WHITEPAPER

Selling in 2026: From Intent to Emotion-Aware, Agentic Commerce

How AI is reshaping how people buy and how businesses must respond

May **2026**

Drafted by
Intueri Consulting LLP

Contents

Section 1: Executive Summary 2
Section 2: Introduction - The Signals Have Changed 3
Section 3: Literature Review 4
Section 4: The New Selling Landscape: Commerce Becomes Perceptual 5
Section 5: Agentic Commerce - What It Actually Is & Why It Changes Everything 6
Section 6: Reading the Room - Microexpressions, Behavioural Cues, and the New Data Layer 11
Section 7: The Interface Layer - Building an Emotionally Intelligent Commercial Presence 12
Section 8: The Shifting Consumer - Loyalty, Trust, and the Fragmentation of Attention 13
Section 9: The SENSE-SIGNAL-SELL-LEARN Framework: Intueri's Strategic Model 15
Section 10: Strategic Imperatives for Business Leaders 17
Section 11: From Strategy to Implementation 19
Section 12: Conclusion 21
Section 13: References 22

Section 1: Executive Summary

Selling is undergoing its most significant structural shift in a generation. For D2C brands, SMEs, and established local businesses looking to modernise, this shift is both an urgent challenge and a genuine opportunity. The question facing these businesses is no longer how to reach customers - it is how to be understood by the AI agents increasingly acting on their behalf, and how to read the human signals those agents cannot fully capture.

This whitepaper argues that three forces are converging to rewrite the rules of commerce: the maturation of affective computing, the rise of autonomous AI agents capable of executing transactions without human intervention, and a fundamental shift in consumer behaviour toward AI-mediated decision-making. Together, these forces are moving selling from a discipline of reach and persuasion to one of signal interpretation - the ability to sense, decode, and respond to human behavioural and emotional cues in real time, at scale.

Key findings of this paper include:

- **Agentic commerce is not a future scenario - it is arriving now.** McKinsey projects up to $1 trillion in US B2C retail revenue orchestrated by AI agents by 2030. Indian platforms, including Flipkart, Tata Neu, and Razorpay, are already building the infrastructure that makes agent-mediated transactions possible.
- **Emotional AI works - but only when calibrated correctly.** Empathic AI interfaces measurably improve conversion, loyalty, and satisfaction. But they fail in high-distress contexts, require cultural calibration, and carry significant ethical risk if deployed without consent safeguards.
- **The traditional customer journey is being compressed and reordered.** Brand influence now operates primarily at the point of agent configuration, not at discovery or purchase. Businesses that are not legible to machines - through structured data, clear intent-matching, and strong trust signals - will be systematically deprioritised.
- **The legal and ethical stakes are immediate, not hypothetical.** India's Digital Personal Data Protection Act (2023), the EU AI Act, and evolving consumer protection frameworks place direct obligations on businesses collecting behavioural and emotional data. Agent-initiated transactions raise unresolved questions of commercial liability that require architecture-level responses today.

The paper introduces the **SENSE-SIGNAL-SELL-LEARN framework** as a practical model for building this capability: a continuous loop in which businesses capture behavioural and emotional signals, interpret them through AI, execute calibrated commercial responses, and feed outcomes back to sharpen future performance. It concludes with a five-phase implementation roadmap and a set of strategic imperatives for leadership teams navigating this transition.

The businesses that will define commercial excellence in the coming years are not those that deploy the most aggressive AI. They are those who use it to understand customers more completely - and act on that understanding with both precision and integrity.

Section 2: Introduction - The Signals Have Changed

In the last three decades, commerce migrated from the physical to the digital. In the next three years, it will migrate from the transactional to the perceptual. What made a salesperson great in 2005 - persistence, charm, product knowledge - and what made a digital marketer effective in 2018 - click-through optimisation, retargeting, A/B testing - are becoming insufficient. The game has shifted.

The central argument of this whitepaper is this: selling in 2026 is a discipline of signal interpretation. The most valuable commercial asset a business can hold is no longer inventory, distribution reach, or even brand - it is the ability to sense, decode, and respond to human behavioural and emotional signals in real time, and to deploy autonomous AI agents that act on those signals at scale.

Three forces are converging simultaneously to produce this shift:

- Affective computing is maturing rapidly, enabling machines to infer emotional states from facial micro-expressions, vocal tone, text sentiment, gaze patterns, and physiological proxies.
- Agentic AI is graduating from prototype to infrastructure - systems that do not merely respond to queries but proactively initiate, negotiate, and complete commercial transactions autonomously.
- Consumer behaviour itself is restructuring around AI-mediated decision-making, with buyers increasingly delegating research, comparison, and even purchase execution to AI agents.

> The question for businesses is no longer 'How do we reach the customer?' It is: 'How do we become legible to the agent that reaches the customer on their behalf - and how do we understand the human behind that agent?'

Section 3: Literature Review

Nonverbal Communication as a Sales Signal

Pauser and Wagner (2019) used wearable sociometric badges to objectively capture nonverbal behaviours - kinesics, paralanguage, and proxemics - across live jewellery retail interactions. Their finding: dynamic, expressive nonverbal cues predicted perceived salesperson charisma and, downstream, commercial outcomes. Crucially, these signals operated below conscious awareness, making them invisible to conventional research but measurable by technology. This establishes the scientific precedent - human-to-human behavioural signals are commercially meaningful and machine-readable.

Emotion AI and the Monetisation of Mood

Raamiz and Reshma (2026), drawing on data from 150 respondents, demonstrate, using regression analysis, that AI-driven emotional personalisation significantly improves purchase intention, impulse buying, and brand loyalty. Trust and perceived accuracy of AI systems act as critical mediators - commercial impact depends on whether consumers believe the system reads and uses their signals responsibly. Their study also surfaces the core ethical tension: emotional data is deeply personal, often involuntarily expressed, and frequently collected without meaningful consent.

