

Forward. Fast.

The Future of AI in Product Discovery



In a sea of product search and discovery vendors quick to hop on the AI bandwagon for their own gain, it's not uncommon for ecommerce decision makers to view the many solutions as somewhat homogenous, ultimately landing on the one that provides the best cost savings.

But as the ecommerce landscape continues to evolve, tech leaders will need to make informed, strategic investments in technology that not only meet their immediate business needs, but also adapt to the breakneck developments of AI. This technology will also need to continue enhancing customer engagement, which involves hyper personalization of the shopping experience.

The ultimate goal is to select a solution that not only has the flexibility to evolve over time with emerging technologies and industry trends, but also helps advance important business KPIs while sparing organizations the burden of replatforming in the near future.



AI Powered Search




Table of Contents

4	Back to the Past Future	Explore how product search and discovery solutions have evolved over the years, thanks to shifting customer expectations and advancements in technology.
5	Keyword search	
6	Vector search	
7	Clickstream-based product discovery	
9	Transformer-based product discovery	
10	Beyond the Hype: Where We Are Now	Unravel the intricacies of emerging technology and learn how they can help drive tangible ecommerce results.
11	Transformers	
12	Large language models (LLMs)	
14	Clickstream	
16	Natural Language Processing (NLP)	
17	Generative AI	
22	The Future Looks Bright	Discover predictions from Constructor CEO Eli Finkelshteyn on what online shopping could look like in five years when enabled by future-fit tech, like Constructor.
26	Future-Fit Product Discovery	Learn how to fairly evaluate vendors' current merit and future-fitness in an AI-driven world.

Back to the ~~Past~~ Future

Search and product discovery has evolved immensely since its early days, continually shifting to keep pace with ever-changing customer expectations and advancements in technology.

A search bar with a light gray background and rounded ends, set against a dark blue background with a pattern of overlapping circles. The text "Search by Context" is in a dark gray font, followed by a vertical line. To the right is a circular button with a blue gradient and a white power icon.

What started with keyword-based engines, which were limited in their ability to return attractive products, evolved to include vector search. While vector widened the net of relevancy, it still didn't solve well for attractiveness (the idea that search isn't just about showing "relevant" products, but about showing products that are most likely to be purchased by a given user in a given context). Thanks to quickly evolving AI and machine learning (ML), clickstream-based approaches came to the forefront, allowing retailers to present highly personalized results. Nowadays, things have gone even further, with the most innovative

companies experimenting with emerging tech like transformers, which allow for an even deeper understanding of what users are looking for when they search, browse, and interact with product recommendations.

Below, we describe each of these technologies in more detail, shining more light upon each solution's overall strategy, challenges, and business outcomes.

Keyword search

They match queries to indexed items — a.k.a. rank products by frequency of matched keywords. The essential logic is that products are considered more relevant based on the number of times the keyword appears, and they often include the ability to set manual rules to break ties when they occur.

When designing an engine meant to work in the most possible use cases — from searching over documentation to social networks and ecommerce sites — it's a lowest common denominator that works OK, but not great.


Take content as an example. Visitors expect an article about hiking to frequently mention the keyword "hiking." Most likely, you'll find some decent matches. But challenges arise when it's applied to use cases like ecommerce that have a specific goal (e.g., showing a shopper something they actually want to buy). The word "relevant"

Keyword search engines are some of the oldest engines in existence. The search companies you've likely heard of most, like Elasticsearch, Solr, or Algolia, are built on these.

starts to lose its meaning. As in, is "peanut butter" relevant to a search for "butter?" If you ask a keyword engine, it might be.

More importantly, even products that do genuinely look relevant aren't always attractive, nor lead to purchases. For example, if you visit an outdoor gear website and search for "hiking boots," the site will return an assortment of hiking boots. But among this product results set will be hiking boots that are out of stock, discontinued, and even poorly rated. Even though they're technically relevant, they're most likely not what you're searching for.

Not to mention, missing from this product results set are hiking boots that lack the keyword in their product data (a problem which [AI can help with](#)).



The end result? Customers drown in unattractive results the system deems to be "relevant" in one way or another, which worsens conversions and overall business outcomes.

Product search and discovery vendors often compensate for these poor results by layering manual rules on top of their keyword-based engines, supplementing with bolted-on AI, or both. These cobbled-together solutions can help re-rank products or expand the result set: changing the order that they appear on search results pages based on popularity, or determining how to weight product page description keywords in relation to title keywords. But it only leads to a [new kind of monster](#).

Vector search

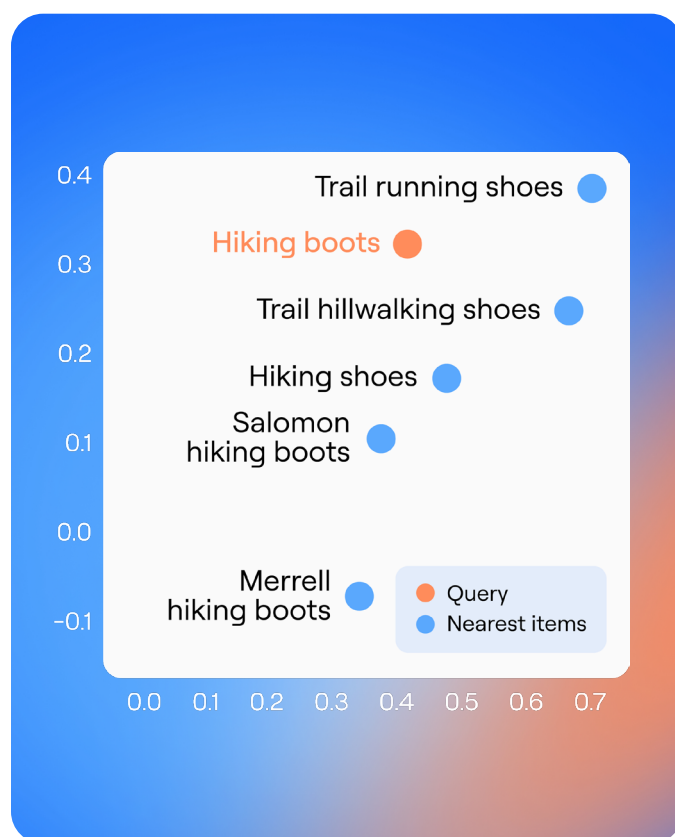
They're often described as representing products as a "sea of stars," where every product is a "star" that's mapped according to similar concepts and associations. From there, the algorithm can cluster, extrapolate, and measure the distance between neighbors to make an inference about what the user actually intended for a given search query.

