

RESEARCH REPORT · HEALTHSPAN MEDIA

The AI Visibility Report 2026: *Longevity Clinics*

How 25 longevity clinics appear across Google AI Overviews, Perplexity, ChatGPT and Claude.

Published by Healthspan Media · 2026

Contents

Executive summary	5
The findings in context	11
The six citation patterns	16
Pattern 1: self-published authority	17
Pattern 2: community citation authority	20
Pattern 3: schema and structural extractability	22
Pattern 4: academic and institutional authority	24
Pattern 5: geographic and contextual fragmentation	26
Pattern 6: evidence-based citation	29
The 25 clinics	32
What is being cited instead	100
The 12-to-24-month forecast	106
Recommendations	112
Methodology, disclosures and notes	117
Colophon	121
THE 25 CLINIC SNAPSHOTS	
01 Fountain Life US · Convergence anchor	39
02 Human Longevity Inc US · Convergence anchor	42
03 Clinique La Prairie Europe · Convergence anchor	46
04 Cleveland Clinic Abu Dhabi Middle East · Convergence anchor	50
05 Biograph US · Convergence anchor	54
06 Chi Longevity APAC · Convergence anchor	58
07 Early Medical US · Narrow visibility	76
08 Sollis Health US · Broad visibility	70
09 Wild Health US · Narrow visibility	78
10 Next Health US · Broad visibility	72
11 Cenegegnics US · Broad visibility	74
12 Lanserhof Europe · Convergence anchor	62

13	SHA Wellness Clinic Europe · Convergence anchor	64
14	Chenot Palace Europe · Convergence anchor	66
15	Conradia Medical Prevention Europe · Zero observations	94
16	Vivamayr Europe · Narrow visibility	80
17	Aman Wellness Global · Zero observations	96
18	AEON Clinic Dubai Middle East · Narrow visibility	82
19	King's College Hospital London Dubai Middle East · Narrow visibility	84
20	Tokyo Midtown Clinic APAC · Zero observations	98
21	CHA Anti-Aging Center APAC · Narrow visibility	86
22	Everlab APAC · Narrow visibility	88
23	Function Health US · Convergence anchor	68
24	Generation Lab US · Narrow visibility	90
25	Healthy Longevity Clinic Prague Europe · Narrow visibility	92

RESEARCH REPORT · AN AUDIT BY HEALTHSPAN *Media*

The AI Visibility Report 2026: *Longevity Clinics*

How 25 longevity clinics appear across Google AI Overviews, Perplexity, ChatGPT and Claude. Most clinics are losing visibility on a competitive surface they are not yet measuring.

Published 2026 · Reading time 38 minutes · By Healthspan Media

SECTION № 01

Executive *summary*

This report measures how 25 leading longevity clinics appear in the four AI search platforms that prospective clients now use to research and compare: Google AI Overviews, Perplexity, ChatGPT and Claude. The study covers a fixed 50-query panel, run across all four platforms, with brand-name, geographic, service-led, comparison and client-research framings. The dataset comprises 200 query-platform observations.

The headline finding is that AI search operates as four parallel markets. Each platform uses a different definition of authority, and those definitions do not agree. A prospective client who asks Google AI Overviews "best longevity clinic in Dubai" sees a different number-one recommendation than the same person asking Perplexity, ChatGPT or Claude. The variance is structural rather than random or model-quality driven. The competitive map of the longevity industry, as constructed by AI search, is a federation of four parallel maps.

Six findings stand out.

FINDING № 01

Google AI Overviews suppresses entire query categories

Google AI Overviews refuses to generate any answer for two query categories that prospective clients commonly use: practitioner-name queries (eleven of eleven tested returned no AIO output) and certain city-anchored queries about longevity medicine in major markets.

Google AI Overviews refuses to generate any answer for two query categories that prospective clients commonly use: practitioner-name queries (eleven of eleven tested returned no AIO output) and certain city-anchored queries about longevity medicine in major markets. ChatGPT and Claude answered every one of these suppressed queries cleanly. The suppression operates at the template level. Different phrasings of the same underlying question can succeed or fail based on the structure of the query, regardless of which clinic is being asked about. The likely mechanism is YMYL (your-money-your-life) conservatism applied before retrieval or ranking.

Clinics whose clients ask "[doctor's name] clinic" or "best longevity clinic in [their city]" cannot win those queries on AIO through content investment, because no AI-generated answer exists to compete for. Conventional SEO remains the only competitive layer.

FINDING № 02

Self-published authority is the strongest investment in AI search visibility

The most reproducible pattern across the dataset is that clinics publishing their own content at the exact long-tail question phrasings prospective clients type into AI search win citations consistently.

The most reproducible pattern across the dataset is that clinics publishing their own content at the exact long-tail question phrasings prospective clients type into AI search win citations consistently. Where this pattern is present, the clinic wins. Where it is absent, even category leaders are routinely outranked by smaller specialty operators.

Fountain Life is the clearest demonstration. Its own blog appears repeatedly across AI Overview citation panels for high-intent category queries, and the reach extends across languages and geographies: in supplementary multilingual testing, Fountain Life's self-published content surfaced as cited authority alongside locally-anchored operators in non-English AI search about geographically distant markets. A US-headquartered clinic's blog is being treated as authority in foreign-market AI search. Clinique La Prairie demonstrates a different version of the same pattern, built on decades of luxury-press citations newer clinics cannot replicate through current content investment.

FINDING № 03

Brand integrity holds, with two exceptions

Across 200 query-platform observations, the four platforms got the factual content right far more often than they got it wrong.

Across 200 query-platform observations, the four platforms got the factual content right far more often than they got it wrong. Claude and ChatGPT produced zero hallucinations across the queries they answered.

Two exceptions stand out. Google AI Overviews produced one oversimplification describing Biograph's ownership structure, implying sole ownership by Peter Attia. Biograph is a separately incorporated venture publicly described as co-founded by John Hering, James Burgess and Peter Attia, with external venture capital backing. Perplexity produced the dataset's clearest brand-integrity failure during the study window, on a Wild Health versus Parsley Health comparison query, citing nine wrong sources including a movie, a dictionary entry and a brand of deodorant. The mechanism in both cases is named-entity disambiguation, not fabrication. Clinics named with common-language phrases (Wild Health, Function Health, Modern Age, Forward Health) are structurally more exposed than clinics with distinctive brand names.

FINDING № 04

Twelve percent of the panel produced no observations in AI search

Three of the 25 clinics in the study sample produced zero observations across 200 query-platform combinations during the study window.

Three of the 25 clinics in the study sample produced zero observations across 200 query-platform combinations during the study window. Not zero anchored citations. Zero observations of any kind, across all four AI platforms and the entire 50-query panel. The clinics are Conradia Medical Prevention (Germany, whole-body MRI specialist), Aman Wellness (multi-location luxury hospitality wellness with longevity programming) and Tokyo Midtown Clinic (Japan, premium longevity service).

Two of the three operate primarily in markets where the dominant content language is not English. The third is a multi-location global hospitality brand whose longevity programming is comparatively recent. All three hold premium pricing tiers and established reputations within their home markets. In our sampled AI-search outputs, none of those qualifications produced visibility during the study window.

In the study sample, premium reputation, premium pricing and home-market authority did not translate into AI search visibility automatically. AI search systems in the study more frequently surfaced English-language content environments than locally authoritative but less internationally indexed sources. A clinic that does not appear in AI search outputs may be less discoverable to international patients using these systems. The pattern may reflect a broader structural discoverability challenge affecting some portion of premium longevity clinics globally, particularly those whose content infrastructure is weighted toward non-English-primary languages.

FINDING № 05

Four-platform convergence on the European luxury quartet and US executive longevity

The most stable category anchors in the dataset are the European luxury wellness comparison (Clinique La Prairie, Lanserhof, Chenot Palace, SHA Wellness) and the US executive longevity comparison (Mayo Clinic Executive Health, Fountain Life, Human Longevity Inc, Biograph).

The most stable category anchors in the dataset are the European luxury wellness comparison (Clinique La Prairie, Lanserhof, Chenot Palace, SHA Wellness) and the US executive longevity comparison (Mayo Clinic Executive Health, Fountain Life, Human Longevity Inc, Biograph). The four AI platforms converge on the same core set for each query, with one or another clinic missing from any single platform but the overall set stable. Clinique La Prairie, Lanserhof and SHA Wellness anchor on all four platforms in the European set; Mayo Clinic Executive Health does the same on the US executive query. These are the brand positions AI search has solidified into category-defining defaults. New entrants to these two categories face a structural moat built on decades of accumulated authority citations.

FINDING № 06

The Gulf has emerged as the global axis of the longevity category

The most consequential industry shift reshaping the 2026 competitive map is the Gulf's emergence as the category's regulatory, funding and expansion axis.

The most consequential industry shift reshaping the 2026 competitive map is the Gulf's emergence as the category's regulatory, funding and expansion axis.

Abu Dhabi's Department of Health established the world's first evidence-based clinical guidelines for healthy longevity medicine in April 2025, and licensed IHLAD (Institute for Healthier Living Abu Dhabi) as the world's first specialised healthy longevity medicine centre under a dedicated regulatory framework, led by Dr. Nicole Sirotin with protocols designed by Chi Longevity's Andrea Maier. Saudi Arabia's Hevolution Foundation has publicly committed to deploying up to \$1 billion annually toward healthspan research from Riyadh, the largest sovereign funding pool for longevity science in the world. The luxury-longevity expansion map increasingly points toward the Gulf: Clinique La Prairie opens a Red Sea coast Saudi outpost in autumn 2026, Lanserhof's Marbella location opens in 2027, and additional Gulf and adjacent-market expansions are in motion across multiple clinics in the panel.

The Gulf now hosts both the world's first regulated longevity facility and its largest research funder. For Western clinics, the strategic question is how to engage the Gulf market, not whether.

The category-evidence anchor

A 2025 survey of longevity clinics conducted by the Detki Family Clinic and Molodost Clinic, widely reported in industry coverage, found that nearly 90 percent of longevity clinics worldwide provide treatments influenced more by marketing than by scientific evidence. Only 10 percent of clinics focus primarily on evidence-based interventions; 25 percent provide experimental but plausible treatments.

This finding informs the report's editorial position. The longevity-medicine sector is not categorically fraudulent, but it is characterized by substantial hype, uneven evidence standards and commercial claims that often outpace scientific validation. Distinguishing legitimate preventive care from weakly

substantiated intervention is the central editorial challenge within the category. Current AI-search systems do not yet reliably distinguish between evidence-based medicine and more speculative longevity offerings.

The report's analytical findings about specific clinics in the study panel address AI search visibility and citation patterns, not the clinical or evidentiary quality of those clinics' offerings.

Implications for clinics, agencies and editorial publishers

For clinic operators: AI search visibility is platform-specific and template-specific. The strongest investment is in long-tail content that answers the exact phrasing of high-intent client questions on the clinic's own domain. Local-market positioning (Google reviews, local longevity directories, regional press) is a different competitive layer from global-brand positioning, and clinics should track each separately. Practitioner-named queries cannot be won on Google AIO; they should be conceded to organic search and to other AI platforms. For clinics operating in non-English-primary markets, the English-language content infrastructure that AI platforms cite needs to be addressed deliberately and at scale. Reputation and home-market credibility do not transfer automatically.

For agencies: the standard SEO playbook helps but is not sufficient. Schema markup, FAQ blocks and topical landing pages for the exact phrasings AI platforms cite are what win self-published authority. Press placement in the publications AI platforms actually treat as authority compounds over time.

For editorial publishers: the editorial lane between geroscience research and clinic marketing is structurally underserved. AI platforms route evidence-evaluation queries through academic journals because there is currently no third-party editorial brand occupying that lane consistently. The clinics that win AI citations are doing so through their own self-published content rather than through any third-party publication acting as a referee. That gap is the opportunity.

This report's findings rest on a defined methodology, transparent disclosures and verifiable source citations.

[View methodology](#) · [Editorial disclosures](#) · [Sources and citations](#)

SECTION № 03

The findings *in context*

IN BRIEF

AI search has fragmented into four parallel discovery markets with limited overlap. The same clinic can be cited everywhere on one platform and nowhere on another, and citation share is dictated by patterns most clinics are not yet measuring.

The same clinic can be the anchored recommendation on one platform and entirely absent on another, on the same query.

From Section 3, the findings in context

12%

of clinics in the study sample produced zero observations across the 200 query-platform combinations measured.

200

query-platform observations form the dataset spine of the report.

AI search for longevity care operates as four parallel markets. The four AI platforms have arrived at four different definitions of authority, and those definitions do not converge on a single canonical answer to most questions. A prospective client asking the same question on Google AI Overviews, Perplexity, ChatGPT and Claude receives, in the typical case, four meaningfully different recommendation sets.

The dataset also includes an empirical extreme: twelve percent of the 25-clinic panel produced no observations across any platform on any query.

The shape of the dataset

Across the study's 200 query-platform observations, the four platforms differ substantially in how many clinics they surface per query and in how they construct their answers. Claude produces the broadest competitive sets, surfacing the largest average number of clinics per query and the highest fraction of queries that surface any of the 25 clinics in the study. Google AI Overviews sits at the opposite end, producing the tightest answers anchored on the fewest clinics. Perplexity and ChatGPT sit between, with Perplexity tilting toward comprehensive competitive sets and ChatGPT tilting toward routing client-research queries through generic medical framing rather than specific clinic recommendations.

Two queries returned zero clinics from the 25 across all four platforms: a query about longevity clinic complaints and a query about whether longevity medicine is evidence-based. The AI search ecosystem routes these entirely through third-party authority (mainstream press, peer-reviewed journals, community discussion). For category-skeptical questions, no clinic in the study, however well-positioned, surfaces as an answer.

The empirical anchor: three of the 25 clinics produced zero observations

Twelve percent of the 25-clinic study panel produced zero observations across the entire 200 query-platform dataset during the study window. The three clinics are Conradia Medical Prevention (Germany, whole-body MRI specialist), Aman Wellness (multi-location luxury hospitality wellness with longevity programming) and Tokyo Midtown Clinic (Japan, premium longevity service). Not zero anchored citations. Not zero primary recommendations. Zero observations of any kind, across all four AI platforms, across all 50 queries in the panel.

Two of the three operate primarily in markets where the dominant content language is not English. The third is a multi-location global hospitality brand whose longevity programming is comparatively recent. All three hold premium pricing tiers and established reputations within their home markets, built through years or decades of local-reputation accumulation (regional press, domestic medical authorities, local review aggregators). In our sampled AI-search outputs, none of those qualifications produced visibility during the study window.

The finding contradicts an assumption widely held in the longevity-care category: that premium positioning, established reputation and home-market authority translate into AI search visibility automatically. In the study sample, they did not. AI search systems in this study more frequently surfaced English-language content environments than locally authoritative but less internationally indexed sources. A clinic that does not appear in AI search outputs may be less discoverable to international patients using these systems for cross-border healthcare research.

The finding may extend beyond the three named clinics. The pattern could reflect a broader structural discoverability challenge affecting some portion of the global premium longevity clinic population, particularly those whose content infrastructure is weighted toward non-English-primary languages. The 22 clinics that did produce observations represent a sample of clinics whose positioning happened to translate into English-language AI visibility within the study parameters; the broader global population may contain a higher proportion of the pattern than this study's scope captures.

Three category anchors hold across all four platforms

A small number of category positions produce strong cross-platform convergence. Of the 50 queries in the study, three returned three or more clinics surfacing on three or more of the four platforms.

The first is European luxury wellness. A query about the best anti-aging clinics in Europe produces the same four-clinic set across the platforms: Clinique La Prairie, Lanserhof, SHA Wellness and Chenot Palace. Clinique La Prairie, Lanserhof and SHA Wellness anchor on all four platforms; Chenot Palace anchors on three of four, with one platform naming a different European wellness operator in its place. This is the strongest cross-platform convergence in the dataset.

The second is US executive longevity. A query about executive longevity assessment produces a stable four-clinic set: Mayo Clinic Executive Health, Fountain Life, Human Longevity Inc and Biograph. Mayo Clinic Executive Health anchors on all four platforms; the other three each anchor on three of four, with different platforms missing different clinics.

The third is full-body MRI screening. A query about where to obtain full-body MRI produces a three-clinic set dominated by Penuvo and Ezra as specialty operators, with Biograph appearing as the integrated-longevity-care alternative. Penuvo anchors on all four platforms; Ezra and Biograph each anchor on three of four.

These three positions are the structural anchors of the AI-defined longevity industry as of mid-2026. A clinic already occupying one of them holds a brand position reinforced every time the AI platforms train on each other's citation patterns. A clinic attempting to enter through near-term content or PR investment faces decades of accumulated authority citations. Displacement on these three positions is effectively prohibitive for most challengers.

On the majority of queries, the platforms disagree

Outside the three category anchors, the four platforms produce different competitive sets on most queries. The typical pattern is that one or at most two clinics overlap across all four platforms, while each platform surfaces its own additional set of recommendations.

A query about the best longevity clinic in Los Angeles illustrates the pattern at full strength. Google AI Overviews anchors on one clinic; Perplexity anchors on a different one; ChatGPT produces an answer dominated by local LA operators with no overlap to the other three platforms; Claude produces an answer that overlaps partially with Perplexity and Google AI Overviews but shares no clinics with ChatGPT. Same query, same week, same study configuration. Four different competitive contexts.

The executive longevity query, despite being one of the three category anchors, shows the divergence pattern operating within the convergence. Mayo Clinic Executive Health surfaces on all four platforms, confirming anchor status. The other anchor clinics each surface on three of four, with different platforms missing different clinics. The platforms agree on a core set but disagree on the surrounding competitive set. A clinic's visibility in the supporting cast of a category-anchor query is a different competitive game from anchoring the query itself.

The pattern repeats across the dataset. Geographic queries diverge most. Service-specific queries diverge less. Comparison queries diverge based on whether the comparison names a category anchor. The general rule: the more specific the query, the more the platforms diverge; the more general, the more they converge on a small canonical set.

Four editorial personalities

The platforms differ in which clinics they surface and in how they construct their answers. Four distinct editorial postures emerge.

Google AI Overviews favors institutional authority paired with clinic-owned content. Its citation panels are densely populated with clinic blog content, paired with selected mainstream press for context. This is the platform where self-published authority operates most powerfully. Google AI Overviews also performs the cleanest template suppression behavior, refusing to generate answers for certain query categories regardless of which clinics might otherwise compete for them.

Perplexity favors review-aggregator and comparison-blog ecosystems. Its citation patterns lean heavily on TripAdvisor for European luxury clinics, Google review counts on geographic queries, Reddit threads on practitioner-name and brand-skeptical queries, and longevity trade press on category comparisons. Perplexity is the most willing of the four to attempt a comprehensive answer, which sometimes produces depth and sometimes produces error. The dataset's clearest brand-integrity failure was a Perplexity output.

ChatGPT favors generic medical framing on category-definition queries, with selective clinic citations on geographic and comparison queries. ChatGPT is also the most sensitive of the four platforms to geographic IP localization and to prompt framing. The same query, asked from different

geographies or with different framing instructions, produces materially different ChatGPT competitive sets. What ChatGPT shows a New York user is different from what it shows a Singapore user, sometimes substantially.

Claude favors comprehensive, citation-dense, evidence-led framing. Claude surfaces more clinics per query than the other three platforms, leans more heavily on academic credentials and peer-reviewed research as authority signals, and demonstrates the strongest sensitivity to recent industry shifts in its dataset. Claude was the only platform of the four to capture certain category-defining 2026 events within the study window. Its framing tilts toward credentials over media prominence, which produces a different competitive map from the platforms that weight media presence more heavily.

What this means for clinics

AI search visibility operates as four separate competitive games. A clinic that wants visibility on all four platforms cannot rely on a single content-strategy bet. Each platform favors different authority signals, and a clinic optimizing for one without considering the others will systematically miss visibility on the others.

The three category anchors where cross-platform convergence holds represent positions worth defending aggressively for any clinic already occupying one. Displacement attempts on these three positions are existential. Continuous content investment on the queries that anchor each category is the minimum required to defend the position.

For clinics not occupying a category anchor, the strategic clarification is sometimes more about which queries to concede than which queries to compete for. Certain categories of question are routed through specialty operators, directories or third-party authority sources that no clinic in this study is winning.

For clinics in markets where the dominant content language is not English, or operating in regions whose local-reputation systems are well-developed but disconnected from the English-language content infrastructure AI platforms cite, the finding from the three zero-observation cases is direct. In the study sample, premium reputation, premium pricing and home-market authority did not translate automatically. The English-language infrastructure is its own competitive layer and requires its own investment. Clinics whose international client acquisition pipeline depends on AI search visibility cannot rely on local-market success as a proxy for global discoverability.

SECTION Nº 04

The six citation *patterns*

IN BRIEF

Six citation patterns explain most of the visibility outcomes in the study. Two of them (self-published authority, and schema and structural extractability) are immediately actionable. The other four shape what is structurally possible and what is not.

CITATION PATTERN № 01

Pattern 1: *self-published authority*

50

queries on each of four platforms, run during the study window.

The single strongest predictor of whether a clinic appears in AI search results is whether the clinic publishes content that does the work an intermediation site would otherwise do. The clinics that consistently surface as authoritative citations across AI platforms are not the ones with the most polished service pages. They are the ones writing comparison guides, ranked lists, cost breakdowns and definitional explainers under their own brand. AI platforms treat that content as authority. Most clinics do not produce it.

