

# Do personas change what AI recommends?

## **A study on how user role context shifts brand visibility in large language models**

Authors: **Yannick Reihs, Niklas Buschner** – Radyant

Study period: **March 18 – May 10, 2026**

Total chats analyzed: **17,929**

Published: **June 2026**

## Abstract

This study examines whether and to what extent adding persona context to prompts, such as “I’m a CEO” or “I’m an Agency Founder”, changes which brands AI-tools recommend. Over 54 days, we tracked 64 prompts with eight buyer personas across five AI models in the funnel builder software category. Based on our observations, persona context appears to meaningfully influence brand visibility and share of voice, with absolute swings of up to +24 pp in visibility and +9 pp in share of voice for individual brands. The effect appears strongest in Google AI Overview and Gemini, but largely absent in ChatGPT. Brands with well-positioned, persona-specific content could gain systematically, while those without could lose ground precisely when the user signals intent through role context. The implication for B2B brands is meaningful: optimizing for AI visibility likely requires thinking not just about keywords and topics, but about whether AI models understand who your ICP is.

## 1. Why we did this study

Does an AI give different brand recommendations depending on who is asking?

When a Head of Performance Marketing asks ChatGPT for the best funnel builder, do they get the same answer as a CTO? Or an Agency Founder? And if the answers differ, do those differences reflect actual relevance, or arbitrary variation?

In March 2026, we at Radyant wondered whether persona-tagged prompts would produce meaningfully different brand recommendations. So we decided to run our own research on the subject.

We ran the study using [Peec AI](#), an AI search tracking platform, across the funnel builder software category, a well-populated B2B vertical with brands that vary significantly in their positioning and target audience specificity.

## 2. Research question

**Primary question:** Does adding buyer persona context to a prompt change which brands large language models recommend, and by how much?

Secondary questions:

- Are the effects consistent across different AI models, or do some models respond to persona context more than others?
- Does persona context change how many total brands are recommended (i.e., does it make responses more focused)?
- Does persona context affect how brands are described (sentiment), or only which brands appear?

## 3. Study setup & methodology

### 3.1 Platform

All tracking was conducted through Peec AI, which executes prompts across multiple AI models daily and measures brand visibility, share of voice, sentiment, average position, and citation data. Prompts are run once per day (three times on the first day) per model, with each result constituting one “chat”.

### 3.2 Study period

March 18 – May 10, 2026 – 54 days of continuous tracking.

### 3.3 Brands tracked

We tracked 10 brands in the funnel builder/interactive lead generation software category. These 10 brands were the 10 most mentioned brands, which is why we chose to track them:

Brand	Domain
Heyflow	heyflow.com
GoHighLevel	gohighlevel.com
Involve.me	involve.me
ClickFunnels	clickfunnels.com
Unbounce	unbounce.com
Systeme.io	systeme.io
Perspective	perspective.co
ConvertFlow	convertflow.com
Typeform	typeform.com
Jotform	jotform.com

Brand detection in Peec AI relies on rule-based string matching: each brand is identified via a case-insensitive tracked name plus optional aliases and RegEx patterns.

### 3.4 AI models tracked

Five individual AI models were tracked, plus an all-models aggregate:

#	AI Platform
1	ChatGPT 5 (search-enabled, via API)
2	Google AI Overview
3	Google Gemini
4	Microsoft Copilot
5	Perplexity

“All Models” figures in this paper represent the aggregate across these five platforms.

### 3.5 Prompt design

We tracked 8 base prompt themes representing real user journeys in the funnel builder evaluation process, constructed following our established prompt creation methodology for client AI visibility tracking: prompts are derived from real buyer pain points rather than short-tail keywords, structured to include a role and a specific problem, and validated through manual testing across all tracked AI platforms before setup (as documented in Radyant's guide on [prompt tracking for AI search visibility](#)). The final set covers TOFU and

MOFU intent stages to reflect the range of queries where persona context is realistically introduced by real users. Below are the exact prompts as tracked for the P0 Control Group (no persona prefix). For all other personas, the same prompts were used with the corresponding persona prefix added at the start.

