



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

TECHNICAL

Structured Data & Knowledge Graphs: Making Your Business Machine-Readable

By Scott Tischler · July 28, 2026 · 8 min read

Here is the uncomfortable truth I share with most executives I advise: an AI answer engine cannot recommend a business it cannot understand. Before ChatGPT, Gemini, or Perplexity will name you, they have to resolve who you are, what you do, and whether you are the same entity that appears on your website, your LinkedIn, your reviews, and a dozen directories. If those signals conflict, the model hesitates. And a model that hesitates recommends someone else.

Structured data is how you remove that hesitation. It is the difference between hoping a machine infers what you mean and telling it plainly, in a format built for machines to read.

What does "machine-readable" actually mean?

A human reading your About page understands from context that "TSC" and "Tischler Strategy Consulting" and the phone number in the footer all refer to one company. A machine does not get that for free. It sees strings of text and has to guess whether they connect.

Machine-readable means you have encoded those facts explicitly, in a structured format, so there is nothing left to guess. The dominant standard for this is **schema.org**, a shared vocabulary maintained by Google, Microsoft, Yahoo, and Yandex. You deliver it using **JSON-LD** (JavaScript Object Notation for Linked Data), a small block of code in the head of your page.

Here is the core idea in one sentence you can quote: ***structured data turns your web page from prose a machine has to interpret into facts a machine can simply read.***

You do not need to be a developer to grasp the strategy. You need to know which facts matter, that they must be consistent everywhere, and that your web team (or your CMS plugin) is emitting them correctly.

Which schema types matter most for AI visibility?

Not every schema type moves the needle. For most businesses, a handful carry the weight. Here is how I prioritize them.

Schema type	What it declares	Why it matters for AI
Organization	Your company as an entity: name, logo, URL, contact, social profiles	The anchor for your entire entity; ties everything together
Person	A named individual: role, employer, credentials, affiliations	Critical for founder-led and expert authority (E-E-A-T)
Product / Service	What you sell: name, description, price, availability	Lets engines match you to purchase-intent questions
FAQPage	Question-and-answer pairs on a page	Feeds directly into how answer engines phrase responses
Review / AggregateRating	Ratings and testimonials	Supplies the trust signals models weigh before recommending
LocalBusiness	A physical or service-area business with location and hours	Essential for "near me" and city-specific queries

If I had to pick two to start with, they would be **Organization** and **Person**. They establish the identity that every other signal hangs from. Get those wrong and the rest is built on sand.

What is a knowledge graph, and why should you care?

A **knowledge graph** is a database of entities and the relationships between them. Instead of storing pages, it stores things: people, companies, products, places, and the connections among them. Google's Knowledge Graph is the reason a search for a public figure returns a panel with their bio, employer, and related people rather than just ten blue links.

Three knowledge graphs matter for AI visibility:

- **Google's Knowledge Graph** powers Google's answer surfaces and AI Overviews.
- **Wikidata** is an open, structured database that feeds Wikipedia and is heavily used in AI training and retrieval.
- **LLM training corpora**, the vast text collections models learn from, function as an informal knowledge graph. When a model has seen your entity described consistently across many trusted sources, it "knows" you.

Your goal is to exist as a clean, consistent **entity** in all three. When you do, an answer engine does not have to reconstruct who you are on the fly. It already has a resolved node to draw from, and resolved nodes get cited.

How do sameAs links tie your identity together?

The single most underused property in structured data is **sameAs**. It is an explicit statement that "this entity is also the entity found at these other URLs." You use it inside your Organization or Person schema to point at your authoritative profiles.

A strong sameAs set typically includes:

- Your LinkedIn company and personal profiles
- Your verified social accounts (X, Facebook, Instagram, YouTube)
- Your **Wikidata** entry, if one exists
- Crunchbase, industry association pages, or reputable directory listings
- Wikipedia, if you are notable enough to have a page

Each sameAs link is a thread stitching your scattered web presence into one recognizable entity. The more authoritative threads you connect, the more confidently a machine resolves you as a single, real, trustworthy thing. This is **entity disambiguation**, and it is quietly one of the highest-leverage things you can do.

Why entity consistency (NAP) makes or breaks you

NAP stands for **Name, Address, Phone number**, but the principle extends to every fact about your business. Entity consistency means your core details are identical everywhere they appear, character for character.

This sounds trivial. In practice it is where most businesses leak trust. "Suite 200" on the website and "Ste. 200" on Google. An old phone number lingering in a directory. A slightly different legal name on one profile. Each mismatch is a small reason for a machine to doubt that these listings describe the same entity.