For example, we can understand the relationship between the search query "hiking boots" and the result "trail hillwalking shoes" by training models on product data, such as categories, product names, and text descriptions.

The algorithm maps out the relations from "hiking boots" to "hiking shoes" to "trail running shoes" to finally "trail hillwalking shoes." It narrows down the search engine's understanding of "hiking" by mapping a specific term to neighboring concepts. Since there is always a closest "star" to a "point in space," we are able to serve results for every search query.

Unfortunately, vector search presents challenges. Even though it tends to improve overall product recall and lessen the likelihood of zero-result searches, it can also lead to lots of irrelevant and unattractive results getting into the mix. For example, "hiking pants" might also be conceptually related to "hiking shoes" when putting the two into vector space, but they are not a good result to return. Plus, more results doesn't necessarily equal more attractive results — nor tangible business value. And they don't tell you much about how

[Starting around 2013](#), vector search started to appear in the product search and discovery landscape. It was an exciting breakthrough at the time. Instead of simply returning products based on keywords, vector search algorithms transform a product catalog into numerical vectors.




to rank the newly added results. Finally, similar to keyword-based engines, personalization is an add-on at best.

While Constructor also uses vectors (we were the first ecommerce product discovery company to do so), we rely on clickstream associated with conversions to generate conceptually similar products that are actually purchased for a given query.

Clickstream-based product discovery

The next iteration to the evolution of product search and discovery encapsulates a more holistic approach that's based on [clickstream data](#).



Clickstream data is the trail of your customers' activity on your website, from the time they enter to the time they leave. Each interaction during their session is recorded and stored in the clickstream — anonymously.

Clickstream-based product search and discovery engines like Constructor ensure products returned site-wide are both attractive to users and advance business KPIs, like conversions, revenue, inventory balancing, and more. Originally, we used vector embeddings in tandem with clickstream to find the most attractive product for a user. A few years ago, we upgraded to using transformers with clickstream instead (there's more on those and why we made the switch later on).

The reason pairing clickstream with vector embeddings is so critical is that the clickstream serves to anchor suggestions given to a user to only those that have conceptual similarity to

products that are actually purchased in a given context.

For example, with pure vector search, someone searching for "chips" might have a sack of potatoes returned as a result. Conceptually, there are similarities. After all, chips are often made from potatoes. But no one searching for "chips" ever buys a sack of potatoes.

Clickstream allows us to infuse our algorithm with this key point: you can see via the clickstream that sacks of potatoes are never purchased for the query "chips." So, this prevents them from appearing as a result. Conversely, the clickstream can show that products similar in concept to "chips" — like "vegetable crisps" — actually are purchased often for this query or queries similar to it. So, they may be prioritized in the product result set.

Similar to previous technologies, clickstream also has its challenges. For one, clickstream is domain-specific. The actions you can take when searching or browsing on an ecommerce site are not the same as those you can take on a blog or social network. The engine works very well if clickstream is similar (i.e. solely ecommerce clickstream, where you can assume that actions like clicks, add-to-carts, and purchases are standard, along with product attributes like cost and margin). This is because in this situation, the system can learn from the clickstream in a standard way. On the other hand, if an engine is built as one-size-fits-all and is geared for every search use case under the sun, you can't make many assumptions about the clickstream at all, rendering it much less useful.

Also, clickstream requires a high volume of data to function well and for the ROI to make sense. This is why Constructor chose to focus on partnering with large ecommerce companies that share similar clickstream from day 1, a domain focus that still allows us a leg-up on legacy engines and unfocused competitors with solutions for all use cases.

The outcome is that by using clickstream within enterprise ecommerce, you get more attractive, hyper-personalized results that are tailored to customers, optimized for tangible business value, and lead to more purchases and revenue. The best part? It can be proven on an ecommerce company's own data before they ever need to sign a contract.

The Proof Schedule™: Part of Constructor Since Day 1

Constructor's platform was built by data scientists and engineers who were dissatisfied with the status quo of search engines. They envisioned an ecommerce world where products returned for a query were hyper-personalized to the user and products people actually wanted to buy in the context of that query, a world in which the product search and discovery experience was a win-win for both shoppers and companies alike. So, they got to work on building Constructor, an AI-native solution that prioritizes learning from clickstream and customer behavior. Using customer behavioral data ranks products better and also allows large ecommerce companies to test our platform risk-free with their own data and see the difference for themselves. Since day 1, this proved to be our most successful differentiator and the reason why some of the biggest ecommerce

companies in the world power their product search and discovery experience with Constructor. It's something we still do to this day and has become a core pillar of who we are, spawning the phrase "don't just trust us, make us prove it."



Transformer-based product discovery



The future of product discovery as molded by AI and ML is happening now, and it revolves around transformers (the “T” in ChatGPT).

