

Beyond Taxonomies: Re-Envisioning Product Discovery in E-Commerce for Industrial Goods

March 2023

nexwise is building the conversational commerce AI agent for industrial companies. It's core capability they've rolled out in January 2025 is a new approach to product search and selection.

1. Introduction

Industrial B2B companies often manage **complex product portfolios**. For example, major electronic component suppliers stock hundreds of thousands of SKUs, covering everything from sensors and IoT modules to relays. While industrials understand the importance of digitizing sales and marketing, few have embarked on extensive efforts in this area, and those that do often struggle with execution.¹ To help customers navigate this breadth, products are usually classified along multiple axes: by product family (e.g., sensors vs. actuators), by industry or sector served (e.g., automotive, aerospace, industrial automation), by application or use-case (e.g., temperature monitoring, vibration sensing), and even by business challenges addressed. With only 13 percent of industrial OEMs saying they can offer effective digital solutions with their current capabilities, effectively organizing these complex product catalogs for digital self-service has become a critical business imperative.²

In theory, this rich classification should make it easier to find the right product. In practice, traditional **static taxonomies** struggle to accommodate such multi-dimensional categorization. Each additional classification axis multiplies the complexity, and a static hierarchy cannot easily capture all the ways a customer might think of the product. Companies like SICK emphasize their extensive product portfolios that offer solutions for virtually any application or industry³ – illustrating how products span many contexts, extending far beyond a single taxonomy branch.

1.1 The Limitations of Static Taxonomies in Search

A static taxonomy typically presents a fixed tree of categories and subcategories that a user must traverse. This approach assumes the user's mental model matches the site's categorization. But

¹ [McKinsey \(2021\)](#)

² [McKinsey \(2021\)](#)

³ [FPE Automation \(2024\)](#)

industrial buyers may not know the supplier's classification terms – they think in terms of the problem they need to solve or the application at hand, not in which siloed category a product resides. The result is often a frustrating experience: users click through endless menus or apply numerous filters, only to miss relevant items because they didn't navigate the predetermined path.

In large portfolios with overlapping categories, a static structure either forces duplication of products in multiple places or makes the user choose one facet at a time, often leading to dead-ends. Moreover, maintaining such a taxonomy is labor-intensive; every time the product range expands or new applications emerge, the hierarchy must be manually updated. In fast-evolving fields like IoT or advanced sensors, static menus simply cannot keep pace.

It's important to note that taxonomies themselves remain valuable organizational structures – in fact, they're essential inputs for modern AI search systems. The limitation lies not in the taxonomies per se, but in the rigid, hierarchical navigation methods traditionally built around them. While taxonomies provide the crucial structured knowledge that our conversational commerce agent relies on, we've developed search methods that use these taxonomies as a foundation rather than as a constraint. This approach maintains all the benefits of well-organized product information while overcoming the user experience limitations of conventional taxonomy-based navigation.

1.2 Multi-Dimensional Classification Challenges

Industrial suppliers often carry extensive product portfolios spanning thousands of items and categories. Traditional static product taxonomies – fixed hierarchical category trees – struggle to accommodate this scale and diversity. As a result, both the theoretical complexity of classifying products and the practical costs to users and businesses are driven upward. A static taxonomy defines a fixed structure – typically a top-down hierarchy where each product is assigned to one category branch. While this works for small, stable catalogs, it presents several limitations in an industrial e-commerce context:

Inflexibility and Scalability: Static taxonomies struggle to scale when the product range is broad or frequently changing. Introducing new products or attributes often requires inserting new categories or restructuring the tree, a challenging task for retailers. In industrial markets, where new product lines, custom configurations, or updated standards emerge regularly, a static hierarchy can quickly become outdated.⁴

Misalignment with User Search Behavior: Industrial buyers often have varied ways of searching for the same product. Different users might categorize a product differently – one may look under its functional category, another by a specific attribute or application. Static taxonomies impose one 'correct' path to find an item, which might not match a given user's mental model. In essence, forcing users to manually traverse an exploding decision tree is diametrically opposed to good UX. They don't want to check ten boxes and drill down six levels; they want to quickly get to the handful

⁴ [Luigi's Box \(2025\)](#)

of products that meet their need. When confronted with a long list of filters, “*many users find filling out multiple filters frustrating and time-consuming.*”⁵ In the industrial context, where buyers may already be pressed for time and making high-stakes purchasing decisions, this frustration is amplified.

