Schema For AI Citations: How To Become A Trusted Source via @sejournal, @lorenbaker

Schema markup won’t make AI systems cite you directly, but it does help search engines like Google, Bing, and answer engines like ChatGPT understand who you are, verify your claims, and decide whether to feature you. I’ve supported schema for years, and although the benefits took time to show, the efforts I made on client sites before the rise of AI are now paying off. During my first SEJ Pro Keystone session on content structure and schema, I focused on validation because it’s a step many overlook.

Think of your schema as your online dating profile; it’s your first chance to introduce yourself to search engines about your content, your company, and the people behind it. If what’s written in your profile doesn’t match the reality when someone visits, it’s like catfishing. And nobody wants to deceive Google or ChatGPT.

Schema Doesn’t Create Trust, It Makes Trust Verifiable

Schema is far less a ranking switch than a trust builder. Think of it as a tool that tells you apart from similar entities, confirms your claims, and helps you qualify for rich results. Platforms are often more nuanced about this than the SEO industry. For example, Microsoft has mentioned that schema helps its LLMs understand your content, while Google’s guidance states that structured data isn’t necessary for AI features but still recommends using it as part of a comprehensive SEO strategy. OpenAI also emphasizes that structured product data feeds what ChatGPT presents to

shoppers.

So, it’s helpful, but not a magic solution. Having markup alone doesn’t guarantee results. Search engines compare your schema to the actual page, your business feed, reviews, and what people are saying online. The most effective schema implementation isn’t about using the most properties but about ensuring that all key properties are consistent and say the same thing everywhere they appear.

The 4 Surfaces That Have To Agree

Every business model I covered came back to the same four-part structure:

- **The webpage:** The facts a person can see and read.
- **The schema:** The machine-readable version of those facts.
- **The platform of record:** Your Google Business Profile if you are local, your Merchant Center feed if you sell products.
- **Third-party corroboration:** Reviews, directories, professional profiles, publications, studies.

When those four sources agree, you provide a single, reliable reference point for the engines. When they don’t, complex

markup can undermine your credibility instead of enhancing it.

This is why paying attention to small details is crucial. For example, using “Suite” on the website but “Ste” in the markup creates a mismatch, which can be costly. The same applies to inconsistencies in a business name, phone number, opening hours, SKU, price, job title, or author name that change across surfaces without explanation.

Local Schema

In any location, everything hinges on the entity. Without it, there is no local schema because other properties need a reference point. During the session, I broke this down into nine pillars, starting from the entity and its stable ID, then covering NAP, geo coordinates, hours, services, conversion actions, reviews, sameAs links, and overall consistency.

The most common source of confusion involves the service area. On a Google Business Profile, “service area” indicates where you dispatch or deliver, not all the locations you operate. In schema, areaServed means all the areas you serve. These are different fields that answer different questions. In a city with multiple locations and neighborhoods, reusing generic markup across all can eliminate the distinctions that help an engine recommend the appropriate one.

I asked ChatGPT for a physical therapist in Hoboken who was open that day, and it returned a client’s brand with hours and services such as orthopedic rehab, sports injuries, and post-op rehab. These services weren’t on that location’s landing page; they were in the schema and on the Business Profile. Two out of three sources contained the facts, enough for ChatGPT to decide.

This doesn’t mean you should hide useful information in your structured data. Visible content remains the foundation. When your sources align, an engine can accurately identify your

business and match it to specific needs without guesswork.

Commerce Schema

The same core structure still applies, just with a different platform of record. For ecommerce, think of your Merchant Center feed as your main source of truth, and try to match it as closely as possible, field by field. Pay special attention to key product details like name, description, brand, images, SKU, MPN, GTIN, and a consistent ID. After that, include offer details such as price, currency, condition, availability, shipping costs, delivery time, return policies, and product variants.

Availability issues can really impact your revenue. When your schema flips to OutOfStock the moment inventory hits zero, it tells Google you no longer sell that item, and you can lose the ranking you built, even if you restock it soon after. Using the values from Schema.org for temporary statuses helps communicate this more clearly, and doing this correctly is often more important than other property details.

Deep attributes are becoming even more important because when people use AI to search, they tend to be very specific, unlike traditional keyword searches. Customers aren’t just looking for a certain shoe brand anymore; they might also want size 15, sky blue color, delivery by Friday, and comfort for a bad ankle. Details like material, weight, waterproofing, color, size, shipping speed, and return policies help search engines connect these detailed queries to your product, making it easier for customers to find exactly what they need.

Entity And Author Schema

Adding a Person entity does not make anyone an expert. Schema can only help a machine verify authority that already exists.

Here’s an example of when it exists. I have a client specializing in indoor gardening technology. Three months prior to this, we launched a content campaign focused on produce recalls, emphasizing that growing your own food can help avoid exposure. During a foodborne illness outbreak that caused recall searches to spike over a weekend, their March post about lettuce recalls and food safety was highlighted at the top of the AI Overview, alongside the FDA and CDC. It ranked second in organic search under the FDA and received about 1,300 clicks from that news cycle.

This was not due to a single property but because we established a strong digital footprint for the brand. The author’s credentials, published papers, and previous work were all included in the markup, both on the author page and within the post. When an engine is deciding whether to cite a commercial site alongside a government agency, it needs a reason not to discount you.

A generic brand name is another scenario where this proves useful. I worked with a client whose company name is shared with several others, and searching for the company’s CEO did not return them at all. We created a complete person schema on the executive bio page and linked it to correct profile and company details. In just a few days, the AI Overview identified the correct individual.

What schema will not do is compensate for thin content, invented credentials or a neglected profile. It will not turn a dead social account into corroboration, nor will it replace your visible HTML. Earn the authority first, express it clearly, and make it easy to verify.

Build One Valid Template, Then

Scale

Schema pays off at scale, and scaling before you validate just multiplies the same error across a thousand pages. Start with one page and treat it as a model. Pick the most specific entity type that honestly describes it, assign a stable ID, then pull the facts from the visible page and your platform of record. Check those against trusted third-party sources, add only values you can support, and validate before anything goes to development.

Once that page is clean, it is your template.

Segmenting like this also gives you a business case. A location group, a product family, or a set of author pages within striking distance can move faster than a full site rebuild, and you can measure it.

What Came Up In The Q&A

The questions were good enough that I want to point at them rather than summarize them badly.

The group asked how quickly AI systems recognize an entity change when someone moves to a new role and updates their markup, which tools I actually use to build and scale schema across a client site, and how to reinforce E-E-A-T in a sector like branded pharma where there is often no named author to mark up at all.

The answers are in the recording. If you thought of a question after the fact, drop it in the comments under the event in SEJ Pro. Katie, Heather, and I are working through the ones we did not get to.

Watch The Full SEJ Pro Session

The recording goes deeper on all three nine-pillar frameworks, with the markup examples on screen and the full Q&A. The slide deck is going out to members, along with three checklists: local schema, a product and commerce one-sheeter, and an E-E-A-T checklist covering everything you need to get set up.

Join SEJ Pro to watch the full recording and access the supporting community discussion.

More Resources:

Featured Image: michnik101/Shutterstock

https://www.searchenginejournal.com/schema-ai-citations-trusted-source/586785/