Emotionally Intelligent Conversational Agents

Rohden and Espartel (2026) find that empathic chatbots are perceived as simultaneously warmer and more competent than non-empathic equivalents - with direct effects on consumer satisfaction and word-of-mouth. Hildebrand and Bergner (2019) add a sharp commercial edge: humanised chatbots with natural turn-taking mechanics nearly doubled upsell conversion compared to standard web interfaces, and even basic personalisation, such as matching a user's name pattern or language style, meaningfully increased persuasion. Meng and Dai (2021) add an important boundary condition: a chatbot's emotional self-disclosure only improves outcomes when paired with genuine emotional support; without it, self-disclosure alone worsens outcomes relative to saying nothing. Zhang et al. (2022) demonstrate that even well-designed emotional AI strategies become ineffective in high-distress contexts. Additionally, gender-differentiated responses to different strategies underscore that calibration to context, not generic warmth, is what works.

Agentic AI and the Architecture of Future Commerce

Kshetri (2025) traces the evolution from predictive AI through generative AI to agentic AI - systems that autonomously initiate and execute multi-step actions. A 2024 Capgemini survey of 1,100 executives found 10% had already implemented AI agents and over 80% planned to within three years. McKinsey (2025) projects up to $1 trillion in US B2C retail revenue orchestrated by AI agents by 2030. Bain (2026) finds that 30-45% of US consumers already use generative AI to research and compare products, with early adopters completing purchases directly from AI platforms.

Section 4: The New Selling Landscape: Commerce Becomes Perceptual

Three Eras, One Direction

Commerce has passed through two discrete technological revolutions and is entering a third. The brick-and-mortar era was relationship-intensive but data-poor - skilled sellers operated on intuition, reading body language and building personal rapport. The e-commerce era was data-rich but relationship-thin - optimisation was algorithmic, personalisation was demographic, and the consumer became a data point navigating a funnel.

The third era - agentic, emotion-aware commerce - synthesises both. It is simultaneously more data-intensive than any previous model and more relationship-like in its interface. AI agents know more about a consumer's preferences, history, and real-time emotional state than any human salesperson could hold in memory; they interact through natural language; and they operate at a speed and scale no human salesforce can match. What varies across each era is who - or what - performs the act of understanding and response.

The Consumer Has Already Changed

This transformation is not purely supply-side. Consumers are changing in response to the same technological forces:

- AI-assisted discovery is becoming the default. McKinsey finds 44% of users who have tried AI-powered search now consider it their primary method. Bain places 30-45% of US consumers as active users of generative AI for product research.
- Purchase delegation is emerging. About half of consumers say they are not yet ready for fully autonomous agent-executed transactions - but that number will shift rapidly as trust accumulates through repeated positive experience.
- Emotional transparency with AI is increasing. Consumers disclose more candidly to AI interfaces than to human representatives, partly because the social risk of judgment feels lower, making AI-captured emotional signals potentially richer than human-gathered insight.
- Context-dependence is supplanting brand loyalty. Consumers are increasingly moment-based rather than brand-based - selecting the best option for their current context rather than defaulting to a familiar name.

Section 5: Agentic Commerce - What It Actually Is & Why It Changes Everything

From List to Transaction: The Consumer Experience of Agentic Commerce

Imagine a scenario where an anniversary gift for one's parents is to be purchased. A query is entered into ChatGPT: 'I want to buy my parents an anniversary gift. They enjoy gardening and are in their sixties. Budget is around ₹3,000.' A year ago, the response would have been presented as a listicle - 'Here are five gift ideas: 1. Garden tools, 2. Seed kits…' While useful, multiple browser tabs would still need to be opened, options compared, delivery timelines checked, and the checkout process completed manually.

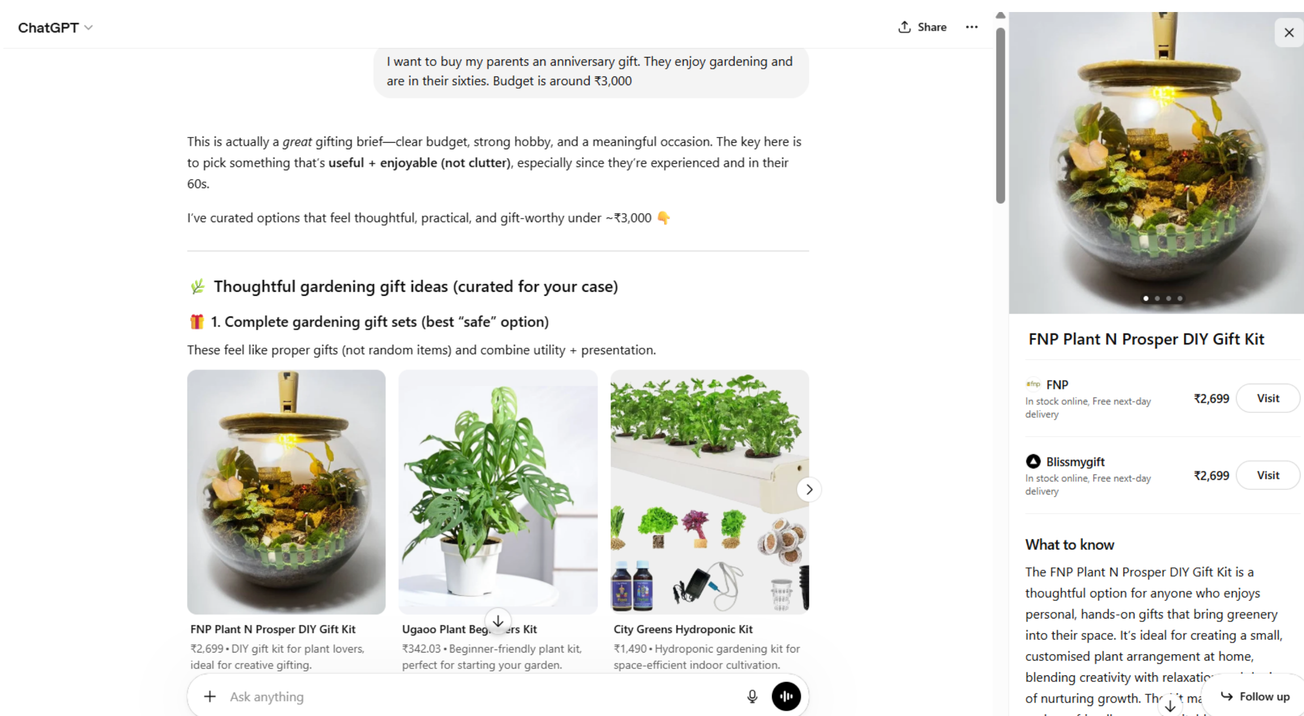

Figure 1: ChatGPT surfacing curated product recommendations from a natural language query - agentic commerce in action

Today, the experience is fundamentally different. User intent is interpreted beyond keywords, and curated, specific product recommendations are presented, complete with images, prices, and estimated delivery timelines. In some cases, a direct tap-to-purchase pathway is enabled, routing the user seamlessly to marketplace checkout. A search is no longer conducted; instead, a commerce transaction is initiated through conversation. This represents agentic commerce in its early, consumer-facing form.

