

Outline	 Outline - Human-in-the-loop AI
KW	Human in the loop AI (390) HITL (2,900) Human in the loop (1,300)
Meta description	This article explains how customer teams can integrate human-in-the-loop AI (HITL) into workflows, and the benefits this has over automating tasks.
Slug	/human-in-the-loop-ai
Image files	 human-in-the-loop-ai

Human-in-the-loop AI (HITL): AI With Oversight for Customer Teams

AI is everywhere, sure. But a stream of high-profile goof-ups really highlight the tradeoff between AI efficiency and human judgment. AI can do wonders for productivity, but it's still not ready to do it on its own — especially when it comes to business-critical situations.

So, how can we use AI to handle the headaches that come across our desks — without accidentally setting the world on fire?



It's simple: get humans involved. Human-in-the-loop AI (or HITL AI for short) gives you the best of both worlds: AI automation, but with humans tagged in at critical decision points to approve the AI's actions.

Let's take a closer look at HITL AI and how you can integrate it into workflows — and why it's sometimes better than fully autonomous systems that plow ahead without bothering to check in how they're doing.

What is human-in-the-loop AI (HITL)?

With human-in-the-loop AI, humans maintain oversight over AI's decision making. Unlike standard AI agents that automate tasks without human intervention, HITL AI adds control steps in AI workflows where a human needs to weigh in. It's like getting your intern to check with you before issuing a \$20,000 refund, instead of letting them go hog wild with your bank account.

By approving the decisions that an AI automation tool recommends, humans can ensure they're getting consistent and reliable outcomes from their AI workflows.



Caption: In a human-in-the-loop workflow, humans can support AI by validating and correcting processes before generating an output.

Alt Text: Flow chart showing a standard HITL AI interaction, where humans validate and authorize AI's decision making.

The HITL process can also improve how the AI itself functions: by incorporating human feedback to approve responses, HITL AI systems can adjust how they work — improving data labelling and decision making, and narrowing down what a “good” response looks like.

When should you use HITL?

HITL AI workflows and fully autonomous AI are both great — but they have different situations in which they’re most useful.

Fully automated AI systems work well for routine tasks that are both well-defined and low-risk — things like logging notes from a sales call in a CRM or sending knowledge base links to a customer with a question.

HITL, on the other hand, is better suited for edge cases and complex decisions with big repercussions — ethical, financial, or reputational. If you’re negotiating million-dollar deals, calming down an angry customer, or approving a large loan, you may want AI to help you out, but you don’t want it calling all the shots.

Why is a human in the loop necessary?

Today’s large language models (LLMs) often struggle to interpret a user’s real intentions. Instead of giving them an accurate answer, it returns a response that it “thinks” the user wants — whether that answer is correct or not.

If you’ve spent any time interacting with LLMs, you know how prone AI is to errors, hallucinating responses with high confidence and perpetuating biases stored in its data. Without a proper check on it, the results can cause real-world problems — such as when [Air Canada’s](#) chatbot gave a passenger non-existent discounts, leading to legal ramifications.

Not only that, but AI models aren’t always truthful. A study by [Palisade Research](#) revealed that advanced AI models like those from DeepSeek and OpenAI engaged in deceptive strategies when losing games of chess. This shows that AI is willing to behave unethically to achieve an outcome — which isn’t something you want from your support chatbot or AI agent you’ve tasked with cold emailing prospects.

Right now, the most effective way to handle these issues is by screening the responses from AI with a human in the loop. By having humans contribute their intuition, emotional intelligence, and industry-specific expertise (things that AI still struggles with), the overall system gets a better outcome, and the AI develops better judgement through additional context. All together, the whole setup is a safer and more inclusive AI ecosystem.

The benefits of HITL to AI in your business

HITL can help your sales, support, customer success, and other teams make the most of AI without introducing a lot of potential risks.

- **Improved accuracy and reliability:** Human oversight reduces AI errors and ensures responses are in line with your intentions. AI can flag uncertain responses for human review, which can improve the quality of future output and reinforce positive AI behavior.
- **Handling edge cases:** AI can struggle with rare, complex, or nuanced situations. For example, public AI tools may lack expertise in specialized fields for industry-specific applications. This makes human input essential to resolve cases that AI can't reliably handle.
- **Compliance and safety:** Who's accountable if an AI agent bankrupts you by making bad trades on your behalf? In regulated industries like finance, healthcare, and legal fields, AI outputs must meet strict compliance standards. Safety needs to be regulated on two fronts: by not trusting false outputs, and not exposing sensitive proprietary data to public AI platforms.
- **Cost optimization and streamlined workflows:** HITL lets businesses use AI to serve more customers and scale without overburdening human teams, while still keeping control over important decisions. You can automate processes like updating licenses or generating invoices, then run it by a human to avoid the risk of AI mistakes making it to customers.


