From search to understanding

The new age of product discovery in agentic commerce

/thoughtworks

Design. Engineering. AI.

Executive summary 3

Search and control 5

A new model for product discovery 8

Modes of discovery 10

Evolving technical foundation 12

Managing risk and observability in agentic commerce 15

Charting a path to agentic commerce 17

Where to begin 18

Why Thoughtworks? 21

Authors 21

The role of agents 23

Executive summary

For decades, winning in digital commerce meant looking for an edge in a familiar model: users search, filter, compare and purchase. While interfaces have become faster and more personalized, the underlying journey has remained largely unchanged.

But AI is changing that. Indeed, the emergence of agentic AI introduces a fundamentally different paradigm, one where systems don't simply respond to queries, but interpret a consumer's intent, adapt in real time and actively guide them toward a purchase. Discovery is no longer a linear path from search to checkout. It is becoming fluid, shaped continuously by context, behavior and interaction.

This shift has significant implications for where competitive edge is created — and where control resides.

If early indications are right, AI-driven discovery platforms are set to become the dominant entry points for commerce, influencing what consumers see and buy. In this scenario, traditional retail platforms risk being reduced to transactional endpoints, while discovery — and with it, customer influence — moves elsewhere.

How to capitalize on this opportunity? Retailers retain critical structural advantages: rich first-party data, deep domain knowledge, established trust and control over the end-to-end customer experience. These strengths put them in pole position — provided they evolve their discovery capabilities beyond static search and navigation.

In this paper, we'll show how to introduce adaptive, intent-aware experiences within familiar interfaces, which enable retailers to deliver more relevant, responsive and engaging discovery — while maintaining ownership of the customer relationship. Achieving this requires new foundations: structured knowledge, real-time decisioning and composable systems that support continuous interaction between user and platform.

We'll provide a blueprint for how that transformation can be approached in practice, exploring:

- How retailers can unlock greater value from their current discovery stack by strengthening data.
- How to refine experience layers and introduce targeted AI capabilities.
- The emerging architecture of agent-driven commerce.
- And the capabilities required to remain competitive in a landscape where discovery increasingly happens both within and beyond retailer-owned environments.

The retailers that act now by building on what already works while introducing intelligent, adaptive layers, will define the next generation of product discovery.

Search and control

For years, e-commerce has optimized around a stable assumption: that product discovery happens within retailer-controlled environments. Search engines, category navigation, filters and recommendations have all evolved to make this journey faster and more efficient. But despite these improvements, the structure of discovery has remained fundamentally the same — users arrive, express intent and navigate toward a product within a defined system.

The emergence of large language models has changed search. Discovery is taking place outside of your retail platforms. These environments are not just directing traffic. They are shaping how intent is formed, refined, and acted upon.

The emergence of AI-driven platforms, particularly those leveraging agentic AI, is enabling experiences that feel far more human, contextual and intuitive. These platforms embed commerce into those digital environments where users already spend time. And it's changing how we shop. In Nov 2025, Black Friday online sales hit $11.8 billion, up 9.1% from previous year. Early signals are already showing levels of engagement that should command every retailer's attention.

What changes with agentic AI is not just the interface, but the locus of decision-making.

Instead of users actively navigating options, intelligent systems are beginning to interpret needs, filter choices and guide decisions before your platform's ever reached. The point at which customer decisions are influenced is moving upstream.

For retailers, this raises a critical question: if discovery happens elsewhere, what remains under your control?

A high-stakes shift with an open outcome

Some perspectives suggest a future where retailer platforms are reduced to fulfillment layers; handling transactions while discovery is owned by external AI ecosystems operating at internet scale.

In this scenario, AI aggregators become the primary interface for commerce, shaping demand, controlling visibility, and determining which products surface and why. Retailers compete within these systems rather than through their own.

For organizations built on curated digital journeys and direct customer relationships, this would represent a profound shift in power. According to McKinsey analysis, agentic commerce, where AI agents autonomously assist discovery and purchase, could unlock retail revenues of $1 trillion in the US by 2030 and $3–$5 trillion globally.

