



Faster, smarter, stronger

Bringing the transformative
power of generative AI to retail

 **thoughtworks**

Strategy. Design. Engineering.

Introduction	3
Generative AI for merchandizing	5
Generative AI for marketing	11
Generative AI for customer service agents	17
How you can make the most of the generative AI opportunity in retail	24
Find the right use cases and seize the retail generative AI opportunity today	27
Authors	30

Introduction



Over the last few decades, we've seen numerous digital and data-driven shifts that have transformed how consumers buy, retailers sell, and marketers reach their audiences.




Today we stand on the precipice of one of the most significant evolutions in recent memory: the growth and proliferation of generative AI (genAI).

Intelligence from Bloomberg forecasts the generative AI market to grow at an astounding compound annual growth rate (CAGR) of 42% over the next 10 years. It's set to become a \$1.3 trillion market by 2032, and transform virtually every industry in the process — and retail is certainly no exception.

Current predictions indicate that for organizations in the retail and packaged goods industry, generative AI will increase productivity by between 1.2 and 2 percent, creating between \$400 and \$660 billion in value.

In retail, generative AI is creating incredibly valuable opportunities across the value chain — from new ways to dynamically optimize product ranges and store layouts, to enabling unprecedented levels of personalization in shopper experiences.

But some functions and roles are likely to be more heavily affected by the introduction and growth of generative AI than others:

-  **Merchandizers** will gain a deeper understanding of customer needs across diverse geographies and develop dynamic merchandizing strategies that continuously drive growth and improve shopper experiences.
-  **Retail marketing teams** will see huge shifts in how they generate, distribute, and personalize their content — enabling a transformation in customer experiences and how retail brands, and the products they sell, are marketed.
-  **Customer service agents** will engage with customers in entirely new ways and gain valuable time to focus on resolving the most complex and demanding customer cases.

In this white paper, we'll take a detailed look at how generative AI can augment and empower the people in those roles. And we'll explore how you can apply AI capabilities in those areas to unlock huge efficiency gains, competitive advantage, and value creation opportunities.

Generative AI for merchandizing



In catering to today's channel-agnostic customers, merchandizing teams have their work cut out for them. They have to continuously manage and optimize product assortment and placement across a huge range of channels, in ways that both enable consistent experiences and make the most of the unique landscape of each channel.

To do that effectively, they must extract insights from vast quantities of data. Sales data, customer preferences and trends, and supplier and procurement insights all have a role to play in enabling rapid, high-value merchandizing decision-making. But traditionally, the merchandizing function has been monolithic, which has made applying and learning from that data at scale challenging.

AI and Large Language Models (LLMs) have changed that. They've made it possible to process and learn from huge retail data sets at speed. Now, generative AI is taking those capabilities a step further, by automatically translating AI-generated outputs into meaningful, value-driving merchandizing actions.

From generating role-specific data summaries that teams can use to inform purchasing decisions and negotiations, to enabling deep analysis of product line performance, generative AI is making it possible for merchandizers to finally harness the full value of the data available to them.

Key moments in the life of a generative AI-empowered merchandizer

Moment 1: Designing products

Product design is a long and complex process. Teams go through numerous design iterations to reach a final version that can meet customer and market demands — each carrying a long lead time as designers iterate on product concepts.

But with generative AI, the design process becomes much faster. Rapidly generating design options for products can significantly cut time to market.

Imagine a scenario where generative AI algorithms can analyze vast amounts of customer data, including purchase history, social media interactions, and even product reviews, to generate innovative product designs that perfectly align with consumer preferences.

Teams can rapidly evaluate a huge range of design options and lock in on a concept that works quickly. Then, human design teams can focus their efforts on making that concept as strong as possible, rather than creating concepts that will ultimately go unused.

Moment 2: Optimizing product assortment and merchandizing

Manually sifting through data to gain product range and assortment insights is a major drain on merchandizing teams' time. Pre-configured analytics dashboards have their own biases and are unable to keep up with ever-exploding channels and rapidly changing consumer behavior.

But with the help of generative AI, teams can use natural language querying to access the information that's most relevant to them, and spend their time acting on it, rather than searching for it.

