COMPETENCE CENTER
for on-the-go Consumption

Generative AI Use Cases in Retailing

How retailers leverage GenAI to enhance customer experience, increase store and operational efficiency

Professor Sabine Benoit,
Director, Competence Center for On-the-go Consumption
Professor of Marketing (Practice), Singapore Management University

Director, Competence Center for On-the-go Consumption
Professor of Marketing (Practice), Singapore Management University

Foreword

This white paper showcases selected, upcoming, lesser-known generative artificial intelligence (GenAI) use cases across the retail value chain, from consumers' homes to retail stores and corporate decision-making. It **focuses** on GenAI, meaning Large Language Models (LLMs) or nowadays more precisely Large Multimodal Models (LMMs), that generate text, pictures, videos, or other content. GenAI differs from discriminatory AI and traditional algorithms in two main ways. GenAI models produce content of some sort, whereas discriminatory AI mainly aims to classify cases, e.g., age classification of customers purchasing restricted products or basket capture in autonomous stores. Beyond this, GenAI models learn through their usage.

The study focuses on emerging GenAI applications and is targeted at practitioners. This complements my other, more academic contributions on AI and retail technology in world-leading journals in retailing and business, for instance, here, here, and here. This white paper is based on extensive desk research of publicly available information and my evaluation and selection of upcoming use cases. I have created this white paper because, in the past, retailers have found it helpful when I share my views as an academic specialising in retail technology and observing the sector.

Retailing, but even more so GenAI, is developing rapidly, and I'm fully aware that any study on GenAI use cases will be outdated on the day it is published. However, that does not mean we should not try to showcase upcoming use cases that are worth investigating as a retailer. My white paper aims to inspire and inform.

I'm sharing my observations on upcoming GenAI cases in retailing, hoping you, as a reader, find them useful in developing your business. If you think I have missed a very innovative and interesting use case, please share this with me. If it's suitable, I'd be happy to add this to the selection.

Sabine Benoit, Professor of Marketing (Practice),
Singapore Management University (SMU),
Singapore, October 2025

TABLE OF CONTENTS

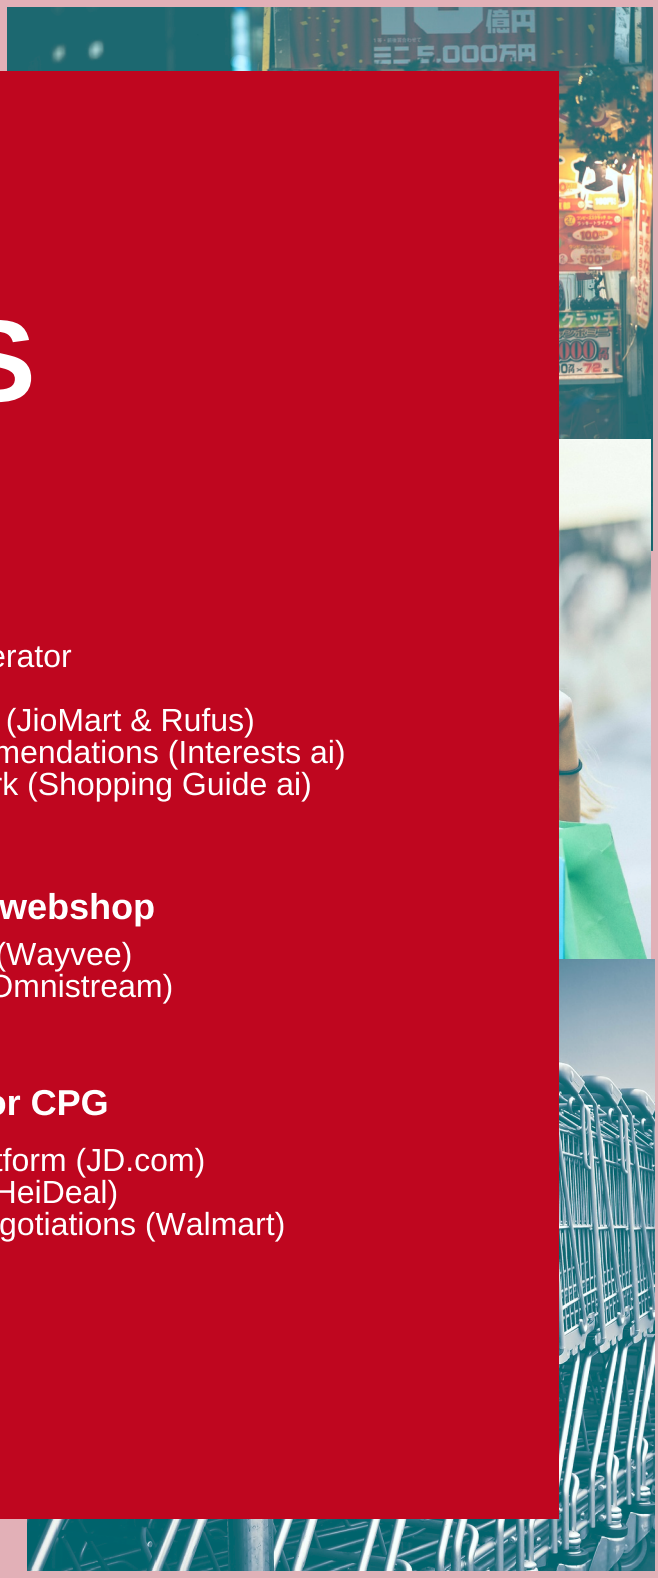

GenAI applications across the retail value chain

This white paper showcases selected, innovative use cases of how generative artificial intelligence (GenAI) use cases are transforming the retail value chain across three stakeholders and domains: customer-facing applications, in-store operations, and retailer/manufacturer business processes.

1. **Applications for Customers**

   AI tools embedded in consumers' lives support better purchasing decisions and personalized shopping experiences. Photo-based calorie counters and recipe generators for leftovers help consumers make healthier, more sustainable choices. Conversational commerce platforms like JioMart on WhatsApp and Amazon's Rufus enable complete shopping journeys with AI-powered guidance. JD.com's virtual avatars, based on their ChatRhino model, conduct live shopping sessions, while Amazon's AI Shopping Guides simplify complex purchase decisions by learning from user interests to tailor recommendations in real time.

2. **Applications for the Store**

   Computer vision and AI optimization enhance both customer experience and operational efficiency. Wayvee Analytics' emotion-tracking monitors shopper responses in real time via privacy-compliant radar sensing, enabling retailers to optimize layouts and merchandising. Omnistream's platform, deployed by 7-Eleven Philippines and others, generates optimized store-level planograms in minutes rather than weeks or months, improving space utilization and category sales while reducing inventory costs.

3. **Applications for retailers and manufacturers**

   At the corporate level, AI automates complex business processes. JD.com's content management platform based on the ChatRhino model generates marketing content such as images, livestream scripts and videos for livestreaming in hours rather than days, while substantially reducing costs. Heineken UK's HeiDeal models supply contract scenarios, optimizing margins and profitability while automating approval workflows. Walmart's partnership with Pactum enables fully autonomous contract negotiation with tail-end suppliers, freeing procurement teams for strategic work while achieving savings and improved payment terms.

GenAI applications across the retail value chain

GenAI applications

for customers

1. Leftover recipe generator
2. Photo calorie estimator
3. Conversational commerce (JioMart, Rufus)
4. AI agents (Interests ai, Amazon)
5. Decision-making framework (Shopping Guides ai, Amazon)

for the store

6. In-store emotion tracking (Wavee)
7. Planogram optimization (Omnistream)

for the retailer or manufacturer

8. Content management platform (JD)
9. Sales contract modelling (Heineken)
10. Autonomous contract negotiations (Walmart)

Use case selection criteria

The selection for showcasing the use cases for this white paper was based on the following criteria:

1. **emerging, less well-known use cases from around the world to inspire retailers**
2. **use cases, where GenAI models generate content, rather than predictive AI based optimization**
3. **use cases that are particularly relevant for the grocery sector**
4. **use cases that have the potential to change consumer behavior and be adopted by the mass market**

Use Case 1

Fridge Leftover recipe generator

Fridge leftover recipe generator: the idea

What is the idea of a **fridge content recipe generator**?

A picture of a customer's his or her fridge, the content is predicted via image recognition and is then transformed into a recipe that includes the ingredients that are in the fridge.

PHOTO

RECIPE

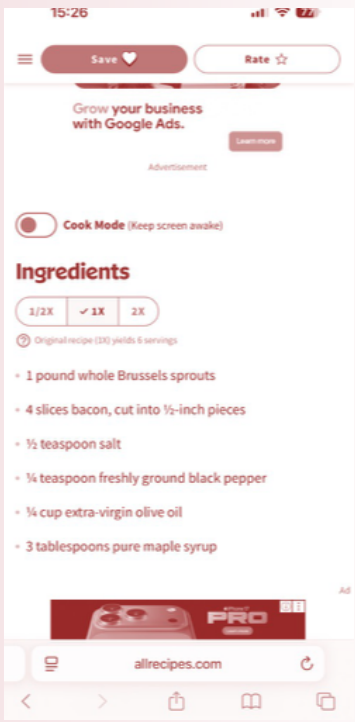

Technical architecture: The photo is analysed using image recognition to predict the contents of the fridge. These contents are then analysed by a model generating a suitable recipe that can be cooked with the available ingredients.

Technology Test Leftover Recipe Generator: set up

During the desk search for this study, we were first made aware of this use case through a feature within the retailer app from Albert Heijn. Although the idea has great potential, upon testing the app, the implementation showed substantial room for improvement. Therefore, for this white paper, four different GenAI applications were tested that convert a picture of fridge ingredients into a recipe based on the available ingredients.

We benchmarked the application offered by the retailer (Albert Heijn) against another by a standalone, specialised app provider and two general GenAI models.

Retail App: Dutch supermarket chain Albert Heijn has integrated the Scan & Cook service into its grocery shopping app, enabling customers to photograph their fridge contents and receive recipe suggestions.

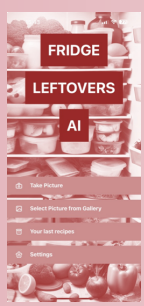

Specialised App Provider: Fridge Leftover AI positions itself as a solution to combat food waste, providing a stand-alone app for transforming fridge photos into recipes.

General GenAI Tools ChatGPT (Open AI): A general GenAI model with image recognition and content generation features was also tested for its capability to identify ingredients from an image and generate corresponding recipes.

General AI Tools – Claude (Anthropic): Similar to ChatGPT this is a general GenAI model, which was also evaluated for its ability to process fridge photos and create recipes.

Technology Test: Result Details

First winner: The general GenAI model Claude took first place in the test. It allows customers to generate a recipe with the basic, free version and demonstrates very accurate identification capabilities. Notably, it was honest about its limitations when it couldn't identify all fridge items. The technology also displayed a sense of humor and showed creativity when suggesting potential other ingredients that might be present in the kitchen, but not in the fridge (such as rice).

Second winner: The specialised app provider is placed second in the technology test. It provides the most comprehensive features, detailed instructions, and also includes nutritional value information. It also offers users the ability to choose from different ingredients after the initial identification. This means that if one ingredient in the fridge is not available for cooking, for instance, because it is needed for the kids or tomorrow's breakfast, it can be excluded. The main drawback is that it comes with a cost of $9 per annum, but this seems reasonable for an app; however, it needs to be seen in light of the fact that the basic general GenAI models are free of charge.

The **shared third place** is held by ChatGPT and the AH retailer app.

ChatGPT also offers recipe generation with the basic free-of-charge version. However, its main weakness was that it appeared too confident when identifying ingredients. It mistook plastic cheese packaging for wraps and suggested a recipe with wraps. This reduces its practical value.

The **Albert Heijn app** performed well in terms of accessibility and integration. It is free to use and seamlessly embedded within the retailer's grocery shopping app. However, the app reverted to Dutch midway through the process, creating usability issues for non-Dutch speakers, who are likely less relevant to the retailer. Furthermore, it generated a recipe presumably from their recipe catalogue that was likely based on the best match of individual ingredients, which meant that a fish recipe was proposed, even though the main ingredient (the fish) was not currently available in the consumer's fridge. This reduces its practical value.

