# AI for Prior Authorization Appeals: Templates and Prompts

Published July 20, 2026 40 min read

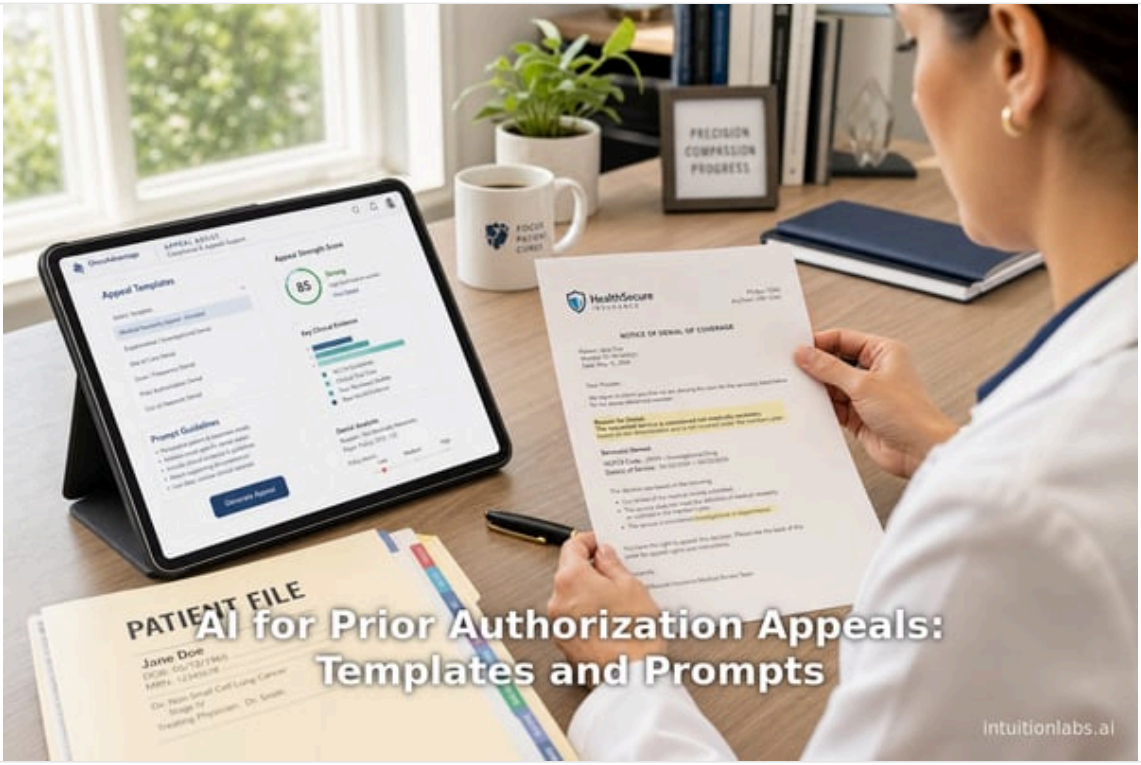


## Executive Summary

Prior authorization, the requirement that a health plan approve a treatment, test, or prescription before it is delivered, has become the most contested administrative process in American health care, and artificial intelligence (AI) now sits on both sides of the fight. In 2024, Medicare Advantage insurers alone processed nearly 53 million prior authorization requests and denied 4.1 million of them, a 7.7 percent denial rate (Source: www.kff.org). Only 11.5 percent of those denials were appealed, yet when patients or providers did appeal, 80.7 percent of the denials were fully or partially overturned (Source: www.kff.org). That gap, a low appeal rate against a high success rate, is the commercial opening that AI-assisted appeal tools such as **Claimable**, **Counterforce Health**, and **Fight Health Insurance** are built to close, alongside general-purpose chatbots like ChatGPT that patients and physicians increasingly use to draft appeal letters.

The physician side of the ledger is equally strained. The American Medical Association's (AMA) 2025 survey of 1,000 practicing physicians found that doctors complete an average of 40 prior authorizations per week, spend roughly 13 hours weekly on the process, and that 74 percent report denials have increased over the past five years (Source: www.ama-assn.org). Sixty percent of surveyed physicians say they are concerned that AI is already increasing, or will increase, denial rates (Source: www.ama-assn.org), a concern reinforced by litigation alleging UnitedHealthcare deployed an algorithm called **nH Predict** with a 90 percent error rate to override physicians' own determinations (Source: www.cbsnews.com), and by ProPublica's investigation showing Cigna doctors used an algorithm called **PXDX** to deny more than 300,000 claims in two months while spending an average of 1.2 seconds per case (Source: www.propublica.org).

This report catalogs the practical landscape of AI for prior authorization appeals as of July 2026: the taxonomy of tools (consumer AI appeal generators, general chatbots, EHR-embedded physician assistants, and payer-side automation), the anatomy of an effective appeal letter, concrete ChatGPT prompt structures, and the regulatory guardrails now constraining how insurers may use AI in coverage decisions. California's **Physicians Make Decisions Act (SB 1120)**, effective January 1, 2025, requires that any AI-assisted denial, delay, or modification of care be reviewed and decided by a licensed physician or qualified health care provider (Source: sd13.senate.ca.gov), and by 2026 four more states (Arizona, Maryland,

Nebraska, and Texas) had passed similar restrictions (Source: klrd.gov). At the federal level, the CMS Interoperability and Prior Authorization final rule (CMS-0057-F) requires impacted payers to issue decisions within 72 hours for expedited requests and seven calendar days for standard requests, and to publish a specific denial reason beginning in 2026 (Source: www.cms.gov).