Generative AI can derive insights from sales, customers, and other metadata across channels — and correlate them with core product attributes to derive patterns and generate optimal merchandizing and product assortment strategies across physical and online storefronts.

Equipped with a continuous stream of AI-derived insights, merchandizers can validate trends fast, and make rapid, frequent decisions to optimize product assortment, positioning, and range composition.



Moment 3: Negotiating with vendors

In negotiations with vendors, merchandizers need reliable product performance data to strengthen their position. This data helps them negotiate favorable prices and contract terms, and ensures that no potential value is left on the table.

By automatically creating summaries of product performance and correlating them with other important commercial data, generative AI can accelerate vendor negotiations and ensure that negotiators are equipped with verifiable data points to back up their perspective.

Empowered with a clear, concise and contextualized summary of relevant data, merchandizers can go into negotiations with a complete view of what they need from a vendor, and the exact price they need to be able to secure it at to make a profit. So, in addition to helping them make a robust case for their desired terms, generative AI also helps prevent them from making any decisions that negatively impact value or profitability.

Moment 4: Analyzing failing product lines

Product lines can fail for a huge range of reasons. To understand exactly why a product line hasn't delivered its projected or intended value, teams must correlate numerous data sets and sources.

Generative AI helps merchandizing teams draw insights from their diverse data and perform far more accurate root cause analysis of failed product lines. Equipped with that insight, teams can quickly remediate issues with product fit or design.

With a deeper understanding of why a line has failed, teams are empowered to tweak strategies and make improvements that minimize loss and drive revenue growth. They can understand when a product is trending towards failure, before it actually happens — giving them the chance to turn things around and transform lagging product lines into leading ones.

Moment 5: Selecting new products

Choosing which new products to buy from suppliers is a tough task. Even products that appear strong and most relevant to customers can fail to deliver value if the conditions aren't right.

Generative AI supports product selection by analyzing the past sales performance of thousands of similar SKUs alongside customer behavior attributes and feedback to help build a correlation model. Also, by leveraging generative AI, merchandizers can analyze unstructured data in real time, extracting valuable insights from various sources. For example, generative AI can rapidly aggregate and conduct sentiment analysis on social media posts or identify patterns and trends from online reviews or customer emails. Using that model, merchandizers can make data-driven purchasing decisions instead of basing their choices on intuition.



Moment 6: Personalizing customer shopping journeys

Most e-commerce sites typically follow a fairly standardized and inflexible layout, presenting the same content, images, and advertisements to most users irrespective of their individual preferences and selections.

Generative AI empowers e-commerce retailers to deliver highly customized site experiences tailored to each customer's unique interests. In essence, each shopper encounters a personalized website where products are automatically populated based on data stored about that customer within the backend system.

This enables an unprecedented degree of personalization, and gives merchandizers the opportunity to deliver experiences that are meticulously tailored to a shopper's behavioral and demographic data.

Before and after: Generative AI in merchandizing

Before generative AI	With generative AI	Enabling technologies	Business benefits
Product design is a long, complex process involving a huge range of stakeholders.	AI helps rapidly generate product ideas and designs that merchandizing teams can validate and build on.	Image generation, content generation, data aggregation.	Time to market for new products is massively accelerated.
Product assortments are tough to optimize, and physical and online retail channels are siloed.	AI helps teams spot patterns across diverse data and uses that insight to automatically generate optimal merchandizing strategies.	AI summarization, purpose-built insight solutions, tagging capabilities. Natural language querying of data bases.	Teams can continuously optimize and improve merchandizing strategies to drive revenue across physical and online retail channels.
Vendor negotiations are tough and inconsistent.	AI automatically generates a clear view of merchandizing and vendor performance data that can be used to inform vendor negotiation.	AI summarization, automatic reporting.	Merchandizing teams are empowered with data and quick summaries to back up their negotiations, helping them secure what they need from vendors at the right cost.

