

FROM THE NAIC CONSUMER REPRESENTATIVES

To: National Association of Insurance Commissioners (NAIC)

Date: November 14, 2024

On behalf of the undersigned NAIC Consumer Representatives¹, we offer the attached report examining the increasing use of artificial intelligence (AI) by health insurance plans to augment utilization management (UM) review and approval processes. The speed of technological advances in AI is far outpacing the changes in state and federal health insurance regulation and oversight needed to protect consumers. While we believe that AI presents opportunities for plan efficiency, unregulated use of AI could also exacerbate existing bias and discrimination, particularly for marginalized and disenfranchised communities who already experience disparate health care access challenges.

Through an extensive environmental scan and a series of in-depth informant interviews, the report uncovers the areas in AI use for health insurance UM activities that pose the most risk of inappropriate or discriminatory denial of medically necessary care. Key findings from the research include the following:

- AI is regularly used by health insurance plans to conduct UM activities. Proponents cite reductions in administrative burden and expedited approvals. However, there are risks that must be considered, such as the exacerbation of existing biases, prioritization of misaligned incentives, and use of technologies outside their original use case or design leading to unintended harm.
- Relevant stakeholders that researchers spoke with, including plans, providers, technologists, regulators, and consumer advocates, see immense opportunities for the use of AI, but warn that proper safeguards need to be in place to protect patient access to needed care, and to ensure that both privacy and safety are appropriately considered.
- Some states have begun to regulate the development and use of AI in health insurance, but for the most part, have not been able to keep up with the proliferation of the use of AI itself. Regulation in this space is in the nascent stages of development and the speed at which change occurs makes this both a challenging and essential problem to solve.

Given these findings, we make the following recommendations to state regulators:

1. Meaningful transparency is critical

- It must be clear, to both regulators and consumers, when AI is being used by health insurance plans for the purposes of UM and what role the AI plays in making determinations about coverage for care.
- Transparency must extend to disclosures about the data used to develop, train, and test the AI tools (with an emphasis on consent for use and representativeness of the population), and the extent to which any AI tool can begin to train itself.

¹ The NAIC works closely with Consumer Representatives to assist state regulators in their primary objective of protecting insurance consumers. The Consumer Representatives do not represent or work for the NAIC, and the views expressed here do not necessarily reflect the views of the NAIC. More information about consumer participation at the NAIC is available [here](#).

- Existing nondiscrimination and data regulation laws should be assessed for their applicability to AI in UM.

2. Health insurance plans must be held accountable for inappropriate or discriminatory use of AI in UM

- Regulatory standards must clearly identify which parties are accountable (e.g., health plans, technology developers, etc.) when AI tools are used in UM decisions that lead to consumer harm, including discrimination, breaches of privacy, and incorrect adverse determinations.
- Regular audits, conducted on behalf of state regulatory agencies by parties with specialization in testing AI technologies, can be an effective way to both understand the ways AI is used in making UM decisions and hold the plans accountable for its use.
- AI tools intended for UM decisions should be built on standards of care that aim to achieve the highest level of quality, and penalties for non-compliance need to be significant enough to have influence.
- Governance structures that measure and prevent harm to historically marginalized and disenfranchised populations must be required.

3. Regulators should ensure that health insurance companies require oversight by humans with the appropriate clinical training regarding decisions that impact patient outcomes

- Robust and accessible appeals processes for coverage denials need to be established and considered a guaranteed right for all health insurance consumers.
- Human oversight must be embedded into UM when AI is used and those reviewers must have the authority and ability to overturn decisions made by the AI without undue consequences.
- AI regulation must to be considered an evolving practice that relies on collaboration between regulators, technical experts, industry stakeholders, consumers, and consumer advocates.

We are eager to discuss our research, findings, and recommendations with regulators and other stakeholders, and work with all parties to ensure they are realized for the benefit of consumers. Thank you for your attention to this important matter.

Sincerely,

Kellan Baker
Stephani Becker
Ashley Blackburn
Bonnie Burns
Jalisa Clark
Laura Colbert
Brenda Cude
Lucy Culp
Adam Fox
Eric Ellsworth
Stephanie Hengst
Claire Heyison

Anna Schwamlein Howard
Janay Johnson
Amy Killelea
Erin Miller
Carl Schmid
Jennifer Snow
Deborah Steinberg
Christa Stevens
Harry Ting
Wayne Turner
Caitlin Westerson
Silvia Yee

CONSUMER

HEALTH

ADVOCACY

AT THE NAIC

FINAL REPORT

November 2024

Artificial Intelligence in Health Insurance:

The Use and Regulation of AI in
Utilization Management

ACKNOWLEDGEMENTS

On behalf of the Consumer Representatives to the National Association of Insurance Commissioners (NAIC), The Leukemia & Lymphoma Society contracted with NORC at the University of Chicago to research the issues summarized in this report. The following Consumer Representative Work Group members contributed to the development of this report: Lucy Culp, Adam Fox, Amy Killelea, Eric Ellsworth, Carl Schmid, Wayne Turner, Caitlin Westerson, and Silvia Yee.

We would like to thank the seven organizations and their representatives that took the time to participate in the key informant interviews. Their candor and openness during the interviews greatly shaped our understanding of both the benefits and challenges AI brings to consumers' experiences with health insurance utilization management.

We would also like to thank the NAIC for supporting opportunities to provide consumer-focused input to insurance regulators on consumer access issues. The NAIC works closely with the Consumer Representatives to assist state regulators in their primary objective of protecting insurance consumers. The Consumer Representatives do not represent or work for the NAIC, and the views expressed here do not necessarily reflect the views of the NAIC. More information about consumer participation at the NAIC is available at content.naic.org/consumer/consumer-participation-program.

Table of Contents

Executive Summary 3

Introduction..... 6

Oversight & Regulation 17

Considerations for Regulators 26

Conclusion..... 32

Glossary..... 33

Executive Summary

The use of Artificial Intelligence (AI) is transforming how health insurance utilization management (UM) decisions are made, directly impacting consumer experiences and access to health care services. This report, compiled by NORC at the University of Chicago on behalf of the **Consumer Representatives to the National Association of Insurance Commissioners (NAIC)** examines the current use of AI in UM and related regulations.

The report found:

- Health insurance plans regularly use AI to conduct UM activities. Utilization management is the process by which health care benefit administrators (most often, health insurance plans) determine whether to approve payment for the course of care recommended by the patient’s medical team.
- AI proponents cite reductions in administrative burden and expedited approvals. However, there are risks that must be considered, such as the exacerbation of existing biases, prioritization of misaligned incentives, and use of technologies outside their original use case, or design, leading to unintended harm.
- Relevant stakeholders—including plans, providers, technologists, regulators, and consumer advocates—see immense opportunities for AI’s use. But, they warn that proper safeguards need to be in place to protect patient access to needed care and to ensure that both privacy and safety are appropriately considered.
- Some states have begun to regulate the development and use of AI in health insurance but, for the most part, have not been able to keep up with the proliferation of AI’s use. Regulation in this space is in the nascent stages of development and the speed at which change occurs makes this both a challenging and essential problem to solve.

Based on these findings, the Consumer Representatives to the NAIC call on state insurance regulators and lawmakers to consider the following insights when developing regulatory and legislative frameworks on the use of AI in UM decisions:

Transparency & Data

- Meaningful transparency is critical. Both regulators and consumers must clearly know when health insurance plans are using AI for the purposes of UM, and what role that AI plays in determining coverage for care.
- AI processes must be transparent, and that transparency must be multi-dimensional. Transparency must extend to disclosures about the data used to develop, train, and test the AI tools (with an emphasis on consent for use and representativeness of the population), the extent to which any AI tool can begin to train itself, and the criteria used for UM decisions. It must also create pathways for data to become public for use by researchers.
- Existing laws that are used to regulate data and take action against discrimination should be assessed for their applicability to AI in utilization management.