On outcomes, vendor-reported figures should be read as marketing claims rather than independent research: Claimable states that over 80 percent of the appeals it generates succeed, most within 10 days (Source: www.getclaimable.com), while Counterforce Health cites a roughly 70 to 75 percent success rate for its AI-generated letters (Source: www.counterforcehealth.org). These figures are not directly comparable to the government's own appeal-outcome data, but they are broadly consistent with the pattern that appealed denials succeed far more often than the low rate of appeals would suggest. For life-science and health-technology organizations evaluating this space, the operative lesson is that the technology question (which AI tool to use) is inseparable from the compliance question (what human physician oversight and documentation the law now requires), a distinction that consultancies advising on responsible AI adoption in regulated health environments, including IntuitionLabs, treat as a first-order design constraint rather than an afterthought (Source: intuitionlabs.ai).

# Introduction and Background

Prior authorization is a cost-control and utilization-management process under which a health plan requires a physician to obtain the insurer's approval before delivering a prescription, procedure, or service in order for that service to be covered (Source: www.ama-assn.org). Virtually all Medicare Advantage enrollees, 99 percent, are required to obtain prior authorization for at least some category of service, most commonly high-cost items such as inpatient hospital stays, skilled nursing facility admissions, and chemotherapy (Source: www.kff.org). Outside Medicare Advantage, Affordable Care Act (ACA) marketplace insurers denied 19 percent of in-network claims and 37 percent of out-of-network claims in 2024, a combined average of 20 percent of all submitted claims (Source: www.kff.org). A separate KFF tracking poll found that two-thirds of insured adults, 66 percent, believe delays and denials by health insurance companies are a "major problem," and one-third say they personally experienced a denial of a physician-prescribed service or medication in the past two years (Source: www.kff.org).

Appeals exist precisely because insurers acknowledge their own initial determinations are frequently wrong.Of the marketplace claims denied in 2024, fewer than 1 percent were ever appealed, and of those that were, insurers upheld their own original decision 66 percent of the time, meaning roughly a third were reversed (Source: www.kff.org). In Medicare Advantage the pattern is starker still: only 11.5 percent of denials were appealed in 2024, yet 80.7 percent of those appeals were overturned, a rate that has stayed above 80 percent in every year examined since 2019 (Source: www.kff.org). A 2018 to 2022 review by the Department of Health and Human Services' Office of Inspector General (OIG) found that 13 percent of denied Medicare Advantage prior authorization requests actually met Medicare coverage rules and likely would have been approved under traditional fee-for-service Medicare (Source: oig.hhs.gov). KFF's own analysis notes plainly that rapidly developing AI tools "may reduce administrative errors that can lead to improper denials, predict whether a claim will be paid, and assist providers and patients in appealing a denial," while cautioning that federal oversight of this use case remains a work in progress (Source: www.kff.org).

This dynamic, an asymmetry where appeals rarely happen but usually work when they do, is what has drawn a wave of AI products into the prior authorization appeals space over the past two years. This report examines that landscape in depth: how the tools are built, how appeal letters should be structured, what prompts work when using general chatbots, what the law now requires of insurers' own AI systems, and what the available data shows about outcomes.

# Taxonomy of AI Tools Used in Prior Authorization Appeals

AI now touches prior authorization appeals from at least four distinct angles, each with a different user, business model, and risk profile. Understanding which category a tool belongs to matters more than any single feature, because the categories differ sharply in cost, oversight, and data-handling practices.

## Consumer-Facing AI Appeal Generators

A cluster of startups builds AI systems specifically to draft patient-facing appeal letters. **Claimable**, founded by Dr. Warris Bokhari, generates a customized appeal letter for roughly $40 that incorporates clinical research on the denied drug or treatment and the outcomes of other patients' appeals for that same therapy (Source: www.nbcnews.com). The company reports that one in five insured adults face a denied claim, that fewer than 1 percent of those denials are appealed, and that over 80 percent of Claimable's own appeals succeed, most resolved within 10 days (Source: www.getclaimable.com) (Source: www.getclaimable.com). **Counterforce Health**, funded through grants from the National Institutes of Health (NIH) and the University of Pennsylvania, offers its appeal-drafting service free to individual patients and reports having produced more than 10,000 appeals (Source: www.counterforcehealth.org), with a majority of those appeals ultimately approved, a figure discussed further below. **Fight Health Insurance**, an open-source-adjacent nonprofit-style project, states it has generated more than 10,000 appeals and lets users pay what they can,

including $0, for its AI-drafted letters (Source: www.fighthealthinsurance.com). All three follow a similar workflow: the patient uploads a denial letter and clinical documentation, the AI system matches that information against relevant clinical guidelines and prior successful appeals, and it produces a formal letter the patient can submit themselves or, in some cases, have the service fax on their behalf.

### General-Purpose AI Chatbots Used Ad Hoc

A second, much larger and less structured category consists of patients and clinicians simply prompting ChatGPT or a similar large language model (LLM) directly, without a purpose-built appeals product. This approach costs nothing beyond a chatbot subscription but places the entire burden of prompt construction, fact-checking, and citation accuracy on the user. Appeal-assistance vendors themselves publish guidance on how to do this better, cautioning that a vague instruction such as "write an appeal letter" produces a generic, low-quality draft, and that a detailed prompt specifying the denial reason, policy language, and clinical evidence performs markedly better (Source: www.counterforcehealth.org).

### EHR-Embedded, Physician-Facing AI

A third category integrates directly into clinical workflows rather than patient-facing portals. **Doximity Ask**, a clinical AI product built for physicians, generates insurance appeal letters, prior authorizations, and other administrative documents in seconds and markets itself as HIPAA-compliant and secure by default (Source: www.doximity.com) (Source: www.doximity.com). One physician user reported that after entering a patient's active diagnoses and treatment goal, "within 15 seconds my letter of clinical necessity was written, describing all the relevant conditions perfectly" (Source: www.doximity.com). This category matters because it addresses the physician-time problem directly: unlike consumer tools that operate after a denial has already occurred, EHR-embedded AI can draft the letter of medical necessity that is itself the strongest evidence in any subsequent appeal.