High Maintenance and Governance Costs: As product lines expand, the effort to maintain a static taxonomy grows significantly. Mapping thousands of products into the correct categories, auditing misclassifications, and updating entries becomes a continuous process.

2. Combinatorial Explosion & Search Costs

At the heart of the static taxonomy problem is a combinatorial explosion of categories. Each additional facet or filter dimension multiplies the number of ways products can be organized.

2.1 The Mathematics of Combinatorial Explosion

A defining characteristic of industrial product catalogs is the multitude of attributes each item can have (e.g., size, material, voltage, capacity, etc.). When attempting to classify products by every possible attribute combination, the number of categories grows multiplicatively – a classic combinatorial explosion. In mathematical terms, if a product line can be described by k independent attributes and the i th attribute can take n_i possible values, the total number of distinct product categories required to cover every combination is:

$$N_{\{\text{categories}\}} = \prod_{i=1}^k n_i .$$

This means the category count increases exponentially with the number of attributes. For example, if an industrial component comes in $k = 5$ attributes (say, length, diameter, material, color, and grade) and each attribute has about 10 possible values, the total number of unique combinations would be $10^5 = 100,000$.⁶ Even a moderately complex product line can yield tens or hundreds of thousands of potential categories, an impractical number to manage or present.

Consider a simple illustration: if an online catalog classifies products by just 3 facets (for example, industry, application, and product family), and each facet has, say, 10 options, that yields $10 \times 10 \times 10 = 1,000$ unique category combinations. Add a fourth facet with 10 options and the combinations jump to 10,000; a fifth facet pushes it to 100,000, and so on – an exponential growth.

In one real-world scenario, a site with just a few filters (5 choices in one, 7 in another, 11 in a third) would have required generating over 21,000 distinct static pages to cover every combination. This combinatorial explosion makes fully pre-defining all navigation paths impractical. More importantly,

⁵ [Hackernoon \(2025\)](#)

⁶ [Tzitzikas et al. \(2024\)](#)

it overwhelms users with complexity – the interface may present too many choices or an incomprehensible hierarchy of options.

2.2 User Search Costs

Such complexity directly translates to higher **search costs** for the user. In economic terms, search cost represents the time and effort expended by a customer to find a desired product. A complex or unintuitive taxonomy directly raises this cost. If a website forces a buyer to click through many levels of categories or wade through irrelevant options, each extra step is a cost in terms of user effort.

We can think of search cost C_{search} as a function of the number of categories or the complexity of the taxonomy, denoted T . In the worst case, a user might have to examine categories one by one to find the right placement – a linear scan through N possibilities – giving a worst-case cost that grows on the order of N (proportional to the number of categories). Formally, one could imagine $C_{search}(N) \sim c \cdot N$ in a naive scenario where c is the effort per category examined.

In a slightly better scenario, the user leverages a well-designed hierarchical structure that prunes large portions of the search space at each step, yielding complexity on the order of $\log N$ (since each choice cuts the remaining options multiplicatively). Even in that ideal best case, the search cost increases as the taxonomy expands: $C_{search}(N) \sim c \cdot \log(N)$.

In reality, user search behavior is not perfectly efficient; buyers may not know which branch of a taxonomy to follow, leading to backtracking or multiple searches. Thus, actual effort often lies between linear and logarithmic growth with respect to taxonomy size. Cognitive psychology and UX research offer clear warnings about giving users too many choices. **Hick's Law** predicts that the time required to make a decision increases logarithmically as the number of options grows. In simple terms, every additional filter or category the user has to consider adds to their cognitive load and slows their decision-making. Industrial catalogs often present dozens of filter criteria (voltage, size, sensitivity, certification, etc.), which can paralyze users.

