zingtree

Hallucination-Proof AI for Complex CX

The CX Leader's Guide to
Deterministic, Trusted Automation

AI is no longer optional, and accuracy is non-negotiable

CX leaders **face pressure to "add AI" everywhere** — from chatbots and call centers to back-office operations and analytics.

Investment in AI pilots is rising fast; however, successful graduation to production has significantly lagged.

Many teams are sprinting to test AI without the same governance and implementation discipline they apply to other systems, exposing the business to new risks.

Herein lies the disconnect: AI has yet to earn the CX leader's trust. Particularly in complex and highly regulated industries like healthcare, finance, and insurance, excitement for AI and AI Agents must be tempered with pragmatic caution.

DEGREE OF PRESSURE TO ADOPT AI

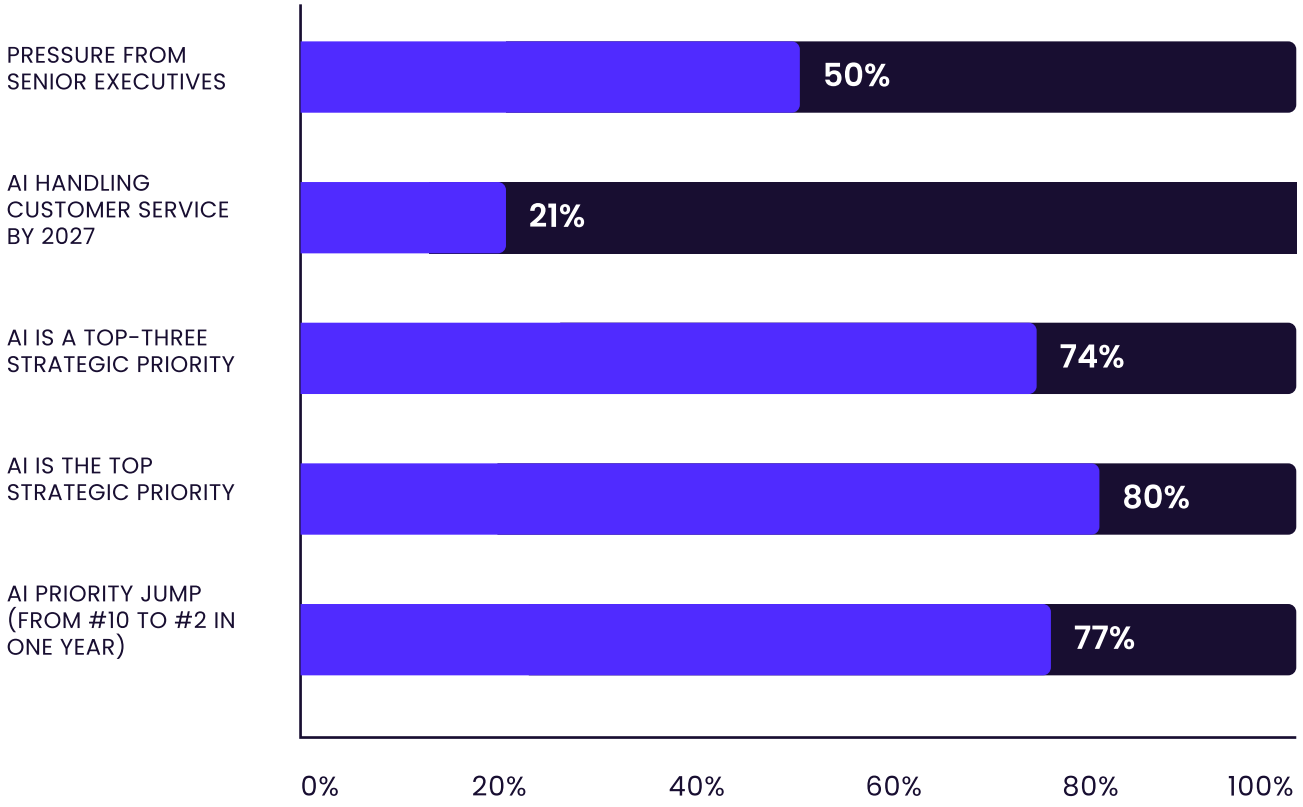

SOURCES: GARTNER | SALESFORCE | BAIN | IDC | TIMES NOW NEWS

When AI gets it wrong, enterprises pay the price

In February 2024, Air Canada's customer-service chatbot gave a grieving passenger false information about bereavement refunds. Relying on the AI's confident answer, the passenger booked full-price tickets and later learned refunds weren't allowed after purchase.

When the case reached the British Columbia Civil Resolution Tribunal, the airline argued that the chatbot was a separate legal entity.

The tribunal rejected the defense and ruled that Air Canada was liable for negligent misrepresentation.

The company was ordered to refund the fare and shut down its chatbot.

