


Key Terms of AI Brand Visibility & Selection

What this means for brands



BEFORE YOU START

Key Terms behind AI Brand Visibility & Selection

This gloassary explains the vocabulary behind AI brand visibility and selection.

Each page gives a person a clean concept to hold onto, and shows what it means for the brand.

The goal is to explain these terms clear enough that teams can speak consistently, make better decisions, and keep brand signals aligned across channels.

TERM FAMILIES

AI Systems How AI systems produce answers and what they already know when they do.

Brand Signal Quality How clearly a brand is represented in public sources and how likely AI systems are to understand it correctly.

Technical Readability Whether brand facts are machine-readable, structured, and accessible to systems that retrieve information.

Retrieval and Selection How AI systems decide what to cite, include, or leave out.

Governance Ownership, maintenance, and review of brand facts over time.

USE THIS WIKI AS

As a reference Look up individual terms when a concept comes up in a meeting, brief, or audit.

As a map Read by family to understand how the concepts connect rather than treating them as isolated definitions.

As a bridge Link back to the thesis page for the strategic argument behind why these terms matter operationally.

PLEASE NOTE

How This Wiki Was Made



ABOUT THE AUTHOR

Kerstin Ludwig runs [going→fwd](#), a consultancy for brand systems, operations, and AI-enabled transformation.

She spent 25+ years directing brand infrastructure at enterprise scale (Shopware AG, Scout24/ImmoScout24, DaWanda) building the standards and operating models that let large organizations move fast without losing coherence.

She also inspects the intersection of brand infrastructure and AI enablement, testing platforms like Rankscale, Semrush, and Otterly firsthand to understand how AI systems actually select and cite brands.

METHODOLOGY

Working on AI visibility and machine readability without using AI tools yourself doesn't hold up. Nothing here leaves the house unchecked.

Everything in this wiki was learned and tested firsthand:

- SEO fundamentals as the starting point, to understand where GEO plays by different rules
- deep dives with Rankscale, Semrush and Otterly to understand what each tool actually does differently
- positioning three brands (in different categories and stages) with the explicit goal of earning AI citations and ensuring each brand was represented correctly and not confused with others across platforms
- multiple variants tested, evaluated, discarded.

Arguments, source selection and final editorial decisions are Kerstin Ludwig's own.

AI SYSTEMS

AI Answer System



An AI answer system generates a direct response to a question rather than returning a list of links.

The user types a question; the system produces an answer built from what the model learned during training, from sources it retrieves at the moment of the query, or both.

The user may never click through to a source. If a source is cited, it was chosen by the system, not earned through ranking.

What this means for brand

Your brand may be described, compared, or recommended without anyone visiting your website. What the system says about you is shaped by what it found reliable enough to use, not by what you intended to communicate.

What it is not

A search engine, a chatbot, or a customer service tool. AI answer systems operate across contexts such as research, comparison, recommendation, and decision support.

Which part of brand is affected

Positioning, messaging, category definition, and factual accuracy of public-facing claims.

What guardrails are needed

A stable, consistently stated brand definition that can be found and used by the system. Inconsistent or missing information will be filled by the model, not always correctly.

KNOWLEDGE LINKS

→ [Google Search Central: AI Overviews](#)

→ [Kevin Indig: State of AI Search 2026](#)

FURTHER READING

→ [Search Engine Land: Mastering GEO 2026](#)

→ Tool: [Perplexity: How it works](#)

SEE ALSO

LLM · Retrieval · Citation Readiness · Grounding

AI SYSTEMS

Large Language Model (LLM)



A large language model is a type of AI system trained on large volumes of text. It learns statistical patterns in language and uses them to generate coherent, contextually appropriate responses.

It does not look things up in real time unless specifically designed to do so. It works from what it learned during training, and that training has a cutoff point.

What this means for brand

An LLM knows about your brand only to the extent that accurate, clear information about it appeared in training data before the cutoff. What was vague, inconsistent, or absent will remain absent, or be invented.

What it is not

A database, a fact-checker, or a live index of the web. An LLM generates plausible text; it does not verify.

Which part of brand is affected

Any claim the system makes about who you are, what you do, and how you compare to others.

What guardrails are needed

Factual consistency across all public-facing sources. The more your brand is described the same way in verifiable sources, the less the model has to fill in gaps.

KNOWLEDGE LINKS

- [Anthropic: What is Claude?](#)
- [Simon Willison: Using LLMs](#)
- [Learn Prompting: Open guide](#)
- Tools: [IBM A list of large language models \(LLMs\)](#)

SEE ALSO

Training Data · Knowledge Cutoff · Model Knowledge

AI SYSTEMS

Training Data



Training data is the body of text (articles, websites, documents, books, and other sources) that an AI system was exposed to during its development.

The model learns language patterns, facts, and associations from this data. It does not store individual documents; it absorbs patterns and encodes them internally.

What this means for brand

If your brand appeared clearly and consistently in sources before the training cutoff, the model has a basis for describing you. If it appeared rarely, inconsistently, or not at all, the model either won't know you or will construct something from adjacent patterns.

What it is not

A database you can update or a crawl you can control. Training data is fixed at the point of training; you cannot send a correction after the fact.

Which part of brand is affected

Brand recognition, category placement, factual accuracy, and competitive positioning.

What guardrails are needed

A sustained, consistent public presence in citable sources. Press coverage, third-party references, award records, and partner pages contribute to what a model may eventually learn.

KNOWLEDGE LINKS

- [Common Crawl: The open dataset behind many LLMs](#)
- [The Markup: What goes into LLM training data and why it matters](#)
- Tool: [Hugging Face: Datasets overview](#)

SEE ALSO

Knowledge Cutoff · Model Knowledge

AI SYSTEMS

Knowledge Cutoff



Every AI model has a point in time after which it has no knowledge of events, publications, or changes.

This is the knowledge cutoff.

A model released in 2026 may have a training cutoff of mid-2024 or earlier. Anything that happened after that date, e.g. a rebrand, a new product, a leadership change, does not exist for the model unless it retrieves that information live.

What this means for brand

A rebrand completed after the cutoff will not be reflected in the model's base knowledge. The model may still describe the old brand, the old name, or the old positioning as if the change never happened.

