

From discovery to influence:

A guide to **AEO** and **GEO**

Practical data strategies to
empower retailers for AI search,
AI assistants and AI browsers

Contents

Executive foreword

from Author Jennifer Myers, and Executive Sponsor, Paul Longo

Dear retailer,

If you're reading this, you already feel both the pressure and promise of AI in retail. Every week, I hear from leaders like you asking:

How do we ensure our products are truly understood by AI?

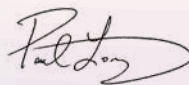
How do we keep our brand's story clear as the rules keep changing?

These aren't abstract challenges—they're our shared reality. In this paper, we explore how Answer Engine Optimization (AEO) and Generative Engine Optimization (GEO) are reshaping digital commerce. With practical strategies and real examples, our goal is to help you and your teams thrive as data quality, context, and credibility become the new currency.

No matter your role—marketing, digital, or tech—you're not facing these changes alone. My hope is these insights help you make real progress. We're building this future together—connecting your data, your brand, and the AI-powered journeys your customers are already on.



Jennifer Myers,
Principal Product Manager,
Microsoft Shopping and Copilot



Paul Longo,
General Manager, AI in Ads,
Microsoft Advertising



**We're building this
future together**

Executive summary

AI-driven shopping is transforming discovery and purchase journeys. Traditional SEO focused on clicks; now Answer Engine Optimization (AEO) and Generative Engine Optimization (GEO) determine visibility in LLM-powered ecosystems.

Why does this matter to your business?

CMOs need to ensure AI correctly understands and elevates brand differentiators. Growth and performance leaders must adapt to AI-led. Digital and e-commerce leaders should plan new metrics for AI-intermediated journeys. CTOs must make the stack AI-readable and AI-accessible. Data and analytics leaders need strategies for the 'invisible' early-funnel research happening inside AI conversations.

Across all disciplines, retailers who invest in consistent, enriched data that LLMs can trust and act on are best positioned to influence their brands positioning as shopping behaviors evolve.

Answer/Agentic Engine Optimization (AEO):

Optimizes content for AI agents and assistants (like Copilot or ChatGPT) so they can find, understand, and present answers effectively.



Generative Engine Optimization (GEO):

Optimizes content for generative AI search environments (like LLM-powered engines) to make it discoverable, trustworthy, and authoritative.



How products surface in conversational and generative ranking

To understand how retailers show up in AI ecosystems, we need to understand how these systems interpret content.

The AI shopping ecosystem is more interconnected than it first appears.

AI browsers—think Edge, Atlas, or Chrome with built-in intelligence—can ‘see’ the page you’re on in real time. They interpret content and surface helpful context while you browse.

AI assistants—like Copilot, ChatGPT, or Gemini—meet you in conversation to answer questions and help with tasks.

Assistants translate your intent:



Find the best trail shoes under \$150

into **useful, actionable guidance.**

AI agents go a step further: they don’t just advise, they act. Agents can navigate sites, fill forms, click buttons, and even complete purchases end-to-end.

These aren’t three separate worlds—they’re overlapping capabilities. A browser may include an assistant; an assistant may include agent behaviors; an agent may rely on an assistant’s reasoning.

The practical question isn’t:

“Which box does this live in?”

It’s:

“What data or content can this capability access and use—product feeds, structured markup, inventory and pricing APIs, reviews, images—and how do we make that data accurate, comprehensive, and trustworthy to deliver great shopping outcomes?”

Competition is shifting from discovery to influence (SEO to AEO/GEO)

In shorthand, the industry is talking about shifting from SEO to AEO and GEO. At Microsoft, we believe SEO and catalog investments built a foundation to expand upon in LLM-based search.

The base remains the same: up-to-date product feeds and clear, crawlable, structured content. But retailers and brands must now treat their entire catalog and site architecture as content, ensuring every product detail, benefit, and price signal is machine-readable, up to date, and context-rich.

If SEO focused on driving clicks, AEO is focused on driving clarity with enriched, real-time data. GEO helps establish credibility through authoritative voice.

Together, AEO and GEO establish your brand as a relevant, trusted shopping partner.

SEO

“Waterproof rain jacket”.

AEO

“Lightweight, packable waterproof rain jacket with stuff pocket, ventilated seams and reflective piping”.

