



## Rezolve Ai Solves the 26% AI Distortion Crisis Stalling Global Retail; Peer-Reviewed Research Validates Near-Perfect Accuracy

May 19, 2026 12:00 PM EDT

***New Research Showing 26.5% AI Product-Review Distortion and 60% Hallucination Rates Proves Commerce Needs Verified Intelligence***

***Rezolve Ai's Dual Layer Architecture - brainpowa and TraceWare - Secures the Future of Agentic Commerce with Scientifically Validated, Hallucination-Free Reliability***

NEW YORK, May 19, 2026 (GLOBE NEWSWIRE) -- Rezolve Ai (NASDAQ: RZLV), a global leader in AI-powered commerce technology, today announced that peer-reviewed research accepted for presentation at the 34th ACM International Conference on User Modeling, Adaptation and Personalization, ACM UMAP 2026, has validated its proprietary TraceWare technology as a breakthrough verification layer for reliable agentic commerce.

The announcement comes as the retail industry faces growing evidence that generic AI is not yet reliable enough for live commerce. Independent UC San Diego research has found that LLM-generated product review summaries changed the sentiment of the original reviews in 26.5% of cases, while the same study found hallucination rates above 60% when LLMs answered post-knowledge-cutoff fact-checking questions. The research also found that consumers were 32% more likely to purchase the same product after reading an LLM-generated summary rather than the original human-written review.

For commerce, the implication is clear: generic AI does not merely risk making mistakes at checkout; it can distort the discovery journey that leads to the purchase.

At the same time, major retail players continue to confront the limits of generic AI in real-world shopping. **Amazon** has brought together Rufus and Alexa+ to create Alexa for Shopping, replacing the Rufus-branded assistant with a more integrated shopping experience across Amazon's app, website and devices. **Walmart's** reported experience with OpenAI-powered Instant Checkout has also highlighted the difficulty of turning generic AI conversations into reliable transactions, with WIRED reporting that conversion rates for products sold directly within ChatGPT were three times lower than for products requiring shoppers to click through to Walmart's own website.

Rezolve Ai believes these developments confirm a fundamental point: generic AI can talk about shopping, but commerce requires verified intelligence.

### **The Commerce AI Reliability Gap**

Commerce AI must do more than generate plausible answers. It must understand products, inventory, customer intent, basket state, checkout status, payment flows, order completion and business rules. When an AI assistant gets any of those wrong, the result is not simply an inaccurate response; it can mean a broken basket, a lost sale, a failed checkout or a damaged customer relationship.

Rezolve Ai's Brain Suite addresses this problem through a dual-layer architecture:

**brainpowa**, Rezolve Ai's proprietary commerce-specific AI architecture, grounds product discovery, search, conversational guidance and personalisation in the retailer's own data, catalogue, inventory and commercial rules.

**TraceWare**, Rezolve Ai's proprietary verification layer, reconstructs the user's journey from real execution logs and verifies every AI claim against factual invariants. Rather than allowing a language model to infer what happened in a transaction, TraceWare traces the facts from each step of a transaction.

Together, brainpowa and TraceWare are designed to give retailers the accuracy, control and reliability required to deploy AI inside live commercial environments.

### **Peer-Reviewed Research Validates TraceWare**

The accepted paper, "**User-State Verification in Conversational Commerce: Detecting Journey Hallucinations via Trace Invariants**," introduces the concept of "journey hallucinations" as a distinct class of AI failure in commerce.

A journey hallucination occurs when an AI commerce agent claims that a transactional event has happened, such as an item being added to a basket, checkout being started, or an order being completed, when the actual transaction record shows otherwise.

Unlike general factual hallucinations, journey hallucinations directly affect commercial outcomes. An AI assistant that incorrectly believes a cart is empty, falsely states that checkout has begun or assumes an order has been completed will trigger the wrong action at the wrong time. That can damage trust, interrupt the buying journey and reduce conversion.

In rigorous testing across 90 real eCommerce sessions, 2,000 events and two merchants, the research evaluated leading foundation models including **OpenAI GPT, DeepSeek-V3, Kimi-K2 and Qwen3-32B** under three operating conditions.

Unconstrained foundation models produced unsupported user-state assertions at rates of up to 8.5%, equivalent to roughly one error in every twelve customer interactions. Conservative prompting reduced those errors only by refusing to answer, with coverage falling as low as 15%, creating agents that were technically safer but commercially ineffective.

By contrast, Rezolve Ai's TraceWare technology achieved 99.5% to 100% user-state accuracy, 84% to 99% coverage and near-zero false positives across every model tested, without retraining a single foundation model.

### **Generic AI Distorts Discovery. TraceWare Verifies Execution. brainpowa Grounds the Journey.**

Independent research shows that the risks from generic AI begin before checkout. If generic AI can change the sentiment of product information in more than a quarter of cases and hallucinate above 60% when information falls outside its training data, then retailers cannot rely on generic model

inference to guide customers through product discovery, comparison and purchase decisions.