What does a human-in-the-loop AI workflow look like?

A workflow is just a sequence of operations for accomplishing a goal. AI workflows lean on AI technologies to automate and optimize the tasks that make up the workflow, so you can reach your goal more efficiently.

While it may seem counterintuitive to add humans back into an automated workflow, HITL still reduces the overall work and adds clear benefits for certain use cases.

HITL AI workflows can be as simple as a one-step approval in an AI automated chain, or they can involve several steps back and forth between humans and AI until a final response is delivered. Multi-step HITL AI workflows can even involve multiple humans and AI agents. But every workflow hinges on clearly defining where human control is needed — and which parts you can automate.

The example below shows an AI workflow that assists support teams in approving refunds by having the AI generate initial responses based on company policies, evaluating the responses and refining future responses based on feedback from representatives.



Caption: How HITL can be integrated into a sales workflow

Alt Text: An HITL AI workflow diagram for sales representatives to review the responses to sales leads.

As seen above, a core feature of HITL workflows is confidence-based intervention, where response patterns are analyzed to automate some HITL points in the system, if similar cases were previously handled smoothly by the AI. By automating straightforward cases and escalating uncertain cases to humans, your AI system maintains efficiency, while enjoying the reliability offered by HITL.

How does HITL AI help customer teams?

The specific way you incorporate HITL AI into your workflows will vary across customer-facing teams.

Sales teams can improve their lead qualification workflows by incorporating HITL AI. They can use AI to find leads and score them based on selected initial qualifications, then loop in your human sales reps to close well-qualified leads. Alternatively, you can use AI agents to

automatically generate emails, then have human sales representatives refine them based on their engagement rate.

Support teams can benefit from HITL AI-assisted ticket management, where autonomous agents automatically handle low-risk cases to reduce overall resolution time. Teams can use humans to check whether or not tickets are correctly categorized based on risk levels. AI agents can resolve “low risk” tickets on their own, leaving more difficult and sensitive issues for human operators. Similarly, in escalation workflows, an AI sentiment analysis model can automate basic issues, and humans can handle queries flagged as too emotional or complex.

Customer success teams can prevent churn by following the same approach: AI can proactively perform sentiment analysis on usage data, then prompt humans to take action before a customer becomes frustrated — or even before they know there is an issue.

The key to HITL being effective in all these scenarios is to [consolidate data from different sources](#) so that your AI tools and HITL operators have complete, up-to-date information and context, and to make sure your team has [user-friendly, optimized interfaces](#) for assessing AI-generated output and guiding its responses.

Potential challenges HITL faces

By now, we know that assuming AI can handle any task without human oversight can lead to errors that are detrimental to your business. HITL keeps you from falling prey to an overreliance on AI.

That said, if implemented carelessly, HITL AI can create its own problems.

Poorly defined [AI workflows](#) with unclear steps and responsibility can do more harm than good. For example, if the escalation paths between AI and humans aren't well defined in a support ticket workflow, how can anyone know when AI has control — or when humans should step in? This scenario can create inconsistent feedback loops that make the AI responses difficult to refine and compromise the whole workflow's effectiveness.

At the same time, if the humans in a HITL workflow are unqualified, automated processes can also slow down, as they don't have the domain knowledge and or skills to annotate, validate, and work with AI appropriately.

The long and short of it is that AI workflows need to be carefully thought out and implemented with flexible tools that allow for [iteration and improvement](#).

How HITL is set to change with AI in 2025

HITL is going to find a lot more use throughout 2025 and the next few years as technologies continue to improve and develop — but not overlap with the ability of humans in many of the situations they're needed.

- Current AI trends are indicating a surge in more adaptive models with the help of [autonomous agents](#). The self-correctability of AI will improve, requiring less frequent human intervention.
- Smarter routing will enable automated agents to incorporate HITL by analyzing human feedback patterns and predicting when human oversight is necessary.
- The creation of HITL workflows is also shifting from static 2-D interfaces to more dynamic software that integrates voice and canvas-based workflows to interact with agents.
- HITL AI can now be integrated with enterprise workflows using [embedded AI agents](#) in CRM, customer support, and decision making platforms.

Appsmith is at the forefront of this change with Appsmith Agents. Our agentic AI platform lets you create agents for your most used SaaS tools and databases that can execute automated workflows with humans in the loop.

Unifying AI automation, HITL, and business apps

Appsmith's new features enable you to integrate HITL oversight into your sales, support, and customer success [AI workflows](#) in just minutes.

Appsmith Agents lets you [create your own AI agents](#) that can safely query all [your existing databases and datasources](#), while retaining relevant context about your organization. It deploys via a Chrome extension, letting you use AI agents that are tailored to your business anywhere in your toolchain.

Ready to get started? Try Appsmith Agents yourself today at <https://appsmithai.com>.