However, this outcome is far from inevitable. Retailers still possess significant structural advantages: access to rich first-party data, deep understanding of their categories, established trust with customers and control over end-to-end experiences. In many domains, particularly complex or high-consideration purchases, replicating this depth of understanding within generalized AI platforms remains a significant challenge.

The reality is more nuanced than a simple transfer of control. The competitive landscape is being reshaped, not decided. What matters now is how retailers respond.

Those that treat AI as an overlay, adding conversational interfaces or incremental search improvements, risk ceding ground to more adaptive systems. In contrast, those that rethink discovery as an intent-driven, continuously learning process have an opportunity to reclaim influence, even as the ecosystem evolves.

The emergence of fluid commerce

One of the fundamental shifts that agentic discovery introduces is not just how systems process information, but how users behave as a result.

Consider a shopper preparing for a weekend hiking trip. In a traditional journey, the buyer might begin with a search like "hiking shoes," apply filters for size, brand and price. Maybe they check out some reviews, see what's recommended on TikTok, compare options, and gradually narrow down their choice. The burden of refining intent — figuring out what matters, what fits their needs and what to prioritize — sits largely with the user.

In an agent-driven experience, that starting point looks similar but what follows is radically different. The same shopper might begin with a broad query, or even a loosely formed thought: "I need some comfortable footwear for a two-day hike." Instead of returning a static list of products, the system begins to interpret intent in context. It recognizes signals — terrain, duration, past preferences, weather conditions — and responds by shaping the journey.

In this model, discovery moves beyond traditional product-to-product navigation toward a richer understanding of broader themes, goals, and motivations driving the user's current context. Agentic systems continuously interpret intent, adapt the experience, and orchestrate outcomes and shift the role of the platform from a passive catalog to an active participant in helping users achieve their objectives.

This is ‘fluid commerce’, where discovery becomes less about navigating a map and more about co-creating a digital journey as it happens.

A defining characteristic of fluid commerce is its ability to adapt within familiar user journeys, while quietly infusing them with the benefits of agentic intelligence. It doesn’t force a shift, it evolves the experience in place.

At the same time, it can seamlessly extend beyond these traditional flows, guiding users into newer paradigms such as conversational commerce or latent-feature driven discovery, doing so progressively, based on user comfort.

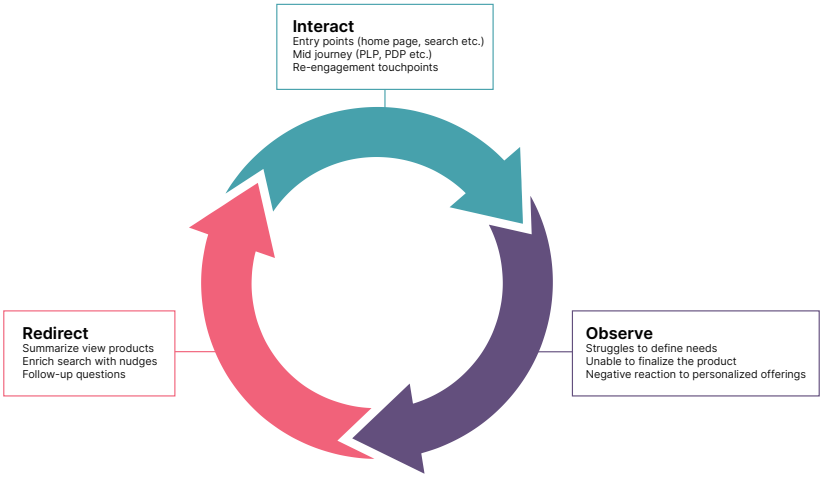

A new model for product discovery

Fluid commerce requires a design thinking built around dynamic states, where users move seamlessly across modes as their intent, confidence, and context evolve.