The mechanism

AI platforms answering a category question need to construct an answer from somewhere. The structurally preferred source is content that has already done the comparison work: ranked lists, cost guides and category explainers. Where a recognizable third party publishes this work (industry directories, trade publications, academic sources), AI platforms cite that source. Where it does not exist at scale, they cite whatever does the work most credibly, which is increasingly content published by clinics on their own domains.

The longevity-clinic intermediation layer is real but thin. The directories that have emerged to fill it are uneven, often outdated and visibly biased toward clinics that have purchased placement. Into that gap, a small number of clinics have published their own intermediation-style content and AI platforms reward them disproportionately.

The dominant case

Fountain Life is the clearest demonstration in the dataset. Its own blog appears as an authoritative citation source across category-ranking, comparison, cost and definitional queries. On Google AI Overviews specifically, Fountain Life's blog appears with citation density approaching that of academic and institutional sources.

The reach extends beyond what most clinics in the study can match. In supplementary multilingual testing, Fountain Life's blog surfaced as cited authority in non-English AI search about distant markets, alongside locally-anchored operators. No other clinic in the study demonstrates comparable cross-cultural citation reach.

What makes the content function as authority is that it survives the editorial scrutiny AI platforms apply when cross-referencing sources. The category-ranking guide largely agrees with what other authoritative sources say about the same competitive set; where it diverges, it reads as reasonable editorial framing rather than promotional distortion. The cost guide is internally consistent and externally verifiable against pricing visible on competitor sites. Platforms treat the content the way they treat trade-press analysis.

This is a high bar. Most clinic-published content fails it. Clinics that publish self-promotional content disguised as comparison content are detected and discounted.

The asymmetric case

Cenegenics demonstrates a more instructive version of the pattern. Its service pages do not win citations in their core commercial category, yet the clinic wins citation share on educational and cost queries through two specific pieces of content: a "what is a longevity clinic" explainer and a longevity clinic cost breakdown. On client-research queries about cost and what a longevity workup includes, Cenegenics is cited as a primary source alongside mainstream press.

Cenegenics has been publishing this content for years. Its citation share in core commercial queries has shifted; the citation weight on its educational content has held. AI citation share, once earned, persists across competitive cycles in ways conventional SEO ranking does not.

Platform sensitivity

The pattern is platform-sensitive. Self-published authority operates most powerfully on Google AI Overviews, somewhat less so on Perplexity and less reliably on ChatGPT and Claude. A clinic optimizing for this pattern is in practice optimizing primarily for Google AI Overviews.

Strategic implications

Pattern 1 is the report's strongest finding for clinic operators because it is one of the few authority mechanisms a clinic can build deliberately. Most of the report's other patterns are structurally inherited or require resources clinics cannot acquire.

The pattern is a long-cycle compound asset. Citation share earned through self-published authority compounds over years rather than months: each citation cycle reinforces the platform's perception of the source as authoritative. Clinics that commit early build moats that clinics committing later cannot reach within their planning horizons. A clinic that begins this work in 2027 rather than 2026 is not losing a year of citations; it is permanently competing against compounded positions it can no longer match. Clinics already doing the work face one question: is the discipline holding. Clinics not doing it face a harder one: begin now, or accept that the position has closed.

The pattern changes what the clinic is, not just what its content team does. A clinic that publishes ranked lists including genuine competitors, names where competitors are stronger, presents accurate pricing and writes category-defining definitional content has become, in part, an editorial publisher. Most clinic operational structures are not built for this. A marketing function built around promotional content cannot make the transition with internal resources alone, not because the writing is hard but because the organizational disposition that produces credible editorial work is genuinely different from the one that produces promotional work. Fountain Life's cross-cultural reach is the empirical demonstration of what this transition, sustained over years, can produce.

The pattern has limits, and recognizing them is itself a strategic asset. Pattern 1 carries category-defining and comparison queries; it does not carry evidence-evaluation queries (which route through peer-reviewed journals), practitioner-name queries (which Google AI Overviews suppresses entirely) or single-test queries (which route through specialty operators). The strongest editorial content cannot win categories AI platforms have routed elsewhere by architectural decision.

CITATION PATTERN № 02

Pattern 2: *community citation authority*

Reddit is the second-most-cited source layer in this study's dataset, behind only the clinics' own domains. It appears in the citation lists of a majority of the queries analyzed, and where it appears, it is rarely a single throwaway link. Specific subreddits anchor specific query categories systematically enough that clinics should think of Reddit not as a monolithic source but as a hierarchy of community-authority pages.

The pattern holds across all four AI platforms studied. Perplexity and Claude carry the highest absolute citation rates, both favoring citation-dense, community-inclusive answers. ChatGPT and Google AI Overviews cite Reddit more selectively, primarily on practitioner-name queries, brand-skeptical queries and client-research queries about category legitimacy.

What the citation behavior looks like

AI platforms do not cite Reddit as a generic source. They cite specific posts within specific subreddits, often by direct URL, and the subreddit selection is systematic. Longevity-medicine subreddits anchor executive and biological-age queries. Biohacking subreddits anchor service-led queries about specific protocols. Brand-specific subreddits anchor cost and value questions about the brands they discuss. Futurology and skeptical subreddits anchor category-evaluation queries about whether longevity medicine is legitimate.

The hierarchy is not formal. It emerges from how AI platforms have learned to weight comment quality, discussion depth and the absence of promotional content. Subreddits that aggressively moderate against promotional posts are cited more heavily. Subreddits that allow promotional content are discounted or absent from citation panels entirely.

Why this matters strategically

Self-published authority is a content investment a clinic can make. Reddit citation share is structurally different. The community has the discussion. Clinics can participate as members but cannot publish their way to Reddit citation share.

The clinics that perform well on Reddit-cited queries share one structural attribute: a consumer-facing model that produces real client experiences which real clients then discuss. Function Health, Tally Health and Fountain Life all benefit here. Clinics with smaller, more discreet client bases

(medical-wellness retreats, concierge urgent-care, institutional executive health programs) have less community discussion to draw on and less Reddit citation share. This is a business-model dynamic, not a content gap, and it is not addressable through marketing investment.

The clinics that perform poorly cannot fix the problem by hiring agencies to manage their Reddit presence. The community detects promotional accounts within hours and downvotes them. AI platforms appear to weight upvote-sorted comment threads, which means agency-managed Reddit activity produces negative-signal content rather than positive citations. Agency-style Reddit influence campaigns are actively counterproductive.

The skeptical layer

A specific subset of Reddit citations carries skeptical framing by default. Threads on community-evaluation subreddits about whether longevity medicine is legitimate, whether specific clinics are worth the cost or whether parts of the category lack evidentiary support are cited by AI platforms answering client-research queries about category legitimacy. When these threads are cited, the AI response inherits some of the skeptical framing.

A clinic otherwise well-cited but appearing in a Reddit-skeptical thread can see its integrity score drop because the citation pulls the AI's framing toward doubt. There is no fix from the clinic's side. What can be managed is whether the clinic is operating in a way that produces or attracts skeptical discussion, which is upstream of the citation behavior.

What this means for clinic visibility strategy

Reddit citation share is structurally weighted toward consumer-facing clinics with real client communities. Clinics with this attribute should understand it as a competitive moat that compounds as more client discussion accumulates. Clinics without it should not invest in manufacturing Reddit presence; the investment will not produce the citation share intended.

Query selection is the other strategic move. Certain query categories (client-research, brand-evaluation, service-comparison) are structurally Reddit-weighted across multiple AI platforms. Clinics competing for visibility on these queries are competing in a citation ecosystem they cannot directly control. The strategic question is whether to compete at all, or to redirect content investment to query categories where the citation ecosystem is more responsive to clinic effort.

CITATION PATTERN № 03

Pattern 3: *schema and structural extractability*

A semantically generic brand name is a structural liability in AI search, regardless of content quality.

From Section 4, Pattern 3

Human Longevity Inc was co-founded by J. Craig Venter, who co-led the original human genome sequencing effort. The company has invested over \$600 million in longitudinal genomics and launched a \$599 AI-enhanced genomic insight platform in May 2026. Yet on the query "whole genome sequencing for longevity," three of the four AI platforms studied do not surface Human Longevity Inc as a primary recommendation. The clinics that do surface include specialty operators with far smaller research investments and far less founder credibility.

The query rewards content built for that specific phrasing, and the specialty operators have published it where Human Longevity Inc has not. The platforms filter for parseable structure before they weigh credentials.

What the pattern measures

AI platforms cite what they can parse cleanly. Pages with descriptive URLs, semantic header structures, extractable data formats (tables, lists, named price tiers) and clearly disambiguated brand entities consistently outperform prose pages covering the same ground. Citation share routes to whichever source has done the structural work.

The Fountain Life biological age testing page is the clearest positive case. It uses Q&A blocks, named protocols, numeric thresholds and schema markup. AI platforms read these signals as reference material and extract elements directly into their answers.

The brand-disambiguation problem

During the study window, a Perplexity query about Wild Health returned a citation list mixing a movie, a dictionary entry and a deodorant brand. Perplexity could not reliably identify Wild Health as a medical clinic. The same failure mode hits every clinic in the study with a generic name.

Distinctive names (Clinique La Prairie, Cenegenics, Lanserhof, Biograph, Fountain Life) get disambiguation for free. Generic names have to earn it through schema, external citations and consistent brand context. Publishing more content compounds the ambiguity. The fix is medical vocabulary, schema architecture and ongoing tuning as platforms shift their parsing. That is specialist work.

The format-over-content insight

AI platforms favor pages where data is presented as data. A pricing table that lists annual cost, included tests and visit frequency as named values wins citation share against a prose paragraph covering the same ground. A 2,000-word essay on hormone optimization underperforms a 600-word page with a pricing table, a structured service list and a comparison against other providers. The second looks like reference material. The first looks like marketing.

Building reference material is information architecture work. The skill set sits closer to product than to copywriting.

Strategic implications

Pattern 3 rewards clinics with distinctive brand names and content infrastructure already in place. The work compounds, and few clinics in the category make the investment.

Clinics with generic names face a different proposition. Without active disambiguation work, the failure mode demonstrated in the Wild Health case is permanent, and content volume does not compensate.

Schema and structured data typically sit with dev teams, but the gap between routine implementation and AI parsing requirements is substantial. Schema markup for AI platforms needs medical vocabulary, platform-aware weighting, entity resolution and recalibration as parsing behavior shifts.

CITATION PATTERN № 04

Pattern 4: *academic and institutional authority*

On queries that ask about evidence, scientific basis or the legitimacy of longevity medicine as a field, AI platforms route citations differently. Peer-reviewed journals and the educational content of major medical institutions carry the answer. The clinic-authored content, Reddit threads and directory sites that anchor other query categories appear much less often, and on some queries not at all.

What the pattern looks like

The query "is longevity medicine evidence-based" returns a citation panel of peer-reviewed journals (the British Medical Journal, The Lancet Healthy Longevity), institutional research centers like the Max Planck Institute for Biology of Ageing, mainstream press health verticals and the educational arms of institutions like Cleveland Clinic and Mayo Clinic. Clinic content, Reddit and directory sites are absent.

For broader educational queries (what biological age testing measures, what a longevity workout includes), academic sources surface alongside clinic content. As query framing shifts from product toward evidence, the academic share rises. On pure legitimacy queries, the citation panel is effectively closed to commercial content.

The structural ceiling

This is the pattern clinics cannot win through marketing investment. Publishing in BMJ or The Lancet requires a research arm. Earning a Max Planck citation requires academic conference participation and peer-reviewed publication. Coverage in the New York Times health vertical requires either original research or a story angle that survives the publication's editorial scrutiny. These are clinical and research functions that marketing cannot substitute.

The clinics that operate a research arm benefit. Cleveland Clinic's Center for Functional Medicine and its consumer education content both feed citations on evidence queries. Mayo Clinic operates similarly. Human Longevity Inc earns occasional academic share through Craig Venter and the Health Nucleus phenotypic database. The rest of the panel carries no institutional weight of this kind.

The middle authority layer

AI platforms weight statistical density even in non-academic sources. Content that cites specific percentages from dated studies, distinguishes strong evidence from preliminary evidence and names uncertainty where it exists earns higher citation rates than the same material written as marketing prose.

This is the middle layer of Pattern 4, and it is reachable on a marketing timeline. A clinic that cannot publish in peer-reviewed journals can write content that reads as reference material: dated sources, structured intervention comparisons, honest acknowledgment of where the evidence is thin.

Few clinics commit to this kind of editorial discipline. Those that do gain a citation advantage that compounds and that competitors cannot displace without making the same commitment.

What this means strategically

Concede the legitimacy queries. Content designed to argue whether longevity medicine is evidence-based is marketing investment routed to a contest clinics cannot win; that question gets settled by academic and mainstream press.

Compete on the queries clinics can win: specific protocols, specific clinics, specific service categories, specific geographies. The legitimacy of any of those is a separate question. A marketing function that spreads investment across both layers spends most of its budget on the layer it cannot influence.

The middle layer is reachable but requires editorial discipline. Moving from promotional writing to reference content is a genre shift. The clinics that make the transition compound citation authority in ways their competitors cannot match.

CITATION PATTERN № 05

Pattern 5: *geographic and contextual fragmentation*

AI search visibility is not uniform across geographies or query phrasings. The same clinic can be the top recommendation on one AI platform from one location and absent from another platform on the same query from a different location.

Two extreme cases sit in the dataset. Three clinics produced zero observations across all 200 query-platform tests. Google AI Overviews refuses to answer certain query templates at all, regardless of which clinics or practitioners they name.

Geographic anchoring on geo-named queries

When a query names a city or region, the four AI platforms return different answers. A query about the best longevity clinic in Dubai produced four different top recommendations across the four platforms with no overlap. ChatGPT favored evidence-based diagnostic positioning. Claude favored content from clinics' own domains. Perplexity defaulted to Google review counts and local aggregator citations. Google AI Overviews favored institutional brand recognition and global names operating local outposts.

Same question, four different versions of the Dubai longevity market. The prospective client sees only one version, returned by whichever AI they happen to use. A clinic dominant on one platform may be invisible to clients arriving through another. Geographic positioning becomes platform-specific work that has to be solved four times.

Implicit geographic localization on non-geo queries

The same fragmentation operates on queries that don't name a location. A query about NAD+ IV therapy, full-body MRI screening or hormone optimization returns one recommendation set from a US IP, a different one from an Asia-Pacific IP and a different one again from a Gulf IP. In the more extreme cases, the local and global answer sets do not overlap at all. None of the global clinics in this study surface for an Asia-Pacific user asking about full-body MRI; none of the local clinics surface for the same query from a US IP.

The empirical extreme: zero observations

At the extreme end, fragmentation produces zero AI-search observations. Three clinics in the study produced zero observations across all 200 query-platform tests during the study window: Conradia Medical Prevention (Germany), Aman Wellness (multi-location global) and Tokyo Midtown Clinic

(Japan). All three hold substantial credibility in their home markets and among their international peer groups. Two of the three operate primarily in markets where the dominant content language is not English.

In our sampled AI-search outputs, AI platforms constructed their recommendation sets from English-language content infrastructure: clinic blogs, English-language press, directories, review aggregators, community discussion. Clinics whose content infrastructure was weighted toward a non-English primary language sat outside that ecosystem in the study sample. Premium pricing did not compensate. Established reputation did not compensate. Regional medical press did not transfer.

The English-language content layer appears to require its own investment, separate from the home-market reputation work. The pattern may reflect a broader structural discoverability challenge affecting more premium longevity clinics globally than this panel's 25 captures.

Template-level suppression on Google AI Overviews

Google AI Overviews suppresses certain query templates regardless of who is named. Two classes of queries fail: requests for the best or top longevity clinic associated with a named practitioner, and the same request anchored to certain major cities. The suppression operates on the shape of the query; the named entity is incidental.

Across eleven named-practitioner queries using a consistent template, AIO refused to generate any answer. The set covered clinic founders with public consumer-facing practices and researchers without public clinics. The same practitioners, queried with templates that did not request endorsement, produced substantive AIO responses. The filtering happens at the query-structure layer before retrieval or ranking.

The likely explanation is YMYL conservatism. Queries requesting endorsement of named medical practitioners or clinics fall into a category Google does not answer with AI-generated content.

City-anchored queries face similar suppression, though not uniformly across geographies. The failing phrasings share an endorsement frame ("best," "top," "leading"); removing it sometimes restores the response.

Clinics whose prospective clients ask AIO to endorse a named practitioner or to rank clinics in a named city cannot win those queries through content investment. Conventional Google search still answers them; the AIO panel does not.

What this means strategically

Track visibility separately by platform, geography and query template. A single global or home-market visibility report works from a partial map.

The local intermediation layer (regional review aggregators, local longevity press, geo-specific directories) is a different competitive game from global authority, and few clinics invest in it. For geo-concentrated clinics, the local layer is more reachable and less defended.

Clinics in non-English-primary markets face an infrastructure decision: whether to build a parallel English-language content presence, at what scale and to what editorial standard.

Google AI Overviews' template suppression places certain queries permanently outside the AIO competitive landscape. Budget allocated to winning them is misallocated. The available work is identifying which queries fall inside the suppression and which adjacent queries remain winnable.

Geographic and template fragmentation is a structural feature of how AI platforms now construct, and decline to construct, answers. Clinics that adapt treat platform, geography and template as separate problems.

CITATION PATTERN № 06

Pattern 6: *evidence-based citation*

A Reddit user posting on the longevity subreddit reported that three commercially available biological age tests had measured them at 56, 68 and 77 years against a chronological age of 71. Same person, same week, three different tests, a 21-year spread. The data point now appears in citation panels across multiple AI platforms answering questions about biological age testing accuracy. AI platforms reach for it because it is specific, verifiable and corrosive to overclaiming.

AI search systems in the study more frequently surfaced content that aligns with established evidence framings, and they discount or penalize content that contradicts them. Clinics whose positioning aligns with the evidence base AI platforms have absorbed gain citation share. Clinics that overclaim against it lose citation share.

What the pattern measures

A 2025 longevity-clinic survey, widely reported in industry coverage, established what has become the canonical statistic for the field: roughly 90 percent of longevity clinics worldwide provide treatments influenced more by marketing than by scientific evidence; 10 percent focus primarily on evidence-based interventions; 25 percent provide experimental but plausible treatments. AI platforms have absorbed this framing. When evidence-evaluation queries are asked, the Aging-US statistic appears in citation panels alongside peer-reviewed journals, institutional research centers and mainstream press, and the framing it establishes propagates into AI-generated answers about specific clinics, specific protocols and specific service categories.

Clinics that acknowledge which of their interventions are evidence-supported, which are experimental but tracked and which they do not offer because evidence is insufficient gain citation alignment with the framing AI platforms have already absorbed. Clinics that present uniform claims of effectiveness across all interventions compete against that framing, and AI platforms discount their content accordingly.

The same dynamic operates at the level of specific evidence statistics. The Mandsager 2018 Cleveland Clinic observational study of 122,007 patients, which found that Elite cardiorespiratory fitness was associated with approximately 80 percent lower all-cause mortality hazard relative to Low fitness, has become a canonical citation for the longevity and exercise evidence base. Clinics that publish outcome data referencing this evidence gain citation alignment on cardiovascular and longevity-fitness queries. Clinics whose fitness programming runs without reference to outcome evidence surface less consistently on those queries.

What evidence-aligned content looks like

Evidence alignment is an editorial discipline. The clinics that practice it write content that cites specific evidence sources by name and date, with appropriate epistemic markers (this study showed, this finding has been replicated, this is preliminary). It distinguishes interventions with strong evidence (Zone 2 training, VO2 max optimization, smoking cessation) from interventions with preliminary or contested evidence (rapamycin, senolytics, NAD+ precursors). It acknowledges where evidence is insufficient, including where the clinic does not offer a treatment because the evidence does not support it.

This is uncomfortable content to publish. It requires acknowledging that some of the interventions the category sells are weakly supported. The clinics that publish it anyway gain a citation advantage on evidence-evaluation queries that competitors making uniform effectiveness claims cannot match. AI platforms weight epistemic honesty as an authority signal with direct commercial consequences.

The asymmetric penalty for overclaiming

When AI platforms encounter content that overclaims against the established evidence base, the content is discounted and sometimes actively flagged in the AI's framing of the clinic. The 21-year-spread finding now functions as a default citation when AI platforms answer questions about biological age testing accuracy. Clinics whose marketing presents biological age as a precise measurement inherit some of that skepticism in the AI's framing.

A clinic that frames biological age testing accurately (a tracking tool subject to test-to-test variance, useful for within-person trend analysis) aligns with the evidence AI platforms have absorbed and gains citation share. Marketing that presents it as a precise measurement does the opposite, competing against the 21-year-spread finding and surfacing less consistently on related queries. The marketing language a clinic chooses about biological age testing has direct AI citation consequences that most marketing teams have not considered.

The same dynamic applies to other commonly overclaimed interventions: stem cell therapy framed as proven rather than experimental, NAD+ IV framed as effective rather than preliminary, hormone optimization framed as universally beneficial rather than appropriately indicated. In each case, AI platforms have absorbed evidence framings that contradict the marketing language commonly used in the category, and clinics that align with the evidence gain citation share that clinics overclaiming cannot recover through marketing investment.

What this means strategically

The editorial position a clinic takes on evidence is a citation strategy. Clinics that acknowledge evidence uncertainty where it exists, and distinguish strong evidence from preliminary evidence, align with the framing AI platforms have already absorbed and gain citation share. Clinics that publish uniform effectiveness claims compete against that framing and lose citation share, often without their marketing teams understanding why.

The strongest single intervention available to most clinics for improving AI citation on evidence-evaluation queries is to audit their own content for overclaims and bring the language into alignment with the established evidence base. This is editorial review work, not content production. It requires understanding both the current evidence base and the specific framings AI platforms have absorbed. Clinics whose clinical and editorial functions can collaborate are best positioned to do the work.