The next evolution lies in a full Agent-to-Agent (A2A) model, where transactions are executed directly between a consumer's AI agent - configured with individual preferences and spending parameters - and a retailer's AI agent, without the user needing to be actively involved in the process.

> Agentic commerce is not a better search engine. It is the replacement of the entire browse-compare-decide-checkout loop with a single intentful conversation - or, at its most autonomous, with no human interaction at all.

The Plumbing Behind the Experience: MCP and A2A Protocols

Briefly understanding the infrastructure that makes this process technically possible is important, since it clarifies why the commercial implications are so significant.

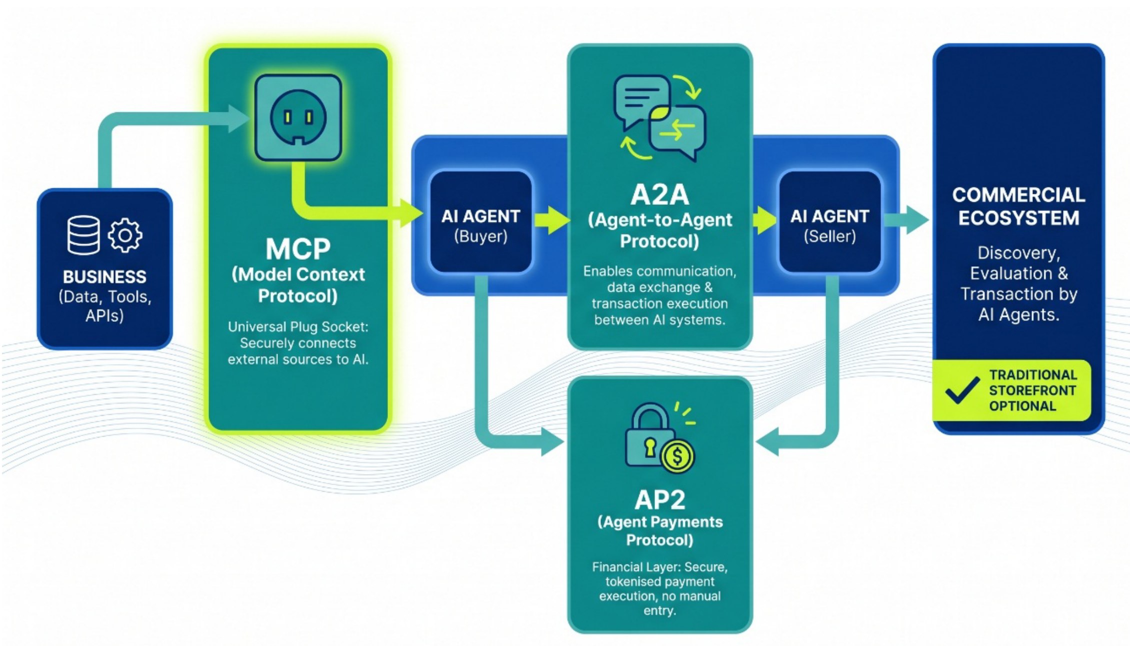

Figure 2: The infrastructure stack powering agentic commerce - MCP, A2A, and AP2 protocols

Anthropic's Model Context Protocol (MCP) is an open standard that allows AI agents to securely connect to external data sources, tools, and APIs. Think of it as a universal plug socket: a business that builds an MCP-compatible product feed or inventory system can instantly make itself accessible to any AI agent that supports the standard. Similarly, the Agent-to-Agent (A2A) Protocol enables different AI systems - a buyer's agent and a seller's agent, for instance - to communicate, exchange structured data, and execute transactions between themselves. The Agent Payments Protocol (AP2) handles the financial layer: secure, tokenised payment execution without requiring the consumer to enter card details manually.

Together, these protocols create an infrastructure layer on which entire commercial ecosystems can be built - ecosystems in which the traditional consumer-facing storefront is, increasingly, optional. A business that is MCP-ready and A2A-compatible can be discovered, evaluated, and transacted with by an AI agent even if no human ever visits its website.

The Agentic Customer Journey: Fundamentally Different by Design

The traditional customer journey - awareness, consideration, intent, purchase, loyalty - assumed a human moving through stages over time, encountering touchpoints managed by the brand. The agentic customer journey looks nothing like this.

In an agentic journey, the consumer makes one primary decision upfront: configuring their agent. They specify their preferences, constraints, values, and trusted brands. From that point, the agent handles everything else - monitoring for relevant products, comparing options across suppliers, applying loyalty credentials, negotiating on price or delivery terms, and executing the transaction within the consumer's approved parameters. The consumer's next touchpoint may be an approval message or a delivery notification.

This compresses what was once a multi-stage, multi-day journey into a near-instantaneous computation. It also radically repositions where brand influence operates. Brands no longer compete primarily for attention at the point of discovery or for persuasion at the point of purchase. They compete for preference at the point of agent configuration - and for the structured data legibility that allows an agent to evaluate and select them at all.

Indian platforms are already beginning to reflect this shift. Flipkart's AI-powered shopping assistant, integrated into its search and discovery layer, interprets natural language queries and surfaces curated product recommendations rather than raw search results - a step toward the intent-driven experience that full agentic commerce will deliver. Tata Neu, Tata Group's super-app, represents a broader architectural bet on the same principle: that consumers will increasingly want a single intelligent interface that navigates across categories - retail, travel, financial services - on their behalf, with accumulated preference data making each subsequent interaction more accurate.

