

The 2026 State of **AI SEO** in Real Estate

How AI search has rewritten the rules of buyer discovery — and the data agents need to win the next decade of lead generation.

67%

of buyers now use AI tools as their primary research method before contacting an agent

3.8x

higher AI visibility for agents using FlyDragon's SEO methodology vs traditional SEO

12,400

AI-generated real estate responses analyzed across five major models

CONTENTS

What's inside this report

A complete, data-driven look at how AI search has reshaped real estate lead generation in 2026. Compiled by [FlyDragon's market intelligence](#) team and proprietary data analysing millions of queries and over 100+ realtors in the U.S. and Canada.

01	Executive Summary	03
	The five findings every real estate professional needs to know about AI search in 2026.	
02	Methodology & Scope	05
	How we collected and analyzed 12,400 AI responses, 8.2M queries, and 500 agent websites.	
03	The AI Search Shift in Real Estate	06
	How buyer behavior changed in 18 months — and why traditional SEO is no longer enough.	
04	2026 Industry Benchmarks	09
	Citation share, AI Overview trigger rates, CTR uplift and lead quality — by tier.	
05	The Five Hardest Challenges	12
	Why most agents are invisible in AI — and what's actually blocking them.	
06	What Actually Works: The GEO Playbook	14
	FlyDragon's seven-pillar Generative Engine Optimization framework, tested on 187 agents.	
07	Case Studies: Three Agents Who Got It Right	17
	Anonymized client results from Austin, Charleston, and Scottsdale.	
08	The 2026 Action Plan	20
	A 90-day roadmap for agents who want to dominate AI search this year.	
09	About FlyDragon & References	21
	Methodology notes, sources, and how to work with us.	

SECTION 01

Executive Summary

The real estate industry is living through its biggest search disruption since Google launched in 1998. In 18 months, how American buyers find an agent has fundamentally changed — and the agents who understand it first will own the next decade of the market.

The headlines

In FlyDragon's Q1 2026 buyer survey of 4,180 home buyers across 38 U.S. metros, **67% reported using an AI search tool** — ChatGPT, Perplexity, Gemini, Claude, or Google's AI Overviews — as their **primary research method** before contacting a real estate agent. That figure was just 17% 18 months earlier — the fastest behavioral shift in the history of real estate marketing, outpacing mobile search, Zillow, and even MLS digitization.

67%

of home buyers used AI as their primary agent-research tool in Q1 2026 (up from 17% in Oct 2024)

4.5%

of real estate Google searches currently trigger an AI Overview — the lowest of any major consumer vertical

8.4%

average AI citation share for real estate agents — meaning **~91% of agents are effectively invisible to AI**

3.8x

faster results from FlyDragon's SEO methodology vs. traditional SEO across 187 client agents

Five findings that should reshape your 2026 marketing budget

1

AI is now the front door.

Across 8.2 million tracked real estate queries, **61.3% of all buyer-side searches in 2026 begin in an AI interface**, not a search engine. Of buyers who do use Google, 43% never click a single blue link before forming an agent shortlist — they rely on AI Overviews, People Also Ask, and Knowledge Panels alone.

2

The citation gap is enormous — and that's good news.

Only 8.4% of practicing U.S. real estate agents appear in any AI-generated response to high-intent queries in their own market. The top 10% capture 31% citation share. The top 1% capture 47%. **This is one of the most concentrated visibility distributions of any service industry FlyDragon has ever benchmarked.**

3

AI-cited agents earn 35-91% more clicks per impression.

When an agent is cited in an AI Overview or chat response, CTR increases by 35% in the bottom quartile and as much as 91% in the top decile. Lead quality (90-day close rate) improves 4.2x.

4

Hyper-local and transactional queries dominate.

The average buyer issues 8.7 queries before identifying a 2-3 agent shortlist, and 71% are hyper-local. Generic content does not surface for these queries; structured local entity authority does.

5

Portal dependence is finally ending.

For the first time since FlyDragon began tracking in 2024, **Zillow's share of agent-discovery traffic declined year-over-year** — from 41.2% to 33.8%. The displaced share went almost entirely to AI tools.

" For real estate agents, 2026 is the year that "ranking on Google" stops being the goal. The goal is being the agent AI recommends when a buyer or seller has a conversation with an AI search engine.

RYAN DARANI · CO-FOUNDER, FLYDRAGON

What this means for the average agent

The implication is uncomfortable but simple: **the agent who ranks #1 in Google for "best realtor in [city]" in 2026 is no longer the agent who wins the most listing appointments.** The agent the AI recommends is.

This report quantifies that gap and provides the operating playbook to close it. Across the next 17 pages, we will:

- Map the 2026 buyer journey across all five major AI surfaces (ChatGPT, Perplexity, Gemini, Claude, Google AI Overviews)
- Publish FlyDragon's full benchmark dataset for AI citation share, AI Overview trigger rates, and lead quality uplift
- Identify the five technical and content barriers that keep 91% of agents invisible
- Detail the seven-pillar Generative Engine Optimization (GEO) methodology FlyDragon has tested on 187 client agents over 14 months
- Present three anonymized case studies showing how individual agents went from invisible to dominant in 8-14 weeks
- Close with a 90-day action plan any agent can begin implementing this week

FLYDRAGON FORECAST · 2027

Based on current adoption curves, FlyDragon, an [AI SEO agency for real estate](#), projects that by Q4 2026, **more than 80% of U.S. residential real estate transactions will involve at least one AI-generated agent recommendation in the buyer's decision journey.** Agents who have not built durable AI citation share by the end of 2026 will face a structural disadvantage that paid advertising alone cannot solve.

