sprout.ai

2026

AI FOR HEALTH INSURANCE CLAIMS

Automate health claim decisions from policy and claim context, not rules

OVERVIEW

Process more claims per adjuster, at a lower cost per claim, with outcomes that stay consistent as volume increases

Health claims volume increases faster than the teams that handle them. Every claim needs someone to read the policy, review the evidence, and form a view, so output is capped by headcount, and cost per claim rises with volume rather than falling.

Sprout.ai interprets the policy alongside the claim, then recommends the coverage outcome, citing the relevant clause and supporting evidence. Clear claims resolve on an automated recommendation. The rest reach an adjuster with the analysis done and the specific uncertainty named. Insurers running Sprout.ai handle more claims per adjuster, process them faster, and hold accuracy steady as volume grows.

Sprout.ai recommends. Your team decides. Every recommendation carries the evidence, the clause, the rationale, and a confidence score, so an adjuster can review, override, or explain any outcome - and you configure which recommendations require review before they take effect.



80%

of claims processed from intake to outcome, without manual intervention

The remaining

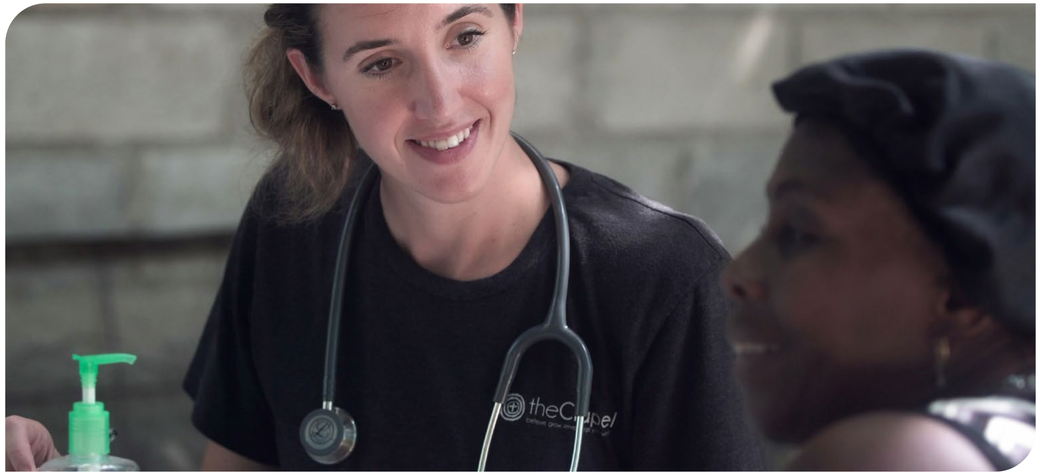
20%

arrives with the confidence gap already identified.



THE CHALLENGE

Cost per claim rises with volume when every policy decision needs an adjuster



Generic automation tools provide rules engines for handling simple claims. Everything else falls to manual review.

Health claims require careful assessment: does this treatment, for this member, under this policy, qualify, and for how much? Generic automation tools handle the narrow version of that question, but break down as soon as an accurate answer depends on context about the condition, the treatment, or the coverage itself. Those claims land on a desk, and turnaround, consistency, and cost move in the wrong direction together.

Policy complexity beyond what rules can express

A single treatment can map across 150+ coverage benefit options, each carrying its own limits, sub-limits, waiting periods, and cost-sharing. That structure can be modelled, but the benefit name alone rarely settles the claim. Whether a particular admission qualifies as a hospital stay, or a procedure counts as day-case rather than outpatient, depends on the wording behind the benefit and the clinical detail in front of you. Rules need those calls made in advance; in practice they are matters of interpretation, and they change with the schedule.

Medical knowledge that rules can't hold

Deciding a claim means knowing what a diagnosis implies, which treatments follow from it, and how the two map to a coding space running to tens of thousands of ICD, CPT, pharmacy, and surgery codes. Then it means reading how those codes were billed, whether the procedure code matches what was documented, whether line items have been unbundled, and whether the modifiers stack up. That understanding sits with experienced clinical, coding, and billing staff. It isn't a lookup, and it can't be written down as rules.

Medical evidence overload

Dense clinical documentation hides the facts that determine coverage. Adjusters routinely work through tens to hundreds of pages per claim, and the question is rarely answered on any single page. Tracing how a condition developed, when symptoms first presented, whether this episode is a recurrence or a new event, which treatment followed which diagnosis - means reconstructing a history across years of records written by different clinicians for different purposes.

Fragmented, high-volume claim data

Claims and supporting documents arrive from hundreds of providers in inconsistent formats, channels, and languages. Some documents are handwritten, some machine-printed, many a mix of both on the same page. With up to 500+ different invoice types, conventional extraction struggles to scale.

Manual review caps throughput

When claims require review inside the core claims management system, adjusting time per claim sets the ceiling. Growing volume means growing headcount just to hold turnaround steady.

