-set-hedis>

thousands of charts for the nurse reviewer and highlight areas that need attention, keeping a human in the loop to ultimately decide if a HEDIS-related preventative service occurred or if there is exclusionary evidence to rule out eligibility for a particular intervention. With AI guiding reviewers to the measures that need the most attention, manual effort is directed to the records that need additional documentation or a correction, rather than reading every record word for word.

In development: Predictive claims optimization for healthcare providers

Nearly 1 in 5 healthcare claims is denied⁶ and 87% of all denials are attributed to front-end workflows⁷, making denials a top priority for providers across care settings and RCM roles. From CFOs to billing managers, claims denials keep many revenue cycle professionals up at night.

Unlike legacy claim scrubbers, Inovalon's model predicts denials across payers, flagging likely causes before submission. Our predictive claims optimization helps providers overcome [the challenge of denials](#). It goes beyond being a claims scrubber and is broader in its analysis by considering various payers and procedures, while also focusing on the specific mistakes on a claim that may cause a denial. This supervised machine learning model identifies patterns in an organization's denied claims to predict the likelihood of future denials. It then predicts a potential denial before a new claim is submitted and guides billers in correcting the claim – which can lead to reduced denials, higher first-pass acceptance, and faster reimbursements.

The model does this in real time, based on an organization's unique claims. It can account for nuances in care settings – claims for emergency care are much different than those for a routine annual checkup – and keeps up with payer rules. It enables faster, more proactive denials management by helping prevent denials before they occur, and with guided correction prompts, this model saves billers time in correcting the mistakes on a claim prior to submission.

6. "HealthCare.gov Insurers Denied Nearly 1 in 5 In-Network Claims in 2023, but Information About Reasons is Limited in Public Data," KFF, January 27, 2025, <https://www.kff.org/private-insurance/healthcare-gov-insurers-denied-nearly-1-in-5-in-network-claims-in-2023-but-information-about-reasons-is-limited-in-public-data/>

7. "Claims Denials: Causes, Impact, and AI-Enabled Prevention Strategies," Inovalon, October 2024, <https://www.inovalon.com/resource/claims-denials-causes-impact-and-ai-enabled-prevention-strategies/>

Foundational Pillars of AI in Healthcare

As exciting as the growing use of AI is, there are also serious considerations that come with every potential new model. The following pillars of AI in healthcare help ensure thoughtful, appropriate development and use of AI by the organizations bringing models to life and the individuals using AI-powered solutions.

Data quality and interoperability

Artificial intelligence in healthcare is only as effective as the quality, consistency, and completeness of the data it relies on. And in an industry where patient safety and clinical decision-making are at stake, the bar for accuracy must be higher than anywhere else.

This reality is underscored by the fact that fewer than [5% of medical large language models](#)⁸ are trained on real-world healthcare data, and of all healthcare data, 97% of data is unused because it's unstructured⁹. Such a low proportion highlights both the challenge and the opportunity: without clean, standardized, and representative inputs, AI cannot deliver reliable, fair outcomes – which is of particular concern given the rise of model collapse, a yet unsolved threat to AI models that are trained on non-real-world data, rather, on AI-generated content. Inovalon's enormous proprietary dataset – the industry's largest real-world dataset – allows us to avoid the threat of model collapse in the AI tools we deliver, therefore empowering those who use our models with confidence in their solution's accuracy and reliability. Ensuring high-quality data requires robust governance policies, automated validation checks, adherence to standards such as FHIR® and CQL, and careful management of master data.

Additionally, interoperability remains one of the industry's most persistent challenges. Healthcare has long operated within siloed systems that collect vast amounts of data but rarely share it in ways that are timely or actionable. The result is redundant

Less than 5% of medical LLMs are trained on real-world healthcare data. Inovalon's AI models are trained by our primary source data, the industry's largest primary source real-world dataset. Our data allows us to introduce more models trained on real-world data to the healthcare ecosystem, to collectively garner greater insights from data and offset the threat of model collapse.