What it is not

The model's release date. Models are released months or years after training was completed. A recent release does not mean recent knowledge.

Which part of brand is affected

Any brand change, positioning shift, new product, new leadership, or corrected information that post-dates the cutoff.

What guardrails are needed

If your brand has changed significantly, ensure current information is available in retrievable, citable sources so retrieval-enabled systems can surface the updated version.

KNOWLEDGE LINKS

- [Simon Willison: LLM knowledge and cutoffs](#)
- [MIT Technology Review: The problem with AI knowledge cutoffs](#)
- [Stanford HAI: AI Index Report](#)
- Tool: [OpenAI Playground](#)

SEE ALSO

Retrieval · Grounding · Governance · Selection Loop

AI SYSTEMS

Model Knowledge



Model knowledge is what the AI system carries from its training: facts, associations, and patterns encoded during the learning process. It is not a lookup table.

When a model describes your brand, it is expressing what it statistically expects to be true based on training patterns, not reporting a verified fact.

What this means for brand

Model knowledge shapes what the system says about you without any live source being consulted. If the model absorbed weak, wrong, or absent signals about your brand, that becomes its working assumption.

What it is not

A record, a citation, or a fact. Model knowledge is inference, not retrieval.

Which part of brand is affected

Default descriptions, category associations, competitive framing, and perceived reputation.

What guardrails are needed

Strong, consistent signals in public sources — especially press, industry directories, published case studies, and third-party references.

KNOWLEDGE LINKS

- [Anthropic: How Claude works](#)
- [Simon Willison: Understanding LLMs](#)
- [AI Snake Oil: Princeton researchers on AI limitations](#)
- Tools: [Profound](#), [Otterly](#), [Peec](#)

SEE ALSO

LLM · Training Data

AI SYSTEMS

Prompt



A prompt is the input a user gives to an AI system: a question, instruction, or request. The system generates a response based on the prompt, its model knowledge, and any retrieved sources.

The specific wording of a prompt can significantly affect what the system produces, which means the same brand can be described differently depending on how the question is asked.

What this means for brand

Different users ask about your brand in different ways. A system that cannot anchor on a clear entity definition will produce different answers to “who is Brand,” “what does Brand do,” and “compare Brand to Competitor.” Inconsistency in the output reflects inconsistency in the signals the system has to work with.

What it is not

A search query, command, predictable trigger. The same prompt, in different systems or days, can produce different answers.

Which part of brand is affected

Consistency of brand representation across different query types and use contexts.

What guardrails are needed

A canonical definition strong enough to anchor multiple prompt types. External sources that confirm the same description regardless of the angle from which the question is asked.

KNOWLEDGE LINKS

- [Anthropic: Prompt engineering guide](#)
- [OpenAI: Prompting best practices](#)
- [Ethan Mollick: One Useful Thing on prompting](#)
- Tools: [Learn Prompting: Open guide](#), [Google AI Studio](#)

SEE ALSO

Canonical Definition · Entity

AI SYSTEMS

RAG / Retrieval-Augmented Generation



Retrieval-Augmented Generation is a method that combines a language model's trained knowledge with live retrieval from external sources. When a user asks a question, the system first searches for relevant content, then uses what it finds to generate a grounded answer.

This allows AI systems to produce responses informed by current, specific, or verifiable sources — rather than drawing on training data alone.

What this means for brand

If your brand has clear, retrievable, well-structured pages, a RAG-enabled system is more likely to ground its answer in accurate information. If your pages are thin, vague, or machine-unreadable, the system may still generate an answer drawn from inference.

What it is not

A search engine, a database lookup, or a guarantee of accuracy. RAG improves grounding but depends entirely on the quality of what it retrieves.

Which part of brand is affected

All content that could be retrieved and used to inform a generated answer.

What guardrails are needed

Pages that state facts clearly, are technically crawlable, and prioritize factual precision over persuasive language. The quality of what RAG retrieves determines the quality of what it says.

KNOWLEDGE LINKS

- [Lewis et al.: RAG for Knowledge-Intensive NLP Tasks](#)
- [NVIDIA: What is retrieval-augmented generation?](#)
- [Microsoft Learn: RAG and AI search](#)
- Tool: [You.com](#)

SEE ALSO

Retrieval · Grounding · Citation Readiness · LLM

BRAND SIGNAL QUALITY

Brand Signal



A brand signal is any piece of public information about a brand that AI systems can detect, interpret, and use when constructing an answer or making a recommendation.

Signals can be structured (Schema.org markup, Wikidata entries, JSON-LD) or unstructured: press coverage, third-party mentions, consistent descriptions across channels. Their value is determined by quality: how consistent, specific, and verifiable they are.

What this means for brand

Every public-facing piece of content is a potential signal. The stronger, clearer, and more consistent those signals are across sources, the more accurately AI systems can describe, classify, and include the brand in relevant answers. Weak, inconsistent, or absent signals leave the model to infer.

What it is not

A campaign message or creative asset. Brand signals work because they are factual and consistent, not persuasive.

Which part of brand is affected

All external-facing content: about pages, press materials, partner pages, structured data, social profiles, awards, external references.

What guardrails are needed

Audit signals regularly. Identify where they are inconsistent, absent, or contradictory. Prioritize consistency of name, category, description, and facts across all public sources.

KNOWLEDGE LINKS

- [Google Search Central: How Search works](#)
- [Kalicube: Brand signals for AI systems](#)
- [Grounding Page Standard v1.6](#)
- Tool: [Scrunch.ai](#)

SEE ALSO

Entity · Canonical Definition · Grounding · Citation Readiness

BRAND SIGNAL QUALITY

Entity



In the context of AI systems, an entity is a clearly identifiable thing: a brand, organization, person, product, service, method, event, or concept.

AI systems build associations around entities. The clearer and more consistent the signals around an entity, the more reliably a system can represent it.

What this means for brand

Your brand needs to be recognizable as a distinct entity, not just a set of keywords. That means a consistent name, a defined category, a clear description, and enough signals to separate you from competitors, generic concepts, or similar names.

What it is not

A keyword, a tag, or a product category. An entity is a specific, identifiable thing with its own properties, relationships, and boundaries.

Which part of brand is affected

Brand naming, category definition, about page, structured data, and all external references.