Evidence alignment compounds over time. Once a clinic's content establishes a pattern of evidence-aligned framing, AI platforms cite it across an expanding range of queries, including queries the clinic did not specifically optimize for. The compound effect is one of the strongest citation moats currently available in the category, and it is available to any clinic willing to make the editorial commitment.

SECTION Nº 05

The *twenty-five* clinics

The snapshots that follow cover each of the 25 clinics in the study panel: visibility profile, authority pattern, brand integrity assessment and strategic implication.

ORDERED BY SNAPSHOT NUMBER, WITH VISIBILITY TIER INDICATED ON EACH CARD

Nº 01

Fountain Life

US

● Convergence anchor

Nº 02

Human Longevity Inc

US

● Convergence anchor

Nº 03

Clinique La Prairie

Europe

● Convergence anchor

№ 04

Cleveland Clinic Abu Dhabi

Middle East

● Convergence anchor

№ 05

Biograph

US

● Convergence anchor

№ 06

Chi Longevity

APAC

● Convergence anchor

№ 07

Early Medical

US

● Narrow visibility

№ 08

Sollis Health

US

● Broad visibility

№ 09

Wild Health

US

● Narrow visibility

№ 10

Next Health

US

● Broad visibility

№ 11

Cenegenics

US

● Broad visibility

№ 12

Lanserhof

Europe

● Convergence anchor

№ 13

SHA Wellness Clinic

Europe

● Convergence anchor

№ 14

Chenot Palace

Europe

● Convergence anchor

№ 15

Conradia Medical Prevention

Europe

○ Zero observations

№ 16

Vivamayr

Europe

- Narrow visibility

№ 17

Aman Wellness

Global

- Zero observations

№ 18

AEON Clinic Dubai

Middle East

- Narrow visibility

№ 19

King's College Hospital London Dubai

Middle East

- Narrow visibility

№ 20

Tokyo Midtown Clinic

APAC

○ Zero observations

№ 21

CHA Anti-Aging Center

APAC

● Narrow visibility

№ 22

Everlab

APAC

● Narrow visibility

№ 23

Function Health

US

● Convergence anchor

№ 24

Generation Lab

US

- Narrow visibility

№ 25

Healthy Longevity Clinic Prague

Europe

- Narrow visibility

CLINIC SNAPSHOT № 01

Fountain Life

● Convergence anchor · US

Clinic identity

Fountain Life is a US-based premium longevity clinic founded in 2020 by Peter Diamandis and Tony Robbins, headquartered in Naples, Florida. The clinic operates an expanding multi-location footprint (Naples, Dallas, NYC and 2026 openings in Houston, Miami and Los Angeles), combining AI-driven diagnostic screening with a membership-based annual care model. Its current programs anchor around two tiers: a CORE membership at \$6,500 per year and an APEX Longevity program at \$19,500 per year, positioned by the clinic as the "gold standard" tier. The clinic completed an \$18 million funding round in August 2025.

Within the study, Fountain Life is the most-cited clinic and the canonical demonstration of self-published authority. Its content strategy is the single clearest illustration of how a US-based clinic can establish category-level authority through its own blog rather than through external press accumulation.

Visibility profile

During the study window, Fountain Life surfaced across all four AI platforms studied with anchored or featured citations on dozens of queries, including its own branded queries, executive longevity queries, US preventive medicine queries, biological age testing queries and cost queries. The clinic is one of the panel anchors of the US executive longevity convergence (alongside Human Longevity Inc and Biograph, with Mayo Clinic Executive Health as the non-panel fourth anchor), a position the study identifies as one of the most structurally protected competitive anchors in the entire AI search landscape for longevity care.

The most distinctive feature of Fountain Life's visibility profile in the study sample is its citation density on Google AI Overviews specifically. On queries about biological age testing, longevity clinic cost and Fountain Life versus its direct competitors, fountainlife.com itself appears as a cited

authority source multiple times within the same AI Overview citation panel. The clinic has effectively published its own pages into the position of being among the canonical references AI platforms use to construct answers about the longevity category.

A second distinctive feature, less widely recognized than the AIO citation density, is the cross-cultural reach of Fountain Life's content. In supplementary multilingual testing, Fountain Life's blog appeared as a cited source in non-English-language AI search testing about adjacent markets, alongside locally-anchored operators. No other clinic in the study demonstrated comparable cross-cultural citation reach. The Fountain Life blog is functioning as global longevity-category authority in a way that exceeds the clinic's actual operational footprint.

Authority pattern

Fountain Life's citation moat operates almost entirely through self-published authority. The clinic publishes content at the exact long-tail question phrasings prospective clients type into AI search boxes, structured with the schema and structural extractability markup that allows AI platforms to parse it cleanly. The two mechanisms reinforce each other: the content exists, and it is extractable. Both halves are required, and most clinics in the study have neither.

What makes this work is the editorial standard rather than the content volume. Fountain Life's category-ranking guides include genuine competitors. The cost content is verifiable against external pricing. The clinical claims are consistent across the clinic's own pages and across the third-party sources AI platforms cross-reference. The content survives the editorial scrutiny AI platforms apply, which is precisely why it functions as authority rather than as marketing.

The academic and institutional authority layer is not where Fountain Life wins. The press and community sources AI platforms cited on Fountain Life included skeptically positioned material (MIT Technology Review coverage of unproven treatments in the longevity category, Reddit discussion threads, employee-review aggregator content). On Google AIO specifically, this skeptical framing is partially preserved in the AI's summary of Fountain Life, alongside the clinic's own self-positioning. The result is that Fountain Life's AI search presence is high-volume and substantively accurate but not uniformly positive.

Brand integrity

Fountain Life has one of the cleanest brand-integrity profiles in the study sample. In our sampled AI-search outputs, AI platforms correctly attributed founder identity (Diamandis, Robbins and Dr. Bill Kapp), location footprint, pricing tiers and service stack during the study window. The clinic's distinctive positioning around AI-driven diagnostics, comprehensive screening and Diamandis's preventive-medicine framing is consistently captured across platforms without significant distortion.

One nuance worth noting: In our sampled AI-search outputs, AI platforms surfaced community-discussion content alongside the clinical content. These citations are accurate representations of what AI platforms drew from public discourse, not hallucinations, but they illustrate an asymmetric integrity dynamic: highly cited clinics with active consumer communities can see their AI framing pulled toward the balance of community sentiment, regardless of how strong the clinical content itself is.

Strategic implications

Three structural observations follow from Fountain Life's position in this study's dataset.

First, Fountain Life's self-published authority strategy is the single clearest demonstration in the longevity category of how a clinic can establish AI search visibility without relying on external press accumulation. The strategy has compounded over four years and now produces, in the study sample, visibility that competitors substantially older and more institutionally credentialed do not achieve.

Second, the clinic's visibility profile has structural gaps even within its strongest categories. Single-test specialty queries (whole genome sequencing for longevity, VO2 max testing for longevity) route to specialty operators rather than to Fountain Life despite the clinic's membership including these tests. The academic-authority ceiling on evidence-evaluation queries operates against Fountain Life just as it operates against every clinic in the study; on questions about whether longevity medicine is evidence-based, the citation panels route through peer-reviewed sources rather than through clinic content. These are structural limits of the self-published authority strategy, not failures of execution.

Third, the cross-cultural citation reach of Fountain Life's content is not currently understood by any other clinic in the dataset. A US-headquartered clinic's blog is being cited as authority in AI search about geographically and linguistically distant markets, including in non-English-language testing. No other clinic demonstrates this reach. Understanding the mechanism that produces it, and whether it can be replicated by other clinics willing to make the same editorial investment, is the open question this finding raises.

Snapshot 01 of 25.

CLINIC SNAPSHOT № 02

Human Longevity Inc

● Convergence anchor · US

Clinic identity

Human Longevity Inc was founded in 2013 in San Diego by Dr. J. Craig Venter, the geneticist who co-led the original sequencing of the human genome. Co-founders include Peter Diamandis (who subsequently co-founded Fountain Life) and Dr. Robert Hariri, a stem cell pioneer who also advises Fountain Life. The company operates from San Diego and South San Francisco. Its flagship offering is the Executive Health Screening, starting at \$8,000 and including whole genome sequencing, full-body MRI, brain MRI, cardiac imaging and analysis of more than 120 biomarkers in a single visit. Over the company's first decade, more than \$600 million has been invested in building one of the most comprehensive longitudinal genomic and phenotypic health datasets in the longevity-medicine category. The platform's most distinctive operational guarantee is a \$1 million cancer-care commitment: members who develop late-stage cancer while under HLI's care receive up to \$1,000,000 in treatment contribution.

Two material events in 2026 reshape HLI's competitive position. On April 29, 2026, Dr. J. Craig Venter died at age 79. In May 2026, HLI launched a \$599 AI-enhanced genomic insight platform, framed by Executive Chairman Wei-Wu He as a corporate effort to honor Venter's legacy and broaden access to the underlying genomic infrastructure. The \$599 platform represents the explicit conversion of HLI's founder-credibility moat into a product-credibility moat at a price point an order of magnitude below the \$8,000 Executive Health entry tier. The strategic and editorial implications of this transition are substantial and are the central forward-looking dynamic the snapshot addresses.

Visibility profile

During the study window, Human Longevity Inc surfaced consistently across all four AI platforms studied, with anchored visibility on the US executive longevity category and on the head-to-head comparison against Fountain Life. This positioning as one of the two flagship US executive longevity brands is one of the most structurally protected category anchors in the entire dataset. The four AI platforms converge on the same competitive set for executive longevity assessment (Mayo Clinic

Executive Health, Fountain Life, Human Longevity Inc, Biograph), with Mayo Clinic Executive Health anchoring on all four and the other three each surfacing on three of four. HLI's place within that set is durable.

The visibility profile beyond executive longevity is uneven. HLI surfaces consistently on full-body MRI queries as the integrated-care alternative to specialty operators like Prenuvo and Ezra, on cost-tier queries as the canonical premium-concierge reference and on certain comparison queries (versus Fountain Life specifically). On geographic queries about cities where HLI does not operate, the clinic does not surface; its San Diego and South San Francisco footprint produces narrow geographic-query coverage relative to clinics with multi-city footprints.

One platform-specific finding deserves direct attention. Claude is the only AI platform of the four to have captured Venter's April 2026 death and the May 2026 \$599 platform launch within the study window. The other three platforms answered queries about HLI without acknowledging either event. This is a freshness gap rather than an integrity failure, and the gap will likely close as the platforms update their content over the following weeks. For the moment, Claude is the only AI source giving prospective clients the current corporate context.

Authority pattern

Human Longevity Inc is the dataset's clearest demonstration of academic and institutional authority operating through founder-credibility-via-research-pedigree. Craig Venter's role in the human genome sequencing effort, the peer-reviewed publications on the Health Nucleus cohort (showing clinically significant previously-unknown findings in a meaningful fraction of asymptomatic participants) and the more than \$600 million invested over a decade in longitudinal genomics constitute an authority signal AI search systems in the study more frequently weighted on executive-longevity and genomics queries. HLI surfaces consistently on these queries because the founder-credibility signal is unusually strong even by the standards of the longevity-clinic category.

The founder-credibility flavor of this signal is durable but bounded. It carries HLI on questions about institutional credibility, comprehensive executive assessment and genomic depth. It does not, by itself, carry HLI on long-tail service-specific queries where the question is structured around a specific protocol rather than around institutional credentialing.

One of the most analytically important findings for HLI in the study sample is its absence from the whole genome sequencing for longevity query on three of four platforms. The clinic founded by the geneticist who co-sequenced the human genome, with \$600 million invested in longitudinal genomics over a decade, does not surface as a primary recommendation when prospective clients ask the AI platforms which clinics provide whole genome sequencing for longevity. The AI platforms route the query instead to specialty operators with substantially smaller research investments and substantially less founder credibility, because those operators have done the schema and structural

extractability work that HLI has not. They have published long-tail landing pages targeting the exact phrasing of the query, with the structured data and schema markup that AI search systems in the study more frequently parsed. HLI has not.

This is the cleanest demonstration in the report of schema and structural extractability operating as a filter even against the strongest possible founder-credibility moat. The strategic implication is that schema extractability is not optional for clinics whose core service categories are well-defined and where competitors with stronger schema infrastructure can capture citation share regardless of clinical credentialing. The implication applies to HLI specifically and with unusual force given the asymmetry between the clinic's clinical depth and its current structural-extractability visibility.

Brand integrity

Human Longevity Inc has one of the cleanest brand-integrity profiles in the executive-longevity tier in the study sample. In our sampled AI-search outputs, AI platforms across the dataset consistently and correctly attributed the \$600 million invested-over-decade figure, the 10,000-individual longitudinal dataset, the \$8,000 Executive Health pricing, the \$1 million cancer-care guarantee, the founder lineage (Venter, Diamandis, Hariri), the San Diego and South San Francisco operations and the 150 GB per-client data scale. No hallucinations have been documented in HLI's data across the platforms studied.

Two specific integrity-positive findings deserve attention. First, Claude's capture of Venter's April 2026 death includes accurate attribution to Wei-Wu He's quote on the corporate transition (drawn from PR Newswire coverage). Second, the founder attribution across all four platforms preserves the distinction between Venter as primary founder and Diamandis and Hariri as co-founders, which is the historically accurate framing and is preserved correctly even on platforms that do not yet acknowledge Venter's death. The integrity discipline reflects the strength of HLI's published institutional materials more than it reflects discipline within any individual AI platform.

Strategic implications

Three structural observations follow from HLI's position in the dataset.

First, the schema gap on whole-genome-sequencing queries is the highest-leverage citation-visibility intervention available to HLI in the current state of its content infrastructure. A single well-structured whole-genome-sequencing-for-longevity landing page on humanlongevity.com, built with the FAQ schema markup and the explicit founder-attribution metadata the schema-extractability mechanism rewards, would likely shift the canonical AI search answer on the WGS-for-longevity query from specialty operators to HLI on multiple platforms within a content-indexing cycle. The work is not large in absolute terms. The competitive cost of not doing it is substantial: every query

that currently routes to specialty operators is a prospective client who reaches a less-credentialed provider when HLI is structurally the most credentialed option. Among the structural fixes the report identifies, this one is unusually self-contained and unusually quick to execute.

Second, the \$599 platform launch is the most consequential corporate move of HLI's last decade for AI search visibility. At \$599, the addressable audience is substantially larger than the \$8,000 Executive Health tier's, and the platform creates a citation surface for HLI on consumer-tier diagnostic queries where HLI has not historically competed. Translating this expanded addressable audience into AI citation share requires deliberate content distribution at scale, structured around the schema-extractability and self-published-authority mechanisms the report identifies. Executed well, the 12-to-24-month citation buildout reshapes HLI's visibility profile from executive-only to executive plus mainstream genomics. Executed poorly or not at all, the platform launch produces near-term revenue without proportionate visibility gains.

Third, the founder-credibility moat is now in transition in a way that warrants explicit strategic acknowledgment. Venter's name will continue to carry citation weight for some period; the half-life is bounded but not short. HLI's content strategy should accelerate the transition from founder-name dependence toward institutional-credibility framing: the \$600 million dataset, the peer-reviewed Health Nucleus cohort research, Wei-Wu He's executive leadership and the partnerships with downstream clinical-care providers that translate the genomic platform into clinical action. The transition is partly underway in the May 2026 platform launch communications, but the work is multi-year and the cadence of post-Venter content matters more than any single corporate communication. The clinics best positioned to sustain founder-credibility transitions are those whose institutional infrastructure can carry the citation weight when the founder's individual authority recedes. HLI has more of this infrastructure than most clinics in similar transitions but less than its closest peers (Mayo Clinic, Cleveland Clinic) that operate within fully institutional parent systems.

Snapshot 02 of 25.

CLINIC SNAPSHOT № 03

Clinique La Prairie

● Convergence anchor · Europe

Clinic identity

Clinique La Prairie is a Swiss medical wellness clinic founded in 1931 and headquartered in Montreux on Lake Geneva. It is the original luxury longevity clinic, with a 90-plus-year operating history that predates the modern longevity-medicine category by several generations. The flagship Revitalization program runs as a multi-day inpatient stay starting around 25,000 euros, and the clinic operates a multi-location network that includes Montreux, Madrid, Doha, Dubai (the Longevity Hub at One&Only Za'abeel), Bangkok, Beijing and a forthcoming AMAALA Saudi Arabia outpost opening in autumn 2026. The clinic also operates a consumer-facing supplement and wellness product line under the Holistic Health brand.

For AI search purposes, Clinique La Prairie occupies a structurally different competitive position from most of the clinics in this study. The visibility moat is built not on what the clinic publishes (Fountain Life's strategy) and not on the academic standing of a single founder (Chi Longevity's strategy) and not on institutional brand inheritance from a parent system (Cleveland Clinic Abu Dhabi's strategy). The visibility moat is built on nearly a century of mainstream luxury press writing about Clinique La Prairie as the canonical European luxury wellness destination. The press has done the work of establishing the clinic's authority over decades, and AI platforms cite that accumulated press as authority when constructing answers to category-defining queries.

Visibility profile

Clinique La Prairie was the strongest cross-platform validator in the entire study sample. It is one of the few clinics in the panel that surfaced on all four AI platforms with anchored or featured citations across multiple query categories during the study window. It anchors the brand-direct query, anchors the European luxury wellness category query (alongside Lanserhof, SHA Wellness and Chenot Palace as the European luxury convergence quartet, with Clinique La Prairie, Lanserhof and SHA Wellness anchoring on all four platforms and Chenot Palace on three of four), anchors the head-to-head

comparison against Lanserhof, surfaces with strong visibility on the Dubai geographic query through the Longevity Hub outpost and surfaces as a featured comparison reference on stem cell therapy and longevity queries, executive longevity queries and full-body MRI queries.

Where other strong-visibility clinics (Fountain Life, Human Longevity Inc, Cleveland Clinic Abu Dhabi) lead in citation count on specific query categories, Clinique La Prairie leads in citation breadth. It appears across more category contexts than any other clinic in the study. The accumulated authority operates across luxury wellness, longevity medicine, anti-aging treatment, executive health and adjacent service categories without requiring category-specific content investment. This is the distinguishing feature of the accumulated-press flavor of academic and institutional authority: once established, the citation reach extends across categories without proportionate effort within each category.

The framing language AI platforms use to describe Clinique La Prairie is notably consistent across platforms: "gold standard," "the most famous longevity clinic in the world," "Rolls-Royce of medical wellness centers," and similar institutional-prestige language. This framing is not generated independently by each AI platform. It is drawn from the source material the platforms cite, where the framing has been established and reinforced through luxury-press coverage in publications like WWD, Net-a-Porter, Vogue, Healing Holidays and dedicated wellness-travel intermediation sites.

Authority pattern

Clinique La Prairie demonstrates the accumulated-press flavor of academic and institutional authority. The clinic does not need to publish its own content at the volume Fountain Life publishes because mainstream luxury press has already published the canonical answer pages for queries about European luxury wellness, anti-aging treatment and longevity. The Healing Holidays expert FAQ pages, the WWD "Top Six Longevity Retreats" article, the Net-a-Porter "Most In-Demand Wellness Retreats" feature, the dedicated luxury-wellness intermediation sites: all of these treat Clinique La Prairie as the default reference for European luxury longevity and AI platforms cite them as authority when constructing answers.

The competitive implication is uncomfortable for newer clinics. A clinic founded in 2022 would struggle to compete with Clinique La Prairie in 2026 even with a substantially better content strategy, because the citation infrastructure that anchors Clinique La Prairie's AI search visibility took ninety years of luxury-press coverage to build. The accumulated-press authority moat is real, structurally protected and effectively prohibitive for displacement attempts on a near-term timeline. New entrants to the European luxury wellness category face exactly this barrier, and recognizing it is the highest-priority strategic clarification for any clinic planning to enter the category.

There is a structural fragility worth naming, however. The same accumulated-press citation infrastructure that anchors Clinique La Prairie's visibility is also the surface most exposed to potential erosion. A critical mainstream press feature or an unfavorable academic commentary would directly

affect Clinique La Prairie's AI search visibility because those publications are the same authority sources currently anchoring its citation share. The exposure is real even though it has not materialized. The strategic implication is that PR continuity and editorial relationship maintenance are higher-priority defensive investments for Clinique La Prairie than additional content production. The visibility moat is real but its foundation is in the press archive, not on the clinic's own domain.

Brand integrity

Clinique La Prairie has one of the cleanest brand-integrity profiles in the study, comparable to Fountain Life's record on the US side. AI platforms across all four consistently and correctly attribute the Montreux flagship, the multi-location expansion (which the AI search captures helped surface as a current operational fact during the study window), the Revitalization program structure and pricing, the 90-year operating history and the Lake Geneva positioning. The brand-name string "Clinique La Prairie" is structurally distinctive enough that named-entity-disambiguation failures of the Wild Health variety do not occur for this clinic. No hallucinations have been documented in the Clinique La Prairie data.

One small operational nuance is worth flagging. On the rare queries where Clinique La Prairie is mentioned only in passing as a single-line comparison, the AI response sometimes references the brand without specifying which location. For a prospective client researching Dubai care specifically, the Longevity Hub Dubai outpost at One&Only Za'abeel is the actual local service, not the Montreux flagship. This is a low-priority refinement rather than an integrity issue, and the clinic's own location-disambiguation infrastructure (the longevity-hub.cliniquelaprairie.com subdomain) is already surfacing as a citation source on relevant queries, which suggests the clinic is partway through addressing this.