Pitfalls and points of caution

To realize this vision of generative AI-enabled merchandizing, your organization must implement AI with care, and take crucial steps such as:

- **Implementing quality control guard rails:** Ensure everyone generating AI outputs can spot hallucinations, flag them, and make sure they don't end up in front of the wrong audiences.
- **Cleaning up the data that will be used to train LLMs:** Ensure that errors in data sets are not propagated and inaccuracies are not automatically fed into outputs.
- **Ensuring fit:** Carefully evaluate whether widely available LLMs such as ChatGPT are fit for your intended use cases, or if you'd be better off creating a custom LLM that's tailored to generate the results your business specifically needs.

Generative AI for marketing

Great marketers understand what their customers and audiences want, and work tirelessly to put the right content in front of the right prospects at the right time.

That's always been challenging. But in recent years, it's gotten a lot harder. Trends and preferences are evolving faster than ever, disruptive conditions and technologies have upended entire markets overnight, and customer expectations have risen sharply.

Marketers are being challenged to move quicker and adapt their strategies with unprecedented speed. But for many teams, their processes and capabilities haven't evolved to keep up. Until now.

We are at the cusp of Marketing 5.0, a new era that blends the technology-driven focus of Marketing 4.0 with the human-centric approach of Marketing 3.0. This shift signifies a powerful convergence of technology and human understanding to create, communicate, deliver, and enhance value across the entire customer journey.

It's against this backdrop that generative AI emerges, offering a transformative opportunity to marketers. Generative AI mimics and amplifies human creativity and intelligence, enabling it to accelerate tasks previously reserved for humans. This "perfect storm" of human-mimicking technology and a holistic marketing vision unlocks a myriad of possibilities for marketers. Generative AI co-pilots can curtail many of the biggest bottlenecks to the timely delivery of deeply personalized, high-quality marketing materials and campaigns.

Let's explore some concrete ways generative AI can create value in an omnichannel, digital e-commerce, or marketplace commerce environment.

Key moments in the life of a generative AI-empowered marketer

Moment 1: Creating new content

Marketing content creation can take a long time. Teams must liaise with category owners to gather information, factor in timely customer insights, work with suppliers and agencies who provide creative services, and constantly ensure the content they publish aligns with brand guidelines and strategic and campaign-level goals.

Generative AI co-pilots can streamline that process significantly. From quickly combining background information into creative briefs, to generating visuals and content from approved knowledge assets, generative AI empowers marketers to create and publish content at speed and scale — and accelerate time to market.

For instance, Einstein Generative AI in Salesforce Marketing Cloud lets marketers craft custom brand personalities, generate subject lines in line with their brand identity, and craft body copy for emails, accelerating key aspects of campaign creation.

In addition, large language models also enable reusability at scale, by replicating campaigns and assets for new audiences and markets. So, when the marketer needs to internationalize their content, for example, they can quickly generate new localized versions of assets they've previously created.

Moment 2: Personalizing content and experiences

According to McKinsey, personalization can deliver five to eight times the ROI on marketing spend. The journey to personalization really begins with marketers knowing enough about a prospect or a customer to start personalizing content and propositions for them, and delivering highly relevant experiences tailored to their exact needs.

Building on the foundation of AI and the marketing automation tools of today, generative AI co-pilots can enable marketers to do that effectively at scale. When a prospect has been fully identified — and data about their journey and needs has been attributed to them — marketers can use generative AI tools to generate personalized content for that individual almost instantly, with effective prompts and guardrails.

Marketers' generative AI co-pilots can also help them better understand who their prospects really are — providing unique slices and cuts of customer cohorts. By segmenting customers based on AI-generated profiles and microsegments, marketers can get granular with their campaigns, and make every prospect's content experiences as relevant as possible from the beginning.

Moment 3: Connecting with prospects and converting them through conversation

AI-powered conversational commerce enables marketers to get closer to their audiences than ever before, without having to engage in conversations themselves. Using AI chatbots, prospects can ask questions, get immediate answers, and accelerate their own discovery and buying journeys without the need for direct intervention from marketers or customer service agents.

This enables teams to break away from search- and browsing-led journeys, and plan out a new conversational path to purchase, where prospects are guided towards a product variant that's right for them. And because the bulk of conversations are handled by AI, marketers can spend their time focusing on the engagements and decisions that need their attention the most. Indian online fashion retailer Myntra's MyFashionGPT is an early promising example of this, and is already changing how people search and shop on the platform.