8. "Testing and Evaluation of Health Care Applications of Large Language Models: A Systematic Review," Suhana Bedi, Yutong Liu, Lucy Orr-Ewing, JAMA Network, October 15, 2024, <https://jamanetwork.com/journals/jama/fullarticle/2825147>

administrative work, delayed decisions, and frustration for patients and their care teams. AI can only reach its potential if it is applied within a connected, transparent data ecosystem and with legislation in place to help move the needle. Both the Centers for Medicare & Medicaid Services (CMS) and the [Office of the National Coordinator for Health IT](#) (ONC) are now requiring interoperability and accessibility between patients and their data, as well as across different data sources – via the [Interoperability and Patient Access Final Regulation](#) and the [Interoperability and Information Blocking Final Regulation](#), respectively.

The implication is clear. Data quality and interoperability are not simply technical issues relegated to IT teams; they are foundational enablers of AI's promise. Without them, AI will only add to the noise of disconnected systems and inconsistent insights. But when data flows securely and seamlessly across the care continuum, AI will become a powerful force for reducing administrative burden, accelerating decision-making, and giving clinicians more time to focus on their patients – all by tapping into the massive amounts of unstructured data going unused in healthcare, and making sense of it at scale via specific, purpose-built models.

Privacy and security

Protecting patient data is paramount in the development and deployment of any AI model. In healthcare, where data carries not only operational, but also deeply personal significance, privacy and security extend far beyond the technical step of de-identification. Safeguards must be embedded throughout the entire lifecycle of AI – from data ingestion and model training to deployment and ongoing monitoring – to ensure compliance, maintain trust, and safeguard sensitive information.

Ultimately, the success of AI in healthcare depends on trust. Payers, providers, pharmacies, life sciences teams, and the patients they serve will only embrace AI solutions if they are confident that privacy is preserved and security is uncompromising. **Safeguarding sensitive data is not a compliance checkbox; it is a strategic foundation on which all other benefits of AI – from administrative efficiency to improved patient outcomes – depend.**

Transparent, human-in-the-loop decisions

No matter the use case or care setting, AI should always be used as a tool, not a deciding factor. It augments clinical judgement; humans retain the final say. At Inovalon, the expert “human-in-the-loop” is an imperative. AI informs healthcare professionals in their decision-making, and they can always decide whether to accept or decline the technology's suggestions. Inovalon's AI systems provide transparent context on how their assessment was determined. Models are built to show their work – to surface the data that informed a particular recommendation so that the human expert has all necessary context readily available.

These provisions increase the speed and efficiency of a decision, while preserving the influence of a human expert's judgement. For example, in clinical rounds, multidisciplinary teams review patient presentations and associated data. Each team member has a unique perspective and expertise to share. Agentic solutions are in development that may soon participate in such teams, though ultimately, it will be up to the experts to decide whether to take one of the agent's suggestions or a different course of action.

Ethical and accountable AI governance

As AI systems continue to advance, the guardrails that uphold accuracy, consistency, and performance should also evolve. These guardrails form the basis of an effective AI governance program, ensuring that innovation moves forward in lockstep with clinical, ethical, and organizational values.

Ethical governance requires a clear framework that defines how AI is built, evaluated, and maintained. At its core, this framework should outline roles and responsibilities, oversight bodies, usage policies, and escalation paths.



Roles and responsibilities establish accountability, clarifying who owns each step of the process from data preparation to model development, validation, deployment, and monitoring.



Oversight bodies provide cross-functional review and decision-making authority, ensuring that compliance, ethics, clinical safety, and business goals are all weighed before an AI system is released. They ensure each model follows ethical principles, is clinically relevant, and is built with definition-of-done checklists, model/data cards, red-team gates, and auditability designed for HIPAA, HITRUST, and specific function audits that may be required.