Strategic implications

Three observations follow from Clinique La Prairie's position in the dataset.

First, the accumulated-press citation moat is one of the most durable authority assets in the longevity-care category and one of the most structurally protected from displacement. It is also one of the least replicable, because the underlying mechanism is time. A clinic that wants to occupy comparable visibility in the European luxury wellness category cannot do so through content investment alone within a near-term horizon. Recognizing this is the most important strategic insight competitors should extract from this study.

Second, the cross-platform consistency of Clinique La Prairie's citations means it functions as a structural canonical reference for European luxury wellness. When AI platforms need to anchor an answer about premium European longevity care, Clinique La Prairie is the default selection on all four platforms. This is a substantial moat that newer European clinics must work around rather than

confront directly. The strategic play for new European entrants is category creation (positioning within an emerging sub-category where Clinique La Prairie does not occupy the anchor) rather than category competition (attempting to displace Clinique La Prairie within established categories).

Third, the Gulf expansion underway in 2026 is strategically consequential beyond its commercial implications. Clinique La Prairie's AMAALA Saudi Arabia outpost, opening in autumn 2026 on the Red Sea coast, positions the clinic at the geographic axis the report identifies as the longevity category's emerging center. The Gulf is not a peripheral market for a European luxury clinic; it is the market where the category's regulatory infrastructure (Abu Dhabi's Department of Health framework), funding infrastructure (Saudi Arabia's Hevolution Foundation) and luxury-expansion footprint are converging. Clinique La Prairie's Gulf positioning is, in effect, a continuation of the same accumulated-press authority strategy applied to the geography that will likely produce the next generation of luxury-wellness press coverage. The strategic question is whether the AI search visibility infrastructure built on European luxury press transfers cleanly to Gulf luxury press, or whether the Gulf citation ecosystem operates on different rules and requires parallel investment. The empirical evidence from the broader dataset suggests the transfer is meaningful but not automatic.

Snapshot 03 of 25.

CLINIC SNAPSHOT № 04

Cleveland Clinic Abu Dhabi

● Convergence anchor · Middle East

Clinic identity

Cleveland Clinic Abu Dhabi is the UAE flagship of the Cleveland Clinic health system, headquartered in Abu Dhabi and operating as one of the region's most institutionally credentialed academic medical centers. Its Longevity Medicine Department, launched in January 2024 within the Fatima bint Mubarak Center (the clinic's Cancer Institute), is structurally distinct from the executive-longevity model most other clinics in this study represent. The department is led by Dr. Fawad Khan and operates three core programs: a Hereditary High-Risk Clinic for genetic counseling and enhanced screening of patients with elevated familial cancer risk; a Cancer Survivorship Program for post-treatment surveillance and quality-of-life optimization; and a Lifestyle Medicine Clinic providing nutrition, exercise, sleep and stress management for both cancer-risk patients and the broader preventive-health client base.

This positioning is unusual within the longevity-medicine category and deserves direct treatment. Cleveland Clinic Abu Dhabi's Longevity Medicine Department is not a biohacking clinic, not an executive-optimization service and not a luxury wellness retreat. It is hospital-integrated preventive medicine with a cancer-survivorship anchor, operating from inside a tertiary-care academic medical center. The Executive Health Program (a one-day comprehensive assessment with advanced diagnostics) is the clinic's most direct competitor to the executive-longevity offerings of clinics like Fountain Life or Biograph, but it sits within a fundamentally different institutional structure. Pricing for the Executive Health Standard tier is approximately AED 13,500, with the Premium tier around AED 20,500 (roughly USD 3,500 to 5,500). This is substantially below the pricing of European luxury wellness clinics or US premium executive-longevity programs, and the cost gap is editorially relevant: Cleveland Clinic Abu Dhabi competes on institutional credibility and clinical depth rather than on luxury positioning or biohacking framing.

Visibility profile

During the study window, Cleveland Clinic Abu Dhabi surfaced consistently across all four AI platforms studied, with anchored visibility on two distinct query patterns: its own brand-direct query and the head-to-head comparison with King's College Hospital London Dubai for institutional Middle East longevity care. Across the broader panel, the clinic also appears as comparison reference on multiple US executive longevity queries, on the European luxury versus institutional positioning queries and on practitioner-driven queries (Dr. Fawad Khan surfaces as a named practitioner anchor on at least one platform).

The visibility profile is qualitatively different from the executive-longevity clinics in the study. Where Fountain Life surfaces because AI platforms cite its own blog content, and where Clinique La Prairie surfaces because AI platforms cite decades of accumulated luxury press, Cleveland Clinic Abu Dhabi surfaces because the Cleveland Clinic name itself carries institutional authority that AI search systems in the study more frequently weighted. The brand recognition is the citation engine. The Longevity Medicine Department's own content is not the dominant authority source; the parent institution's reputation is.

This produces an unusual citation pattern. On the head-to-head query against King's College Hospital London Dubai, all four AI platforms anchor on Cleveland Clinic Abu Dhabi with explicit framing as the institutionally stronger option ("No.1 in UAE, top smart hospital," "academic medical center," "evidence-led winner"). The framing language across platforms is remarkably consistent, and the cited source mix in the study shifted toward institutional press and academic citations rather than the wellness-travel publications that anchor European luxury queries.

Authority pattern

Cleveland Clinic Abu Dhabi is the clearest demonstration in the dataset of academic and institutional authority operating through brand-inheritance rather than through direct research publication. The parent Cleveland Clinic institution's accumulated peer-reviewed research, mainstream-press coverage and institutional reputation function as authority signals that AI platforms apply to the Abu Dhabi flagship without requiring the Longevity Medicine Department to have built its own independent citation footprint.

This is structurally important. The Longevity Medicine Department launched in 2024 and is, in conventional content terms, a young service line. A standalone clinic of similar age would face an authority deficit on most of the queries Cleveland Clinic Abu Dhabi anchors. The brand inheritance closes that deficit. AI platforms treat the Abu Dhabi facility as carrying Cleveland Clinic's full institutional authority, which produces visibility that newer Middle East operators (even well-funded ones) cannot match without parallel institutional infrastructure.

The pattern has limits. Cleveland Clinic Abu Dhabi's visibility is strongest where the institutional framing matters most (head-to-head against other hospital-integrated options, executive health, comprehensive assessment). It is weakest where the question is specifically about longevity-as-optimization (the executive longevity assessment category dominated by Fountain Life, Human Longevity Inc and Biograph) because the clinic's positioning explicitly is not optimization-focused. The brand inheritance helps where the question is "which institution should I trust"; it helps less where the question is "which longevity protocol should I adopt." Cleveland Clinic Abu Dhabi's positioning aligns with the first framing and the visibility profile reflects that orientation.

Brand integrity

Cleveland Clinic Abu Dhabi has one of the cleanest brand-integrity profiles in the study. AI platforms across all four consistently and correctly attribute the January 2024 launch date, the Fatima bint Mubarak Center location, Dr. Fawad Khan as department lead, the three-program structure (Hereditary High-Risk, Cancer Survivorship, Lifestyle Medicine) and the Executive Health Program pricing tiers. Specific factual content (clinician names, program names, pricing in AED) is verifiable across platforms with no notable discrepancies.

A specific integrity-positive finding deserves attention. Claude's response to the brand-direct query about Cleveland Clinic Abu Dhabi includes an explicit editorial disambiguation: the clinic's Longevity Medicine Department is "structured differently from typical anti-aging or executive longevity clinics" and "not a biohacking clinic." This kind of category-correction framing is unusual in AI search responses and is editorially valuable for clinics whose positioning is genuinely distinct from the dominant category framing. Cleveland Clinic Abu Dhabi benefits from AI platforms doing the editorial work of clarifying its category positioning rather than misframing it as a peer of biohacker-optimization clinics. Few clinics in the study receive this kind of platform-side editorial support.

Strategic implications

Three structural observations follow from Cleveland Clinic Abu Dhabi's position in the dataset.

First, the brand-inheritance authority mechanism is one of the most powerful citation moats in the study, available only to clinics operating within institutionally credentialed parent systems. For Cleveland Clinic Abu Dhabi, the visibility is real, durable and difficult for non-institutional competitors to match. For other clinics in the longevity category, the implication is that the institutional-credibility ceiling is genuinely high and that attempts to compete with brand-inheritance authority through marketing investment alone will systematically underperform. The asymmetry is structural.

Second, the cancer-survivorship anchor is a strategic differentiator that AI platforms have correctly captured but that most prospective clients searching for longevity care will not initially understand. A high-net-worth executive searching for executive longevity assessment will not naturally route to Cleveland Clinic Abu Dhabi because the clinic's positioning is not optimization-focused. A patient with familial cancer risk, post-cancer surveillance needs or institutional-medical preferences will route to Cleveland Clinic Abu Dhabi correctly. The clinic's challenge is not visibility (it has visibility); it is whether the AI-search audience reaching the clinic matches the clinical service the clinic actually provides. The fit between citation visibility and intended client base is its own strategic consideration that few clinics in the dataset have addressed deliberately.

Third, the Cleveland Clinic Abu Dhabi case demonstrates that the Middle East regulatory and institutional infrastructure is producing AI search visibility outcomes that match the region's emerging strategic position in the longevity category. Abu Dhabi's Department of Health licensing of IHLAD, the Cleveland Clinic Abu Dhabi Longevity Medicine Department and the regional clustering of institutional medical infrastructure are reinforcing each other in AI search citations. This is a positive flywheel for clinics operating within the Gulf institutional layer and is one of the most strategically consequential dynamics for clinics planning regional or international expansion.

Snapshot 04 of 25.

CLINIC SNAPSHOT № 05

Biograph

● Convergence anchor · US

Clinic identity

Biograph is a US-based premium longevity diagnostics platform founded in 2023, co-founded by John Hering, James Burgess and Peter Attia, headquartered in the San Francisco Bay Area with a New York City flagship. The clinic operates a members-only model with two pricing tiers: Core membership at \$7,500 per year and Black membership at \$14,500 per year. The core offering is a single-day comprehensive diagnostic assessment combined with ongoing concierge longevity care, integrating MRI, DEXA, VO2 max, biomarker testing and Attia-aligned clinical protocols. Biograph has external venture capital backing from Vy Capital, Human Capital, Alpha Wave Global and WndrCo, with angel investors including Balaji Srinivasan. The clinic's most-cited outcome statistic, surfacing across multiple AI platforms, is that approximately one in six members have uncovered urgent or life-altering findings during assessment, including early-stage cancers.

Biograph's market positioning is structurally important for understanding its AI search visibility profile. The clinic occupies a distinctive position in the US executive longevity tier: not the largest, not the most institutionally credentialed, not the oldest, but arguably the most contemporary in framing and the most directly associated with Peter Attia's clinical philosophy. The Attia association is the single most important fact about Biograph's positioning, and it is also the source of both the clinic's strongest visibility asset and its most analytically interesting integrity vulnerability.

Visibility profile

During the study window, Biograph surfaced consistently across all four AI platforms studied, with anchored visibility on the US executive longevity category (as one of the four members of the convergence set, alongside Mayo Clinic Executive Health, Fountain Life and Human Longevity Inc), on the NYC geographic query, on the head-to-head comparison against Fountain Life, on hormone optimization queries and on whole genome sequencing for longevity queries. The visibility profile is unusually broad given the clinic's recent founding date. Most clinics with two years of operating history do not surface across this many query categories with anchored framing.

Biograph's place in the executive longevity convergence set is the most strategically important visibility fact for the clinic. The four AI platforms converge on a stable four-clinic set for the category (Mayo Clinic Executive Health, Fountain Life, Human Longevity Inc, Biograph); Mayo Clinic Executive Health anchors on all four, and Biograph, like Fountain Life and Human Longevity Inc, anchors on three of four. This is one of the most structurally protected category positions in the entire dataset, and Biograph occupies it within two years of founding. The achievement is unusual and reflects the strength of the Attia association as an authority signal that bypasses the long content-buildout cycle most clinics require.

A specific platform-by-platform variance is worth noting. Perplexity surfaces Biograph more prominently than other platforms on the Peter Attia clinic query, prioritizing Biograph above Early Medical (Attia's primary clinical practice). The other three platforms route the Attia clinic query to Early Medical first, with Biograph as the secondary or Attia-adjacent reference. This is a platform-specific framing difference rather than an accuracy variance, but it has commercial implications: Perplexity traffic on Attia-related searches converts toward Biograph at a higher rate than other platforms' Attia-related traffic.

Authority pattern

Biograph demonstrates two patterns operating in tandem: academic and institutional authority via Attia-anchored founder credibility, and schema and structural extractability via the clinic's structured content and clearly defined service-tier architecture. The institutional-authority mechanism is the dominant visibility driver. AI search systems in the study more frequently weighted the Attia association because Attia's content empire (podcast, book, blog, professional network) functions as a credibility anchor that propagates to clinics he is associated with. Biograph inherits the citation share that flows to Attia's broader category authority, in a structurally similar way to how Cleveland Clinic Abu Dhabi inherits authority from its parent system and how Chi Longevity inherits authority from Andrea Maier's individual academic credentialing.

The schema-extractability mechanism reinforces the institutional-authority visibility. Biograph's website is structurally clean, the pricing tiers are explicit (\$7,500 Core, \$14,500 Black), the service stack is named precisely (MRI, DEXA, VO2 max, biomarker testing, longevity-medicine protocols) and the outcome statistic is presented in extractable form (one in six members with urgent or life-altering findings). AI platforms can parse the content cleanly and cite it accurately on category, geographic and comparison queries. The schema-extractability work is doing exactly what the report identifies as the highest-leverage technical investment in AI visibility.

The combination produces visibility that is unusually strong for a two-year-old clinic, but it also creates a specific structural vulnerability that is worth naming directly.

Brand integrity

Biograph is the subject of one of the two notable brand-integrity findings in this report. Google AI Overviews, on a query about clinics resembling Peter Attia's practice, described Biograph as "Dr. Peter Attia's own personalized medicine clinic in NYC." The framing oversimplifies the actual ownership structure: Biograph is a separately incorporated venture co-founded by John Hering, James Burgess and Peter Attia, with external venture capital backing, not a wholly-owned Attia practice. The report frames this as an oversimplification rather than a hallucination because Biograph is genuinely Attia-co-founded; the imprecision is in the "Attia's own" framing, which implies sole ownership and elides the co-founder, the venture-capital structure and the company's independent corporate identity.

The integrity finding has implications beyond the specific framing error. The mechanism that produced the oversimplification is the same mechanism that produces Biograph's strong visibility: In our sampled AI-search outputs, AI platforms absorbed the Attia association as the dominant frame and treated Biograph as an extension of Attia rather than as an independent entity. The visibility benefit and the integrity vulnerability are two faces of the same dynamic. As long as the Attia association is the primary authority signal AI platforms weight for Biograph, the framing imprecision will recur in some form and Biograph's independent corporate identity will be intermittently underrepresented in AI search.

The other 41 query-platform observations where Biograph surfaces show clean integrity. The Core and Black pricing tiers are correctly attributed across platforms. The co-founder relationship between Attia and John Hering is correctly attributed where it surfaces. The venture-backing structure is accurately described. The one-in-six urgent-findings statistic is consistently attributed. No hallucinations beyond the single AIO oversimplification have been documented. The integrity profile is strong, with the single named vulnerability being a structural function of the visibility mechanism rather than a content-quality failure.

Strategic implications

Three observations follow from Biograph's position in the dataset.

First, the Attia-anchored visibility mechanism is one of the most powerful authority signals available in the longevity-care category and Biograph has captured it more cleanly than any other clinic in the dataset. The strategic asset is real and the citation share it produces is durable as long as the Attia association remains current. The question worth examining deliberately is what proportion of Biograph's AI search visibility would survive a meaningful change in the Attia association: Attia stepping back from active media, Attia's clinical practice shifting or a market evolution that reduces the citation weight of individual practitioner credentialing relative to other authority signals. The current visibility looks like an asset, and it is; the durability question is the unasked strategic question.

Second, the oversimplification finding is editorially important beyond Biograph specifically. The mechanism that produced it (AI platforms absorbing the strongest single association as the dominant frame and underweighting independent corporate identity) is structurally relevant to any clinic whose visibility flows primarily through a founder, practitioner or parent-brand association. The recurring framing imprecision is not a Biograph problem; it is a structural feature of how AI platforms compress complex corporate structures into citation-friendly framings. For Biograph specifically, the implication is that explicit content investment in the John Hering co-founder relationship, the venture-backing structure and Biograph's independent operational identity would reduce (though not eliminate) the recurring oversimplification and would also begin diversifying the authority signals AI platforms weight for the clinic.

Third, Biograph occupies one of the four canonical positions in the US executive longevity category convergence set, alongside Mayo Clinic Executive Health, Fountain Life and Human Longevity Inc. This position is structurally protected and would be very difficult for a new entrant to displace. The strategic question for Biograph is not how to enter the category (it has done that), but how to deepen the position. The visibility moat is real but it sits primarily on the Attia association. Building parallel authority infrastructure (independent peer-reviewed research, named clinicians beyond Attia, self-published content under the Biograph brand rather than under Attia's content empire) converts the current visibility from association-dependent to institutionally-distributed. This is the same structural recommendation Chi Longevity and Human Longevity Inc receive in their respective snapshots, applied to a different specific mechanism. The pattern is generalizable: clinics whose visibility flows through a single anchor are stronger in the near term and more vulnerable in the long term than clinics with diversified authority infrastructure.

Snapshot 05 of 25.

CLINIC SNAPSHOT № 06

Chi Longevity

● Convergence anchor · APAC

Clinic identity

Chi Longevity is a Singapore-based premium longevity clinic founded in March 2023 and co-founded by Professor Andrea Maier, who serves as the clinic's clinical and scientific anchor. The clinic operates from two Singapore locations: a flagship at Camden Medical Centre and a second location at the Four Seasons Hotel, which opened in early 2025. Its core offering is the Maier Method, a personalized longevity assessment and intervention framework structured across three tiers: Vital Start (introductory), Momentum (intermediate) and Pinnacle (comprehensive). The clinic offers a non-resident program structure with 10-month engagement requiring two physical visits supplemented by teleconsultation, designed for international clients who cannot relocate to Singapore for the full program. Pricing starts at approximately SGD 1,800.

Andrea Maier's professional positioning is structurally important for understanding Chi Longevity's AI search visibility profile. She holds appointments at the National University of Singapore, has authored over five hundred peer-reviewed publications in the longevity-medicine field, serves as founding president of the Healthy Longevity Medicine Society and designed the protocols underlying the Abu Dhabi Department of Health's clinical guidelines for healthy longevity medicine. She is, by a meaningful margin, the most academically credentialed individual practitioner anchoring any of the 25 clinics in this study. Chi Longevity's positioning is built on that credentialing in a way that distinguishes it from every other Asia-Pacific clinic in the dataset and from most clinics globally.

Visibility profile

During the study window, Chi Longevity surfaced consistently across all four AI platforms studied. Its brand-direct query anchored on all four platforms with full visibility (Andrea Maier named as scientific anchor, Four Seasons location captured, Maier Method framework described, pricing range cited). On the broader Singapore market query, three of four platforms anchor Chi Longevity as the canonical Singapore longevity recommendation; Google AI Overviews refuses to generate an answer at all for the "best longevity clinic in Singapore" query template, a manifestation of the geographic-template suppression behavior documented elsewhere in the report.

The clinic surfaces consistently as the APAC anchor on category queries (executive longevity assessment, biological age testing, hormone optimization, whole genome sequencing for longevity, stem cell therapy and longevity) across platforms. This positioning as the canonical Asia-Pacific recommendation is one of the more stable cross-platform anchors in the study. Where European luxury queries anchor on the Clinique La Prairie quartet and US executive longevity queries anchor on Fountain Life and Human Longevity Inc, APAC queries route through Chi Longevity with unusual consistency across the four platforms.

Andrea Maier herself surfaces on practitioner-led queries. AI platforms answering questions about top longevity doctors in 2026 surface Maier alongside Peter Attia, Mark Hyman and Craig Venter, with Chi Longevity attached as her clinical practice. This practitioner-to-clinic resolution is unusually clean for the APAC region, where most clinics in the broader longevity-care category do not have a single named practitioner serving as the public credentialing anchor.

Authority pattern

Chi Longevity is the cleanest demonstration in the dataset of academic and institutional authority operating through individual practitioner credentialing rather than through institutional brand inheritance. Andrea Maier's peer-reviewed publication record, her role as founding president of the Healthy Longevity Medicine Society, and her authorship of the Abu Dhabi regulatory framework are the citation engine. AI search systems in the study more frequently weighted these signals when answering questions about scientifically credible longevity care, and Chi Longevity inherits the citation share that flows to Maier.

This is distinct from self-published authority, which Fountain Life demonstrates through volume and structure of its own blog content. It is distinct from the institutional-inheritance variant, which Cleveland Clinic Abu Dhabi demonstrates through brand inheritance from its parent system. Chi Longevity is the practitioner-credentialing variant of academic and institutional authority: the founder's academic standing is sufficient to anchor the clinic's authority without the volume of self-published content, the institutional parent or the accumulated luxury press that the other variants require.

The mechanism is durable but creates a specific vulnerability worth naming. Chi Longevity's AI search visibility is heavily dependent on Andrea Maier's continued academic visibility and credentialing. Should Maier's institutional position shift, or should her authoring of the Healthy Longevity Medicine Society and the Abu Dhabi standards become historical rather than current credentialing, the citation flow that anchors Chi Longevity would attenuate. The clinic's strategic interest is in building parallel authority infrastructure (institutional partnerships, additional credentialed practitioners, self-published content) that does not depend exclusively on its founder's individual credentialing. This is a longer-cycle investment than the current visibility might suggest, and the urgency of beginning it now is greater than the current strong visibility might imply.