Discovery becomes adaptive rather than sequential. It's built on a simple three-step loop:

1. **Interact.** The user engages with the system, via search bars, micro-prompts, conversational interfaces, visual nudges, comparisons, bundles and recommendations.
2. **Observe.** The system continuously learns from user behavior and signals. Signals may include edits, pauses, scroll depth, clicks, repeated rephrasing and dwell time, along with context like personalization and session intent.
3. **Redirect.** The experience responds in real time, either by seamlessly redirecting or by subtly nudging the user toward better outcomes. The system may escalate into conversation, simplify the interface, surface additional context, alter the ranking, generate thematic bundles or route refined intent to the retrieval engine.

This loop underpins how modern product discovery works. It treats discovery not as a linear funnel but as an ongoing exchange between the user, the system and the surrounding context.

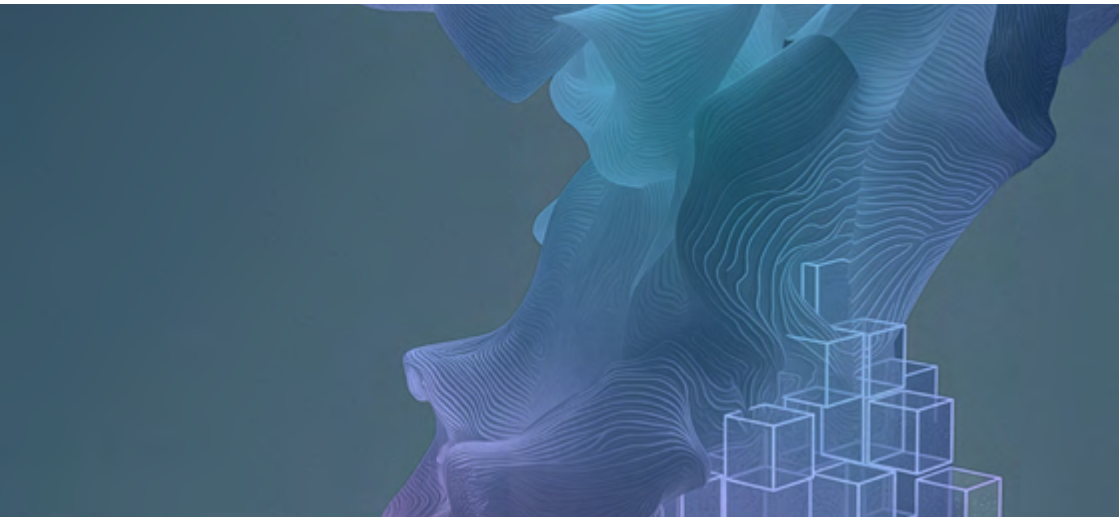

Modes of discovery

Tomorrow's commerce systems address needs rather than just products, interpreting keywords, intent, or feelings. Whether a shopper searches "Cola" or thinks "I need something refreshing," the system translates these moments into optimized choices using conversational, context-aware interactions.

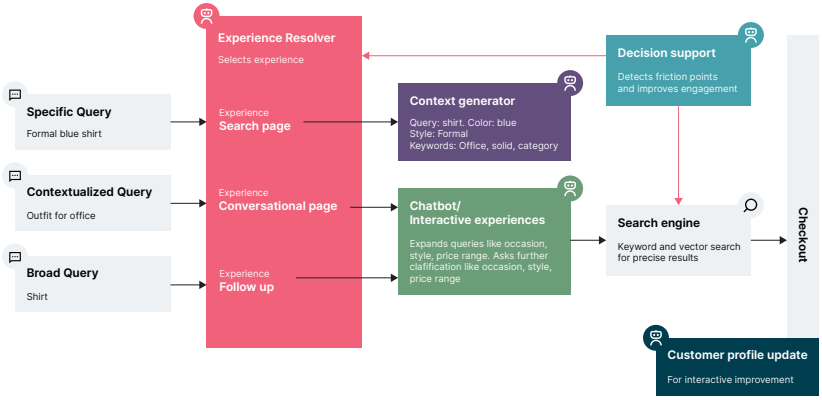

The sections that follow frame these journeys not as fixed paths, but as flexible modes that can be adapted across categories and contexts, allowing commerce to align more closely with how people actually discover and decide. Modern e-commerce no longer needs to choose between static page flows and text-only chat interfaces. It can seamlessly integrate both.