Usage policies set boundaries for where and how AI can be applied, giving clinicians, payers, pharmacists, and research teams clear guidance on what decisions AI may inform and which must remain entirely human-led.



Escalation paths define what happens when concerns arise, from technical errors to ethical questions, ensuring that issues are raised quickly and addressed with transparency.

Strong governance also provides the foundation for fairness. Ethical AI tools are ones that work equitably across populations and avoid reinforcing or amplifying disparities. The best defense against bias is twofold. First, models must be built on diverse, representative datasets that reflect the realities of race, gender, age, language proficiency, and socioeconomic status. Data breadth and depth determine model quality, while intentional checks for gaps in representation ensure that the outputs are not skewed. Models should also be trained to surface their own limitations, flagging missing or incomplete inputs so that human users can interpret outputs in the proper context.

Second, effective governance emphasizes the ongoing role of human experts. A human-in-the-loop is not optional; it is a necessary safeguard during both development and deployment. Diverse teams of clinicians, data scientists, and technology experts working together provide a richer, more reliable perspective than any single reviewer could achieve. Their collaboration helps ensure that AI systems are both technically robust and clinically appropriate.

Additionally, monitoring for bias and fairness is never complete. It continues throughout the lifecycle of every model, supported by continuous auditing and fairness-focused evaluation metrics to strengthen trust and safeguard against harm.

Ethical and accountable AI governance therefore creates the conditions under which AI can be safely scaled across payers, providers, pharmacies, and real-world data use cases, delivering innovation that is as trustworthy as it is transformative. Similarly, users have a responsibility to verify the models they use. From understanding how to use a new AI solution to making sure teams don't get too comfortable with AI's outputs, and always choosing the best course of action, ensuring appropriate AI use in a trust-but-verify approach is the responsibility of all healthcare professionals.

These fundamental pillars of AI in healthcare guide AI development and use at Inovalon. Additionally, we're making sure that the AI solutions we create are made within clear constraints to drive accuracy and to ensure the tools we're building are helpful.

To get there, the importance of a human-in-the-loop approach is paramount. Our vision is never to stray from our commitment to incorporate human experts in the oversight of accurate, quality-driving AI models at every stage of development, in use, and in ongoing performance enhancements of these products.

Regulatory and Ethical Considerations

While the potential for AI in healthcare is clear today, and will continue to reach new functions of the care continuum as models and adoption expand, there are fundamental regulatory and ethical considerations that cannot be overlooked.

Regulatory considerations

The adoption of AI in healthcare is inseparable from compliance with regulatory frameworks that define how data, models, and workflows must operate. These requirements establish the technical baseline for safe, reliable deployment across the care continuum. Therefore, developers should design systems with regulatory alignment in mind from the earliest stages of model development through production.

Federal guidance is central to this landscape. The Office of the National Coordinator for Health Information Technology (ONC) has outlined principles for trustworthy AI that place a strong emphasis on transparency, accountability, and reliability. Technically, this means AI systems must document the provenance of training data, provide auditable outputs, and incorporate mechanisms for human oversight. In addition, ONC guidance underscores the importance of continuous monitoring and lifecycle management, requiring developers to establish processes for validating accuracy, updating models as data shifts, and preventing unintended performance degradation.

Beyond ONC, industry expectations have crystallized around three core technical requirements.



Explainability

Models must produce outputs that can be interpreted by end users, whether through feature attribution in structured data models or source-cited evidence in natural language models.



Real-world validation

AI must prove performance not only in development environments, but also on representative, production-grade datasets that reflect the complexity of payer, provider, pharmacy, and life sciences workflows.



Bias auditing

Models are expected to undergo regular, quantitative evaluations of subgroup performance, with established methods for detection, reporting, and remediation of disparities.

Taken together, these regulatory considerations define the technical bar for AI in healthcare. They are not optional features; these are required capabilities for systems that seek to integrate into high-stakes environments such as clinical care, health plan operations, pharmacy workflows, clinical trials, and research studies. Meeting these requirements ensures that AI tools can be trusted to operate effectively while satisfying the compliance demands of healthcare's highly regulated environment.