Accountability

- Regulatory standards must clearly identify which parties are accountable (e.g., health plans, technology developers, etc.) when AI tools are used in UM decisions that lead to consumer harms, including discrimination, breaches of privacy, and incorrect adverse determinations.
- Regular audits, conducted on behalf of state regulatory agencies by parties that specialize in testing AI technologies, can be an effective way to both understand the ways AI is used in making UM decisions and hold the plans accountable for its use.
- AI tools intended for UM decisions should be built on standards of care that aim to achieve the highest level of quality, and penalties for non-compliance need to be significant enough to have influence.
- Regulators must require governance structures that measure and prevent harm to historically marginalized and disenfranchised populations.

Ensuring Effective Oversight

- State regulators should establish robust and accessible appeals processes for coverage denials.
- Human oversight must be embedded into UM when AI is used. Those reviewers must have the authority and ability to overturn decisions made by AI.
- AI regulation must be considered an evolving practice that relies on collaboration between regulators, technical experts, industry stakeholders, consumers, and consumer advocates.

The rapid expansion of AI tools in health insurance demands immediate regulatory attention to protect consumers from potential harm and discrimination when AI is used in UM decisions. While this report outlines some key considerations, it is not exhaustive. Instead, it offers a foundation for understanding current AI-use cases in UM and highlights the urgent need for state policymakers to examine and regulate these practices.

AI presents opportunities to reduce administrative burdens, enhance providers' abilities to fully utilize their expertise, and improve patient care. However, it also carries significant risks that could negatively impact health outcomes and deepen discrimination against historically marginalized and disenfranchised populations. Decisions about coverage for essential medical services can have life-altering consequences. The importance of acting now cannot be overstated. Without immediate safeguards, the risks posed by unchecked AI in health insurance processes will only continue to grow.

Important Terminology

Utilization management (UM) is the process by which health care benefit administrators (most often, health insurance plans) determine whether to approve payment for the course of care recommended by the patient's medical team. While there may be a range of practices within this definition, the most common are prior authorization, concurrent review, and retrospective review.

Artificial intelligence (AI) is a catch-all term referring to technologies that enable computers and machines to mirror human learning and decision-making. Within AI, there are many different models and capabilities.

For this report, we are primarily focused on applications of **natural language processing (NLP)** and **machine learning (ML)**. NLP is a form of AI that allows computers to understand, interpret, and generate human language. ML refers to the ability of computer systems to learn and adapt beyond their initial instructions.

Introduction

As AI continues to evolve and expand into nearly every aspect of daily life, it is crucial that AI developers, organizations that use AI tools, and regulators with authority over health care entities work to protect consumers from the consequences that AI tools may have on their health care. Regulators have begun to close the regulatory gap on the use of AI in certain sectors of the insurance industry but those efforts to increase oversight have primarily been focused on life insurance rather than health care.¹

Some of these efforts have cross-industry implications that can be applied to health insurance (e.g., marketing, rate setting). However, the potential harm of AI's unregulated use in health care decision-making, and the direct impact it has on human health, requires additional safeguards and attention.

Methodology

To examine AI's current use in health insurance decision-making processes, NORC conducted a thorough literature review focusing on prior authorization as a form of utilization management (UM) and preliminary efforts to regulate AI. This report summarizes those findings and provides insights from structured interviews with seven organizations and their representatives who spoke to the implications and future of this largely unregulated space, including:

- 1) The current use of AI in UM
- 2) The evolution and future use of AI in UM
- 3) Regulatory and oversight recommendations

To preserve anonymity and encourage an open dialogue, all interview insights and quotes are attributed to a high-level descriptor, as outlined below. The intent was to gather diverse viewpoints from across the health insurance industry. None of the interviewees serve as NAIC Consumer Representatives.

Perspective	High-Level Descriptor
Health Plan	Analytics Executive at a Regional Health Plan
Thought Leader	Health Policy Professor
Consumer Advocate	Attorney for Underserved Patients and Families

Perspective	High-Level Descriptor
Consumer Advocateⁱ	Leaders at a Patient Advocacy Organization
Regulator	Representative from a State Department of Insurance
Technical Expert	Algorithmic Design and Measurement Consultants
Provider	Representatives from a Trade Group for Physicians

Examples of AI in UM

To help illustrate the use cases for AI in UM, the research team developed three examples of how health plans are using this technology today. While not all-inclusive, these examples are intended to provide plain language descriptions that reflect the differing levels of impact this technology can have on UM, and subsequently, how it may shape a patient’s ability to access needed care. Just as there is no one-size-fits-all approach to developing and using AI technologies, regulation to mitigate the potential harms of this technology needs to be responsive and sensitive to its range of capabilities.

Example	Scans Large Datasets	Uses Fixed Inputs to Make Determinations	Evolves Algorithm Based on Data
Administrative-Only (Example 1)	✓		
Decision-Making (Example 2)	✓	✓	
Learning Model (Example 3)	✓	✓	✓

ⁱ The second consumer advocate provided written responses to the structured interview questions.

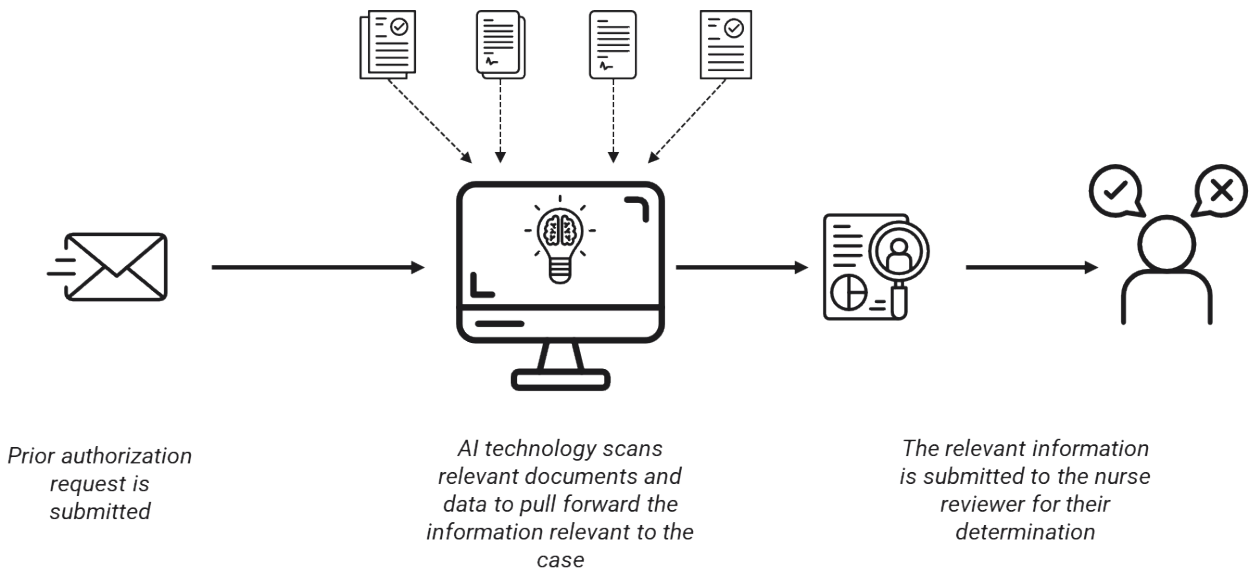
Example 1: Administrative-Only

One way that health plans currently use AI in UM is for sifting through a multitude of documents, to collate the relevant information for a human case reviewer. When a patient's physician submits a prior authorization request, that case is added to a queue for approval by a nurse reviewer. While in the queue, the AI program is instructed to pull a checklist of vital information to help the reviewer make their determination.

The AI technology has access to information such as the patient's medical claims history, details about the type of coverage they have and what is included in their benefits, the physician's prior authorization request, and guidance documentation from the health plan related to clinical appropriateness. It will sift through all of these documents and compile the relevant information into a single report.

In this case, the AI is focused only on collecting information. It alleviates the need for the nurse reviewer to spend time doing so, allowing the human reviewer to focus their time on evaluating the clinical elements of the case. The ultimate decision whether to approve or deny the claim lies solely with the human reviewer.