In regulated sectors, one wrong AI answer can turn into:

- **A compliance breach** — as seen when Cigna's AI auto-denied claims after 1.2-second "reviews."
- **A public-trust crisis** — like UnitedHealth's denial algorithm overriding physicians.
- **A data-liability event** — such as Blue Shield of California's "shadow AI" pixels leaking patient data.
- **A discrimination lawsuit** — as with Earnest Operations' AI underwriting, that penalized graduates of historically Black colleges.

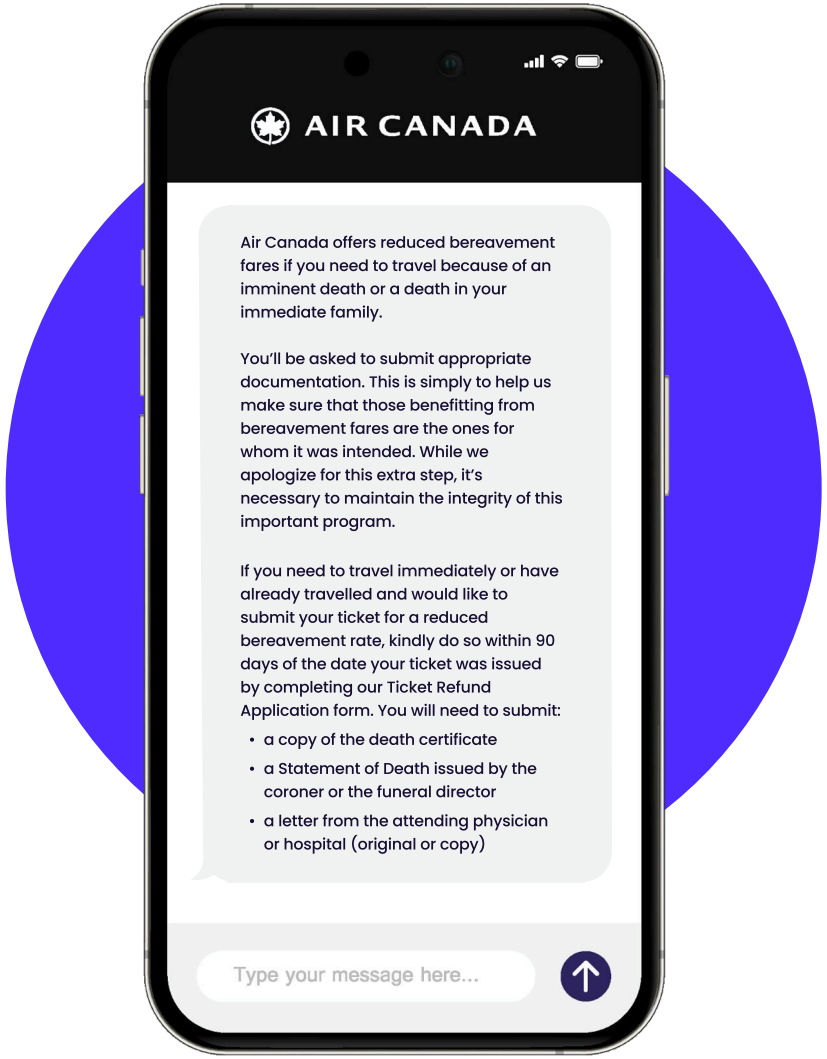

Not all automation should be left to AI

Generative AI isn't one-size-fits-all.
Different CX tasks demand different levels of control, context, and oversight.

AI exists on a spectrum — from fully deterministic with low autonomy and output variance, to probabilistic with high autonomy and output variance.

The higher the business or compliance risk, the more structure you need.

Automation Type	Business Risk	Use Case Example
Agentic AI	Low	Password reset
AI Agents	Low-to-Medium	Insurance Policy Recommendations
Hybrid (AI + Rules)	Medium	Technical Troubleshooting
Rule-Based Automation	High	Denied Claim Dispute