Moment 4: Analyzing performance and planning out the next campaign

AI co-pilots monitor ongoing campaign performance and give marketers a holistic view of what's working and what's not, so they can intervene at the right time and ensure their efforts deliver the right results.

With intuitive natural language summaries of individual customers' journeys and overall campaign performance, marketers can quickly see which aspects of their campaigns had the greatest impact. Using that insight, teams can ensure that every campaign they plan is stronger and more relevant than the last — continuously improving conversions and ROI.

Google Analytics 4, for instance, is upping the game with AI-driven natural language narratives and insights. The conversational aspect of generative AI in BI tools like PowerBI provides a querying interface or co-pilot that can help teams drill down to the right level of detail with descriptive queries without the need to wade through complex data correlations and charts.



Moment 5: Building insight-driven customer journeys

As customer needs and habits change, it's important for marketers to ensure that the journeys they guide prospects through remain aligned with how, when, where, and why people buy.

By analyzing customer journeys at scale and contextualizing them with campaign performance and sales data, generative AI co-pilots can help marketers map out optimal journeys.

At every touchpoint, AI can help marketers understand the next best action for every one of their prospects and ensure they're engaging them with the right offers at the right time, whether that's a nudge to come back and purchase a product in their cart, an exclusive preview to an upcoming collection, or a personalized offer aligned with a recent life event.

Before and after: generative AI in marketing

Before generative AI	With generative AI	Enabling technologies	Business benefits
Creating marketing content for broad audiences is time-consuming and requires tight coordination between dispersed teams.	Marketers can use AI to generate imagery and content, accelerating the content creation process.	Image generation, LLMs, content generation, tagging capabilities.	Content creation is significantly accelerated, creation costs are lower, and time to market is reduced.
Personalization is complex and individualisation / hyper-personalization is tedious and expensive to do at scale.	AI helps teams spot patterns across diverse data and use that insight to automatically generate optimal marketing strategies.	AI summarization, classification, clustering, LLMs, content generation with variations, plus integration with other systems including CRMs.	Customer experiences are improved, marketing experiences are more relevant, and lead generation is accelerated, creating positive impact on upper funnel metrics.
Marketing is a one-way process because it's so hard for marketers to connect with customers at scale.	Conversational commerce experiences powered by AI close the gap between marketers and customers, enabling customers to drive their own journeys.	Conversational AI, LLMs, chatbots.	Customers are empowered to ask questions throughout their journey, deepening engagement and helping marketers better understand and meet their needs.

Pitfalls and points of caution

To ensure that generative AI leads to the right outcomes for marketers and positively impacts customer journeys and experiences, teams must:

- **Carefully and continuously monitor the quality of AI-generated content and visuals:** Ensure that they're in line with branding, and up to the quality standards of the business. The downside of AI generation is that it's prone to hallucination where reference data lacks customer and organizational context. Hence, purposeful intervention is a must to steer clear of bias and hallucination.

- **Build new approval processes and workflows:** Ensure that as marketers are empowered to generate their own content, outputs are still verified and reworked by content and design experts before going to publishing.
- **Scrutinize any customer-facing capabilities such as chatbots closely:** Ensure they're consistently able to deliver exceptional customer experiences that are aligned with the organization's brand promise.

Generative AI for customer service agents



As retailers seek new ways to stay ahead of their competitors, customer service has quickly become a key differentiator.

Today's customer service agents are tasked with serving more customers, to a higher standard, through more channels than ever before. It's a tough job. And its high volume, high repetition, and high churn nature make it even harder.

Exceptional customer service demands three key things:

- Being present when and where customers may need service
- Understanding the need of the customer and being able to build a meaningful connection at that micro moment
- Having a comprehensive view of knowledge systems and fast access to information to be able to address customer needs quickly, without compromising on accuracy or service quality

Generative AI co-pilots can help human agents perform better in all three areas. With the help of intelligence-augmenting capabilities, there's no limit to the volume or variety of touchpoints through which they can deliver exceptional customer service.

AI-generated summaries empower every agent with a clear view of the customer's issue and journey so they can provide the right experiences instantly. Plus, by capturing information directly from agent-customer interactions, generative AI can support faster creation and updating of knowledge assets — helping every agent learn and share knowledge with ease.