Brand integrity

Chi Longevity has one of the cleanest brand-integrity profiles in the study sample. In our sampled AI-search outputs, AI platforms across all four consistently and correctly attributed Andrea Maier as co-founder and clinical lead, the Camden Medical Centre and Four Seasons Hotel locations, the Maier Method framework, the three-tier program structure, the non-resident 10-month program option and the pricing range. Andrea Maier's academic credentials, her role at NUS, and her connection to the Abu Dhabi regulatory framework are accurately captured across platforms. No hallucinations have been documented in Chi Longevity's data.

The integrity profile reflects the clinic's positioning discipline. Chi Longevity has made specific positioning claims (evidence-based biological age reversal, Maier Method as a structured framework, premium pricing tied to specific protocols rather than to luxury experience) and the AI platforms have absorbed those claims accurately. The clinic does not overclaim across the categories where evidence is preliminary, and the absence of overclaims is itself an integrity-positive signal that distinguishes Chi Longevity from clinics whose marketing language outpaces evidence support.

Strategic implications

Three observations follow from Chi Longevity's position in the dataset.

First, practitioner-credentialing authority is one of the most efficient citation moats currently available in the longevity-care category, provided the clinic has a founder with sufficient academic standing to anchor it. Chi Longevity has achieved cross-platform visibility comparable to clinics with substantially larger operational footprints and substantially older content infrastructure, because the practitioner credentialing carries citation weight that takes years of content investment or decades of institutional accumulation to replicate. For clinics without a similarly credentialed founder, this mechanism is structurally unavailable. For clinics with one, it is the highest-leverage authority asset they have.

Second, the founder-dependency built into the current visibility profile is a structural risk that warrants explicit mitigation. The clinic's strategic position would strengthen materially through investment in parallel authority infrastructure: additional credentialed clinicians whose own academic standing supports citation share, institutional partnerships that produce separate authority signals and self-published content that builds visibility independent of the practitioner-credentialing layer. The work is multi-year and does not produce immediate visibility gains, but it converts the current visibility from founder-dependent to institutionally-distributed, which is the durable form.

Third, Chi Longevity is the canonical APAC anchor in this study and the clinic's positioning as the regional reference for serious longevity medicine has commercial implications that extend beyond Singapore. Andrea Maier's authoring of the Abu Dhabi Department of Health framework means the clinic is positioned at the intersection of Asia-Pacific clinical operations and Gulf regulatory

infrastructure. This is unusual in the dataset and structurally valuable for clinics planning expansion across both regions. The strategic question for Chi Longevity is whether to build on this cross-regional positioning explicitly through expansion, partnership or licensing or to maintain the current Singapore-centered operating model and accept that the broader citation reach is currently unmonetized.

Snapshot 06 of 25.

CLINIC SNAPSHOT № 12

Lanserhof

● Convergence anchor · Europe

Identity

Lanserhof is a heritage Austrian medical-wellness brand founded in 1984, combining the Mayr-medicine tradition of gut health with sleep and movement science and luxury retreat operations. It runs four locations, the original at Lans in Austria plus Tegernsee and Sylt in Germany and a London satellite, with a fifth opening in Marbella in 2027; base programs run from roughly 5,900 to 8,175 euros a week before add-ons. Its distinctive position in the panel is as a member of the European luxury wellness set that AI search treats as a category anchor, consistently framed as the most medically rigorous of that group.

Visibility profile

Lanserhof was one of the strongest cross-platform performers in the study. It anchored its own brand-direct query across all four platforms, anchored the head-to-head against Clinique La Prairie on all four and anchors the European anti-aging category query as a member of the four-clinic set that produces the dataset's single strongest cross-platform convergence. It also appears as a comparison reference across executive, cost and pricing queries.

The strongest single finding is that Lanserhof holds one of the three category-anchor positions in the entire study, the European luxury wellness convergence set and within that set it consistently carries the medically rigorous framing.

Authority pattern

Lanserhof demonstrates academic and institutional authority in its accumulated-press variant, the same family as Clinique La Prairie. In the study sample, its citations were built on the wellness-travel review and directory ecosystem, luxury press and its own domain on brand-direct queries, rather than on a single founder or a research arm.

The pattern is durable: convergence-set membership is reinforced every time the platforms train on each other's citations, and the position would be effectively prohibitive for a new entrant to displace. The structural caveat is the one the Clinique La Prairie snapshot also names: the foundation sits in third-party press and aggregators rather than on the clinic's own domain.

Brand integrity

Integrity was clean in the study sample. In our sampled AI-search outputs, AI platforms described Lanserhof's Mayr-medicine methodology, four-location footprint, programs and pricing accurately, and several produced balanced "worth it" and "not worth it" framing, including notes on aspects of the Mayr diet that are not supported by current evidence. No hallucinations surfaced in its data.

Strategic implication

Lanserhof occupies one of the three structurally protected category-anchor positions the study identifies, holding the most medically rigorous framing within the European luxury set. That position is a genuine moat, reinforced by cross-platform citation reinforcement and very difficult to displace. But the citation infrastructure under it sits in the wellness-travel and luxury-press ecosystem rather than on Lanserhof's own domain, the same structural fragility the Clinique La Prairie snapshot identifies for the accumulated-press variant of academic and institutional authority: the moat is real, but its foundation is in the press archive. Snapshot 03 covers the closest parallel.

Snapshot 12 of 25.

CLINIC SNAPSHOT № 13

SHA Wellness Clinic

● Convergence anchor · Europe

Identity

SHA Wellness Clinic is a Spanish luxury medical-wellness destination founded in 2008 by Alfredo Bataller Parietti and still family-owned, with its flagship at L'Albir on the Alicante coast. It operates a three-location footprint, Spain since 2008, Mexico's Riviera Maya since 2024 and a UAE site opening in 2027, running structured longevity programs across detox, cellular regeneration and performance, with a 300-strong professional team and a scientific advisory board. Its distinctive position in the panel is as the most expansion-oriented member of the European luxury wellness set, framed by AI platforms as the category's "hottest" destination.

Visibility profile

During the study window, SHA anchored its own brand-direct query across all four platforms, anchored the European anti-aging category query as a member of the four-clinic convergence set and anchored the head-to-head against Chenot Palace on all four platforms. It also appears as a comparison reference on executive, cost and service queries.

The strongest single finding is that SHA's brand-direct query produces unusually specific and accurate pricing across all four platforms, because the platforms are drawing on the clinic's own structured pricing pages.

Authority pattern

In the study sample, SHA demonstrated academic and institutional authority in its accumulated-press variant for category visibility, the same family as Clinique La Prairie and Lanserhof and a self-published authority element on brand-direct and pricing queries, where its own structured pricing content is what the platforms cite.

The combination is durable. Convergence-set membership is the protected category asset, and the "World's Best Wellness Clinic" award and scientific advisory board are credentialing signals that reinforce it, while the own-domain pricing content gives SHA a second, clinic-controlled channel that most of the European luxury set does not show as clearly.

Brand integrity

Integrity was clean and notably precise in the study sample. In our sampled AI-search outputs, AI platforms described SHA's program structure, three-location footprint and pricing with high specificity and cross-platform consistency, and the pricing tables were verifiable against the clinic's own published rates during the study window. No hallucinations surfaced in its data.

Strategic implication

SHA is the clearest case in the European luxury set of own-domain content discipline reinforcing an accumulated-press position. It holds a convergence-set category anchor through the press ecosystem, and on brand-direct and pricing queries its own structured content produces unusually accurate representation, which means the two channels are complementary rather than alternatives. The open question is its multi-continent expansion: whether the European-press-anchored visibility transfers to Mexico and the UAE, or whether those markets run on different citation ecosystems and require their own investment, the same transfer question the Clinique La Prairie snapshot raises for the Gulf. Snapshots 03 and 12 cover the closest parallels.

Snapshot 13 of 25.

CLINIC SNAPSHOT № 14

Chenot Palace

● Convergence anchor · Europe

Identity

Chenot Palace is the newer "palace" brand built on the Chenot Method, distinct from the heritage Espace Henri Chenot, with its flagship opened on Lake Lucerne in 2020 and additional locations in Azerbaijan and Italy. The method itself, established by Henri Chenot roughly fifty years ago, anchors the offering: a fasting-mimicking dietary protocol of around 850 calories a day, a signature seven-day Advanced Detox and traditional-Chinese-medicine integration. Its distinctive position in the panel is as the youngest member of the European luxury wellness set, a recent brand carrying an established method's authority.

Visibility profile

In the study sample, Chenot Palace's visibility was concentrated in two contexts. It anchored the head-to-head against SHA Wellness on all four platforms, and it appeared in the European anti-aging category query as a member of the convergence set on three of the four platforms, missing only ChatGPT, where the fourth slot goes to a different clinic. Outside those category and comparison contexts it surfaces little.

The strongest single finding is that Chenot Palace holds a category-anchor position on three of four platforms while carrying the narrowest overall query footprint of the four convergence-set members.

Authority pattern

Chenot Palace demonstrates academic and institutional authority in its accumulated-press variant, with a specific twist: the accumulated authority belongs to the Chenot Method rather than to the young Palace brand. In the study sample, luxury press and the clinic's own domain were the visible citation sources, and the brand inherited the method's standing.

The pattern is genuinely strong but slightly indirect, because it is method-inherited rather than brand-built, and the position is held less unconditionally than the rest of the set: the narrowest footprint of the four and a category-set slot already lost on one platform, indicate a position that is secure but not fully settled.

Brand integrity

Integrity was clean in the study sample. In our sampled AI-search outputs, AI platforms correctly attributed the Chenot Method's origins to Henri Chenot, the fasting-mimicking dietary protocol and the signature program structure, and handled the method's heritage accurately where it surfaced. No hallucinations surfaced in its data.

Strategic implication

Chenot Palace is the panel's clearest case of a young brand holding a category-anchor position through inherited method authority. The "palace" brand is only a few years old, but the Chenot Method's long heritage carries it into the European luxury convergence set. The strategic observation is that this is accumulated-press authority where the accumulation belongs to the method rather than the brand, which makes the position both real and slightly indirect and the narrowest query footprint of the four members, plus the slot already lost on one platform, suggests the position is held but less settled than Clinique La Prairie's or Lanserhof's. Snapshots 03, 12 and 13 cover the rest of the set.

Snapshot 14 of 25.

CLINIC SNAPSHOT № 23

Function Health

● Convergence anchor · US

Identity

Function Health is a membership-based diagnostic platform founded in 2023 and headquartered in Austin, Texas, operating nationally through telehealth and the Quest Diagnostics network. Its model is comprehensive lab testing, more than 160 biomarkers twice a year, on a membership reduced during the study window from 499 to 365 dollars a year; Mark Hyman is chief medical officer and it raised a Series A in 2024. Its distinctive position in the panel is as the clearest mass-market diagnostic disruptor, democratized longevity diagnostics at a low annual price, distinct from the concierge and luxury clinics.

Visibility profile

Function Health was one of the strongest and cleanest performers in the study sample. It anchored its own brand-direct query across all four platforms, anchored the head-to-head against InsideTracker on all four and was the canonical entry-tier value pick across the cost and value frameworks in the dataset. It also surfaced via Mark Hyman on practitioner-led queries and as a comparison reference elsewhere.

The strongest single finding is that, in the study sample, Function Health's brand-direct citation share was essentially entirely its own domain across all four platforms, the cleanest single demonstration of self-published authority in the study.

Authority pattern

In the study sample, Function Health demonstrated self-published authority in its cleanest form: on its brand-direct query, the citation panel was dominated by its own content on every platform. A community citation authority element reinforced it through an active consumer community that platforms cited for pricing and critical context, and an academic-and-institutional-authority founder signal ran through Mark Hyman on practitioner queries.

The pattern is durable and demonstrably responsive: the mid-study membership price change propagated accurately across all four platforms because Function Health published it and the platforms picked it up.

Brand integrity

Integrity was clean in the study sample. The pricing change propagated consistently, and biomarker counts, the membership model and founder attribution were accurate across platforms, including on the InsideTracker comparison. Notably, Function Health has a semantically generic name, structurally exposed to the disambiguation failure documented for Wild Health, but its own-domain authority is strong enough that entity resolution held: strong self-published authority is, among other things, a defense against the schema-and-structural-extractability disambiguation problem.

Strategic implication

Function Health is the dataset's cleanest demonstration of self-published authority done well. Its brand-direct citation panel is almost entirely its own domain across all four platforms, its pricing change propagated accurately because it published it and an active consumer community reinforces the position. It has built, in a short time, the self-published-authority position most panel clinics are still trying to reach, through content discipline and a clear, low-priced consumer proposition rather than press accumulation or a research arm. Snapshot 01 covers Fountain Life, the other strong self-published-authority case; Snapshot 09 covers Wild Health, the generic-name counter-case.

Snapshot 23 of 25.

CLINIC SNAPSHOT № 08

Sollis Health

● Broad visibility · US

Identity

Sollis Health is a concierge medicine network founded in 2016, backed through a Series C round, with membership-based locations across several premium US markets including New York, Los Angeles and the Hamptons. Its core offering is 24/7 concierge urgent care and private-emergency-room access, on a membership starting in the low single-thousands of dollars per year, with a top tier that adds longevity-adjacent benefits such as vitamin infusions. Its distinctive position in the panel is that it is fundamentally a reactive-care brand, not a dedicated longevity clinic: the clearest case in the study of a clinic present in the longevity conversation by adjacency rather than category fit.

Visibility profile

During the study window, Sollis had strong, clean visibility on a narrow set of query types and little elsewhere. It anchored its own brand-direct query on all four platforms, and anchors the head-to-head against Next Health on all four, where the platforms frame it as the reactive, urgent-care counterpart to Next Health's proactive optimization. It does not surface on the open category, service or geographic queries that define the longevity market.

The strongest single finding is that AI platforms classify Sollis accurately and consistently as concierge urgent care, and that accurate classification is exactly what bounds its longevity-category visibility to near zero.

Authority pattern

Sollis demonstrates self-published authority on its brand-direct surface, where its own domain dominates the citation panel, plus a community citation authority footprint from a real concierge-client community that platforms cite for member feedback. The more important observation is that this authority is bounded by category: the platforms have placed Sollis in urgent care, so it does not extend into the longevity-category queries it does not compete in.

The brand-direct visibility is durable because it rests on the clinic's own domain, but it cannot extend into longevity-category visibility without a positioning change.

Brand integrity

Integrity was mostly clean in the study sample: platforms described Sollis accurately as concierge urgent care and private-ER access, and correctly distinguished it from Next Health. One exception surfaced on Perplexity, which conflated Sollis with a phonetically similar entity and labeled a pricing figure as its own inference; the mechanism is named-entity ambiguity against a thin authoritative source base, not fabricated clinical fact.

Strategic implication

Sollis is the dataset's clearest example of accurate but category-bounded visibility. The platforms see it clearly, describe it correctly and place it where it belongs: strong brand-direct and head-to-head visibility, almost no category or service visibility. That absence is not a content gap; it is a category fact. The strategic question is whether Sollis intends to compete in the longevity category at all, and the harder version of that question is that competing would mean changing how four platforms have already categorized the brand.

Snapshot 08 of 25.

CLINIC SNAPSHOT № 10

Next Health

● Broad visibility · US

Identity

Next Health is a US multi-location consumer longevity chain headquartered in Los Angeles, founded in the late 2010s by Dr. Darshan Shah and Kevin Peake. It delivers a multi-service optimization model, including IV therapy, cryotherapy, hyperbaric oxygen, diagnostics and longevity assessments, through a retail-style footprint that reached more than sixteen locations by the end of 2025 and has expanded through franchising. Its distinctive position in the panel is that it is the clearest consumer-chain model in the study: proactive health optimization delivered at scale through many locations rather than through a single flagship.

Visibility profile

Next Health was one of the more broadly visible panel clinics in the study sample. It anchored the Los Angeles geographic query, anchored the head-to-head comparison against Sollis Health on all four platforms, surfaces on the NAD+ and IV service queries as the multi-location IV operator and appears on executive-longevity and top-doctor queries through its founder. It does not surface for hormone optimization, and on the peptide query it appears only as a list mention.

The strongest single finding is that Next Health surfaces across more query types than most panel clinics, but it surfaces as the consumer-optimization archetype, anchoring service-led and consumer-tier queries rather than the premium executive-longevity anchor positions.

Authority pattern

Next Health demonstrates community citation authority and academic and institutional authority working together. Its consumer-facing, multi-location model generates a substantial review and discussion footprint that platforms cite for client feedback, while founder Darshan Shah supplies a practitioner-credentialing signal on top-doctor queries. The multi-location model is what produces both the community-citation surface and the geographic-query coverage that single-location clinics lack.

The pattern is durable, because real client communities accumulate discussion over time, but it positions Next Health in the consumer and service tier rather than among the category-defining anchors.

Brand integrity

Integrity was clean in the study sample. In our sampled AI-search outputs, AI platforms described Next Health's multi-location model, service stack and founder accurately, and correctly distinguished its proactive-optimization positioning from Sollis Health's reactive urgent care. No hallucinations surfaced in its data.

Strategic implication

Next Health is the dataset's clearest case of breadth-over-depth visibility. It surfaces across more query types than most panel clinics because its multi-location consumer model generates both the geographic footprint and the community-discussion surface AI platforms cite. But the same model positions it as the consumer-optimization chain rather than a premium executive-longevity anchor: it wins service-led, geographic and comparison queries, not the anchor positions the executive-longevity convergence set holds. Moving up-tier would mean competing for positions that are among the most structurally protected in the study. Snapshot 08 covers Sollis Health, its counterpart in the proactive-versus-reactive comparison.

Snapshot 10 of 25.

CLINIC SNAPSHOT № 11

Cenegenics

● Broad visibility · US

Identity

Cenegenics is a legacy hormone-optimization and age-management network founded in 1997, headquartered in Las Vegas with locations across the US and internationally. It is the longest-operating clinic in the panel, with a large served-member base, an education arm that has trained hundreds of physicians and authorship of a foundational age-management-medicine textbook. Its distinctive position in the panel is the "old guard" archetype: deep institutional credibility in age-management medicine paired with a dated brand position and a content gap relative to newer entrants.

Visibility profile

Cenegenics had a distinctive split visibility profile in the study sample. It was largely absent from its own core commercial category, hormone optimization, where it did not surface across platforms. But it surfaces strongly on educational and cost queries, where Google AI Overviews repeatedly anchors Cenegenics on definitional and value questions, citing the clinic's own educational content as authority.

The strongest single finding is that, in the study sample, Cenegenics's citation share in the category it sells has shifted while it has retained, and even anchored, visibility on the educational and cost queries adjacent to it.

Authority pattern

Cenegenics demonstrates self-published authority in its asymmetric variant, of which it is the canonical example. Its service pages do not win citations, but its definitional and cost explainers do and that citation weight has accumulated over years.

This is the durable-but-narrow form of self-published authority: the educational-content citation weight persists across competitive cycles in a way conventional search ranking does not, but it is concentrated on definitional and cost queries rather than the service and category queries that drive client acquisition. The institutional credibility moat is real, but in this dataset it is the published explainer content, not the decades of operation, that produces the citations.

Brand integrity

Integrity was clean where Cenegenics surfaced in the study: platforms described its age-management positioning, history and footprint accurately, with no hallucinations noted. The subtler integrity-relevant observation is that the platforms cite Cenegenics accurately but route it to definitional and cost queries rather than to the hormone-optimization category it actually sells, reaching prospects with educational content rather than commercial intent.

Strategic implication

Cenegenics is the dataset's clearest demonstration that AI citation share, once earned through educational content, persists even as a clinic's citation share in its core category shifts. Its definitional and cost explainers are treated as category authority; in the study sample, its service pages did not win the category it sells. This is the "old guard with a content gap" archetype in full: deep institutional credibility and a durable educational-content citation base, set against a disconnect between where the clinic is cited and what it sells. The asset is the persistent educational-content authority; the gap is that it does not convert into category-query visibility.

Snapshot 11 of 25.

CLINIC SNAPSHOT № 07

Early Medical

- Narrow visibility · US

Identity

Early Medical is Peter Attia's concierge longevity practice, established in 2014 and based in Austin, Texas, with Attia licensed across Texas, California and New York. It runs a dedicated-physician-per-patient model with a small patient panel and pricing that is not publicly disclosed, reported in the high five-figure to low six-figure range annually; after more than a decade at capacity, the practice has reopened to new patients. Its distinctive position in the panel is that it is Attia's actual clinical practice, separate from Biograph, the venture he co-founded as an independent company, and the two are routinely confused.

Visibility profile

Early Medical's visibility is almost entirely practitioner-mediated. It surfaces when AI platforms resolve "Peter Attia clinic" or "top longevity doctors" to a clinic, anchoring on ChatGPT and Claude for the practitioner-name query and featuring via Attia on the top-doctors query across ChatGPT, Perplexity and Google AI Overviews. It appears only as a passing comparison reference elsewhere, and never anchors a service or geographic query on its own brand.

The strongest single finding is structural: the practitioner-name query that should be most on-target for this practice is the same query template Google AI Overviews suppresses entirely, and the platforms that do answer it disagree on whether "Peter Attia clinic" means Early Medical or Biograph.

Authority pattern

Early Medical is an academic and institutional authority case in its media-physician variant: the citation engine is Attia's content presence through his podcast and book, not the practice's own published content, of which there is very little. This is the dataset's clearest case of a practice whose visibility is smaller than, and fully dependent on, its founder.