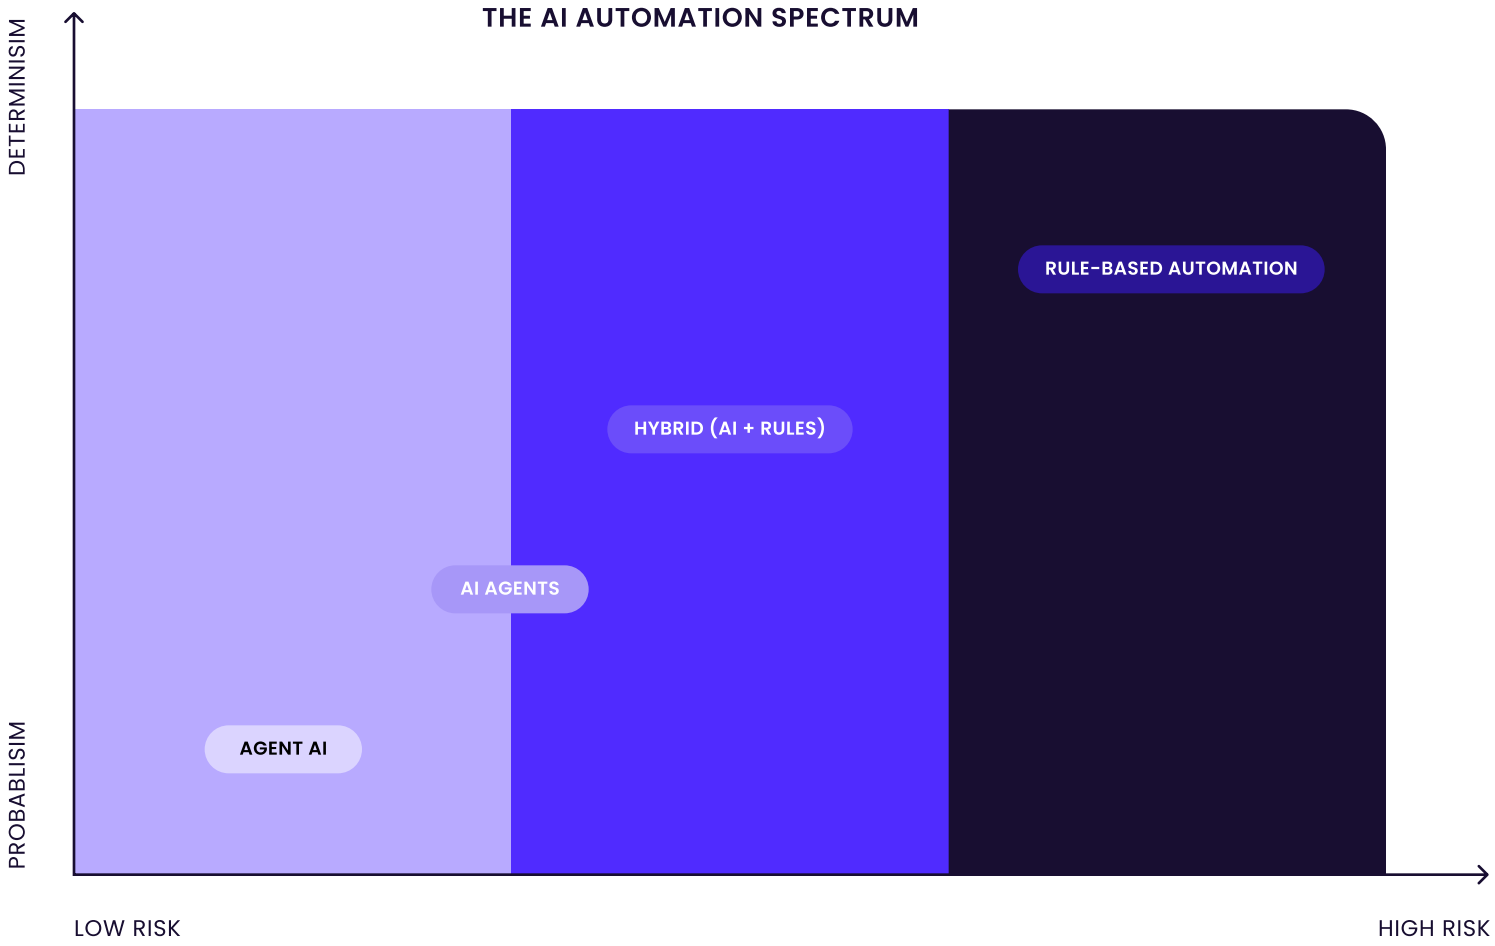

Generative AI predicts words, not truth

After a year of headlines about chatbots inventing facts and executives apologizing for “AI errors,” one term defined 2024: hallucination.

AI hallucinations are events in which LLMs (Large Language Models), like ChatGPT or Gemini, produce outputs that sound confident and correct but aren’t true.

To understand why hallucinations happen, it helps to look at how LLMs actually work.

LLMs are probability engines built to generate answers. They rely on statistical patterns in vast datasets to generate human-like text by predicting the most likely next word or sequence.

Just like meteorologists predict tomorrow’s weather using historical patterns and satellite data, LLMs predict the next word in a sentence.

A weather forecaster doesn’t know it will rain. They calculate the probability that it might. An LLM works the same way.