Technology Test: Result Overview

Retailer App	Evaluation
Free of charge	**
Integrated into the retailer app	**
App translation was incomplete, the app switched to Dutch in the process	--
The app generated recipe with main ingredients unavailable	--

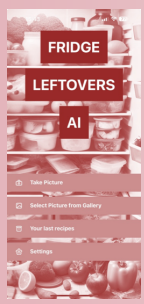

Specialised App Provider	Evaluation
Comprehensive instructions	**
Nutritional value	**
Choice of ingredients identification	~
Cost: 9 Dollar	~

General GenAI model: ChatGPT	Evaluation
Recipe generation available with the basic, free of charge version	**
Too confident in identifying ingredients that were not in the fridge (mistook cheese plastic packaging for wraps)	~
Result (recipe) ok for general model	--

General GenAI model: Claude	Evaluation
Recipe generation available with the basic, free of charge version	**
Very accurate on ingredient identification, honest about not being able to identify some of the ingredients with a sense of humor	**
Outputted various results (recipes) for different variations of the non-identifyable ingredient	**

Positive / very positive: */** Neutral: ~ Negative / very negative: -/--

Fridge leftover recipe generator: summary

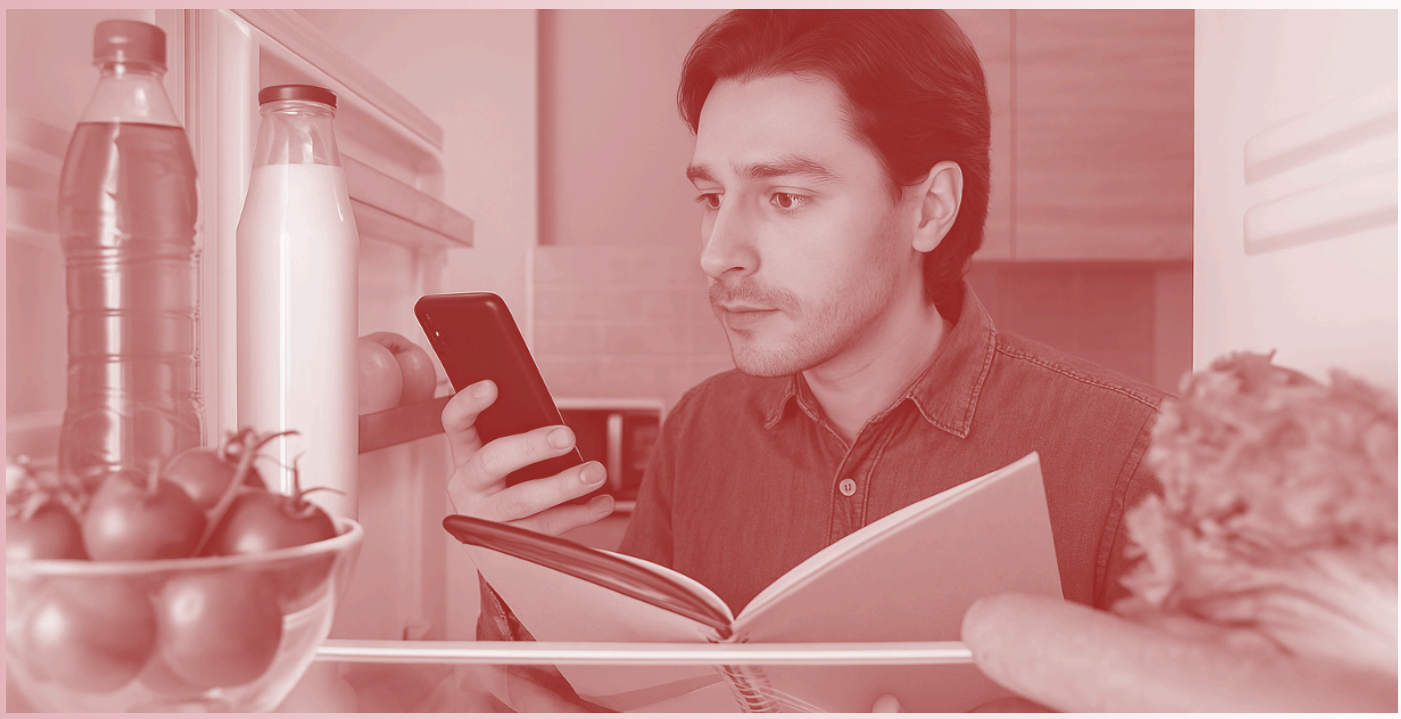

1. **Sustainability impact: Avoids food waste by allowing customers to generate tailored recipes**
2. **Low barriers: General-purpose GenAI tools such as Claude achieve usable outcome at no-costs**
3. **Allows retailers to provide extra customer service, but current models need to be improved**

Use Case 2

Photo based calorie estimator

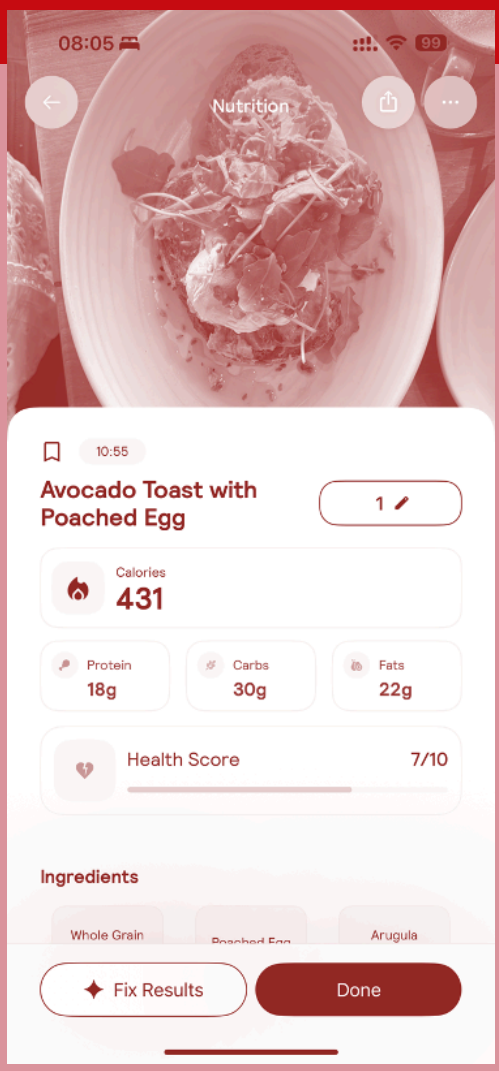

Photo based calorie estimator: the idea

What is the **idea** of a **photo-based calorie estimator**?

Many consumers track their food intake, but currently, this requires substantial effort. The photo-based calorie estimator estimates the nutritional values, including calories, of a meal from a picture of a customer's food on his or her plate. This allows customers to track their food and share this data with third parties (doctors, health coaches).

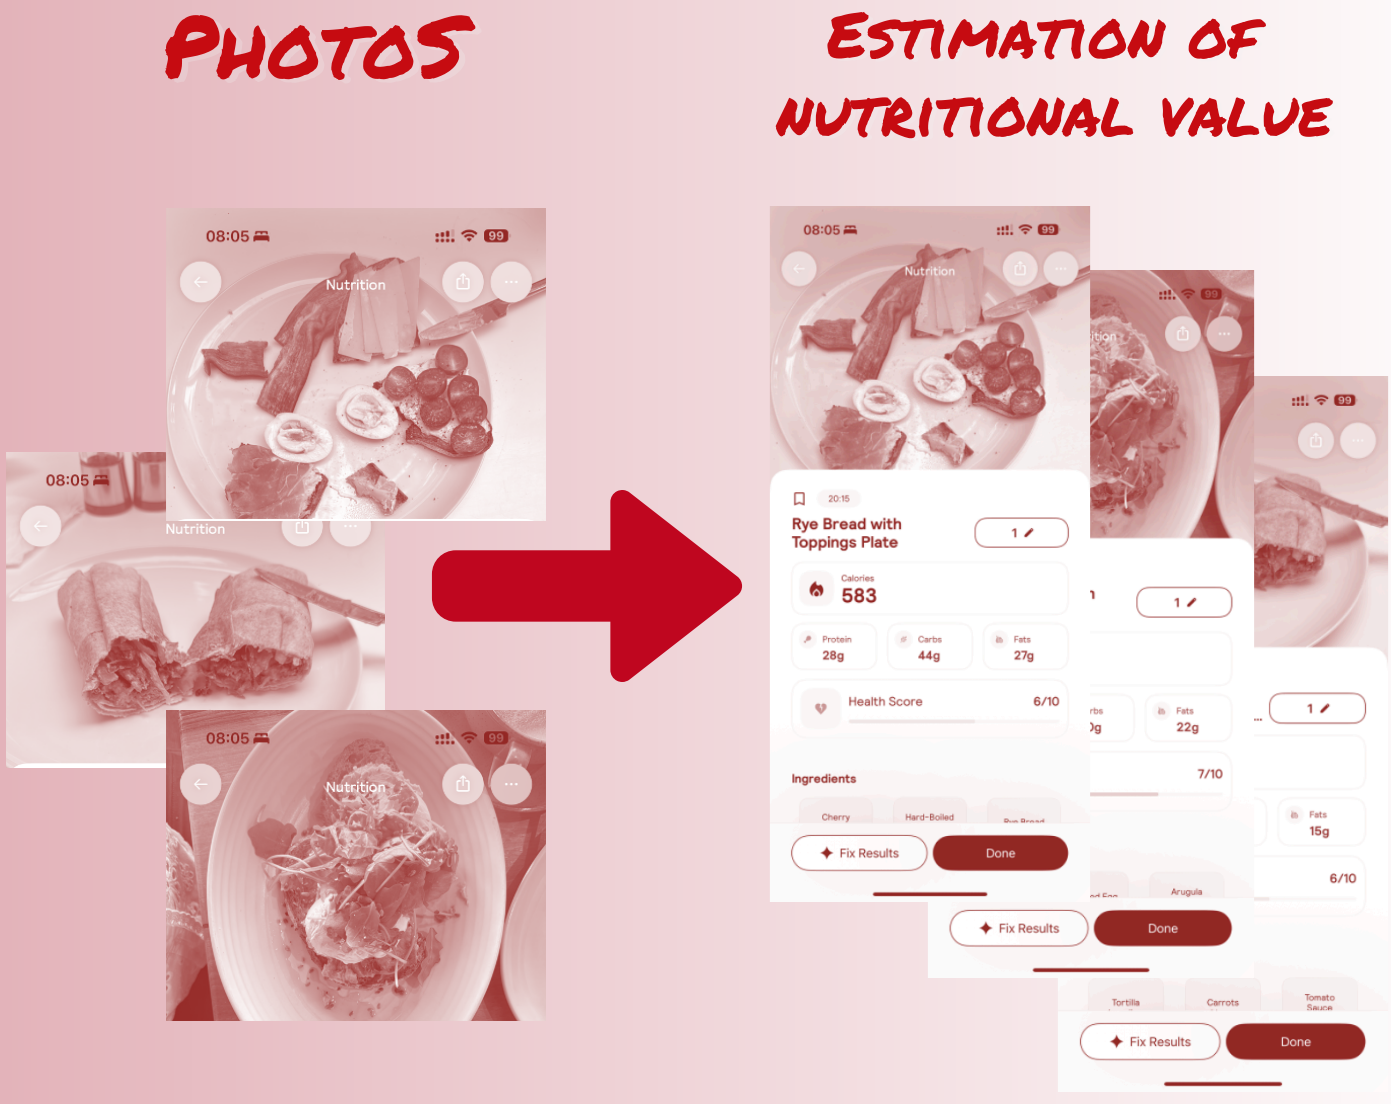

Photo based calorie estimator

The emergence of photo-based calorie estimator applications such as Cal AI offered by Apple represents a significant **advancement in nutritional tracking**. These applications employ computer vision and machine learning algorithms to analyze photos of food, automatically identifying ingredients, estimating their weight, and calculating associated nutritional values. The output typically includes caloric content, macronutrient composition (proteins, carbohydrates, and fats), and micronutrient values such as vitamins. More sophisticated implementations extend this functionality to generate composite health scores, providing users with an at-a-glance assessment of their dietary choices.

The deployment of these applications presents a unique **challenge related to food recognition accuracy across different culinary contexts** and cuisines. Unlike standardized consumer products with uniform presentation, prepared foods exhibit substantial **regional variation** in ingredients, preparation methods, and aesthetics, including how they are presented on the plate and the size of the plates. Consequently, the market has evolved to include multiple specialized providers, each focusing on specific geographic regions or culinary traditions. This regionalization is necessary because the underlying machine learning models require training on representative datasets that accurately reflect local food versions and cultural variations in cuisine. A model trained predominantly on Western dishes, for instance, would perform poorly when presented with Southeast Asian or Middle Eastern dishes.

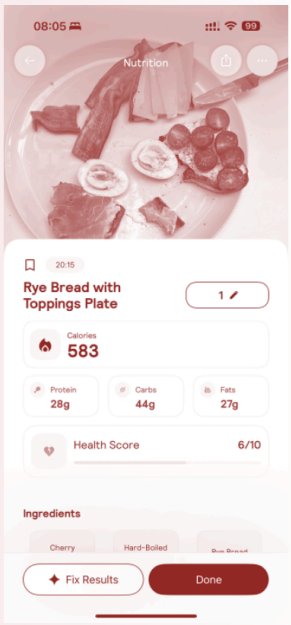

The apps regularly allow users to adjust the predicted food items and the number of grams. For instance, in one photo, red strips of pepper were misidentified as bacon, which could be corrected in Cal AI. This feature is likely used for further **model training**. It can therefore be expected that these apps improve over time.

To evaluate the **comparative performance** of different AI technologies in this application domain, we conducted tests using one leading app: Cal AI with general LLMs such as ChatGPT and Claude. The results showed similar accuracy for the tested (mainly Western) dishes and are not presented here to reduce redundancy.