### Payer-Side AI for Utilization Management

The fourth category sits on the insurer's side of the transaction and is the one attracting the most regulatory scrutiny. **Cohere Health** builds AI-powered prior authorization platforms for health plans; a regional Blue-affiliated plan using Cohere's Unify platform reported achieving 93 percent digital adoption by providers shortly after launch, and the company states its broader customer base averages 96 percent digital adoption, a 64 percent reduction in call volume, and a 60 percent reduction in patient care delays (Source: www.coherehealth.com) (Source: www.coherehealth.com). This category illustrates that "AI in prior authorization" is not a single market: the same underlying technology can be deployed to speed up legitimate approvals or, as later sections describe, to accelerate denials at a scale and speed that draws legal challenge.

Advisory firms working with life-science and health-technology organizations, rather than selling any of these four categories of software directly, occupy a fifth, distinct position: helping regulated organizations decide which category of AI tool fits a given compliance posture and how to document human oversight of it, a role IntuitionLabs describes in terms of building AI that is "compliant by design" for regulated environments (Source: intuitionlabs.ai).

## How Prior Authorization Denials and the Appeals Process Work

Understanding what AI tools are drafting requires understanding the underlying legal process. When a prior authorization request is denied, the CMS Interoperability and Prior Authorization final rule (CMS-0057-F), published January 17, 2024, requires most impacted payers, including Medicare Advantage organizations, state Medicaid and Children's Health Insurance Program (CHIP) fee-for-service programs, Medicaid managed care plans, and qualified health plan issuers on the federally facilitated exchanges, to send a decision within 72 hours for expedited (urgent) requests and seven calendar days for standard requests (Source: www.cms.gov). Beginning in 2026, those payers must also provide a specific reason for any denial, regardless of whether the decision is communicated by portal, fax, email, mail, or phone (Source: www.cms.gov). By January 1, 2027, payers must also implement a Fast Healthcare Interoperability Resources (FHIR) Prior Authorization application programming interface (API) capable of communicating approvals, denials with reasons, or requests for more information electronically (Source: www.cms.gov).

Once a denial is issued, most insurance plans, whether employer-sponsored, marketplace, or Medicare Advantage, offer multiple levels of appeal: an internal first-level appeal, often a second internal appeal or a peer-to-peer review between the treating physician and a health plan medical director, and finally an external independent review conducted by a third party. Deadlines to file typically run from 30 to 180 days from the date of denial depending on the plan (Source: www.counterforcehealth.org), and missed deadlines are a common, avoidable cause of appeal failure. Marketplace enrollees in 2024 filed at least 5,881 external appeals, equal to 4 percent of all internal appeals that were upheld, after exhausting the internal process (Source: www.kff.org).

The peer-to-peer review step has become a particular flashpoint for AI-related trust concerns. As part of a June 2025 pledge by roughly 60 health insurers to streamline prior authorization, insurers committed to ensuring that all medical necessity denials would be reviewed by a licensed and qualified clinician (Source: www.ama-assn.org). Yet the AMA's 2025 follow-up survey found only 24 percent of physicians reported that such reviews were consistently conducted by appropriately qualified clinicians, and only 16 percent of physicians who participate in peer-to-peer reviews said the health plan's representative often or always had the appropriate qualifications (Source: www.ama-assn.org). Only 33 percent of surveyed physicians believed the 2025 insurer pledge would make a meaningful difference at all (Source: www.ama-assn.org). This is precisely the credibility gap that both patient-facing and physician-facing AI appeal tools are designed to route around: rather than trust a peer-to-peer call whose reviewer's qualifications are opaque, an AI-drafted letter forces a documented, evidence-based written record that is harder to dismiss without explanation.

## Anatomy of an Effective Prior Authorization Appeal Letter

Whether drafted by a human, a purpose-built AI tool, or a general chatbot, insurers and patient-advocacy organizations describe a consistent structure for a strong medical necessity appeal letter. A template published by the Patient Advocate Foundation opens with identifying information (the insurer's name, the appeals department address, the member's name, member identification number, and the claim reference number), followed by a direct statement of intent to appeal and a request that the insurer "thoroughly review the provided documents and reconsider the previous adverse decision" (Source: www.patientadvocate.org). A parallel physician-authored template from the American College of Foot and Ankle Surgeons (ACFAS) frames the letter as a formal memo addressed to the plan's appeals department, referencing the patient's policy number, group number, and claim number, and stating explicitly that supporting documentation demonstrates the procedure "was medically necessary, according to the medical necessity definition in my provider contract" (Source: www.acfas.org).

Counterforce Health's published guidance distills seven essential components that any effective appeal letter, human- or AI-drafted, should include (Source: www.counterforcehealth.org):

- **Personal information and claim details**: name, policy number, claim number, and the specific service or treatment denied, so the appeal is routed correctly.
- **A clear statement of intent to appeal**, explicitly referencing the denial letter and its date.
- **Brief patient history and medical necessity**, connecting the diagnosis and treatment history to why the denied service is required.
- **Expert medical opinion**, a statement from the treating physician explaining why the treatment is appropriate.
- **Supporting clinical evidence**, such as peer-reviewed studies or clinical guidelines.
- **Insurance policy references**, citing the specific plan language, medical policy, or coverage criteria the insurer appears to have misapplied.
- **A request for expedited review**, when the denial involves urgent or time-sensitive care.

Two documented mistakes recur across appeal-writing guidance. First, missing deadlines: the National Association of Insurance Commissioners (NAIC) has been cited by appeal-writing services as attributing roughly 15 percent of rejected appeals to late submission alone (Source: www.counterforcehealth.org). Second, vagueness: statements like "I need this treatment" fail because they do not connect the specific medical condition to the specific denied treatment or explain why alternative, cheaper options are clinically insufficient (Source: www.counterforcehealth.org). This is precisely where AI drafting tools argue they add value: a well-built prompt or purpose-built model forces the specific policy-criteria-to-clinical-evidence linkage that a rushed, emotional letter often skips. Claimable's own published sample appeal, for a migraine treatment denial, illustrates this discipline by directly citing a professional medical society's position statement ("the American Headache Society... recognizes CGRP-targeted therapies... as first-line options for migraine prevention") alongside the specific ACA coverage obligation the plan is alleged to have violated (Source: www.getclaimable.com).