2.3 Impact on Transaction Costs

Transaction cost in this context refers to the total cost incurred to complete a transaction. This includes the user's search cost as well as additional costs on the business side to facilitate the sale. If the website and its taxonomy allow a seamless find-and-buy experience, the transaction cost is minimal. However, when the taxonomy impedes discovery, transaction costs mount in several ways:

Longer Sales Cycles or Assistance Cost: A common outcome of high search cost is that the customer cannot self-serve. They might resort to contacting customer support or a sales representative to locate or configure the product. We can express this situation with a simple function: if $P_{fail}(T)$ is the probability the customer cannot find the product due to taxonomy

complexity T , and each failed self-service attempt triggers an assisted sales process costing (on average) C_{assist} , then an expected transaction cost function might be:

$$C_{transaction}(T) = C_{search}(T) + P_{fail}(T) \times C_{assist}$$

As taxonomy complexity grows, both terms $C_{search}(T)$ and $P_{fail}(T)$ drive transaction cost up.

Lost Sales and Opportunity Cost: Not all customers who struggle will seek help; many simply leave. For the business, an abandoned cart or an exit due to search frustration is a hidden transaction cost – the cost of lost opportunity. While not as measurable as labor, it's arguably more significant.

Error Costs: Another subtle cost arises when users think they found the right product in a complex taxonomy, but actually selected the wrong item due to confusion. The resulting return or exchange is a transaction cost for both buyer and seller.

3. The AI-Driven Industrial Solution Search

To overcome the limitations of static categorization, forward-thinking firms are turning to AI-driven search solutions that dynamically understand user needs. B2C shoppers have grown used to Google-like search that can handle natural language; increasingly, B2B buyers expect the same.⁷ Instead of forcing the user to adapt to the website's taxonomy, the idea is to have the search system adapt to the user's query – and intention. We call this approach a **Compound AI System** for industrial solution search, because it combines multiple AI components into an intelligent whole.

3.1 Three-Layer Architecture

The Compound AI System can be thought of in three key layers or modules, working together:

Proprietary Data Pipeline for Intelligent Data Ingestion: At the foundation is a robust data pipeline that pulls in all relevant product data and context. This includes structured data like specifications and inventory, as well as unstructured data such as product datasheets, manuals, application notes, and even past customer query logs. The pipeline ingests this information and enriches it – normalizing terminology, extracting key features, and tagging products with relevant attributes (like industries or use-cases they apply to). High **data quality** is essential: if product information is incomplete or inconsistent, no search algorithm can yield good results. Thus, the AI system might integrate with Product Information Management systems to ensure clean and comprehensive data. The end result is an intelligent index or knowledge graph of the product portfolio, far richer than a conventional product database.

⁷ [Optimizely \(2025\)](#)

Cognitive Architecture for Reasoning and Contextual Understanding: On top of the data layer sits the cognitive engine of the system – essentially the ‘brain’ that interprets user queries and maps them to solutions. This architecture leverages transformer-based large language models (LLMs), domain-specific ontologies, and reasoning algorithms working in concert. Unlike traditional ML approaches that required extensive feature engineering, these modern transformer architectures can understand context, nuance, and intent from natural language queries with unprecedented accuracy. When a user enters a query (which could be a few keywords or a full sentence describing a problem), the AI parses the language to understand the intent. It can decipher, for instance, that *"sensor for extremely cold warehouse"* implies a need for a temperature sensor with a certain operating range and perhaps wireless connectivity (if "warehouse" suggests difficulty in wiring). The cognitive layer uses context: it knows synonyms and industry jargon (e.g., that "PT100" is a type of temperature sensor, or that "intrinsically safe" relates to certain environments). Moreover, it can perform reasoning: if a query mixes concepts (like an industry and a problem), the AI can infer which product attributes are relevant. By employing a compound of AI techniques (transformer-based LLMs, knowledge graphs, ML ranking algorithms, etc.), the system can answer complex queries that would stump a traditional keyword search. The transformer architecture's attention mechanisms allow the system to weigh the importance of different words and concepts within a query, similar to how a human expert would prioritize certain requirements over others. It essentially builds a dynamic, on-the-fly taxonomy tailored to each query, rather than relying on a pre-built static one.