What guardrails are needed

One consistent name used everywhere. A clear, stable description that specifies what the brand is, does, and serves. Explicit disambiguation where confusion is likely.

KNOWLEDGE LINKS

- [Google Search Central: Understanding entities](#)
- [Schema.org: Entity types reference](#)
- [Jason Barnard: Entity SEO explained](#)
- Tool: [Wikidata](#)

SEE ALSO

Entity Confusion · Canonical Definition · Knowledge Graph

BRAND SIGNAL QUALITY

Canonical Definition



A canonical definition is the official, stable, short description of what an entity is.

It specifies what the entity is, what it does, and who it serves — in factual language, without qualification or marketing framing.

It is the sentence or paragraph that defines the entity clearly enough to be used as a reference by any system or reader.

What this means for brand

If you do not define your brand in one consistent, authoritative statement, AI systems will construct a definition from whatever signals they can find. That definition may be inaccurate, incomplete, or blended with competitors.

What it is not

A tagline, a mission statement, or a value proposition. A canonical definition is factual and stable, not aspirational or campaign-specific.

Which part of brand is affected

About page, homepage description, metadata, structured data, and all external brand references.

What guardrails are needed

One agreed definition used consistently across all channels and maintained over time. Any change should be deliberate, documented, and propagated.

KNOWLEDGE LINKS


- [Schema.org: Description property](#)
- [Jason Barnard: How to define your entity](#)
- [Google Search Central: Structured data overview](#)
- Tool: [Otterly](#)

SEE ALSO

Entity · Grounding Page · Structured Data

BRAND SIGNAL QUALITY

Hallucination



Hallucination is when an AI system generates information that is plausible but factually wrong.

It is a structural property of how language models work.

When a model lacks sufficient grounding for a fact, it fills the gap with what is statistically likely based on surrounding patterns.

The result can sound authoritative and be entirely incorrect.

What this means for brand

A model may state the wrong founding year, wrong product category, wrong headquarters, wrong leadership, or wrong ownership. Not as a guess, but as a confident statement! If no one checks, the error persists.

What it is not

A lie, a bias, or a deliberate error. Hallucination is a consequence of insufficient factual grounding, not intent.

Which part of brand is affected

Any factual claim the model makes, including history, structure, offerings, people, awards, and partners.

What guardrails are needed

All key facts made explicit, citable, and consistent across sources. The more clearly a fact is stated in verifiable places, the less the model needs to infer it.

KNOWLEDGE LINKS

→ [AI Snake Oil: On hallucination and AI limitations](#)

→ [Anthropic: Reducing hallucination](#)

FURTHER READING

→ [MIT Technology Review: AI hallucination explainer](#)

→ Tool: [Peec](#)

SEE ALSO

Grounding · Citation Readiness · Brand Sentiment

BRAND SIGNAL QUALITY

Entity Confusion



Entity confusion occurs when an AI system mixes up two entities: a brand confused with a competitor, a person confused with a namesake, or a product confused with a category.

Systems that cannot clearly distinguish one entity from another will blend descriptions, misattribute claims, or apply the wrong associations.

What this means for brand

If your brand name is similar to a competitor, a generic category, or a company in another market, the model may describe the wrong one, or combine both into a single incoherent description.

What it is not

A deliberate misrepresentation. Entity confusion is a disambiguation failure, not an editorial judgment.

Which part of brand is affected

Brand naming, positioning, differentiation, and category definition.

What guardrails are needed

An explicit statement of what the brand is not. Clear distinction from similar entities. Specific language that gives the system enough signal to separate entities reliably.

KNOWLEDGE LINKS

- [Google Search Central: Disambiguation in structured data](#)
- [Kalicube: Entity disambiguation](#)
- [Search Engine Journal: Entity SEO guide](#)
- Tool: [Rankscale](#)

SEE ALSO

Entity · Canonical Definition

BRAND SIGNAL QUALITY

Non-Inclusion



Non-inclusion is when a brand, product, or person is simply absent from an AI-generated answer where they should logically appear. The entity is not confused with another; it is ignored.

This happens when the system does not find sufficient, clear, trusted signals to include the entity in its response.

What this means for brand

You may not be hallucinated or misrepresented. You may simply not exist in the model's answer, even when the question is directly relevant to your category or capability. If you do not appear, you lose not just visibility but measurable attribution.

What it is not

A ranking penalty or a deliberate exclusion. Non-inclusion is a signal-strength problem, not a judgment.

Which part of brand is affected

Discoverability, category presence, and competitive positioning in AI-generated recommendations or comparisons.

What guardrails are needed

Clear, citable, structured brand information in sources the model can retrieve. Third-party coverage that confirms the entity's existence and relevance in its category.

KNOWLEDGE LINKS

- [Kevin Indig: AI visibility and inclusion signals](#)
- [Semrush: AI visibility tracking](#)
- [Search Engine Land: GEO and AI inclusion](#)
- Tool: [Ahrefs Brand Radar](#)

SEE ALSO

Retrieval · Citation Readiness · Grounding · Topical Authority

BRAND SIGNAL QUALITY

English Retrieval Bias



Many AI systems rely disproportionately on English-language sources, both in training data and in retrieval.

Brands and expertise primarily documented in other languages are at a structural disadvantage: underrepresented in model knowledge and less likely to be retrieved or cited.

What this means for brand

A strong German-language presence does not automatically translate into AI visibility in global systems. A brand well-known in DACH but poorly documented in English may be invisible to systems that route through English sources first.

What it is not

A deliberate policy. English retrieval bias is a structural consequence of how training data is assembled and how retrieval systems are weighted.

Which part of brand is affected

Global discoverability, investor and partner perception, and international market positioning.

What guardrails are needed

A parallel English-language version of key brand pages: a functionally complete, citable English description of who the brand is and what it does.

KNOWLEDGE LINKS


- [Common Crawl: Language distribution in training data](#)
- [MIT Technology Review: Language bias in AI](#)
- [JHU: Multilingual artificial intelligence often reinforces bias](#)
- [Stanford HAI: AI Index Report](#)

SEE ALSO

Retrieval · Citation Readiness

BRAND SIGNAL QUALITY

Disambiguation



Disambiguation is the practice of making clear what an entity is and what it is not. AI systems that encounter similar names, overlapping categories, or generic terms need explicit signals to separate one entity from another.