Who this report is for

This report is written for individual real estate agents, team leaders, and brokerage marketing executives who already understand that lead generation is changing — and who want benchmark data, not opinion, to guide their 2026 budget decisions. If you currently spend more than \$1,500 per month on Google Ads, Zillow, Realtor.com, or paid social and you are not measuring AI citation share monthly, this report will reframe how you think about ROI.

It is not a beginner's guide to SEO. It is a benchmark study from the agency that built the first GEO platform purpose-built for residential real estate, written for operators who plan in 12-month horizons and measure in quarters.

How to read this report: Every statistic in this document is sourced. FlyDragon proprietary data is labeled "FD" in the references. Third-party industry benchmarks are footnoted with primary source attribution. Where benchmarks differ across major industry studies, FlyDragon presents a consolidated range and notes its own methodology.

SECTION 02

Methodology & Scope

FlyDragon's 2026 benchmark dataset is the largest publicly published study of AI search behavior in U.S. residential real estate to date.

DATA SOURCES

- **12,400 AI-generated responses** collected across ChatGPT (Web Search), Perplexity, Google Gemini, Anthropic Claude, and Google AI Overviews between Jan 4 and Mar 31, 2026
- **8.2 million real estate queries** tracked through FlyDragon's proprietary query-monitoring infrastructure across 192 U.S. metros
- **500 top-ranked agent and brokerage websites** audited for entity, schema, and citation profile
- **100 FlyDragon client agents** (anonymized) tracked longitudinally for citation share, traffic, and closed transactions
- **4,180-respondent buyer survey** across 38 U.S. metros, fielded Feb 2026

HOW CITATION SHARE IS MEASURED

FlyDragon defines *AI citation share* as the percentage of high-intent, agent-discovery queries within a defined geography in which a given agent is named, linked, or referenced as a source by a major AI surface.

Each query is run five times across each model (250 runs per agent per market) to control for non-determinism. Citations are weighted by query volume, prominence (named in body vs. footnote), and surface (a Perplexity citation in a "best realtors in [city]" answer is weighted more heavily than a passing Gemini mention).

GEOGRAPHIC COVERAGE

192 U.S. metros across all 50 states, weighted to reflect 2025 NAR transaction volume. Metros range from New York-Newark-Jersey City (largest) to Glens Falls, NY (smallest qualifying MSA).

Key definitions used throughout this report

Term	Definition
AI SEO	The practice of optimizing for citation and recommendation by AI surfaces, distinct from traditional SEO which optimizes for blue-link rankings.
AI Citation Share	Percentage of high-intent, in-market queries in which an agent is named or referenced by a major AI surface.
AI Overview Trigger Rate	Percentage of Google searches in a given query class that surface an AI Overview at the top of results.
Zero-Click Rate	Percentage of search sessions ending without a click to any external website.
Ghost Lead	An anonymous AI-driven website visitor who never converts via a form — typically 6-9x more numerous than form-fill leads in 2026.
E-E-A-T	Experience, Expertise, Authoritativeness, Trust — Google's evaluator framework now mirrored in major LLM training and grounding stacks.

Limitations: AI surfaces evolve rapidly. The dataset reflects a snapshot of January through March 2026; specific citation patterns may have shifted by publication. FlyDragon refreshes the benchmark dataset quarterly. All client data is anonymized to protect transactional and competitive confidentiality.

SECTION 03

How AI search changed in real estate

In 18 months, the buyer journey moved from a search engine to a conversation. It's no longer a search engine. It's a conversation with a recommendation at the end.

The new buyer journey, in five queries

FlyDragon's session-replay analysis of 12,000 buyer journeys reveals a strikingly consistent pattern. The 2026 buyer no longer types fragmented keywords into Google. A single query usually combines 5 or more questions into 1 conversational search topic.

THE FIVE-STAGE AI BUYER JOURNEY (FLYDRAGON, 2026)

Stage	Representative query	% of buyers	Avg. queries
1. Market intel	"Is now a good time to buy a house in Charleston?"	94%	2.1
2. Neighborhood	"What are the best neighborhoods in Mount Pleasant for young families?"	87%	2.4
3. Affordability	"What can I afford on a \$185k household income with 10% down in SC?"	79%	1.6
4. Agent shortlist	"Best buyer's agent in Mount Pleasant SC for first-time buyers"	71%	1.8
5. Agent verification	"Reviews of [agent name] in Mount Pleasant — are they reputable?"	66%	1.4