Versatile User Interface for Natural Language Queries: The final layer is what the user interacts with – a search interface that is flexible, user-friendly, and deliberately channel-agnostic. Rather than presenting a rigid filter panel from the start, the interface invites the user to type in their need in natural language (just as they would explain it to a human salesperson) through whichever communication channel they prefer. This could be an **interactive conversational UI** (think Perplexity AI), a traditional channel like email, a mobile app, or even integration with voice systems. The underlying technology functions independently of the communication channel, allowing for consistent intelligence across all customer touchpoints. For instance, a user might type or ask *"Looking for an IoT vibration sensor for predictive maintenance on motors"* in any of these channels. The system might respond with a clarifying question if needed: *"What size or mounting type do you need for the sensor?"* – imitating a helpful sales conversation. The interface handles free-form input gracefully across all channels. It can also provide rich results – not just a list of products, but possibly ranked solutions, technical Q&A, or interactive filters that appear after understanding the query. The UI thus adapts to the user, rather than forcing the user to adapt to the UI, and does so consistently regardless of how the user chooses to engage.

3.2 Benefits of AI-Driven Search

The benefits of this AI-driven ‘industrial solution search’ approach are compelling, particularly when contrasted with the old static taxonomy method:

Reduced Search Effort (Lower Search Costs): The AI system dramatically cuts down the time and clicks needed to find relevant products. Users can jump straight to a short list of suitable items by describing their need, instead of laboriously intersecting multiple filters. This reduces cognitive load and frustration – which in turn means users are more likely to stay on the site and explore. In an ideal scenario, the experience feels like getting an expert's instant recommendation out of a huge catalog, rather than hunting for a needle in a haystack.

Higher Conversion Rates and Sales Efficiency: Making it easier to find the right product inevitably improves conversion. There is evidence across e-commerce that better search leads to better sales. Research studies have shown that optimizing search functionality can lead to substantial increases in conversion rates and revenue for online retailers.⁸ When users find what fits their needs, they buy – and they often buy more. In B2B industrial sales, this could also mean more RFQs (requests for quote) or leads generated via the site. Internal data from case studies show significant uplifts: industry studies have shown that website revamps with intuitive navigation and search can yield significant improvements in order volume and line items per order.⁹ An AI-driven search can also intelligently recommend complementary products, boosting cross-sell and upsell opportunities (something static taxonomies never do).

Enhanced User Experience & Engagement: The overall UX becomes more engaging and satisfying. Instead of a daunting catalog, the site feels like it ‘understands’ the customer. This personalization fosters trust – the user feels confident the recommendations are relevant. A smooth search experience can differentiate a company in a competitive market; it's a form of digital customer service. Satisfied users are more likely to return, and in a B2B context, that could mean a loyal customer account that repeatedly orders through the convenient online channel rather than calling a competitor. As one UX principle states: a small number of proposed options + their relevance = your aim.¹⁰ It also handles nuances like typos, synonyms, or natural language questions, meaning the user doesn't hit dead-ends. This reduces the dreaded "No results found" occurrences which often cause users to bounce. Research on AI search implementations shows significant reductions in null results and corresponding increases in successful searches.

The transformer architecture powering our LLMs enables particularly sophisticated understanding of user queries. Unlike earlier NLP approaches that analyzed queries word-by-word with limited context, transformers process entire sentences holistically, detecting relationships between terms that might be separated by many words. This capability is especially valuable in industrial contexts where technical specifications might be intermixed with application needs in the same query. For example, when a user asks for *"a vibration sensor that can withstand high temperatures for monitoring bearings in steel production,"* the system recognizes both the technical requirements (temperature resistance, vibration sensing) and the application context (bearing monitoring in steel production) simultaneously.

⁸ [BetterCommerce \(2023\)](#)

⁹ [Optimizely \(2025\)](#)

¹⁰ [Shakuro \(2020\)](#)

Ability to Capture Customer Intent Data: An often overlooked benefit is that a conversational or natural language interface can capture rich data on what customers are looking for, in their own words. This can be invaluable feedback for the business – revealing new use cases or language that marketing/sales teams hadn't considered. Over time, the AI system can learn from this data to further improve results (a virtuous cycle). Static taxonomies, by contrast, give very little insight – one can only see clicks on categories, which doesn't tell you why the customer chose or abandoned that path.

4. Implementation & Business Impact

Integrating an AI-driven search system into an industrial e-commerce platform is an undertaking that can yield significant business returns. From an implementation standpoint, this typically involves layering the AI capabilities on top of existing infrastructure.