Ethical considerations

AI can help us do so much more to serve patients. There's immense opportunity to achieve faster patient access, more time in the patient-provider encounter, and more secure information sharing to guide critical care decisions.

However, each AI model needs to be carefully managed.

Healthcare is arguably the most important industry in which ethical considerations need to be part of every AI conversation. At the end of the day, the healthcare system exists to serve human lives – and keeping all patients safe, their information secure, and delivering the timely care they need must always remain the top priority.

Keeping a human-in-the-loop as the decision point for each AI use case is only one part of responsible, ethical AI use. In addition, all healthcare professionals using AI should be transparent with patients about when and where AI tools are supporting their care journey. This is key. As AI becomes a more common aspect of everyday life, patients will either assume or find out that this technology is being used in their healthcare – and trust will be eroded if they haven't been informed up front. The best way to overcome this is for the physician – or nurse, health plan, or whomever along the care journey is interacting with the patient – to clearly state how they are using AI and why.

In time, we predict the patient perspective of AI in healthcare will be very similar to that of clinical trials. Patients are aware of trials and have the opportunity to participate in trials or opt out. Trials are a commonly understood part of the healthcare ecosystem, and both the view that they are necessary and helpful – as well as hesitations on being among the first to try a new therapy – are valid. The same will occur with AI. There will be patients who are eager to have AI support their care journey, who value the use of more advanced technology. Other patients will be cautious; they may either ask for more information or prefer to opt out. Both types of patients – and all those who are somewhere in between – must be equally informed about AI use and their options as adoption continues to grow.

Lastly, the risks and uncertainty around AI use are another necessary ethical consideration. Models will not be perfect. The risk for hallucinations and bias is real – and their potential negative impact on patient care even more so. For both newly developed models and established AI tools, a healthy level of skepticism should remain. It comes down to keeping a human-in-the-loop every time: in development, in using an AI product, and in continuing to train and improve models. Developers must continue to test and verify the tools they've created, and users can regularly audit their AI tools to keep a pulse on their accuracy and effectiveness. By keeping healthcare and technology experts in charge of AI across each use case and care setting, technology and skilled talent can work together to advance patient care.

Exploring the Future of AI

As AI in healthcare continues to evolve every day – with new models supporting outdated workflows, new versions of existing models continuing to learn, and as healthcare professionals' understanding and use of AI deepens – the future of AI is constantly taking shape. The speed, power, and adaptability of this technology is truly transformational.

Agentic AI

Conversational assistants powered by generative AI have the potential to streamline clinical and administrative workflows across care settings, from faster prior authorizations to more accurate claims documentation. An agentic AI model could be the co-pilot to a human expert, helping them progress through long manual processes with greater speed. Inovalon is building these assistants on a foundation of governed, deidentified healthcare data. Our vision is that AI would analyze, or even enter, all the relevant data points for healthcare professionals, so that they can make the most appropriate decision for their use case.

Valuable opportunities for agentic AI in healthcare include:

1. Prior authorization

Prior authorizations are needed for many procedures and radiological tests, specialty and infusion medications, and transfers of care. Yet, the prior authorization process is burdensome and frustrating for physicians, pharmacies, and patients.

Physicians and pharmacists have the administrative load of understanding what programs might be available to reduce costs and provide access for each test, medication, or transfer of care that requires prior authorization. Then, health plans process the authorization requests – and often ask for more information which essentially restarts the process as requestors have to amend their request, resubmit, and then the health plan reviews again. Meanwhile,

patients are waiting weeks for the care they urgently need, while on a broader scale, the burden of prior authorization and other administrative tasks are driving physicians, nurses, and pharmacists out of clinical practice.