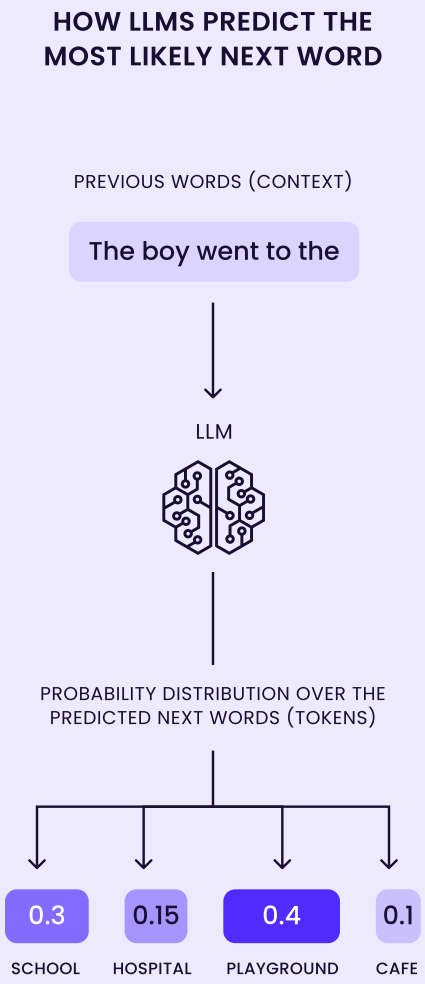

Complex questions introduce more AI variance and errors

The more complex the question, the higher the variance of possible AI answers, further exacerbating hallucination.

Ask an AI assistant a straightforward question like, "How do I view my bill?" The model has one overwhelmingly likely answer. It's factual, direct, and low risk.

But ask something more operationally complex, like, "How do I appeal a denied insurance claim?" Now AI is juggling variables: compliance rules, policy exceptions, tone, escalation paths, and customer sentiment.

LLMs hallucinate when they lack:

- Context: who the customer is, what systems they use
- Business logic: your policies or rules
- Guardrails: compliance and traceability
- System awareness: access to live data or the ability to act

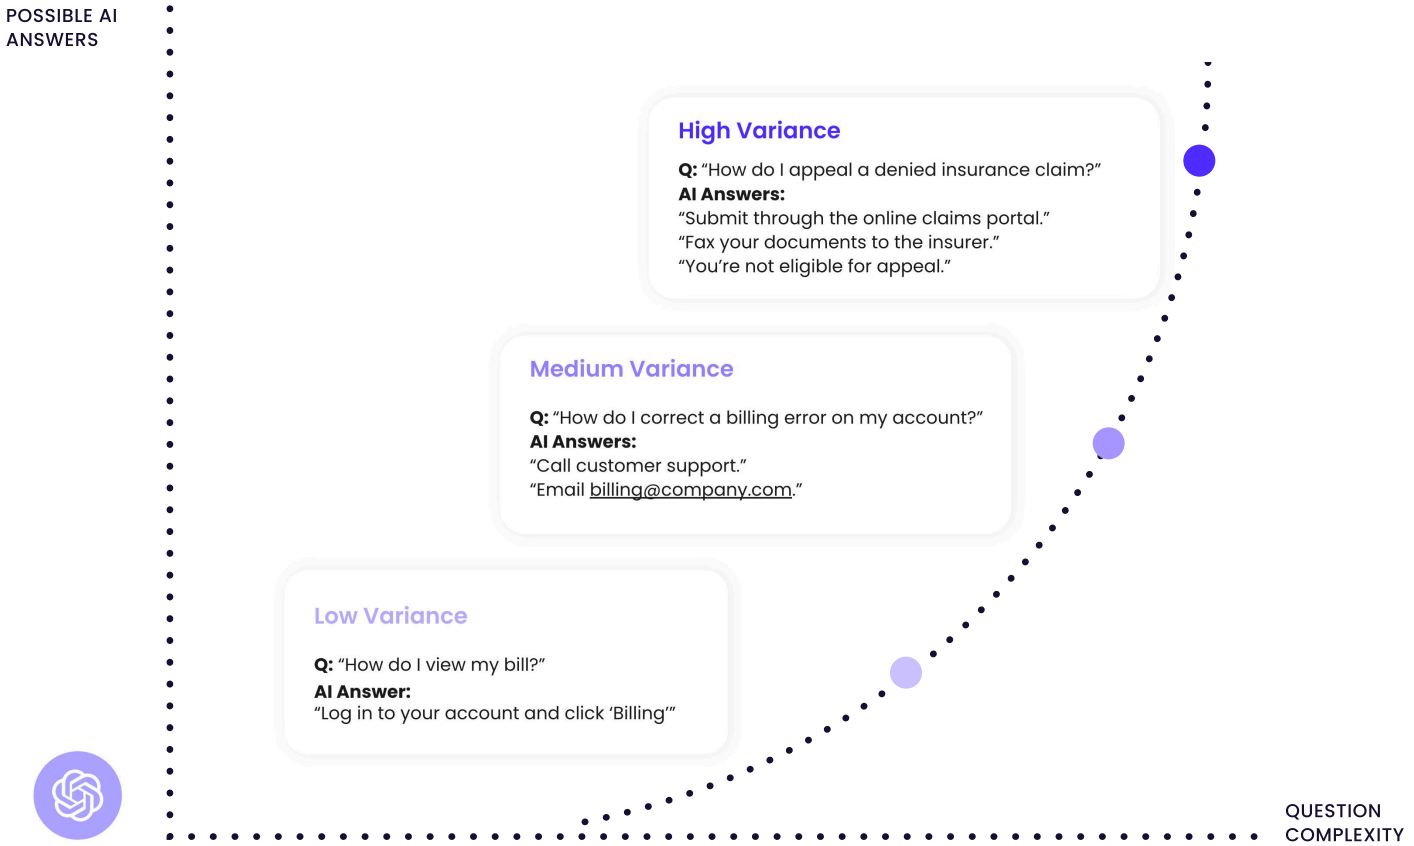

Deterministic guardrails:
Turning prediction into precision

Just like an intern needs a playbook on day one, **LLMs need a business-compliant structure** where they can explore possible answers.