For the Indian market, the payments infrastructure for agentic commerce is already emerging. Razorpay has launched 'Agentic Payments', enabling transactions through conversational AI without traditional checkout flows. Merchants can now accept payments initiated by AI agents, removing a key friction point in A2A commerce. For Indian D2C brands, this means the required checkout infrastructure is already in place - making early adoption a clear competitive advantage.

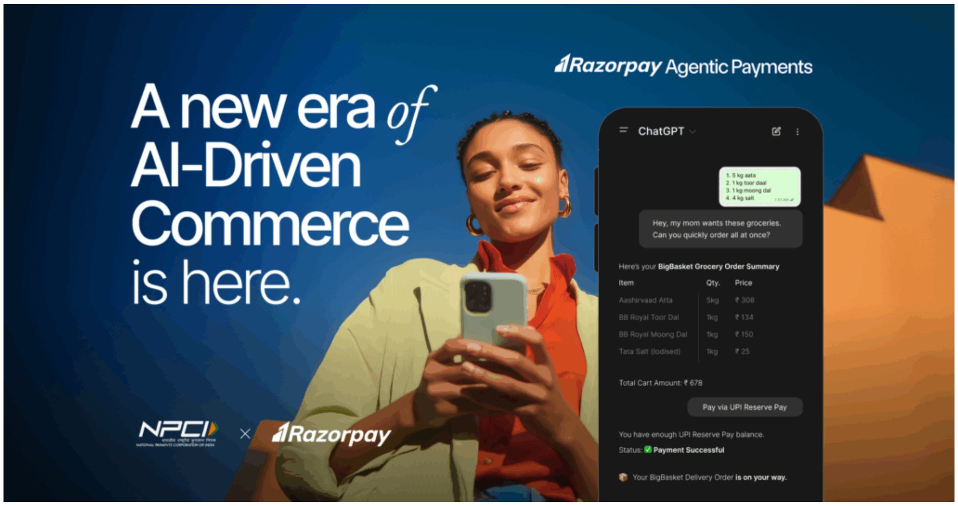

Figure 3: Razorpay Agentic Payments - a grocery order placed and paid via ChatGPT, without a traditional checkout (Source: Razorpay Website)

The Product Listing Problem: From Features to Intent

Here is an underappreciated consequence of agentic commerce that most businesses have not yet confronted: the product listing formats that work for human browsers actively fail for AI agents.

A typical e-commerce product listing is built around features. A pressure cooker listing reads: '5-litre capacity, stainless steel body, ISI certified, three safety valves, compatible with induction and gas.' This is written to reassure a human scanner. An AI agent tasked with finding 'a durable, safe pressure cooker for a family of four that works on induction' needs something different - it needs the listing to be structured around use-cases, context, and intent-matching attributes that allow the agent to confidently evaluate fit.

This is the agentic commerce equivalent of the SEO revolution - and SEO itself is quietly being displaced as a result. In the Google era, businesses invested heavily in Search Engine Optimisation: crafting content to rank on the first page of search results, chasing keywords, building backlinks. That discipline is not disappearing, but it is being supplemented - and in many contexts, superseded - by a new generation of optimisation practices built for AI-mediated discovery.

The emerging landscape of AI-era optimisation goes by several overlapping names - Answer Engine Optimisation (AEO) focuses on structuring content to be directly selected as the answer by AI systems rather than appearing in a ranked list. Generative Engine Optimisation (GEO) targets generative AI platforms - ChatGPT, Perplexity, Gemini - so they cite your brand or product when generating responses. AI Optimisation (AIO) and Large Language Model Optimisation (LLMO) both aim to ensure that LLMs, when generating recommendations or answers, surface your content as a trusted, citable source. Despite their different names, these disciplines share a common logic: AI systems that generate answers prioritise content that is authoritative, clearly structured, intent-matched, and written in natural language that answers real questions - not keyword-dense copy written for a ranking algorithm.

To optimise product listings for agentic discovery, businesses must redefine what 'good' looks like. A strong listing highlights the problem the product solves, the target audience, the context, and why it's the right choice. This replaces the traditional focus on features with structured data, use-case metadata, and natural-language descriptors. Early adopters will be favoured by AI agents, while those who don't adapt risk being deprioritised as their relevance cannot be determined effectively.

Legibility to machines is the new shelf placement - and the optimisation discipline required to achieve it is as significant a shift as SEO was when Google first rewired how the web was structured.

The D2C Liberation Thesis: Could Agents Disintermediate the Marketplaces?

One of the most commercially significant implications of A2A commerce is its potential to fundamentally alter the economics of distribution for direct-to-consumer (D2C) brands.

Today, a D2C brand in India faces a stark commercial reality: genuine scale requires listing on Amazon, Myntra or Flipkart. These platforms charge commissions typically ranging from 15-40%, depending on category, control the customer relationship, own the transaction data, and determine visibility through their own ranking algorithms. The brand manufactures and markets, but the platform owns the customer.

In a mature A2A commerce environment, the dynamics change. A consumer with an AI agent configured to find the best gardening tools within a specified budget and delivery timeframe no longer requires a marketplace as an intermediary. The AI agent can directly discover products from a D2C brand, provided that the brand's data is structured for the agent to access, its checkout infrastructure supports tokenised agent payments, and its reputation signals - such as reviews, return rates, and delivery reliability - are available for the agent to evaluate.

For a D2C brand or SME with strong product quality and a willingness to invest in agent-ready infrastructure, this is a genuine liberation opportunity. Removing the marketplace intermediary could improve contribution margins by 20-35 percentage points, restore ownership of customer data, and enable direct relationship-building with buyers - including emotional memory across interactions. Brands that have built their own digital infrastructure and customer data assets are structurally better positioned to capture this opportunity than brands that have remained entirely marketplace-dependent.

The caveat is significant: this opportunity is most accessible to brands with established trust signals and technical capability. For a new or small D2C brand with no review history and no agent-ready infrastructure, the marketplace's trust halo remains essential. The strategic question is not 'marketplace or no marketplace' but rather: 'At what scale and trust threshold does it become viable to invest in agent-direct infrastructure - and how do we build toward that threshold systematically?'