Differnt regionalised photo based calorie estimators

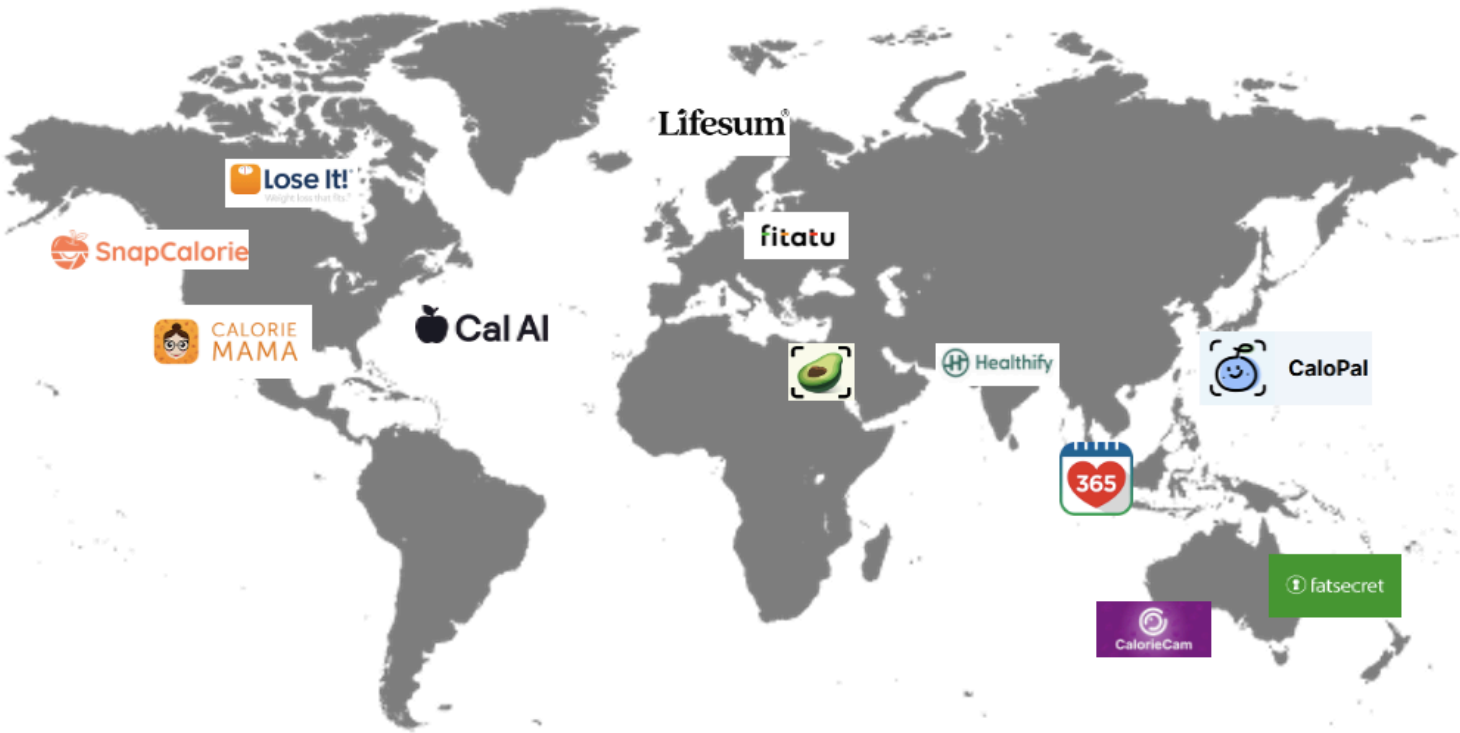

Photo based calorie estimator: summary

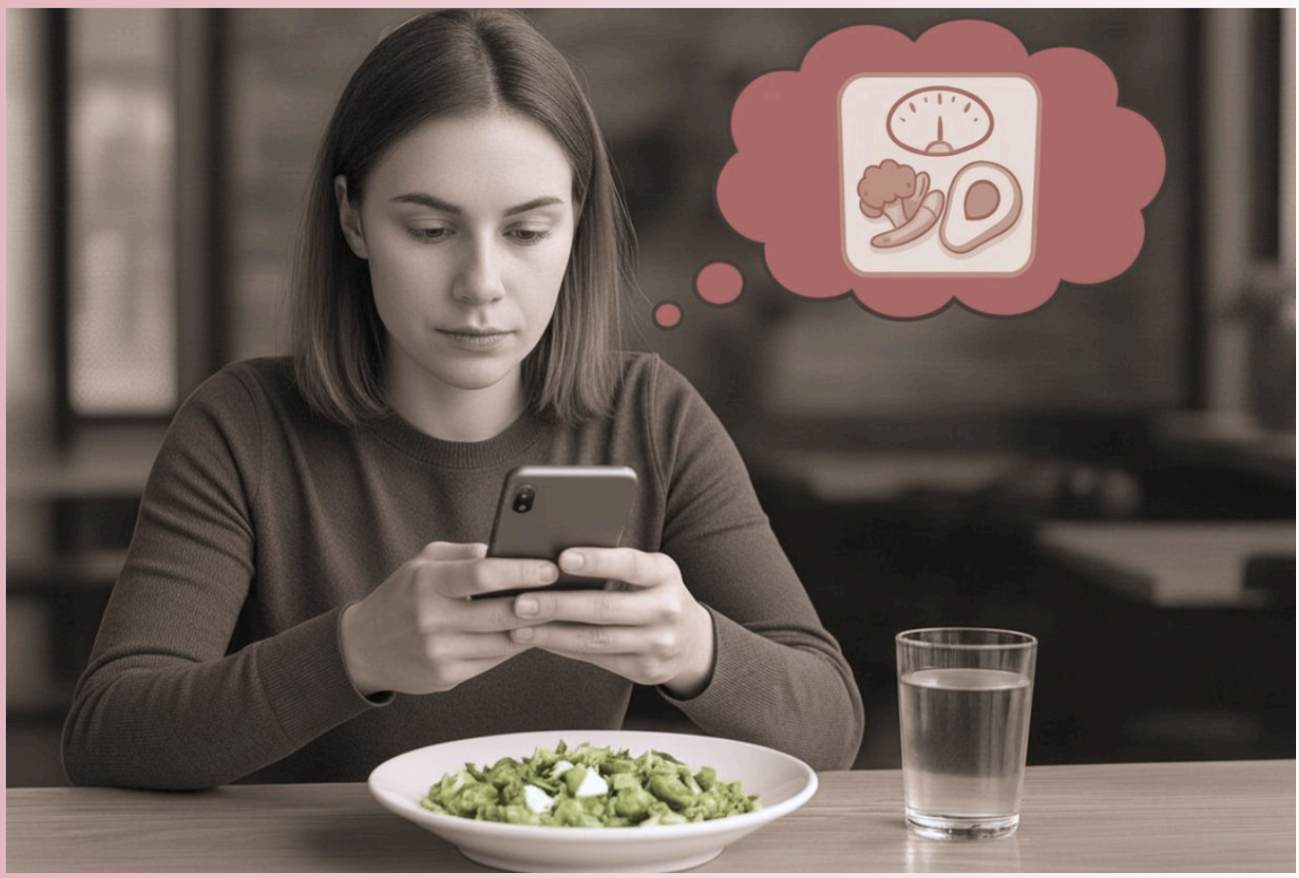

1. **Photo based estimation eliminates manual logging and likely widens the target group of food trackers**
2. **Accuracy challenges persist: Regional food variation and presentation differences require specialized models**
3. **User correction allows continuous model improvements with free general-purpose models being competitive**

Use Case 3

Conversational Commerce

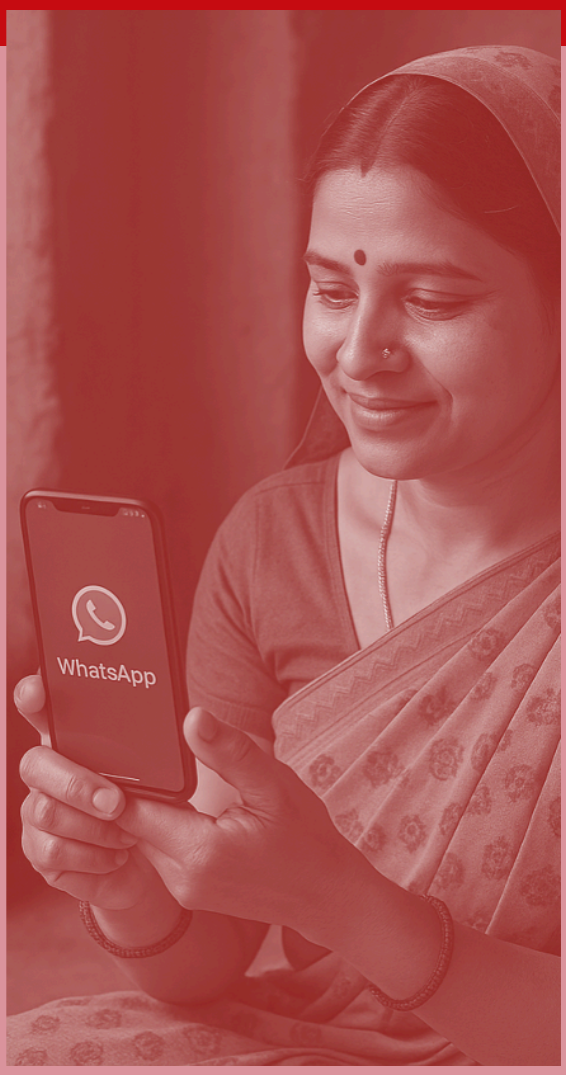

Conversational Commerce: the idea

What is the **idea** of **conversational commerce**?

The use of chat or messaging apps, such as WhatsApp or voice assistants, such as Siri or Alexa to accomplish shopping tasks. The entire shopping process, including the payment process, can be accomplished through the conversational interface.

CHAT

DELIVERY

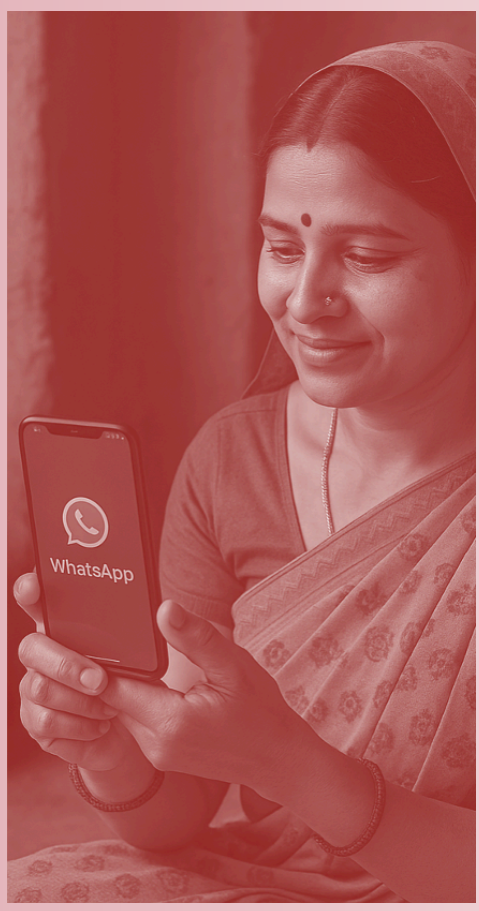

Conversational Commerce: step by step example

Step 1

A customer launches a chat with a messenger service, initiates a chat with a chatbot on an e-commerce website or GenAI application or speaks to a voice assistant, asking for product recommendations. These inquiries may address product specifications, product availability, recommendations, or comparative assessments.

Step 2

The conversational commerce system then generates a contextually appropriate response, drawing on product databases and, potentially, customer purchase history to guide the user toward an informed purchase decision.

Step 3

What distinguishes conversational commerce from conventional chat-based customer service is the seamless integration of all steps, including payment processing, within the conversational interface itself (end-to-end). Users can complete their transactions without exiting the chat window or navigating to separate checkout pages, significantly reducing friction in the purchase path and potentially decreasing cart abandonment.

Step 4

Following transaction completion, the last step mirrors traditional e-commerce fulfillment processes. The key distinction is that order tracking and post-purchase communication can continue within the same conversational interface, maintaining continuity throughout the entire customer experience and providing a unified channel for addressing delivery inquiries or issues.

Note: The conversational interface can be powered by automated chatbots, staffed by live agents, or operated through hybrid models that leverage both approaches. This architectural flexibility allows organizations to optimize the balance between scalability, cost efficiency, and service quality based on business requirements and customer expectations.

Conversational Commerce: JioMart's, Haptik & WhatsApp collaboration

In August 2022, JioMart, a major Indian e-commerce platform, launched a collaboration with Meta (WhatsApp) and Haptik, developing the first **comprehensive end-to-end, chat-based shopping experience** available on WhatsApp. This marked a significant milestone in **conversational commerce**, enabling users to complete the entire transaction, from product discovery through payment, within the messaging application without requiring transitions to external websites or dedicated shopping apps.

The **technical infrastructure** underlying this platform is powered by a conversational AI chatbot developed using proprietary software from Haptik, a company that became part of Jio Platforms Limited following a $100 million majority acquisition in 2019. This vertical integration has enabled JioMart to maintain greater control over the development and customize the technology to address specific market requirements in the Indian retail context.

Haptik's conversational AI capabilities extend beyond JioMart, with the technology provider serving multiple clients, including Mondelez International and Grofers (India). The robustness of their AI models derives from extensive training on proprietary datasets, including conversational data from their Intelligent Virtual Assistants, which have processed over six billion conversations with more than half a billion customers. This substantial data foundation has enabled the development of models that can handle the linguistic diversity and contextual complexity characteristic of the Indian market.

Since its launch, the platform has undergone **significant enhancements to its features**. Notable additions include a dynamic search bar for improved product discovery, pin code verification (representing a first for WhatsApp-based chatbots), enhanced product listing categories, and cart merge features that facilitate consolidated ordering across multiple shopping sessions.

The **strategic vision** for this platform extends beyond its initial scope. JioMart announced plans to expand its conversational commerce offerings to encompass multiple product categories beyond groceries, including fashion, electronics, and additional merchandise categories. This expansion reflects growing confidence in the conversational commerce model and its ability to serve diverse product types within a unified interface.