Without disambiguation, systems blend descriptions, misattribute claims, or default to the more prominent entity with similar name.

What this means for brand

If your brand name resembles a competitor, a generic category, or a company in another market, the system may describe the wrong one or construct a composite that represents neither accurately. Disambiguation is not defensive; it is foundational. With clear boundaries, AI systems represent you more reliably.

What it is not

A legal disclaimer or a footnote. Disambiguation is structural: it belongs in your canonical definition, your about page, your structured data, and anywhere your entity is described publicly.

Which part of brand is affected

Brand naming, category definition, canonical definition, and all pages that describe what the brand is and does.

What guardrails are needed

State explicitly what your brand is not. Name the categories, products, or entities it could be confused with. Keep this consistent across all public sources.

KNOWLEDGE LINKS

- [Grounding Page Standard v1.6: Goal 04](#)
- [Rankscale: Grounding Page module](#)
- [Schema.org: sameAs property](#)
- [Google Knowledge Graph Search API](#)

SEE ALSO

Entity · Entity Confusion · Canonical Definition · Grounding Page

TECHNICAL READABILITY

Retrieval



Retrieval is the process by which an AI system, before or during generating an answer, searches for relevant external sources and incorporates what it finds. Not all AI systems do this.

Those that do (sometimes called retrieval-augmented systems) can draw on current information beyond the training cutoff. What they retrieve and how they weight it depends on the quality, clarity, and structure of what they find.

What this means for brand

If your brand has clear, well-structured, citable pages, a retrieval-enabled system is more likely to find and use accurate information. If your pages are marketing-heavy, vague, or poorly structured, the system may retrieve them but not cite them.

What it is not

A guarantee of inclusion. Retrieval means the system looks, not that it uses what it finds.

Which part of brand is affected

Website structure, about pages, factual content, structured data, and press materials.

What guardrails are needed

Pages that state facts clearly, use consistent naming, are technically accessible to crawlers, and prioritize clarity over persuasion.

KNOWLEDGE LINKS

- [Perplexity: Retrieval-augmented generation](#)
- [Simon Willison: How retrieval-augmented AI works](#)
- [Kevin Indig: State of AI Search 2026](#)
- [Bing Copilot](#)

SEE ALSO

Grounding · Citation Readiness · Machine Readability

TECHNICAL READABILITY

Grounding



Grounding is the use of clear, verifiable sources to reduce the gap between what an AI system generates and what is actually true.

A grounded system draws on explicit facts rather than inference. For a brand, grounding means ensuring that the facts a system needs (name, category, description, key claims) are available in sources the system can find, read, and trust.

What this means for brand

Grounding is not a technology task. It is a content and governance task. The brand team decides what the authoritative facts are and ensures they are published in a form AI systems can use.

What it is not

SEO, prompt engineering, or a product you buy. Grounding is a discipline applied to how facts are structured and maintained.

Which part of brand is affected

All public-facing factual content, especially about pages, press materials, structured data, and external references.

What guardrails are needed

A clear definition of what is true about the brand, where it is published, and who is responsible for keeping it current.

KNOWLEDGE LINKS


- [Microsoft Bing: Evolving role of the index](#)
- [Grounding Page Standard v1.6](#)
- [Grounding Page Project: AI Grounding Playbook](#)
- Tool: [The Stacc AI Content Quality Checklist](#)

SEE ALSO

Citation Readiness · Governance · Grounding Page · Citation Drift

TECHNICAL READABILITY

Citation Readiness



Citation readiness describes how likely a page is to be used as a source in an AI-generated answer.

A page is citation-ready when its facts are clearly stated, verifiable, well-structured, and easy to extract. A page that is persuasive but vague, or visually rich but structurally opaque, may be retrieved but not cited.

What this means for brand

Being findable and being citable are different things. A brand page optimized for human persuasion is often structured in a way that makes it hard for a system to extract a specific, attributable fact.

What it is not

A guarantee of citation. Citation readiness increases the probability of inclusion, but it does not control the output.

Which part of brand is affected

Website copy, about pages, press materials, structured data, and all external-facing factual content.

What guardrails are needed

Clear, factual language. Short, extractable statements. Consistent use of brand name, category, and description. Marketing superlatives kept out of factual sections.

KNOWLEDGE LINKS


- [ALPHAXXXX: AI Citation Readiness Guide](#)
- [Stay Citable: AI Citation Readiness Checklist](#)
- [TheStacc: AI Citation Readiness Checklist: 31 Points](#)
- Tool: [AlvsRank: AI Citation Readiness Checker](#)

SEE ALSO

Grounding · Retrieval · E-E-A-T · Chunking

TECHNICAL READABILITY

E-E-A-T



E-E-A-T stands for Experience, Expertise, Authoritativeness, and Trustworthiness. It is the framework Google's quality raters use to assess whether content deserves to be relied on.

It is not a score, and no system publishes an E-E-A-T number. It describes a set of signals: who wrote this, what do they demonstrably know, who else confirms it, and can the claims be checked. The same signals decide whether an AI system treats a page as evidence or as one more opinion.

What this means for brand

E-E-A-T is the reason a trust gap forms. A brand can be well known and still be judged unsuitable as a source without named author, demonstrable experience, and external confirmation. Systems then use the brand as a topic and someone else as the proof.

What it is not

A ranking factor you configure. A checklist you complete once. E-E-A-T is not implemented; it accumulates through named authorship, verifiable claims, and third-party corroboration over time.

Which part of brand is affected

Author attribution, credentials, case evidence, press coverage, external references, every factual claim published without source.

What guardrails are needed

Named authors with real credentials on substantive content. Claims tied to evidence that can be checked. A defined owner for proof, and a review rhythm for credentials, references, figures etc.

KNOWLEDGE LINKS

- [GSC: Creating helpful, reliable, people-first content](#)
- [Google: Search Quality Rater Guidelines \(PDF\)](#)
- [GSC: Google Search and AI-generated content](#)
- Tool: [Google Search Console](#)

SEE ALSO

Citation Readiness · Trust Gap · Grounding · Governance

TECHNICAL READABILITY

Machine Readability



Machine readability refers to how well a page's content can be accessed, parsed, and interpreted by automated systems like crawlers, indexers, AI retrievers.