Exhibit 1. Flow of information and decision-making when AI is relegated only to administrative tasks

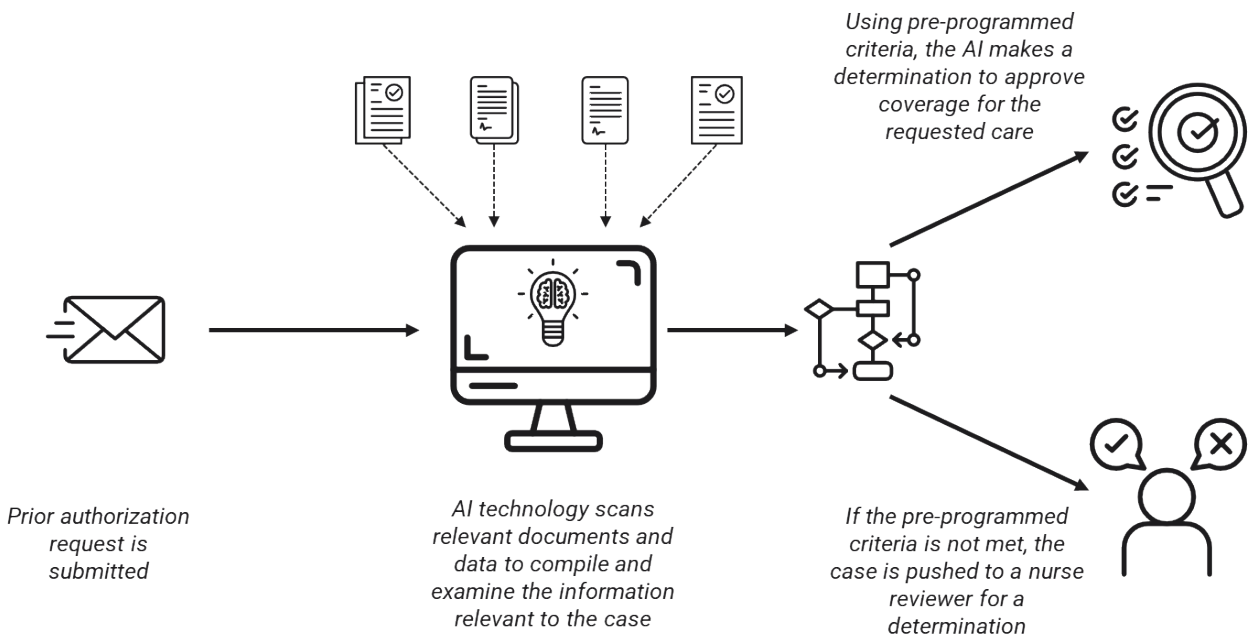


Example 2: Decision-Making

Another current use of AI in UM is to program the technology to do an initial assessment of all prior authorization requests using a specified set of criteria for decision-making and escalation. When a prior authorization request is submitted to the plan via an online portal, it is first evaluated by the AI technology using a set of inputs programmed by the health plan. These inputs may follow criteria to automatically approve certain common, low-cost requests and tag all others for secondary review by a nurse reviewer.

In this scenario, the AI determines the case's outcome (approve or escalate) based on a set of pre-determined criteria. This could mean that select prior authorization requests are never reviewed directly by a case manager. In this design, the AI only has the authority to approve coverage, and all denials are completed by a human reviewer.

Exhibit 2. Flow of information and decision-making when AI is allowed to make determinations about coverage



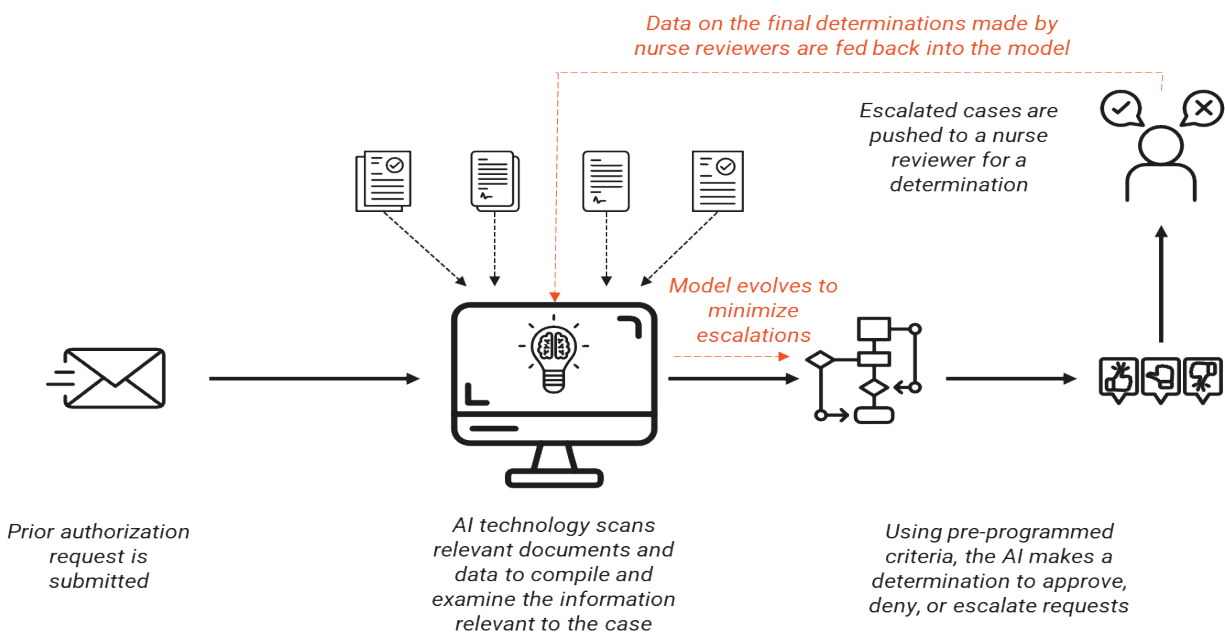
Example 3: AI Learning Model

The third way that health plans may use AI in UM is through a learning model. This model follows many of the same initial steps as the first two examples, collating and examining data inputs (such as the patient's claims history and clinical guidelines). However, in this case, the AI has greater authority to make one of three determinations: approve coverage, deny coverage, or escalate the case to a nurse reviewer.

This model is also programmed with the ability to evolve based on feedback. Specifically, the nurse reviewer's final determinations on escalated cases become data that the AI model uses to evolve its algorithm. For example, if nearly all escalated cases for a particular treatment were approved, the AI model will adapt its criteria to reflect this. Meaning, the AI model will approve future cases of the same type rather than escalating them, mirroring the vast majority of the nurse reviewer's prior determinations for similar cases.

While this type of model holds promise for innovation and increased efficiency, it also poses significant risks. Models that are allowed to evolve based on data from incomplete or unverified datasets can reflect and reinforce biases or inaccuracies carried out by the nurse case reviewers.

Exhibit 3. Flow of information and decision-making when AI is allowed to make coverage determinations and evolve based on feedback



The Current State of Artificial Intelligence in Health Insurance Decision-Making Processes

Artificial intelligence has been around for decades, but more recent technological advancements have led it to permeate nearly all aspects of everyday life, including health care. AI's uses in health insurance processes have a wide variety of potential benefits, including improving care delivery, quality, and overall consumer experience. However, AI also poses a variety of potential threats to consumer access to care, particularly for historically marginalized and minoritized populations.

While important, a broad examination of the impact of AI across health care is not within the scope of this report. This research is focused exclusively on the application and use of AI in health insurance UM processes. Within this one application, there is evidence that AI is already widely used today and that, while it may alleviate the administrative burden on staff, it also has the potential to cause significant harm both in perpetuating existing issues with UM and introducing new ones.

To compile an understanding of the current use of AI in UM and its impact on consumers, the research team first conducted a comprehensive literature review. To add context and detail to that search, the research team also conducted interviews with seven organizations representing diverse perspectives of those directly impacted by the use of AI in UM. The following section of the report summarizes and integrates findings from both modes of data collection.

Today, one of the most common uses of AI in health insurance is for utilization management.