## Writing Appeals with ChatGPT and Other AI Chatbots: Prompts and Best Practices

Because most patients do not have access to a purpose-built appeal generator, general chatbots such as ChatGPT have become a common do-it-yourself substitute, and appeal-assistance organizations have published explicit prompt-engineering guidance for this use case. The core recommendation is to never issue a bare instruction; instead, the prompt should supply the denial reason, the specific service or medication denied, the relevant policy language, and the desired tone in a single structured request (Source: www.counterforcehealth.org). Recommended prompt construction proceeds in five steps:

- **Gather documentation first**: the denial letter, clinical notes, and any relevant policy language before opening the chatbot at all.
- **Draft a detailed prompt** that specifies the denial reason, the treatment history, and quoted policy language, then explicitly asks the model to reference the specific denial reason, provide medical evidence of necessity, cite policy language, address any alternative treatments already tried,

and maintain "a respectful but firm tone" (Source: www.counterforcehealth.org).

- **Review and personalize the draft**, adding specific quotes from the treating physician, exact policy section numbers, and removing generic or repetitive language the model may have inserted (Source: www.counterforcehealth.org).
- **Include critical components**: a clear deadline reminder noting the appeal window and requesting a response timeline.
- **Follow up with supporting documents**, including a formal letter of medical necessity from the treating physician, rather than relying on the AI-drafted narrative letter alone.

This guidance implicitly acknowledges a real limitation of general-purpose chatbots for this task: they have no built-in access to the specific payer's medical policy documents, no verified database of prior successful appeals, and no mechanism to confirm that a cited clinical guideline or study actually exists and says what the model claims. Purpose-built tools attempt to close this gap by grounding output in curated data; Counterforce Health, for instance, states its AI "is built on a deep dataset of successful insurance appeals," drawing on curated clinical and regulatory sources rather than the chatbot's general training data (Source: www.counterforcehealth.org).

A second, less discussed limitation is privacy. Genevieve Kanter, a senior fellow at the University of Southern California's (USC) Leonard D. Schaeffer Center for Health Policy and Economics, has warned that clinicians who paste patient information into a consumer chatbot to draft appeal or clinical-necessity letters may be creating a Health Insurance Portability and Accountability Act (HIPAA) violation, because "once you enter something into ChatGPT, it is on OpenAI servers and they are not HIPAA compliant. That's the real issue, and that is, technically, a data breach" (Source: schaeffer.usc.edu). Kanter notes that 18 specific identifiers are considered protected health information (PHI) under HIPAA, including patient names (and nicknames), dates of birth, and admission or discharge dates, all of which must be scrubbed from any text before it is entered into a general consumer chatbot (Source: schaeffer.usc.edu). For patients drafting their own appeal about their own denial, this is a matter of personal choice rather than legal exposure; for clinicians and health-system staff drafting appeals on a patient's behalf, it is a live compliance question, and it is the primary practical reason vendors of clinician-facing tools such as Doximity Ask market their products explicitly as "HIPAA-compliant and secure by default" rather than leaving clinicians to use a general consumer chatbot, as discussed above.

*Table 1* below compares the major AI appeal-assistance tools discussed in this section along the dimensions most relevant to a patient or clinician choosing between them.

| TOOL | PRIMARY USER | COST | VENDOR-REPORTED OUTCOME | DELIVERY MECHANISM |
|---|---|---|---|---|
| **Claimable** | Patients (currently covers over 50 specialty treatments including migraine, GLP-1s, asthma, and rheumatology drugs) | Approximately $40 per appeal (Source: www.nbcnews.com) | Over 80% of appeals succeed, most resolved in 10 days or less (vendor claim, cited above) | Patient uploads denial and health story; AI drafts letter citing clinical research and policy compliance gaps (Source: www.getclaimable.com) |
| **Counterforce Health** | Patients, caregivers, and physicians (NIH- and university-grant funded) | Free for individuals | Roughly 75% of appeals get approved; over 10,000 appeals generated (vendor claim, cited above) | Upload denial, insurance plan, and documents; AI builds case and legal-grade letter, with an optional voice AI assistant |
| **Fight Health Insurance** | Patients (pay-what-you-want model, including $0) | $0 to user-chosen amount (Source: www.fighthealthinsurance.com) | Over 10,000 appeals generated; multiple AI models used per draft (vendor claim) (Source: www.fighthealthinsurance.com) | Free self-service drafting; user submits appeal themselves or pays a small fee for the service to fax it (Source: www.fighthealthinsurance.com) |
| **General chatbot (e.g., ChatGPT) via manual prompting** | Patients or clinicians with no dedicated tool | Cost of chatbot subscription only | Not independently measured; depends heavily on prompt quality and user follow-up | User manually gathers documents, drafts a detailed prompt following the structure described above, and personalizes the output before submission |
| **Doximity Ask** | Physicians and clinical staff | Included in Doximity's clinician platform | Not independently measured; physician-reported anecdote of a 15-second letter of medical necessity (Source: www.doximity.com) | Physician enters diagnoses and goal directly inside a HIPAA-compliant clinical AI interface, generating a letter in seconds (Source: www.doximity.com) |

The table shows a clear pattern: patient-facing consumer tools compete on price and reported success rate, while the clinician-facing tool competes on speed and HIPAA compliance within an existing clinical workflow. None of these tools independently verify their own success-rate claims through third-party audit, and none of the reported figures should be treated as equivalent to the government-published Medicare Advantage and marketplace appeal statistics discussed later in this report, which are drawn from mandatory insurer reporting to CMS rather than vendor self-reporting.