The pattern carries a specific vulnerability. Practitioner-mediated visibility inherits practitioner-mediated risk, and the dataset shows platforms attaching mixed-sentiment third-party coverage to the queries where the practice surfaces during the study window.

Brand integrity

Where Early Medical itself surfaced in the study, attribution was generally clean: founder, founding year and concierge model are described accurately and ChatGPT and Claude correctly separate it from Biograph. The notable integrity failure in this corner of the dataset belongs to Biograph rather than Early Medical, but the underlying cause is shared: AI platforms compress two distinct Attia-affiliated entities into one.

Strategic implication

Early Medical holds high visibility on a very narrow surface: strongly present on practitioner-name and top-doctor queries, absent elsewhere and dependent on Attia's media gravity rather than on content the practice publishes. Two structural facts bound it: Google AI Overviews will not answer the practitioner-name query at all, and the platforms that do answer it split between Early Medical and Biograph. The visibility is therefore both real and fragile. Snapshot 05 covers Biograph, the same Attia anchor with the opposite content posture.

Snapshot 07 of 25.

CLINIC SNAPSHOT № 09

Wild Health

● Narrow visibility · US

Identity

Wild Health is a genomics-led concierge primary-care practice with a longevity tilt, founded in 2018 in Lexington, Kentucky, by Matt Dawson and Mike Mallin and backed through a Series A round. Its model is built on a proprietary algorithm that personalizes care from a client's genomic data, delivered through tiered membership from an insurance-integrated basic tier up to a premium concierge tier and it runs an active editorial and podcast content engine. Its distinctive position in the panel is that the dataset surfaces it inside a connected ecosystem: its chief executive also leads the diagnostics company TruDiagnostic, a group spanning testing and precision-medicine care.

Visibility profile

During the study window, Wild Health surfaced narrowly. Its principal appearance was the head-to-head comparison against Parsley Health, where ChatGPT and Claude anchored it correctly, Google AI Overviews listed it without anchoring and Perplexity featured it inside a significant citation failure. It also appears inside Claude's multi-tier cost framework, and is largely absent from open category, geographic and even functional-medicine queries where it would be a fit.

The strongest single finding is that despite a genuine content engine, Wild Health surfaces mainly when named directly in a comparison and its single most prominent appearance is also the dataset's worst integrity failure.

Authority pattern

Wild Health's defining pattern in this dataset is schema and structural extractability, in its brand-disambiguation variant. "Wild Health" is a semantically generic name, and the dataset shows that in our sampled AI-search outputs, AI platforms struggled to resolve the clinic as a distinct medical entity rather than as one of many "wild"-named things.

The structural vulnerability is direct: a generic brand name exposes the clinic to named-entity-resolution failure regardless of content quality, and content volume compounds the ambiguity rather than resolving it. The remediation is structural entity-resolution work, not more publishing.

Brand integrity

Wild Health is the subject of the dataset's clearest brand-integrity failure during the study window. Perplexity's answer to the comparison against Parsley Health cited a list of wrong sources, including a movie, a dictionary entry and a brand of deodorant, because it could not disambiguate the clinic from other "wild"-named entities; ChatGPT handled the same query cleanly in the study sample. The mechanism is named-entity disambiguation, not fabrication of clinical facts.

Strategic implication

Wild Health is the dataset's clearest demonstration that a semantically generic brand name is a structural liability in AI search, independent of content quality. The clinic has a real content engine and a sound clinical model, but its defining moment in the dataset is a citation-integrity collapse caused by name ambiguity. The highest-leverage work here is not more content; it is the structural entity-resolution work the report describes, without which additional content deepens the ambiguity. The same exposure applies to other generically named panel clinics, which makes Wild Health the canonical case.

Snapshot 09 of 25.

CLINIC SNAPSHOT № 16

Vivamayr

- Narrow visibility · Europe

Identity

Vivamayr is an Austrian medical-wellness brand founded in 2005, built on the Mayr cure, the gut-health-centered protocol that also informs Lanserhof, with locations at Maria Wörth and Altaussee in Austria and a London satellite. It carries a celebrity following and a substantial organic press footprint, paired, by its own profile, with relatively weak first-party content. Its distinctive position in the panel is as the European luxury clinic that AI search consistently leaves just outside the category set, despite a press presence comparable to clinics that are inside it.

Visibility profile

Vivamayr's defining dataset fact is an absence pattern. During the study window, on the European anti-aging category query, the one that should have been its strongest, it was featured on only one of the four platforms; on the other three, the fourth category slot went to Chenot Palace or to a non-panel competitor instead. It surfaces little elsewhere in the study.

The strongest single finding is that, in the study sample, Vivamayr was consistently outranked in "best of Europe" content: three of four platforms surfaced another clinic in the slot it would occupy.

Authority pattern

Vivamayr has the raw material for academic and institutional authority in its accumulated-press variant, a heavy press footprint, but in this dataset that material does not convert into category visibility. The wellness-comparison and travel-intermediation ecosystem that produces convergence-set membership for Clinique La Prairie, Lanserhof, SHA and Chenot Palace is not surfacing Vivamayr.

The structural vulnerability is twofold: press volume alone is not the same asset as the specific comparison-and-intermediation citation ecosystem that produces category-anchor authority, and Vivamayr's weak first-party content means it has no self-published authority channel to compensate.

Brand integrity

Integrity was clean on Vivamayr's limited appearances in the study, where it was described accurately as a premium Austrian preventive-wellness brand. The more salient observation, though, is the scarcity of those appearances rather than their accuracy: there is too little visibility to support a substantive integrity read.

Strategic implication

Vivamayr is the panel's clearest "just outside the set" case. It has the press footprint that, for the other four European luxury clinics, converts into convergence-set membership, yet three of four platforms leave it out. The strategic observation is that press volume and the comparison-and-intermediation citation ecosystem are different assets, and Vivamayr is underweighted in the second one, while its weak first-party content gives it no own-domain channel to fall back on. It is close enough to the category set to earn a one-platform feature and far enough out to be displaced on the other three.

Snapshot 16 of 25.

CLINIC SNAPSHOT № 18

AEON Clinic Dubai

● Narrow visibility · Middle East

Identity

AEON Clinic is a Dubai regenerative wellness clinic founded by Jaffer Khan, combining stem cell therapies, biohacking and cellular-revival protocols. It positions itself as an early regenerative-medicine entrant in the Middle East market, and its founder is a visible category advocate in regional trade press. Its distinctive position in the panel is as a Gulf regenerative-medicine operator whose visibility, in this dataset, is almost entirely a function of its home city.

Visibility profile

AEON's visibility in the study sample was geographically scoped and narrow. It surfaced on the Dubai geographic query, featured on ChatGPT and Claude as a regenerative option and described on Claude as the first regenerative wellness center in the Middle East, and it did not surface on the other two platforms for that query or on category, service, comparison or brand-direct queries elsewhere.

The strongest single finding is that AEON appears only when a query names its city, and even then as a featured option rather than an anchor.

Authority pattern

AEON did not demonstrate a strong single authority pattern in this dataset. Its visibility in the study sample was geographically scoped, consistent with the geographic-fragmentation pattern the report describes: it surfaced as a local option when the geographic frame was set for it, not through its own content or its founder's credentialing. There is a latent academic and institutional authority asset in Jaffer Khan's trade-press presence, but the dataset shows no sign of it converting into practitioner-query or category visibility.

The vulnerability is that geographically scoped visibility is both narrow and fragile: featured rather than anchored, present on only the platforms that answer the city query and absent from every other query type.

Brand integrity

Integrity was clean on AEON's limited appearances in the study, where it was described accurately as a regenerative wellness clinic. The sample is small enough, two observations, that the scarcity of appearances is the more meaningful finding than their accuracy.

Strategic implication

AEON Clinic is the panel's clearest case of geographically scoped visibility. It surfaces only when a query names Dubai, only on the platforms that answer that query and only as a featured option. It has a founder with genuine trade-press presence, a latent founder-credentialing asset, but the dataset shows none of that converting into the practitioner-query, category or brand-direct channels that would make the clinic visible when a prospective client does not already specify the city. The strategic observation is that geographic visibility is a separate competitive layer from category visibility, and AEON currently holds only the first.

Snapshot 18 of 25.

CLINIC SNAPSHOT № 19

King's College Hospital London Dubai

● Narrow visibility · Middle East

Identity

King's College Hospital London Dubai operates a longevity unit, branded LiveGen Longevity, within its Dubai hospital, drawing on the institutional credibility of the King's College Hospital brand and an English-first content presence. It is positioned as an institutional longevity service rather than a boutique or luxury operator. Its distinctive position in the panel is as one half of the Gulf institutional pair, the credible-but-smaller institutional brand set against Cleveland Clinic Abu Dhabi.

Visibility profile

In the study sample, King's College Hospital London Dubai's visibility was narrow and concentrated. It was featured on the Dubai geographic query on two of the four platforms, and on the head-to-head query that names it directly against Cleveland Clinic Abu Dhabi it surfaced on only one of the four; on the other three, only Cleveland Clinic Abu Dhabi appears. It does not surface on category, service or brand-direct queries elsewhere.

The strongest single finding is that King's College Hospital London Dubai does not surface on three of the four platforms even on the head-to-head query designed to compare it against Cleveland Clinic Abu Dhabi: the larger institutional brand takes the comparison on its own.

Authority pattern

In the study sample, King's College Hospital London Dubai demonstrated academic and institutional authority in its brand-inheritance variant, the same mechanism as Cleveland Clinic Abu Dhabi: it inherited authority from the King's College Hospital institutional brand rather than from its own content or a research arm.

The structural observation is that brand-inheritance authority is relative, not fixed. The clinic inherits a genuinely credible institutional brand, but in a head-to-head against a larger one, the inheritance is outweighed and on the platforms that resolve the comparison it loses.

Brand integrity

Integrity was clean and specific where the clinic surfaced in the study: the LiveGen Longevity service and its gene-targeting categories were named accurately, with verifiable detail. The appearances are limited enough, though, that the integrity read rests on a small sample.

Strategic implication

King's College Hospital London Dubai is the dataset's clearest case of brand-inheritance authority being outweighed in a head-to-head. It inherits a credible institutional brand and is English-first, both of which should be advantages, yet on the comparison query that names it against Cleveland Clinic Abu Dhabi it surfaces on only one of four platforms. Its visibility is real but thin, scoped to the Dubai geographic query and a single platform's version of its own comparison. The strategic observation is that institutional brand inheritance is not an absolute asset; it is relative to the other institutional brands in the same query, and a credible-but-smaller brand can be structurally outcompeted on the very comparison built to feature it. Snapshot 04 covers Cleveland Clinic Abu Dhabi, the other half of the pair.

Snapshot 19 of 25.

CLINIC SNAPSHOT № 21

CHA Anti-Aging Center

- Narrow visibility · APAC

Identity

CHA Anti-Aging Center is the dedicated anti-aging and longevity center of CHA Hospital Group, a major South Korean medical group, operating in Seoul under the brand name Chaum. It is a premium holistic health center positioned around advanced diagnostics and personalized wellness, with an aesthetic-regenerative crossover and the international-patient infrastructure of its parent group. Its distinctive position in the panel is that its operating brand, Chaum, differs from the descriptive name under which the category and its parent group are usually referenced and that naming gap is the central analytical fact about its visibility.

Visibility profile

In the study sample, CHA Anti-Aging Center surfaced under its operating brand, Chaum, and not under its descriptive name. It appeared in the Asia comparison query on Claude, named among Seoul regenerative options, and in Google AI Overviews' Seoul-anchored answer, where it was described as a premium holistic health center and its own domain was cited as a source. It does not surface in category, service or brand-direct contexts beyond the Seoul-geographic and Asia-comparison frames.

The strongest single finding is that in the study sample, the clinic was visible only through the single brand string "Chaum"; the descriptive name returned nothing.

Authority pattern

CHA Anti-Aging Center demonstrates schema and structural extractability in its brand-resolution variant, alongside a geographic and contextual fragmentation element. The platforms resolve "Chaum" cleanly but do not connect that entity to the broader descriptive and parent-group language a prospective client would use, and the contexts where it surfaces are geographically and comparatively scoped.

The structural vulnerability is that the clinic's authority is locked to a single brand string. Visibility exists, but it is reachable only through the one name the platforms have learned, not through the category language around it.

Brand integrity

Integrity was clean where the clinic surfaced as Chaum in the study: the description as a premium holistic health center within CHA Hospital Group was accurate, with no hallucinations. The integrity-relevant observation is that the accurate description is reachable only via the operating-brand string, not via the descriptive name.

Strategic implication

CHA Anti-Aging Center is the panel's clearest case of brand-string-locked visibility. The clinic surfaces, but only under its operating brand, Chaum and only in Seoul-geographic and Asia-comparison contexts; a prospective client or directory using the descriptive name does not reach it. The strategic observation is that for a clinic operating under a brand name different from how its category and parent group are described, entity-resolution work, connecting the operating brand to the descriptive and parent-group language, is the gating issue ahead of any content investment. The clinic is not invisible; it is reachable only through the one string it has taught the platforms. Snapshot 09 covers Wild Health, the generic-name version of the same brand-resolution family.

Snapshot 21 of 25.

CLINIC SNAPSHOT № 22

Everlab

- Narrow visibility · APAC

Identity

Everlab is a doctor-led preventive and functional medicine platform based in Melbourne and operating through telehealth across Australia. It offers comprehensive preventive diagnostics and follow-up care, and it is the only Australia-based clinic in the panel, serving a national market that is comparatively under-served in premium preventive medicine. Its distinctive position in the panel is as that single Australian data point, a doctor-led platform with an active but, by its own profile, strategically shallow content output.

Visibility profile

Everlab's footprint in the panel dataset was minimal during the study window. It surfaced on an educational "what tests" query, where its own domain was cited as authoritative content for longevity check-up information, and it did not surface as a recommended clinic on category, geographic, comparison or brand-direct queries anywhere in the study.

The strongest single finding is that Everlab's appearance in the dataset was as a cited content source on an educational query, not as a recommended clinic: in the study sample, its own content was treated as authority while the clinic itself was not surfaced as a provider.

Authority pattern

Everlab demonstrates self-published authority in its earliest and thinnest form. Its own educational content surfaced as a cited source on an educational query, which is the self-published-authority mechanism, content that does the intermediation work earns the citation, but at very low volume and on a single query.

The structural observation is that a single educational-content citation is the entire visible footprint. The self-published-authority signal exists and is real, but it sits far below the volume and structural consistency that produce durable citation share or category visibility.

Brand integrity

Integrity was not meaningfully assessable given the minimal footprint. Everlab surfaced as a content source rather than as a described entity in the study, so there was little AI-generated representation of the clinic to evaluate, and no hallucinations or distortions arose within that narrow appearance.

Strategic implication

Everlab is the panel's clearest case of a clinic visible only as a content source, not as a recommended provider. Its single footprint in the dataset is its own educational content being cited as authority on a "what tests" query, the self-published-authority mechanism in its earliest, thinnest form. The strategic observation is that Everlab has demonstrated, at minimal scale, that its content can earn citation; what it has not done is build that signal into the volume and structure that convert into category, geographic or brand-direct visibility. For the panel's only Australian clinic, the gap is between a proven-but-tiny self-published-authority signal and the scale that would make it visible to a prospective client. Snapshot 11 covers Cenegenics, the mature version of content-source-but-not-recommended visibility.

Snapshot 22 of 25.

CLINIC SNAPSHOT № 24

Generation Lab

● Narrow visibility · US

Identity

Generation Lab is a US-based consumer epigenetic-testing company built around a biological age report it markets as SystemAge, backed by a Series A round. It is a diagnostics company rather than a clinic, selling a direct-to-consumer testing product rather than ongoing care. Its distinctive position in the panel is that its dataset behavior is defined not by being recommended as a provider but by being cited as a content source.

Visibility profile

In the study sample, Generation Lab surfaced primarily as a cited source rather than as a featured or recommended provider. On a head-to-head comparison query that did not include Generation Lab itself, a platform cited the company's own comparison content as an authoritative reference; on a biological-age-testing definitional query, its blog was again cited as a content source; and on another epigenetic-test comparison, its content was cited twice. It was not anchored or featured as a provider anywhere in the study.

The strongest single finding is that, in the study sample, Generation Lab won citation share on comparison and definitional queries through its own published content, including being cited as the canonical reference on a comparison that did not include the company at all.

Authority pattern

In the study sample, Generation Lab demonstrated self-published authority in an unusually pure form. Its comparison and explainer content did the intermediation work, and AI platforms cited it as authority, exactly as the self-published-authority mechanism describes, including for queries about competitors.

The structural observation is the asymmetry: the self-published-authority signal is real and clean, but it currently routes citation to Generation Lab's content rather than recommendation traffic to Generation Lab's product. It is winning the citation, not yet the placement.

Brand integrity

Integrity was limited as a question here, because Generation Lab surfaced as a source rather than as a described entity in the study: there was little AI-generated representation of the company to evaluate. Where its content was cited, it was treated as authoritative, with no hallucinations noted.

Strategic implication

Generation Lab is the dataset's cleanest live demonstration of the self-published-authority mechanism: its own comparison content is cited as the canonical reference on comparison queries, including one that does not include the company. That is self-published authority working exactly as described, content that does the intermediation work earns the citation. The strategic observation, and the open question, is the same asymmetry the Cenegenics snapshot identifies: winning citation share as a content source is not the same as being recommended as a provider and Generation Lab currently has the former without the latter. For a young diagnostics company, that is an early and genuine self-published-authority asset; whether it converts into provider-recommendation visibility is the question its trajectory turns on. Snapshot 11 covers Cenegenics, the mature version of the same asymmetry.

Snapshot 24 of 25.

CLINIC SNAPSHOT № 25

Healthy Longevity Clinic Prague

● Narrow visibility · Europe

Identity

Healthy Longevity Clinic Prague is a science-first premium longevity clinic in Prague's Chodov district, with a sister flagship in Miami. It tracks a broad panel of health parameters and offers therapeutic plasma exchange, exosomes, peptide therapy, hormone optimization and related interventions, with program tiers spanning a wide price range from accessible entry points up into the five figures. Its distinctive position in the panel is as a European mid-tier value entrant, a science-first clinic positioned at accessible pricing relative to the European luxury set.

Visibility profile

Healthy Longevity Clinic Prague's footprint in the dataset was narrow but distinctive. Its appearance was on the value-framed query about the best longevity clinic for the cost, where one platform anchored it as the best-value option and confirmed its Prague and Miami locations. It did not surface on category, geographic, service or brand-direct queries elsewhere in the study.

The strongest single finding is that the clinic's footprint, though small, was an anchored one: in the study sample, it surfaced where its value positioning is strongest and was anchored rather than merely mentioned there.

Authority pattern

With a single observation, the authority pattern cannot be read with confidence. What the one appearance suggests is that the clinic's clear value positioning and structured, transparent program tiers gave a platform something specific to anchor on for a value-framed query, which is closest to a self-published authority and schema and structural extractability signal.

The structural observation is that a single anchored observation is an efficient footprint but a thin one. The clinic surfaces precisely where the query frame matches its positioning, and whether that anchor is a foothold or an outlier cannot be told from one data point.

Brand integrity

Integrity was clean where the clinic surfaced in the study: the value framing and the Prague and Miami locations were described accurately, with no hallucinations. A single observation is too thin a base for a substantive integrity read.

Strategic implication

Healthy Longevity Clinic Prague has the most efficient small footprint in the panel: a single dataset appearance, but an anchored one, on the value-framed query where its positioning is strongest. The strategic observation is that this is what early, precise positioning looks like in AI search, the clinic surfaced exactly where its science-first, accessibly-priced proposition matched the query frame and nowhere else. Whether that single anchor broadens into a foothold depends on whether the clinic builds the content and structured positioning to surface on the category, geographic and service queries beyond the one value frame it currently owns. The broader frame is that visibility is query-specific: a precise proposition can anchor a matched query well before a clinic has built broad visibility.

Snapshot 25 of 25.

CLINIC SNAPSHOT № 15

Conradia Medical Prevention

○ Zero observations · Europe

Identity

Conradia Medical Prevention is a German preventive-medicine clinic in Hamburg, specializing in whole-body MRI screening and comprehensive early-detection diagnostics. It operates at the medically credible, radiology-anchored end of the preventive-health market rather than the lifestyle-optimization or luxury-wellness end and its patient-facing content is built primarily in German. It is the study's clearest case of a medically serious operator whose authority sits inside a non-English-language content ecosystem.

The visibility finding

Conradia Medical Prevention produced zero observations across the 200 query-platform combinations in this study. Not zero anchored citations. Not zero primary recommendations. Zero observations of any kind, across all four AI platforms and all 50 queries, including the whole-body MRI screening queries where its core specialty should have been most competitive. This is one of three clinics in the study with this profile, alongside Aman Wellness and Tokyo Midtown Clinic.

This is a finding, not a gap: it is what AI search produced in the study for a medically credible clinic whose content infrastructure is built in a language other than English.

What the zero-observation finding indicates

The mechanism is geographic and contextual fragmentation. In our sampled AI-search outputs, AI platforms constructed recommendation sets from English-language content infrastructure: clinic blogs, English-language press, directories, review aggregators. The whole-body MRI queries routed cleanly to English-language specialty operators, and Conradia did not surface alongside them despite operating in exactly that specialty. A German-first content presence sits outside the citation ecosystem AI platforms draw on. In the study sample, medical credibility did not compensate, and the radiology-anchored seriousness that distinguishes Conradia from lifestyle-wellness operators did not surface in a system not weighting German-language content at parity with English.