JioMart's Conversational Commerce: market context

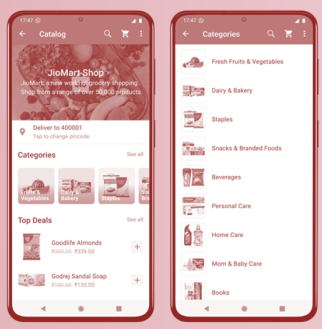

The selection of WhatsApp as the provider for JioMart's conversational commerce implementation reflects several market dynamics and strategic considerations. **WhatsApp** maintains its position as the **world's most widely used messaging application**, providing further reach for retail applications. Despite its acquisition by Facebook (now Meta) in 2014 for $19 billion, WhatsApp has notably abstained from traditional advertising-based revenue models. This positioning made the exploration of alternative revenue streams necessary, with business-facing services and commerce enablement emerging as primary options for monetization.

India represents a **compelling market for WhatsApp-based commerce**. The country hosts one of the largest WhatsApp user bases globally, with approximately 850 million users. This massive customer base provides a ready-made distribution channel that requires no additional application downloads or user education regarding the fundamental interface mechanics.

However, the Indian e-commerce landscape presents a complex opportunity. While **online retail** has achieved significant **penetration** mainly in **urban centers**, online grocery shopping remains relatively **underdeveloped across the broader population**. Current estimates suggest that only one in ten Indians purchases groceries through digital channels. This low adoption rate can be attributed to several factors, including limited digital payment infrastructure in rural areas, concerns about product quality and freshness verification, delivery logistics challenges, and the well-established nature of traditional shopping behaviors.

WhatsApp Shopping represents a strategic approach to bridging this **adoption gap** by targeting a demographic segment that has been underserved by conventional e-commerce platforms: shoppers without prior online shopping experience. By embedding commerce functionality within a familiar messaging interface that users already employ for daily communication, the platform reduces the cognitive barriers and technological intimidation that often prevent first-time adoption of e-commerce services. Users can engage with shopping functionality using conversational patterns they have already mastered, rather than learning to navigate unfamiliar website or application interfaces. This approach effectively democratizes access to online retail by meeting potential customers within their existing digital habitats.

JioMart's conversational commerce: adoption, efficiency & ROI

The **JioMart WhatsApp Shopping platform** reported **performance increases** across multiple dimensions since its launch.

New customers: From a customer acquisition perspective, the platform has proven particularly effective at reaching new market segments. Approximately 37% of WhatsApp shoppers represented first-time customers for JioMart, indicating that the conversational interface effectively lowered barriers to entry for consumers who had not previously interacted with the retailer's traditional e-commerce channels. This acquisition profile supports the hypothesis that embedding commerce within familiar messaging environments can effectively convert non-digital shoppers into online purchasers.

Customer support effectiveness: The retailer's metrics demonstrate improvements. The implementation of conversational AI has yielded a 62% reduction in average response time compared to previous customer service models, enabling faster resolution of customer inquiries and concerns.

The platform has also achieved 88% automation of issue resolution, substantially reducing the burden on human support agents and allowing them to focus on more complex customer needs that require human judgment and empathy.

Sales: While JioMart has not publicly disclosed precise sales figures attributable to the WhatsApp channel, reports from the Economic Times indicate that chatbot-facilitated orders have contributed significant incremental sales volume to the company's overall revenue. This sales contribution represents truly additive value rather than mere channel shifting, as evidenced by the high proportion of first-time customers.

Return on investment: From an ROI perspective, Haptik has reported achieving a 97× ROI from the chatbot implementation. This exceptional return ratio reflects both the relatively modest incremental costs of serving additional customers through automated conversational interfaces and the substantial revenue generation enabled by improved conversion rates and expanded market reach.

Future outlook of conversational commerce via JioMart's WhatsApp channel

- **Language Expansion:** Recognizing the linguistic diversity of the Indian market, Haptik has expanded language support to include Hindi, with plans to incorporate additional regional languages in subsequent releases. This multilingual expansion is critical for reaching consumers in tier-two and tier-three cities where English proficiency may be limited, thereby extending the platform's addressable market beyond urban, English-speaking demographics.

- **Enabling increasing complexity**: Haptik is developing next-generation capabilities powered by large language models (LLMs) to handle increasingly complex customer queries. These generative AI-enhanced assistants are designed to move beyond scripted response patterns, enabling more nuanced understanding of customer intent and the ability to address sophisticated product comparisons, personalized recommendations, and context-dependent inquiries that current rule-based systems struggle to resolve effectively.

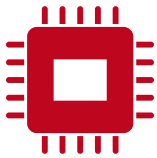

- **Broader use case replication:** The success of the JioMart implementation has generated broader industry interest in replicating this conversational commerce model. Meta has reported rising engagement from other global retailers seeking to deploy similar WhatsApp-based shopping experiences in their respective markets. This pattern suggests that the JioMart case may serve as a template for conversational commerce adoption beyond India, particularly in markets with high WhatsApp penetration and developing e-commerce infrastructure.

- **Strategic importance of conversational commerce**: Meta's leadership has articulated an ambitious vision for the future of mobile commerce in developing economies. In public statements, Meta's CEO Zuckerberg has forecasted that shopping via WhatsApp is positioned to become the default purchasing method in emerging markets, displacing both traditional e-commerce websites and native shopping applications. If this prediction materializes, it would represent a fundamental restructuring of the digital retail landscape, with messaging platforms evolving into primary commerce channels and conversational interfaces becoming the dominant paradigm for online shopping interactions.

Amazon's Rufus - launch, expansion, and set up

Amazon's entry into **conversational commerce** represents a significant development in the sector, given the company's dominant position in global e-commerce and its substantial investment in proprietary AI technologies. Unlike JioMart's partnership model, Amazon has pursued an entirely **in-house development** for its conversational shopping assistant, named **Rufus**, leveraging internal technical capabilities and vast data resources accumulated through decades of retail operations.

Rufus was **initially launched in the United States in February 2024**, followed by a beta release in India in August of the same year. The platform subsequently **expanded** to major European markets in November 2024, including Germany, France, Italy, and Spain, complementing an earlier rollout in the United Kingdom. This phased global expansion strategy reflects Amazon's approach of validating functionality in mature markets before adapting the technology for regions with distinct linguistic, cultural, and shopping behavior characteristics.

The core functionality of Rufus extends beyond simple product search to encompass a comprehensive **shopping assistance paradigm**. Users can engage in general product research, pose specific questions about items or categories, shop organized by occasion or purpose rather than predefined categories, receive updates on recently launched products, and obtain personalized recommendations based on their inquiry context. This breadth of capability positions Rufus as a **shopping concierge** rather than merely a search interface enhancement.

The **technical architecture** underlying Rufus represents a sophisticated integration of **multiple data sources and systems**. The assistant draws upon an extensive information ecosystem including customer reviews, Amazon's complete product catalog, community-generated questions and answers, and real-time data accessed through various Amazon Store APIs. This multi-source approach enables Rufus to provide responses grounded in both structured product data and unstructured user-generated content, creating a more contextually rich interaction than would be possible from catalog data alone.

Amazon's Rufus: data sources & functionality

For product-specific inquiries, **Rufus** offers support and assistance by **synthesizing information from product descriptions, product technical specifications, e.g. battery life, customer reviews, and community Q&A content**. Customers can **ask questions** about compatibility, dimensions, material composition, care instructions, or any other product-specific detail, with Rufus providing consolidated answers drawn from multiple authoritative sources.

Rufus can also help customers **stay informed** about the latest products and emerging trends within categories of interest. Users can inquire about **recent launches**, trending items, or new developments in specific product domains, allowing them to make purchasing decisions informed by current market offerings rather than potentially outdated information.

Amazon's Rufus: paradigm shift in shopping

Rufus operates as a **contextual guidance system** that fundamentally reimagines how customers interact with Amazon's product catalog. Rather than requiring users to navigate predetermined category hierarchies or formulate precise keyword searches, Rufus enables natural language interactions that mirror how shoppers might discuss their needs with a knowledgeable sales associate in a physical retail environment.

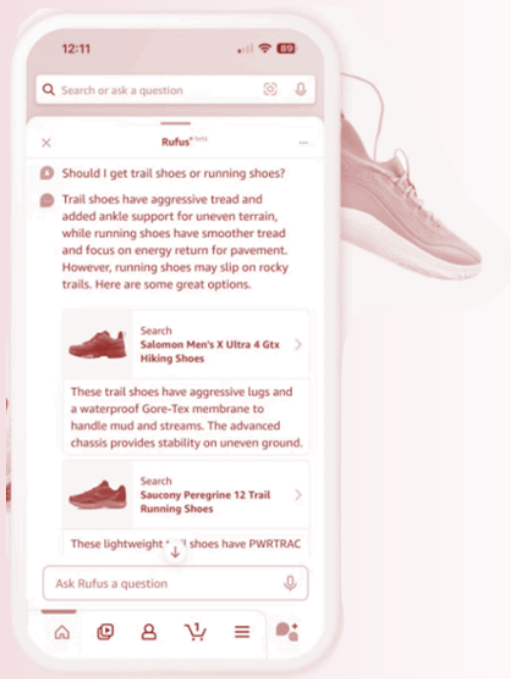

The assistant supports multiple distinct use cases that address **different stages of the shopping journey**. Customers can search for and d**iscover products based on activities**, events, purposes, and use cases, rather than relying on product categories. For example, rather than searching "camping equipment," a user might ask "what do I need for a weekend camping trip with young children?" This intent-based discovery approach allows Rufus to suggest comprehensive product assortments that address the underlying need rather than merely matching keywords.

Rufus facilitates **comparative evaluation** by enabling customers to compare entire product categories based on relevant attributes and trade-offs. Users can pose questions such as "what's the difference between air fryers and convection ovens?" and receive structured comparisons that illuminate functional distinctions, use case appropriateness, and relative advantages. This capability addresses a common pain point in online shopping: the difficulty of understanding product differentiation without hands-on experience or expert consultation.

The **recommendation engine within Rufus** provides personalized product suggestions tailored to users' exact specifications. By processing detailed requirement descriptions in natural language, the assistant can filter Amazon's catalog to surface options that align with stated preferences regarding features, price ranges, brand considerations, and use contexts. This precision matching reduces the cognitive burden of manually filtering large result sets.

Amazon's Rufus: adoption, updates and future

Adoption metrics indicate substantial user engagement, particularly in emerging markets. By November 2024, Rufus had attracted over 10 million users in India alone, suggesting strong market receptivity to conversational shopping interfaces in regions where traditional e-commerce shopping may be less adopted. While Amazon has not disclosed precise sales attribution figures for Rufus, the company's internal projections anticipate that the assistant will generate over $700 million in global profits, indicating confidence in its commercial impact.

Since launch, Amazon has continuously **enhanced Rufus's capabilities** through several major **feature updates**. Notable additions include Rufus assuming control of the primary search bar, enabling conversational queries to replace traditional keyword searches; integration with desktop search experiences, extending functionality beyond mobile interfaces; and expanded visual integration capabilities that allow the assistant to process and respond to image-based queries alongside text.

L**ooking forward,** Amazon is building upon the Rufus foundation with complementary technologies. The company recently introduced **Interests AI**, a natural language-based product discovery system that fundamentally reimagines how shoppers identify relevant products. Rather than navigating through hierarchical category structures or conducting keyword searches, Interests AI enables users to describe their needs, preferences, or use cases in natural language, allowing the system to interpret intent and surface relevant products. This evolution suggests Amazon's vision of conversational interfaces progressively displacing traditional navigation and search paradigms across the entire shopping journey..

Conversational commerce: disruption potential

Chat GPT announced a collaboration with Etsy and Shopify with direct access to shops from the chat window. (Retaildive, Sep. 2025)

Copilot

Microsoft announced the Copilot Merchant Programme (Microsoft, press release, April 2025)

amazon

Amazon announced sponsored ads will be integrated into Rufus AI assistant (Amazon, press release, Jan 2025)

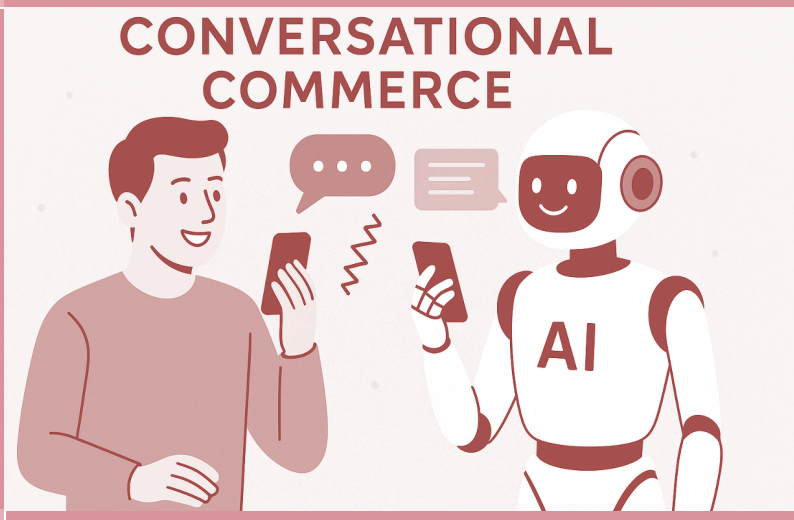

WeChat

WeChat mini-programs facilitate full in-app commerce, with ~240 billion Euro (RMB 2 trillion) in Q3 in 2024 (Tencent 2024)