This emerging technology addresses a key problem with vector search, which is not having context around individual words. For example, vector search falls short for search queries like “flights boston to new york” or “shoes not blue.” A vector search engine wouldn’t be able to derive the meaning behind keywords in both queries: “to” and “not.” It wouldn’t understand that “flights boston to new york” is different from “flights new york to boston,” or that “shoes not blue” is asking for every color except blue. This leads to incorrect result sets for the search intent.

In contrast, transformers have a deep understanding of the query, especially when paired with user-level clickstream data. This allows transformers to rank and filter results by attractiveness, which are further optimized by business goals and LLMs (more on these later).

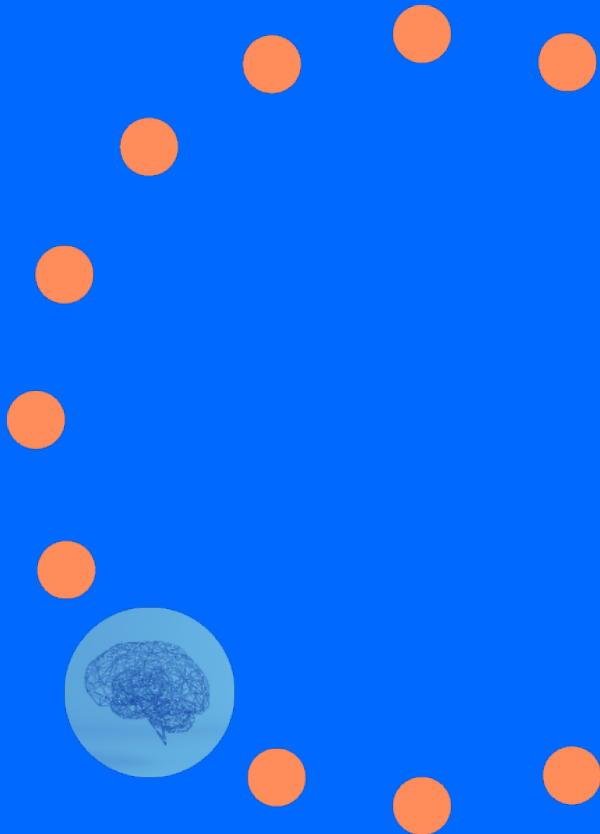
There are a couple downsides to transformers. Firstly, because they’re still cutting edge, experimentation is needed to determine which use cases they serve best. Here at Constructor, we’ve been experimenting with them since 2021 for this exact reason. Transformers are also expensive to train, as they require extremely specific data within the domain you care about (i.e. ecommerce) to work well.

But the outcomes are bright. Because they use text and clickstream as context, transformers return even more attractive and personalized product result sets across ecommerce sites. That’s led to significant A/B test wins throughout the years, as proven by the 2–3% additional marginal lifts in revenue for our customers that we announced about a year ago and have been building on since.



Beyond the Hype: Where We Are Now

Generative AI, transformers, LLMs. Many competitors throw these terms around in their marketing copy in hopes to snag the unsuspecting prospect. But not all their search engines leverage these emerging technologies.



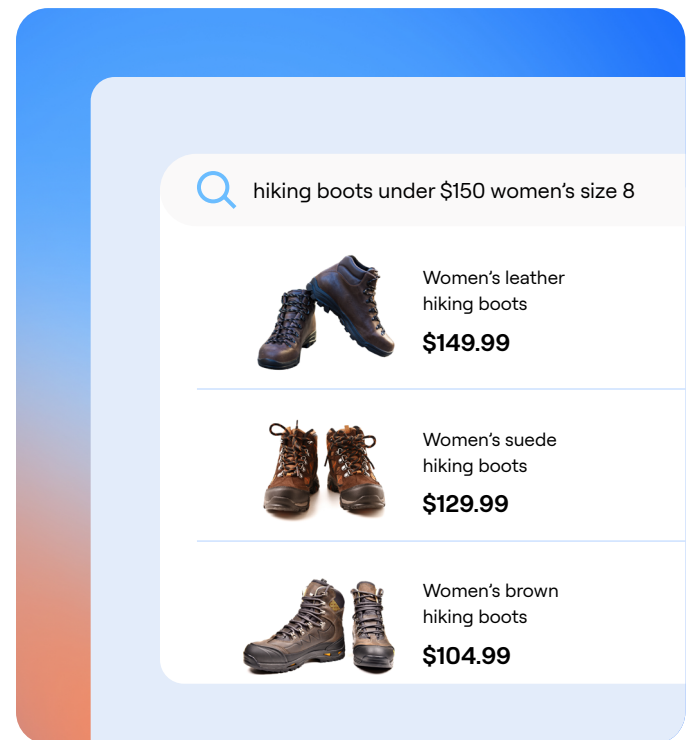
Get beyond the AI hype in this next section as we dive further into the nitty gritty of emerging technologies, including their importance in enterprise ecommerce and — most importantly — how they help you hit important business KPIs.

Transformers

Transformers are advanced algorithmic models trained on large ecommerce data sets. They're built on deep learning architecture with a natural language processing (NLP) model to weigh the importance of different words in a sequence.

Prior to their advent, leading AI translation methods relied heavily on recurrent neural networks (RNNs). These RNNs processed each word sequentially within a sentence. However, significant relationships among words often exist, regardless of whether they are sequential. Transformers use a mechanism known as attention, enabling models to recognize the relationships between words — irrespective of the distances between them — and to identify the critical words and phrases within a passage.

And because of the way they're built, transformers can better comprehend the context of each term in a query, making them useful for processing complex [natural language search queries](#) (i.e. "find me hiking boots under \$150 that come in women's size 8").