A page that looks good visually but encodes its content in ways that are opaque to machines (images, JavaScript rendering, poor HTML structure) is functionally invisible to those systems regardless of how good the content is.

What this means for brand

The brand information you want AI systems to find needs to be in the HTML, not just in the visual layer. A well-designed page with brand facts embedded in graphics or loaded through JavaScript may not be readable at all.

What it is not

Accessibility for screen readers, though the two overlap. Machine readability is about structural clarity for automated systems.

Which part of brand is affected

Website structure, technical implementation, and all pages carrying key factual brand content.

What guardrails are needed

Key facts in HTML text. Semantic headings. A canonical URL. Structured data where useful. No critical content behind login walls or in non-crawlable formats.

KNOWLEDGE LINKS


- [Google Search Central: How Googlebot crawls](#)
- [web.dev: Core Web Vitals and crawlability](#)
- [Ahrefs: Technical SEO and machine readability](#)
- Tool: [Screaming Frog: Technical SEO audit guide](#)

SEE ALSO

Structured Data · Citation Readiness · Chunking

TECHNICAL READABILITY

Chunking



Before a system can retrieve a page, it splits it into pieces. These pieces are called chunks, and retrieval works on chunks, not on whole pages.

A chunk is what gets compared against a query and what ends up quoted in an answer. If a fact only makes sense across three paragraphs, the chunk that contains one of them carries an incomplete version of it.

What this means for brand

A well-argued page can be retrieved and still misrepresent the brand, because the piece that was extracted did not carry the qualification, the context, or the subject of the sentence. Facts that survive extraction are stated completely in one place: subject, claim, and qualifier in the same passage.

What it is not

A reason to write in fragments. Chunking rewards passages that are self-contained, not text that has been stripped of connective reasoning. A section can be well written and still be extractable.

Which part of brand is affected

Page structure, section length, heading logic, and every core fact spread across a build-up rather than stated in one place.

What guardrails are needed

Key facts stated in full in a single passage, with the brand named. Semantic headings that mark where a topic begins and ends. No essential claim that only resolves after two paragraphs of context.

KNOWLEDGE LINKS

- [Anthropic: Contextual Retrieval](#)
- [Pinecone: Chunking strategies for LLM applications](#)
- [Microsoft Learn: RAG and AI search](#)
- Tool: [Screaming Frog SEO Spider](#)

SEE ALSO

Machine Readability · Citation Readiness · Retrieval · RAG

TECHNICAL READABILITY

Structured Data



Structured data is a standardized way of encoding information so machines can read it directly.

Rather than a paragraph saying “this company was founded in Berlin in 2012,” structured data encodes that as a set of labeled fields (organization, location, founding date) that a system can parse without interpretation.

It is typically added to a page’s HTML as JSON-LD.

What this means for brand

Structured data removes ambiguity. It tells the system explicitly what type of entity the page describes and what the key facts are. It does not guarantee citation, but it reduces the work the system has to do to extract and classify information.

What it is not

A replacement for good page content. Structured data and readable HTML reinforce each other; one without the other is incomplete.

Which part of brand is affected

Technical implementation of brand pages, especially about pages, product pages, and event pages.

What guardrails are needed

Structured data must reflect current facts. Out-of-date structured data is worse than none because it states incorrect facts as explicit machine-readable claims.

KNOWLEDGE LINKS


- [Google Search Central: Structured data documentation](#)
- [Schema.org: JSON-LD examples](#)
- [Ahrefs: Structured data guide](#)
- [Moz: Structured data explainer](#)

SEE ALSO

Machine Readability · [Schema.org](#)

TECHNICAL READABILITY

Schema.org



Schema.org is an open vocabulary for structured data, developed collaboratively by Google, Microsoft, Yahoo, and Yandex.

It defines a shared set of entity types, e.g. Organization, Person, Product, Event, and hundreds of others, along with the properties that can be associated with each.

Using Schema.org vocabulary means systems across different platforms can interpret the data consistently.

What this means for brand

Schema.org gives you a standardized language for stating: this is an Organization, its name is X, it is based in Y, it offers Z. Using it correctly helps systems classify your brand and connect it to the right entity type.

What it is not

A requirement, a guarantee, or an SEO ranking factor on its own. Schema.org is a shared vocabulary, not a publishing platform and not an endorsement system.

Which part of brand is affected

Technical metadata on brand pages; primarily about pages, product pages, and team pages.

What guardrails are needed

Schema type must match what the entity actually is. Properties must be accurate and current. Inconsistency between visible page content and structured data undermines the signal.

KNOWLEDGE LINKS


- [Schema.org: Organization type reference](#)
- [Google Rich Results Test](#)
- [Google Search Central: Introduction to structured data](#)
- [Google Schema Markup Validator](#)

SEE ALSO

Structured Data · Entity

TECHNICAL READABILITY

Knowledge Graph



A knowledge graph is a structured database of entities and the relationships between them.

Google maintains one of the most prominent: it connects people, organizations, products, places, and concepts through verified or inferred relationships. When a system draws on a knowledge graph, it works from a pre-structured map of how entities relate, not generating those relationships from scratch each time.

What this means for brand

Inclusion in Google's Knowledge Graph, or similar systems, gives a brand a more stable and consistent representation in AI-mediated contexts. It is harder to hallucinate or confuse an entity that is already explicitly mapped.

What it is not

Something you submit to directly. Knowledge graphs are built and maintained by their operators. Brands influence inclusion through Wikipedia presence, structured data, and consistent external signals; they do not control the graph.

Which part of brand is affected

Entity definition, Wikipedia presence, external references, and structured data.

What guardrails are needed

Consistent entity naming and description across all external sources that knowledge graph operators are likely to draw from.

KNOWLEDGE LINKS


- [Google: Knowledge Graph overview](#)
- [Wikidata: Open knowledge graph](#)
- [Kalicube: Getting into the Knowledge Graph](#)
- Tool: [neo4j](#)

SEE ALSO

Entity · Canonical Definition · Structured Data

TECHNICAL READABILITY

Knowledge Panel



A Knowledge Panel is Google's structured entity card, appearing on the right side of search results when the Knowledge Graph can identify a brand, organization, or person with confidence. It displays name, description, category, location, and related entities.

For AI systems, it is one of the strongest signals that a brand is a recognized, verifiable entity.