High disruption potential

Conversational commerce: summary

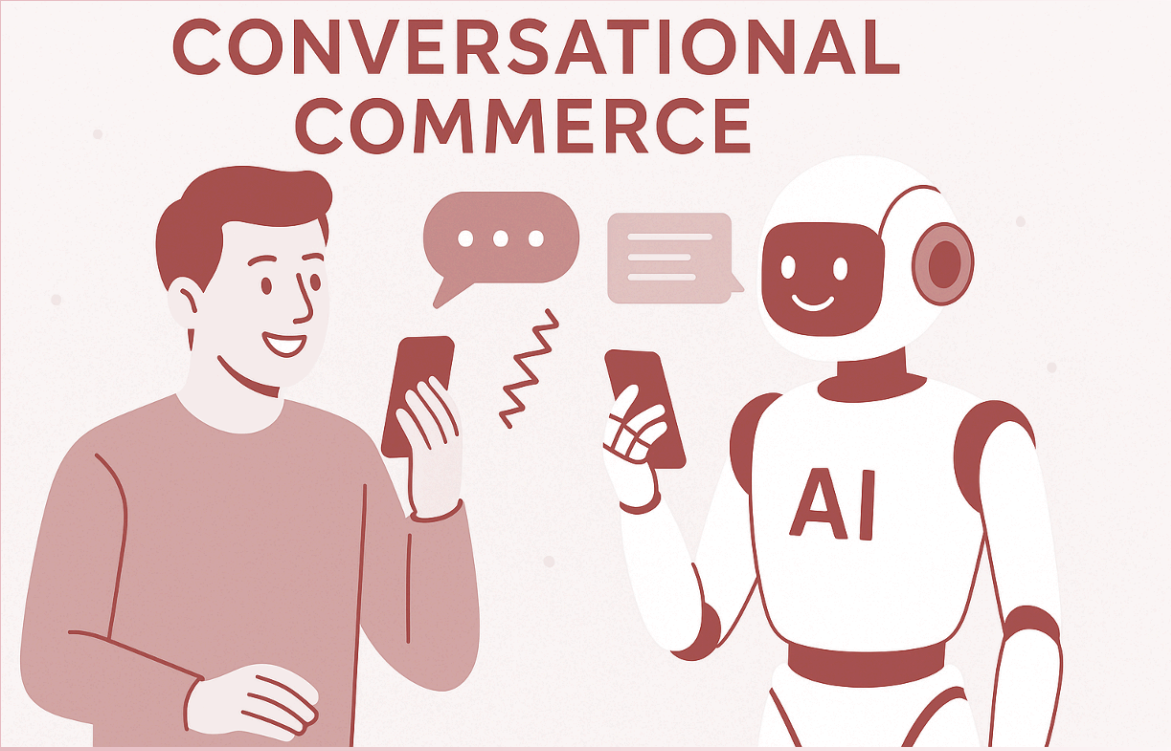

1. **Lowers barriers to e-commerce in emerging markets with high mobile phone penetration**
2. **Consumers increasingly use conversational tools for search, possibly extending to shopping**
3. **Early retail adopters signal high consumer acceptance and ROI**

Use Case 4 + 5

AI Supported Shopping (Interests AI and Shopping Guide AI, Amazon)

Amazon's other AI initatives: Rufus, Interests ai, Shopping Guide ai

While researching the GenAI use cases, specifically Rufus, we noticed **Amazon's other AI initiatives** (especially Interests ai and Shopping Guides ai) and wondered about the differences. Although strictly speaking, shopping guides ai should not necessarily be covered under generative AI, I found it useful to differentiate the three initiatives. In short, whereas Rufus enables conversational commerce, Interests AI keeps customers up to date with ongoing product recommendations, and Shopping Guide AI provides customers with a decision-making framework emphasizing important criteria when making product choices.

Ask Rufus

Rufus enables **conversational commerce**. It is available on the Amazon e-commerce website and mobile app, enabling a fully chat-based shopping experience with an end-to-end shopping journey, not only suggesting products but also including ordering and payment, all within the chat.

Shopping Guides ai provide customers with an **evaluation and decision-making framework** emphasising important criteria of their purchase decisions. The tool is designed to address a pain point in the online shopping experience: the complexity of evaluating feature-rich products, particularly in the technical domain. The shopping guides tool proactively suggests category-specific key criteria based on patterns observed in how other customers approach similar purchases (and/or the known preference of the customer).

Shopping Guides ai

Interests ai beta

Interests ai acts like an AI agent and provides customers with **automated, proactive product recommendations** and is based on the elaborate Amazon recommendation system already utilized for their e-commerce business, considering, for instance, prior purchase history, browsing history, and customer reviews. It is a feature that prompts customers to describe their interests “using their own words,” and then generates a curated selection of products from the Amazon webshop. Once the interest is generated, it continuously monitors the Amazon store for new product launches, sales initiatives, restocks, etc., that match the interest and informs customers accordingly.

Amazon's other AI initatives: Shopping Guides ai

Shopping Guides ai

Example TV

Factors to consider

...when buying a TV: Display type, resolution, size, ...

Factors to consider: Display type

...LED display, OLED, QLED, with explanations of these categories ...

Popular brands

...and their customer reviews...

Shop by type

...for instance an art like TV, an outdoor TV, etc....

Shopping guide ai: summary

1. **Provides customers with a decision making framework by identifying the main criteria**
2. **Aims to reduces complexity of decision-making for multi-feature products**
3. **Acts like an AI agent and mimics the shooping process supported by a salesperson**

Interests ai: summary

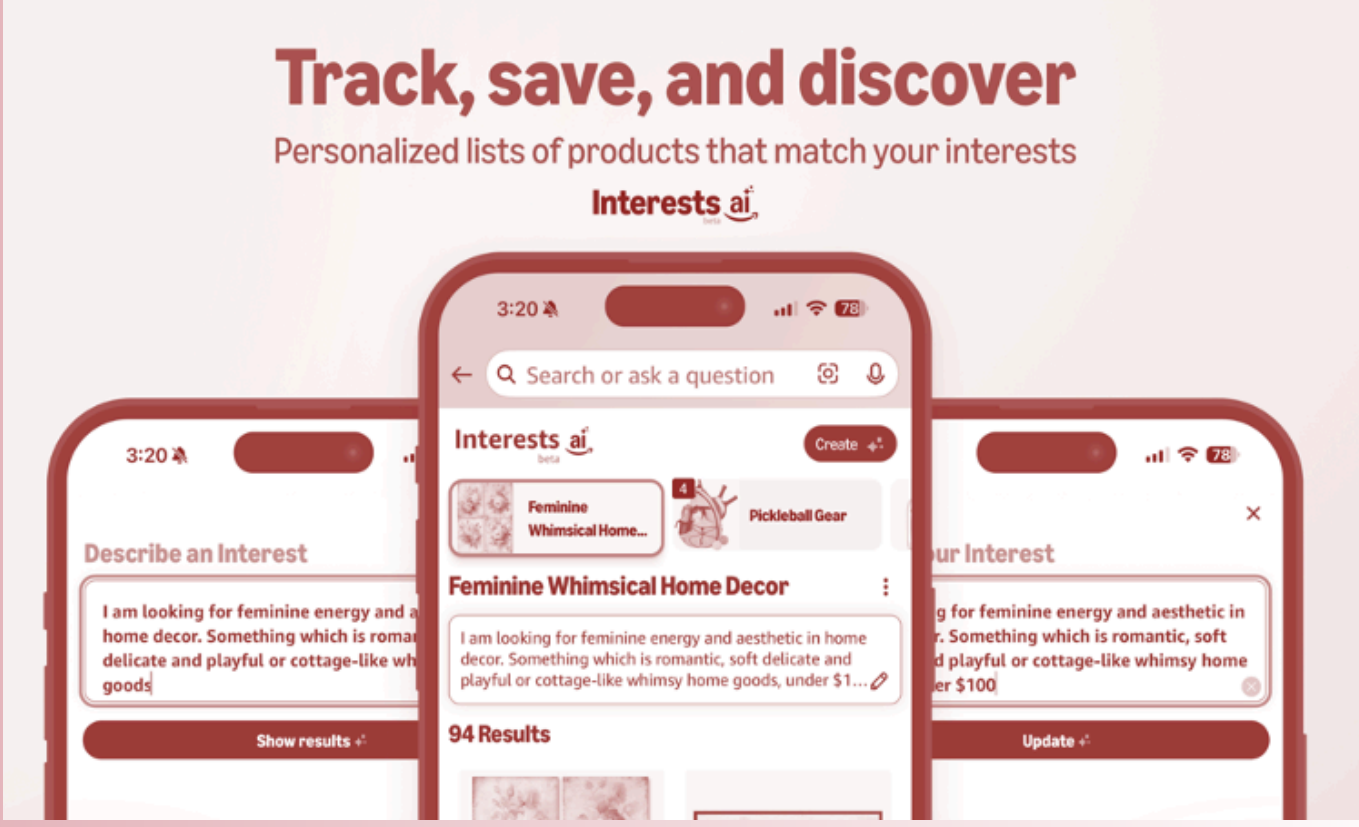

1. **Translates the customer inputted text into an interest**
2. **Provides real time recommendation based on the elaborate Amazon recommendation system**
3. **Keeps customers up to date on their specific interests, e.g. when new products are launched**

Use Case 6

Wayvee: In-Store Emotion Tracking

Wayvee: In-store Emotion Tracking

AI technologies are also transforming physical retail spaces through advanced sensing and analytics capabilities. Wayvee developed an application for an in-store, privacy-compliant, real-time emotion-tracking system based on radio sensor technology, designed to provide retailers with insights into customers' emotional responses during in-store shopping experiences.

Launched in 2024, Wayvee Analytics deploys a class of machine learning algorithms specifically trained to **infer emotional states** and psychological responses **from observable behavioral signals**. Unlike traditional retail analytics that track aggregate foot traffic or dwell time, emotion tracking systems attempt to assess the affective dimension of the shopping experience, generating data on customer satisfaction, frustration, confusion, or delight as shoppers navigate store environments and interact with products on the shelves.

The technical foundation of Wayvee's system relies on **FMCW (Frequency-Modulated Continuous Wave) radar technology** integrated with proprietary AI models. This sensor approach offers distinct advantages over camera-based computer vision systems, particularly in terms of **compliance with privacy regulations**. Rather than capturing and processing facial images through cloud-based services, which raise significant privacy concerns and regulatory complications under frameworks such as GDPR and various state-level privacy laws, FMCW radar detects movement patterns, proximity changes, and behavioral signals and analyzes them on-premises **without creating identifiable biometric data** that is shared with a third party. The AI models process these anonymized sensor inputs to predict outcome variables, including customer satisfaction levels and purchase intent, enabling retailers to assess shopping experience quality without compromising individual privacy.

The system has **progressed beyond pilot testing** into early commercial deployment, with several prominent retailers participating in implementation trials. Current retail partners include Amazon, IROHA Mart, Lenskart, CJ Express, Jim Thompson, and Carrefour, representing diverse retail formats across grocery, eyewear, convenience, specialty retail, and hypermarket categories. The number of deployed stores and sensors has not been publicly disclosed.

Wayvee: In-store Emotion Tracking

From a **cost perspective**, Wayvee's pricing model reflects the company's positioning as an accessible analytics solution for retailers of various scales. The hardware component starts at approximately $100 per sensor unit, with each sensor providing coverage for a range of 3.5 meters. This spatial coverage determines the number of sensors required based on store layout and the specific zones retailers wish to monitor. Beyond the initial hardware investment, the system operates on a Software-as-a-Service (SaaS) model priced at $240 per sensor annually, covering cloud-based analytics processing, dashboard access, and ongoing algorithm improvements.

Since launch, Wayvee has expanded the analytical scope of its platform through iterative **feature development**. The initial deployment focused on **satisfaction prediction**, establishing baseline capabilities for assessing overall customer sentiment via FMCWs.

Subsequent updates added **purchase intent prediction**, enabling retailers to identify moments when customers appear likely to convert or abandon consideration. Most recently, the system incorporated the ability to detect when **price serves as a purchase barrier**, identifying instances where customers demonstrate interest in products but appear deterred by pricing, insights that could inform dynamic pricing strategies or targeted promotional interventions.

Looking forward, Wayvee Analytics has articulated an ambitious geographic **expansion strategy**. The company plans to enter the U.S. market in 2025, transitioning from pilot programs to broader commercial rollout. The expansion approach prioritizes the United States and European Union markets initially, leveraging free pilot programs to reduce adoption barriers and generate case study evidence. Following establishment in these developed markets, the company intends to expand to additional geographic regions, capitalizing on the growing global interest in retail analytics and experience optimization technologies.