AI includes a wide range of capabilities and technologies, all culminating in computer or machine systems that are capable of complex tasks, such as decision-making and processing large datasets. In recent years, health plans have leveraged the abilities of AI to make UM decisions, in particular, to respond to prior authorization requests.²

AI Applications in Prior Authorization Processes

Prior authorization is an administrative process in which a benefit administrator, such as a health plan, requires providers to seek approval for a specific course of treatment on behalf of their patient, to ensure the patient will have coverage.³ This process can safeguard a patient from unnecessary or inappropriate care, or surprise bills for services that their plan doesn't cover. A major concern, however, is that seeking prior authorization can slow down or even prevent access to necessary care due to arbitrary or wrongful denials. Even if the treatment is eventually approved, the administrative burden and the time spent waiting for approval can negatively impact patient outcomes and quality of life. In a survey conducted by the American Medical Association in 2023, the overwhelming majority of providers (94 percent) reported prior authorization delays care and more than half (55 percent) said that such delays often or always slow down care delivery.⁴ Since prior authorization is an inherently administrative task that increases the clerical burden on staffers, it is unsurprising that nearly all interviewees said that one of AI's primary benefits is that it reduces this burden.

"In an ideal world, AI would increase efficiency without posing any additional harms to patients or their access to care."

- Consumer Advocate

Of all interviewees, the health plan executive was the most optimistic about the potential applications for AI in UM practices. They pointed to several other AI benefits beyond reducing the administrative burden of prior authorization requests. They explained that their plan uses natural language processing to collate several sources of information essential to a nurse case reviewer's determination of coverage (similar to Example 1, above). This might include pulling together details related to patient coverage, medical history, and any relevant clinical criteria determined by the plan, to inform medical

"The primary benefit of using AI for utilization management is the ability to reduce the clerical burden, expedite approvals for patients, and enable practitioners to practice at the top of their license."

- Health Plan Executive

necessity. Instead of the nurse case reviewer collating this information—which may be stored across multiple systems or buried within larger documents—AI can take over this task, liberating the human reviewer to focus on pertinent clinical details and decision-making and, ideally, use their time more efficiently.

It is important to note that efficiency cannot be the only goal of integrating AI into UM. The ultimate goal is getting to the correct determination. AI's inclusion

should be welcomed depending on the extent to which it enables organizations to realize this end. However, efficiency simply for the sake of efficiency may only reinforce existing issues within UM rather than alleviate them. The consumer advocate interviewee was particularly concerned by this. On the whole, they regarded the entire premise of UM with skepticism and worried that greater efficiency would simply lead to a faster cadence of incorrect or adverse decisions that could result in delayed or denied care.

While there was some agreement among stakeholders that AI can reduce the clerical burden of prior authorizations, there was also unanimous concern that AI-enhanced processing could lead to harmful delays or faulty rejections.

While this concern predates AI, the increased volume and speed at which decisions may be made while using AI, the rigid clinical criteria used by AI, and the delegated authority to computer-based algorithms pose additional complexities and concerns.

These findings drive home the paramount need to clarify lines of authority and accountability so that patients and providers can effectively advocate for themselves when insurers deny medically necessary care.

As AI tools are developed and deployed to make coverage decisions, concerns arise.

In the absence of a comprehensive regulatory framework for using AI technology in health insurance, stakeholders themselves have begun identifying the potential risks that may adversely impact care delivery and health outcomes. Those risks include:

Tools trained by biased datasets

As AI developers create new tools to augment decision-making related to care provision, they often rely on existing datasets, written descriptions of clinical care standards, or clinical algorithms to inform the decision criteria. The problem is that AI may be pulling from datasets that reflect known or unknown biases. One example of this is the calculation for the estimated glomerular filtration rate (eGFR), which was used to diagnose and manage kidney disease. Historically, the commonly accepted

“Transparency must be meaningful and enable end users to trace a decision back to a specific actor, to accurately determine decision rationale for potential adverse outcomes.”

– Provider Trade Group

standard for this estimate included race as a factor. However, recent research called into question the validity of including race in eGFR and highlighted the resulting negative impact on Black patients, who experienced poorer outcomes than white patients with the same condition.⁵ While this research largely led to the discontinuation of race as a factor in eGFR calculations in 2021, it was only after decades of use. AI developers need to continually update tools to reflect new clinical practice guidelines, especially when the new guidelines address health disparities.

Several interviewees raised the valid concern that AI trained on biased datasets poses potential risks for discrimination against consumers, particularly patients from protected classes.⁶ Developing an AI tool to standardize prior authorization requests that rely on biased data will exacerbate the same biases—in diagnoses, treatments, and provision of services—that pre-dated AI and still exist today.⁷ While this finding is not new, the speed at which these tools are evolving and the lack of transparency surrounding them limits regulators and other stakeholders from deeply understanding the algorithms and processes that underlie these technologies. This makes it ever-more difficult to prevent the potential harm inflicted on historically marginalized and disenfranchised groups.

“The AI tools being used today are based on historically biased data.

It’s one thing to look at a model and say ‘This algorithm is biased based on the data that we use to develop it,’ but there is also a gap in the patients who are able to fight back against the resulting denials.”

– Consumer Advocate

The consumer advocate interviewee emphasized that, beyond worsening existing biases in health care delivery, a lack of transparency around AI tools will disproportionately increase harm to historically marginalized groups by limiting or preventing their ability to appeal denials from their health plan. This is further exacerbated by the fact that AI sometimes uses proprietary datasets, making it even more difficult for consumers and providers to get detailed

information about why a request was denied and what criteria and data were used.

Algorithms developed with misaligned incentives

Prior authorization decisions are expected to be based upon generally accepted standards of care. However, there is evidence that this is not always the case. In *Wit v. United Behavioral Health (UBH)*, the Ninth Circuit Court of Appeals found that UBH—a subsidiary of United Healthcare, the largest U.S. health insurance company by both market share and

revenue—had been denying behavioral health claims in favor of optimizing their financial outcomes.⁸ With AI tools increasingly involved in making these types of coverage decisions, consumer advocates are concerned that these tools will be trained based on financial or otherwise misaligned incentives that are in direct conflict with the generally accepted clinical standards of care.

To address these concerns, stakeholders are calling on AI developers and organizations that use AI tools for decision-making to clearly and meaningfully disclose how their AI tools make prior authorization decisions. This includes full transparency on the inputs and decision-making criteria that are being used.

Another important safeguard is ensuring an appropriate level of oversight for any consequential patient care decision. There is evidence that using AI in decision-making processes creates a circumstance wherein algorithms can make thousands of unwarranted denials at a rapid pace.⁹ While some health plans have policies that promise internal oversight of and accountability for any decision that may result in a denial, it is a largely unregulated and uneven practice across the industry. The health plan executive interviewee said that while their plan uses AI tools to expedite approvals, any claims that could be denied are routed to the appropriate, clinically-trained case manager.

Studies are needed on the impact of this type of tiered decision-making, but it may prove to be one model on which to build future standards. Either way, embedding transparency and accountability throughout the process is essential. Establishing discrete and clearly documented processes for how all UM determinations are made provides a clear perspective of how a decision came to be and who is accountable for any resulting harms.

Machine learning systems developing their own processes

Another significant concern about using AI in health insurance decision-making processes is that machine learning tools are constantly learning and evolving as they get introduced to more data, as outlined in Example 3. This can lead to the tool operating beyond its originally intended use case, increasing potential harm to consumers. A dynamic AI model poses a unique challenge to regulators because it requires continuous monitoring and other processes to ensure that it is consistent and adheres to the use cases for which it was initially trained.¹⁰ The technical experts who were

“The chance to monitor and test AI systems is a chance to test and monitor outcomes to the standard that society expects.”

– Technical Expert

interviewed believe that ongoing monitoring is essential to identifying harms, for holding those who are responsible for them accountable, and for helping dynamic models evolve in a direction that builds towards a better future.