4.1 Implementation Approach

Many companies start by leveraging their current data – product databases, specifications, and support knowledge bases – and feeding it into the proprietary **AI-enabled product data pipeline** described earlier. This might be facilitated by a modern PIM or content management system to aggregate product content.

Once the data is prepared, an AI search engine – whether a commercial solution or a custom-built system – is implemented to enhance product search and discovery on the website through an interactive AI interface¹¹ or via traditional channels like email. Crucially, this doesn't necessarily mean reinventing the entire website. For example, one pragmatic approach is to use AI to parse the user's natural language query and then translate it into filters or database queries behind the scenes. This requires sophisticated transformer-based language models, extensive domain knowledge encoded in taxonomies, and considerable data expertise. The taxonomies that organize product information provide an essential structure that these models leverage to deliver accurate results. Rather than replacing taxonomies, nexwise's system uses them as a critical knowledge foundation while providing a more intuitive interface for users to access this structured information.

4.2 Business Impact

For decision-makers in industrial digital sales and UX, the takeaways are compelling. In an era where B2B buyers expect the convenience of B2C platforms, not adopting advanced search capabilities is increasingly a competitive liability. Over 80 percent of B2B sellers say they now hold their online channel to the same standards as their traditional channels¹² – meaning they aim to offer equally good (or better) product availability, pricing, delivery, and personalized service on the e-commerce

¹¹ [MIT Technology Review \(2023\)](#)

¹² [McKinsey \(2021\)](#)

site. This raising of expectations means that industrial companies can no longer treat their digital presence as secondary to in-person sales.

A significant misconception in industrial e-commerce is that complex or high-value transactions require in-person sales interactions. However, this assumption is being challenged by changing buyer preferences. According to McKinsey, despite the conventional wisdom that big-ticket sales require in-person contact, 20 percent of B2B buyers said they would be willing to spend more than \$500,000 in a fully remote/digital sales model¹³. This indicates that with the right digital experience, even substantial purchasing decisions can be facilitated through well-designed online channels.

From the business perspective, the goal of such integration is to enhance the digital sales funnel: attract, engage, convert, and retain customers more effectively online. For industrial suppliers who have historically relied on human sales teams, an AI-powered site search is a key component of digital transformation – enabling self-service for customers who increasingly prefer it. Research shows 41 percent of B2B buyers now prefer not to interact with a sales representative for research, and 64 percent research online before ever talking to a salesperson.¹⁴ This underscores that if your website can't answer their questions, they will look elsewhere (or revert to costly one-on-one sales calls).

The business impact of a well-executed AI solution search can be measured in several ways:

Increased Online Conversions and Revenue: As noted, better search directly drives higher conversion rates. More users find what they need, and find it faster – leading to more completed purchases or RFQs. Even a small uptick in conversion can translate to huge revenue gains in e-commerce. Industry analyses have found that visitors who use on-site search are far more likely to convert – often twice as much as those who just browse, because searchers exhibit intent. By capturing those intent-driven users and satisfying them, the site can lift overall sales. For instance, Algolia (a search technology provider) noted clients seeing over 350 percent ROI from investing in optimized site search features.¹⁵ These gains come not just from conversion rate, but also from larger average order values and more frequent repeat purchases, as the search system can intelligently promote relevant add-ons and re-engage users.

Improved Customer Satisfaction and Retention: In B2B, a satisfied customer tends to be a repeat customer, given the costs of switching suppliers. If your online channel consistently helps engineers or procurement officers quickly find the exact part that meets their complex requirements, they will return to it. Conversely, if they often hit dead ends and have to call support, they might try a competitor's site next time. An AI search acts like a 24/7 efficient support agent. By enabling successful self-service, you not only reduce support workload but also increase customers' confidence in your digital tools. Moreover, an AI-driven interface can be more engaging – for

¹³ [McKinsey \(2021\)](#)

¹⁴ [Optimizely \(2025\)](#)

¹⁵ [Algolia \(2024\)](#)

example, a conversational tool can make the interaction almost enjoyable, turning what used to be a tedious search into a dialogue. This positive experience builds brand loyalty.