In-store Emotion Tracking by Wavee: summary

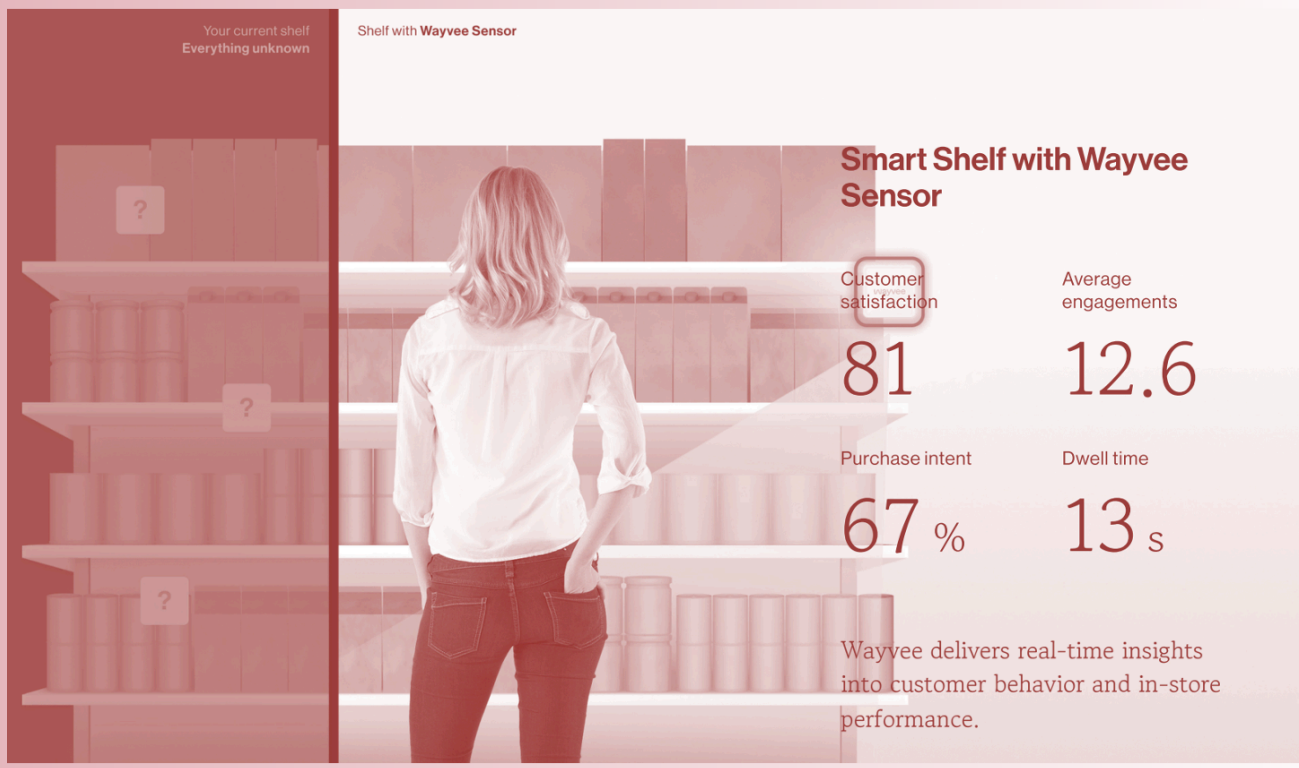

1. **Privacy-compliant detection of customer in-store micro-movements with data staying on-premise**
2. **Low-cost, low-energy sensors cover 3.5 meters also suitable for smaller stores**
3. **Real-time tracking of outputs on a dashboard: satisfaction, purchase intent, price hesitation**

Use Case 7

Planogram Optimization (Omnistream)

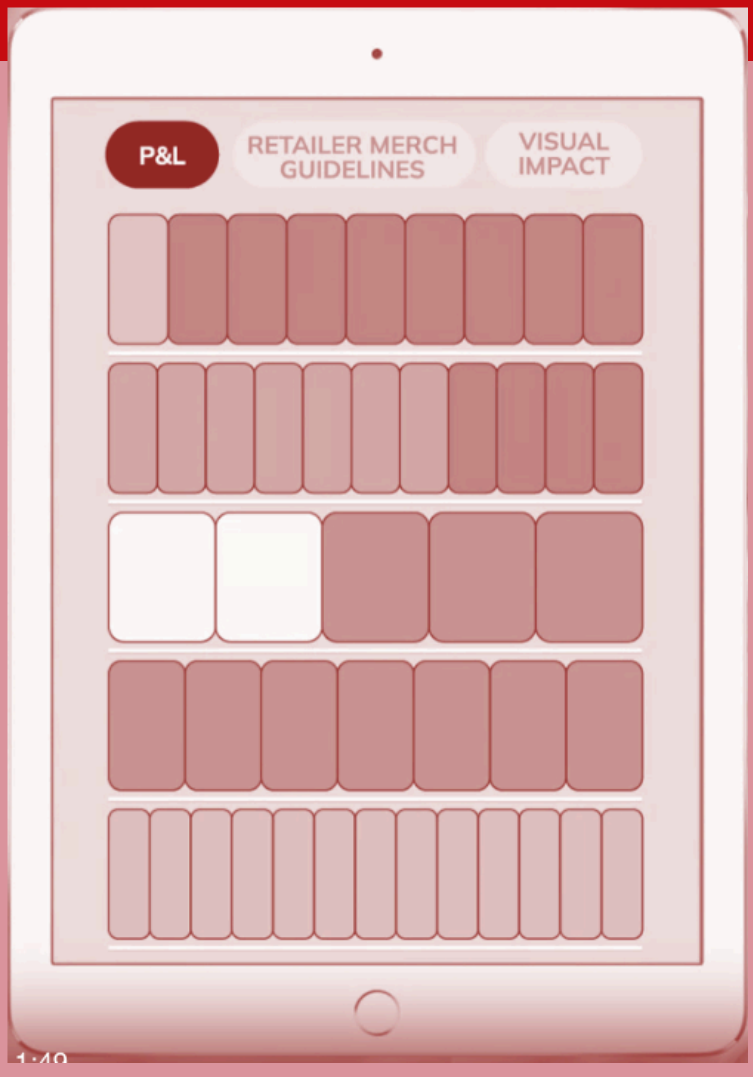

Omnistream's Planogram Optimization: Approach

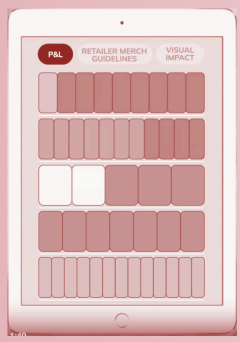

OMNISTREAM

Beyond customer-facing applications, AI technologies are increasingly deployed to **optimize behind-the-scenes retail operations** that directly impact sales performance and inventory efficiency. The collaboration between various retailers and the technology provider Omnistream exemplifies this operational AI application domain through intelligent planogram optimization systems.

Planograms, the detailed visual diagrams that specify product placement, shelf allocation, and merchandising configurations, have traditionally been developed through a combination of category management expertise, historical sales data, and negotiations with manufacturers. Often, companies rely on manual, template-based planogram creation processes that require weeks or months to complete. The Omnistream platform augments this traditional approach with AI-driven predictive capabilities, enabling retailers to **evaluate proposed planogram modifications** before implementation. The core functionality centers on **scenario-based sales forecasting**: retailers can model alternative shelf configurations and receive AI-generated predictions regarding the likely sales impact of each proposed change.

This **predictive approach** represents a significant **departure** from conventional planogram management tools, many of which **focus primarily on compliance monitoring** to ensure that store-level execution matches the intended corporate planogram design. While compliance-oriented systems use image recognition to identify discrepancies between planned and actual shelf configurations, Omnistream's value proposition emphasizes the front-end optimization process itself, helping retailers make better decisions about what the planogram should be rather than merely verifying adherence to existing plans.

Beyond compliance: Planogram modifications can be evaluated with a scenario-based prediction before implementation

Omnistream's Planogram Optimization: model & costs

The **technical architecture** underlying Omnistream's platform integrates **image recognition** capabilities with **proprietary AI algorithms**. Image recognition enables automated capture and digitization of current shelf states, eliminating the manual effort traditionally required to document existing configurations. These digitized planograms then serve as inputs to AI models that have been trained on historical sales data, product attributes, category dynamics, and presumably numerous retailer-specific factors that influence purchase behavior in convenience retail contexts.

Omnistream has achieved substantial **commercial traction** beyond the initial 7-Eleven Philippines deployment in 2021. The platform currently serves various **retail clients across multiple retail formats and product categories,** indicating the applicability of their technology. Clients include 7-Eleven Philippines (convenience), FairPrice (grocery), Supabarn (grocery), Watsons (health and beauty), and Pharmacity (pharmacy). The geographic footprint spans over 8,500 stores across more than 10 markets, suggesting successful adaptation to diverse regional contexts, regulatory environments, and consumer preferences.

Regarding **implementation costs**, Omnistream has not publicly disclosed its pricing structure. This is common in enterprise software solutions where pricing typically varies based on several factors such as store count, transaction volume, customization requirements, and contract duration. The absence of standardized public pricing may also reflect a consultative sales approach where solutions are tailored to individual retailer needs rather than offered as standardized packages. It might also reflect that pricing might be based on improving sales figures.

Information regarding **future developments** and expansion plans similarly remain undisclosed. However, the platform's current multi-market presence and diverse client base suggest that Omnistream has **moved beyond initial market validation** into a growth phase, with ongoing geographic expansion and potential functional enhancements likely under development.

Omnistream case study: Fortune 500 Grocery Chain

In the following, we showcase **one Omnistream implementation case** of a Fortune 500 grocery chain.

Pre-Implementation challenge: Prior to adopting Omnistream, the Fortune 500 grocery chain exemplified many of the inefficiencies characteristic of traditional planogram development approaches in large retail organizations. The company relied on manual, template-based planogram creation processes that required months to complete, with timelines frequently exceeding 20 weeks from initial category review to finalized store-level plans. Large space-planning teams remained mired in spreadsheet-based workflows, manually evaluating product performance data and translating category strategies into physical shelf layouts.

This timeline created operational difficulties. **Store compliance** with centrally developed planograms varied significantly, as generic templates often proved ill-suited to individual store configurations or local demand patterns. **Return on investment** from category resets remained difficult to quantify, as extended development cycles often left planograms **outdated** by the time they reached implementation. Furthermore, the resulting layouts sometimes exhibited suboptimal merchandising characteristics: product categories were distributed across non-contiguous shelf sections, some categories were oriented horizontally while others stretched vertically without a clear logic, and shopper personas struggled to locate related products efficiently. **Space utilization** stood at just 90%, indicating that a tenth of available shelf space remained either unallocated or assigned to unproductive inventory.

Solution Architecture: The Omnistream implementation centered on three interconnected capabilities: 1) the platform **utilised and enriched the retailer's existing data foundation** by integrating point-of-sale transactions, fixture specifications, and promotional calendars with loyalty program data, then augmenting these internal sources with external demographic and mobility datasets to create comprehensive store-level demand profiles. 2) Omnistream's AI-powered optimization engine **transformed the category review process** itself, enabling rapid generation of store-specific planograms from these enriched inputs while optimizing for both accuracy and speed. 3) The platform expanded staff capabilities by **automating manual computational work**, allowing category managers and space planning teams to focus on simulation, strategic evaluation, and vendor negotiations while the AI handled the labor-intensive aspects of planogram creation.

Omnistream case study: Fortune 500 Grocery Chain

The deployment across more than 100 stores yielded **substantial improvements across temporal, financial, and operational dimensions**.

The most dramatic impact was observed in **planning cycle time**, as the retailer compressed planogram development from over 20 weeks to under 4 weeks, representing an approximately 80% **reduction in time-to-implementation**. This acceleration enabled more frequent category reviews and faster response to market dynamics.

Financial performance improvements were significant, with pilot stores achieving a 7% blended category sales increase attributable to optimized assortments and improved product placement. Space utilization increased from 90% to 96%, indicating that the AI-driven approach more effectively allocated shelf space to productive SKUs while minimizing allocation to slow-moving inventory. This 6-percentage-point improvement in range utilization translated directly to revenue-generating capacity previously left unrealized.

Operational benefits extended beyond quantifiable metrics. The introduction of store-relevant planograms improved compliance rates, as store teams found the localized assortments more practical to implement and maintain. The improved layouts enhanced shopability by consolidating discrete but related categories into contiguous, logically organized blocks that aligned with shopper decision trees. Perhaps most significantly, the reduction in manual planning effort enabled workforce redeployment, with category management teams redirected from spreadsheet manipulation toward higher-value activities including vendor relationship management and strategic planning.

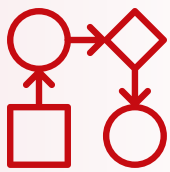

This case study demonstrates that AI-powered planogram optimization can deliver value not merely through incremental performance improvements but by fundamentally **restructuring the economics of space planning**. By compressing cycle times by 80% while simultaneously improving sales outcomes by 7%, the platform resolves what has traditionally been presented as an unavoidable trade-off between speed and quality in retail category management. The ability to generate accurate, store-level plans in weeks rather than months enables a shift from periodic, resource-intensive category resets to continuous optimization responsive to evolving market conditions.

Omnistream's Planogram Optimization: summary

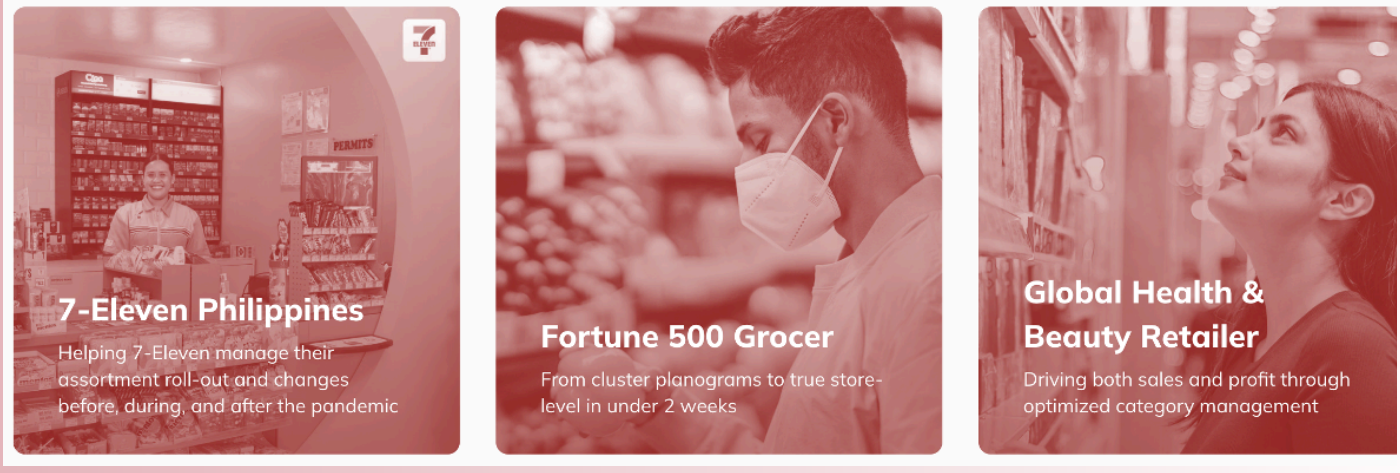

1. **The system is focused on planogram scenario planning rather than compliance checks**
2. **Prompts retailers to restructure and speed up their planogram creation process**
3. **Positive impact on various operational and financial outcomes for the retailer and customers**

USE CASE 8

CONTENT MGT PLATFORM (JD.COM)

JD's content management platform based on RhinoChat

JD.com is China's largest e-commerce retailer with annual revenue of 158.8 billion USD in 2024. It operates primarily as an **e-commerce platform but also manages a substantial marketplace** where retail partners can sell their merchandise. JD.com developed a comprehensive AI-powered content management platform based on its own Large Language Model ChatRhino to create, e.g., product descriptions, photos, videos, and other marketing content.