The health plan executive interviewee reported that their regional plan already has safeguards in place to ensure that its algorithms are not evolving into an entirely new model. Specifically, when an algorithm is ready to be used for UM, this health plan “freezes” the model ensuring that it continues to operate as originally intended no matter what data it is given. See Example 2, above. This particular plan uses its model for one to two years,

“ ‘Good’ AI governance not only requires companies to be aware of what they are doing and what models they are using; they must also have a regular assessment to ensure models behave appropriately.”

- Health Plan Executive

while retraining it on newer data offline, validating its efficacy prior to use on real patient cases. This approach of freezing machine learning or evolving models and routinely monitoring them while in use warrants additional study, but could be a starting point for future regulatory action. This would not prevent health

insurance companies from using parallel processes wherein their models are allowed to evolve using offline data. It just would limit their ability to do so before the evolved model can be tested and validated.

Oversight & Regulation

As the health insurance industry continues to invest in AI tools, regulators have struggled to keep up, leaving an uneven regulatory landscape particularly at the state level resulting in a legislative patchwork across state lines. Whether state or federal, efforts to regulate this ever-changing area have to balance the desire for innovation with the need to protect consumers from any potential adverse effects.

To address the lack of regulatory continuity, the NAIC's Innovation, Cybersecurity, and Technology (H) Committee developed a model bulletin on the use of AI in insurance processes.¹¹ While not specific to health insurance, the bulletin outlines how insurers should govern the development, acquisition, and use of AI technologies. As of October 2024, 17 states had adopted this guidance.¹²

States have begun developing their own approaches for regulating this evolving environment.

In the absence of a comprehensive federal approach, some states have created their own regulations of AI use in the insurance industry.

State Oversight Landscape

In 2024, there was a flurry of consumer protection-oriented state legislation on the general use of AI across multiple industries.¹³ This report does not focus on such broad legislation. Instead, it examines legislation that specifically aims to regulate the use of AI in health insurance decision-making processes, particularly prior authorization, or legislation that may have clear implications in this space. State-level legislative efforts on the use of AI in health care are still nascent and the majority of proposed bills are focused on creating task forces and developing studies on current-use cases of AI in their own state.

In the first half of 2024 (January–June), there were more than 30 state bills that created AI task forces. While each task force has slightly different responsibilities, their overarching goal is to develop recommendations for state government AI use oversight.¹⁴ For example, in April 2024, Connecticut's House of Representatives proposed CT SB2 to establish an AI Advisory Council that would make recommendations on the ethical and equitable use of AI in state government.¹⁵ Similarly, New Jersey's State Senate introduced NJ S3357, to establish an AI advisory council and also focus on the opportunities and risks for state agencies in leveraging AI tools.¹⁶ While neither of these task forces specifically focuses on

AI use in health care, there likely will be implications for AI's use in health insurance decision-making processes.

In addition, some states are making headway in regulating transparency between AI developers, organizations that use the tools, and end users (i.e., consumers) who are impacted by AI use. A common theme throughout the stakeholder interviews was the importance of “meaningful” transparency, which includes patient-facing disclosures being written clearly so that consumers can understand and consider the content and its implications. Patients face innumerable, lengthy disclosures while receiving health care, but very few of them meaningfully engage with this information because it is provided in formats that (or at times when) it cannot reasonably be consumed and thoughtfully considered. The goal across all of health care should be to evolve past this “check the box” disclosure approach. As disclosures related to AI are brought into existence, they should set a new standard for what is meaningful to consumers.

States have taken varied approaches to requiring transparency from AI developers. Quite often this includes compelling the developers to provide documentation to regulators that details intended use cases, what type of data was used to train the model, data-collection practices, and how the developer mitigated risks of discrimination.^{17,18} While transparency between the developer and organization using the tool is crucial to ensuring its ethical use, it is equally important for there to be the same level of meaningful transparency between the organization using the tool and regulators, so that they can monitor use over time and address any adverse outcomes.

Individuals who had their prior authorization request for medical care denied often do not know why their insurance will not cover what they and their provider both believed to be necessary care. According to the consumer advocate interviewee, before AI it was somewhat possible for consumers to appeal claim denials or, at the very least, have a clear answer as to why coverage was denied. Now, AI is an impenetrable “black box,” obscuring the chain of command and making it nearly impossible for consumers to push back on decisions regarding their own care. As health plans continue to use AI tools to reduce administrative burdens, states will need to take legislative or regulatory action to ensure consumers can understand how their health care decisions are being made. Several states—including Illinois, Vermont, Virginia, and California—have introduced legislation that would require organizations that use AI tools to notify the end consumer, at or before the use of AI, that an AI tool is being used in health care decision-making.^{19, 20, 21,22}

While there is a wide range of state legislative action related to the use of AI in insurance processes, few states have taken concrete action to advance regulation in the health insurance space. Following are examples of three states that are addressing AI in insurance broadly and otherwise.

Colorado

Colorado is often touted as the state at the leading edge of AI regulation. In July 2021, Governor Jared Polis enacted SB 169 to protect consumers from discriminatory insurance practices.²³ A primary component of the law is that it holds insurers accountable for testing their big data systems, including algorithms and predictive modeling, to ensure that systems do not go unchecked and perpetuate discrimination against protected classes. The law's implementation is well underway and may serve as a model for other states looking to be more proactive in regulating AI in the insurance industry.

As part of this implementation, Colorado's Insurance Commissioner held several meetings with stakeholders to hear their thoughts on how future AI governance should be structured. While there has been stakeholder engagement across insurance types—including life, auto, and health—as of now, the Commissioner has only adopted framework requirements for life insurance practices. Other states may be able to learn from Colorado's challenges and how its Division of Insurance has overcome them so far, including having limited expertise in how to conduct evaluations and little oversight of how insurers report on AI.

More recently, to protect consumers, Colorado enacted its Colorado Artificial Intelligence Act (CAIA), a comprehensive law on the development and use of certain AI systems which will go into effect on February 1, 2026.²⁴ CAIA takes a risk-based approach, meaning it focuses on regulating high-risk AI systems. High-risk AI systems are those that make consequential decisions, which may include health care decisions. It requires both the developers and users of high-risk AI systems to document and disclose the use of AI in tools or processes, analyze and prevent potential harms, and establish data governance structures to examine potential biases, among other measures.

Utah

Utah has begun to regulate AI use across industries. In March 2024, Utah's state legislature passed SB 0149, the Artificial Intelligence Policy Act, and the state's government signed it into law shortly thereafter. It establishes requirements for AI-use disclosures and creates a state Office of AI Policy, which is billed as the first office dedicated to AI policy, regulation,

and innovation.²⁵ The office works with stakeholders to align on regulations and find the balance between innovation and consumer safety.²⁶ Another primary component of the new law, which is being quickly implemented statewide, is the establishment of disclosure requirements that compel businesses that use AI tools to clearly disclose to the consumer that they are interacting with AI and not a human. While not specific to regulating AI's use in the insurance or health care spaces, there are more restrictive disclosure requirements for regulated entities, including health care.

While Utah's new law may serve as a promising example of how regulators should think about disclosure requirements, it is limited to the regulation of generative AI that interacts directly with consumers. Regulators must be able to build off of legislation such as this to ensure that AI use in health insurance decision-making—such as UM processes—is also properly regulated.

California

California has made recent advancements in regulating AI use in the health sector, passing a law in September 2024 to regulate AI use by health plans.²⁷ Specifically, California amended § 1367.01 of the Health and Safety Code and § 10123.135 of the state's Insurance Code, to require that AI tools for utilization management be open to inspection for audit by the State Department of Health and that they make coverage decisions based on the enrollee's medical history and clinical circumstances, not solely on the group dataset on which the tool was trained. The law also requires physicians to have the final say in determining medical necessity when health plans use AI tools for UM.²⁸ Consumer advocates and other stakeholders, including the California Hospital Association, supported the law's passage and argued that it balances consumer protection with the growing use of innovative tools.^{29, 30} It is likely that there will be more states that propose similar legislation specific to using AI in UM of health care decisions in the coming months.

While this report focuses on state-level approaches such as California's, there have been preliminary efforts to regulate the use of AI at the federal level. Whether these federal efforts are applicable to regulating AI use in the health insurance space, especially related to UM, remains to be seen. However, such efforts may have implications that cross over, as they did at the state level.