Section 6: Reading the Room - Microexpressions, Behavioural Cues, and the New Data Layer

Most commercial interactions generate enormous amounts of data that are never captured or used. In a video call, a salesperson might notice a prospect's slight hesitation - a tightening around the eyes that signals disagreement before it becomes verbal. In a chat interface, an unusually long pause before a response. These signals are real, commercially meaningful, and almost entirely invisible to traditional analytics. Microexpressions - fleeting facial movements lasting between 1/25 and 1/5 of a second - are among the most information-dense signals a human can emit, reliably indicating emotional states even when a subject attempts to conceal them (McKnight, O. et al. 2021). The practical question for 2026 is no longer whether machines can read these signals. Systems entering the market can already do so - integrating facial action units, vocal pitch and pace, linguistic hedge density, behavioural scroll and click patterns, and emerging physiological proxies such as heart rate variability via smartwatches.

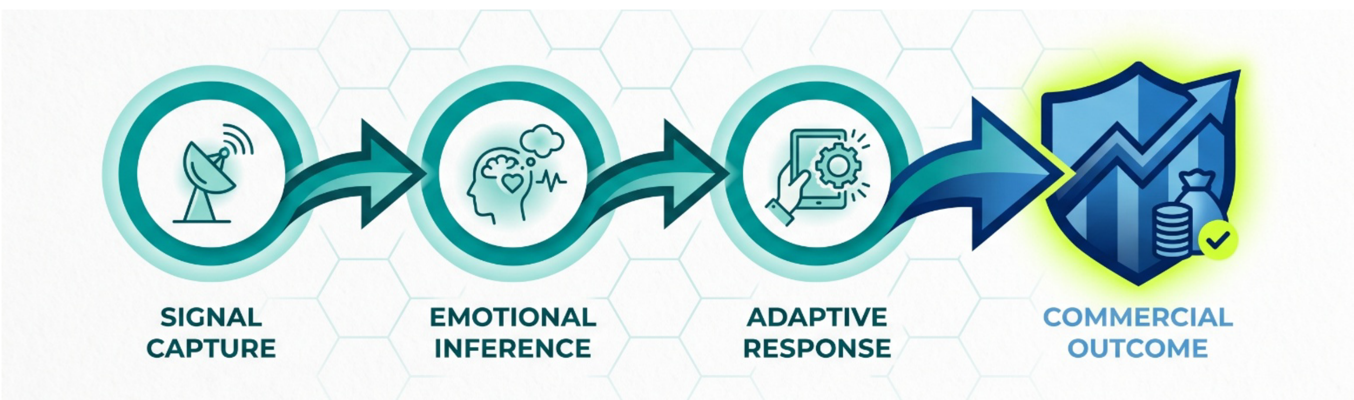

Figure 4: From signal to sale - how emotional AI converts behavioural cues into commercial outcomes

A responsible deployment, however, requires acknowledging limitations. Cultural variation in emotional expression is substantial, individual variability is high, and the commercial incentive to interpret ambiguous signals as buyer readiness creates systematic bias risk. Businesses deploying emotion AI must build interpretive uncertainty into their systems, avoid overconfident inferences, and design for graceful human handoff when signals indicate distress or complexity beyond the system's reliable range.

Section 7: The Interface Layer - Building an Emotionally Intelligent Commercial Presence

The Empathy Architecture

These principles apply to businesses that are independently building their own AI chat interfaces, as well as those deploying third-party conversational AI solutions - whether through white-label platforms, API-integrated chatbot tools, or customised models built on foundation AI. In both cases, the design decisions around emotional intelligence are within the business's control and directly influence commercial outcomes.

Consumers respond to AI interfaces with the same social instincts they apply to human interactions - attributing warmth or coldness, competence or suspicion. These attributions directly determine commercial outcomes. This is not a design nicety. It is load-bearing commercial infrastructure. The architecture of emotional intelligence in an AI commercial interface has four critical dimensions:

- **Tone calibration.** AI interfaces must adapt their linguistic register to match consumer signals. A consumer writing in terse, businesslike language requires a different register than one using casual syntax. Systems that treat all users identically sacrifice measurable persuasion and trust value.
- **Turn-taking discipline.** The rhythm of conversational exchange - when the AI speaks, how long it waits, how frequently it prompts - is structurally important. Interactions with appropriate turn-taking feel more human, build more trust, and generate better commercial outcomes than information-dense, monologue-style responses.
- **Emotional memory.** AI commercial interfaces with persistent memory of interaction history - emotional tone, expressed concerns, past hesitations - replicate the relational continuity of a skilled human seller. Every session that treats the consumer as a new blank slate is a missed commercial and relationship opportunity.
- **Escalation sensitivity.** Knowing not just how to engage emotionally, but when to step back and connect the consumer with human support, is the most important design challenge in emotional AI. Zhang et al. (2022) establish that emotional AI strategies actively fail in high-distress contexts - and in commercial terms, failure in those moments means churn.

Section 8: The Shifting Consumer - Loyalty, Trust, and the Fragmentation of Attention

The End of Static Segmentation

Traditional marketing segmentation grouped consumers by demographic proxies - age, income, location, life stage - and assigned them to personas that were revisited quarterly at best. This model has always been a simplification, but it was a workable one when consumer behaviour changed slowly. In an environment where AI agents act on real-time contextual signals and purchasing decisions can be executed in seconds, static segmentation is not merely imprecise - it is operationally useless.

The consumer of 2026 does not behave like a persona. They behave like a context. The same individual who is highly price-sensitive when buying household consumables may be entirely indifferent to price when purchasing for a gift or a special occasion. The same person who trusts a well-known brand in one category actively seeks lesser-known alternatives in another. AI agents, trained on longitudinal behavioural data, can model this contextual variability with a precision that no static segmentation exercise can replicate - which is precisely why the businesses investing in real-time behavioural intelligence are gaining a structural edge over those still operating from last quarter's persona.

The Trust Paradox

Consumer trust in AI systems is simultaneously high and fragile. Hildebrand and Bergner (2019) find that in a financial advisory context, consumers were three times more likely to accept an objectively incorrect recommendation from a humanised AI than from a standard web interface - even when warned that the advice might be wrong. This suggests well-designed emotional AI can generate trust that exceeds what the underlying system deserves.