## Regulatory Guardrails: What the Law Says About AI in Prior Authorization

The rapid growth of AI-assisted appeal tools has been matched by a rapid tightening of state and federal rules governing how insurers themselves may use AI to make coverage decisions in the first place. California's **Physicians Make Decisions Act (SB 1120)**, authored by state Senator Josh Becker and effective January 1, 2025, was described by its author as ensuring that "decisions about medical treatments are made by licensed health care providers, not solely determined by artificial intelligence (AI) algorithms used by health insurers" (Source: sd13.senate.ca.gov). The law,

sponsored by the California Medical Association, which represents 50,000 physicians statewide, requires that AI or algorithmic tools used in utilization review rely on an individual enrollee's actual clinical record rather than generalized datasets, and that final medical necessity determinations be made by a licensed physician or qualified health care provider competent in the relevant clinical area (Source: sd13.senate.ca.gov) (Source: www.fenwick.com).

Other states followed quickly. During the 2025 legislative session, four additional states, Arizona, Maryland, Nebraska, and Texas, passed bills restricting payors from using AI as the sole basis for medical necessity or prior authorization determinations (Source: klrd.gov). Arizona's HB 2175 requires independent human review of claims and prior authorization requests before any denial; Nebraska's LB 77 prohibits AI output from being the sole basis of a medical necessity determination and requires disclosure of AI use to providers and enrollees; and Texas's SB 815 restricts payors to using AI only for administrative support or fraud detection, subject to inspection by the state insurance commissioner (Source: klrd.gov) (Source: klrd.gov). Colorado had already become the first state to regulate AI in insurance broadly, in November 2023, with parallel rules for health and auto insurers following in October 2025 (Source: klrd.gov).

At the national level, the NAIC's Model Bulletin on the Use of Artificial Intelligence Systems by Insurers was adopted by the Executive (EX) Committee and Plenary on December 4, 2023, and states that it "sets forth the Department's expectations as to how Insurers will govern the development/acquisition and use of certain AI technologies" (Source: content.naic.org) (Source: content.naic.org). A 2025 NAIC survey of 93 insurance companies across 16 states found 84 percent of insurers use AI or machine learning (ML) across their product lines, with 68 percent currently using or exploring AI specifically for prior authorization approval processes, and 12 percent using it for denying prior authorizations outright (Source: klrd.gov) (Source: klrd.gov).

Federally, CMS clarified in a 2024 FAQ memo that Medicare Advantage organizations may use AI or algorithmic systems to help make coverage determinations, but that any medical necessity determination must be "based on the circumstances of each specific individual" and reviewed by "a physician or other appropriate health care professional with expertise in the field of medicine or health care that is appropriate for the services at issue" under 42 CFR Section 422.101(c) (Source: klrd.gov). A proposed 2026 guardrail requiring MA plans to "ensure services are provided equitably, irrespective of delivery method or origin, whether from humans or automated systems" was ultimately not included in the final 2026 rule, though CMS stated it would "continue to consider the extent to which it may be appropriate to engage in future rulemaking in this area" (Source: klrd.gov).

*Table 2* summarizes the principal state and federal guardrails now governing AI use in prior authorization and utilization review decisions.

| JURISDICTION / BODY | RULE | EFFECTIVE DATE | CORE REQUIREMENT |
|---|---|---|---|
| **California** | SB 1120 (Physicians Make Decisions Act) | January 1, 2025 | AI may not be the sole basis for denying, delaying, or modifying care; a licensed physician or qualified provider must make the final medical necessity determination (Source: sd13.senate.ca.gov) |
| **Colorado** | AI insurance regulations | November 14, 2023 (life); October 15, 2025 (health and auto) | Insurers must report AI model governance and risk management frameworks (Source: klrd.gov) |
| **Arizona** | HB 2175 | 2025 | Requires independent human review of claims and prior authorization requests before any denial (Source: klrd.gov) |
| **Maryland** | HB 820 | 2025 | AI utilization review tools must use individualized clinical data and be open to state audit (Source: klrd.gov) |
| **Nebraska** | LB 77 | 2025 | AI output cannot be the sole basis for a medical necessity denial; AI use must be disclosed publicly (Source: klrd.gov) |
| **Texas** | SB 815 | 2025 | AI restricted to administrative support and fraud detection; subject to state insurance commissioner inspection (Source: klrd.gov) |
| **NAIC (multi-state model)** | Model Bulletin on AI Systems | Adopted December 4, 2023 | Requires insurers to adopt a documented AI governance program (Source: content.naic.org) |
| **CMS (federal)** | CMS-0057-F Interoperability and Prior Authorization Rule | Operational provisions January 1, 2026; API requirements January 1, 2027 | 72-hour/7-day decision timelines, published denial reasons, and a Prior Authorization API (Source: www.cms.gov) |

The consistent thread across these rules is that regulators are not trying to ban AI from prior authorization, but to prevent it from being the sole and unreviewable basis for a denial. That distinction is exactly what makes AI-assisted appeals a legally coherent countermeasure: an appeal letter that documents individualized clinical facts and forces human physician review is the mirror image of what these laws now require insurers to provide on the front end.