Rezolve Ai believes that reliable agentic commerce requires three things: grounded discovery, verified execution and retailer-controlled transaction infrastructure.

brainpowa is designed to ground the discovery journey in the retailer's live product data, inventory and rules. TraceWare is designed to verify the execution journey against what actually happened. Together, they enable AI agents to act on the basis of verified facts rather than plausible guesses.

"Retailers do not need another AI experiment; they need an AI commerce platform that works now," said **Daniel M. Wagner, Founder, Chairman and CEO of Rezolve Ai**. "The industry is learning that generic AI can influence shoppers, distort discovery and break down at the point of transaction. You cannot build the future of commerce on systems that guess what happened in the customer journey."

"Recent developments around Rufus and OpenAI checkout illustrate the same issue from different angles. Generic AI can talk about shopping but commerce requires product grounding, catalogue intelligence, basket awareness, business rules, payment integration and verified reliability at the point of purchase."

"That is why Rezolve Ai built brainpowa and TraceWare. brainpowa grounds the shopping journey in commerce-specific intelligence and TraceWare verifies the transaction journey against what is real. Agentic commerce will not scale on hallucination-prone systems. It will scale on verified intelligence."

**Dr. Salman Ahmad, Chief Scientist at Rezolve Ai Labs, added:** "Reliability in commerce AI depends on grounding architecture, not model size. Across every foundation model we tested, TraceWare delivered near-perfect accuracy without retraining. That is the critical point. The answer is not simply a bigger model; it is a commerce-specific architecture that can verify what the AI says against what actually happened."

"The research shows that TraceWare can materially improve the reliability of leading foundation models in transactional commerce environments. Combined with brainpowa, Rezolve Ai is giving retailers the verified intelligence layer required for trusted agentic commerce."

### **A New Standard for Agentic Commerce**

Rezolve Ai believes the market is now moving beyond the first phase of conversational AI, where success was measured by whether an assistant could answer questions. The next phase is agentic commerce, where AI must search, recommend, compare, personalise, transact and complete purchases reliably.

In that environment, accuracy is not a feature. It is the foundation.

The peer-reviewed paper, authored by researchers at Rezolve Ai Labs and the University of Sarajevo, will be presented at ACM UMAP 2026 in Gothenburg, Sweden, from June 8 to 11, 2026.

A copy of the paper is available here: <https://rezolve.com/reports/user-state-verification-in-conversational-commerce-detectingjourney-hallucinations-via-trace-invariants/>

### **About Rezolve Ai**

Rezolve Ai (NASDAQ: RZLV) is an industry leader in AI-powered solutions, specializing in enhancing customer engagement, operational efficiency, and revenue growth. The Brain Suite is the world's first enterprise AI platform built for Agentic Commerce, delivering advanced tools that harness artificial intelligence to power search, transact, fulfill, and personalize at global scale. For more information, visit [www.rezolve.com](http://www.rezolve.com).

### **Media Contact**

Rezolve Ai  
Urmeek Khan - Global Head of Communications  
[urmeekhan@rezolve.com](mailto:urmeekhan@rezolve.com)  
+44 7576 094 040  
[investors@rezolve.com](mailto:investors@rezolve.com)

### **Forward-Looking Statements**

This press release includes "forward-looking statements" within the meaning of the "safe harbor" provisions of the Private Securities Litigation Reform Act of 1995. The actual results of Rezolve Ai plc ("Rezolve") may differ from its expectations, estimates and projections and consequently, you should not rely on these forward-looking statements as predictions of future events. Words such as "expect", "estimate", "project", "budget", "forecast", "anticipate", "intend", "plan", "may", "will", "could", "should", "believes", "predicts", "potential", "continue", "design" and similar expressions are intended to identify such forward-looking statements. You should carefully consider the risks and uncertainties described in the "Risk Factors" section of Rezolve's Annual Report on Form 20-F and its subsequent filings made with the SEC. These filings identify and address other important risks and uncertainties that could cause actual events and results to differ materially from those contained in the forward-looking statements. Most of these factors are outside Rezolve's control and are difficult to predict. Factors that may cause such differences include, but are not limited to: (1) competition, the ability of Rezolve to grow and manage growth profitably, and retain its management and key employees; (2) changes in applicable laws or regulations; and (3) weakness in the economy, market trends, uncertainty and other conditions in the markets in which Rezolve operates, and other factors beyond its control, such as inflation or rising interest rates. Rezolve cautions that the foregoing list of factors is not exclusive and not to place undue reliance upon any forward-looking statements, including projections, which speak only as of the date made. Except as required by applicable law, Rezolve does not plan to publicly update or revise any forward-looking statements contained herein, whether as a result of any new information, future events, changed circumstances, or otherwise.