GEO

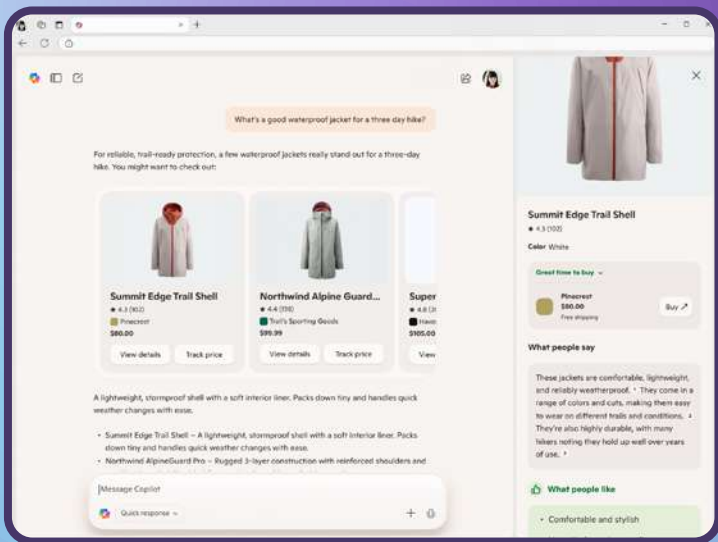
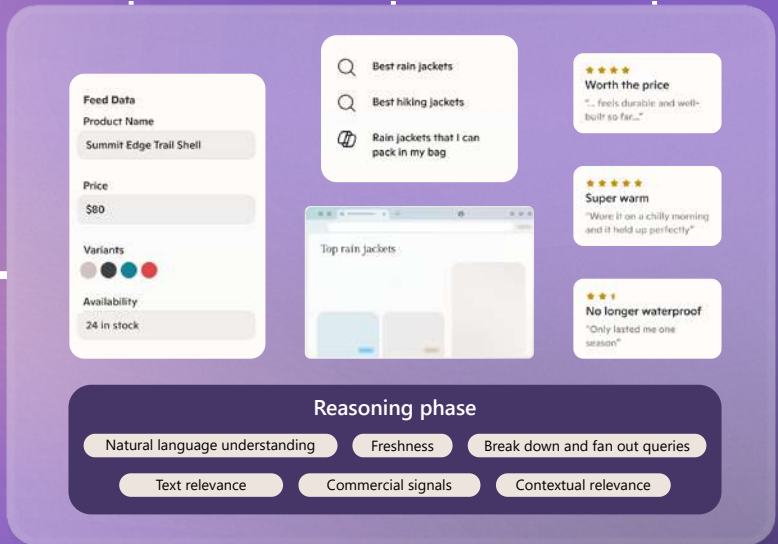
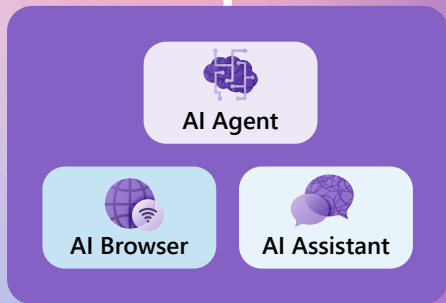

“Best-rated waterproof jacket by Outdoor magazine, no-hassle returns allowed for 180 days, three year warranty, 4.8 star rating”.

Retailers don't need to start from scratch

You can build on existing product feeds and on-site content and treat them as dynamic, data-rich assets. Enriched, real-time data will surface more often in conversational discovery, curated results, and AI summaries (“most durable”, “top value over time”), helping to establish your brand as a relevant, trusted shopping partner.

Hey Copilot, what's a good waterproof jacket for three day hike?

Let's break down how AEO and GEO apply when a user interacts with an AI assistant, browser or agent:



When a user asks for recommendations for a good rain jacket under \$200, Copilot breaks down the query in what we call the reasoning phase, using crawled web data and product feeds.

Crawled data provides:

General knowledge

“Patagonia and North Face make quality rain jackets.”

Category understanding

“Rain jackets need good waterproof ratings. Hiking jackets need to be lightweight.”

Your brand positioning

“Brand X is known for hiking equipment.”

Feeds provide:

Current prices

“Your model is \$179, competitor is \$199.”

Availability

“You have stock, competitor is backordered.”

Key specs

“Waterproof rating above 1500mm, sealed seams, GORE-TEX, PU or PVC fabric.”

AI decision:



Your product makes the top three recommendations because feeds show competitive price and in-stock status.



The user clicks through the product recommendations into your site. Here an AI Agent can provide additional context.

The AI agent sees:



Detailed reviews mentioning "great for hiking trips".



Video showing how they stay dry during rainstorms.



Current promotion: "Free water bottle with purchase".



Real-time delivery estimate: "Arrives by Friday".

The AI Agent lets the user know:



These are in stock and can arrive by Friday.
There's currently a promotion for a free water bottle.

If the user wants to **purchase**, the Agent can do so via the live website:

01

Agent adds to cart (requires functional e-commerce site).

02

Applies user's promo code (only works on live site).

03

Calculates exact shipping to user's address.