Many organizations have developed frameworks on how AI should be used and regulated in health insurance practices.

With the absence of federal regulation, and a patchwork approach at the state level, many organizations have developed their own policy agendas or guiding principles in order to better serve their consumers and be better prepared for future regulation. Similar to the evolution of AI regulation at the state level, many organizations' AI frameworks are not specific to AI in health insurance but rather provide high-level guidance for the ethical use of AI that has cross-cutting implications for health care.

Some advocates in the field have already started to ideate the potential rules and oversight for AI's safe and ethical use. This report examined several select frameworks released by NAIC and other consumer-oriented organizations. While the list that follows is not exhaustive, it is intended to reinforce the concerns raised throughout this report. ***There is consensus among consumer advocates that ethical use, governance, transparency, and human-centered decision-making need to be central to future regulatory or legislative oversight.***

National Association of Insurance Commissioners: Guiding Principles on AI

In August 2020, NAIC adopted a set of guiding principles on AI to help insurance companies, and organizations entering into business with insurance companies, navigate this complex field.³¹ While not limited to health insurance specifically, the principles provide high-level guidance for stakeholders and serve as a basis for any future recommendations that NAIC may make to state regulators.

NAIC consumer representatives subsequently commissioned this report to examine the current status of AI in the health insurance landscape, as a first step toward developing concrete regulatory recommendations.

National Health Law Program: Principles for Fairer, More Responsive Automated Decision-Making Systems

In May 2023, the National Health Law Program (NHeLP) released its guiding principles for automated decision-making systems, explaining that without oversight, big data systems are left susceptible to gaming, which may result in discrimination.³² NHeLP protects and advances the health care rights of low-income and traditionally underserved individuals. Its principles largely align with those of other organizations and include:

1. **Transparency across the tool's life cycle**, which starts at development and continues through implementation and any potential revisions
2. **Protection of civil rights**, to ensure that benefits programs, such as Medicaid, continue to adhere to all civil rights and due process requirements and that tools do not discriminate on the basis of any legally protected characteristic
3. **Tools should be user-focused**, to ensure that the benefit of the end-user is prioritized throughout the entire life cycle of the tool
4. **Tools must be regularly validated**, to ensure that they accurately and consistently measure what they claim to measure for their intended use case
5. **AI developers and users must regularly review how their tools are being used**, to prevent any potential biases and minimize potential adverse outcomes against any protected community
6. **Human oversight and feedback loops**, to protect against potential errors at the individual or systemic level

The National Health Council: Position on the Promise and Pitfalls of AI and Health Care

The National Health Council (NHC) was created by and for patient organizations to be a voice for people living with chronic diseases and disabilities, and their family caregivers.³³ NHC has more than 170 active member organizations ranging from patient- and disease-specific advocacy groups, professional societies, and cross-industry businesses. In February 2024, NHC released a statement on AI's potential benefits for patients as well as its risks, all of which legislators must consider as they begin to regulate this space.³⁴

NHC's statement provides a wide range of key components for using AI responsibly, including:

- **Human oversight**, to guarantee the algorithm's safety and accuracy and ensure continuous improvements as conditions change, including AI-tool advancements
- **Robust and continuous feedback loops**, to identify and prevent the risk of harm
- **Pre-deployment testing** that includes testing on a diverse range of real-world settings

Along with CVS Health, the American Heart Association, and other health systems and organizations, NHC is one of the founding partners of the Coalition of Health AI (CHAI). CHAI represents a wide array of stakeholders to drive development, evaluation, and the appropriate use of AI in health care.³⁵

The American Medical Association: Principles for AI Development, Deployment, and Use

In November 2023, the American Medical Association (AMA) released its set of principles for the development and use of AI in health care decision-making settings.³⁶ In the absence of a comprehensive federal policy that oversees the use of AI, the AMA developed the following principles to protect and inform both physicians and patients:

- **Oversight of AI use in health care** should include a “whole government” approach in order to mitigate any potential risks.
- **Transparency around the use of AI is critical for physician and patient relationships**; there should be laws in place to mandate transparency around the design, development, and deployment processes of AI tools.
- **Documentation and disclosure of the use of AI** should be required when AI may impact patient care, which may include access to care, decision-making processes, and communications.
- **Generative AI requires additional safeguards to protect patients**, so governance policies must be in place prior to the adoption and use of any generative AI tools.
- **AI developers must be held accountable for data privacy when developing tools**, and both developers and health care organizations must create safeguards to protect patient privacy and data when AI tools are being used.
- **AI developers and health care organizations should be proactive in identifying and mitigating biases**, to ensure equitable health care outcomes among patients.
- **Physician liability related to the use of AI** should be limited to existing legal approaches to medical liability.

AHIP:

Priorities for AI to Transform Health Care and Administrative Processes

In July 2023, AHIP (formerly known as America's Health Insurance Plans), an advocacy and trade association that represents health insurance companies, responded to a request for information (RFI) for potential regulatory action to protect consumers and prevent biases in the development of a national AI strategy.³⁷ In response, AHIP outlined several policies to supplement existing industry regulations, including:

- **Increase consumer education around the uses of AI**, so that consumers are able to make informed decisions and are aware of potential benefits and harms when AI tools are used in relation to their health care
- **Continue to monitor AI advancements**, so that U.S. policy and advancements can remain competitive and protect consumers from potential nefarious activities that endanger data privacy and security
- **Balance innovation with consumer safeguards**; to protect consumers, efforts to improve disclosures, transparency, and audits of AI tools should be tailored to the specific algorithm rather than use a one-size-fits-all approach
- **Oversight frameworks should be risk-based**, to ensure that risk-mitigation techniques are right-sized to the potential for adverse outcomes and are not overly restrictive
- **Regulatory frameworks should be flexible, based on the use case of the tool**, so that any general AI regulations can be tailored to the specific respective needs, such as health care
- **Mitigate the perpetuation and introduction of biases** in both underlying data and the development of algorithms

Considerations for Regulators

In addition to describing how AI is currently being used in health insurance decision-making, and how stakeholders view the potential opportunities and risks of AI, NAIC Consumer Representatives call on regulators to keep the following considerations top-of-mind while they work to address the ever-evolving AI space.

Transparency & Data

Transparency—to consumers, providers, and regulators—is a crucial component of AI oversight. This can be seen in regulatory and legislative action to date and is reflected in the guiding principles for AI put forward by health care advocates.

Meaningful transparency is critical. Both regulators and consumers must be fully aware when AI is being used by health insurance plans for UM purposes and be clear on what role the AI plays in determining coverage for care.

- There is overwhelming consensus—in interviews conducted for this report, existing research, and across organizational AI frameworks—that meaningful transparency must be included as one of the first steps of any AI regulation. Meaningful transparency must go beyond a one-size-fits-all approach and instead be tailored to the use case for which the AI tool is being deployed. What this means is that disclosures will differ based on the audience, be it consumer, provider, or health insurance plan.
- To ensure that AI-use disclosures go beyond simply adding another form to the multitude that patients already receive at the doctor’s office, some states—including Utah—have said that businesses using AI tools are responsible for providing a clear and plain disclosure every single time a consumer interacts with the tool.
- Transparency between technology developers and regulators is also particularly important in this rapidly evolving space. To adequately protect consumers, regulators must be able to review the use cases, implementation, and data being used by respective tools.

Transparency must be multi-dimensional. It must extend to disclosures about the data used to develop, train, and test the AI tools (with an emphasis on consent for use and representativeness of the population), the extent to which any AI tool can

begin to train itself, and the criteria used for UM decisions. It must also create pathways for data to become public, for use by researchers and others.

- Along with requiring transparency around the use cases of AI tools, transparency between the technology developer and regulators must include what data is being used relative to the tool. This transparency will help regulators prevent discrimination and potential adverse outcomes, as outlined in Colorado’s SB 169.
- Beyond requiring disclosures around what data is used to develop, train, and test AI tools, data outputs related to the use of AI tools by health plans must be made publicly available. This transparency will enable relevant stakeholders, including academics and consumer advocates, to aggregate trends and identify any potential bad actors.

