



INDUSTRY STUDY | SEPTEMBER 2026

The 2026 Citation Gap Report

What AI answer engines can and cannot
read on orthodontic practice websites



THE CORE FINDING

About a third of the practices in this study
have given AI engines nothing at all to read.
Not one was ready to be recommended.

Summary of Findings



KaleidoscopeAI has audited more than 120 orthodontic practice websites since March 2026. This report draws only on those with a genuine baseline reading, taken before any optimization work was carried out, which gives an unusually clean view of how the profession looks to an AI answer engine when nobody has prepared the site for one.

About a third have no schema markup at all

Schema is a short block of code beneath the surface of a website that states plainly what the site represents: that this is an orthodontic practice, that this is the doctor, that these are the treatments offered, that this is where it can be found. Where it is absent, an AI engine must work all of that out by reading the page, in much the way a passer-by might deduce what a building is used for by looking through the windows.

About a third of the practices examined have none of it whatsoever. A further third hold so little that it does nothing for them. Only a minority carry enough for it to function as intended.

No practice reached the AI Ready standard

Practices are scored out of 100. A score of 85 or above indicates that AI engines can locate a practice, understand what it does, and recommend it with confidence. No practice examined cleared that bar, and the strongest performer in the study fell well short of it.

The typical practice holds its position by default

Most practices fell into the band we describe as At Risk. Practices in this band do appear in AI answers today, but not because of anything they have done to earn it. They appear because nothing else nearby is any easier for an engine to read. It is the position of the only lit shop on a dark street: perfectly good while it lasts, and entirely dependent on what the neighbors decide to do next.

Strong where patients look, weak where machines look

Reviews, Google Business Profile, photography, contact details: all strong across the board. The profession has spent a decade on this and the investment is visible. The layer underneath, the one only a machine reads, is largely absent. The distinction is roughly between a building's frontage and its wiring. One persuades the visitor who has already arrived. The other determines whether anything works.

The difficulty is not that these practices are invisible to AI. Many are appearing in AI answers right now. The difficulty is that a third of them appear without having given AI anything at all to read.

What This Report Does Not Claim



Orthodontists do not need persuading that AI matters. Most are already seeing patients arrive who found them through ChatGPT or a Google AI answer, and the conversation in the profession has moved on from whether this is real to what ought to be done about it. This report is written for that audience and does not spend its opening pages establishing the obvious.

Neither is this a warning about an approaching deadline. Nothing expires on a particular date, and no practice will wake to find itself erased from AI results. Any vendor presenting this as an emergency is selling a deadline that does not exist.

What the report does is establish where practices actually stand, explain what the remedial work involves, and set out the order in which it is worth doing.

Method



Who is in the study

More than 120 practices have been audited since March 2026. This report uses those measured in their original condition, before any work was carried out on the site.

Every practice counted here has a genuine pre-intervention baseline. Where the only record available was an audit run after a practice had already been through an optimization program, the practice was left out entirely, because a site that has already been worked on says nothing about the starting condition of the profession. Test environments, repeat audits of the same practice, and audits with incomplete records are excluded on the same principle.

The practices are spread across 37 states, with a small number in Canada and the Caribbean. Every one was measured before anything was changed.

How these practices were selected

These practices approached us for an audit rather than being drawn at random from a directory. The findings therefore describe the practices examined, and while they are a reasonable indication of the profession, they are not a national survey. A randomized study is planned.

What is measured

The audit examines 31 individual checkpoints across seven areas and combines them into a score out of 100. The seven areas do not carry equal weight, and the distribution of points reflects how much each one contributes to whether a practice gets cited.

Area	Points	What it covers
Content quality	29	Whether there is enough substantial, useful writing on the site
Proof of expertise	20	Whether an engine can establish that the doctor is genuinely qualified
Answer readiness	19	Whether the writing is arranged so an engine can lift an answer from it
Schema	12	Whether code exists declaring what the site represents
Local signals	9	Google Business Profile, directories, geographic information
Patient experience	7	Reviews, responses, and how the site handles a visitor
Competitive position	4	How the practice compares with others in the same market

How the bands work

Band	Score	Interpretation
AI Ready	85 and above	Engines locate the practice, understand it, and recommend it
Developing	65 to 84	Practice appears intermittently rather than reliably
At Risk	45 to 64	Visible today, but only while nothing nearby is easier to read
Critical	Below 45	Largely unreadable to AI engines

Why the scoring has changed, deliberately

We have revised the scoring more than once since we began measuring this in 2025, and we expect to revise it again. AI engines continually adjust what they reward. A scoring system that never changes is not stable; it is simply measuring an earlier version of the thing in front of it.

This matters for how the report should be read. The study period spans more than one version of our scoring, so every finding presented here was tested against both. Anything that held under only one version has been left out.