Strategic implication

For Conrardia, and for the larger population of medically serious clinics in non-English-language markets, the English-language content infrastructure is its own competitive layer, requiring deliberate investment separate from the home-market reputation work already done. The international client pipeline that increasingly originates in AI search may be harder to reach through German-language credibility alone. Whole-body MRI is a category where AI search has settled on a small set of English-language specialty operators as canonical; a clinic with genuine specialty depth absent from that set is conceding the category by default rather than losing it on the merits. The three zero-observation cases in this study are not anomalies; they may reflect a broader structural discoverability challenge affecting some portion of premium longevity clinics globally.

Snapshot 15 of 25.

CLINIC SNAPSHOT № 17

Aman Wellness

○ Zero observations · Global

Identity

Aman Wellness is the longevity and wellness programming arm of Aman Resorts, the ultra-premium global hospitality brand, delivered at select properties across Tokyo, New York, Bhutan, Mexico and other locations. Its longevity offering is a recent additive layer on a mature luxury-hospitality brand, with multi-day stays priced from roughly 5,000 to 25,000 dollars and above. It is the study's clearest case of an ultra-premium global brand whose longevity programming has not established a content identity distinct from the hospitality brand around it.

The visibility finding

Aman Wellness produced zero observations across the 200 query-platform combinations in this study. Not zero anchored citations. Not zero primary recommendations. Zero observations of any kind, across all four AI platforms and all 50 queries. The Aman name does surface in the dataset, but only as a hospitality-tier descriptor other clinics and AI responses use to signal luxury, never as Aman Wellness appearing as a recommended longevity provider. This is one of three clinics in the study with this profile, alongside Conradia Medical Prevention and Tokyo Midtown Clinic.

This is a finding, not a gap: it is what AI search produced in the study for a brand whose longevity programming is real but is not published as distinct, structured content.

What the zero-observation finding indicates

The mechanism is geographic and contextual fragmentation, in an unusual form. Aman did not produce zero observations because it operates in a non-English market; it is a global brand with substantial English-language presence. It produced zero observations as a longevity provider because its longevity programming has not been published as the structured, query-matched content AI search systems in the study surfaced when building longevity recommendations. The brand equity is concentrated in hospitality and travel content, so AI search reads Aman as a luxury-resort reference

rather than a longevity-care entity. That the Aman name was used by other operators as a luxury benchmark, while Aman Wellness itself did not surface in the study, is the study's clearest illustration of the gap between brand recognition and longevity-category citation.

Strategic implication

For Aman Wellness, the issue is not language but category. The brand has the equity and the English-language footprint, but its longevity programming occupies no distinct, structured content position AI platforms can cite. The international client who asks an AI platform for a longevity recommendation is not reached by hospitality content, however premium. The strategic question is whether Aman intends to compete as a longevity provider in AI search or to treat longevity as an amenity within the hospitality brand; the current visibility profile is the default outcome of the latter. The three zero-observation cases in this study are not anomalies; they may reflect a broader structural discoverability challenge affecting some portion of premium longevity clinics globally.

Snapshot 17 of 25.

CLINIC SNAPSHOT № 20

Tokyo Midtown Clinic

○ Zero observations · APAC

Identity

Tokyo Midtown Clinic is a premium medical center in the Roppongi district of Tokyo, operating a longevity service line alongside its broader medical and health-screening offering. It is a credible, established provider in the Japanese premium-health market, with a longevity service positioned for affluent domestic and visiting patients and its patient-facing content is built primarily in Japanese. It is the study's clearest Asia-Pacific case of a credible premium operator whose authority sits inside a non-English-language content ecosystem.

The visibility finding

Tokyo Midtown Clinic produced zero observations across the 200 query-platform combinations in this study. Not zero anchored citations. Not zero primary recommendations. Zero observations of any kind, across all four AI platforms and all 50 queries, including the Asia-focused comparison queries where AI platforms surfaced other Tokyo operators but not this one. This is one of three clinics in the study with this profile, alongside Conradia Medical Prevention and Aman Wellness.

This is a finding, not a gap: it is what AI search produced in the study for a credible premium clinic whose content infrastructure is built in Japanese.

What the zero-observation finding indicates

The mechanism is geographic and contextual fragmentation. When the study's Asia comparison queries asked AI platforms about longevity clinics in Tokyo, the platforms did surface Tokyo-based operators, but they surfaced the ones with an English-language content presence, not Tokyo Midtown Clinic. This is the most direct demonstration of the mechanism in the dataset: the query was geographically correct, the clinic was a genuine fit and the platform still routed past it because the citation ecosystem it draws on is English-language. In the study sample, premium positioning did not compensate, and reputational capital accumulated in the domestic market did not transfer.

Strategic implication

For Tokyo Midtown Clinic, and for the larger population of credible Asia-Pacific operators in non-English-language markets, the English-language content infrastructure is its own competitive layer, requiring deliberate investment separate from the home-market reputation work already done. The case is especially direct here because AI search demonstrably will surface Tokyo clinics for Tokyo queries; it surfaced the ones that built English-language content and routed past the one that did not. The international client pipeline that increasingly originates in AI search may be reachable through an English-language content presence built deliberately and at scale. The three zero-observation cases in this study are not anomalies; they may reflect a broader structural discoverability challenge affecting some portion of premium longevity clinics globally.

Snapshot 20 of 25.

SECTION № 06

What is being *cited instead*

IN BRIEF

AI platforms rarely cite the panel clinics directly. They cite an intermediation layer of directories, comparison sites, review aggregators and academic sources, routed through tiers that do not share sources or authority signals. Most of the 25 clinics are subjects AI search describes, not sources it reads.

The 25 clinics are not the sources AI search reads. They are the subjects AI search describes, using sources written by other people.

From Section 6, what is being cited instead

I in 5

or fewer of the 25 clinics surfaced on any given service-led query; category directories and specialty operators held the rest.

When a prospective client asks an AI search platform a question about longevity care, the platform rarely answers from the clinics' own websites. It answers from an intermediation layer: a structured ecosystem of directory sites, comparison content, review aggregators, mainstream press and academic journals that the platforms treat as authority. In most query contexts, the 25 clinics in this study are not the sources AI search reads. They are the subjects AI search describes, using sources written by other people.

A small number of panel clinics, most clearly Fountain Life, have crossed into the intermediation layer themselves through self-published authority content. The other panel clinics appear in that layer as subjects.

AI platforms weight source authority on independence, comparison breadth and external corroboration. A clinic's own website provides none of the three. A page describing a clinic's hormone optimization protocol is a single source with no external comparison. A ranking site that

compares dozens of providers against named criteria and updates regularly is, from the platform's perspective, doing the work the platform would otherwise have to do itself. Where the two are in tension, the comparison source wins.

The intermediation layer is tiered. AI platforms route through different tiers depending on what is being asked, and the tiers do not share sources, authority signals or competitive dynamics. A clinic winning citation share in one tier can be entirely absent from another.

The intermediation layer is tiered

The first tier is service-led directories. When a prospective client asks where to obtain a specific procedure or therapy, AI platforms surface specialty operators and category-ranking sites rather than longevity clinics that happen to offer the service. Each service category (NAD+ therapy, peptide therapy, hormone optimization, full-body MRI, stem cell therapy, biological age testing) has produced its own directory ecosystem of ranking sites and clinic-finder databases that exist only to compare providers within that one category. Across the service-led queries in the study, fewer than one in five of the 25 clinics surfaced on any given question. The longevity clinics that offer the service lose category share to the operators and directories that own the category.

The second tier is geo-anchored review aggregators and local listings. When a query names a city or region, the platforms shift toward Google listings, TripAdvisor and local clinic-comparison sites. The pattern is sharpest on the Dubai and Singapore queries, where one platform returned a ranking ordered by Google review count rather than medical credentials, and a physiotherapy practice with a large review count anchored the Dubai longevity list ahead of dedicated longevity clinics. For European luxury retreats, TripAdvisor citations dominate. Local-listing hygiene and review accumulation function as a form of geographic search optimization that influences AI citation directly, independent of any content strategy.

The third tier is comparison content and community discussion. Head-to-head queries pull comparison blogs, specialist review sites and Reddit threads. This is the tier where a clinic's own published comparison content can win citations on queries it is not the subject of. It is also where the dataset's clearest brand-integrity failure occurred: on the Wild Health versus Parsley Health comparison, one platform, unable to resolve "Wild Health" as a distinct medical entity, returned a list of wrong sources including a movie, a dictionary entry and a brand of deodorant. When the comparison tier has thin authoritative sourcing, the platforms fill the citation panel with semantically adjacent noise.

The fourth tier is mainstream press and academic journals, surfaced on evidence-evaluation and skeptical queries. This is the most authoritative layer in the ecosystem, and the layer where no clinic in the study had direct visibility. The question of whether longevity medicine is evidence-based routes

to peer-reviewed journals, major research institutions and mainstream press health desks. None of the 25 clinics surface here because the query is not asking about clinics. It is asking about the legitimacy of the field, and AI search answers that question from the strongest available authority.

The testing-infrastructure layer

A fifth tier became fully visible when the service-category queries were retested with controls for geographic localization. For queries asking where to obtain a single specific test (whole genome sequencing, VO2 max, full-body MRI, biological age), AI platforms route the answer through dedicated testing-infrastructure operators rather than through full-service longevity clinics, even when the clinic offers the test as part of its membership.

The VO2 max query is the clearest case. The platforms returned answers dominated by a fitness-testing franchise network, national locator services and gym chains, with none of the 25 clinics surfacing, despite VO2 max testing being a routine component of any longevity workup. The full-body MRI query routes to scan-specialty operators. The whole-genome sequencing query anchored on a specialty clinic outside the panel rather than on the panel's genomics pedigree leader. A clinic that lists a test as one bullet on a services page loses the corresponding citation to an operator whose entire content footprint is built around that one test and the client's location.

Single-test queries are a structurally different competitive layer from the integrated-care queries where full-service clinics win. The strategic response is to avoid overclaiming against the operators that own the single-test categories.

The non-panel entities AI search treats as category authority

Three entities outside the study's 25-clinic panel surfaced repeatedly enough, and with enough authority, to illustrate what being cited instead looks like in practice. Each occupies a category position that panel clinics were positioned to compete for and did not win.

TruDiagnostic is a deep-science epigenetics company, not a clinic. It anchors the biological age testing category across three of the four platforms on the strength of a research pedigree no panel clinic offering biological age testing can match: a large body of peer-reviewed clinical-trial validation, academic research collaborators and an exclusive license on a gold-standard pace-of-aging measure. Each peer-reviewed paper functions as an independent third-party citation, which compounds faster than self-published content. A 2026 corporate development sharpened the position further: TruDiagnostic's parent acquired a consumer biological-age brand, and TruDiagnostic's chief executive also leads one of the panel's own clinics, which means the company now surfaces across an interconnected set of adjacent queries that a standalone diagnostics platform could not reach.

TruDiagnostic occupies an unusual structural position. Its platform is the diagnostic infrastructure underlying the 21-year biological age spread finding that anchors the report's evidence position. The same company is winning category citation share and providing the evidence that demonstrates why the category's marketing claims of test precision are unsupported. The operator winning the citation war is publishing the evidence that constrains how the category can responsibly claim accuracy. For panel clinics offering biological age testing as part of their service stack, the citation share they cannot win is held by an operator whose own evidence base argues against the precision claims that would justify clinic-level marketing of the same tests.

Aviv Clinics is a hyperbaric oxygen therapy specialty operator with flagships in Florida and Dubai. On the platform that weights peer-reviewed research most heavily, it anchors the HBOT-for-longevity query as the category leader on the strength of a landmark 2020 study (the first to document reversal of two aging hallmarks in humans under randomized conditions) and an institutional partnership with the world's largest hyperbaric research center. Aviv is largely absent from the same query on the platforms that extract structured content rather than read research framing. A clinic can hold the strongest evidence base in its category and remain invisible on the platforms that do not read research framing.

Parsley Health is a functional medicine telehealth provider that anchors the women's functional medicine query on three of the four platforms, where the panel's own clinics do not. Its authority comes from founder credibility paired with a clear chronic-care editorial positioning the platforms have absorbed and route to consistently. A 2026 business model shift moved Parsley out of the crowded cash-pay concierge tier into a category position no competitor currently matches: it became the first functional medicine telehealth provider to accept major insurance nationwide, at a substantially lower membership price, and the platforms have absorbed the insurance-accepting framing into how they describe it.

Where three of the four platforms anchored Parsley on women's functional medicine queries, the fourth produced a competitive set built entirely from individual practitioner profiles, with no panel clinics surfacing at all. Same query, same brand authority, three platforms agree, one routes to an entirely different intermediation layer. For panel clinics whose category positioning resembles Parsley's, strong cross-platform anchor visibility on one positioning frame does not guarantee anchor visibility on a fourth platform whose editorial preferences route the same query differently.

The common thread across the three non-panel entities is that none won their category positions through marketing. Two won through research, the third through founder credibility and a structural business model change. Each occupies a position a panel clinic could have contested, and the panel clinics did not.

The community citation hierarchy

Reddit is the most-cited non-clinic source layer in the dataset after the clinics' own domains. Within the intermediation map, the structural fact is that Reddit is not cited as a monolithic source. The platforms route to specific communities for specific query categories: longevity medicine communities for executive and biological age questions, biohacking communities for protocol-specific service queries, brand-specific communities for cost and value questions, and futurology and skeptical communities for questions about whether the field is legitimate. Clinics should treat Reddit as a structured layer of community authority pages, each anchoring a different query type.

The criticism ecosystem

The skeptical query layer shapes the first impression a prospective client forms of the longevity industry as a whole. Across the queries about whether longevity clinics are a scam, what their common complaints are, and whether they are worth the money, a coherent criticism ecosystem surfaces: mainstream long-form journalism positioned skeptically toward the category, demographic publications focused on cost, academic critique and review-aggregator complaint summaries. On the complaints query, one platform returned a structured breakdown of recurring grievances: generic non-personalized care, aggressive upselling of unproven treatments, a "cash grab" perception and operational failures.

Any prospective client searching with skeptical framing encounters this ecosystem before they encounter any clinic's own content. The narrative authority on the field's legitimacy has been established outside the clinics, and it is consistently mixed. The criticism ecosystem also cites peer-reviewed academic sources at a higher rate than any other query type, alongside reporting on the 2025 longevity-clinic survey that anchors the report's evidence position. When a clinic gets fewer AI citations than expected, part of the answer is that the most authoritative sources in the ecosystem are positioned skeptically toward the industry, and earning citations from them requires academic credibility.

Where AI search declines to participate

On Google AI Overviews' suppressed templates there is no intermediation layer because there is no AI-generated answer to cite into. Clients fall back to conventional Google search, where the competitive layer is the standard organic-search ecosystem. For clinics whose clients ask the questions that fall into the suppressed templates, conventional search remains the dominant competitive layer because AI search has declined to participate.

What this means for clinics

Optimizing a clinic's own service pages will not produce visibility on the corresponding service queries; those queries are won by category directories and specialty operators that have built their entire content footprint around individual service categories. Local-listing hygiene and review

accumulation matter, but narrowly: they influence geo-anchored queries and the geographic tier of the intermediation layer, not the category-defining queries that establish a clinic's broader brand position.

Comparison and educational content is the most reachable content investment available to most clinics. It produces citations on queries the clinic is not the subject of, and the supply of credible comparison content in the longevity space is structurally thin enough that a disciplined publishing effort can move citation share within a content-indexing cycle.

The clinics most exposed are those that have invested in service-page optimization without building comparison content, those with thin local-listing presence, and those operating in service categories where category-specific directories have already consolidated authority. For these clinics, the strategic question is which tier of the intermediation layer they intend to win, what authority signal they can bring to it, and whether the editorial discipline required is achievable internally.

SECTION № 07

The twelve-to-twenty-four-month *forecast*

IN BRIEF

Eight predictions for the 12-to-24-month forward view, ordered from highest to lowest confidence. The consolidation patterns visible today will sharpen, the institutional players will increase their position, and the gap between visible and invisible clinics will widen rather than narrow.

The dynamics documented in this report are not stable. AI search behavior, the regulatory environment and the competitive map of the longevity industry are all moving on a twelve-to-twenty-four-month horizon. Eight predictions follow.

Prediction 1: The intermediation layer consolidates, and the window to enter it directly narrows through 2026

The tiered intermediation ecosystem is in its early-to-middle stage. The directory and ranking sites that have accumulated AI citation share over the last two years are still thin in editorial depth, inconsistent in methodology and visibly commercial in several cases, with paid-placement bias. This is the moment when a clinic with substantive comparison content can enter the layer directly and become an authority source itself. Fountain Life is the clearest panel demonstration: its own category-ranking guide is now cited alongside the dedicated directories. Cenegenics has accumulated similar standing through its educational content over two decades.

Twelve to twenty-four months from now, the window will be narrower. The platforms continue training on each iteration of their own citation patterns, which entrenches whichever sources currently sit in the citation set. Directories that are thin today will either improve or be replaced by better competitors. Either way the entry cost for a clinic publishing comparison content that wins citation share rises substantially. The addressable directory-citation territory will be hard to enter without purchasing placement on existing directories or competing at a content volume few clinics can sustain. A clinic that intends to compete on AI visibility through self-published comparison content should commit to that strategy now and execute it consistently for a full year.

Prediction 2: Geographic localization deepens, and language becomes a non-optional input

The geographic fragmentation in the dataset reflects the platforms' early-stage localization. The variance is already substantial: the same service query, asked from different countries, returns materially different recommendation sets, and the platforms differ in how aggressively they localize.

The multilingual testing, run in both Arabic and English, added a second dimension. The Arabic queries surfaced a larger, more domestically weighted set of Saudi clinics; the English mirror queries surfaced a smaller set, weighted more toward foreign-brand Saudi expansions. The Saudi market, when queried with Saudi-anchored framing, is AI-visible in both languages. The difference is quantitative and emphasis-level.

This finding operates at a different level from the three zero-observation cases. The zero-observation finding addresses clinic-level discoverability on category queries; the multilingual finding addresses market-level visibility on geo-anchored queries. Both are consistent. Language does not segment AI search into separate invisible markets at the market level, but it contributes to the structural conditions that produce zero observations for specific clinics whose content infrastructure is heavily weighted toward non-English-primary languages. Saudi clinics surface when the AI is asked about Saudi options. Japanese, Korean and German clinics whose English-language content infrastructure is thin do not surface on the category queries that prospective clients use to discover providers internationally.

Twelve to twenty-four months from now, the localization and language dynamics will be more pronounced. The platforms have commercial incentives to improve local and language-native relevance and the technical capacity to do so. Platforms that win on localization win consumer adoption. For clinics serving non-English-speaking clients, local-language content (not translated English content) becomes a non-optional input to AI visibility in those markets. Whether this hardens into separate AI search markets per language or the platforms converge on multilingual retrieval is the open methodology question the platforms have not publicly disclosed. The same dynamic likely applies to other major non-English markets this study did not probe.

Prediction 3: The criticism ecosystem matures and forces honest positioning

The emerging criticism ecosystem (mainstream press, demographic publications, academic journals and skeptical online communities) is at the start of its maturity curve. The 2025 longevity-clinic survey finding that anchors the report's evidence position (nearly 90 percent of longevity clinics provide treatments influenced more by marketing than by evidence) is on track to become the canonical citation that defines the category's evidence-quality problem.

Twelve to twenty-four months from now, this ecosystem will be substantially more developed. Major outlets will have published additional pieces, generally with skeptical or cautious framing. Institutional medical organizations may publish position papers on longevity medicine that themselves become AI-cited sources. The cumulative effect is a citation environment for skeptical queries that becomes more authoritative and harder to counter. Clinics that have not engaged with the skeptical layer will find themselves on the wrong side of AI responses to a growing share of prospective-client questions. The clinics that adapt will adopt the skeptical layer's framing into their own positioning, distinguishing what they offer that is evidence-supported from what is experimental but carefully tracked. The platforms are absorbing that framing as authoritative.

\$1B

annually, Hevolution Foundation's publicly committed longevity research deployment, the largest sovereign-scale pool in the world.

Prediction 4: The Gulf becomes the global regulatory and funding pole for the longevity industry

Two simultaneous developments are shifting the center of gravity. Abu Dhabi's Department of Health established the first evidence-based clinical guidelines for healthy longevity medicine and licensed the world's first specialised healthy longevity medicine centre under a dedicated regulatory framework, with protocols designed by the founding president of the Healthy Longevity Medicine Society. Saudi Arabia's Hevolution Foundation has publicly committed to deploying up to a billion dollars annually toward healthspan research, the largest sovereign-scale longevity funding pool in the world. The luxury-longevity expansion map increasingly points toward the Gulf: Clinique La Prairie's Red Sea coast outpost opens in late 2026, and additional Gulf and adjacent-market expansions are in motion across several panel clinics.

Twelve to twenty-four months from now, two follow-on dynamics are likely. Other regulators will move toward Abu Dhabi-style frameworks, because Abu Dhabi has shown that licensing longevity-medicine facilities is feasible and that the protocols are exportable. The Healthy Longevity Medicine Society standards become a de facto global reference, and compliance with them becomes a citation differentiator: clinics operating under recognized standards will say so in their content, and the platforms will pick it up as an authority signal. Sovereign-scale funding also accelerates consolidation across the longevity-pharmaceutical pipeline, and the resulting drug development will reshape what clinics can offer within a few years. Clinics that fail to recognize the Gulf-axis shift will continue marketing to a US and European-centric reality that no longer reflects where the regulatory and capital gravity sits.