What this means for brand

Incorrect or outdated panel information (e.g. wrong description, wrong category) is presented as fact and spreads faster than corrections reach it.

What it is not

Something you create or submit directly. Google builds Knowledge Panels from its own analysis of available signals. Brands can influence them through consistent structured data, Wikipedia presence, and verified social profiles — but cannot control them.

Which part of brand is affected

Brand entity recognition, factual accuracy in AI-generated descriptions, and competitive positioning in knowledge-based search results.

What guardrails are needed

Monitor your Knowledge Panel regularly. If information is wrong, use Google's feedback mechanism to request corrections. Maintain consistent entity data across structured data, Wikipedia, Wikidata, and verified social profiles.

KNOWLEDGE LINKS

- [Google: About Knowledge Panels](#)
- [Kalicube: How to get and fix your Knowledge Panel](#)
- [Google: Claim your Knowledge Panel](#)
- [Tool: Google Search](#)

SEE ALSO

Knowledge Graph · Entity · Canonical Definition · Structured Data

RETRIEVAL AND SELECTION

Query Fan-Out



When an AI system receives a question, it does not run a single search. It runs multiple queries in parallel: breaking the original prompt into sub-queries, retrieving evidence for each, and synthesising the results into one answer.

This process is called query fan-out. The user sees one question and one answer. Behind that, the system may have run five, ten, or more searches the user never sees.

What this means for brand

Being findable for your brand name is not enough. AI systems also retrieve evidence for adjacent queries: category, competitors, market, use case. If your brand does not appear in the results of hidden sub-queries, it will not appear in the final answer — even when the original question was directly relevant to what you do.

What it is not

A single keyword search. Query fan-out means the system is building an evidence base. Optimising for the visible prompt alone misses most of what determines inclusion.

Which part of brand is affected

Category presence, thought leadership content, third-party coverage, and all content that positions the brand in its market.

What guardrails are needed

Think beyond branded search. Ensure your brand is present and citable in content covering category, competitors' territory, and the questions your buyers ask before they know your name.

KNOWLEDGE LINKS


- [Microsoft Bing: Evolving role of the index](#)
- [Aiso: What query fan-out means for brands](#)
- [Kevin Indig: State of AI Search 2026](#)
- Tool: Self-Check category query without brand name

SEE ALSO

Retrieval · RAG · Non-Inclusion · GEO · Topical Authority

RETRIEVAL AND SELECTION

Topical Authority



Topical authority describes whether a domain covers a subject coherently and completely, or touches it once.

It is built through connected content across a subject area: the category, the adjacent questions, the objections, the comparisons. Not through one page that names the brand.

What this means for brand

Query fan-out means a system retrieves against sub-queries the user never typed: the category, the alternatives, the use case, the objection. A brand with content only about itself has nothing to return for them. It is not that the brand loses the comparison. It is not in the retrieval set the comparison is built from.

What it is not

Content volume. Publishing frequency. A domain authority score. Topical authority is coverage of a subject with substance, not output measured by quantity.

Which part of brand is affected

Category-level content, thought leadership, FAQ coverage, comparison and alternative pages, and everything that describes the market rather than the company.

What guardrails are needed

Content that covers the category, not only the brand. The questions buyers actually ask, answered in full. Coherent internal linking so the coverage reads as one connected body of work rather than scattered posts.

KNOWLEDGE LINKS

- [GSC: Creating helpful, reliable, people-first content](#)
- [Semrush: Topical authority](#)
- [Ahrefs: Topical authority](#)
- Tool: [Semrush](#)

SEE ALSO

Query Fan-Out · Non-Inclusion · GEO · Citation Readiness

RETRIEVAL AND SELECTION

AI Overview / AI-Generated Answer



An AI Overview is the generated summary that appears at the top of some search results, or as the primary response in an AI assistant, before any list of links.

The user reads the generated answer. The sources it draws from may appear as citations or may not be visible at all. This is increasingly the first point of contact between a question and an answer.

What this means for brand

If someone searches for a category your brand operates in, the AI Overview may name competitors, explain the category, and make recommendations without mentioning you. What appears in that response is determined by what the system found, trusted, and chose to include.

What it is not

A featured snippet, an ad slot, or a place you apply to appear. AI Overviews are generated, not curated by humans.

Which part of brand is affected

Category positioning, competitive visibility, and brand recognition in AI-mediated discovery.

What guardrails are needed

A clear, citable definition of what your brand does in the categories most relevant to your buyers. Presence in external sources the system would consider authoritative.

KNOWLEDGE LINKS

- [Google: AI Overviews](#)
- [Bing Copilot: How answers are generated](#)
- [Kevin Indig/AirOps: The Impact AI Overview on SEO](#)
- Tool: [Two Squares How AI Answers Are Built](#)

SEE ALSO

Retrieval · Citation Readiness

RETRIEVAL AND SELECTION

AI Citation / AI Mention



An AI citation occurs when an AI system links directly to your URL as a source inside a generated answer. An AI mention occurs when your brand name appears in the answer text without linked source.

Both signal that the system recognizes your brand, but they carry different weight. A citation indicates the system trusts your content enough to surface it as evidence. A mention indicates awareness, not necessarily trust.

What this means for brand

You can appear in an AI answer without your content being cited. Citations mean your pages were retrieved, evaluated, and deemed credible enough to reference. Building citation readiness means closing the gap between being known and being sourced. A citation is a potential Click-Through that can be measured.

What it is not

The same thing. A mention without a citation is not a citation. The distinction matters for how to measure AI visibility and what to do.

Which part of brand is affected

Website credibility signals, structured data, factual content quality, and external references; all factors that influence whether a system cites or mentions you.

What guardrails are needed

Track both separately. A high mention rate with a low citation rate signals a trust gap. Prioritize grounding, citation readiness, and external source quality to shift mentions into citations.

KNOWLEDGE LINKS

- [Rankscale: AI Citation](#)
- [Rankscale: AI Mention](#)
- [Grounding Page Standard v1.6: Goal 05](#)
- Tool: [Otterly.ai](#)

SEE ALSO

Citation Readiness · Trust Gap · Grounding · Non-Inclusion · Brand Sentiment

RETRIEVAL AND SELECTION

Brand Sentiment



Brand sentiment describes how a brand is characterised when an AI system mentions it, not whether it is mentioned.