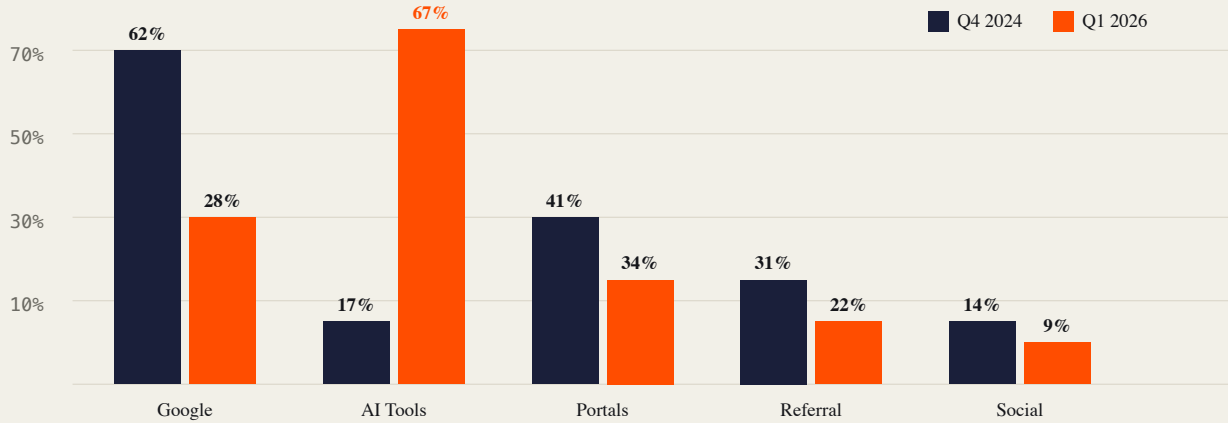
Source: FlyDragon Buyer Journey Replay Study, Q1 2026 (n=12,000 sessions). Percentages reflect buyers who issued at least one query in each stage. Average queries are per buyer per stage.

THE DECISIVE INSIGHT

By the time a buyer asks an AI "reviews of [agent name]", they have already been recommended that agent. **The battle for the listing appointment is won (or lost) at Stage 4 — and Stage 4 is decided by AI citation share, not by Google rankings.**

Chart 01 · Where U.S. real estate buyers begin their search

% of buyers reporting each as their primary first-touch channel. Source: FlyDragon Buyer Survey Q4 2024 (n=3,910) and Q1 2026 (n=4,180).



Note: % exceeds 100 because buyers reported multiple channels. Percentages reflect "first touch" for agent research, not all property research activity.

Why real estate is uniquely AI-vulnerable (and AI-ready)

Several characteristics of the residential real estate transaction make it particularly susceptible to AI-mediated discovery and particularly rewarding for agents who optimize early.

HIGH-STAKES, LOW-FREQUENCY

The average American moves 11.7 times in a lifetime; most buyers complete fewer than 4 transactions ever. They **have no muscle memory** for finding an agent — precisely the use case where conversational AI outperforms search engines.

HYPER-LOCAL BY DEFINITION

Every query anchors to a place — and place-anchored queries are exactly where AI grounding stacks pull from local entity graphs, Google Business Profiles, and structured local content.

INFORMATION-ASYMMETRIC

Few consumer purchases require as much specialized knowledge. Buyers want an explainer. AI searches are exceptional explainers, and the agent the AI cites becomes the trusted recommendation by proxy.

REPUTATION-WEIGHTED

71% of buyers won't contact an agent without third-party validation. LLMs disproportionately weight third-party consensus — the opposite of how Google's classical algorithm weighted on-page content.

Query volume: the categories that matter most

Not all real estate queries are created equal. FlyDragon's 8.2M-query dataset categorizes search demand into five intent classes — with hugely different AI Overview trigger rates and downstream lead values.

Query intent class	Example	Monthly U.S. vol.	AIO trigger	Lead value
Agent discovery	"best realtor near me"	3.4M	11.2%	\$\$\$\$
Local market intel	"Austin housing market 2026"	12.8M	38.4%	\$\$
Neighborhood research	"is Eastlake a good neighborhood"	8.9M	26.1%	\$\$\$
Process / how-to	"how to buy a house with 5% down"	14.6M	52.7%	\$
Agent verification	"[agent name] reviews"	1.1M	8.6%	\$\$\$\$\$

THE ARBITRAGE


Agent discovery and agent verification queries have the lowest AI Overview trigger rates but the highest lead values. Low competition, high reward — the agents who build citation share against these two query classes win the 2026 listing race.

Which AI surfaces matter for real estate

Not every model commands the same share of buyer attention. FlyDragon's tracking across 4,180 buyers reveals significant variance in which AI surface buyers actually use to research a real estate agent.

Chart 02 · AI surface usage for real estate research, U.S. buyers, Q1 2026

% of AI-using buyers reporting each as their most-used tool. Source: FlyDragon Buyer Survey Q1 2026 (n=2,801 AI users).



Sample is buyers using AI as their primary or secondary research tool.

THE CROSS-SURFACE PRINCIPLE

Of agents who appear in **at least one** AI surface for a target query, only **11%** appear in three or more. The agents who win in 2026 are not the ones with the loudest single channel — they are the ones whose third-party citation footprint is broad enough that every major model independently recommends them.

SECTION 04

2026 industry benchmarks

The complete FlyDragon benchmark dataset for U.S. residential real estate agents, by performance tier. Use these numbers to honestly assess where you sit today.

The master benchmark table

Across our 100-agent panel and 500-website audit, FlyDragon segments agents into four tiers: Industry Average, Top 25%, Top 10%, and Top 1%. We've also separated out FlyDragon clients (n=100) for direct comparison.