Emotional AI can build genuine value by establishing trust in competent recommendations. However, if it merely masks poor quality or price issues, it amounts to manipulation. The line is thin, and regulatory scrutiny is increasing. Businesses that exploit trust asymmetry risk regulatory and reputational fallout as consumer awareness rises. Sustainable success lies in using well-designed AI to provide trustworthy recommendations that truly earn trust.

Attention Has Fragmented - and So Has Brand Power

For decades, brand power was built through repeated exposure across consistent channels - television, print, outdoor, and later social media. The advertising model was predicated on capturing attention at scale and converting it to brand recall that influenced future purchase decisions. Both halves of this model are under pressure.

Attention is fragmenting across an expanding number of AI-mediated surfaces - chatbots, voice interfaces, AI search summaries, personalised agent feeds - where traditional advertising has no obvious place and brand messaging has no guaranteed reach. Brand recall matters less in an agentic purchase where the agent is evaluating structured data attributes rather than emotional brand associations built through repeated creative exposure.

This does not mean brand building is irrelevant. A brand with strong trust signals, clear quality associations, and a reputation that feeds positively into AI evaluation models (through reviews, return rates, and third-party verification) will be systematically preferred by agents. But the mechanism by which brand equity is built and converted into commercial preference is changing. Businesses that continue to measure brand health only through awareness metrics and brand tracking surveys, without also measuring their AI discoverability and agent evaluation performance, are flying partly blind.

Section 9: The SENSE-SIGNAL-SELL-LEARN Framework: Intueri's Strategic Model

The research reviewed in this paper points consistently toward a four-phase model of competitive advantage. We term this the SENSE-SIGNAL-SELL-LEARN framework - not a sequential funnel but a continuous real-time loop operating across every customer interaction.

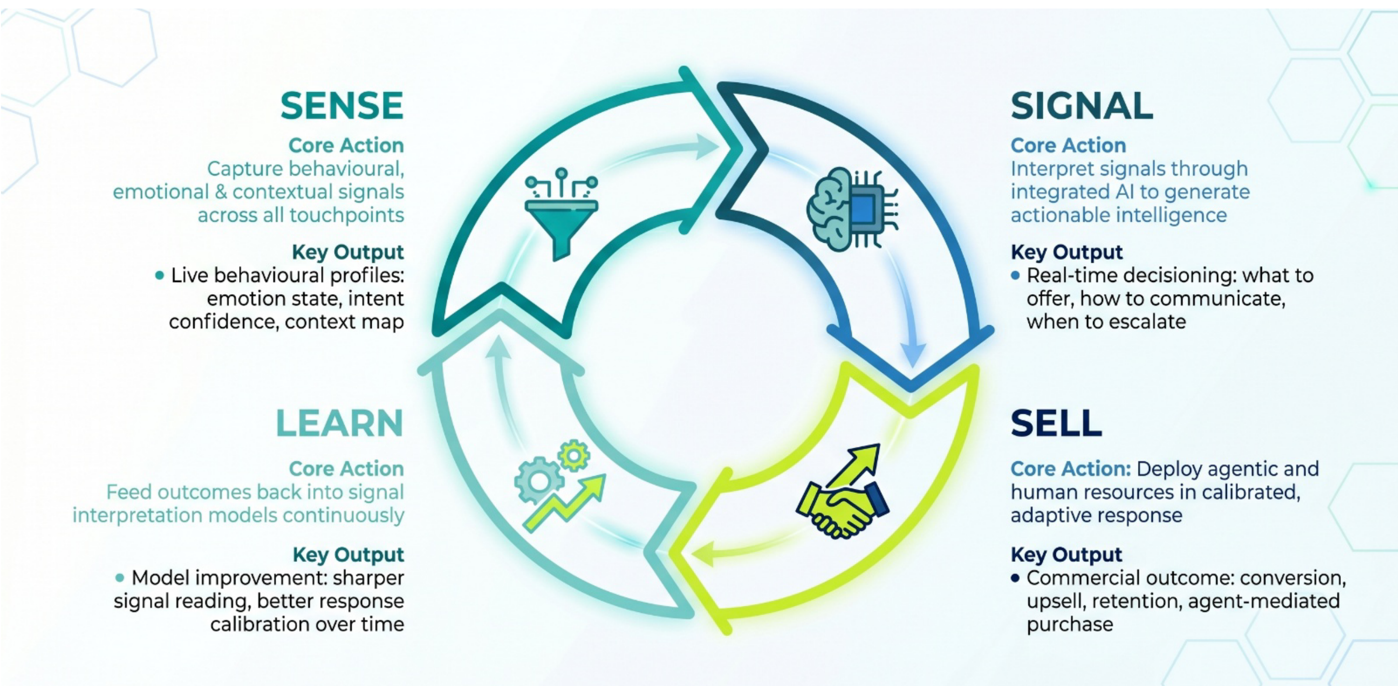

Figure 5: Intueri's SENSE-SIGNAL-SELL-LEARN framework

SENSE: Building the Behavioural Intelligence Layer

The SENSE phase is the data infrastructure layer - all mechanisms through which a business captures signals from current and prospective customers across channels, devices, and interaction types. In 2026, a comprehensive SENSE infrastructure spans digital behavioural data (clicks, scrolls, cart events), conversational signals from AI chat and voice (tone, vocabulary, response latency), video and expression signals from Zoom and Teams (facial action units, gaze patterns), and cross-session emotional continuity that builds a longitudinal emotional profile over time.

The SENSE phase requires investment in data architecture before investment in AI. Many businesses are data-rich in transactional records but signal-poor in behavioural and emotional dimensions. The first strategic task is auditing what signals you currently capture - and identifying the commercially significant gaps.

SIGNAL: The Interpretation Engine

SIGNAL is where captured data becomes actionable intelligence - the analytical and inferential layer that translates raw signals into commercial decisions. A well-designed SIGNAL phase integrates multimodal signals in real time, quantifies rather than suppresses uncertainty, recalibrates personalisation dynamically across a consumer's shifting contexts, and routes to human support when signals indicate distress or complexity beyond the AI's reliable range.

SELL: Matching Execution to Complexity

The SELL phase matches execution mode to the nature of the interaction. Low-complexity, low-stakes transactions - a repeat consumable order, a straightforward product search - are agent-appropriate: fully autonomous, high-speed, low-cost. Moderate-complexity, relationship-significant transactions work best with AI-augmented human selling: the agent prepares, the human executes, armed with the prospect's emotional history and current inferred state. High-complexity, high-stakes transactions - enterprise deals, strategic partnerships - remain human-led with AI as analyst, memory, and scribe.