1. Clear intent: Direct results

When a consumer expresses a precise query, for example, "blue formal shirt", the system goes straight to results instead of asking follow-up questions. It retrieves and ranks products based on the user's intent.

The experience remains fast, familiar and predictable. In high-intent scenarios, the system reduces friction by responding directly to the request.

2. Ambiguous intent: Clarification prompt

When a query is incomplete or broad, for example, shirts, the system does not move into a full conversation. Instead, it offers a short prompt to help clarify intent.

The system may ask the user to specify preferences such as:

- Color preferences
- Formal or casual
- Sleeve length
- Budget range
- Brand preferences

Once clarified, the experience returns to the search results page, now aligned with the user's intent.

3. Open-ended intent: Guided discovery

Queries such as "I'm going to my daughter's graduation and need help picking a dress" shift the experience into a more guided mode.

This can appear as:

- A chat sidebar
- A guided panel
- An assistant overlay

In this mode, the system:

- Asks for relevant context.
- Narrows down user goals.
- Suggests products within the flow.
- Highlights trade-offs where needed.
- Builds a clearer shopping request over time.

The conversation then transitions to the search results page with a well-defined query.

Evolving technical foundation

To make fluid commerce work in practice, both on retailer-native channel and third-party AI platforms, retailers must move beyond microservices architectures made up of many discrete components. Instead they need to operate as a coordinated system built around an interact-observe-redirect loop, where each buyer action dictates what happens next.

This loop works only when supported by three foundational components:

- A **knowledge fabric**, which ensures product data is structured, connected, and interpretable beyond simple attributes.
- A **factual backbone**, which grounds every response, recommendation and comparison in accurate, catalog-aligned data.
- An **adaptive experience layer**, supported by specialized task agents that decide how and when to respond within the journey.

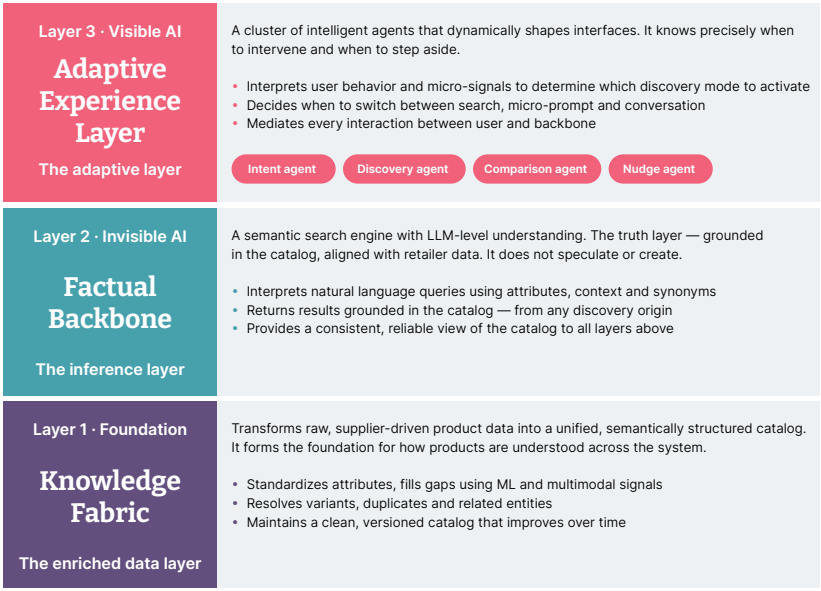

Together, these components allow the system to do something traditional architectures struggle with: remain responsive in the moment without losing alignment with the underlying catalog. What the user sees may change dynamically, but what it's based on remains consistent and trustworthy.