## Data Analysis and Evidence

The quantitative record on prior authorization denials, appeals, and AI's growing role spans several independent data sources that, taken together, paint a consistent picture of a high-denial, low-appeal, high-overturn system. *Table 3* presents CMS-reported Medicare Advantage prior authorization trends compiled by KFF across four selected years.

| YEAR | TOTAL PA REQUESTS | DENIAL RATE | SHARE OF DENIALS APPEALED | APPEAL OVERTURN RATE |
|---|---|---|---|---|
| **2019** | Not separately reported in this dataset | Not separately reported | 7.5% (Source: www.kff.org) | Over 80% in every year examined (Source: www.kff.org) |
| **2022** | Not separately reported in this dataset | 7.4% (Source: www.kff.org) | Not separately reported | Over 80% |
| **2023** | 49.8 million (Source: www.kff.org) | 6.4% (Source: www.kff.org) | 11.7% (Source: www.kff.org) | Over 80% |
| **2024** | 52.8 million (Source: www.kff.org) | 7.7% (4.1 million denials) (Source: www.kff.org) | 11.5% (Source: www.kff.org) | 80.7% (Source: www.kff.org) |

The table demonstrates two stable trends across five years of data: the denial rate has fluctuated in a narrow band around 6 to 8 percent, while the overturn rate for appealed denials has remained remarkably consistent above 80 percent in every year measured. This consistency is what makes the low appeal rate, just 11.5 percent in 2024, the central inefficiency that AI appeal tools target: if appeal rates rose toward 100 percent while overturn rates held near 80 percent, several million additional Medicare Advantage beneficiaries each year would receive care that was already, by the insurer's own later admission, medically appropriate.

Beyond Medicare Advantage, additional data sources corroborate the pattern of under-appealed but often-successful appeals. In traditional fee-for-service Medicare, CMS completed just over 625,000 prior authorization reviews in fiscal year 2024, of which 22.9 percent, fewer than 150,000 requests, were denied, a substantially higher denial rate than Medicare Advantage despite far fewer total requests, reflecting the narrower set of services subject to prior authorization in traditional Medicare (Source: www.kff.org). A 2026 OIG review of skilled nursing facility (SNF) admissions across 19 Medicare Advantage organizations found a 12 percent denial rate in June 2024, with only 18 percent of those denials appealed, yet 95 percent of appealed SNF denials were overturned in the enrollee's favor (Source: oig.hhs.gov) (Source: oig.hhs.gov). Notably, the contractor naviHealth, a subsidiary of UnitedHealth Group, processed half of all SNF admission requests reviewed and denied 14 percent of them, a higher rate than either internal MAO review (11 percent) or other contractors (9 percent), and MAOs later overturned 97 percent of naviHealth's SNF denials on appeal (Source: oig.hhs.gov).

On the cost side of automation, the CAQH Index, an annual industry benchmark, found in its 2024 report that fully automating prior authorization submission and response could save the health care industry $515 million annually and 14 minutes of provider and staff time per authorization by adopting electronic standards (Source: www.caqh.org). Physician-reported time costs remain substantial even before appeals begin: the AMA's survey found practices complete an average of 39 to 40 prior authorizations per physician per week and spend 13 hours weekly on the process, with 40 percent of physicians employing staff who work exclusively on prior authorization (Source: www.beckershospitalreview.com) (Source: www.beckershospitalreview.com) (Source: www.ama-assn.org). Only 20 percent of physicians said they always appeal an adverse prior authorization decision, while roughly two-thirds said they do not appeal if they believe the appeal will not succeed, and more than half do not appeal if they lack sufficient resources or time (Source: www.beckershospitalreview.com). That resource and time constraint is the exact friction AI appeal tools are engineered to reduce.

# Case Studies and Real-World Examples

## Claimable: AI-Drafted Appeals for Specialty Drug Denials

Bryan Kopsick, a 30-year-old professional caddy on the PGA Tour, had received an infusion of the drug Remicade every eight weeks since age 12 to control his Crohn's disease. When his health insurance changed to UnitedHealthcare on January 1, the new insurer required him to try cheaper alternatives before continuing Remicade, and he missed a dose while the dispute was unresolved, developing new inflammatory cysts within weeks (Source: www.bloomberg.com). Kopsick began working with Dr. Warris Bokhari, the physician and entrepreneur who co-founded Claimable three years earlier specifically to help patients appeal health insurance denials using AI (Source: www.bloomberg.com). In a separate Claimable case, Jason Nixdorf sought help after his wife Stephanie's insurer, Premera, denied a prescribed treatment for nine months; Claimable's chief executive personally helped draft an AI-assisted appeal letter, and two days after submission Premera approved the drug, with the approval letter itself apologizing that "you have been waiting to receive treatment for nine months" (Source: www.nbcnews.com). Claimable's co-founder Zach Veigulis, a

former chief data scientist at the Department of Veterans Affairs, said the company built its AI model beginning with rheumatology and migraine coverage and now supports over 50 treatment categories, with roughly 1,000 denials overturned through the platform as of the reporting period (Source: www.nbcnews.com).

## Counterforce Health: Free AI Appeals Funded by Research Grants

Counterforce Health, built on grant funding from the National Institutes of Health and the University of Pennsylvania, operates as a free service for individual patients and reports having generated more than 10,000 appeals (Source: www.counterforcehealth.org). One user, a healthcare professional referred to as Lee, reported that using the Counterforce software helped her overturn denials with unusual speed: "I've gotten back approvals on the same day and the day after... It definitely limits the time we have to spend formulating each denial letter and allows me to work more on initial prior authorization and patient assistance" (Source: www.nbcnews.com). In a separate patient testimonial published on Counterforce's site, a caregiver named Robert Chen described the platform guiding him through his mother's stroke rehabilitation therapy denial "identifying exactly which medical records to include and which insurance policy clauses to reference," resulting in the denial being overturned within 21 days (Source: www.counterforcehealth.org).

## The UnitedHealth nH Predict Litigation

In November 2023, the families of two deceased Medicare Advantage beneficiaries filed a federal class-action lawsuit in Minnesota alleging UnitedHealth Group knowingly used a faulty AI algorithm, developed by its subsidiary NaviHealth and known as **nH Predict**, to deny elderly patients coverage for extended care their own physicians had deemed medically necessary (Source: www.cbsnews.com). The complaint alleged the company deployed an AI model known internally to have a 90 percent error rate, overriding the treating physicians' own determinations of medical necessity, and that elderly beneficiaries were "prematurely kicked out of care facilities nationwide or forced to deplete family savings" because the AI model disagreed with their doctors (Source: www.cbsnews.com). A NaviHealth spokesperson countered that the AI-powered tool is "a guide to help [UnitedHealth] inform providers... about what sort of assistance and care the patient may need" rather than the sole basis for coverage decisions, adding that final decisions are "based on CMS coverage criteria and the terms of the member's plan" (Source: www.cbsnews.com). The plaintiffs additionally alleged the scheme depended on elderly beneficiaries lacking "the knowledge and resources to appeal the erroneous AI-powered decisions" (Source: www.cbsnews.com), a claim that directly foreshadows the market gap AI appeal-generation tools now target. As of February 2025, a federal court in Minnesota allowed the breach-of-contract portions of the case to proceed and waived the requirement that plaintiffs exhaust Medicare's internal appeals process first, citing the potential for irreparable harm (Source: klrd.gov).