Why they matter in ecommerce

The more ecommerce data you train your transformers with, the more fine tuned your base model becomes to your specific use cases, regions, unique retailer attributes, etc. The outcome is an iterative process that continually evolves based on the depth of insights ingested, enhancing the precision of the overall product search and discovery experience.

And thanks to having a deep understanding of a query in the context of clickstream, the use of

transformers allows ecommerce companies to return more attractive and personalized results across their entire suite of product search and discovery tools — every interaction, every time.


This is different from vectors, which increase the number of products returned for a query, but don't have much understanding of whether a product is attractive nor ensure user-level attractiveness.

Large language models (LLMs)

Large language models (LLMs) are a category of AI models that have been trained on massive amounts of text data to perform natural language understanding and generation tasks.

These models are also typically based on deep learning techniques, particularly transformer architectures. This allows them to mimic an understanding of the world from which the data

comes, which has great implications for large ecommerce companies that send LLMs large amounts of domain-specific data.



Why they matter in ecommerce

Training LLMs on massive amounts of ecommerce queries and clickstream creates a model that can mimic what people really want in a given situation.

As an example, through the clickstream, the model can start to understand that when someone searches for "butter," they mean a product like salted butter, not peanut butter. Or, when they search for "pants" in the U.S., they mean trousers, but when they search it in the UK, they may mean underwear. This can all be learned from the

clickstream, which allows you to see what people buy versus what they scroll past in each of these contexts, to fine tune the system.

LLMs also have many other applications in the product discovery space. At Constructor, we're constantly testing new ways to incorporate LLMs into our products, most recently using them to improve how Constructor filters irrelevant results for customers.

For each search query, we first rate product result sets on our base models, like Cognitive Embedding Search (CES). Those results then immediately pass through an LLM that essentially acts as a critic, or a “second-layer” model that further refines and personalizes products for the specific user.

Customers are able to track all model scoring from their merchant tools dashboard, providing transparency into how and why products are ranking certain ways across their ecommerce sites.

Lightweight Flannel Shirt (Navy Atlas Check) ×

Item data **Ranking** Recommendations API response

ATTRACTIVENESS SCORE

Based on popularity, quality of match and ML factors for the query.

Score
8.5727

RANKING FACTORS

Understand how ML optimizes attractiveness.

- Group attractiveness & Searchandizing 3.2
- Item attractiveness 3.8
- Base score 1.5



Clickstream

Clickstream data is first-party behavioral data that can be used to understand shopper behavior onsite along with how the popularity of products changes over time.

It includes, but is not limited to:

- Which pages shoppers visit first, so you can pinpoint popular entry points and fine tune the shopping experience.
- What search terms they use, so you can gauge how well your site is optimized for specific keywords
- Which products visitors click on (or don't), so you can theorize why some products are scroll-stopping, and others aren't.
- Which products visitors add to or remove from their cart, to open the door to cross-sell or upsell opportunities.
- Which products visitors finally convert on and which they abandon in cart, to keep track of best sellers and under-performers.
- How all of this changes over time, thanks to consumer shopping behaviors influenced by macroeconomic trends, evolving personal preferences, etc.

Timestamp	Event	Value
December 20, 2023 01:27:13 PM -07:00	Select	women's hiking boots
December 20, 2023 01:27:13 PM -07:00	Search submit	women's hiking boots
December 20, 2023 01:27:13 PM -07:00	Search results	297 results for women's hiking boots
December 20, 2023 01:27:20 PM -07:00	Click through	Salomon Outpulse Mid GTX Hiking Boot - Women's

With clickstream, you go beyond simple segmentation, which relies on third-party data and retargeting. Because you're collecting customer data anonymously in real-time, you provide shoppers a hyper-personalized product search and discovery experience that surfaces products that are attractive to them as early as the first session.

For example, when shoppers click through to Salomon and Mammut hiking boot PDPs, they show affinity for them. These signals boost these brands' products sitewide — to appear at the top

of further search results pages, within Collections, among recommendation pods, etc.

Clickstream data also simplifies the near-impossible task of layering your business KPIs on top of user-level personalized search results, allowing your team to further optimize product rankings sitewide. As in, not only can you display more attractive search results based on each customer's affinities, but also rank them according to the KPI you're trying to optimize for — such as revenue, profit margin, or inventory rebalancing.

Why it matters in ecommerce

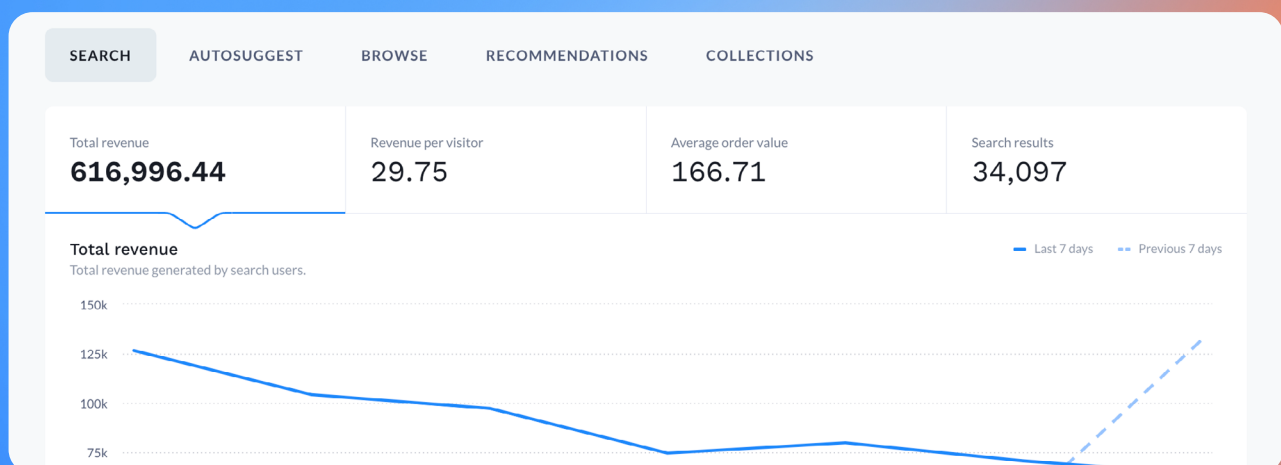
Enjoying fully automated personalization on a user-level across your entire ecommerce site has tangible business benefits, directly raising your conversion rate, AOV, and RPV. Large ecommerce companies like Petco have seen double-digit conversions from baking clickstream data into their sites, while simultaneously saving their merchandising teams countless hours of manual work.