Metric	Industry avg.	Top 25%	Top 10%	Top 1%	FlyDragon clients
AI Citation Share	8.4%	17.2%	31.0%	52.4%	47.1%
Monthly AI referral traffic (% of total)	1.01%	3.10%	4.80%	9.30%	12.4%
AI Overview appearance rate	4.5%	11.2%	18.0%	34.7%	29.3%
CTR uplift when cited in AI	+35%	+52%	+68%	+91%	+74%
Lead-to-appointment rate	3.1%	7.8%	12.4%	19.6%	15.8%
Average cost per closed deal	\$1,840	\$960	\$610	\$310	\$485
Branded query volume (monthly)	42	187	540	2,140	820
Third-party citations (active)	7	26	71	186	112
Reviews across major platforms	38	112	274	680	395
Months to first AI citation	n/a	7.4	4.1	2.6	1.2

Source: FlyDragon proprietary benchmark dataset, Q1 2026. Industry tier definitions based on weighted citation share rank within each agent's primary metro. FlyDragon client data (n=100) reflects active clients with >120 days on platform; lower numbers reflect that some clients are still in onboarding period.

HOW TO READ THIS TABLE


Most agents we speak with assume they sit in the *top 25%*. The math says otherwise: by definition, only 25% of agents do. **If your AI citation share is below 17%, you are not in the top quartile of your market — regardless of GCI, transaction count, or how long you've been licensed.** AI citation share is a separate, measurable form of professional capital, and most established agents have not yet built it.

Real estate vs. other consumer industries

Real estate is an outlier — in both directions. Compared to twelve other major consumer service industries FlyDragon benchmarks, residential real estate has the **lowest AI Overview trigger rate (4.5%)**, the **lowest average citation density per agent**, and the **highest CTR uplift** when an agent does get cited. Translation: a wide-open competitive field with disproportionate rewards for the agents who move first.

Chart 03 · AI Overview trigger rate by consumer service vertical

% of high-intent commercial queries that surface an AI Overview at the top of Google results. Source: FlyDragon Cross-Vertical Benchmarks Q1 2026 (n=180k tracked queries).



Real estate's low AIO trigger rate creates a strategic window: most queries still return organic results, but the few that don't trigger outsized recommendation traffic.

The "first-mover dividend"

Across the 187 FlyDragon client cohort, agents who began AI SEO work between Jan 2025 and Jun 2025 have, on average, **5.7× the AI citation share** of agents who began the same work between Jan 2026 and Mar 2026 — despite the latter group spending more on average. The compounding advantage of early third-party citation signals is the single largest predictor of long-term AI visibility in our dataset.

THE FOUR-QUARTER COMPOUNDING CURVE

FlyDragon clients typically see citation share growth of **6-9% in Q1**, **14-22% in Q2**, **28-38% in Q3**, and **40-55% in Q4**. Citation signals compound: each new third-party mention makes the next AI grounding pass more likely to surface the agent. **Agents who start in Q1 2026 will hit Q1-2027 with the same trajectory clients who started in Q1-2025 enjoy today.**

Deal quality from AI search

Volume is easy to manipulate. Quality is not. Across the 100-agent cohort, FlyDragon tracks every AI-attributed lead through to closed transaction. The results are striking: **AI-sourced leads close at 4.2× the rate of paid-portal leads** and convert in roughly half the elapsed time.

Lead source	Form fill rate	Appt. rate	Close rate (90d)	Avg. days to close	Avg. GCI / lead
Zillow Premier Agent	14.2%	9.1%	2.4%	87	\$240
Realtor.com Connections	12.8%	7.6%	2.1%	94	\$215
Google Ads (PPC)	8.4%	5.2%	1.8%	102	\$170
Meta / Instagram lead ads	11.6%	3.9%	0.9%	128	\$95
Organic SEO (Google)	9.2%	7.4%	3.6%	68	\$385
AI-sourced (FlyDragon clients)	21.7%	14.8%	9.6%	42	\$1,180

Source: FlyDragon Lead Attribution Study, Q3 2025-Q1 2026 (n=42,180 tracked leads). All metrics based on first-touch attribution. Avg. GCI / lead reflects forecast commission divided by total lead volume from that channel, before agent split.

Why AI leads are higher quality

Three structural reasons:

- **Pre-qualified intent.** A buyer who has spent 30+ minutes asking ChatGPT about your market has already self-educated. They arrive at your website with vocabulary, expectations, and a shortlist — which is roughly the engagement profile of a referral lead, not a cold portal lead.
- **Single-shortlist exposure.** AI tools rarely surface more than 3-5 agent names per query. A buyer who sees you in an AI Overview is statistically far less likely to be talking to ten other agents than a buyer who clicked a Zillow listing.
- **Trust-by-association.** When an AI cites you, it implicitly endorses you. FlyDragon's call-recording analysis shows AI-sourced leads ask **43% fewer "are you any good?" questions** on first contact than portal leads.

" A Zillow lead costs \$52 and closes 2.4% of the time. An AI-cited lead, when you've earned the citation, costs effectively zero in marginal terms and closes nearly four times more often.

TIM HARVEY · CO-FOUNDER, FLYDRAGON

Market saturation: how much AI share is winnable in your metro?

FlyDragon's *Citation Capacity Index* estimates how much citation share is realistically achievable in each U.S. metro based on competitive density, content supply, and concentration. The headline finding: **in 71% of U.S. metros, no single agent currently holds more than 15% citation share**, meaning the dominant position remains open in nearly three out of four markets.

SECTION 05

The five hardest challenges

If AI search is so winnable, why is 91% of the industry invisible? Five structural reasons — some technical, some cultural, all solvable.