## Cigna's PXDX Algorithm and the Limits of Automated Denial

A separate and earlier case illustrates the payer-side use of algorithmic denial at scale. ProPublica and The Capitol Forum's investigation of Cigna's internal claims-review system, known as **PXDX**, found that over a two-month period, Cigna doctors denied more than 300,000 requests for payment using an algorithm-driven review process, spending an average of 1.2 seconds on each case (Source: www.propublica.org). One Cigna medical director, Dr. Cheryl Dopke, denied roughly 60,000 claims in a single month, and rejected 121,000 claims in the first two months of 2022 alone, according to an internal corporate scorecard reviewed by the news organizations (Source: www.propublica.org) (Source: www.propublica.org). A former Cigna doctor described the process bluntly: "We literally click and submit. It takes all of 10 seconds to do 50 at a time" (Source: www.propublica.org). Dave Jones, California's former insurance commissioner, said "it's hard to imagine that spending only seconds to review medical records complies with the California law," referring to a requirement that insurers conduct a "thorough, fair and objective investigation" of claims (Source: www.propublica.org). Cigna disputed the characterization of PXDX, stating the system was designed to "accelerate payment of claims for certain routine screenings" and that reviews "occur after the service has been provided to the patient and does not result in any denials of care" (Source: www.propublica.org). This case, alongside the nH Predict litigation, is the direct regulatory catalyst behind laws such as California's SB 1120 and the state legislation summarized in Table 2.

## Cohere Health: Payer-Side AI Aimed at Reducing, Not Accelerating, Denials

Not every payer-side AI deployment has drawn the scrutiny that PXDX and nH Predict have. Cohere Health, describing itself as a leader in clinical intelligence solutions for health plans, announced in February 2025 that a regional Blue-affiliated health plan using its Cohere Unify platform achieved 93 percent digital adoption by providers shortly after launch, reducing time-consuming fax- and phone-based prior authorization submissions (Source: www.coherehealth.com). The plan reported reducing daily provider phone calls by one third as a direct result (Source: www.coherehealth.com).

Cohere's chief experience officer, Elif Eracar, stated the company's broader customer base sees averages of "96% digital adoption, 64% reduction in call volume, and 60% reduction in patient care delays" (Source: www.coherehealth.com). This case demonstrates that AI deployed on the payer side of prior authorization is not inherently adversarial to patients; the same underlying capability, automated matching of a request against coverage criteria, can be built either to speed up legitimate approvals and reduce the friction that leads to appeals in the first place, or, as the Cigna and UnitedHealth cases show, to accelerate denials at a volume that invites litigation and legislative response.

## Implications and Future Directions

Several converging trends will shape how AI is used in prior authorization appeals over the next several years. First, the CMS Prior Authorization API, mandatory for impacted payers by January 1, 2027, will for the first time give providers and, indirectly, appeal-assistance vendors a standardized electronic channel to both submit prior authorization requests and receive structured, machine-readable denial reasons (Source: www.cms.gov). A structured, coded denial reason is considerably easier for an AI system to parse and respond to than the current patchwork of fax, portal, and phone notifications, which should make automated appeal drafting materially more precise once the rule takes full effect.

Second, the described dynamic is increasingly framed by industry observers as an "arms race" between payer-side and patient-side automation. Counterforce Health's own marketing materials quote outside coverage describing the situation as "a battle of AI versus AI in the insurance world... and if the patient doesn't have a tool, 'you lose'" (Source: www.counterforcehealth.org). Whether or not that framing proves durable, it captures a real asymmetry: as payer-side AI adoption for prior authorization approaches 68 percent of surveyed insurers (Source: klrd.gov), patient- and provider-facing AI adoption for appeals is likely to grow correspondingly, not because AI appeals are inherently superior to human-drafted ones, but because the volume and speed of automated initial denials increasingly outpaces what an unaided patient or overburdened physician practice can manually contest.

Third, the regulatory settlement between AI efficiency and human oversight, evident in SB 1120, the four 2025 state laws, and the NAIC Model Bulletin, is likely to become the template other states and eventually federal regulators adopt: AI may assist with document review, evidence-matching, and drafting, but a licensed clinician must remain the accountable decision-maker for medical necessity determinations. For AI appeal-generation tools, this same principle cuts the other way and is arguably a feature rather than a limitation: the appeal-writing guidance surveyed in this report consistently recommends that a treating physician review, sign, and where possible personally author the medical necessity component of any AI-assisted appeal, precisely because insurers, courts, and independent reviewers are likely to weigh a documented physician judgment more heavily than an AI-generated narrative alone.

Fourth, for life-science and health-technology organizations building or evaluating AI in this space, whether internally for a health system's revenue-cycle function or as a vendor to providers or plans, the practical lesson from the regulatory record is that governance documentation (data provenance, human-review checkpoints, and audit logging of AI-assisted determinations) is no longer optional infrastructure but a direct compliance requirement under multiple state laws and the NAIC framework. Consultancies that specialize in regulated life-science and health-technology environments, including IntuitionLabs, frame this requirement not as a constraint bolted onto AI deployment after the fact, but as a design principle from the outset, an approach consistent with the broader pattern of building AI systems that are compliant with FDA, EMA, and equivalent regulatory frameworks by design rather than by retrofit (Source: intuitionlabs.ai).