Key moments in the life of a generative AI-empowered service agent



Moment 1: Assisting customers with queries in real time

When a customer reaches out with a query, AI can analyze the sentiment and content of a customer's initial message to determine whether their question could be best handled by a service agent or a chatbot. By offloading simple, routine queries to chatbots, human service agents are able to be a lot more responsive to complex queries — the ones that demand their human expertise.

When a customer needs direct help from an agent, that agent is immediately empowered with an AI-generated overview of the customer's journey up to that point. That enables them to engage with the right response, without asking repetitive questions.

Customer and query analysis, subject/issue understanding, customer sentiment analysis, and real-time language translation all come together to significantly accelerate response times for customers, and increase satisfaction. Plus, by eliminating repeated questions about their need from the interaction, the service experience is significantly improved for the customer, and the agent can progress the customer's engagement quickly. Another key advantage of generative AI is language translation. With such capabilities, a customer agent can service customers across different languages even with basic language proficiency.

As the agent prepares responses to the customer's query, generative AI can turn the basics of their replies into eloquent, empathetic, and informative messages. This cuts response times even further, enabling the agent to support more customers throughout their day, without sacrificing the quality of service that each customer receives.



Moment 2: Post-call processing

When a service agent ends a customer interaction, they must complete post-call processing. That typically involves uploading a contact summary and disposition code to the CRM system.

With capabilities like automated note-taking, summarization, and language translation, generative AI can automate much of that process — cutting the time agents have to spend on post-call processing.

As information is automatically captured from customer interactions, it can also be automatically tagged, making it simple for service agents and other teams across the business to search for, learn from, and report on.



Moment 3: Strengthening the knowledge base

The contact center knowledge base is an important resource for all service agents. But, there isn't always an answer to every agent's question. When an agent needs information that isn't currently in the knowledge base, they can flag the gap, so that as soon as the knowledge is gained, it can be translated into a new asset that forms the basis of AI-generated responses related to that topic.

During and after customer engagements, generative AI capabilities can analyze customer intent, assess how agents and supervisors have successfully handled similar queries, and use that information to develop a new knowledge article, or update an existing one. This can then be reviewed and edited by an appropriate agent or manager before publishing.



Moment 4: Helping agents perform better

With high churn in the department, onboarding and education are a continuous challenge for customer service teams. But, thanks to the library of knowledge assets created by generative AI, new agents are empowered to answer most of their questions themselves.

AI can detect patterns and gaps in agent interactions and use those insights to generate tailored learning and development plans. Agents can quickly be routed to relevant learning resources, closing knowledge gaps in minutes.

That enables the continuous improvement of performance across the service team. Any gaps in agent responses can be flagged immediately, and relevant agents are empowered with the insight and knowledge they need to improve their performance and deliver better customer outcomes.

These moments are just the beginning of the possibilities that can be offered by generative AI capabilities. When applied holistically with a strong data value chain, these capabilities can empower agents to provide effective service to customers across the length and breadth of their journeys, and revolutionize the customer service function.

Before and after: Generative AI in customer service

Area	Before generative AI	With generative AI	Example LLM capabilities	Illustrative benefits
Handling customer queries/ interactions	<ul style="list-style-type: none"> • Time-consuming • Long and repetitive • Disconnected information sources • High expectations without enablement 	<ul style="list-style-type: none"> • Agents are better equipped with the information they need to respond quickly and appropriately to customer needs. • Implicit or supporting activities e.g. note taking, post call processing, follow ups etc are automated 	<ul style="list-style-type: none"> • Content creation & augmentation • Sentiment, Tone of content • Content classification • Summarization • Language translation • Live transcription • Conversational AI/ Smart Response 	<p>Improvement in agent performance indicators:</p> <ul style="list-style-type: none"> • Resolved cases • Customer Satisfaction (CSAT) • Avg. handling time • Cust. interactions • Time to first response <p>Improvement in influenced metrics such as up-sell, cross sell, revenue, ROI, product usage, and reduced churn</p>
Knowledge Management	<ul style="list-style-type: none"> • Disconnected, heavy and dated • Time-consuming • Less coverage 	<ul style="list-style-type: none"> • Knowledge is more accessible and consumable within agents' operational constraints. • The knowledge base is continuously updated and improved. 	<ul style="list-style-type: none"> • Content creation & augmentation • Discovery and simplification • Content classification • Summarization • Language translation • Live transcription • Conversational AI/ Smart Response 	<ul style="list-style-type: none"> • Faster onboarding • Better performance • Reduced churn • Improved CSAT and employee satisfaction (ESAT)