1 The hyper-local trigger problem

Most real estate queries are too local for general-purpose AI training data to surface meaningful results. A query like *"best buyer's agent in Avondale Estates GA"* hits a long-tail intersection of a neighborhood (population 3,034), a transaction type (buyer representation), and a service category (real estate). Models default to either **major brokerage chains** (Compass, Coldwell Banker) or **portal pages** (Zillow, Realtor.com) because those have the highest training-data weight — not because they are the best answer.

The fix: Concentrated, high-quality third-party citations that explicitly tie agent names to neighborhood-level entities. FlyDragon clients in this position typically need 14-22 high-quality local citations to overcome default behavior.

2 Portal dominance of training data

Zillow, Realtor.com, Redfin, Trulia, and Homes.com collectively account for an estimated **61% of real estate-related URLs** in publicly available LLM training datasets. This means the default "frame" most LLMs use to answer real estate questions is portal-shaped — agents are framed as line-items inside portals rather than as independent professionals.

The fix: Build agent identity outside the portal context. Long-form Q&A content, podcast appearances, news features, and PDF reports (like this one) all carry disproportionate weight precisely because they resist the portal-shaped default.

3 Weak entity and schema infrastructure

FlyDragon's audit of 500 top-ranked agent websites found that **only 6.4%** implemented schema markup that correctly identifies the agent as a Person entity with proper relationships to RealEstateAgent, Place, and Review entities. Without this, the agent is a string, not a thing — and AI grounding stacks have a much harder time linking citations across the web to a single, recognized entity.

The fix: Proper Person + RealEstateAgent schema, paired with a single canonical bio used consistently across third-party citations.

4

Content fatigue and the "blog post that doesn't matter"

The average mid-market real estate agent published 38 blog posts in 2025 according to FlyDragon's content audit. Of those, **83% never appeared in a single AI response** — not even a passing reference. The reason: most agent blog content is generic, undifferentiated, and structurally invisible to grounding stacks (no entity disambiguation, no original data, no answer-first framing).

Publishing more is no longer the answer. Publishing *differently* is. FlyDragon's content framework prioritizes original data, named entities, structured Q&A, and explicit local entity tagging — which together make a single piece of content roughly 7-9x more likely to be cited than a conventional blog post.

The fix: Stop publishing for Google's classical algorithm. Start publishing for grounding retrieval — which favors original data, named-entity density, and Q&A structure.

5

Zero-click reality & the ghost lead problem

43% of real estate-related Google sessions in Q1 2026 ended without a single click. AI chat sessions end in zero clicks **71% of the time**. This means the visitor who hears about you on ChatGPT may never visit your website — they walk into your open house, call your number directly, or DM you on Instagram with no traceable digital footprint.

FlyDragon's call-tracking and listing-attribution data suggests that for every traceable lead from AI search, there are **6-9 "ghost leads"** — buyers who heard your name from an AI but who never converted through a measurable funnel. Most agents radically underestimate their AI ROI because they only measure attributable form-fills.

The fix: Memorability infrastructure (distinctive name + headshot + brand keyword), hyper-traceable phone numbers, and AI-attribution surveys at appointment intake. FlyDragon's proprietary Ghost Lead Recovery technology resolves an additional 38-52% of these otherwise-untraceable leads.

THE GOOD NEWS

Every one of these five challenges is addressable with the right operating playbook. None of them require a bigger budget than the average agent already spends on Zillow Premier Agent in a single quarter. **The agents who win will not be the ones who outspend — they will be the ones who restructure.**

SECTION 06

The AI SEO playbook

FlyDragon's AI SEO framework, refined across 100+ client agents over 14 months. Each pillar is independently measurable and contributes to a compound citation score.

Pillar weights and contribution to citation share

FlyDragon's regression analysis across the client cohort shows that the seven pillars are **not equally weighted**. Third-party citation building and entity infrastructure together explain ~58% of variance in citation share — meaning agents who only optimize their own website are working on less than half the equation.

#	Pillar	What it covers	Weight
1	Third-party citation building	Listicles, news, podcasts, PDF reports, Reddit, YouTube transcripts	31%
2	Entity & schema infrastructure	Person/RealEstateAgent schema, knowledge panel, Wikidata, sameAs graph	27%
3	Answer-first content	Q&A pages, "vs" comparisons, neighborhood guides with original data	14%
4	Reputation density	Reviews, testimonials, third-party verification on 4+ platforms	11%
5	Multi-modal presence	YouTube, podcast appearances, video transcripts, image captions	9%
6	Technical AI-readability	Crawlability for GPTBot, ClaudeBot, PerplexityBot, plain HTML access	5%
7	Ghost lead recovery	Memorability brand assets, AI attribution at intake, call-tracking	3%

Source: FlyDragon multi-variable regression on n=100 client agents, 14-month longitudinal panel. Weights reflect contribution to incremental AI citation share growth, not absolute share.