Ranking products in the context of clickstream is also extremely important as enterprise ecommerce companies continue to grow sales channels, expanding into marketplaces. Since marketplace products aren't always proprietary, companies don't always have direct control over third-party catalog data. There's also often misalignment pre-purchase, as merchandisers take different approaches with recommendation

pods and collect clickstream only from the website, disregarding the store and other channels. Pair this with the fact that shoppers are getting better at describing what they want, moving beyond simple keywords to search.

Leveraging AI-native product discovery tools that position you to excel in this new reality and allow you to merge data across channels is paramount for business success. With Constructor, companies can do just that.

Powered by our Native Commerce Core™, our platform offers a holistic approach to product discovery that helps ecommerce companies track omnichannel user behavior from one dashboard, acting as a single source of truth to solve inputs and outputs of the entire product discovery experience.



Whether customers are searching, browsing, or getting recommendations, all products learn from each other, leveraging the same

clickstream data to provide one unified experience for site visitors and ecommerce teams alike.

Natural Language Processing (NLP)

Natural Language Processing (NLP) is a broad multidisciplinary field in AI that focuses on how to understand, interpret, and generate human language. NLP tasks can include things like speech-to-text, semantic analysis, sentiment analysis, and more.

Some examples of NLP tasks include:

- Distinguishing the meaning of the verb 'make' in 'make the grade' (achieve) vs. 'make a bet' (place)
- Determining the person or object to which a certain pronoun refers (e.g., 'she' = 'Mary')
- Extracting subjective qualities — i.e., attitudes, emotions, sarcasm, confusion, suspicion, etc.
- Understanding queries in multiple languages

Enhanced with NLP, Constructor's Native Commerce Core™ is language-agnostic, serving clients outside of English. This includes German, French, Italian, and Dutch for furniture and home living customer [home24](#). We even have a customer for whom we power search in 50 different languages across their global ecommerce footprint. Thanks to being able to identify new relationships between words, the platform's learnings evolve as the language does.

Why it matters in ecommerce

Aside from working across countries, one of the most important use cases of NLP is to allow shoppers to search in everyday language. Natural language search isn't based on keywords like traditional search engines, and it picks up on intent better since people can use connective language

to form full sentences and queries. For instance, the query "hiking boots size 8" can be formatted as "find me a pair of hiking boots in size 8 that are waterproof and mid height," and the shopper will be returned products that are attractive to them (thanks to clickstream).

Showing results for

"find me a pair of hiking boots in size 8 that are waterproof and mid height"

Sort By : Featured ▼

20 results

Categories ▼

Brands ▼

Size ▼

Sale ▲




This hyper personalization boosts consumer engagement and [helps ecommerce businesses drive the KPIs that matter.](#)

Generative AI

At a high level, generative AI refers to the use of AI models to generate content, responses, or suggestions. It goes beyond traditional rule-based systems by leveraging machine learning algorithms — often based on neural networks like transformers — for language understanding.

In the context of large amounts of clickstream data, it provides an even more tailored product search and discovery experience, where shoppers can quickly and easily find attractive products in one session — something that sounds easy in theory, yet is difficult in practice. (A [recent study](#) found that less than a third of U.S. consumers describe their most recent experience finding products on a retail website as “enjoyable.”)

This is where product search and discovery technology based on generative AI becomes so appealing.



Why it matters in ecommerce

There are a mix of customer-facing and back-end generative AI use cases in enterprise ecommerce that provide tangible business value while automating redundant merchandising tasks, ultimately saving them time, improving productivity, and allowing them to focus on strategic initiatives.

One such use case is to create dynamic landing pages based on customer on-site behavior, such as a collection of gifts for runners. Every time a customer clicks and interacts with your site, they're essentially voting for products they want to see. An AI-native search engine pays attention and dynamically re-ranks products to show the most attractive items site-wide.

A well-known outdoor retailer uses Constructor to do just this. When a user searches for "snowboarding gear," a well-rounded list of products is returned on the results page to encapsulate more than just

helmets and snowboards. This keeps the user on-site and prompts them to engage with more products, which is good for business as well — decreasing bounce rates and increasing AOV, RPV, and overall revenue.

Showing results for

"snowboarding gear"

Sort By : Featured

234 results

Categories

+ Snowboard (229)

+ Ski (154)

+ Kids (126)

+ Travel (18)

+ Hike & Camp (8)

+ Surf (6)

+ Paddle (6)

+ Overlanding (6)

+ Bike (6)

+ Climb (5)

+ Men's Accessories (1)

Brands

Search Brands...