Personalization runs across this stack as both long-term and short-term memory. Applications such as search, home and merchandising continue to draw on historical customer data such as purchases, browsing behavior and profiles, using AI tools to add insights. This forms the system's long-term memory: what it already knows about the user before the session begins.

The funnel, however, becomes reactive. As users interact by refining queries, pausing on products, comparing options, the system builds a short-term understanding of their immediate intent. This session-level memory continuously updates what the system believes the user is trying to achieve, without requiring them to restate it explicitly.

The result is that every session starts informed, and then adapts in real time. This is the adaptive intelligence of the system. It's what makes the experience feel responsive and contextual; not by replacing search or navigation, but by knowing when to guide, when to simplify, and when to stay out of the way.

Agents join the stack

To execute this adaptivity at scale, the system relies on small, specialized task agents — each performing a narrow, well-defined job:

- Generating comparison summaries.
- Producing on-the-fly bundles or lookbooks.
- Crafting personalized tags or micro-nudges.
- Validating promo logic.
- Extracting or enriching product attributes.
- Translating user expressions into structured queries.

These agents do not operate as a monolith, deterministic components with rigid communication protocols.

They are modular, composable components that the experience layer can call, combine, or dismiss based on the user's real-time journey whether originating within retailer platforms or from external discovery environments.

Woven through all three is a shared agent fabric of specialized task agents that continuously validate against truth, enrich context, and dynamically shape the experience with reasoning, comparisons, bundles, and next-best actions.

With these three layers in place, orchestrated by atomized task agents, the architecture delivers the intelligence, adaptability and personalization required for modern retail discovery.

Please see Agents Appendix on Page 23-26

But intelligence alone isn't enough. For this system to drive business value, it must be governed with precision and engineered for scale. Two cross-layer priorities complete the picture: enterprise control and operational performance.

Managing risk and observability in agentic commerce

Even with a well-orchestrated, multi-agent discovery architecture, risk doesn't disappear. In fact, a new breed of risks comes into play — ones that are harder to detect precisely because they are buried inside automated decision chains. Retailers need to treat evaluation and observability as first-class capabilities: continuously measuring agent behavior, monitoring for drift and degradation, and maintaining the human oversight necessary to catch what automation cannot. AI-driven product discovery systems are inherently non-deterministic.

That flexibility is what enables better experiences, but it also introduces new ways for systems to drift, misfire or degrade quietly. These risks don't always show up in traditional KPIs. Conversion and relevance may still appear strong, even as discovery quality declines. AI evals and observability help teams see what is happening and understand why.

Here is a set of practical risks that surface directly in real discovery journeys and specific metrics to build observability around.

Risk 1: Confident output, incorrect outcome

Generative systems can produce fluent, confident responses that are subtly misaligned with the catalog or unsupported by data.

- Shows up as: incorrect attributes, misleading comparisons, explanations that go beyond what the catalog supports.
- Diagnostic signals: grounding to catalog data, retrieval precision, hallucination or unsupported-claim rates.

Risk 2: Help that arrives at the wrong moment

Adaptive systems constantly decide how to engage. When those decisions are mistimed, guidance becomes friction rather than value.

- Shows up as: premature prompts, unnecessary conversational turns, drop-offs after intervention.
- Diagnostic signals: intent classification accuracy, mode-selection precision, over- or under-intervention signals.

Risk 3: Correct agents, degraded journey

Even when individual agents behave correctly, poor coordination can slow or fragment the overall discovery flow.

- Shows up as: repeated refinements, mode switching without progress, longer time-to-product.
- Diagnostic signals: mode transitions, hesitation loops, journey-level time-to-clarity and assisted conversion paths.

Charting a path to agentic commerce

Traditional product discovery KPIs such as relevance, click-through rate and conversion will continue to be essential. They provide continuity, comparability and a clear view of commercial performance.