What the SELL phase eliminates in all modes is uninformed selling. Whether execution is agentic, augmented, or human-led, the commercial interaction is rooted in real signals rather than demographic assumptions.

LEARN: The Compounding Advantage

The LEARN phase separates a one-time AI deployment from a compounding commercial advantage. Every interaction generates outcome data - did the offer land? Did the emotional calibration improve or damage trust? Did the agent-executed purchase satisfy the consumer? This feeds back into the SIGNAL phase, refining interpretation models continuously.

Businesses that build robust LEARN infrastructure accumulate a proprietary behavioural and emotional intelligence asset that becomes more accurate and effective over time. In a competitive environment where multiple players are deploying similar tools, the depth of the LEARN infrastructure becomes the primary differentiator.

Section 10: Strategic Imperatives for Business Leaders

Rethink What 'Knowing Your Customer' Means

Annual surveys and quarterly NPS scores are the demographic equivalent of starlight - they tell you where the customer was, not where they are. In the SENSE-SIGNAL-SELL model, customer knowledge is live, continuous, and behavioural. The strategic question for leadership teams is whether their current customer intelligence infrastructure supports real-time decision-making or is constitutionally oriented toward retrospective analysis.

Redesign the Sales Function Around Orchestration

The most significant threat to sales organisations in the agentic AI era is not displacement - it is irrelevance. Human sellers who continue performing tasks that AI agents now execute more effectively will become a cost overhead. Human sellers who evolve into orchestrators of AI-augmented commercial experiences - reading complex emotional signals, managing high-stakes relationships, designing the experience architecture within which agents operate - will become indispensable.

The 2026 sales leadership agenda must include: identifying which selling activities are agent-appropriate, designing human-AI handoff protocols that preserve trust at escalation points, and retraining or recruiting for a skill profile that is more behaviourally sophisticated, more data-literate, and more strategically oriented than the traditional seller profile.

Establish Your Agent Posture - But Know Your Starting Point

Every business with a commercial presence must determine its strategic posture toward AI agents - Embrace (open your infrastructure to third-party agents and optimise product data for machine ingestion), Build (develop a proprietary agent that makes your brand the starting point for the consumer journey), or Fortify (invest in making the direct channel uniquely valuable through native agentic capabilities and exclusive experiences).

The posture is not binary - most businesses will operate across more than one mode simultaneously. But the crucial caveat is that this framework assumes roughly equivalent access to capital, data, and technical talent. Reality is more asymmetric. Amazon's 'Buy for Me' agent is already live, trained on billions of purchasing interactions, integrated into a checkout infrastructure that most retailers cannot match, and actively routing consumers away from competing storefronts. For the majority of businesses - particularly SMEs - the strategic question is not 'which of these three postures do we choose?' but rather 'how do we remain discoverable and commercially viable in an environment increasingly mediated by platforms that actively disadvantage us?'

Practical responses for SMEs and challenger brands include: investing in agent-readable product data (low cost, high impact), building trust signals - verified reviews, transparent return policies, delivery reliability data - that AI agents can evaluate and weight; finding categories or segments where the major platform agents are weak or biased; and exploring collective infrastructure solutions, such as industry consortia that build shared agent-ready product databases, that individual SMEs could not fund alone.

Build Ethics Into Architecture - and Establish Legal Circuit Breakers

The commercial use of emotional data is on a collision course with regulatory frameworks across multiple jurisdictions. The EU AI Act specifically addresses AI systems processing biometric and affective data. In India, the Digital Personal Data Protection Act (2023) - which imposes consent, purpose limitation, and data minimisation obligations on any entity processing the personal data of Indian residents - applies directly to businesses that capture behavioural and emotional signals from consumers. DPDP compliance is not a future concern for Indian businesses deploying emotion AI: the Act is in force, enforcement is forthcoming, and the obligations it places on data fiduciaries handling sensitive behavioural data are substantive. But beyond data regulation, the emergence of autonomous commercial agents raises a legal liability question this industry has not yet seriously confronted: when an AI agent makes an error, who is responsible?

If a buyer-side AI agent executes a ₹4 lakh purchase that the consumer did not explicitly authorise in that transaction, or if a seller-side negotiation agent inadvertently commits the business to terms that generate a loss, existing consumer protection and contract law frameworks in most jurisdictions are not equipped to handle these cases clearly. 'Human oversight protocols' are not a legal framework. Businesses deploying commercial AI agents need explicit contractual circuit breakers: defined spending limits per agent session, transaction categories requiring explicit human confirmation, audit trails for every agent decision, and clear liability assignment in agent-to-agent commercial agreements. This is not a future concern - it is a present operational risk.

The constructive path is to build ethical and legal constraints into the technical architecture itself: granular consent management, signal capture proportionate to commercial purpose, inference systems that surface uncertainty, and consumer-facing transparency about when and how emotional data is being used. This is not merely a defensive posture - it is a commercial differentiator in a market where consumer trust in AI is simultaneously high and volatile.

> Companies that weaponise emotional AI will win in the short term. Companies that deploy it with genuine care for consumer experience - and genuine legal accountability for agent behaviour - will win in the long term.

Section 11: From Strategy to Implementation

The five phases below are practical next steps - a structured path from signal awareness to full agentic readiness, sequenced so that each phase builds the foundation the next one needs. Whether it is a D2C brand scaling fast, an SME exploring AI for the first time, or an established local business beginning to modernise, the starting point remains the same: the current stage must first be identified, that phase must then be fully completed, and only then can further progress be built upon. The businesses that gain a durable edge are not those that deploy the most AI the fastest - they are those that build the underlying capability in the right order.

Phase 1 - Signal Audit (Months 1-3)

Before deploying any new AI capability, map what behavioural and emotional signals you currently capture and what you are missing. This audit should cover every customer touchpoint - website, app, chat, email, voice, video calls, and in-person interactions. The output is a signal inventory that identifies your richest existing data sources and your most commercially significant gaps.