☐ Burton (19)

☐ DAKINE (18)

☐ ARVA (12)

☐ Osprey Packs (12)

☐ Backcountry (2)

☐ Giro (11)

Show 44 more

Size

One Size

XXS

XS

S

M

L

XL

1

1.5

2

2.5

3

3.5

4

4.5

5


Show 29 more

Sale

☐ 10% and more (46)

☐ 20% and more (46)

Exclusive



Backcountry

Ski & Snowboard Boot Bag


\$159.20

~~\$199.00~~

20% off

★★★★★ (10)

Exclusive



Backcountry


Double Ski & Snowboard Rolling Bag

\$209.30

~~\$299.00~~

30% off

★★★★★ (16)




Marker

Bino Helmet - Kids'

\$45.46

~~\$64.95~~

30% off




Marker

Vijo Helmet - Kids'

\$59.46

~~\$84.95~~


30% off



WHITESPACE

Freestyle Shaun White Pro Youth Snowboard - 2024 - Kids'

\$279.00




DAKINE


Freestyle Snowboard Bag


\$52.00 - \$80.00

Up to 35% off

★★★★★ (1)



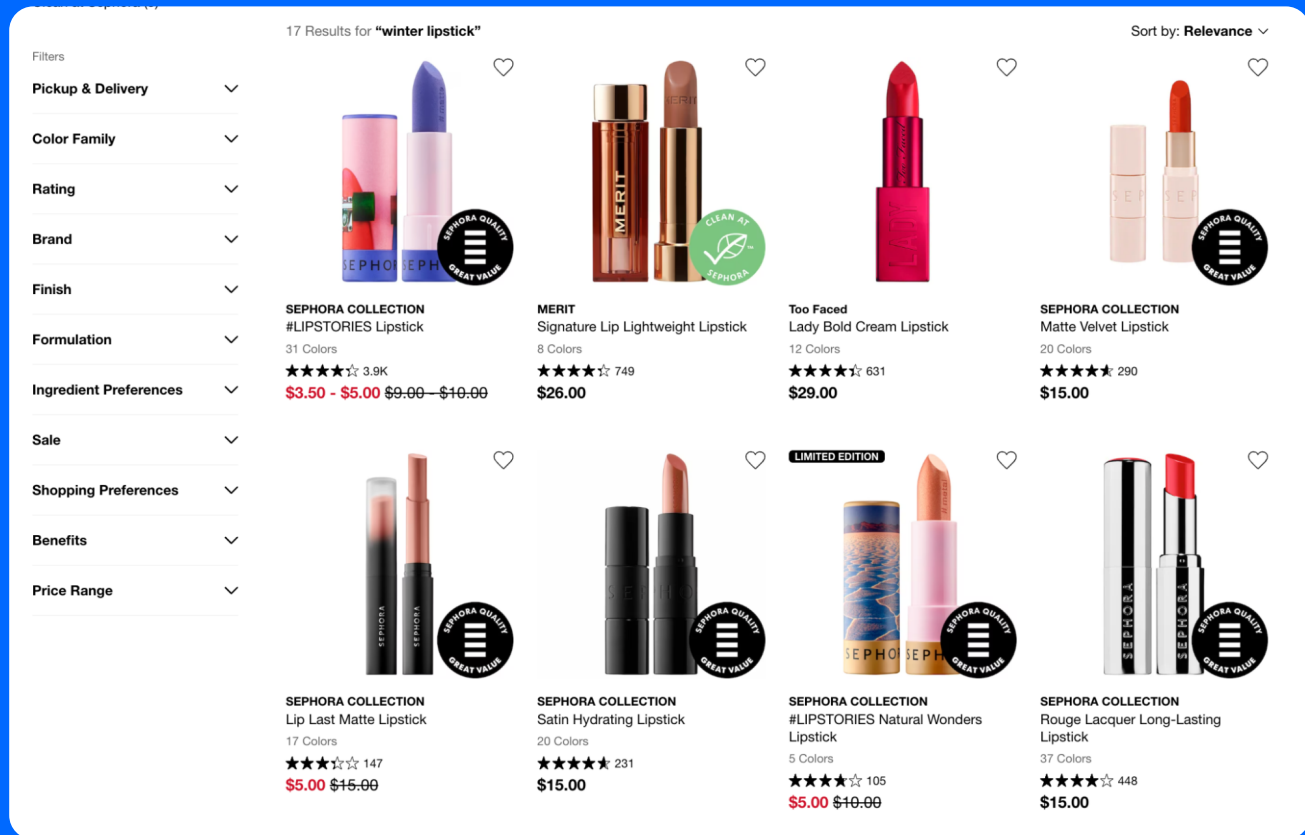





18

In the same respect, Sephora also uses Constructor. When a consumer searches for

"winter lipstick," the engine returns products that focus on dryness and richer deep colors



This technology can be integrated into existing search, such as having a pop-up with options to search intelligently, or an intent hub for [conversational commerce](#) interactions. You can also have an AI search toggle or generate AI search suggestions within the search bar.



Generative AI can also be connected to your product catalog and used to improve your search infrastructure. This is possible with [Attribute Enrichment](#). Using a mix of AI plus new innovations in machine vision and text classification, our models tag your products with new, relevant attributes and categories automatically, leveraging clickstream to prioritize the attributes that are most important to your buyers in real-time.

In the image below, Constructor's AI-generated attributes are identified with a lightning bolt. This same set of attributes may remain the same for some time, or they may change daily. They update dynamically based on individual behavioral data in conjunction with broader consumer trends.