Enter: Deterministic AI. Because an LLM doesn't "figure it out" on its own. You tell it what's allowed, what data to use, and when to escalate in order to shrink the variance for greater control and reliability

Think of Deterministic AI as a funnel of control. Before responding, AI passes through three guardrails:

- **Intent** – Detect what the user wants
- **Business rules** – Define what's allowed
- **Context** – Injects real-time truth

Each layer narrows the model's freedom until only valid, compliant outcomes remain.

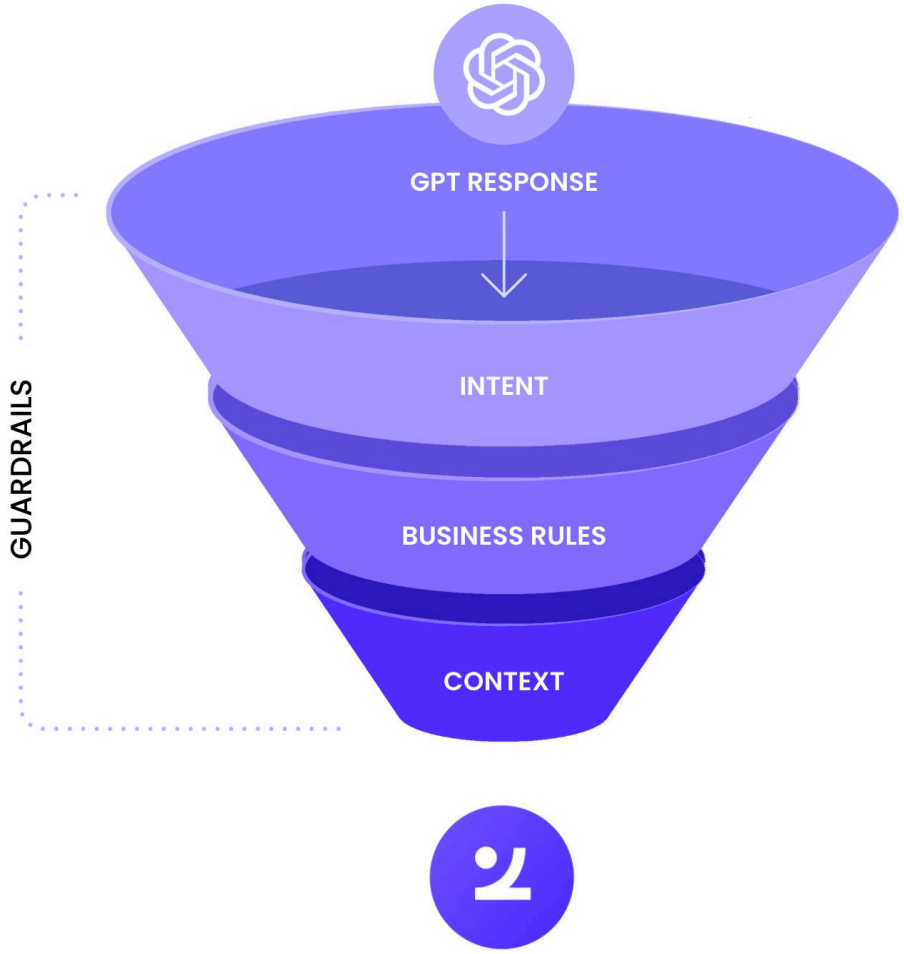

Intent: AI starts with a clear goal, not a guess

Intent detection is the first gate. AI analyzes the user's input and maps it to a predefined intent. It tells AI what the user wants and what actions are allowed.

Think of intent detection as setting a destination in Google Maps. Before you do, the map can't chart a course. Once you say "Go to the airport," it knows where to head.

For example, when a user texts "Cancel Subscription," Zingtree interprets the intent as cancel_and_refund. That single trigger activates the right logic flow.