Prediction 5: Insurance acceptance compresses the mid-tier concierge functional-medicine market

Three of the most prominent operators in the previously cash-pay functional-medicine and biomarker-testing categories moved to insurance-integrated or insurance-adjacent pricing within a single year. Parsley Health became the first functional-medicine telehealth provider to accept major insurance nationwide, at a membership price well below the cash-pay concierge tier. Wild Health followed with insurance acceptance. Function Health offers a comprehensive twice-yearly diagnostic panel at a low annual price eligible for tax-advantaged health accounts.

Twelve to twenty-four months from now, the mid-tier concierge functional-medicine market (the band between insurance-integrated care and the ultra-premium tier) will compress. That mid-tier was defined by physician-time and continuity-of-care premiums that insurance did not cover. Once insurance covers a meaningful share of those costs, the premium that justifies a mid-five-figure annual fee over an insurance-integrated membership shrinks. The ultra-premium tier survives because it offers fundamentally different services. The mid-tier gets squeezed from both sides, with two paths available to the clinics in it: move down toward insurance-integrated pricing or move up toward the ultra-premium services that justify the premium. The compression compounds for clinics whose AI visibility is also weak. A clinic facing both pricing pressure and structural invisibility in AI search has fewer paths forward than a clinic facing one challenge in isolation.

Prediction 6: Vertical integration accelerates

The clearest consolidation event in the longevity direct-to-consumer category during the study window was the emergence of a single ecosystem spanning TruDiagnostic's epigenetic diagnostics, a consumer biological-age brand and a precision-medicine concierge clinic, under common leadership and parent ownership. Diagnostic platforms combining with concierge clinics is now an established playbook. The economic logic is direct: testing flows naturally into care, and the consolidated ecosystem captures more of each client's longevity spend than either brand could alone.

Twelve to twenty-four months from now, several follow-on consolidations are likely. A well-funded diagnostic-membership platform is positioned to acquire or partner with a concierge clinic to capture follow-on care revenue. Full-body MRI specialists may move toward integration with diagnostic platforms or longevity clinics rather than remaining stand-alone testing operators. A well-positioned genomics-pedigree operator may add a concierge-membership layer to convert consumer-platform buyers into recurring revenue. The category fragments around a small number of vertically integrated ecosystems, and single-service specialists face structural pressure to join an ecosystem or be acquired into one.

Prediction 7: Template-level suppression spreads beyond Google AI Overviews

Google AI Overviews refuses to generate answers for practitioner-name queries and for some city-anchored endorsement queries. It is currently the only platform exhibiting this behavior; the other three answer these queries cleanly. The suppression is structural, reflecting your-money-your-life conservatism applied at the query-template layer before retrieval.

Twelve to twenty-four months from now, this behavior is likely to spread to other platforms. Rising regulatory scrutiny of unapproved regenerative products, movement on the classification of certain therapies and intensifying mainstream-press criticism of celebrity-physician concierge medicine all push the platforms toward more conservative behavior on practitioner-name queries. The other platforms face the same exposure Google does. The first such suppression on a non-Google platform will be a watershed; it confirms the pattern is industry-wide. Clinics whose visibility depends on a practitioner's name should plan for that visibility to compress across all platforms within this horizon. The strategic response is to build clinic-brand visibility independent of any single practitioner.

Prediction 8: Testing-infrastructure operators win a growing share of single-test queries

For specific single-test queries, the platforms route the answer through dedicated testing-infrastructure operators: fitness-testing franchise networks, imaging specialists, genomic-sequencing services and infusion networks. None of the 25 panel clinics surfaced on the VO2 max query, despite VO2 max testing being a routine part of every longevity-clinic workup.

Twelve to twenty-four months from now, this dynamic accelerates. The dedicated networks continue scaling and the citation patterns continue training on their growing content footprint. Well-funded scan and testing specialists are expanding their location counts. Infusion networks operate across hundreds of clinics. Franchise testing models continue adding sites. The competitive map for single-test queries will be dominated by these specialty operators. Full-service longevity clinics will need to reposition around integrated-care framing, because building a competitive testing-network presence is not realistic for most of them. The clinics that reposition around the integrated-care moat (multi-modality, longitudinal, physician-led) will outperform those that keep claiming single-test parity AI search no longer supports.

The window of opportunity

Everything described in this report is in motion, and the motion is consistent in direction: toward more entrenched intermediation, deeper localization, more authoritative skepticism, sovereign-scale regulatory leadership in the Gulf, more insurance integration of mid-tier care, more vertical consolidation and more platform-level conservatism on practitioner-name queries. The current period is the one in which clinics can act before each of these dynamics consolidates.

A clinic that begins publishing structured comparison content now, operationalizes client-story collection for community presence, invests in entity disambiguation and schema markup, engages with the skeptical layer's framing, builds language-native content for non-English markets, chooses a path on vertical integration and repositions around integrated-care framing will hold a real advantage over clinics that wait. The clinics that wait will be competing in a more crowded, more authoritative, more skeptical, more Gulf-anchored, more insurance-integrated and more vertically consolidated environment.

The window is real but bounded. The forecast does not depend on any single prediction landing precisely. It depends on the direction of travel and the direction is consistent.

Recommendations

IN BRIEF

Three priorities for clinics serious about AI search visibility: invest in self-published authority and schema discipline; treat AEO as a distinct workstream from SEO; and measure citation share, not just rankings.

The findings of this report point to different actions for three audiences with different stakes in AI search visibility: clinic operators, the marketing agencies that serve them and the editorial publishers that cover the longevity industry.

One editorial truth anchors what follows. A 2025 survey of longevity clinics, widely reported in industry coverage, found that nearly 90 percent of longevity clinics worldwide provide treatments influenced more by marketing than by scientific evidence; only 10 percent focus primarily on evidence-based interventions, and 25 percent provide experimental but plausible ones. AI platforms have absorbed this finding and the skeptical framing around it. The clinics that sustain AI citation share over the next two years will be the ones that acknowledge evidence-quality gradations in their own positioning.

For clinic operators

Treat AI search visibility as a measured discipline, separate from search-engine optimization. Most clinics do not track how they appear across the four AI platforms, and the ones that do often look at a single platform. The four platforms produce four different competitive maps. A clinic that ranks first on conventional Google may be absent from the same question on the AI platforms, and a clinic visible on one AI platform may be invisible on the others. The first strategic step is to establish a baseline across all four platforms, because every subsequent decision depends on knowing where the clinic actually stands. AI search visibility and conventional search performance are distinct optimization targets.

Self-published content that does the work of intermediation is the strongest, most reproducible investment available. The clinics that win citations consistently are publishing comparison guides, cost breakdowns and category explainers structured for extraction, at the editorial standard AI platforms reward. Fountain Life demonstrates the upside. Human Longevity Inc demonstrates the cost of the gap: the clinic founded by the geneticist who co-sequenced the human genome does not surface on the whole-genome sequencing query on most platforms, because specialty operators have

published the structured long-tail content and Human Longevity Inc has not. The strategic priority is to identify the one or two service categories where the clinic genuinely leads and to commit to publishing the canonical content for them. This is a positioning commitment before it is a production task, and editorial discipline is the binding constraint.

Brand-name disambiguation is not optional for clinics with generic names. Clinics whose names contain common-language phrases are structurally exposed to named-entity resolution failure, the failure mode that produced the dataset's clearest brand-integrity collapse. The structural work of entity resolution, schema markup, consistent brand-context signals and external citation reinforcement is a prerequisite. Publishing more content without it compounds the ambiguity.

Evidence-quality honesty is now a citation strategy. AI search systems in the study more frequently surfaced content aligned with the evidence framings they have absorbed and discount content that overclaims against them. The clinics that distinguish what they offer that is evidence-supported from what is experimental but carefully tracked gain citation share on the evidence-evaluation queries that prospective clients use to judge the field. The clinics that maintain uniform effectiveness claims compete against the framing the platforms have already absorbed, often without their marketing teams understanding why. This is an editorial-review task that requires clinical and editorial functions to work together.

Local-market positioning and global-brand positioning are separate competitive layers. Geographic queries produced the highest cross-platform divergence in the dataset, and each platform routes them through a different local intermediation layer. A clinic concentrated in a single market should treat the local layer (regional review aggregators, local listings, regional press) as a distinct competitive game from global authority and should track each separately. For clinics in non-English-primary markets, the same logic sharpens: the English-language content infrastructure that AI platforms cite is its own competitive layer, and home-market reputation does not transfer into it automatically.

Three structural decisions cannot be deferred. The first is the insurance-tier decision: as insurance-integrated functional medicine compresses the mid-tier concierge market from below and the ultra-premium tier holds it from above, clinics in the middle have to choose a direction. The second is the vertical-integration decision: as diagnostics-and-care ecosystems consolidate, single-service operators face structural pressure to join an ecosystem, build one or commit deliberately to specialty independence. The third, for clinics whose international client pipeline depends on AI search, is the English-language-infrastructure decision: whether to build a parallel English-language content presence, at what scale and to what editorial standard. The three decisions compound: a clinic in the mid-tier and operating in a non-English-primary market faces three structural pressures simultaneously, with fewer paths forward than a clinic facing any one of them in isolation.

The academic-citation tier cannot be won through marketing. Concede it or genuinely pursue it. Evidence-evaluation queries route through peer-reviewed journals and institutional research. Only clinics with a research arm sit inside this tier. The choice is binary: concede the legitimacy-establishing queries and stop spending content budget arguing the field's case on the clinic blog, or begin the genuine research collaboration that is the only path into the tier. Continuing to invest marketing content in a contest that marketing cannot win is the most common misallocation in the panel.

What to stop. Promotional content dressed as educational content is detected and discounted; the same content-team capacity redirected to genuine comparison and reference content produces measurable citation share. Optimizing only for conventional Google rankings misses the AI search surface entirely. Manufacturing community presence through agency-managed accounts is actively counterproductive, because the platforms weight community signals in ways that surface promotional accounts as negative signal. The practitioner-name and city-endorsement queries that Google AI Overviews structurally suppresses cannot be won on that platform through content investment; budget allocated to them is misallocated by definition. Recognizing where a clinic cannot win matters as much as recognizing where it can.

For marketing agencies serving the longevity category

The standard search-optimization playbook helps but is not sufficient. Several shifts are worth making deliberately.

Rebuild measurement. An agency reporting on conventional Google rankings is reporting on a different competitive surface from the AI search surface that increasingly drives client discovery. Tracking AI search visibility across all four platforms, separately from organic search, is the precondition for demonstrating value on the dimension that now matters.

Treat structured data as a core advisory deliverable, not a technical handoff. Schema markup, FAQ structure and extractable formatting are what make a clinic's self-published authority legible to the platforms. Most agencies route schema work to the client's development team as a technical implementation task. This is the wrong staffing model for AI search work. Schema architecture for AI citation requires medical-vocabulary precision, platform-specific weighting decisions, entity-resolution strategy and ongoing recalibration as platforms shift their parsing behavior. None of that is generalist developer work.

Build fluency in the industry shifts, not only in marketing tactics. Insurance-tier compression and vertical-integration consolidation reshape the competitive set in ways tactical search work does not address. An agency that can advise on positioning against those dynamics operates at a different level from one that stays narrow on keyword work.

Advise clients toward the skeptical-press ecosystem. The mainstream outlets and peer-reviewed journals that AI platforms cite on legitimacy-establishing queries are the authority sources. Clients that contribute sober, evidence-led work into those publications accumulate authority-tier citations no other channel produces.

For editorial publishers covering the longevity industry

The editorial lane between geroscience research and clinic marketing is structurally underserved.

Produce evidence-led category coverage. AI platforms route evidence-evaluation queries through academic journals and mainstream press because no third-party editorial brand consistently occupies the lane between the research and the clinics. A publisher producing evidence-led coverage that translates what longevity science is funding and proving into what it means for the people seeking care today would occupy a citation position the platforms currently have no native source for.

Build genuinely parallel coverage across languages. The multilingual testing confirmed that the same market surfaces differently in different languages, richer in domestic operators in one and heavier on foreign-brand expansions in the other. A publisher producing parallel coverage in more than one language would surface across AI-language markets simultaneously and become the cross-language reference for the category. The logic applies to all major non-English longevity markets.

Earn placement in the publications AI platforms already treat as authority. A publisher that becomes a cited contributor to the mainstream press and peer-reviewed venues the platforms weight becomes, itself, an AI-cited authority. The category-defining influence follows from the citation.

Produce independent comparison content. The first rigorously researched, genuinely independent set of regional "best longevity clinics" guides, with verified pricing, evidence-quality flags and named-criteria comparison, would displace clinic-owned content from those citation positions over a year-or-two window. The barrier is editorial rigor, the discipline to publish nothing that reads as advertorial. The gap between the research and the marketing, occupied by nobody and cited by the platforms by default, is the opportunity for editorial publishers.

The single highest-leverage move for each audience

A clinic operator should commit to self-published authority on the one or two categories where the clinic genuinely leads and hold that commitment long enough for it to compound. Every other recommendation compounds more slowly.

An agency should rebuild client measurement to track AI search visibility as a discipline distinct from search-engine optimization. Reporting on conventional rankings while clients are losing share on the AI search surface is the misalignment the next twelve months will expose.

An editorial publisher should occupy the independent-comparison and academic-bridge lane that AI search is currently filling with clinic-owned content and academic journals by default. That lane is both unoccupied and actively cited, and the window in which it can be entered is bounded.

SECTION 02

Methodology, disclosures & notes

This report measures one thing: how 25 leading longevity clinics appear in AI search results across the four AI platforms prospective clients are most likely to use in 2026. It does not measure organic search rankings, conversion rates, web traffic or clinic revenue. It measures who AI search platforms recommend, name, cite and link to when asked about longevity care. This is the metric clinic marketing teams increasingly need to track and have had no systematic way to measure.

Disclosure of editorial interest

Healthspan Media publishes editorial research and analysis covering the global healthspan economy. The company also provides paid advisory services to clinics, agencies and publishers operating in this sector, which may include organizations named in this report.

Commercial advisory activities are managed separately from the editorial research process. Clients do not receive editorial control, pre-publication review rights or guaranteed favorable coverage, inclusion or ranking outcomes. No clinic in the study panel paid for inclusion, exclusion or specific framing within the report.

The findings reflect documented observations of AI-platform behavior collected during the defined study window under the methodology described in this publication. Inclusion in the report should not be interpreted as endorsement, certification or commercial recommendation.

Data collection

All queries were conducted manually through publicly available user interfaces during the study window. The study did not use automated scraping, API extraction or circumvention of platform controls.

AI assistance in production

The report was researched, drafted and revised with substantial AI assistance, using Claude (Anthropic), ChatGPT (OpenAI) and Perplexity for research synthesis, drafting and editorial review. All findings, framings, evidence anchors and final editorial decisions are the responsibility of the editor and Healthspan Media as publisher. AI assistance was used as an editorial tool throughout the process; AI outputs were not published without human editorial review.

What was measured

The study covers a fixed query panel of 50 questions executed across four AI platforms (Google AI Overviews, Perplexity, ChatGPT and Claude), producing 200 query-platform observations. The query panel was designed to cover the full range of prospective-client research behavior: brand-direct queries (asking about a specific clinic or practitioner), category queries (asking for recommendations without naming a clinic), service and protocol queries (asking about specific diagnostic or therapeutic offerings), comparison queries (pitting clinics or platforms against each other) and client-research queries (the questions a person asks before booking, including skeptical framings). The structure surfaces both how AI platforms handle direct brand inquiries and how they construct recommendations when no brand is named.

Supplementary panels were added during the study window to test specific anomalies and to extend the scope into markets the primary panel did not reach. An Arabic-language probe was conducted as a deliberate scope test of the English-language methodology's limits. Where supplementary findings materially changed conclusions, the change is documented in the relevant section.

The 25 clinics

The 25 clinics were selected against four criteria: geographic diversity across the markets where premium longevity care is currently provided, business-model diversity across medical-led, resort-clinical-hybrid, membership and disruptor models, sufficient public content presence to be plausibly visible to AI search, and editorial defensibility against the question a clinic CMO would reasonably ask: why these and not those? The final list balances established category leaders (whose visibility patterns define the competitive ceiling) with newer entrants (whose visibility patterns reveal what works and what does not for a clinic without decades of accumulated authority).

Geographic distribution: United States (10), Europe (6), Middle East (4), Asia-Pacific (4), with a disruptor tier representing newer business models within those geographies. The geographic weighting is a methodological choice rather than a definitive map. AI platforms are trained predominantly on English-language content, and clinic visibility in English-language AI search is partly a function of English-language web presence. Supplementary panels tested the limits of this English-language scope.

Scoring

Each clinic's appearance in each query response was scored on two independent axes: a visibility score capturing whether and how prominently the clinic appears in the AI response, and a brand integrity score capturing whether the clinic is represented accurately when it does appear. The two scores are independent because the underlying problems they measure are independent. A clinic can be highly visible but misrepresented (the most actionable diagnostic finding, because the clinic is reaching prospects with the wrong message), or accurately represented but rarely visible (a content-scale problem with a different remediation path).

Visibility was scored from the AI response itself on a graduated scale capturing whether the clinic was absent, mentioned in passing, described in detail, cited with a source link or anchored as a primary recommendation. Modifiers adjusted scores for position, sentiment and the structural difference between substantive citation and list-mention.

Integrity was scored manually against each clinic's own published positioning. Integrity cannot be scored by the AI platforms themselves without asking them to grade their own output. The Healthspan Media team scored integrity by comparing each AI response against the clinic's homepage, services page, founder biography and most recent twelve months of press, producing a judgment of whether the AI representation matched, partially matched, materially diverged from or fabricated the clinic's actual positioning.

Working scoring tags maintained during data collection were verified against the raw query-platform captures before final analysis. Where working tags and raw captures diverged, the raw captures were treated as authoritative.

Controlling for platform variance

Two methodological challenges shaped how data was collected. The first is geographic localization: AI platforms with IP-based geolocation return different recommendation sets depending on where the user appears to be searching from, and a study run from a single location risks producing geographically biased results. The second is platform tier variance: AI platforms produce different output on different subscription tiers, and a study conducted on a single tier risks confounding architectural findings with tier-specific behavior.

Both challenges were addressed through documented controls developed during the study. Findings on platform-architectural behavior (template suppression, cross-platform divergence, citation patterns, integrity case studies) are independent of these controls and hold regardless of how they are configured. Findings on relative answer quality between platforms are partly sensitive to tier and are reported with that caveat.

What this report does not measure

The report measures AI search visibility today. It does not measure organic search rankings, conversion rates, traffic, revenue or any clinic-side operational metric. It does not exhaustively measure AI search behavior in languages other than English; the Arabic-language probe was a scope test, and broader multilingual coverage is a future-work item. It does not attempt to predict future visibility quantitatively or to model the marketing investment required to shift current visibility. These are downstream questions the report's findings should inform without answering directly.

The methodology measures what AI platforms surface in response to queries from the panel, not whether clinics exist in the underlying training data or could be returned by specific brand-name searches. Where clinics in the study panel produced no observations across the 200 query-platform combinations, the finding is that these clinics did not surface in response to the queries clients are likely to ask, not that the AI platforms are unaware of the clinics' existence.

Limitations

The 25-clinic scope is a sample. Conclusions about the most-cited clinics in this sample generalize within the sample but should not be extrapolated to clinics or markets the study did not cover without verification.

The English-language weighting means non-English-language longevity ecosystems are underrepresented. Markets with substantial domestic longevity provision in non-English languages may be AI-visible in those languages even where they appear invisible in English-language AI search.

Single-query results may reflect ordinary platform variance rather than stable patterns. Findings throughout the report distinguish between patterns that replicate across multiple queries or platforms and single-query observations, which are noted as such.

AI platforms cite community-aggregation sources (Google reviews, TripAdvisor, social platforms) at meaningful rates. The competitive layer those sources operate on is real and worth tracking, but is partially outside the content-strategy scope of this report's recommendations.

AI search is a moving target. The patterns documented in this report reflect platform behavior during the study window. Some patterns will hold over time; others will shift as platforms update their retrieval, ranking and citation systems.

What is held back

The complete operational layer behind this analysis is documented internally: the full query panel, the scoring rubric with modifiers and edge-case rules, the inter-rater calibration procedures, the platform-control protocols, and the ongoing tracking as the platforms shift. The report presents the structural findings and the strategic implications; the per-clinic application of those findings is the work of audit engagement.

COLOPHON

Published by Healthspan Media, 2026. Methodology, data collection and analysis by the Healthspan Media editorial team, with AI-assisted research and drafting throughout. See Section 2 for full editorial disclosures.

AI Visibility Report 2026 reflects fieldwork conducted across the four major AI search platforms (Google AI Overviews, Perplexity, ChatGPT and Claude) during 2025-2026, anchored on a 25-clinic panel and a 50-query test set. This is the first annual edition; the next is planned for 2027.

Healthspan Media is an editorial studio covering the global healthspan economy. It works with clinics, agencies and editorial publishers on the operational application of the report's findings.

Inquiries regarding the audit methodology, engagement or commentary on the report's findings: hello@healthspanmedia.com.

Healthspan Media | healthspanmedia.com | © 2026

Where does your clinic *stand*?

This report maps where 25 clinics stand in AI search. Healthspan Media does the same for individual clinics, and then closes the gap.

We begin with a per-clinic AI visibility audit: where you surface across Google AI Overviews, Perplexity, ChatGPT and Claude, where your citation share is being lost, and which of the six patterns you can realistically build. From there, we do the editorial and structural work the report describes, the self-published authority content, the schema and extractability, the answer-engine optimization that earns citations rather than buys placement.

The clinics that establish citation share now are building positions later entrants cannot reach. To discuss where your clinic stands, get in touch.