Consistency degrades under pressure

Your most experienced adjuster and your newest adjuster should reach the same answer on the same claim. Under caseload pressure they don't, and inconsistencies across documents and claim history are hard to spot without cross-claim analysis.

THE CHALLENGE

From claim data to decision

Claim dimension	What Sprout.ai contextualises	What it returns
Coverage	Policy documents, riders, endorsements, and benefit schedules, against the treatment claimed	Whether the treatment is covered, partially covered, or excluded, and the clause that says so
Eligibility	Membership records, policy dates, waiting periods, network status, and cost-sharing terms	Whether the member was covered on the date of treatment, and the portion they owe
Clinical evidence	Doctors' reports, discharge summaries, treatment records, and prior authorisation letters	The facts that decide the claim, plus any conflicts or gaps in the evidence
Coding and billing	Medical invoices and bills, itemised line items, diagnosis and procedure codes	Whether the codes match the treatment documented, and which lines are payable
Claim integrity	The full claim file, alongside claim and provider history	Duplicates, unbundling, and patterns worth investigating
Completeness	Everything submitted, against what the policy requires	What is missing, and who needs to supply it
Recommendation	All of the above, taken together	The recommended outcome, a confidence score, and the automation path - straight through, or adjuster review

HOW SPROUT.AI IMPROVES CLAIMS PERFORMANCE

More claims per adjuster Lower cost per claim. Faster turnaround. Same standard of decision.

How can health insurers automate claim decisions without building rules?

Sprout.ai runs a set of specialist AI agents on every claim, each assessing a different dimension, and summarises their findings into a single recommendation. Where confidence is high, the claim is processed without adjuster intervention. Where it isn't, the claim routes to an adjuster with the analysis attached and the uncertainty named.

Sprout.ai recommends processing up to 80% of health claims straight through, from intake to outcome, without manual intervention. For the remaining claims, adjuster review time is 67% faster. The calculation is based on hundreds of thousands of claims, with no filtering or exclusion based on claim type, complexity or eligibility.

Extraction and evidence AI agent

Reads every document in the file including claim forms, provider invoices, clinical records, and establishes what the claim actually says. Confirms the file is complete against what the policy requires and identifies what is missing.

Fraud AI agent

Examines patterns across documents in the file and across claim history, surfacing duplicates, unbundling, and anomalies that single-claim review cannot see.

Coding AI agent

Works from the claim context, the claim form, the invoice, the clinical record, to determine the ICD, CPT, and pharmacy codes that represent the treatment actually given, drawing on external clinical and coding references where the file is ambiguous. Those codes are what the coverage assessment then runs against. Where the provider has already coded the claim, it checks their coding against what the documentation supports.

Policy and coverage AI agent

Maps the treatment to the benefit schedule and applies the policy terms - limits, sub-limits, exclusions, waiting periods, cost-sharing, network status - including any riders or endorsements attached to the member's cover, which routinely change the answer the base schedule would give. Returns the coverage position, the clause that governs it, and the reasoning that connects the two.

Tuned by your adjusters, not by developers

Every override an adjuster makes is feedback. Sprout.ai learns from the decisions your team actually makes, so accuracy improves against how you interpret your own policies, rather than how a rule was written. Configuration handles the thresholds, benefit interpretation, and new schedules, without a development cycle or the need for continuous improvement.

The recommendation

The AI agents' findings resolve into one output per claim: the recommended outcome, the clause applied, the evidence considered, and a confidence score. Every element is grounded in the claim file and the policy wording, traceable back to the document and clause it came from.

Straight-through processing on clear claims

Sprout.ai automatically processes straightforward claims without adjuster intervention, while claims that need a human decision are routed to an adjuster with the relevant information already assessed and a clear explanation of what still needs reviewing. You can define which types of claims and outcomes can be handled automatically, while ensuring specific scenarios always receive adjuster review.

THE KPIS SPROUT.AI MOVES

- STP rate
- Touch time per claim
- Claims turnaround time
- Claims per adjuster
- LAE / ALAE
- Cost per claim
- Recommendation accuracy
- Decision consistency
- Override rate

19%

increase in customer satisfaction (CSAT)



67%

reduction in adjuster handling time



>23%

faster turnaround times



sprout.ai

About Sprout.ai

Sprout.ai is a London-headquartered global technology company transforming insurance claims with AI-powered automation and decision intelligence. Since 2018, Sprout.ai has partnered with leading insurers, MGAs, and TPAs across major markets to deliver real-time, accurate, fair, and compliant decision-making at scale. Customers include AXA, MetLife, AdvanceCare, and Total Scholastic Solutions (TSS).

With proven and rapid enterprise deployment, Sprout.ai helps insurance teams boost efficiency, strengthen compliance, and deliver exceptional customer experiences. Sprout.ai injects intelligence into insurance claims - changing insurance, for good.

Changing insurance, for good.

Contact our claims experts for more information

Contact us

Book a demo

© 2026 Sprout.ai Limited. All rights reserved.