ChatRhino, known as Yanxi in Chinese, is **JD.com's proprietary LLM model** on which the **AIGC Content Marketing Platform** is based. The platform was unveiled at the company's JD Discovery tech summit in July 2023. The platform represents a significant advancement in JD.com's AI capabilities, building upon earlier AI models launched in 2021 and 2022. ChatRhino was specifically designed for retail applications, combining general knowledge with JD.com's extensive data on products, supply chains, and customer behavior. The platform was developed through a strategic **partnership** involving **JD Cloud, JD Technology, and DeepSeek**. It was initially deployed for internal use across JD.com's e-commerce, logistics, and marketing operations. Following this internal rollout, the platform became available to external business clients in early 2024, with ongoing enhancements continuing through 2024-2025.

JD's content management platform: step by step example

Step 1

Example: Haruto, a winemaker from Japan, decides to sell his wine on the JD.com platform. He logs into the JD merchant portal and navigates to the AIGC Content Marketing Platform. The platform is available free of charge to JD merchants as part of JD.com's commitment to supporting seller success, particularly for small and medium-sized enterprises.

Step 2

Haruto uploads a single, representative photo of his product, taken with an ordinary smartphone. The system analyses the photo and extracts product characteristics and features, specifications relevant to the product category, visual elements that will resonate with potential customers, and appropriate styling and presentation approaches for the product type. This process occurs automatically in the background, requiring no additional user input.

Step 3

The Content Marketing Platform generates an extensive range of marketing material customized to his product, such as e-commerce product images optimized for JD.com's listing requirements, with appropriate sizing and formatting for different display contexts (search results, product detail pages, mobile views). Product images can present the product from multiple perspectives and in various contexts, with different backgrounds, lighting treatments, and compositional approaches. It can also produce posters with compelling visual design, appropriate text overlays, and styling consistent with successful campaigns in your category

Step 4

Haruto can now browse through the generated assets presented in a gallery view in the Content Marketing Platform. The platform offers multiple options for each asset type, allowing users to select images that best represent the brand and its product positioning. The user can preview how each asset will appear in different JD.com contexts before making final selections.

JD's content management platform: livestreaming

Beyond static marketing content, ChatRhino also powers **JD.com's digital presenter solution** for **livestreaming commerce**, addressing one of the most resource-intensive aspects of modern e-commerce: live shopping events. The platform automates both script generation and presentation delivery, fundamentally transforming the economics and scalability of livestream shopping.

The livestreaming tool leverages ChatRhino's understanding of product attributes and consumer psychology to automatically **generate engaging livestreaming scripts**. Utilizing JD Cloud's e-commerce knowledge model, K-PLUG (Knowledge-injected pre-trained language model for natural language understanding and generation in e-commerce), merchants simply upload their product URL. The system extracts product details, key specifications, and selling points, transforming this technical information into reader-friendly, conversational livestreaming scripts designed to drive engagement and conversions. Traditionally, creating livestreaming scripts required hours of work from experienced content creators familiar with both the products and effective livestreaming techniques.

Once scripts are generated, the platform creates **lifelike digital presenters** to deliver the content. These AI-powered avatars can be customized to match brand aesthetics or, in a particularly compelling application, replicate the appearance, voice, and mannerisms of real individuals. The technology captures not only visual likeness but also vocal patterns, accents, and characteristic gestures, creating remarkably authentic digital representations.

JD's content management platform: livestreaming

To demonstrate the sophistication of this technology, JD.com orchestrated a **high-profile launch** featuring **digital avatars of prominent business leaders**. The company cloned its own CEO, Richard Liu Qiangdong, along with 17 other CEOs and business celebrities, to deliver virtual messages showcasing the platform's capabilities.

Liu's AI avatar made its public debut in April 2024, hosting two livestreaming sessions that promoted consumer electronics and groceries. The **virtual CEO** successfully replicated Liu's distinctive communication style, including his well-known habit of gesticulating while speaking. These inaugural sessions generated over 20 million combined views, validating the technical quality of the avatar technology and indicating consumer acceptance.

While the **avatar primarily adhered to prepared scripts** rather than engaging in spontaneous interaction with viewers, the demonstration showed that virtual presenters could command audience attention and drive commercial results comparable to human hosts. The success of this launch positioned ChatRhino's digital presenter technology as a viable alternative to traditional livestreaming approaches, with significantly **lower operational costs and scaling potential.**

JD's private label brand J.Zao demonstrated the platform's capabilities by **generating a suite of marketing materials** for a new product line **within minutes**. The brand uploaded product photos and received professionally designed marketing materials covering all required touchpoints, such as product listings, promotional campaigns, and social media assets, in a **single automated workflow**.

As retailers use the platform, **ChatRhino learns from their selections and preferences**. Over time, the AI refines its output to better match your brand aesthetic and the preferences of your specific customer base. The platform also provides performance analytics showing how AI-generated assets perform compared to traditionally created materials, enabling data-driven decisions about your visual merchandising strategy.

Content Marketing Platform based on RhinoChat: Summary

1. **Support of smaller retailers with low volumes, removing the need for costly professional photos**
2. **Professional-looking photographs optimized for the JD.com platform**
3. **Faster development allows more adaptability to market trends**

Use Case 9

Sales Contract modelling (HeiDeal)

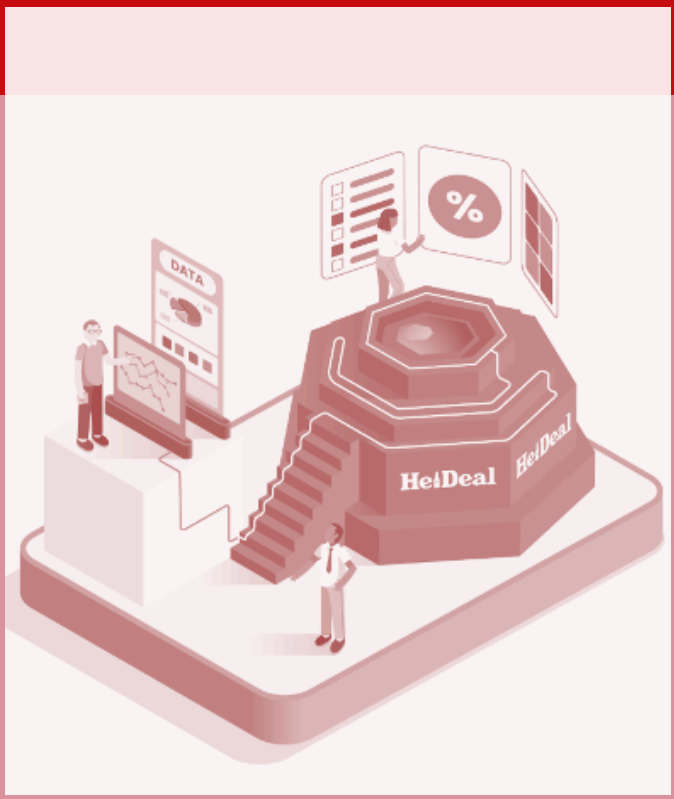

HeiDeal: Sales Contract Modelling

Business challenge and context: This case study is from Heineken UK, which operates in a complex sales environment, supplying beer, cider, and other beverages to approximately 2,400 company-owned pubs throughout the United Kingdom, as well as serving independent venues, including bars and restaurants. The company's sales representatives negotiate customized supply contracts, colloquially termed "deals" within the organization, that must balance multiple competing considerations: customer profitability, competitive pricing, supply chain logistics, and internal approval workflows.

Prior to HeiDeal's implementation, contract development required **substantial manual effort**. Sales representatives needed to assess customer purchase history, evaluate pricing structures across multiple product lines, model profitability implications under various discount scenarios, navigate multi-stage approval processes, and coordinate with operations teams regarding supply logistics and billing arrangements. This manual approach consumed **significant time**, introduced potential for **calculation errors**, and limited representatives' ability to explore **alternative deal structures** that might better serve both customer needs and company profitability objectives.

HeiDeal addresses these challenges through a data-driven and AI-supported system that **integrates** directly with the existing **SAP enterprise resource planning infrastructure**. The platform was designed around a modular architecture accessing application programming interfaces (APIs) coupled with a developed AI-supported deal-building feature. This architectural approach enables the system to pull real-time data from SAP while maintaining sufficient flexibility to accommodate the nuanced requirements of contract negotiations in the on-premises beverage trade.

HeiDeal: Sales Contract Modelling

The platform's core functionality is sales contract **scenario modeling**. Sales representatives input customer parameters and proposed contract terms, and HeiDeal on the spot **creates multiple alternative configurations**, assessing each configuration against key performance dimensions, including customer profitability projections, pricing analytics benchmarked against comparable accounts, supply chain feasibility given production and distribution constraints, and alignment with corporate profitability thresholds and strategic priorities.

The system **automates numerous manual steps** of deal approval and execution. Routine elements of ordering, billing configuration, and standard approval workflows execute automatically when deals meet predefined criteria, freeing sales representatives to focus on relationship management and strategic account planning rather than administrative processing.

The intelligence layer enabling HeiDeal's scenario modeling and optimization capabilities is built on **Microsoft's Azure AI platform**. This cloud-based artificial intelligence infrastructure provides the computational resources and machine learning frameworks necessary to process Heineken's substantial historical transaction data, identify patterns in successful deal structures, and generate recommendations that optimize across multiple objectives simultaneously. The Azure integration also positions the platform for continuous improvement, as the AI models can be refined over time based on actual deal performance and evolving market conditions.

The platform represents a **substantial investment**, requiring **~18 months of development** in partnership with AND Digital. This extended timeline reflects not merely software engineering complexity but also the necessary work to deeply integrate HeiDeal with Heineken's existing SAP systems, encode business rules and approval workflows accumulated over years of operational experience, and design **user interfaces appropriate for field sales environments** where representatives may be accessing the system during on-site customer visits via mobile devices. There is no publicly available information on the financial investment for HeiDeal.

Heineken

AND Digital

HeiDeal: Sales Contract Modelling

Heineken UK pursued a **phased rollout**, initially deploying HeiDeal to a subset of its sales force to **validate functionality** and gather user feedback before broader implementation. Even though not explicitly mentioned, it is likely that the 2400 Heineken-owned UK pubs were a focus of the initial rollout. By January 2025, six months after the August 2024 launch, ~150 sales representatives were actively using HeiDeal. This phased rollout allows the organization to refine workflows and address usability concerns while building internal capability and confidence with the new system.

The transaction volume processed through HeiDeal during this initial **six-month period** reached **3,000 deals**, indicating substantial engagement from the deployed user base and demonstrating that the platform has moved beyond pilot testing into **operational reliance.** This volume represents a significant portion of Heineken UK's contract negotiations during the period, suggesting that HeiDeal has successfully integrated into representatives' daily workflows rather than serving merely as an optional tool.

Heineken has identified **HeiDeal as a potential template for international deployment.** The platform's modular architecture and API-based design facilitate localization, as country-specific pricing rules, regulatory requirements, and customer data can be integrated without requiring fundamental restructuring of the core deal-building engine. This adaptability positions HeiDeal as a scalable solution applicable across Heineken's global operations, which span numerous markets with distinct commercial practices and competitive dynamics. However, international rollout will require careful consideration of local market conditions, such as different pricing structures, customer relationship norms, distribution models, and regulatory frameworks.

By embedding analytical capabilities directly into the workflow of sales representatives, the platform elevates negotiations without requiring individual representatives to possess deep expertise in profitability modeling or supply chain optimization. This **democratization of analytical capability** enables more consistent deal quality across the sales organization while simultaneously improving efficiency through process automation.

HeiDeal: Sales Contract Modelling: summary

1. **Business customers gain access to multiple contract options tailored to their specific needs**
2. **Democratized sales expertise through standardized contract modeling, localisation options**
3. **Manual labor and approval processes were removed from contract modeling, increasing efficiency**

Use Case 10

Autonomous Contract Negotiation (Walmart)

Walmart's Autonomous Supplier Negotiation Platform

Walmart & PACTUM

Walmart's collaboration with Pactum, creating a fully autonomous business-to-business contract negotiation platform, also represents a non-customer-facing use case. It takes the Heineken case one step further by offering contract negotiations to suppliers in a self-service manner rather than supporting the sales representatives. This implementation demonstrates how agentic AI systems can act independently within defined parameters and, in turn, transform procurement processes at enterprise scale.