#	Prompt (P0 – Control Group, exact wording)
1	Which platforms let me replace static lead forms with interactive funnels, CRM integration, and paid campaign performance tracking?
2	What are the best funnel builder tools for lead generation with strong design flexibility, built-in A/B testing, and CRM integration?
3	Which is the best no-code funnel builder?
4	I need a no-code tool to build interactive multi-step lead funnels and run A/B tests independently. What do you recommend?
5	Which funnel builder is best for high-value lead gen with strong qualification logic, conversion tracking, and no-code flexibility?
6	I want to move from a basic page builder to a funnel tool with conditional logic, pixel-accurate design, and A/B testing. What are the best options?
7	What's the best alternative to static HubSpot forms for interactive lead funnels that are GDPR-compliant with full CRM integration?
8	How can I improve conversion rates on landing pages and lead funnels that receive paid traffic from Meta and Google campaigns?

### 3.6 Personas

Each prompt was run with eight persona variants, producing 64 tracked prompts in total (8 themes × 8 personas):

Persona ID	Label	Prompt prefix
P0	Control Group	(no persona prefix — baseline)
P1	Head of Performance Marketing (HoPM)	"I'm Head of Performance Marketing."
P2	Agency Founder	"I'm the founder of an agency."
P3	Head of Growth	"I'm Head of Growth."
P4	CTO (Chief Technology Officer)	"I'm CTO at a company."
P5	CEO (Chief Executive Officer)	"I'm CEO of a company."
P6	CMO (Chief Marketing Officer)	"I'm CMO of a company."
P7	Marketing Manager	"I'm a Marketing Manager."

The eight buyer role personas were selected to span the typical decision-making unit in B2B SaaS purchases, from individual contributors to C-suite roles. P0 (no persona) serves as the baseline condition for all comparisons.

**Key design decision:** Persona prefixes contain only role information, with no industry, company size, or vertical context. This was intentional, to isolate the persona effect from company-context effects. Differences observed are attributable to role framing alone.

### 3.7 Total data volume

We tracked and analyzed a total of 17,929 unique AI responses. This substantially exceeds the minimum threshold for statistically reliable visibility measurement established by Graphite's research on LLM response sampling (Smith & Druck, 2025), which identified 10 responses per prompt as sufficient for a reliable estimate. Our study averages approximately 280 responses per prompt-persona combination, providing a robust baseline.

### 3.8 Metrics & delta format

For each brand per prompt per model per day:

- Visibility (%): Share of responses in which the brand was mentioned
- Share of Voice/SOV (%): Brand's relative share of all brand mentions
- Average Position: Where the brand appears in the response (1 = first mentioned)
- Sentiment Score (0–100): Tone of brand mentions in responses

All delta values in this paper are **absolute changes** measured against P0 (Control Group). A delta of +9 pp for share of voice means the brand's SOV increased by 9 percentage points (e.g., from 13 % to 22 %).

### 3.9 Statistical significance

For Visibility and Share of Voice, we applied two-proportion z-tests comparing each persona's observed brand mention rate against the P0 control group. Each cell rests on approximately 432 observations for individual models (54 days × 8 prompts) and approximately 2,160 for the All Models aggregate. Because 840 tests were run simultaneously (420 Visibility + 420 SOV), we applied Benjamini–Hochberg (BH) False Discovery Rate correction at  $q < 0.05$  across all tests combined. BH is preferred over Bonferroni here because the tests are correlated (the same brand appears across all 7 persona variants), making Bonferroni overly conservative. In the accompanying delta tables, cells marked ❖ are statistically significant at  $p < 0.05$  (uncorrected); the interactive table also allows BH correction to be toggled on.