Existing laws that regulate data and take action against discrimination should be assessed for their applicability to AI in utilization management.

- Several states—including Colorado, Utah, and California—have made significant headway in regulating the use of AI in the broader insurance space. The policies, approaches, challenges, successes, and other lessons learned by these pioneering states may all be used by other states as frameworks for developing their own AI regulations.
- There has been substantial state regulation of the AI space beyond the insurance sector. Reviewing these laws and regulations should also be one of the first steps in developing frameworks to regulate the use of AI in UM processes specific to health insurance.

Accountability

Relying on proprietary AI technologies for critical decisions about patient care is problematic because it can obscure who is accountable for decisions that may harm a patient. Transparency is therefore critical to holding health insurance plans accountable and, when appropriate, liable for the harm caused by the integration of AI into UM activities.

Regulators must also consider how AI tools are trained, to ensure that (a) such tools are not developed by bad actors with misaligned incentives and (b) there are strategies in place to ensure that the tool operates, throughout its entire life cycle, as initially intended.

Regulatory standards must clearly identify which parties are accountable (e.g., health plans, technology developers, etc.) when AI tools are used in UM decisions that lead to consumer harms, including discrimination, breaches of privacy, and incorrect adverse determinations.

- Beyond increasing transparency to enable potential enforcement of regulatory standards, regulations must clearly delineate accountability when AI tools are used for UM decisions.
- Utah has held businesses accountable for failing to disclose the use of AI. However, there is currently limited regulatory action elsewhere regarding which parties involved in using AI in UM are to be held accountable for poor outcomes.

Regular audits, conducted on behalf of state regulatory agencies by parties that specialize in testing AI technologies, can be an effective way to both understand the ways AI is used in making UM decisions and hold the plans accountable for its use.

- In Colorado, SB 169 requires insurers to review their big data systems, including algorithms and predictive models, and report the status of their review and findings back to the Colorado Department of Insurance for evaluation. This process aims to identify and correct any discrimination against protected classes and ensure that the insurer's AI tools are not perpetuating bias.
- The technical expert interviewed for this report expounded on the importance of having regular audits, of both dynamic and static models, to ensure that models are acting within their use case and not increasing harm to any consumer.

AI tools intended for UM decisions should be built on the highest-quality standards of care and penalties for non-compliance need to be substantial enough to have influence.

- AI tools need to be built upon clinical standards that seek to achieve the highest quality of care and not to be in service of other outcomes, including financial savings. There is evidence that AI is being used by health plans to expedite denials

in order to increase savings and/or to create structures that disincentivize clinical reviewers from spending adequate time on reviewing the clinical merit of individual cases.

- Requiring the consideration of the highest-quality clinical standards is an essential component of UM that protects consumers from harm. This is true regardless of whether or not AI is used in the decision-making process. As AI grows in influence, it is essential that clinical criteria are embedded in any algorithm that has the ability to deny coverage.

Regulators must require governance structures that measure and prevent harm to historically marginalized and disenfranchised populations.

- Just as it does across society, bias exists in the medical field and can be measured both statistically and socially.³⁸ Statistically, we know that there is a persistent lack of representation of certain groups—including women, people of color, and people with disabilities—in large datasets commonly used for clinical research. As a result, AI tools that are trained on such datasets will continue to reinforce biases that lead to outcome gaps among historically marginalized and disenfranchised populations.
- Potential ways of preventing harm may include requiring developers to test and train emerging tools on diverse datasets, mandating the establishment of internal governance structures that determine the metrics for bias, conducting regular external measurements, and creating regulatory structures that require the resolution of measured bias in UM outcomes.
- Regulators can also craft regulations that require AI governance structures to create policies that compel health plans to document and publicly disclose the detection of biases and how it was addressed.

Ensuring Effective Oversight

When it comes to oversight, regulators need to ensure that health insurance companies place humans who have the appropriate clinical training in positions that are responsible for determining patient care. Those individuals must have the level of authority and support needed to override determinations made by AI algorithms, whenever appropriate. There also must be a clear pathway for consumers to appeal denials, as a backstop when the wrong determination is made by either a human reviewer or AI.

State regulators should establish robust and accessible appeals processes for coverage denials.

- Patients and providers need to be given a rationale for every denial, so that they can determine whether or not to appeal the decision or resubmit the request with modifications. There should also be up-front disclosures about the role AI plays in the determination process. The technical experts interviewed for this report highlighted the need for a clear definition of how AI was involved in the process, including the level to which it assisted in the determination (e.g., substantially assisted, minimally assisted, etc.).
- There needs to be an identification and resolution process in place to address tools that are used beyond their intended use. This may require health plans to establish and maintain a registry of AI-enabled actions that is embedded in their UM process. Making this registry transparent to regulators may further aid accountability and understanding of how the field is evolving.

Human oversight must be embedded into UM when AI is used, and those reviewers must have the authority and ability to overturn decisions made by the AI without undue consequences.

- While AI tools may be useful in expediting approvals, at minimum, human practitioners with the appropriate clinical training must be accountable for any denials of coverage. If an AI tool recommends that a claim be denied, then a health plan needs to have a clinically trained case manager available to directly support that patient should they choose to appeal the decision.
- Human reviewers must be appropriately supported and empowered to overturn AI algorithms' recommendations. This means that (a) expectations and safeguards must be in place to protect clinical reviewers from retribution and (b) financial incentives based solely on the volume of claims processed should be eliminated.

AI regulation must be considered an evolving practice that relies on collaboration between regulators, technical experts, industry stakeholders, consumers, and consumer advocates.

- Because this is a complex and ever-evolving field, regulatory approaches need to be sufficiently prescriptive to protect consumers but flexible enough that regulations

do not quickly become outdated or unenforceable. States should consider leveraging technical experts, and allocating funding to study the use of AI by health plans that operate in their state, as important first steps in this process.

- To achieve the dual objectives of encouraging innovation while also ensuring protection from harm, all parties in this space need to closely collaborate and co-design solutions. This includes operating from a place of good faith and transparency.

Conclusion

Given the rapid proliferation of AI tools in health care, regulators must engage more proactively in both the development and enforcement of regulations that protect consumers from potential harm and discrimination by AI. Within the current technical environment, change happens at a rapid pace and immediate action is necessary to stay ahead of the potential harms that can emerge from the unregulated use of these technologies. The considerations laid out in this report are not exhaustive, but they may provide a baseline understanding of current use cases of AI in health insurance processes and opportunities for regulation by stakeholders who see the impact of unregulated AI use in health care today.

AI provides opportunities to reduce administrative burdens, increase practitioners' abilities to practice at the top of their license, and improve patient care. However, as outlined above, there are also real and significant risks to using AI tools in health care that must be safeguarded against to prevent poor health outcomes and further discrimination against historically marginalized and disenfranchised populations.

Determinations of whether provider-recommended clinical items and services should or should not be covered can be life-altering, thereby cannot be left merely to AI tools or the developers or plans that use them. The time to act is now.

Glossary

Artificial intelligence (AI) is a catch-all term referring to technologies that enable computers and machines that have the ability to mirror human learning and decision-making. There are many different models and capabilities within AI.

Utilization management (UM) is the process by which health care benefit administrators (most often, health insurance plans) determine whether to approve payment for the course of care recommended by the patient's medical team. While there may be a range of practices within this definition, the most common are prior authorization, concurrent review, and retrospective review.

Machine learning (ML) refers to the ability of computer systems to learn and adapt beyond their initial instructions.

Natural language processing (NLP) is a form of AI that allows computers to understand, interpret, and generate human language.

Use case refers to a purposeful application of an AI model for a discrete business case or outcome. Use cases often outline the expected series of interactions between a computer system and a select set of data inputs to achieve a predetermined goal.

Prior authorization is an administrative process in which a benefit administrator, such as a health plan, requires providers to seek approval for a specific course of treatment on behalf of their patient to ensure the patient will have coverage.