Attribute Enrichment







AI generated attributes that make your products more discoverable for shoppers. [Learn more](#)

FILTERS

Search Item name or ID

Newly added

EDIT

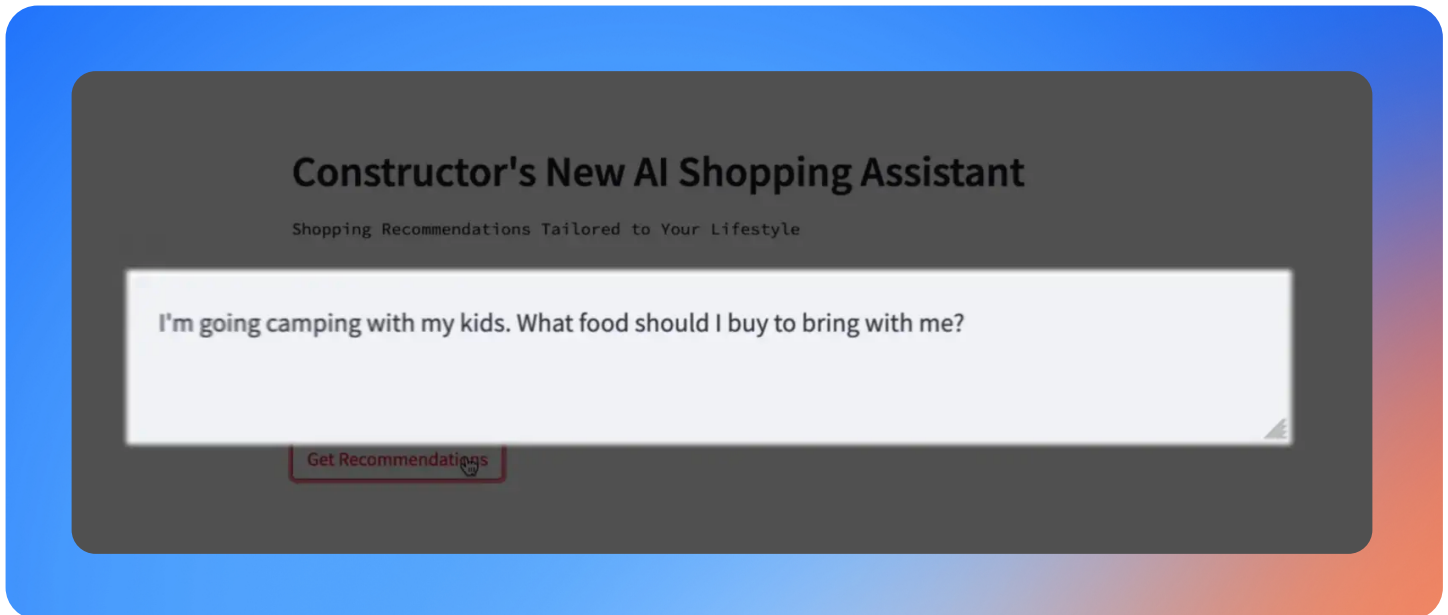
Item	Description	Category	Unit vol	Color
 Billie womens razor kit... id12345	meet the internets favorite razor winner...	Health & beauty > ⚡ razors	⚡ 1 unit	⚡ pink ⚡ minty
 Skechers men's relaxe... id12345	walk and work out in comfort and sporty sty...			⚡ grey ⚡ black
 Essie gel couture nail... id12345	my sources say this sheer tea rose pink wit...	Health & beauty > ⚡ nail polish		⚡ pink
 Fila women's v11... id12345	synthetic upper for durability; rubber...	Shoes > ⚡ casual		
 Nature valley peanut... id12345	nature valley™ crunchy bars are made with the...		⚡ 8.94oz ⚡ 6-pack	
 El Monterey signature... id12345	start your day off right with el monterey...	Food > ⚡ nail polish	⚡ 36oz ⚡ 8 pack	

This saves precious time for your team as they no longer need to manually manage your product catalog and results in customer experiences that will not only drive conversions, RPV, and AOV, but — more importantly — keep buyers coming back for more.

And finally, you can also leverage generative AI to create personalized and engaging experiences for different types of users through conversational commerce. This is seen through Constructor's recently released [AI Shopping Assistant](#).

Shoppers enter retail sites with an intent. They might not know exactly what they want to search for or where to browse, but they often have an

intention — starting a new hobby, figuring out what to cook for dinner, planning for a camping trip with kids, etc.



Below are some strategies ecommerce companies can consider when leveraging Constructor's AI Shopping Assistant to help shoppers get started:

- **Item to item.** Given an item as context, what other items might they be interested in? For alternatives, look for what you could substitute. For complementary, look for what you would add. This can solve for queries like "what are some good dips to go with some chips I just bought."
- **Shopper to item.** Given what you know about a shopper based on their affinities (clickstream), what products might they be interested in purchasing next? This can solve for queries like "show me some outfits you think I may be interested in."

- **Intent to items.** Intent-based recommendations aim to answer (most completely) which items, content, or other materials shoppers may need to meet their goal. Here you can recommend content, groups of items, expert advice, products, etc. This can solve for queries like "I want to make spaghetti carbonara tonight. What ingredients do I need?"

Overall, use of conversational commerce shortens the time it takes for shoppers to go from goal to reality, allowing them to shop all in one place and effectively addressing a fundamental problem in the way people shop (taking hours to do something simple).

The Future Looks Bright

We sat down with Constructor CEO Eli Finkelshteyn to discuss predictions on what online shopping could look like in five years when enabled by future-fit tech, like Constructor.



What does product discovery look like five years from now?

"One of the more exciting things we're starting to dig into is this idea that customers need to go offsite to conduct research for certain needs. Take, for example, someone searching for a shirt with a unique design — something that has a modern style, but is still comfortable and is great for going-out in. It's challenging to describe that need using traditional search with keywords. Or similarly, when a parent wants to select the right Star Wars Lego set for a 7-year-old child they're hoping to delight on their birthday. In these cases, customers have trouble having their needs met via traditional search on an ecommerce website, and most likely turn to Google for inspiration, reviews, and recommendations instead. This highlights a key limitation of the current on-site search experience. Ecommerce companies can hope the user comes back to the site after Googling, but it's unclear whether they will, and often requires ad dollars to remind the user to come back.