Sentiment scores and average position are treated as descriptive metrics. Applying t-tests to these metrics would require within-cell standard deviations across individual daily observations, which are not available in the exported aggregate data. With  $n = 432$  per cell and any plausible SD estimate, even a 1-point sentiment delta produces a t-statistic above 10 ( $p \approx 0$ ), meaning virtually every non-zero delta would appear significant, a sample-size artifact rather than a meaningful signal. Changes of  $\geq 3$  points in sentiment or  $\geq 0.3$  in average position are considered directionally meaningful in this study.

### 3.10 Citation pool and source vocabulary

Across the eight persona variants, Peec AI retrieved and cited sources drawn from a combined pool of 201 unique URLs across 88 domains (derived from the top-100 citation lists across all eight personas, which total 800 slots before deduplication). The majority of sources appear across multiple persona pools simultaneously: 47 URLs appear in all eight, establishing the brand labels that hold regardless of persona context. A subset of 68 URLs

appears in only one persona's top-100. These persona-unique sources are where per-persona vocabulary shifts concentrate.

Appears in N persona top-100 lists	URL count
1 (persona-unique)	68
2	23
3	20
4	5
5	12
6	12
7	14
All 8 (universal)	47
<b>Total unique URLs</b>	<b>201</b>

Persona	Persona-unique URLs
P0 – Control	8
P1 – Head of Performance Marketing	10
P2 – Agency Founder	9
P3 – Head of Growth	10
P4 – CTO	8
P5 – CEO	5
P6 – CMO	11
P7 – Marketing Manager	7
<b>Total</b>	<b>68</b>

No source in the 201-URL dataset uses the exact persona job titles defined in this study. Terms like “Head of Performance Marketing”, “Agency Founder”, “Head of Growth”, or “Marketing Manager” do not appear as editorial audience labels in any retrieved source. The correlating vocabulary is functional equivalents: “performance marketers”, “agencies”, “paid acquisition teams” and “marketing teams”. Whatever mechanism drives the visibility shifts, it operates through semantic proximity rather than verbatim title matching.

## 4. Model overview: persona sensitivity by platform

The table below summarizes, for each AI model and across all models combined, the average visibility change per persona (across all 10 brands and all 7 non-control personas), the average absolute movement, and the single largest positive and negative brand-level shift. Delta values are in percentage points (pp).

Model	Avg Net $\Delta$ Vis	Avg Abs $\Delta$ Vis	Avg Abs $\Delta$ SOV	Biggest Gain in Vis (Brand   Persona)	Biggest Loss in Vis (Brand   Persona)
<b>ChatGPT</b>	+0.8 pp	2.2 pp	0.7 pp	GoHighLevel   Agency Founder: +14 pp	Heyflow   Agency Founder: -6 pp
<b>Google Gemini</b>	-1.1 pp	5.0 pp	2.0 pp	Unbounce   HoPM: +24 pp	ClickFunnels   Head of Growth: -16 pp
<b>Google AI Overview</b>	-7.6 pp	8.9 pp	2.9 pp	Perspective   HoM: +17 pp	ClickFunnels   CTO: -24 pp
<b>Microsoft Copilot</b>	-0.3 pp	2.2 pp	0.9 pp	GoHighLevel   Agency Founder: +7 pp	Systeme.io   Head of Growth: -8 pp
<b>Perplexity</b>	-0.6 pp	2.4 pp	1.1 pp	ConvertFlow   Marketing Manager: +7 pp	Unbounce   Head of Growth: -10 pp
<b>All Models</b>	-1.8 pp	2.6 pp	0.9 pp	GoHighLevel   Agency Founder: +10 pp	ClickFunnels   Head of Growth: -9 pp

Avg Net  $\Delta$  Vis: mean signed visibility change across all 10 brands and all 7 persona variants vs. P0 (positive = net gain, negative = net loss). Avg Abs  $\Delta$  Vis / SOV: mean of absolute (unsigned) changes, i.e. average movement regardless of direction. Biggest gain/loss: single largest brand-level delta observed across any brand-persona combination on that model.