Area	Before generative AI	With generative AI	Example LLM capabilities	Illustrative benefits
Personal reach and growth	<ul style="list-style-type: none"> • Similar and repetitive roles • Bounded context to operate • Time constraint • Resource unavailability 	Customer service agents can access coaching opportunities, improve their skills quickly, and service customers with broader reach.	<ul style="list-style-type: none"> • Content creation & augmentation • Discovery and simplification • Summarization • Language translation • Live transcription • Conversational AI/ Smart Response 	<ul style="list-style-type: none"> • More proactive reach • Better agent performance • Reduced churn • Improved agent experience and motivation

Pitfalls and points of caution

To ensure that generative AI capabilities deliver the right results for customer service agents and the customers they serve, teams must pay attention to a few key points:

- **Data quality and quantity:** Generating accurate and contextually relevant responses relies heavily on the strength of the data value chain. Poor data quality or low data quantity can lead to inadequate answers, potentially frustrating customers and damaging brand reputation. Hence, companies need to pay attention to continuous data quality and enrichment as well as have a “human in the loop” approach to bake in a human review system.
- **Maintaining brand voice and tone:** This is extremely important in customer service to deliver a cohesive customer experience. Hence, all AI-generated customizations, templates, and communications need a strong mechanism to ensure they stay aligned with brand voice and tone.
- **Handling complex customer interactions:** Even though most customer queries fall into the same broad categories, there will always be complex and unique outliers. You’ll need to define clear escalation and classification mechanisms to handle those cases, and ensure that your knowledge base is continuously updated with learnings from such interactions.

- **Ethical guidelines and biases:** Attention should be given to ensure ethical guardrails are put in place, and mechanisms are created to avoid biases in responses or information
- **Privacy and security:** Since customer service is the front line for customer interactions and data, strong security and privacy policies should be built into your generative AI plans from their conception. data lacks customer and organizational context. Hence, purposeful intervention is a must to steer clear of bias and hallucination.

How you can make the most of the generative AI opportunity in retail

Generative AI has the potential to accelerate workflows, empower human teams, and drive better results across virtually every function within the modern retail organization. But those benefits don't just materialize by default; they must be implemented strategically and iteratively.

To make the most of the generative AI opportunity in retail, we recommend following these five tips:

1

Strategically implement generative AI to get the most from your current capabilities

Generative AI shouldn't be approached as a new technology silo. Rather, it's a new capability that can help you get the most from your existing technology and data estate.

It should build on a strong data foundation, and evolve in conjunction with your existing analytical and AI/ML capabilities across the full spectrum of descriptive and prescriptive capabilities to make the most of all the opportunities generative AI has opened up.

AI is only as good as the data that fuels it. So, start your journey by focusing on generative AI use cases where you have a lot of high-quality data. For example, if you've gathered vast quantities of data about your customers — who they are, when and why they buy — you could use your success in marketing use cases to expand your generative AI efforts into other functions.

2

Work closely with teams to design the perfect generative AI co-pilots

Just like a human co-pilot, generative AI co-pilots should support how your teams work and make their working lives easier — not force change on them. As you identify generative AI use cases across your organization, consult with teams to discover the real points of friction in their workflows.

When AI is applied to solve the challenges faced by human teams, they adopt it more quickly. They are motivated to embrace new co-pilots and learn how to make the most of them, which is essential if you want to drive lasting operational change across diverse functions.

You'll need to ensure that the people using your co-pilots are supported with relevant training and upskilling. Without that, adoption will be patchy, and the results will be inconsistent across teams.