Context of the development: Walmart US faces a procurement challenge common to large retailers: managing contractual relationships with approximately 100,000 suppliers across diverse product categories, service providers, and geographic regions. While Walmart's procurement resources are appropriately allocated to negotiating high-value contracts with strategic suppliers, the "tail end" of the supplier base, typically smaller vendors providing specialized goods or services, has historically received less attention, despite representing substantial aggregate expenditure.

Walmart initiated its exploration of autonomous negotiation through a measured **rollout** beginning in 2022. Rather than immediately applying the technology to merchandise procurement, Walmart started with **non-customer-facing product categories** such as fleet services and logistics support. This approach **reduced risk** by testing the platform in contexts where negotiation failures would not directly impact product availability or customer experience. Additionally, Walmart piloted the system at Walmart Canada before broader U.S. deployment, further validating functionality in a smaller, controlled market environment.

The negotiation platform operates as an **autonomous negotiation agent that covers the entire process:** a chatbot that conducts complete contract negotiations from initial outreach through final agreement, without human intervention for routine cases. The system is built on what the industry terms "agentic AI," meaning it possesses decision-making authority within predefined boundaries rather than merely providing recommendations for human review.

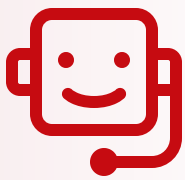

Walmart's Autonomous Supplier Negotiation Platform

Prior to engaging suppliers, Walmart's procurement team establishes **negotiation parameters** that encode the company's strategic priorities and **acceptable trade-offs**. These parameters might, for example, specify that Walmart prioritizes extended payment terms over marginal price reductions, or that it values supply reliability guarantees more highly than volume discounts. The **AI agent** then **operates within these constraints,** adapting its negotiation tactics based on supplier responses while ensuring all agreements remain within Walmart's acceptable outcome range.

The platform primarily targets tail-end suppliers, which are smaller vendors that **typically operate under standard contract terms rather than bespoke agreements** negotiated by dedicated account managers. For these suppliers, the autonomous system offers an opportunity to optimize contracts that would not otherwise justify Walmart's procurement team's resources. Suppliers interact with the chatbot through a conversational interface that presents proposals, responds to counteroffers, and articulates the rationale behind Walmart's positions, mimicking the dynamic exchange that typically occurs in human-to-human negotiations.

Adoption: Walmart's approach to **engage suppliers** balanced encouragement with voluntary participation. The company approached its tail-end supplier base with invitations to engage with the autonomous negotiation platform, positioning it as an opportunity for suppliers to customize their contract terms rather than operating under standard agreements. The **response rate** was substantial: 68% of approached suppliers engaged with the negotiation chatbot, indicating their receptivity to the concept despite its novelty.

This **high engagement rate** suggests that suppliers recognized value in the process beyond mere compliance with a major customer's request. For small suppliers who rarely have direct access to Walmart's procurement decision-makers, the platform provided a mechanism to advocate for terms better suited to their business, for instance, achieving more favorable payment schedules, adjusted minimum order quantities, or modified delivery requirements, all within a framework that Walmart had predetermined as acceptable.

Walmart's Autonomous Supplier Negotiation Platform: Business Impact

The platform's **performance** validated the autonomous negotiation concept across multiple dimensions. Approximately **two-thirds of initiated negotiations culminated in mutual agreement,** indicating that the AI agent successfully identified contract modifications acceptable to both parties in the majority of cases. This success rate compares favorably to human-conducted negotiations, particularly considering the volume of concurrent negotiations the platform could manage.

Walmart's **strategic priority** in the initial deployment phase focussed on **improving payment terms,** extending the duration between invoice receipt and payment obligation, thereby improving the company's working capital position. This focus reflected a calculated approach: payment terms represent a zero-sum negotiation dimension where Walmart's gain directly translates to supplier cost, making it an appropriate testing ground for autonomous negotiation capabilities before applying the technology to more complex, multi-dimensional contract structures.

The financial impact proved meaningful even in the pilot phase. Walmart achieved approximately **1.5% in procurement savings** through the renegotiated contracts, a figure that, while seemingly modest in percentage terms, translates to substantial absolute savings given Walmart's procurement expenditure scale. Beyond direct savings, the platform **accelerated negotiation timelines** significantly compared to manual processes, enabling faster contract finalization and reducing the administrative burden on both Walmart's procurement team and supplier account management.

Perhaps most strategically valuable, the autonomous system **freed Walmart's human procurement** professionals to **redirect** their expertise toward **high-value, complex negotiations** with strategic suppliers. By automating routine tail-end supplier negotiations, the Walmart effectively expanded its procurement capacity without proportional headcount increases.

Walmart's Autonomous Supplier Negotiation Platform: Business Impact

The platform relies on **Pactum's proprietary agentic AI** specifically designed for enterprise procurement applications. Unlike general-purpose conversational AI systems, Pactum's technology incorporates domain-specific knowledge about procurement practices, contract structures, and negotiation dynamics. This specialization enables the platform to navigate the nuanced exchanges characteristic of business negotiations, including understanding when to make concessions, how to package multiple terms into balanced proposals, and when to escalate to human review for exceptional cases.

The specific **financial arrangements** between Walmart and Pactum remain **undisclosed**, though enterprise AI platforms of this sophistication typically employ either **licensing models** based on transaction volume or **performance-based pricing** where compensation correlates with measurable value creation. The latter approach aligns vendor and client incentives, as the technology provider benefits directly from successful negotiations that generate savings or other measurable improvements.

Following successful validation, Walmart has **continued rolling out the platform**, expanding both the supplier segments addressed and the breadth of contract terms subject to negotiation. This expansion reflects **organizational confidence** in the technology's reliability and strategic value, as well as supplier base adoption to engaging with AI agents in contractual contexts. This suggests that autonomous negotiations are a permanent evolution in procurement. As the platform accumulates negotiation data and outcomes, machine learning capabilities enable **continuous refinement** of negotiation strategies, potentially improving success rates and value capture over time.

Walmart's autonomous negotiation implementation represents a category of AI application fundamentally different from the pattern recognition, content generation, or recommendation systems that dominate retail AI discussions. **Agentic AI systems** act with **delegated authority rather than merely advising human decision-makers**. This approach introduces new considerations around governance, accountability, and organizational trust. Walmart's phased approach, starting with lower-risk categories, validating the system in smaller markets, targeting routine negotiations, and maintaining human oversight for complex cases, provides a template for organizations considering similar transformations in domains where autonomous AI agents negotiate, transact, or commit organizational resources.

Walmart's Autonomous Supplier Negotiation Platform: summary

1. Enables smaller suppliers who usually have standard contracts to design it to their needs
2. The platform had a high adoption and continually improves with new datapoints
3. The platform had a measurable business impact for Walmart

Main Sources

Case 1

Albert Heijn (2025): Albert Heijn zet volgende stappen in Generative AI, Press releases,
Albert Heijn (2025): App store, AH app
ESM Managzin (2024): Albert Heijn Adds New Gen AI-Based Feature 'Scan & Kook' To Its App
Fridgeleftoversai.com (2025): Features, How it Looks, FAQ

Case 2

Cal.ai (2024): Homepage, April/May 2024.
Morris, Marc (2023): The 3 Best AI Calorie Tracking Apps in 2025, YouTube.

Case 3, Jio Mart

Economic Times (2023): JioMart witnesses 7x growth in monthly orders via WhatsApp, 29. September 2023.
Facebook/Meta (2022): Shop on WhatsApp with JioMart in India, Meta News, August 2022.
Haptik (n.d.): JioMart - First Ever End-to-End Shopping Experience on WhatsApp, Case Study.
JioMart (2020): Homepage, Launch May 2020.
Meta Chat Center (2024): The Rise of Conversational Commerce in 2024: JioMart's Success Story on WhatsApp by Gonzalo Baez.
TechCrunch (2019): India's Reliance Jio unveils video call assistant to help businesses automate customer support, in: TechCrunch, 14. October 2019 by Manish Singh.
Vertex Suite (n.d.): Why only JioMart WhatsApp has Pin Code Search Bar and Category Features.
WA Notifier (2025): WhatsApp Statistics 2025: Usage Trends, Demographics & More. By Ram Shengale 4. November 2025.
YouTube (2018): How Haptik Augmented Mondelez India's Seasonal Marketing Strategy for Cadbury, by Haptic on 30. July 2018.

Case 3, Rufus

Amazon India (n.d.): Amazon announces the launch of Rufus, a new generative AI-powered conversational shopping assistant, in India.
Amazon Science (n.d.): The technology behind Amazon's GenAI-powered shopping assistant Rufus.
Amazon (n.d.): Artificial Intelligence Amazon Features Interest.
Amazon India (n.d.): Rufus - AI Shopping Assistant.
Amazon EU (2024): Amazon announces the launch of Rufus, a new generative AI-powered conversational shopping assistant in beta across Europe, by Amazon Team 19. November 2024
Amazon (2025): New Year, New Rules: AI-Powered Insights & Policy Shifts for 2025, by Arishekar N. January 2025.
Azoma.ai (n.d.): Amazon Rufus Actual Usage.
Business Insider (2025): Amazon predicts 700 million potential gain AI assistant Rufus, in: Business Insider, April 2025.
Ecommerce AI (n.d.): The Rufus Update: All the Latest New Changes on Amazon.
India AI (n.d.): Amazon launches AI shopping companion Rufus to enhance user experience.

Case 4 + 5

About Amazon (2025): Artificial Intelligence Amazon Features Interest, by Daniel Lloyd, 27. March 2025.
Baeldung (2024): How Does the Amazon Recommendation System Work? By Safa Bouguezzi and Milos Simic, 18. March 2024.

Case 6

Timashev, A. (2025): From NRF straight to Green Week in Berlin, LinkedIn-Post, 30. January 2025.
Unite.AI (2024): Alex Ovcharov, Founder & CEO of Wayvee Analytics – Interview Series, by Antoine Tardif, 8. November 2024.
Wayvee (n.d): Pilot Platform Real-time CSAT Analytics.
Wayvee (2025): Wayvee Analytics- Anonymous Emotion AI Technology posted by CONNECTpreneur on YouTube.
Wayvee (2024): Real-time CSAT Analytics - Apply for Free, LinkedIn-Post von Wayvee, 7. November 2024.

Case 7

Omnistream (2025): NRF 2025 Retail Innovation, LinkedIn-Post, Omnistream, 16. January 2025.
Trax Retail (n.d.): Revolutionizing retail through AI-driven data, Homepage, 2010.

Case 8

Intellias: (2025): Generative AI in Retail: Use Cases, Examples, Implementation 24. February 2025
PingWest (2023): Beyond the Buzz Chinese e-commerce platforms use generative AI for boosting sales during 618 shopping festival.

Case 9

AND Digital (n.d.): Heineken UK - Our Work.
Dunlop, A. (2024): AND Digital partners with Heineken for digital innovation, LinkedIn-Post, 22. August 2024.
Heineken (2025): Heineken launches global GenAI Lab in Singapore, Newsroom, 19. March 2025.
YouTube (n.d.): Heineken Digital Innovation, Video.

Case 10

Pactum AI (2024): How Pactum AI is advancing negotiation science, posted on YouTube.
Pactum (n.d.): The Leading Agentic AI Platform for Enterprise Procurement.
PYMNTS (2023): Walmart Finds 75 Percent Vendors Prefer Negotiating with Chatbot, in: PYMNTS, 2023.
Van Hoek, Remko, Michael DeWitt, MaryLacity and Travis Johnson (2022): How Walmart Automated Supplier Negotiations, in: Harvard Business Review, 8. November 2022.
Van Hoek, Remko and Mary Lacity (2023): Procurement in the Age of Automation, MIT Sloan Review, 31. August 2023

Declaration

Conflict of interests: As mentioned in the foreword, I'm sharing my observations on upcoming GenAI cases in retailing, hoping the reader finds them helpful. The choice was independent. I have **no conflict of interest** relating to any of the named companies. I have not received anything for naming them in this white paper, and there are no relevant personal contacts or financial relations to these firms.

Contracted research: This white paper has been developed as contracted work for Lekkerland SE, Headquartered in Cologne, Germany, a firm I have been working with since 2018 and have regularly conducted contracted studies on retailing topics. Lekkerland had influence on the topic but none on the process or the outcome of this white paper.

Third-party contributors: This white paper aims to identify less well-known GenAI applications, which is why it places particular focus on Asia. First, Asia is often ahead of the curve in developing GenAI applications. Second, because of language barriers, developments might go more unnoticed in the global retailing world. I speak two languages fluently: English (my day-to-day language) and German (my mother tongue), which is why I sought support from native speakers for identifying the use cases. Specifically, I was supported by Yuanchin Lee and TongTong LI (China), Zeffa Aprilasani (South East Asia), Sandeep Munasinghe (India), and Tomoko Yoshimi (Japan and Korea). I thank them for their valuable contributions in the search for use cases.

GenAI use: This white paper explores the incredible capabilities of GenAI and how it is transforming retail. As a scholar, I'm in awe of the advancement in the quality of the outcomes. It would be remarkable to write a white paper on GenAI without using GenAI. Herewith, I disclose that GenAI was used throughout the process as support, not a replacement for human intelligence. For my desk research, I mainly used Undermind, Perplexity, and Claude. Claude and Grammarly supported my writing. For image generation, I primarily used ChatGPT. The white paper was produced with support from Venngage and its proprietary AI tools for icon and image generation.

Despite the above, I take full responsibility for the report. If you like to contact me about this, please reach out via LinkedIn.
October 2025

Sabine Benoit, Professor of Marketing (Practice),
Singapore Management University