Google AI Overview is the most persona-sensitive platform by a wide margin: its average per-persona shift in visibility is -7.6 pp, nearly seven times larger than the next most affected (Gemini at -1.1 pp). The extremes also stand out: AI Overview hosts the single largest negative swing in the study (across all platforms) (ClickFunnels, CTO persona: -24 pp), while Gemini produces the largest positive swing (Unbounce, HoPM persona: +24 pp). ChatGPT remains the most stable, with an average change of just +0.8 pp and no individual brand exceeding  $\pm 15$  pp. Copilot and Perplexity sit in the middle range, with average per-persona shifts close to zero and moderate individual outliers driven primarily by GoHighLevel (Agency Founder) and Unbounce/ConvertFlow, respectively. Across all models combined, the average net shift in visibility is -1.8 pp with a 2.6 pp absolute visibility movement and 0.9 pp absolute SOV movement.

### 4b. Average delta per persona by model

The tables below show, for each model (and all models combined), the per-persona average visibility delta vs. P0. Net  $\Delta$  Vis is the signed average across all 10 brands (gains and losses cancel). Abs  $\Delta$  Vis and Abs  $\Delta$  SOV are the mean absolute movements regardless of direction.

#### ChatGPT

Persona	Net $\Delta$ Vis	Abs $\Delta$ Vis	Abs $\Delta$ SOV
HoPM	+0.1 pp	1.5 pp	0.5 pp
Agency Founder	+0.6 pp	3.5 pp	1.2 pp
Head of Growth	-0.3 pp	1.0 pp	0.4 pp
CTO	+0.7 pp	2.1 pp	0.7 pp
CEO	+0.9 pp	1.9 pp	0.6 pp
CMO	+2.2 pp	2.8 pp	0.7 pp
Marketing Manager	+1.5 pp	2.6 pp	0.7 pp

### Google Gemini

Persona	Net $\Delta$ Vis	Abs $\Delta$ Vis	Abs $\Delta$ SOV
HoPM	-0.5 pp	6.7 pp	2.6 pp
Agency Founder	-1.4 pp	4.6 pp	2.0 pp
Head of Growth	-1.2 pp	6.9 pp	2.8 pp
CTO	-2.4 pp	5.7 pp	2.3 pp
CEO	-1.7 pp	2.4 pp	1.0 pp
CMO	-0.9 pp	5.0 pp	2.0 pp
Marketing Manager	+0.5 pp	3.5 pp	1.3 pp

### Google AI Overview

Persona	Net $\Delta$ Vis	Abs $\Delta$ Vis	Abs $\Delta$ SOV
HoPM	-5.9 pp	8.9 pp	3.4 pp
Agency Founder	-7.0 pp	8.8 pp	3.3 pp
Head of Growth	-9.5 pp	10.4 pp	2.9 pp
CTO	-11.4 pp	11.4 pp	3.3 pp
CEO	-6.0 pp	6.7 pp	2.7 pp
CMO	-6.7 pp	7.9 pp	2.8 pp
Marketing Manager	-7.2 pp	8.5 pp	2.2 pp

## Microsoft Copilot

Persona	Net $\Delta$ Vis	Abs $\Delta$ Vis	Abs $\Delta$ SOV
HoPM	+0.3 pp	1.9 pp	0.8 pp
Agency Founder	+0.1 pp	2.6 pp	1.1 pp
Head of Growth	-0.5 pp	2.9 pp	1.1 pp
CTO	-0.9 pp	2.4 pp	1.1 pp
CEO	-1.6 pp	2.7 pp	0.9 pp
CMO	-0.0 pp	1.4 pp	0.6 pp
Marketing Manager	+0.7 pp	1.8 pp	0.7 pp