## Frequently Asked Questions (FAQs)

**Is it legal for an insurer to let AI deny my prior authorization request without a doctor's review?** In California, no: SB 1120 requires that any denial, delay, or modification of care based on medical necessity be reviewed and decided by a licensed physician or qualified health care provider (Source: sd13.senate.ca.gov). Arizona, Maryland, Nebraska, and Texas passed similar restrictions in 2025 (Source: klrd.gov), and CMS guidance requires Medicare Advantage medical necessity determinations to be reviewed by a qualified clinician (Source: klrd.gov). Coverage in states without such laws is less clear, which is part of why 84 percent of insurers surveyed by the NAIC in 2025 reported using AI or ML across their product lines with varying internal governance (Source: klrd.gov).

**Can I use ChatGPT to write my own prior authorization appeal letter?** Yes, and appeal-assistance organizations publish specific guidance for doing so, recommending a detailed, structured prompt rather than a vague request (Source: www.counterforcehealth.org). Patients drafting their own appeal about their own denial face no HIPAA exposure, since HIPAA governs covered entities and their business associates, not individuals acting on their own behalf; the privacy risk described by USC's Genevieve Kanter applies specifically to clinicians and staff entering other patients' protected health information into a non-HIPAA-compliant consumer chatbot (Source: schaeffer.usc.edu).

**What is the actual success rate of AI-assisted prior authorization appeals?** Vendor-reported figures vary and should be treated as marketing claims rather than independently audited statistics: Claimable reports over 80 percent of its appeals succeed, while Counterforce Health reports roughly 75 percent, as detailed earlier in this report. These are broadly consistent with, though not directly comparable to, the government-reported

figure that 80.7 percent of appealed Medicare Advantage prior authorization denials were overturned in 2024 (Source: www.kff.org).

**How long do I have to file a prior authorization appeal?** Deadlines typically run from 30 to 180 days from the denial date, depending on the specific health plan, as discussed above, and most plans offer multiple levels of appeal, including an internal appeal, a peer-to-peer review, and an external independent review. Under the CMS Interoperability and Prior Authorization final rule, most impacted payers must issue a decision within 72 hours for urgent requests and seven calendar days for standard requests, and provide a specific denial reason beginning in 2026 (Source: www.cms.gov).

**Do AI-generated appeal letters still need a doctor's involvement?** Every structured appeal-letter template reviewed for this report, whether from a patient-advocacy organization, a physician specialty society, or an AI appeal-generation vendor, recommends including a statement of medical necessity from the treating physician as a core, non-optional component (Source: www.acfas.org).

## Conclusion

AI has entered prior authorization appeals from every direction at once: patients now draft appeal letters with consumer chatbots or dedicated platforms such as Claimable and Counterforce Health, physicians use EHR-embedded tools such as Doximity Ask to produce letters of medical necessity in seconds rather than hours, and insurers themselves increasingly rely on AI both to process prior authorization requests efficiently and, in cases that have drawn litigation and new state law, to deny them at a scale that has alarmed regulators, physicians, and courts alike. The underlying data explains why this technology has found a market: appeals are rarely filed, in the low single digits or low double digits of denials depending on the payer type, yet when they are filed they succeed at rates consistently above 70 to 80 percent across both government and vendor-reported figures. That gap between rare appeals and frequent success is not a marketing artifact; it is documented in mandatory CMS reporting, HHS OIG audits, and KFF's independent analysis of federal transparency data alike.

For patients and physicians deciding how to respond to a denial, the practical guidance converges regardless of which tool is used: gather the denial letter and clinical documentation first, build a specific rather than generic case connecting the individual's medical history to the denied treatment, cite the specific policy language and clinical guidelines the insurer's own criteria reference, secure a physician's statement of medical necessity, and track every deadline carefully, since a substantial share of failed appeals fail on procedural grounds rather than the underlying medical merits. AI tools, whether a general chatbot prompted carefully or a purpose-built appeal generator, can meaningfully compress the time this process takes, but none of the evidence reviewed in this report suggests they substitute for the physician judgment and specific evidentiary linkage that insurers, external reviewers, and the growing body of state law all now require as the accountable center of any medical necessity determination.

The regulatory trajectory, from California's Physicians Make Decisions Act through the CMS Interoperability and Prior Authorization rule to the NAIC's Model Bulletin, points toward a settled architecture rather than an unresolved fight: AI may assist at nearly every step of the prior authorization and appeals lifecycle, but a licensed human decision-maker must remain accountable for the medical necessity determination itself. Organizations building, deploying, or advising on AI in this space, across payers, providers, and the life-science and health-technology vendors that serve them, will find that architecture, not any single tool's feature set, is the durable constraint worth designing around.

---

Tags: ai for prior authorization appeals, prior authorization appeal letter template, chatgpt prompts for insurance appeals, medical necessity appeal letter, prior authorization denial appeal, ai in healthcare, claimable, counterforce health, cohere health, insurance claim denial appeal

## IntuitionLabs - Industry Leadership & Services

**North America's #1 AI Software Development Firm for Pharmaceutical & Biotech:** IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

**Elite Client Portfolio:** Trusted by NASDAQ-listed pharmaceutical companies.

**Regulatory Excellence:** Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

**Founder Excellence:** Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

**Custom AI Software Development:** Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

**Private AI Infrastructure:** Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

**Document Processing Systems:** Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

**Custom CRM Development:** Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

**AI Chatbot Development:** Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

**Custom ERP Development:** Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

**Big Data & Analytics:** Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

**Dashboard & Visualization:** Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

**Leading Claude & ChatGPT AI Enablement and Training:** Help biotech and pharmaceutical teams adopt Claude and ChatGPT through practical training, governed workflows, use-case development, and implementation designed for regulated environments.

**AI Consulting & Training:** Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at intuitionlabs.ai/contact for a consultation.

## DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by Adrien Laurent, a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.