A brand can be named accurately, cited correctly, and still be positioned as the expensive option, the legacy choice, or the one with the difficult onboarding. The facts hold. The framing does not.

What this means for brand

What a brand says about itself carries little weight against the aggregate of what others said. Systems draw that characterisation from review sites, forum threads, comparison articles, and press coverage, weighted by whatever they found most often.

What it is not

A reputation score, and not the same as accuracy. A false statement is a hallucination. An unfavourable but factually correct statement is sentiment, and it cannot be corrected. It can only be outweighed.

Which part of brand is affected

Reviews, third-party comparisons, forum and community discussion, press tone, customer support visibility, and every public account of what it is like to work with the brand.

What guardrails are needed

Monitor characterisation alongside presence and accuracy. Treat recurring negative framing as a product or service signal, not a communications problem. Build the countervailing evidence in public, in sources the systems already draw on.

KNOWLEDGE LINKS

- [Semrush: AI brand sentiment and perception analysis](#)
- [Search Engine Land: Mastering GEO 2026](#)
- Tools: [Profound](#), [Otterly](#), [Peec](#)

SEE ALSO

Citation Readiness · Trust Gap · Grounding · Non-Inclusion

RETRIEVAL AND SELECTION

Zero-click



Zero-click describes the growing pattern of users receiving a direct answer from an AI system or search result without clicking through to any external source.

The answer is generated or summarized at the point of query. No site visit occurs. This is now standard behavior across AI answer systems and an increasing share of traditional search results.

What this means for brand

Your website may never be visited even when your brand is mentioned in an AI-generated answer. Visibility, reputation, and category positioning are shaped by what AI systems say. No clicks means no GA4 signal. That is exactly where attribution breaks.

What it is not

A temporary trend, a technical failure, or a problem that more SEO will fix. Zero-click is a structural shift in how audiences access information, accelerated by the rise of AI answer systems.

Which part of brand is affected

All brand touchpoints that previously relied on clicks to create impressions. Brand equity is now also built (or damaged) in AI-generated answers that no one clicks through from.

What guardrails are needed

Invest in citation readiness and grounding alongside click optimization. Ensure AI systems have accurate, structured, citable information so that zero-click answers about your brand are correct — even when no visit occurs.

KNOWLEDGE LINKS


- [SparkToro: Zero-click search research](#)
- [Microsoft Bing: Evolving role of the index](#)
- [Microsoft Bing: Optimizing content for inclusion in the era of AI](#)
- Tool: [Semrush AI Toolkit](#)

SEE ALSO

Retrieval · Citation Readiness

RETRIEVAL AND SELECTION

Trust Gap



The trust gap is the distance between being mentioned and being cited. An AI system that names your brand in an answer but does not link to or draw from your content knows you exist; it does not yet trust your pages as a source.

This gap is not a visibility failure; it is a credibility failure. The system has found you but judged your content insufficient to cite.

What this means for brand

A high mention rate with a low citation rate is not a good sign. It means AI systems recognize your brand but do not consider your pages reliable enough as evidence. You are present in the model's world but absent from its sourcing. A gap between mention and citation is a gap between visibility and measurable impact.

What it is not

Same as non-inclusion. Non-inclusion means you don't appear at all. The trust gap means you appear only as a name, not as source.

Which part of brand is affected

Website content quality, factual precision, structured data, and external authority signals. All factors that determine whether content crosses the threshold from known to trusted.

What guardrails are needed

Audit the gap between mentions and citations regularly. Prioritize factual, well-structured content on key brand pages. Build external authority through press coverage and third-party references that confirm what your brand claims about itself.

KNOWLEDGE LINKS


- [Rankscale: Trust Gap](#)
- [Kevin Indig: The ghost citation problem](#)
- [Implementation: Grounding Page Standard v1.6: Goal 05](#)
- Tool: [Peec.ai](#)

SEE ALSO

AI Citation · Citation Readiness · Non-Inclusion · Grounding · E-E-A-T

RETRIEVAL AND SELECTION

GEO / Generative Engine Optimization



Generative Engine Optimization is the practice of structuring content so that AI answer systems are more likely to include, cite, and accurately represent an entity in generated answers.

It extends the logic of traditional search optimization into a context where the output is a generated response, not a ranked list of links.

What this means for brand

SEO optimizes for clicks. GEO optimizes for inclusion. The goal is not to be the top link; it is to be the source an AI system trusts enough to draw on when constructing an answer about your category.

What it is not

A replacement for SEO, a manipulation technique, or a guarantee of citation. GEO is a discipline, not a workaround. AEO and GEO are used interchangeably by most practitioners; where they are distinguished, AEO refers to direct-answer formats and GEO to generated, synthesised answers.

Which part of brand is affected

All external-facing content, especially factual pages, category descriptions, and thought leadership content.

What guardrails are needed

Content that prioritizes clarity and factual precision. GEO and marketing copy serve different functions; conflating weakens both.

KNOWLEDGE LINKS

- [Princeton University: Generative Engine Optimization](#)
- [Search Engine Land: GEO — How to win AI mentions](#)
- [Clairon AI: How to Do GEO: Complete Implementation Guide](#)
- Tool: [AI Citation Readiness Checker | AlvsRank](#)

SEE ALSO

Citation Readiness · Grounding · Retrieval · Topical Authority · Selection Loop

RETRIEVAL AND SELECTION

Selection Loop



The Selection Loop is a going→fwd framework for how brands become selectable, by people and by AI systems. It starts before SEO, GEO, or content optimisation.

Two cascades run in parallel. In the brand cascade, relevance creates substance, substance enables structure, structure creates visibility. In the source cascade, connection generates proof, proof requires order, order creates machine readability. They meet in selection. People trust brands when things make consistent sense. Machines select brands when evidence is clear, ordered, readable.

What this means for brand

A brand can rank and still fail to appear in AI-generated answers. It can be cited and still be represented wrongly. The Selection Loop locates the break: in the offer, in the proof, in the sources coherence, or in the technical layer. Each break needs a different repair.

What it is not

A funnel, a customer journey, a recommendation engine, or a piece of software. Also no process to run or product to install.