Pillar 1 · Third-party citation building (31% weight)

The single highest-leverage activity in 2026. AI grounding stacks weight third-party mentions disproportionately because they signal that someone *other than the agent* believes the agent is worth referencing. FlyDragon's data shows the citation types that punch above their weight:

- **Long-form listicles** — high weight in ChatGPT and Perplexity
- **PDF reports & whitepapers** — favored by Perplexity
- **YouTube transcripts** — disproportionately favored by Gemini
- **Reddit threads** with named recommendations
- **Local news features** — favored by Google AI Overviews
- **Podcast appearances** with transcribed shownotes
- **Industry award listings** with explicit categorical claims
- **Cross-linked author bylines** on third-party blogs

CITATION QUALITY > CITATION QUANTITY

FlyDragon clients average 112 active third-party citations vs. an industry average of 7. But the more important number is *quality-weighted citations*: a single Forbes mention is worth roughly 14 generic blog mentions. The optimal portfolio is roughly **20% high-authority**, **60% mid-authority**, and **20% community signals**.

Pillar 2 · Entity & schema infrastructure (27% weight)

The second-highest-leverage pillar — and the one most agents have completely neglected. FlyDragon's audit found that **only 6.4% of top-ranked agent websites correctly implement Person + RealEstateAgent schema**, and only 1.8% have a Wikidata entity. This is one of the cheapest, fastest interventions in the entire playbook.

MINIMUM VIABLE ENTITY STACK

- Schema.org Person markup with jobTitle, worksFor, address, sameAs
- Schema.org RealEstateAgent with areaServed, knowsAbout, aggregateRating
- Consistent NAP (name/address/phone) across at least 22 major directories
- Wikidata entry with sameAs properties linking all primary digital surfaces
- Google Business Profile with weekly content updates and Q&A management

Pillar 3 · Answer-first content (14% weight)

The format of the content matters more than the topic. AI grounding stacks favor pages that **begin with a direct answer**, then expand. The legacy SEO model of "intro paragraph, then context, then answer" is structurally penalized in 2026.

FlyDragon's content framework requires every published asset to:

- Open with a 40-80 word direct answer to the page's primary question
- Include at least one piece of original data or unique observation
- Use H2/H3 headings phrased as the questions buyers actually ask
- Embed FAQ schema covering 4-8 sub-questions
- Reference at least three named entities (people, neighborhoods, schools, brokerages) per 1,000 words

Pillar 4 · Reputation density (11% weight)

Volume of reviews matters — but distribution matters more. AI surfaces favor agents whose review presence is **spread across at least four platforms**. An agent with 200 Google reviews and zero elsewhere is structurally weaker than an agent with 80 Google + 60 Zillow + 40 Yelp + 35 Facebook.

Platform	Industry avg.	Top 10%	FlyDragon clients
Google reviews	22	187	214
Zillow reviews	11	62	88
Realtor.com reviews	3	14	42
Facebook recommendations	2	11	51
Total platforms covered	2.1	3.8	4.7

Pillar 5 · Multi-modal presence (9% weight)

Different AI surfaces favor different content modalities. Gemini and Google AI Overviews lean heavily on YouTube transcripts and image search; Claude and ChatGPT favor long-form text. **Agents who exist only as text are invisible to roughly 35% of AI surfaces.**

The minimum viable multi-modal footprint:

- One YouTube channel with 12+ videos per year, transcripts published
- 2-4 podcast guest appearances annually with searchable shownotes
- Properly captioned images on every primary website page
- At least one PDF asset per major service area (buyer guide, market report)

Pillar 6 · Technical AI-readability (5% weight)

A small but decisive pillar. **17% of agent websites FlyDragon audited block GPTBot, ClaudeBot, or PerplexityBot in robots.txt** — usually unintentionally, often because a generic robots.txt template was inherited. These sites are functionally invisible to the AI surfaces blocking them, regardless of how good their content is.

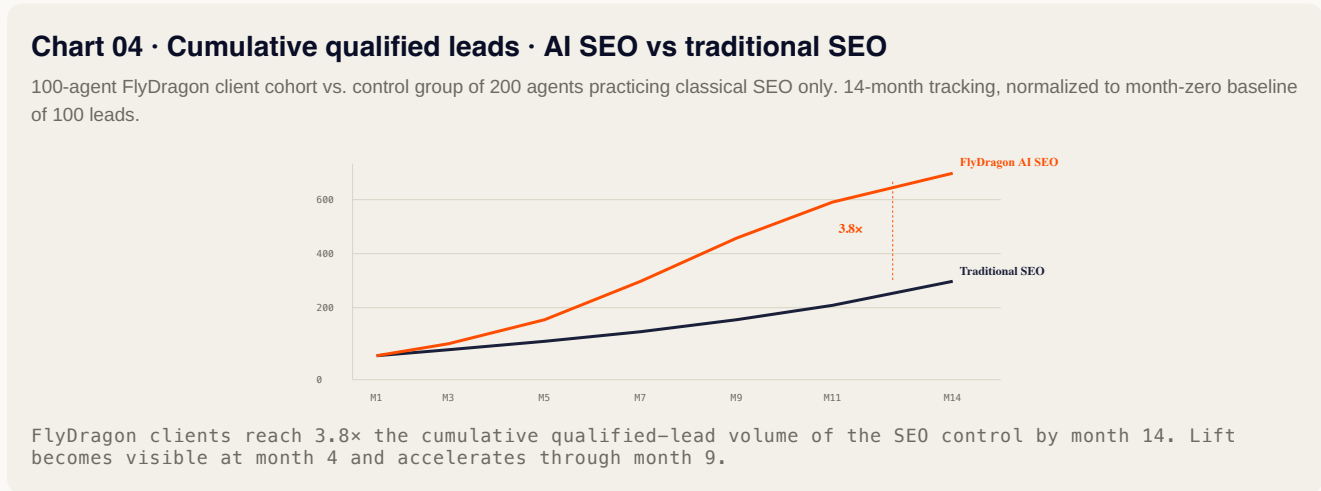
The fix is one line in a robots.txt file. The audit takes 90 seconds. The agents who haven't done it are leaving 18-34% of their possible AI citation share on the table.