## Perplexity

Persona	Net $\Delta$ Vis	Abs $\Delta$ Vis	Abs $\Delta$ SOV
HoPM	-2.4 pp	3.5 pp	1.4 pp
Agency Founder	+0.1 pp	2.9 pp	1.5 pp
Head of Growth	-1.1 pp	3.0 pp	1.5 pp
CTO	-0.6 pp	1.6 pp	0.7 pp
CEO	-0.5 pp	1.6 pp	0.7 pp
CMO	+0.2 pp	1.2 pp	0.5 pp
Marketing Manager	+0.1 pp	2.8 pp	1.4 pp

## All models

Persona	Net $\Delta$ Vis	Abs $\Delta$ Vis	Abs $\Delta$ SOV
HoPM	-1.7 pp	3.3 pp	1.3 pp
Agency Founder	-1.5 pp	3.2 pp	1.1 pp
Head of Growth	-2.5 pp	3.0 pp	1.0 pp
CTO	-2.9 pp	3.1 pp	0.9 pp
CEO	-1.8 pp	1.9 pp	0.7 pp
CMO	-1.0 pp	1.8 pp	0.7 pp

<b>Marketing Manager</b>	-0.9 pp	1.7 pp	0.7 pp
--------------------------	---------	--------	--------

Several patterns stand out. First, AI Overview is the only platform where every persona produces a notably negative average delta, ranging from -5.9 pp (HoPM) to -11.4 pp (CTO). This suggests a behavior of narrowing brand recommendations aggressively when role context is present. The CTO persona triggers the strongest overall compression on AI Overview, suggesting that technical buyer framing prompts the most selective filtering.

Second, ChatGPT is the only model where the average delta is positive for most personas, reflecting its different approach to persona context. On ChatGPT, persona context slightly broadens brand coverage. CMO and Marketing Manager prompts average +2.2 pp and +1.5 pp respectively, making ChatGPT the only model where most personas push brand mentions upward. This is consistent with Finding 3: ChatGPT appears largely persona-indifferent at the aggregate level, redistributing share rather than shrinking the field.

Third, Gemini and Perplexity show consistent but moderate negative deltas across most personas, with Gemini slightly more sensitive (peaking at -2.4 pp for CTO) and Perplexity being the most stable outside of HoPM (-2.4 pp). Copilot is broadly neutral, rarely moving more than -1.6 pp in either direction. Across all models combined, the CTO and Head of Growth personas produce the largest average absolute movements (3.1 pp and 3.0 pp), while CEO (1.9 pp), CMO (1.8 pp) and Marketing Manager (1.7 pp) are the most benign.

## 5. Key findings

### Finding 1: Personas appear to change what AI recommends, and the observed effect is substantial

Adding a persona to a prompt produces measurable and often substantial shifts in brand visibility and share of voice. Across all models and all brands, we observed absolute changes of up to +24 pp in visibility (Unbounce in Gemini, Head of Performance Marketing persona) and +9 pp in share of voice (GoHighLevel in AI Overview, Agency Founder persona) compared to the control group. Several of the largest headline deltas are statistically significant at  $p < 0.05$  (two-proportion z-test vs. P0); all survive BH/FDR correction across 840 simultaneous tests. See Section 3.9 and the accompanying delta tables for full significance detail.

Between the best-performing and worst-performing persona for a single brand, spreads of over 20 pp in visibility are observed in some model-brand combinations. Selected headline deltas vs. P0:

Brand	Model	Persona	Metric	P0	Pers ona	Δ
GoHighLevel	AI Overview	Agency Founder	SOV	13 %	22 %	<b>+9 pp</b>
GoHighLevel	AI Overview	Agency Founder	Visibility	45 %	59 %	<b>+14 pp</b>
GoHighLevel	ChatGPT	Agency Founder	SOV	14 %	20 %	<b>+6 pp</b>