Broader Market Reach and Solution Selling: A static catalog often caters best to customers who already know exactly what product they want (e.g., entering a part number). AI search, however, opens the door to a wider spectrum of customers, including those who know their problem but not the solution. By capturing vague queries like "*need to reduce energy consumption in my factory*" and guiding the user to a set of solutions (which might include sensors, software, services), the website can act as a consultative salesman. This enables cross-selling across product families and even up-selling complete solutions rather than single components. For the business, this means higher-value sales and the ability to market integrated solutions online – something that traditionally required a human expert to assemble. It effectively turns your catalog into a solutions advisor. In the competitive field of industrial IoT and components, this can be a major differentiator: selling value and outcomes, not just parts.

Internal Efficiency and Resource Allocation: While the primary focus of this whitepaper is on digital user experience, it is also essential to consider the impact of AI-driven search systems on internal efficiency. By automating repetitive inquiries that are traditionally addressed through the same static and manual workflows that customers experience, AI reduces the resource burden on sales engineering teams. This enables these teams to allocate more time to higher-value activities, such as developing project business, engaging in complex customer interactions, and contributing to strategic initiatives. The AI system can even surface trends (e.g., "We have a lot of searches for IoT sensors in datacenter applications in Germany"), informing R&D and product management teams where demand is emerging.

4.3 Implementation Challenges and ROI

Implementation of AI-driven search systems presents substantial technical challenges that shouldn't be underestimated. Decision-makers should be aware that success with AI search requires deep expertise in multiple domains:

Advanced data engineering for processing and harmonizing disparate product data sources.

Deep learning expertise in model orchestration and optimization techniques.

Domain knowledge in industrial product taxonomies and technical specifications.

User experience design specialized for complex search interfaces.

These challenges explain why nexwise has built a team combining AI researchers, data engineers, and industry experts focused exclusively on solving this complex problem.

From a return on investment perspective, research indicates that AI-driven search is a high-ROI initiative for digital commerce. By enhancing conversion and customer satisfaction, it drives revenue. By reducing manual effort (for both users and sales teams), it cuts costs in the long run. And by differentiating the customer experience, it helps protect and grow market share. In an era where B2B buyers expect the convenience of B2C platforms, not adopting advanced search capabilities is increasingly a competitive liability.

5. Conclusion

The industrial sales landscape is at a digital inflection point. This whitepaper has outlined how the traditional approach of static product taxonomies is ill-suited for today's vast and complex B2B portfolios. The **problem** is clear: when you have thousands of products and multiple ways to categorize them, a rigid hierarchy leads to a combinatorial overload that hampers user experience. We saw how this complexity increases cognitive load and search costs for customers, often pushing them away just when they're most ready to find a solution.

The **solution** is equally clear: harness Artificial Intelligence to create a more flexible, intelligent search experience that can handle this complexity behind the scenes. A Compound AI System for industrial solution search offers a way to have the best of both worlds – a rich, extensive product offering and an easy, intuitive way for customers to navigate it. By ingesting and understanding data, reasoning like a domain expert, and interacting naturally with users, such a system transforms the online catalog into a dynamic, responsive sales agent. The result is an online shop that can genuinely guide each buyer to the right solution, no matter how they phrase their requirement or how many products are in the back-end inventory. For decision-makers in industrial digital sales and UX, the **takeaways** are compelling. Adopting AI-driven search is not just a UX improvement; it's a direct lever for increasing conversions, customer satisfaction, and ultimately revenue, while reducing resource commitments in their technical sales functions.

Research consistently shows that improving product findability leads to significant improvements in conversion rates and order volumes. Moreover, providing a superior search experience can become a competitive advantage – it positions your brand as easier to do business with, which in B2B can be a deciding factor for customers choosing long-term suppliers.

In conclusion, static taxonomies belong to yesterday's web. The future of industrial e-commerce is **dynamic, AI-powered discovery**. Companies that embrace this will not only reduce the friction in their digital sales channels, but also unlock the full potential of their complex product portfolios. Every sensor, filter, or IoT device in your catalog has a potential buyer out there with a specific need – the challenge is connecting the two. AI can make that connection with speed and precision that humans alone cannot match at scale.