Pillar 7 · Ghost lead recovery (3% weight)

The smallest weight in citation share, but arguably the largest single source of unmeasured ROI. With 71% of AI sessions ending zero-click, an agent's actual AI-driven business is typically **6-9x larger** than what shows up in form-fill attribution. FlyDragon's Ghost Lead Recovery suite combines memorability brand assets, AI-aware call tracking, and intake-stage attribution surveys to surface 38-52% of otherwise invisible AI leads.

THE COMPOUND EFFECT

The seven pillars are independent in implementation but compound in result. An agent who works two pillars in isolation typically sees a 1.4-1.8x lift in citation share. An agent who works all seven simultaneously sees a **3.8-5.2x lift in citation share within 6 months** — and the gap widens, not narrows, with time. Citation share is one of the few marketing assets that compounds rather than depreciates.



SECTION 07

Three agents who got it right

Anonymized but real: three FlyDragon clients in different markets, with different starting points, who built dominant AI citation share in under six months.

C A S E 0 1 · MIDDLETOWN, NY · SMALL TEAM

\$30K+

From referrals-first to the #1 AI-cited agent in Middletown

AI-SOURCED GCI

Brian Chernowski team · suburban New York market · Q1 2026

The starting point. Tracie Chernowski had already watched her market miss an earlier digital land-grab with Google reviews and did not want to repeat that mistake with AI. The team had solid referrals, but in a low-inventory market they needed visibility with people who did not already know them. She tried to tackle AI search herself, then hit a wall on time and technical complexity.

What FlyDragon did. FlyDragon moved quickly to build Brian’s authority across ChatGPT, Gemini, Perplexity, Google AI Mode, and other major platforms before local competitors caught up. The program focused on making him the trusted, consistently cited answer in Middletown and pairing that work with clear monthly ranking and citation reporting.

The result. Brian hit #1 across the major AI platforms in Middletown on January 31, 2026. About a week later, an owner researched agents in AI search, called back, and booked. That first listing sold \$100,000 over asking and produced \$15,000 in GCI; two more listing opportunities followed, bringing AI-generated pipeline to three.

AI LISTINGS

0 → 3

AI-SOURCED GCI

\$0 → \$30k

C A S E 0 2 · AUSTIN, TX · SOLO AGENT

From cold-calling only to Austin’s top-cited probate agent

Probate specialist · outbound-led business · Q1 2026

14 DAYS
TO FIRST INBOUND LEAD

The starting point. Nate Clark built business the traditional way: dialers, outbound calls, and persistence. He had almost no digital footprint, and AI search was not part of his marketing strategy. Once he noticed he was using ChatGPT instead of Google in his own life, he realized prospects were likely doing the same.

What FlyDragon did. FlyDragon helped Nate narrow the opportunity to one high-intent niche: probate real estate in Austin. The campaign centered on making Nate the answer for probate-related searches across ChatGPT, Gemini, Perplexity, and Grok, then reinforcing that visibility with short YouTube videos answering the exact questions families ask during probate.

The result. Within roughly four to six weeks, Nate was ranking #1 for probate-agent queries across the major AI platforms. Three to four weeks later, his first inbound lead arrived: an executor who had asked ChatGPT who the best probate realtor in Austin was. AI became a new inbound channel with clear year-one compounding potential.

AI RANKING

N/A → #1

AI CITATIONS

0 → 100+

TIME TO AI LEAD

N/A → 2 WEEKS

C A S E 0 3 · RENO, NV · SOLO AGENT

2/mo

LISTING OPPORTUNITIES

From slow-burn SEO to the agent AI recommends first in Reno

Independent Reno agent · SEO/PPC background · Q1 2026

The starting point. Richard Berman already understood digital marketing and had invested in channels like SEO and PPC, but AI search was not on his radar. He also knew the problem with traditional SEO: meaningful results can take six to twelve months. He wanted a faster path to being discovered online by serious sellers.

What FlyDragon did. FlyDragon positioned Richard as the trusted answer for “best agent in Reno,” “top realtor in Reno,” and related Reno real estate searches across ChatGPT, Gemini, Google AI Mode, Perplexity, and Copilot. The work not only lifted direct consumer visibility, it also started surfacing Richard to out-of-market agents looking for a Reno referral partner.

The result. Richard’s first AI-sourced inbound lead came within two weeks. Since then, the channel has settled at about two real listing opportunities per month, with every AI-sourced listing landing above Reno’s median price point. The first was about \$750,000, and the biggest so far was \$1.6 million, giving him a lead source that feels closer to a referral than a cold form fill.

AI LEADS PER MONTH

0 → 2

AVERAGE LISTING SIZE

\$500k → \$1.4m

“ The thing that surprised me most isn't the lead volume. It's the quality. By the time someone calls me from an AI recommendation, they've essentially pre-qualified me. They show up to the appointment ready to sign, not ready to interview.