Brand	Model	Persona	Metric	P0	Pers ona	Δ
Perspective	AI Overview	Head of Perf. Mktg.	Visibility	46 %	63 %	+17 pp
Unbounce	Gemini	Head of Perf. Mktg.	Visibility	17 %	41 %	+24 pp
Heyflow	AI Overview	CTO	SOV	35 %	41 %	+6 pp
ClickFunnels	AI Overview	Head of Perf. Mktg.	Visibility	41 %	17 %	-24 pp
Systeme.io	AI Overview	Head of Perf. Mktg.	Visibility	21 %	7 %	-14 pp
Systeme.io	Gemini	Head of Perf. Mktg.	Visibility	13 %	1 %	-12 pp

## Finding 2: Overall brand mentions decline with personas as LLMs get more focused

When no persona is given (P0), the average visibility across all 10 tracked brands is 30.0 % (all-models aggregate). With any persona, average visibility drops to between 28.3 % and 29.3 %.

Persona	Avg Visibility	Δ vs P0	Avg Position (all brands)	Δ vs P0
P0 – Control	30.0 %	—	3.14	—
P1 – HoPM	29.2 %	-0.8 pp	3.09	-0.06
P2 – Agency Founder	29.3 %	-0.7 pp	3.12	-0.03
P3 – Head of Growth	28.7 %	-1.3 pp	3.13	-0.01
P4 – CTO	28.5 %	-1.5 pp	3.19	+0.04
P5 – CEO	28.3 %	-1.7 pp	2.97	-0.18
P6 – CMO	28.5 %	-1.5 pp	3.05	-0.10
P7 – Marketing Manager	29.3 %	-0.7 pp	3.04	-0.10

### Supporting observation: Aggregate average position as a signal for response density

A lower aggregate average position across all brands means those mentioned tend to appear earlier in responses, consistent with shorter, more focused lists. In 6 of 7 personas, lower visibility corresponds to a lower aggregate position. The CEO is the clearest example: lowest visibility (-1.7 pp) and lowest aggregate position (-0.18). The CTO is the one exception (+0.04 position despite -1.5 pp visibility).

### Persona type grouping: C-Suite vs. Marketing practitioners

Persona group	Avg Visibility	Δ vs P0	Avg Position	Δ vs P0
P0 – Control	30.0 %	—	3.14	—

Persona group	Avg Visibility	Δ vs P0	Avg Position	Δ vs P0
Marketing practitioners (HoPM, Growth, Marketing Manager)	29.1 %	-0.9 pp	3.09	-0.06
Agency Founder	29.3 %	-0.7 pp	3.12	-0.03
<b>C-Suite (CTO, CEO, CMO)</b>	28.4 %	<b>-1.6 pp</b>	3.07	<b>-0.08</b>

C-Suite personas seem to produce more focused responses than marketing practitioners (-1.6 pp vs. -0.9 pp total visibility).

### Finding 3: ChatGPT appears largely unresponsive to persona context, other models are more responsive

ChatGPT shows the smallest persona-driven shifts of any model tested. Across all brands and personas, ChatGPT SOV deltas rarely exceed  $\pm 2$  pp in absolute terms. Google AI Overview and Gemini exhibit the largest persona-driven swings, with individual brand visibility shifting by 10–24 pp compared to the control. Copilot and Perplexity fall between these extremes with moderate persona sensitivity.

### Finding 4: AI Overview is the most persona-sensitive interface and the most strategically important

Google AI Overview shows the largest absolute swings of any platform tested, with some brand-persona combinations seeing visibility changes of more than 20 pp. Perspective is the consistent positive outlier, gaining visibility across every single persona (+5 pp to +17 pp). ClickFunnels shows the opposite: 14–24 pp visibility losses across all personas.

The strategic weight of this finding is amplified by AI Overview's scale: with 2.5 billion monthly users, it is the largest AI search surface in the world, larger than ChatGPT, Claude, and Perplexity combined. According to research from Peec AI, AI Overviews appear in 87 % of Google searches, including 88.5 % of bottom-of-funnel queries (Rudzki, 2026).