3

Don't cut humans out of the process

Everything your generative AI capabilities generate needs to be validated by a human. The technology has seen huge advances in recent years, and some capabilities are more mature than others. Areas like text generation have become very consistent, but if you branch out into more novel use cases such as video generation, where tools are still prone to hallucinations and even IP issues, the results will be less reliable.

A good example can be seen in autonomous LLM agents. Today, even the best GPT-4 agents can only perform around 15% of tasks successfully, compared to the average human success rate of 92%. If an organization were to pivot to exclusively use autonomous LLM agents today, it could have a catastrophic impact on their business and customer experiences.

By keeping humans in the loop and building human verification into your generative AI workflows, you can ensure that the results presented to customers and internal stakeholders are always human-ready and consistent with the content your human teams produce.

4

Appoint new roles to support your generative AI transformation

While generative AI co-pilots are designed to be highly intuitive, you'll want to get a few new skills in place before your organization starts using them. You'll need people who understand potential intellectual property issues that arise when generating content, and analysts with a deep understanding of generative AI outputs.

To get those skills in place, you can either develop them internally or hire new people with them. Whichever path you take, it's a good idea to create some new roles that are responsible for overseeing your AI efforts, and that are ultimately accountable for the generated outputs your people share with the world.

5

Make security a priority from day one

Like any data-driven technology, security should be considered and prioritized at every stage of your generative AI journey. The information security risks introduced by generative AI fall into two major categories: information leakage and vulnerability introduction.

Information leakage comes in three main forms.

- Public data leakage happens when information from your website or public content is scraped by someone else's AI model.
- Prompt leakage happens when public AI models learn from the prompts your people give them.
- Personal data exposure happens when personal data — such as information you hold on your customers — is fed into models to generate outputs, and it starts to appear elsewhere.

Vulnerability introduction arises from quality issues prevalent in AI-generated code. While that's not as relevant in retail use cases, if your teams start using generative AI to generate new code, you'll need to stay on top of these security issues.

Find the right use cases and seize the retail generative AI opportunity today

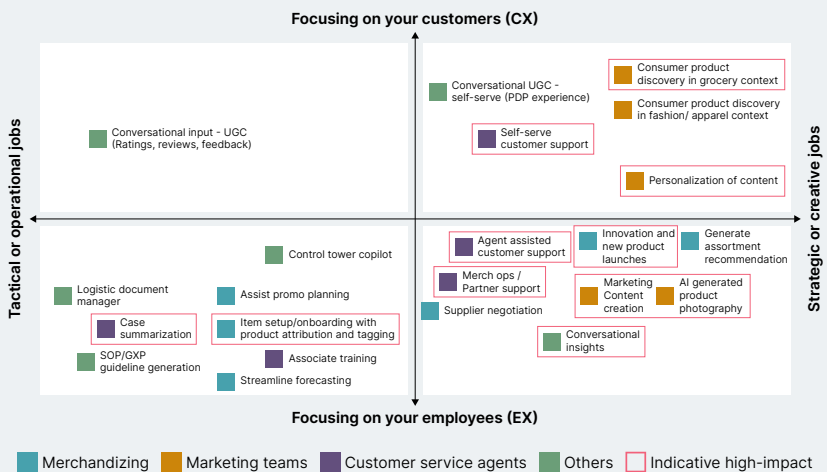
Thoughtworks advocates for a human-centric approach to generative AI use case selection. At its core, generative AI creates ways of augmenting human intelligence, capabilities, and experiences. So, your journey should start with the identification of which capabilities, processes, and experiences you want to transform — whether that's for your customers, or your employees.

To maximize the value generative AI can deliver, it's a good idea to start out with a balanced set of employee and customer-centric use cases, leaning first towards employee-centric experiments. Within those, it's also worth balancing more tactical use cases that can be pursued to demonstrate value quickly, with strategic or creative use cases that will take longer to come to fruition, but generate more significant long-term ROI.

The matrix below shows how use cases can be identified and mapped across those criteria to build a balanced generative AI use case portfolio.