There are nascent solutions to this problem, like [Quizzes](#) — which allows an ecommerce site to ask a shopper questions about their needs to better understand them and to be able to better address them — and [AI Shopping Assistant](#), which lets users express complex needs in their own words, using LLMs and transformers to understand them in a way that was never possible before. These solutions are not yet as wide-spread or standardized as more traditional product discovery solutions like Search and Recommendations, but they are an exciting frontier that hint at how we'll be able to shop more conveniently in the coming years.

As these solutions become more standardized, we expect to see user interfaces (UIs) for them become more standardized as well. Similar to how a shopper intuitively knows they can find a search bar on an ecommerce website, the shopper will also know where they can express their longer form needs, or have an interface where the ecommerce site asks questions to help understand their needs better. We also expect to see more retail scaffolding become created around these solutions, to merchandize within them, help them play a role in SEO, and more.

Looking ahead, we want to see a continued marriage of interests between retailers and shoppers so shoppers feel like their favorite retailer is like their digital home, and they feel less need to research on sites like Google or Amazon before making a purchase. Shoppers seek a seamless experience where they know they're getting an advantage from buying from the retailer they're loyal to, while retailers aim to keep users on their site, avoiding the need for them to go off site for research, which could incur costs in winning them back.

The future landscape is likely to see a holistic integration of these technological advancements, shaping a more intuitive, personalized, and streamlined ecommerce experience for both shoppers and retailers."

What are the best ways that retailers can prepare for this transition?

"Proactively experiment with new product search and discovery technology now. Don't wait until it's a must five years down the line because then you'll be playing catch-up to your competitors.

And as you're experimenting, remember that product adoption should take a user-centric approach. Encourage customers to try out new technology via an opt-in. Don't force it on them. Let them try new experiences when they're ready

and come back only when these experiences are genuinely helpful to them. By nature, if someone visits your website, the easiest thing for them to do is what they're used to. Offer them a new and enticing way to shop — but only if they're interested.

Once they opt in, it's up to you to deliver an engaging, personalized experience. If the experience is great, they'll keep coming back, and they'll tell their friends.

As the retail landscape gets shaped, these cool experiences will cause a brand halo effect, where companies driving them are viewed as forward-thinking and ahead of customer needs. This only further expands your customer base and loyalty organically."

What will definitely NOT be a part of product discovery in five years?

"I'm not sure five years is the exact timeline, but in the foreseeable future, we'll see a shift in the way people express themselves while searching, mirroring historical shifts we've already seen.

Much like the transition from handwritten shorthand and cursive to the efficiency of typing on laptops, the advent of generative AI, transformers, and other emerging tech allow users to search using more natural, human-like language.

This progression will cause older, more rigid speech patterns and keyword-centric approaches to become a thing of the past.

When we think back to how we learned to do research online on sites like Google, there was a learning curve to learning to search with keywords. It feels natural now because we've been doing it for so long, but it's like writing in cursive: people learned to do it and kept doing it until typing — a better, easier way to write — came along. Online research and discovery will be the same way, and as people realize they can express themselves in more fine-grained and nuanced language now and that systems built on LLMs and transformers will understand them, I expect it will change the pattern of how we find things online."

After transformers, what is the next iteration of AI? Is there any nascent tech on the horizon?

"This can be taken in a couple of directions, but I'll stick to AI specifically. In essence, transformers excel at leveraging all the knowledge they have to predict a reasonable next set of words. This gives them the ability to mimic a profound understanding of and ability to predict various contexts within the universe, which works well when there's no definitive answer. The downside is that it often does not work well when something does have a definitive answer. Transformers simply generate responses that sound like they could be true, but unfortunately, sometimes are not. This has given rise to terms like "hallucination" when referring to technology like ChatGPT sometimes confidently stating answers that make great grammatical sense, but are also provably false.

A number of companies providing Foundation models like OpenAI are working on this problem, focusing on enabling AI to reason through scenarios with clear, definitive answers. Shifting

from being able to understand and generate human-sounding text that's grammatically correct to a more structured form of reasoning — like solving math problems — sounds less impressive. (We've had calculators for years. So, people are used to computers being able to solve this kind of problem.) But, this evolution lays the foundation for not only providing answers that sound human, but also providing the exact accurate response in situations where a single correct answer exists.

As AI enters this new phase, ecommerce companies will be challenged with finding the right application of them to commerce. The possibilities in minimizing returns (imagine being able to have a shopper find out everything they need to know before they buy) or in personalized, conversational commerce are pretty huge, for example. Figuring out how merchandising and searchandising keeps up with the new paradigms will also be full of exciting challenges and opportunities. A new Wild West is opening up before us, and those of us who seize the opportunities to take advantage of this new frontier today will be the ones who define the shopping and overall product discovery experience of tomorrow."

Future-Fit Product Discovery: Built for Ecommerce. Powered by AI.

The use of AI and emerging technologies within product discovery solutions makes it possible to personalize shopping experiences at scale across ecommerce sites. But not just any vendor will do. And certainly not those duct-taping solutions together to look like they're keeping up with technological advances.

Constructor is the only product discovery solution that's actually built for the new world — on a foundation of machine learning and AI rather than keyword matching. Our proprietary Native Commerce Core™ is purpose-built for ecommerce use cases, leveraging advanced algorithms, transformers, and LLMs in the context

of clickstream to power our holistic suite of AI-native merchant tools. Together, they deliver real-time, personalized experiences that decode complex patterns and understand user intent.

But don't just take our word for it. See how we measure up against other leading vendors with The Ecommerce Product Discovery RFP Template. Save hours of time with this ready-to-use spreadsheet that will help you fairly evaluate vendors' current merit and future-fitness in an AI-driven world.

[Get the template](#)