However, as discovery experiences become more adaptive and interactive, these metrics alone are no longer sufficient to explain how users arrive at outcomes or why they sometimes don't. Hence, we recommend augmenting existing KPIs that double-click into users' state of mind as they're being nudged, guided, and assisted through a more one-to-one shopping journey.

When combined with established KPIs, these metrics offer deeper insight into the customer journey and help guide the evolution of multi-agent e-commerce flows.

KPI	Indicates	How to measure
Intent clarity progression	Whether the user's intent becomes clearer over the session.	Change in query specificity, reduced refinements or fewer conflicting signals over time.
Meaningful engagement rate	If users are genuinely exploring, not just clicking.	Sessions with multiple product views, comparisons or guided interactions.
Assisted conversion share	How often support mechanisms contribute to conversion.	Percentage of conversions that involved prompts, recommendations or conversational steps.
Early exit with context	Why do users leave before converting?	Exit events tagged with last interaction type (search, browse, guidance).

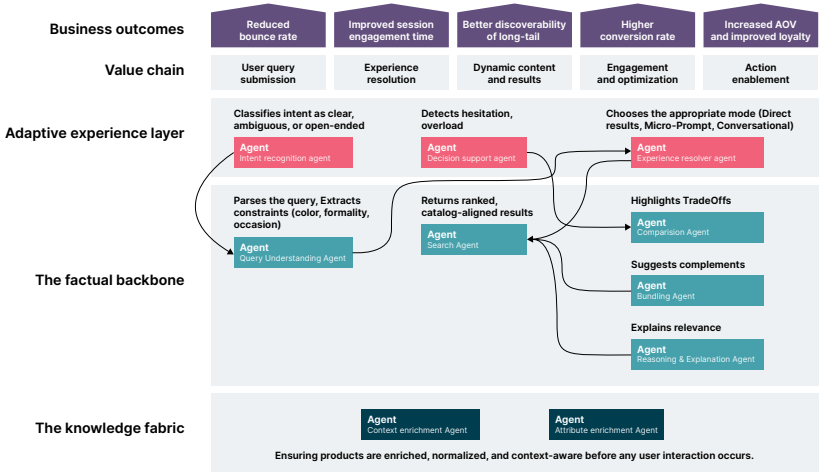

Instead of counting events, these KPIs treat discovery as a sequence of evolving probabilities, showing how confidence, risk and influence change as users interact.

Where to begin

The effective path forward combines strong invisible foundations with thin, targeted moments of visible AI — each reinforcing the other. The winning approach is not maximum AI visibility.

Invisible AI focuses on strengthening the underlying capabilities that enable intelligent discovery, such as improving product data quality, enabling dynamically composed interfaces, establishing orchestration layers, and capturing meaningful intent signals. These capabilities operate behind the scenes but ultimately determine how effectively future experiences can scale.

Visible AI, in contrast, introduces customer-facing intelligence such as guided discovery, conversational assistance, and adaptive workflows that directly improve how customers explore and decide.

Rather than pursuing large-scale horizontal transformations that delay outcomes, the most effective path forward is to introduce focused, job-driven capabilities that enhance existing discovery journeys while building the foundations for future intelligence. The visible portion builds momentum, stakeholder confidence, and measurable business value, while the larger invisible foundation ensures scalability, adaptability, and long-term sustainability.

These capabilities can also co-exist with legacy systems, operating as additive layers that improve decision-making without disrupting core transaction flows. By starting with one well-defined user need, retailers can validate new discovery models in production while continuing to rely on existing systems for browsing, product detail, and checkout.

Begin with a specific moment in the customer journey where AI can meaningfully close a gap

Identify a customer task where better guidance improves confidence and reduces friction: for example, selecting the right tools for a project, replenishing frequently purchased items, or assembling compatible product bundles. Begin with a specific moment in the customer journey where AI can meaningfully close a gap.