### Finding 5: Persona-driven volatility concentrates in high-visibility brands while low-visibility brands remain largely unaffected

Grouping the 10 tracked brands into a top 3 (highest P0 visibility: Heyflow 65 %, ClickFunnels 46 %, GoHighLevel 45 %) and bottom 3 (lowest P0 visibility: Jotform 7 %, ConvertFlow 14 %, Typeform 16 %) reveals a structural asymmetry.

**Top 3 brands experience twice the persona-driven volatility of the bottom 3.** Average absolute visibility change per persona: 3.1 pp for the top tier vs. 1.6 pp for the bottom tier. The top 3 brands generate 36 % of all persona-driven visibility change across the 10 brands, while representing only 30 % of the set. This comes even though the most visible brand (Heyflow) only shows little volatility. The bottom 3 brands are largely persona-indifferent at the all models level.

Brand	P0 Vis (all models)	Max Vis Δ (all models)	Max SOV Δ (all models)
Heyflow	65 %	$\pm 1$ pp	$\pm 2$ pp

Brand	P0 Vis (all models)	Max Vis Δ (all models)	Max SOV Δ (all models)
ClickFunnels	46 %	-9 pp	-2 pp
GoHighLevel	45 %	+10 pp	+5 pp
—	—	—	—
Typeform	16 %	±2 pp	±1 pp
ConvertFlow	14 %	±2 pp	±1 pp
Jotform	7 %	-3 pp	-1 pp

### Finding 6: AI Overview is 2–9x more persona-sensitive than ChatGPT, depending on the brand

The sensitivity gap between AI Overview and ChatGPT is not uniform across brands. For some brands, it is modest (GoHighLevel: 2.1×, Unbounce: 2.1×), while for others it is considerably larger: Systeme.io moves 9.0× more on AI Overview than on ChatGPT (10.8 pp vs. 1.2 pp average absolute delta), Heyflow 8.0× (12.7 pp vs. 1.6 pp), and ClickFunnels 6.3× (21.2 pp vs. 3.4 pp). The brands with the highest AIO-to-ChatGPT ratio tend to be those with the weakest or most generic positioning, suggesting that AI Overview applies stronger persona-based filtering where there is the least role-specific content signal to work with.

Brand	ChatGPT abs Δ	Gemini abs Δ	AIO abs Δ	Copilot abs Δ	AIO / ChatGPT
Heyflow	1.6 pp	5.6 pp	12.7 pp	2.4 pp	8.0×
GoHighLevel	4.1 pp	4.5 pp	8.6 pp	2.7 pp	2.1×
Involve.me	2.3 pp	4.9 pp	12.9 pp	2.5 pp	5.6×
ClickFunnels	3.4 pp	11.2 pp	21.2 pp	1.8 pp	6.3×
Unbounce	3.3 pp	8.3 pp	6.8 pp	2.5 pp	2.1×
Systeme.io	1.2 pp	5.1 pp	10.8 pp	2.1 pp	9.0×
Perspective	1.3 pp	1.8 pp	4.5 pp	1.5 pp	3.5×
ConvertFlow	1.3 pp	2.1 pp	3.0 pp	1.7 pp	2.4×
Typeform	1.0 pp	1.4 pp	2.3 pp	1.2 pp	2.3×
Jotform	2.7 pp	3.3 pp	6.6 pp	1.3 pp	2.5×

Avg abs Δ = mean absolute visibility change (in pp) across all 7 personas vs. P0. AIO / ChatGPT ratio = how many times more sensitive AI Overview is than ChatGPT for that brand. Gemini and Copilot abs Δ included for reference.

### Finding 7: More than 1 in 4 brand-platform-persona combinations show a visibility shift of 5 pp or more

Across the 350 brand-persona-model combinations tracked (10 brands × 7 personas × 5 models), 97 show an absolute visibility change of 5 pp or more vs. P0. That is 28 % of all combinations. The distribution is heavily skewed toward AI Overview and Gemini: AI Overview accounts for the majority of the largest individual swings, while ChatGPT and Copilot together contribute fewer than 15 % of the ≥5 pp moves. This suggests that more than one in four persona-platform combinations represents a material visibility shift rather than a marginal one.