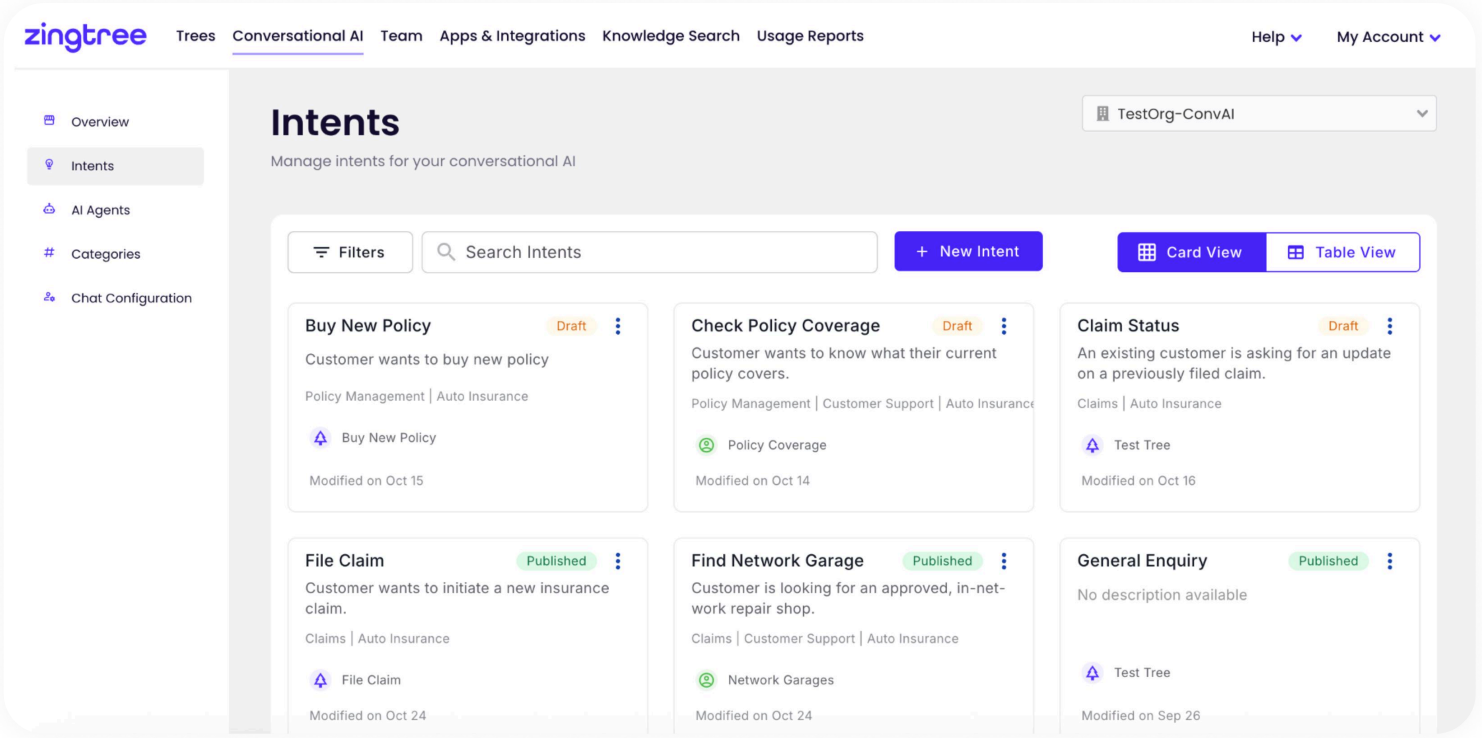

Business rules: Define how AI behaves before and after it speaks

When ChatGPT writes a response, it predicts one word at a time based on probabilities and context. That's why the same question can yield different answers.

To make outputs consistent and accurate, agentic orchestration platforms like Zingtree apply pre-generation and post-generation controls.

Pre-generation controls shape how the model behaves. Before the AI writes a single word, we calibrate its "creativity" and decision space.

Once an AI answer is generated, Zingtree applies a second layer of control to measure how reliable the response is. When AI isn't sure what the user meant, it asks clarifying questions.

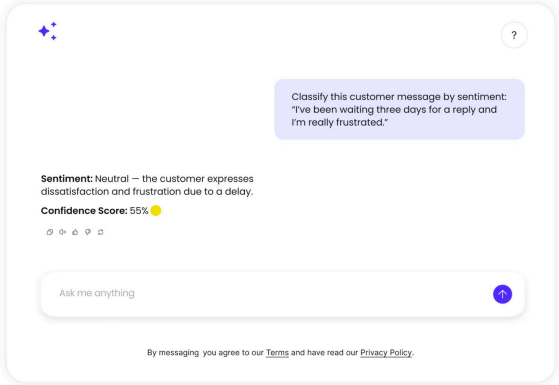

Action triggered: Follow non-AI path

Control Phase	What It Does	Mechanisms
Pre-generation	Shapes how the AI writes its response, controlling creativity, tone, and scope before generation.	**Temperature:** Low = consistent and repeatable answers. High = increased creativity **Top-p (P-score):** Narrows word choices to the most probable ones.
Post-generation	Evaluates and corrects responses after generation, ensuring accuracy and compliance in real time.	**Confidence Scoring:** Reruns low-confidence answers. **Re-Engagement Logic:** Clarifies uncertainty before proceeding.