[Agentic AI](#) can turn traditional prior authorization into a smoother, more effective process, and ultimately, eliminate this from being a barrier to care. It can predict what information is going to be required to get the submission approved by a health plan and prompt the physician to complete the form in that way from the time of the initial order. This helps ensure data completeness while improving speed and reducing the burden on physicians and pharmacists. On the health plan side, agentic AI can review prior authorization forms to drive approvals – saving time and labor costs in comparison to the traditional approval process.

In the event of a rejected authorization, agentic AI can offer support in the prescribing physician's office during the appeals process. It can interpret the rejection and create the first draft of an appeal letter for the physician or nurse to verify.



Potential for impact:

Patients start therapies sooner and improve treatment and medication adherence; lower pharmacist and physician burnout; and stronger payer-provider-pharmacist relationships.

2. Accurate claims documentation

Risk adjustment for health plans helps them understand their member population and that plans are paid fairly for the members they cover. However, risk adjustment heavily depends on physicians' documentation. Health plans have to review the coded conditions on member claims to ensure risk score accuracy – but what if more claims were more accurate from the point of documentation?

That's where agentic AI can greatly improve risk adjustment. With an AI-powered guide to help physicians during documentation, the healthcare professionals who are working with patients and coding their risk



Potential for impact:

Clinical professionals save time and improve documentation accuracy; health plans save labor costs in medical record review while increasing risk score accuracy across their member populations.

adjustable conditions can increase clarity and accuracy in every claim, ensuring that all validly documented diagnoses and only validly documented diagnoses make it to the claim and ultimately to submission for risk score determination.

3. Quality gap review

For health plans, quality measurement is never complete. It's an ongoing process of reporting on members' conditions and the care they've received to manage their health needs. An agentic AI tool can help health plans improve their documentation accuracy – which in turn can drive better member care outcomes at a faster rate and lower costs for the health plan, as well as better quality measurement performance and ratings. Agentic AI can proactively identify and flag documentation that doesn't meet specific numerator or exclusion criteria for quality measurement, guiding a health plan's quality team to focus on the measures and members that need the most attention.



Potential for impact:

Health plans save time and labor costs, while improving member care outcomes by better understanding member needs and planning appropriate interventions faster – which ultimately, help improve quality measurement performance and ratings.

AI-powered data de-identification for research teams, data aggregators, and beyond

In addition to the growth of agentic AI, there's strong potential to transform healthcare data de-identification, a critical process to safeguard sensitive data before it's used to inform a particular question, study, or decision point.

However, de-identification is often manual and can miss key data points, especially in clinical text, which puts PHI at risk. Effective de-identification also accounts for variances in data. For example, a patient's name may appear as "Jane Doe", "J. Doe", or "Janie Doe" – all for the same individual. De-identification requires all versions of a name to be removed, and this is just one data type; other patient identifiers and their potential variations must also be removed.

AI for data de-identification can empower teams working with large amounts of data to move faster and more confidently in the de-identification process. We envision a model that would scan both structured and unstructured healthcare data to identify and redact sensitive elements, supporting compliant data sharing at scale. It would then flag when a name is entered on a form

(structured) or written into a clinical note by hand (unstructured), informing a human expert to make the proper decision on how to manage all flagged data points.

This can greatly help teams optimize efficiency – however, it's not the final step in de-identification. While AI can reach a point of completing much of the de-identification, the final step must be third-party verification to prove proper de-identification through a secure method (such as Safe Harbor or Expert Determination).

Conclusion

AI is an essential resource to help advance the healthcare system, which is in dire need of reform. From giving overburdened physicians more time back in their day, to streamlining outdated, broken processes across the care continuum and creating a better patient experience, the impact of ethical and scalable AI in healthcare will be transformational.

AI is the latest era in a long-brewing cultural shift. We use technology every day, and as AI increasingly becomes the norm in finance, travel, and other aspects of daily life, healthcare will adapt and benefit, too. We're just beginning to truly see the impact of AI, ML, and NLP in our industry and in our personal lives as patients – and the future of AI is quickly developing before our eyes every day.