Phase 2 - High-Impact Intervention Identification (Months 2-4)

Using the signal inventory, identify the three to five moments in the customer journey where emotional signal is most commercially consequential - where a customer's emotional state most strongly predicts conversion, churn, or upsell. These are your priority AI deployment zones. The output is a ranked opportunity map with estimated commercial impact per intervention.

Phase 3 - Incremental Agentic Layering (Months 4-12)

Deploy conversational and agentic AI capabilities progressively, starting with the highest-impact, lowest-risk interventions from Phase 2. Build in human oversight at every escalation point. Run controlled comparisons - identical journeys with and without AI augmentation - to build internal evidence for the commercial impact of specific capabilities. Do not treat deployment as completion.

Phase 4 - Feedback Loop Engineering (Months 6-18)

Establish the LEARN infrastructure: outcome tracking tied to specific AI interventions, model performance monitoring, and a regular cadence of model retraining informed by real commercial data. This phase requires data engineering investment and organisational discipline - the tendency is to treat a deployed AI system as done. The deployment is the beginning of an iterative improvement process, not the end of a project.

Phase 5 - Agent Posture and Infrastructure (Months 12-24)

Based on the evidence accumulated in Phases 3 and 4, make the strategic determinations around agent posture: whether to embrace, build, or fortify your position relative to third-party AI agents. Invest in the technical infrastructure required by your chosen posture - MCP-compatible product data, tokenised payment systems, or proprietary agent development. Establish governance: spending limits per agent session, legal liability frameworks for agent-initiated transactions, data use policies, and ethical review processes for new emotional AI capabilities.

Section 12: Conclusion

The history of commerce, viewed at sufficient altitude, is a history of progressively better understanding of the customer. The general store owner knew every family in town - their preferences, their constraints, their rhythms. Mass marketing replaced that intimacy with reach. E-commerce replaced that reach with data. Agentic, emotion-aware AI commerce is attempting something more ambitious: to recover the depth of individual understanding that characterised the general store, but at the scale and speed of a global digital platform.

Whether it succeeds - and whether its success serves consumers or merely extracts from them - depends on choices that businesses, technologists, and regulators are making right now. The evidence is broadly optimistic: empathic AI interfaces genuinely improve consumer experiences; agentic systems genuinely reduce friction, and emotional personalisation genuinely drives commercial outcomes that benefit both parties. But these benefits are conditional on systems that are honest about their limitations, calibrated to context, legally accountable for their decisions, and designed with consumer trust as a primary rather than residual value.

The businesses that will define commercial excellence in 2026 and beyond are not those that deploy the most aggressive emotional AI, automate the greatest proportion of the sales function, or squeeze every marginal conversion from every behavioural signal. They are those that use AI to understand customers more completely - and use that understanding in service of genuine value creation rather than extraction.

Selling, at its best, has always been about understanding. The tools have changed. The principle has not.

Section 13: References

1. Alam, M., Bhattacharjee, K., & Kumar, A. (2025). The role of artificial intelligence in e-commerce: A comprehensive review of emerging trends, applications, and challenges. International Journal for Research in Applied Science & Engineering Technology, 13(6). https://doi.org/10.22214/ijraset.2025.73212
2. Cheris, A., & Miller, J. (2026, March). Agentic AI commerce: The next retail revolution is here. Bain & Company. https://www.bain.com/insights/agentic-ai-commerce-the-next-retail-revolution-is-here/
3. Davitaia, A. (2025). From scripted replies to intelligent assistance: A study on artificial intelligence chatbots in user support. World Journal of Advanced Engineering Technology and Sciences, 15(3), 1160-1164. https://doi.org/10.30574/wjaets.2025.15.3.1036
4. Hildebrand, C., & Bergner, A. (2019). AI-driven sales automation: Using chatbots to boost sales. NIM Marketing Intelligence Review, 11(2), 36-41. https://doi.org/10.2478/nimmir-2019-0014
5. Kshetri, N. (2025). From predictive and generative to agentic AI: Shaping the future of marketing operations and strategies. Computer, 58(4), 121-128. https://doi.org/10.1109/MC.2025.3530304
6. McKinsey & Company. (2025, October). The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants. McKinsey & Company.
7. McKnight, O., Mahar, C., Paugh, R., Meredith, J., & Moore, J. (2021). Capturing Micro-Expressions on Zoom: A Promising Sales Opportunity. Journal of Marketing Development and Competitiveness, 15(4).
8. Meng, J., & Dai, Y. (2021). Emotional support from AI chatbots: Should a supportive partner self-disclose or not? Journal of Computer-Mediated Communication, 26(4), 207-222. https://doi.org/10.1093/jcmc/zmab005
9. Pauser, S., & Wagner, U. (2019). A wearable sales assistant: capturing dynamic nonverbal communication behaviors using sensor technology. Marketing Letters, 30, 13-25. https://doi.org/10.1007/s11002-019-09483-x
10. Lonegren, U. (2025). AIO, AEO, GEO, AISO, LLMO: The great SEO rebrand explained. Roketto. https://www.helloroketto.com/articles/aio-aeo-geo-aiso-llmo
11. Raamiz & Reshma, S. (2026). The business of emotions: How AI is learning to monetize your mood. International Journal of Advanced Research in Science, Communication and Technology, 6(3), 298-306. https://doi.org/10.48175/IJARSCT-31344
12. Razorpay. (2025). Razorpay agentic payments: India's first AI-powered conversational payment experience. https://razorpay.com/agentic-payments/
13. Redegal. (2025). What are AEO, GEO, and LLMO: The evolution of SEO with the rise of AI. https://www.redegal.com/en/blog/what-is-seo-geo-aeo-llmo-evolution-seo-ai/
14. Rohden, S. F., & Espartel, L. B. (2026). Emotional artificial intelligence: The impact of chatbot empathy and emotional tone on consumer satisfaction and word of mouth. International Journal of Human-Computer Studies, 210, 103764. https://doi.org/10.1016/j.ijhcs.2026.103764
15. Zhang, T., Feng, C., Chen, H., & Xian, J. (2022). Calming the customers by AI: Investigating the role of chatbot acting-cute strategies in soothing negative customer emotions. Electronic Markets, 32, 2277-2292. https://doi.org/10.1007/s12525-022-00596-2