Context: Ground every answer in real-time truth

Even with the right intent and business rules, **AI can still hallucinate without access to live data.**

For example, a customer asks, “Am I eligible for a refund?” AI knows the intent (“refund request”) and the rule (“refunds apply within 30 days”). But without access to billing data, it can’t verify the purchase date or payment method.

Context fixes that. By connecting directly to your live systems (CRM, ERP, billing, policy), the AI can see what’s real before it responds. It knows who the customer is, what plan they’re on, and what actions are allowed.

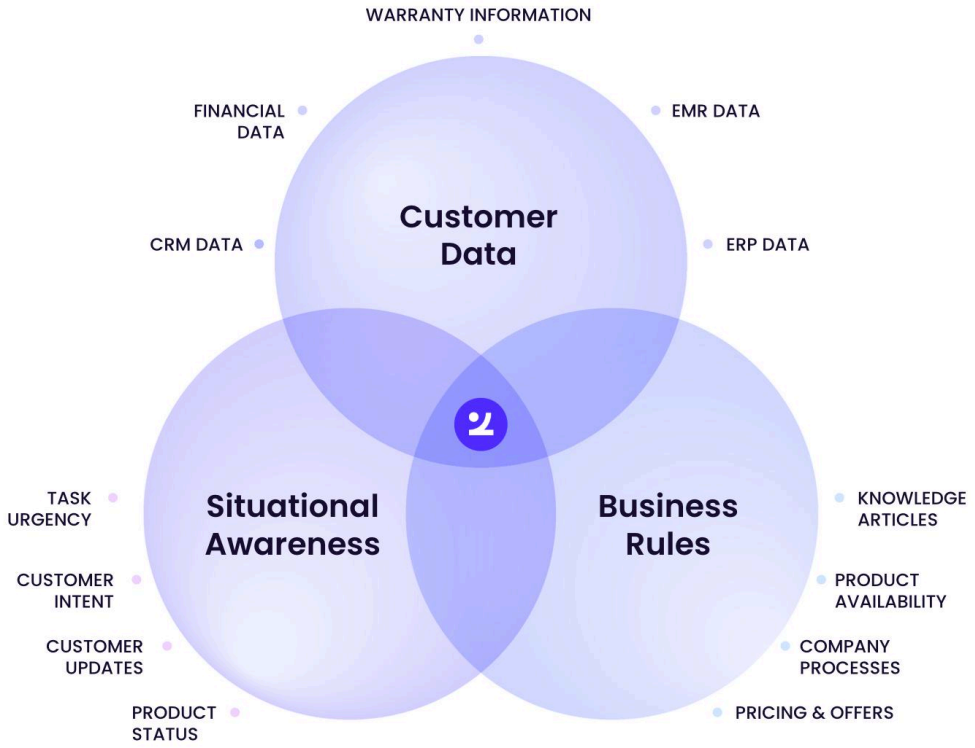

From there, AI can confidently take the safest:

- **Knowledge Assist** – pulls verified information from your Knowledge Base to answer simple questions.
- **Workflows** – launches structured, rule-based flows for complex or regulated scenarios.
- **Human agent escalation** - Routes to an agent when the task exceeds the model’s authority or confidence threshold.

By grounding every answer in live data, you remove the model’s need to guess, making each interaction factual, compliant, and explainable.

The architecture of control

Deterministic AI isn't just about adding rules. It's about engineering predictability into every layer of automation.

Take Zingtree's architecture, for example. It turns open-ended language into controlled action.

From the moment a user asks a question, each system component works in sequence to interpret, constrain, verify, and resolve, balancing intelligence with oversight.

It's an adaptive decision engine, continuously analyzing risk, confidence, and data integrity before an answer ever reaches the customer.