Prepare minimum viable foundations that the job requires

Before building, assess where you stand. Evaluate current capabilities across experience flexibility, data quality, and decisioning infrastructure — then identify the minimal set of product attributes, behavioral signals, and orchestration components needed to support the target customer problem, rather than attempting a broad data transformation program.

Deliver one guided capability

Introduce visible enhancements within existing journeys that redirect users to the right experience, such as chat, guided conversations, contextual nudges or intelligent prompts.

Make it reusable

Expose the capability as a reusable service, one that works across channels today and can be progressively extended as new customer needs emerge. Let observable performance data drive what gets built next.

The architecture outlined in this paper, the knowledge fabric, the factual backbone, an experience layer, and specialized task agents, provides the foundation for this new era of intelligent discovery within retailer platforms, while also enabling readiness for participation in external, AI-driven discovery ecosystems.

Together, these layers move search beyond a fixed interface and support a system that can respond to intent wherever it begins, whether on the retailer's site or outside it. This helps keep products discoverable and consistent across AI agents, extensions and off-platform journeys.

Cross-cutting capabilities such as enterprise governance, AI evals and observability ensure that this system remains controlled, reliable and scalable. The result is a discovery system that is responsive and personalized, while still aligned with business rules and operational requirements.

The mandate for retailers is clear: discovery must evolve from pages and filters to adaptive, agent-powered journeys, experiences that understand customers. The retailers that invest now in this architecture will own the next generation of discovery: faster, more intuitive, more grounded, and ultimately more aligned with how consumers actually shop in an AI-mediated world.

Why Thoughtworks?

We empower our retail and e-commerce such as Etsy, Falabella and Reece Group to forge strong customer connections, streamline operations and outperform their competitors. Our end-to-end retail expertise spans large-scale commerce and marketplace transformation, omnichannel merchandising, agile supply chain and logistics, customer loyalty, media, and emerging technologies. We build scalable and evolvable solutions, which drive value and growth for your business.

Authors

Vanitha Kumar

Vanitha is a Market Technology Director and Technical Partner at Thoughtworks, specializing in platform modernization for Retail, CPG, Food Services and Hospitality industries. With over 25 years of experience spanning engineering, product, and organizational change, she has led the delivery of cloud-native platforms, digital payment solutions, and enterprise automation frameworks — and is now focused on helping organizations design and deliver agentic AI experiences responsibly.

Vanitha is a recognized thought leader in advancing AI adoption that balances innovation with consumer trust and ethical responsibility, ensuring AI-enabled commerce systems create meaningful value without compromising transparency or equity. She is known across business domains as a trusted advisor who translates complexity into clarity, bridges technical and business strategy, and advocates for humans building and scaling real systems every day.

Madhubala Phani

Madhubala is a seasoned Product Strategy leader who partners closely with CXOs and leaders across enterprise organizations to shape strategies that transform customer and associate experiences. She specializes in evolving legacy ecosystems into scalable, platform-driven architectures that enable intelligent digital experiences, streamlined operations, and measurable business outcomes.

Her current focus is on designing agentic AI–powered experiences and building intent-driven workflows that guide users, automate decisions, and continuously improve operations. With deep expertise in platform product thinking and enterprise delivery, she helps organizations transition from static systems to adaptive, intent-driven ecosystems that learn, optimize, and scale with evolving customer needs.

Kumar Gaurav

Gaurav is a Search and AI engineer with over a decade of experience improving how people discover information and products across complex platforms. He has worked on large-scale systems that enhance relevance, personalization, and overall user experience, with a focus on practical, real-world impact.

His approach emphasizes aligning technology with business goals and continuously refining how systems understand and respond to user intent. Currently, he is exploring how newer AI approaches can make search more intuitive, adaptive, and useful over time.

Reviewers: Sincere thanks to
Chris Ford, Sanjeev Seshagiri Athreya and Jay Amburkar.

Appendix

Knowledge fabric agents

Convert raw product data into clean, enriched, semantically consistent inputs.