References

- ¹ AI regulation faces test on life insurance in Colorado. (2023, March 28). *Bloomberg*. <https://www.bloomberg.com/news/articles/2023-03-28/ai-regulation-faces-test-on-life-insurance-in-colorado>
- ² AI ushers in next-gen prior authorization in healthcare. (2022, April 19). *McKinsey & Company*. <https://www.mckinsey.com/industries/healthcare/our-insights/ai-ushers-in-next-gen-prior-authorization-in-healthcare>
- ³ What is prior authorization? (2022, July 12). *American Medical Association*. <https://www.ama-assn.org/practice-management/prior-authorization/what-prior-authorization>
- ⁴ American Medical Association. (2023). *2023 AMA Prior Authorization Physician Survey*. <https://www.ama-assn.org/system/files/prior-authorization-survey.pdf>
- ⁵ Ahmed, S., Nutt, C. T., Eneanya, N. D., Reese, P. P., Sivashanker, K., Morse, M., Sequist, T., & Mendu, M. L. (2020). Examining the potential impact of race multiplier utilization in estimated glomerular filtration rate calculation on African-American care outcomes. *Journal of General Internal Medicine*, 36(2), 464–471. <https://doi.org/10.1007/s11606-020-06280-5>
- ⁶ Agarwal, R., Bjarnadottir, M., Rhue, L., Dugas, M., Crowley, K., Clark, J., & Gao, G. (2022). Addressing algorithmic bias and the perpetuation of health inequities: An AI bias aware framework. *Health Policy and Technology*, 12(1), 100702. <https://doi.org/10.1016/j.hlpt.2022.100702>
- ⁷ *Ibid.*
- ⁸ *WIT v. United Behavioral Health I*. (2022). American Psychological Association. <https://www.apa.org/about/offices/ogc/amicus/wit>
- ⁹ Rucker, P., Miller, M., & Armstrong, D. (2023, March 25). How Cigna saves millions by having its doctors reject claims without reading them. *ProPublica*. <https://www.propublica.org/article/cigna-pxdx-medical-health-insurance-rejection-claims#:~:text=Over%20a%20period%20of%20two%20months>

-
- ¹⁰ Meskó, B., & Topol, E. J. (2023). The imperative for regulatory oversight of large language models (or generative AI) in healthcare. *Npj Digital Medicine*, 6(1).
<https://doi.org/10.1038/s41746-023-00873-0>
- ¹¹ National Association of Insurance Commissioners. (2023, December 2). *NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers*.
https://content.naic.org/sites/default/files/inline-files/2023-12-4%20Model%20Bulletin_Adopted_0.pdf
- ¹² National Association of Insurance Commissioners. (2024, October 3). *Implementation of NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers*.
<https://content.naic.org/sites/default/files/cmte-h-big-data-artificial-intelligence-wg-ai-model-bulletin.pdf.pdf>
- ¹³ National Conference of State Legislatures. (2024, September 9). *Artificial Intelligence 2024 Legislation*. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation#:~:text=Establishes%20a%20plan%20for%20the%20use>
- ¹⁴ Augenstein, J., Seigel, R., Fox, A., & Shashoua, M. (2024, July 31). *Manatt Health: Health AI Policy Tracker*. Manatt. <https://www.manatt.com/insights/newsletters/health-highlights/manatt-health-health-ai-policy-tracker>
- ¹⁵ An Act Concerning Artificial Intelligence, Connecticut S.B. 2 (2024).
- ¹⁶ Establishes NJ Artificial Intelligence Advisory Council, New Jersey S.B. 3357 (2024).
- ¹⁷ Artificial Intelligence Developer Act, Virginia S.B. 747 (2024).
- ¹⁸ An Act Relating to Regulating Developers and Deployers of Certain Artificial Intelligence Systems, Vermont S.B. 710 (2024).
- ¹⁹ Creates the Automated Decision Tools Act, Illinois S.B. 5116 (2024).
- ²⁰ An Act Relating to Regulating Developers and Deployers of Certain Artificial Intelligence Systems, Vermont S.B. 710 (2024).
- ²¹ Artificial Intelligence Developer Act, Virginia S.B. 747 (2024).
- ²² Automated Decision Systems, California S.B. 2930 (2024).

-
- ²³ Colorado Division of Insurance. *SB21-169: Protecting Consumers from Unfair Discrimination in Insurance Practices*. Colorado Department of Regulatory Agencies. <https://doi.colorado.gov/for-consumers/sb21-169-protecting-consumers-from-unfair-discrimination-in-insurance-practices>
- ²⁴ Levi, S. D., Kumayama, K. D., Ridgway, W. E., Ghaemmaghami, M., & Neal, M. M. (2024, June 24). *Colorado's Landmark AI Act: What Companies Need to Know*. Skadden, Arps, Slate, Meagher & Flom LLP. <https://www.skadden.com/insights/publications/2024/06/colorados-landmark-ai-act>
- ²⁵ *Office of Artificial Intelligence Policy*. Utah Department of Commerce. <https://ai.utah.gov/>
- ²⁶ Artificial Intelligence Amendments, Utah S.B. 149 (2024).
- ²⁷ Robles, C. (2024, September 11). *California Legislature Passes First-Ever Bill Specifically Regulating Payer AI Use*. Inside Health Policy. <https://insidehealthpolicy.com/daily-news/california-legislature-passes-first-ever-bill-specifically-regulating-payer-ai-use>
- ²⁸ California Code, HSC 1367.01. (2024, September 28). https://leginfo.legislature.ca.gov/faces/codes_displaySection.xhtml?lawCode=HSC&ionNum=1367.01
- ²⁹ Scott, K. (2024, September 5). *SB 1120 (Becker) – Request for Signature*. <https://calhospital.org/wp-content/uploads/2024/09/SB-1120-Request-for-Signature-final-09052024.pdf>
- ³⁰ Robles, C. (2024, September 11). *California Legislature Passes First-Ever Bill Specifically Regulating Payer AI Use*. Inside Health Policy. <https://insidehealthpolicy.com/daily-news/california-legislature-passes-first-ever-bill-specifically-regulating-payer-ai-use>
- ³¹ National Association of Insurance Commissioners. (2020). *National Association of Insurance Commissioners (NAIC) Principles on Artificial Intelligence (AI)*. https://content.naic.org/sites/default/files/inline-files/AI%20principles%20as%20Adopted%20by%20the%20TF_0807.pdf
- ³² Machledt, D., & Edwards, E. (2023). Principles for fairer, more responsive automated decision-making systems. In *National Health Law Program*. https://healthlaw.org/wp-content/uploads/2023/05/NHeLP_PrinciplesforFairerMoreResponsiveAutomatedDecisionSystems_05052023.pdf

-
- ³³ *About us*. National Health Council. <https://nationalhealthcouncil.org/about-us/>
- ³⁴ National Health Council. (2024, February 8). *Statement on Artificial Intelligence and Health Care: Promise and Pitfalls*. <https://nationalhealthcouncil.org/letters-comments/nhc-statement-on-artificial-intelligence-and-health-care-promise-and-pitfalls/#:~:text=AI%20applications%20in%20health%20care%20must%20be%20trustworthy%2C,oversight%2C%20and%20safeguarded%20by%20robust%20privacy%20and%20security.>
- ³⁵ *Coalition for Health AI*. CHAI. <https://chai.org/>
- ³⁶ American Medical Association. (2023, November 14). *Principles for Augmented Intelligence Development, Deployment, and Use*. <https://www.ama-assn.org/system/files/ama-ai-principles.pdf>
- ³⁷ AHIP. (2023, July 7). *Request for Information (RFI) Response: National Priorities for Artificial Intelligence (OSTP-TECH-2023-0007-0001) —AHIP Comments*. <https://ahiporg-production.s3.amazonaws.com/documents/OSTP-AI-RFI-Response-for-Posting.pdf>
- ³⁸ Norori, N., Hu, Q., Aellen, F. M., Faraci, F. D., & Tzovara, A. (2021). Addressing bias in big data and AI for health care: A call for open science. *Patterns*, 2(10), 100347. <https://doi.org/10.1016/j.patter.2021.100347>