GUARDRAILS FOR GENAI

DETERMINISTIC + PROBABILISTIC IN ACTION

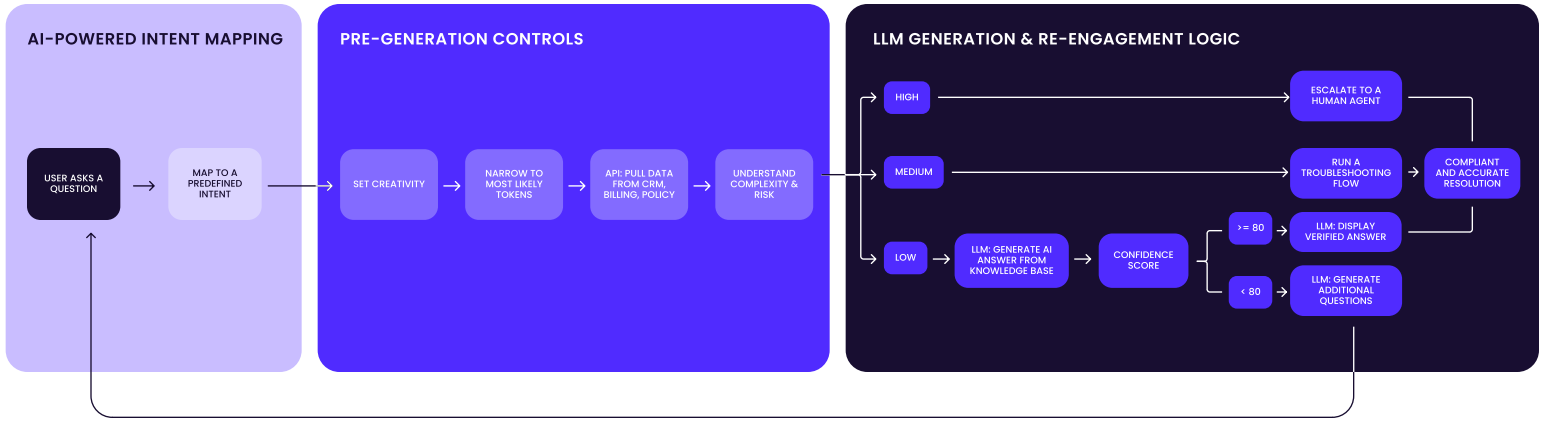

Real-world CX use case:
Optimizing FCR for 1st Central

1st Central, a leading UK motor and home insurer, manages over 3 million customer interactions annually.

Faced with rising customer expectations, tightening regulatory standards, and increasingly complex workflows, the organization identified first contact resolution (FCR) as its north star for optimizing contact center performance across the board.

With Zingtree's deterministic agentic AI workflow automation seamlessly guiding human agents through complex resolutions, 1st Central saw:

+10%
INCREASE IN FCR

30%
HIGHER QA SCORES ACROSS KEY JOURNEYS

3x
REDUCED AGENT ERRORS

FASTER ONBOARDING AND HIGHER AGENT CONFIDENCE

"We built our business case on First Contact Resolution. Because if you get the first call right, the customer doesn't need to call back. That's a better experience for them, and operationally, it means fewer repeat contacts and less burden on our team."

Andrew Smith

Director of Customer Support, 1st Central

Zingtree: Trusted AI Automation for Complex CX

We give support teams control over how AI automates complex resolutions, connecting the dots between data, logic, and compliance.

- 400+ customers
- Solving the most complicated support cases in the most complex industries: Finance & Insurance; Healthcare; B2C SaaS & Marketplaces; Industrial & Consumer Products
- 70,000 authors use Zingtree to deliver over 260,000 workflows powering 160M+ agent and self-service interactions - from guided troubleshooting to compliance-critical decisions
- 10+ years in business

12x
INCREASE IN DEFLECTIONS

30%
FCR IMPROVEMENT

30%
CSAT BOOST

40%
INCREASED REVENUE

SONY · experian · INTUIT · Shark NINJA · United Healthcare

Sage · sleep number · GROUPON · Cigna

LINKEDIN · YOUTUBE · CONTACT

AI Glossary

zingtree

LLM (Large Language Model)

An AI system trained on text datasets to generate human-like responses by predicting the most likely next word in a sequence.

Hallucination

When AI produces information that sounds correct but is factually false. It happens because the model is predicting language, not retrieving verified data.

Probabilistic AI

Traditional AI models (like LLMs) that generate outputs based on statistical likelihoods. Useful for creativity and language generation, but prone to inconsistency without control.

Deterministic AI

Probabilistic AI operating within deterministic guardrails to deliver compliant and repeatable outputs. It ensures every answer is explainable and grounded in business logic.

AI Agent

A goal-oriented AI system that can take actions, make recommendations, or complete tasks on behalf of a user — for example, summarizing tickets or pulling CRM data.

Agentic Orchestration

A structured framework where multiple AI agents work together, coordinated by business logic and guardrails, to complete multi-step tasks across systems safely

Intent Detection

The process of identifying what a user wants to do based on their input. It ensures AI understands the goal before generating or triggering a response.

Temperature

A setting that controls how "creative" or "random" an AI's response can be. Lower values make outputs more predictable; higher ones make them more varied.

Top-p (P-score)

A control that limits AI's vocabulary to the most probable next words. Lower P-scores focus on accuracy; higher ones increase diversity of expression.

Confidence Score

A metric that measures how certain the AI is about its answer. Low confidence can trigger re-checks, clarifications, or human escalation.

Re-engagement Logic

Rules that guide AI to clarify, confirm, or retry when it's uncertain about the user's intent or the accuracy of its answer.

Knowledge Assist

A tool that retrieves verified information from a company's knowledge base, ensuring answers are grounded in trusted data.

Workflow Logic

A structured sequence of business steps that guide the AI through compliant, repeatable actions.

Context Injection

The process of feeding real-time business data (from CRM, ERP, billing, or policy systems) into the AI so it can personalize and validate its responses.

Deterministic Guardrails

Zingtree's 3-layer control system — Intent, Business Rules, and Context — that ensures every AI decision is accurate, compliant, and explainable.