Models weigh **corroboration**. When ten independent sources agree on your name, location, and offering, that agreement is itself a trust signal. When they disagree, the model has to pick a version or discount you entirely. I have watched businesses climb in AI visibility for no reason other than cleaning up inconsistent listings that had drifted over years.

The practitioner rule is simple: **decide the canonical version of every fact once, then enforce it everywhere.**

A practical playbook for business owners

You do not have to do this all at once. Sequence it.

1. **Define your canonical entity.** Write down the exact legal name, DBA, address, phone, and one-sentence description. This is your single source of truth.
2. **Implement Organization and Person schema** on your homepage and About page, with a complete **sameAs** list. Most modern CMS platforms and SEO plugins support this without custom code.
3. **Audit your citations.** Check that your NAP matches across Google Business Profile, LinkedIn, industry directories, and review sites. Fix the drift.
4. **Add FAQ and Review schema** where you have genuine questions and genuine testimonials. Never fabricate either; models and platforms increasingly detect and penalize fake markup.
5. **Claim a Wikidata entry** if you qualify, and connect it via sameAs. This is often the cleanest on-ramp into the structured knowledge that AI systems trust.
6. **Validate.** Use a schema validation tool to confirm your markup is error-free, and re-check after any site redesign, because redesigns are where structured data silently breaks.

Do this and you have done something most of your competitors have not: you have made yourself legible to the machines now deciding who gets recommended.

The bottom line

Structured data is not a technical nicety you delegate and forget. It is the infrastructure of being recommended in an AI-mediated market. The businesses that win the next few years will be the ones a machine can identify without hesitation, resolve to a single trustworthy entity, and cite with confidence. That legibility is buildable, and it starts with telling the machines plainly who you are.

Key takeaways

- 1 AI answer engines can only recommend businesses they can unambiguously resolve as a single, trustworthy entity.
- 2 Schema.org markup delivered via JSON-LD turns interpretable prose into machine-readable facts; Organization and Person schema are the highest-priority starting points.

- ② Knowledge graphs (Google's, Wikidata, and LLM training corpora) store entities and relationships, not pages; your goal is to exist cleanly in all three.
- ② The sameAs property stitches your scattered profiles into one recognizable entity and is one of the most underused, highest-leverage properties available.
- ② NAP consistency across every listing is a trust signal; drift and mismatches give models a reason to doubt and discount you.
- ② Never fabricate FAQ or review markup; validate your structured data and re-check it after every site redesign.

Frequently asked questions

What is structured data in plain terms?

Structured data is code you add to your web pages that states facts about your business in a format built for machines to read. It uses the shared schema.org vocabulary, usually delivered as JSON-LD, so search engines and AI systems can understand who you are without having to guess from context.

Do I need to be a developer to use schema markup?

No. Most modern content management systems and SEO plugins can generate schema.org markup for you without custom code. As a business owner, your job is to define the correct facts and enforce their consistency; implementation is a task your web team or a plugin can handle.

What is a knowledge graph?

A knowledge graph is a database that stores entities such as people, companies, and places along with the relationships between them, rather than storing web pages. Google's Knowledge Graph and the open Wikidata database are prominent examples, and AI models effectively build an informal one from the text they are trained on.

Why does NAP consistency matter for AI recommendations?

NAP stands for Name, Address, and Phone number, and consistency means these details are identical everywhere they appear. AI systems treat agreement across many independent sources as a trust signal, so inconsistent listings give a model a reason to doubt you are a single, legitimate entity.

What does the sameAs property do?

The sameAs property explicitly tells machines that the entity on your page is the same as the ones on your other authoritative profiles, such as LinkedIn, Wikidata, or verified social accounts. It connects your scattered web presence into one recognizable entity, which helps AI systems resolve and cite you confidently.

Which schema types should a business implement first?

Start with Organization and Person schema, because they establish the core identity every other signal hangs from. From there, add Product or Service, FAQPage, Review, and LocalBusiness schema where they genuinely apply to your business.

Will structured data alone get me recommended by AI?

Structured data is necessary but not sufficient. It makes you legible and resolvable, which is the foundation, but you still need genuine authority, consistent citations, and real reviews for a model to choose you over a competitor. Think of it as the infrastructure that lets your other signals count.

How does Wikidata fit into AI visibility?

Wikidata is an open, structured database that feeds Wikipedia and is widely used in AI training and retrieval. A clean Wikidata entry, connected to your site via sameAs, is often the cleanest way to place your business into the structured knowledge that AI systems already trust.

About the author. Scott Tischler is the Founder & Chairman of Alrecommend.ai and a practitioner-authority on AI search and Answer Engine Optimization. With 20+ years in marketing technology — including American Express, MetLife, and UBS — and executive study at Wharton, Harvard, Yale, and Oxford, he helps businesses become the ones AI recommends.