If a marketing agency is still grading orthodontic practices against the same checklist it used last summer, the resulting numbers say rather more about the agency than about the practices.

The Citation Gap




Every practice website operates on two levels at once.

The visible level is the frontage: photographs, reviews, telephone numbers, doctor biographies, the Google listing. It is what persuades a patient who has already arrived.

The readable level is the address plate and the name on the door. It is the code declaring what the practice is, writing arranged so a machine can extract an answer from it, and clear connections between doctor, credential, treatment and location. It is how anyone gets sent to the building in the first place.

Practices are strong on the first and weak on the second. That distance is the Citation Gap.

The Citation Gap: The Visible Layer vs the Readable Layer

Strong where patients look. Weak where AI looks.

The practical consequence is easily pictured. A parent searches and sees a practice with four hundred reviews and a photograph of a bright, welcoming office. An AI engine reads the same page and must work out for itself that this is an orthodontist, that the practice offers clear aligners, and that the doctor is board certified, because at no point does the site actually say so in a form the engine can read.

What being recommended requires

- Material worth quoting. Writing with sufficient substance that an engine sees value in citing it.
- A declaration rather than an implication. Code that states what things are: this is a practice, this is a physician, this is a treatment.
- An arrangement an engine can lift from. Clear headings, direct answers to the questions patients actually ask, sections that stand on their own.

Across the study, practices are adequate on the first, largely absent on the second, and inconsistent on the third.

The Case Against Urgency

There is a serious argument against reports of this kind, and it deserves a direct answer rather than a footnote.

The argument runs as follows. Large language models read plain English perfectly well. Where no code exists to identify a site, an engine falls back on the signals that have always driven search: established authority, relevance, links, local prominence. Practices with sound traditional optimization are appearing in AI answers today with no schema whatsoever. The schema case,

in this reading, is a vendor manufacturing a problem.

Most of that argument is correct, and this study is evidence for it rather than against it.

A third of the practices examined have no schema at all, and they are still receiving patients. Were the absence of schema genuinely disqualifying, those practices would already have disappeared from AI results. They have not. Engines are demonstrably falling back on traditional signals, precisely as the critics describe.

There is no cliff edge. A practice with a solid search history and no schema will continue to surface for some time yet.

Where the argument actually falls apart

It falls apart on a single phrase: falling back.

When an engine falls back, it is inferring rather than reading. The distinction is the one between a curriculum vitae that states “Orthodontist, board certified” at the top and one that describes twenty years of work from which the reader is expected to draw the same conclusion. The inference is usually correct. It costs the reader effort. And it holds only until a second curriculum vitae arrives with the title printed plainly at the top.

The question is therefore not whether AI can find a practice without schema. It is what happens the first time an engine weighs that practice against another in the same town, one that has stated plainly what it is, what it treats and who works there, while the first left the matter to be deduced.

At that point the practice is not invisible. It is second. And it will have no way of knowing, because nothing on its own website will have changed.

Falling back works until the engine has something better to fall forward to.

Why this is a question of order rather than urgency


Schema is closer to wiring a building than to renting one. It is installed once, it does not wear out, it carries no monthly cost, and it costs much the same whether it is done this quarter or a year from now. There is no compounding penalty for waiting and no prize for haste.

What does exist, at present, is the plain fact that a third of the profession has not done it at all, and a third more have done it so partially that it makes no difference. That opportunity is not closing on a schedule. It closes one practice at a time and one town at a time, as competitors get around to it. Being early costs the same as being late and leaves a practice in a materially better position.

It is also why the central finding of this report concerns schema rather than scores. Whether a website carries schema is a matter of fact rather than judgement. Anyone can establish it in a browser in under a minute, without our scoring system, our thresholds, or our permission. Agreement with anything else in this report is not required for it to be true.

The Seven Areas

The areas are set out below from strongest to weakest. Because the scoring was revised during the study period, each is described by where it ranks rather than by an average. The ranking held under both versions.



Strongest at the top, weakest at the bottom.

Patient experience — strongest

Reviews, responses to reviews, and Google Business Profile completeness. Nearly every practice scores well here, and the profession has earned that result through a decade of sustained attention to reputation.

What it does not do is secure a recommendation. A flawless Google listing is a strong local signal and no substitute for telling an engine what the practice

actually is.

Local signals — strong

Google Business Profile, local directories, and geographic information. Second strongest, and for the same reason as above.

Scores fall away on the more technical elements: location code, the structure of service-area pages, and whether the practice name, address and telephone number agree wherever they appear.

Content quality — middle

The largest area in the scoring at 29 of the 100 points, and it ranks third. Most practices have content, and a good deal of it is perfectly respectable.

What is missing is content built for extraction: direct answers to the questions parents genuinely ask, treatment descriptions with substance, pages deep enough to be worth quoting. This is usually a matter of arrangement rather than authorship, which makes it considerably cheaper to remedy than it first appears.