Which part of brand is affected

Positioning, category definition, evidence layer, source coherence across owned and third-party channels, technical readability on top.

What guardrails are needed

Four questions, answered honestly. Does the brand solve a clear, relevant job. Can its core claims be supported by public, citable evidence. Do website, profiles, policies, reviews, and external mentions tell the same truth. Can systems identify the entity, category, offer, and proof without guessing.

KNOWLEDGE LINKS

- [going→fwd: Selection Loop](#)
- [Whitepaper 01: From AI in Marketing to AI-Relevant Brands](#)
- [Princeton University: Generative Engine Optimization](#)

SEE ALSO

GEO · Grounding · Citation Readiness · Entity

GOVERNANCE

LLMS.TXT



llms.txt is a file that a website can publish to give AI systems explicit instructions about how their content should be treated.

It works analogously to robots.txt for crawlers: indicating what content the site owner considers suitable for AI use, what should be excluded, and how content should be attributed.

It is not a universally enforced standard and compliance varies by system.

What this means for brand

llms.txt gives brand teams a mechanism to signal intent: what they want AI systems to use and how. It does not guarantee compliance, but it is part of a responsible governance posture.

What it is not

A blocking mechanism, a legal instrument, or a replacement for robots.txt. llms.txt communicates preference; it does not enforce it.

Which part of brand is affected

Content governance, legal and compliance considerations, and technical brand management.

What guardrails are needed

A clear internal decision about which content the brand wants AI systems to use, and a governance process for keeping llms.txt current as content changes.

KNOWLEDGE LINKS

- [llmstxt.org: Specification](https://llmstxt.org/)
- [Cloudflare: AI crawlers and llms.txt](#)
- [Simon Willison: On llms.txt](#)
- [WordLift: Generate llms.txt](#)

SEE ALSO

Citation Readiness · Grounding · Retrieval

GOVERNANCE

Citation Drift



Citation drift describes the way the set of sources behind an AI-generated answer changes over time, without the question changing.

A brand cited in an answer this month may be absent from the same answer next month. Nothing was penalised. The system re-retrieved, weighted differently, or found something newer.

What this means for brand

AI visibility is an outcome that gets recalculated at every query. A single check confirms nothing beyond that one moment. What matters is whether the brand appears repeatedly, across systems and over time, and the sources carrying it are stable enough to keep being found.

What it is not

Not the same as brand drift. Citation drift concerns the changing source set behind an answer. Brand drift concerns changing meaning, behaviour, or representation. The two reinforce each other, but they have different causes and different remedies.

Which part of brand is affected

Source stability, freshness of published facts, external coverage, and assumption that results observed once will still be there later.

What guardrails are needed

Measure as a time series, not as a screenshot. A monthly review reveals a pattern that isolated checks cannot.

KNOWLEDGE LINKS

- [going→fwd: Citation Drift](#)
- [Kevin Indig: The ghost citation problem](#)
- [Microsoft Bing: The evolving role of the index](#)
- Tool: [Rankscale](#)

SEE ALSO

Governance · Grounding · AI Citation · Trust Gap

GOVERNANCE

GOVERNANCE



Governance, in the context of brand and AI visibility, refers to the rules, ownership, and processes that ensure facts about a brand stay accurate, current, and consistent over time.

It answers who owns the official brand definition, who can update it, when it is reviewed, and how changes are tracked.

What this means for brand

A Grounding Page or about page without governance is a fact at a point in time. Brands change: leadership, product, positioning, ownership... Without governance, outdated facts persist in citable sources, and AI systems continue to use them.

What it is not

Bureaucracy or a legal function only. Brand governance is an operational discipline; the process by which what is true about a brand stays aligned with what is published about it.

Which part of brand is affected

All external-facing factual content, structured data, llms.txt, about pages, partner references, and press kits.

What guardrails are needed

A named owner, a review cadence, a process for propagating changes across channels when facts change, and documentation of what changed and when.

KNOWLEDGE LINKS

- [Frontify: AI for Brand Governance](#)
- [Meltwater: Brand Governance Key Principles & Best Practices](#)
- [Martech Alliance: Brand governance guide](#)
- Tool: [Frontify Brand Audit How-To Guide](#)

SEE ALSO

Grounding · Canonical Definition · Citation Readiness · Citation Drift

The thesis behind these terms

This document defines the vocabulary.
It does not make the argument.
The argument lives in the companion
whitepaper.

It maps the structural shift:

why using AI in marketing is no longer the
hard part, and why becoming a brand that
AI systems find, understand, and recommend
is.

It explains:

- why the consideration and decision loop
has shifted to AI mediation
- what the four structural risks are for
brands that have not addressed this
- why this is not an optimisation problem
- what a different starting point looks like

The terms in this document are the working
vocabulary for that argument.

Read together, they give teams both the
conceptual frame and the operational
language to act on it.

Read the whitepaper:

→ going-fwd.com



From Marketing with AI to **AI-Relevant Brands**



**Why Brands Today Must Be Relevant to
Both Humans and AI Systems**

KI-IM-MARKETING.AT × GOING-FWD.COM



LET'S STAY IN THE (SELECTION) LOOP!

Kerstin Ludwig of going→fwd adds this perspective with more than 25 years of experience in brand systems, operations and governance. Her focus: setting up organisations so that brand, decisions and AI deployment play together consistently.

Anja and Julie Teßmann of SISTERHOOD Berlin & Vienna (aka the KI-Sisters) help organisations strategically embed AI in marketing — from brand management through content and processes to concrete AI workflows. From over 100 AI workshops they know the challenges marketing teams face today.

Together we guide organisations on the path from diagnosis through strategy to implementation — so that your brand remains visible, verifiable and selectable not only for humans, but also for AI systems.



Kerstin Ludwig

Founder, going→fwd
Directing Brand, Systems



Julie Teßmann

Co-Founder SISTERHOOD
AI Innovation & Implementation Specialist



Anja Teßmann

Co-Founder SISTERHOOD
Marketing Expert, AI Manager

KI-IM-MARKETING.AT



GOING-FWD.COM

A whitepaper by going→fwd in cooperation with Ki im Marketing / SISTERHOOD:



KI im Marketing by SISTERHOOD
c/o 1010 Works GmbH
Seitenstettengasse 5/37
www.ki-im-marketing.at