	Attribute enrichment agent	**Context enrichment agent**
What data do they need?	Raw product feeds, missing attributes, unstructured text, images, specification inconsistencies.	Product attributes, usage patterns, search queries, engagement signals.
What decisions do they make?	Which attributes to infer, normalize, or flag based on confidence thresholds.	Relevant contexts (e.g., formal, outdoor, humid climate, compatibility).
What actions do they perform?	Extracts attributes, standardizes units, enriches metadata, updates catalog records.	Tags products with contextual metadata used in discovery and reasoning.
How do they self-improve?	Improves extraction accuracy from corrections, approvals, and downstream retrieval performance.	Refines context mapping based on click-through, conversions, and feedback loops.

Factual backbone agents

Deliver accurate, catalog-grounded results every time.

	Query understanding and rewrite agent	Search agent
What data do they need?	User queries, session history, domain vocabulary.	Structured query, catalog embeddings, availability and business rules.
What decisions do they make?	Intent interpretation, synonym expansion, constraint inference.	Retrieval strategy (semantic vs hybrid), ranking weight.
What actions do they perform?	Rewrites query into structured parameters for retrieval.	Returns ranked, catalog-aligned results.
How do they self-improve?	Improves rewrite accuracy from relevance outcomes and failed searches.	Optimizes ranking models from engagement, conversions and relevance metrics.

Adaptive experience layer agents

Dynamically guide users by switching discovery modes in real time.

	Intent recognition agent	Experience resolver agent	Decision support agent
What data do they need?	Query patterns, rewrites, scroll behavior, filter usage.	Intent state, behavior signals, latency constraints.	Hesitation signals, dwell time, repeated comparisons.
What decisions do they make?	Intent state (clear, ambiguous, open-ended) and confidence score.	Stay in results, trigger micro-prompt, or enter conversation.	Whether to intervene and what form of support is needed.
What actions do they perform?	Signals the appropriate experience mode.	Orchestrates UI mode switching and graceful fallbacks.	Surfaces guidance (comparisons, narrowing suggestions, explanations).
How do they self-improve?	Refines classification using session outcomes and correction signals.	Improves transition timing based on abandonment and engagement patterns.	Tunes intervention thresholds based on effectiveness and user response.

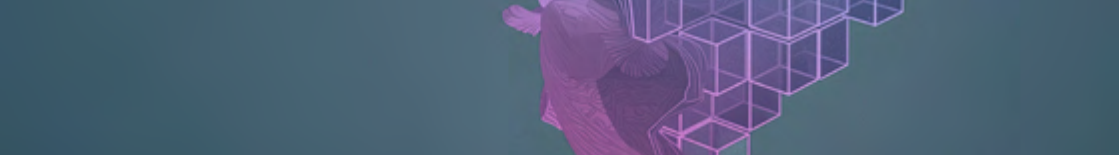

Specialized task agents

Add targeted intelligence — comparisons, bundles, and explanations — exactly when needed.

	Comparison agent	**Reasoning and explanation agent**
What data do they need?	Candidate products, attributes, user intent.	Ranking signals, attributes, user preferences.
What decisions do they make?	Key differentiators to highlight.	Which rationale best supports the recommendation.
What actions do they perform?	Produces side-by-side comparisons tied to catalog data.	Displays concise, grounded explanations (“best match”, “better for hot weather”).
How do they self-improve?	Improves comparison relevance from user engagement and selections.	Improves explanation clarity from user trust signals and feedback.

We are a global technology consultancy that delivers extraordinary impact by blending design, engineering and AI expertise.

For over 30 years, our culture of innovation and technological excellence has helped clients strengthen their enterprise systems, scale with agility and create seamless digital experiences.

We're dedicated to solving our clients' most critical challenges, combining AI and human ingenuity to turn their ambitious ideas into reality.

thoughtworks.com

/thoughtworks

Design. Engineering. AI.

Last updated June 2026
© Thoughtworks, Inc. All Rights Reserved.