Proof of expertise — middle

Doctor biographies, credentials, board certification, professional memberships, awards. The information sits on almost every practice website.

It is simply written as prose. A biography states that the doctor is board certified, but nothing identifies that statement as a credential rather than a sentence. The diplomas are in the building. They are just not on the wall.

Answer readiness — middle

Whether the writing is arranged so an engine can lift a clean answer out of it: clear headings, real questions with direct answers, sections that make sense in isolation.

Most practices hold the material and present it as continuous prose. A recipe and an essay about cooking may contain identical information; only one of them can be followed. Reorganizing what already exists improves this without a word of new writing.

Competitive position — weak


Whether a practice surfaces when an engine compares it with others in the same market.

Practices are not short of things that distinguish them. They are short of stating those things anywhere a machine can find them.

Schema — weakest, and the most straightforward to remedy

Schema is the code that declares what a site is, rather than leaving an engine to deduce it. It is the weakest area in the study by a considerable margin.

Schema Markup: About a Third Have None at All




About a third of practices carry none whatsoever. A further third hold so little that it achieves nothing. Only a minority have enough for it to work as intended.

It is also the most tractable problem in this report. It is a technical deployment rather than an editorial project. A baseline covering the practice, the physicians, the location, the treatments and the reviews can generally be installed in a single pass, and on its own tends to move a practice a full band.

Whether a website carries schema is not a score and not an opinion. It is a matter of fact that anyone can establish in a browser in under a minute.

How the Scores Fell



How Practices Scored: Not One Practice Was AI Ready

No practice reached the AI Ready threshold. About a fifth were Developing. Nearly four in five fell into At Risk or Critical.

The At Risk group is the one that matters

More than half of every practice examined landed here, and the definition is exact: visible today, fragile tomorrow. These practices have completed the patient-facing work and skipped the machine-facing work. They appear in AI answers because nothing in their market is easier to read, which is a position held on someone else's inaction rather than one's own effort.

They are also where remedial work pays back fastest, because the expensive half is already behind them. What remains is the cheaper half.

The Critical group is a different case entirely

A substantial group sits here. There is effectively nothing for an engine to read, and in most instances these practices are not appearing in AI answers at all. They are not at risk of losing ground. They do not presently hold any.

What Practices Should Do

The order matters a good deal more than the speed.

1. Establish where the practice stands

None of this is apparent from looking at one's own website, because none of it was written for a human reader. It requires someone to examine what a machine sees. Any audit reporting area by area will serve; ours is one option among several.

2. Deploy schema first

It is the weakest area by a wide margin and the least expensive to correct. A one-time installation, no ongoing cost, and on its own it generally moves a practice a full band. For most practices reading this report, nothing else on the list returns as much for as little.

3. Reorganize existing content before commissioning any more

Answer readiness is usually held back by arrangement rather than by any shortage of material. Breaking long pages into genuine questions with direct answers and clear headings improves the position immediately. This is worth doing before paying anyone to write more.

4. Label the credentials already present

The biographies already state that the doctor is board certified. Adding the code that marks that as a credential is a technical task rather than an editorial one, and it lifts the expertise score without a single new sentence.

5. Treat it as maintenance rather than a project

Engines continually adjust what they reward. A practice that installs schema once and considers the matter settled will drift as the mechanics move beneath it. The ongoing effort is modest, but it is not nothing.

6. Judge suppliers on method rather than vocabulary

Every marketing firm serving this profession now describes itself as using AI. The more useful question is narrower. Ask what they measure, ask to see it, and ask how those measures have changed as the engines have changed. If nothing has changed, that is itself an answer.

About KaleidoscopeAI



KaleidoscopeAI™ is an AI-driven marketing system built specifically for the orthodontic and dental professions, with 17 years behind it serving this field. More than 120 practices have been audited since March 2026. The company's product is the KAL AI Growth System™, covering audits, content, SEO, AEO, ADA compliance and websites, all with human review.

Most marketing platforms added AI to what they already had. KAL™ was built from the ground up, for one profession, for this particular problem.

For an AI Visibility Audit of your own website, visit kaldental.ai.

Study details

Published	September 2026
Audits conducted	More than 120 practices since March 2026
Period covered	March to early September 2026
Geography	Practices across 37 states, with a small number in Canada and the Caribbean
Scoring	31 checkpoints across seven areas, scored out of 100. The scoring was revised during the study period; every finding presented held under both versions.
Inclusion	Only practices with a genuine pre-intervention baseline. Practices whose only record was a post-optimization audit are excluded, as are test environments, repeat audits and incomplete records.
Sampling	Practices approached us for an audit rather than being drawn at random. Indicative of the profession rather than a national survey.

© 2026 KaleidoscopeAI. All rights reserved.