Identifying use cases for your business

A human-centric approach

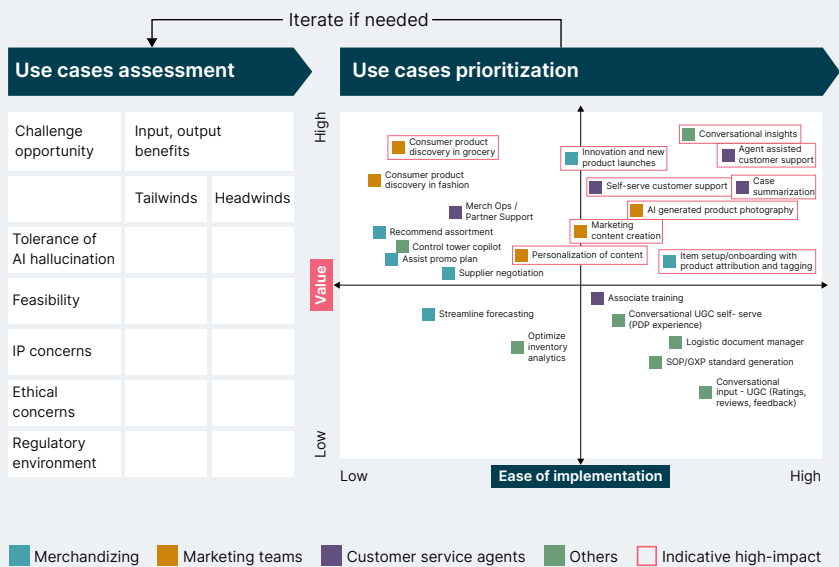


Once you've identified a balanced set of relevant generative AI use cases, you need to carefully prioritize them and choose which to move ahead with first. This is a critical step, as you'll need to identify quick and clear wins that can demonstrate the value of generative AI to the rest of your business in order to secure executive buy-in for your plans, and drive adoption and excitement around future use cases.

Thoughtworks uses the below framework to make value-based assessments of generative AI use cases.

A framework to prioritize use cases for your business

A value-based approach



Faster, smarter, stronger

Assessment criteria including ethical concerns and tolerance for hallucinations can help you identify safe spaces for experimentation that lend themselves well to early generative AI use cases. And by mapping use cases out by their potential value and ease of implementation, you can identify both “quick wins” (use cases with high value and high ease of implementation) and long-term high-impact cases (those with low-to-medium ease of implementation and very high value).



Start your journey to generative AI success today

Thoughtworks works closely with organizations across all industries to help them make the right generative AI decisions, and tightly align generative AI capabilities to their strategic goals. If you're evaluating your AI options and wondering where your journey should begin, talk to us today.

Learn more on [Thoughtworks/Retail-e-commerce](#)

Authors

Sanjeev Athreya

Head of Retail, CPG and Commerce Practice

Sanjeev leads Thoughtworks' Global Retail practice, bringing nearly three decades of industry experience to clients around the world. He has deep expertise in helping retailers reinvent themselves to stay relevant through significant paradigm shifts in technology and consumer mindsets. He enjoys music, traveling and reading on geo politics.



Preetisudha Pandab

Retail and Commerce Industry Principal


Preeti is a retail and commerce expert with more than 16 years experience. She leads Thoughtworks' Digital Commerce center of excellence and helps clients redefine business challenges, simplify complexity and use agile, outcome-driven approaches to solve their most complex business challenges. Preeti specializes in digital commerce, digital marketing, web analytics, personalization, content marketing, social media, loyalty, and more.



Anitha Rajagopalan

Retail and Commerce Industry Principal and Product Strategist

Anitha is a product leader with more than 18 years of retail industry experience. She is focused on driving business outcomes for strategic retail and commerce clients and leads Thoughtworks' Retail Media center of excellence. Anitha specializes in e-commerce, omnichannel, retail media and monetization, merchandizing, pricing and promotions, in-store innovation and distributed order management.



Thoughtworks is a global technology consultancy that integrates strategy, design and engineering to drive digital innovation. We are over 11,500 Thoughtworkers strong across 51 offices in 18 countries. For 30 years, we've delivered extraordinary impact together with our clients by helping them solve complex business problems with technology as the differentiator.

[thoughtworks.com](https://www.thoughtworks.com)



Strategy. Design. Engineering.

© Thoughtworks, Inc. All Rights Reserved.