Inovalon is uniquely positioned to be at the forefront of healthcare's transformation and to empower our customers to advance their AI journeys across the care continuum. Our AI capabilities are informed by the largest primary source real-world healthcare dataset and governed by rigorous ethical frameworks – and built to embed in key healthcare workflows to deliver real-time insights that drive greater efficiency.



About us

Inovalon is a leading provider of data and solutions empowering data-driven healthcare. The Inovalon ONE® Platform brings together national-scale connectivity, real-time primary source data access, and advanced analytics into a sophisticated cloud-based platform empowering improved outcomes and economics across the healthcare ecosystem. The company's analytics and capabilities are used by nearly 50,000 active, licensed customers. For more information, visit www.inovalon.com.

Inovalon
4321 Collington Road
Bowie, MD 20716
(301) 809-4000

www.inovalon.com

Authors

The authors wish to extend our sincere gratitude to all the individuals who contributed their time and insights to this white paper. This is the result of a collaborative effort with our valued technical and clinical experts at Inovalon, who have deep experience in clinical care, quality measurement and risk adjustment, revenue cycle management, healthcare data governance and analytics, specialty pharmacy, and beyond.

Special thank you to [Courtney Breece](#), VP, Product Management; Cameron Breze, MS, Product Manager; [Kathryn Eshelman](#), M.D., MPH, VP, Medical Informatics; Scott Jubeck, SVP, Product; Kate Marta, M.D., VP, Clinical Quality, Training, Oversight, and Innovation; Liz Serie, VP, Product Management; Zach Williams, VP, Product Management.

Paige Kilian, M.D.

Dr. Kilian serves as the Chief Medical Officer (CMO) for Inovalon. In this role, Dr. Kilian is responsible for the oversight of the medical directors of the Company, the clinical content and design of the Company's software and analytics, and the clinical training, quality, policies, oversight, and compliance of associated operations. For almost two decades, Dr. Kilian has led a team of clinical personnel who bring the latest best practices and evidence-based clinical standards to Inovalon's products and services.

Dr. Kilian is well known throughout the healthcare community as an expert on the application of data, software, and analytics to the benefit of clinical protocols, care guidelines, real-world healthcare clinical workflow, and regulatory policy guidelines adherence. Dr. Kilian has been recognized with the Healthcare Business Women's Association 2020 Luminary award.

Prior to joining Inovalon in 2007, Dr. Kilian was in clinical practice with Johns Hopkins Community Physicians, where she served as Medical Director with responsibilities including patient care, clinical training of medical students and residents, and clinical oversight of a multispecialty practice.

Dr. Kilian earned her doctorate in medicine from Tufts University School of Medicine. She completed her residency in Internal Medicine at The Johns Hopkins Hospital in Baltimore, Maryland.

Rachel Stemerman, Ph.D., M.P.H.

Rachel Stemerman, Ph.D., M.P.H. is a Senior Principal Technical Product Manager at Inovalon, where she leads strategy and development for AI/ML and data engineering platforms that support payers, providers, pharmacies, and real-world data applications. She directs cross-functional initiatives developing AI/ML features to reduce administrative burden, improve accuracy of users, and ensure that advanced analytics and AI tools are designed for measurable impact across the healthcare ecosystem.

Dr. Stemerman combines her background as a practicing paramedic with expertise in health informatics and product leadership to bridge the gap between frontline care delivery and enterprise-scale technology solutions. She has authored peer-reviewed publications, delivered presentations at national conferences, and contributes thought leadership through her role as a podcast host for the Prehospital Emergency Care journal. She is also an active member of the National Association of EMS Physicians, where she contributes to advancing the conversation on the role of AI in emergency and prehospital medicine.

Dr. Stemerman holds both a Master of Public Health (M.P.H.) and a Doctor of Philosophy (Ph.D.) in Health Informatics from the University of North Carolina at Chapel Hill. She is also a former National Library of Medicine (NLM) Biomedical Informatics Fellow.