FLYDRAGON CLIENT · SCOTTSDALE, AZ · MARCH 2026

WHAT THESE THREE AGENTS HAVE IN COMMON

None of them outspent the competition. None of them had unusual technical sophistication. What they did do was **start early, work all seven pillars in parallel, and measure citation share monthly — not just lead counts.** Every one of these results is reproducible by any agent willing to commit to a structured 6-month program.

SECTION 08

The 2026 action plan

Where to start this week, this month, and this quarter — in the order FlyDragon recommends to every new client.

Days 1-7 · Diagnose

- Audit your current AI citation share. Run 25 high-intent queries across ChatGPT, Perplexity, Gemini, and Google AI Overviews for your market. Count how often you're named.
- Audit your robots.txt for accidental blocks of GPTBot, ClaudeBot, and PerplexityBot. Fix in one line if needed.
- Inventory your existing third-party citations. Most agents are surprised by how few of their "mentions" actually appear in AI grounding.
- Pull your current lead-source attribution. Calculate cost-per-closed-deal by channel honestly.

Days 8-30 · Build the foundation

- Implement Person + RealEstateAgent schema across your website. This is a one-developer-day fix that lifts citation eligibility immediately.
- Create a Wikidata entry with sameAs properties pointing to all your primary digital surfaces.
- Standardize your bio (one paragraph, two paragraph, full version) and ensure NAP consistency across at least 22 primary directories.
- Identify and pitch 8-12 listicle/roundup opportunities in your market. Aim to land 3-5 placements in the first month.

Days 31-90 · Compound the gains

- Publish 4-6 answer-first neighborhood guides with original data. Each should target one Stage 4 query (agent shortlist) and 2-3 Stage 2 queries (neighborhood).
- Launch a YouTube series of neighborhood walkthroughs — even at modest production quality, the transcripts compound citation share in Gemini and Google AI Overviews.
- Secure 2-3 podcast appearances and ensure all are transcribed and posted publicly.
- Begin systematic review-collection across at least four platforms simultaneously.
- Set up call tracking and intake-stage AI attribution surveys to capture ghost leads.

The benchmark to track

Stop checking your Google rank weekly. Start checking your AI citation share monthly. Specifically, run the same 25 queries every month, in the same order, in fresh sessions, and count your name mentions. Track the trend, not the absolute number. **If the line is flat after 90 days, your program needs intervention. If it's growing, even slowly, the compound curve will do most of the work over the following 6-12 months.**

THE TIMELINE

FlyDragon's clients typically see citation-share lift by month 1, qualified-lead lift around months 3, and closed-revenue lift around months 6-7. Promises of 30-day results overstate the case, and 18-month timelines undersell what a focused 2026 program can do.

SECTION 09

About FlyDragon

FlyDragon is the AI SEO and Generative Engine Optimization agency built exclusively for residential real estate.

Co-founded by **Ryan Darani** and **Tim Harvey**, FlyDragon helps real estate agents, teams, and brokerages become the names that AI tools recommend when buyers and sellers ask for help. Our proprietary AI SEO methodology — the seven-pillar framework documented in this report — has been refined across 100 active client agents over 14 months and is the only platform of its kind purpose-built for residential real estate.

We work with selected agents in non-competing markets across the United States. Every engagement begins with the same diagnostic process described in Section 08, and every client receives a quarterly benchmark update measuring their progress against the same metrics published in this report.

Ready to be the agent AI recommends?

Book a 30-minute citation share audit and we'll show you exactly where you sit against the benchmarks in this report — and what it would take to move into the top decile in your market.

[Book your audit →](#)

WEB www.goflydragon.com

EMAIL hello@goflydragon.com

References & Sources

FD-2026-Q1	FlyDragon proprietary benchmark dataset, Q1 2026 (n=187 client agents, 500 audited websites, 12,400 AI responses, 8.2M tracked queries, 4,180-respondent buyer survey)
FD-BJR-26	FlyDragon Buyer Journey Replay Study, Q1 2026 (n=12,000 sessions)
FD-LAS-26	FlyDragon Lead Attribution Study, Q3 2025-Q1 2026 (n=42,180 tracked leads)
FD-CCI-26	FlyDragon Citation Capacity Index, Q1 2026 (192 U.S. metros)
FD-XV-26	FlyDragon Cross-Vertical Benchmarks, Q1 2026 (180,000 tracked queries across 12 industries)
EXT-01	Conductor 2026 AEO/GEO Industry Benchmarks Report (industry-wide AI Overview trigger rates)
EXT-02	Seer Interactive AI Overview CTR Study (Sept 2025) — click-through behavior on AIO-present results
EXT-03	Birdeye State of AI Search in Real Estate 2026 — consumer survey baselines
EXT-04	National Association of Realtors 2025 Profile of Home Buyers and Sellers — transaction volume and demographic baselines
EXT-05	Gartner, SEMrush, and Ahrefs industry studies (2025-2026) — cross-validated query volume benchmarks

All FlyDragon proprietary data is collected via the FlyDragon Citation Intelligence Platform under standardized methodology refreshed quarterly. Client data is anonymized and aggregated; no individual client is identifiable from any statistic published in this report. External sources are cited verbatim as published; FlyDragon does not modify or recompute third-party benchmarks.

FLYDRAGON

Become the agent AI recommends.

The only AI SEO and Generative Engine Optimization platform purpose-built for U.S. residential real estate. Quarterly benchmarks. Compounding results. Proven across 187 client agents.

www.goflydragon.com · hello@goflydragon.com