04

Completes purchase with saved payment.

05

Provides order confirmation and tracking.

Without your live site working properly, the sale fails even if your feed and crawled data were perfect.

Connecting your data to discovery

In the AI-powered shopping landscape, your business needs to show up in three distinct ways.

First, there's **crawled data**:

The information AI systems learned during training and retrieve from indexed web pages, which shapes your brand's baseline perception and provides grounding for AI responses, including your product categories, reputation and market position.

Second, there's **product feeds and APIs**:

The structured data you actively push to AI platforms, giving you control over how your products are represented in comparisons and recommendations. Feeds provide accuracy, details and consistency.

Third, there's **live website data**:

The real-time information AI agents see when they visit your actual site, from rich media and user reviews to dynamic pricing and transaction capabilities. Each data source plays a distinct role in the shopping journey—traditional SEO remains essential because AI systems perform real-time web searches frequently throughout the shopping journey, not just at purchase time, and your site must rank well to be discovered, evaluated, and recommended.



Take action: Enhance your content to improve discoverability

Now that you understand the signals AI search leverages, implement these **three strategies** to ensure your catalog is readable, your offers are accurate, and your brand is trustworthy to AI systems.



Learn more about the Microsoft Merchant Center feed schema [here](#).




01

Data structure: Make your catalog machine-readable

AI systems require structure, consistent data across all touchpoints.

Implement these technical foundations:



Schema implementation

Deploy Product, Offer, AggregateRating, Review, Brand, ItemList, and FAQ schema types.


Include dynamic fields: price, availability, color, size, SKU, GTIN, and dateModified.

Use ItemList markup for collections and category pages so AI understands product groupings.

For multi-region operations, express localized pricing and language via inLanguage and priceCurrency.

Ship JSON-LD with correct types and attributes to help browsers understand page entities.

Descriptive titles that pair product name and key differentiator (e.g., “TrailMaster 30L Hiking Jacket — Waterproof 3-Season Gear”).



Real-time synchronization

Sync price and inventory in real time between product feeds and on-site schema.

Expose dateModified and availability attributes in structured data.

Include explicit start/end dates for promotions and limited-time offers.

Maintain consistent values across feed, on-site schema, and user-facing displays.

Ensure rendered DOM contains the same facts consumers see — never serve different HTML to bots.

02 Content enrichment: Design for intent and context

AI assistants interpret queries as intents.
Structure your content to answer **real-world questions directly**.



Intent-driven product information

Write descriptive titles pairing product name with key differentiators (e.g., “TrailMaster 30L Hiking Jacket—Waterproof 3-Season Gear”).

Front-load descriptions with benefits: who it’s for, what problem it solves, what makes it better.

Add clear use-case context AI can match to queries (e.g., “best for day hikes above 40 degrees”).

Create headings and copy mirroring real-world queries, including use cases and context.

Modular, Citable Content.

Provide Q&A blocks AI can reason over and cite (“Which size should I pick?” “Is it energy efficient?”).

Display product specs as key/value pairs and feature lists.

Include comparison tables (“Model A vs Model B”) highlighting contextual differences.

Add “goes well with” data for complementary or bundled products.



Multi-modal signals

Write detailed alt text and ImageObject schema describing visuals (“green jacket with reinforced zipper and extended hood”).

Provide video transcripts parsing feature explanations.

Ensure mobile and voice experiences expose identical structured data— not just desktop HTML.



03

Trust signals: Establish authority and credibility

AI systems prioritize trustworthy sources.
Build credibility through **verified, factual content**:



Verified social proof

Include verified reviews marked with Review and AggregateRating schema.

Highlight review volume and verified purchase ratios.

Surface review sentiment that enables natural-language recommendations (“highly rated for comfort and fit”).



Authoritative brand identity

Add brand identifiers and official social/retailer links in structured data.

Link to expert reviews and articles where your products are featured.

Surface certifications, sustainability badges, and partnerships as factual entities (“Certified B Corp”, “Climate Neutral Certified”).



Content Integrity

Avoid exaggerated or unverifiable claims— AI systems penalize low-trust language.

Maintain consistent brand voice across all touchpoints.

Provide structured FAQ content and help resources that ground conversational answers.

Key takeaways

Retailers already hold most of the data signals that influence Copilot and Bing ranking—they're just not surfaced in product feeds. By enriching feeds and content assets with attributes and trust-based data, retailers can help Copilot understand not just what the product is, but why users love it and when it performs best.

This is the foundation of AI ranking readiness—a data discipline that directly impacts **discoverability** in the age of **conversational commerce**.

Have any questions?

Get in touch with your account representative or schedule a consultation with our **team of experts** to learn more. New to Microsoft Advertising? We welcome you to **join us**.