### **Finding 8: Visibility shifts correlate with persona-relevant vocabulary in the citation pool**

For most large visibility deltas in this study, there is a corresponding source-level correlate. A URL that uses the relevant persona's functional vocabulary and either names the brand positively or omits it. Brands whose source ecosystems are labelled for professional buyer roles gain visibility when those personas are active. Notably, no source in the dataset uses the exact job titles from this study's persona definitions; the mechanism appears to operate through functional vocabulary equivalents rather than verbatim role-title matching.

## **6. Strategic implications for B2B brands**

### **6.1 Persona-specific positioning could improve visibility in AI Search**

Brands with ICP-specific content, framing their product for specific roles, use cases, and pain points, could see systematic gains when users provide role context. Brands with generic positioning could lose ground. The important question: when your ideal customer tells an AI who they are, does the AI know your product is for them?

### **6.2 Publish in the vocabulary your ICP uses, not just for the topics they search**

The source analysis suggests that brands do not necessarily need to target exact ICP job titles in their content to benefit from persona-driven visibility effects. What matters is publishing in the vocabulary those roles use about themselves, like “performance marketers”, “paid acquisition”, “ROAS” and “CPL” for HoPM contexts or “white-label”, “client management” and “agency” for agency founder contexts. When that vocabulary appears in high-retrieval sources alongside a positive brand mention, visibility rises for that persona.

### **6.3 AI Overview and Gemini deserve dedicated persona attention**

Google-powered AI Overview and Gemini show the highest persona sensitivity. Brands optimizing primarily for ChatGPT may be missing the highest-leverage opportunity. Google's AI surfaces combine the reach of search intent with LLM recommendation depth, and they are the most responsive to persona signals.

### **6.4 The control group is not the right benchmark for AEO**

Most AEO measurements track prompts without persona context. Persona-tagged prompts can show absolute differences of 10–24 percentage points compared to control group prompts for the same brand and model. Therefore, brands should measure their visibility with role context (covering the most important roles of their ICP) to make sure they show up when it matters.

## **7. Limitations**

## 7.1 Single category and market

All data is from the funnel builder / interactive lead generation category, tracked with German-market IP addresses. Persona sensitivity effects may differ in other categories and geographies.

## 7.2 Prompt design is not fully representative of real user behavior

Real users rarely open an AI chat by stating their job title. The persona prefix is a methodological construct to isolate the persona variable. In practice, LLMs may infer persona context from conversation history, connected accounts, or implicit signals.

## 7.3 No company size or vertical context in personas

Persona definitions deliberately excluded company type, industry, and size to isolate the role effect.

## 7.4 LLM stochasticity

Despite 280+ responses per prompt-persona combination, LLMs generate probabilistic outputs. We consider differences above 5 percentage points in absolute terms to be robust; smaller differences should be treated as directional.

## 7.4 Snapshot in time

The study ran from March 18 to May 10, 2026. LLM training data, ranking behavior, and brand content are all dynamic.

## 7.5 Prompt-level autocorrelation

The 432 observations per model-brand-persona cell are produced by 8 prompt themes repeated daily. These are not fully independent: the same prompt querying the same brand on consecutive days generates correlated outputs. Standard two-proportion z-tests assume independence. A more robust approach would model prompts as a random effect, using a mixed-effects framework. The significance values reported here should be read as approximate estimates rather than exact p-values. The direction and magnitude of the largest effects, those surviving BH correction with p well below  $10^{-10}$ , are not materially affected by this limitation, but smaller significant cells should be treated with additional caution.

# 8. Related research & analysis

## 8.1 Smith, E. & Druck, G. (2025). Demystifying Randomness in AI

Graphite Research. Established the statistical framework for LLM visibility measurement, identifying 10 responses per prompt as the minimum for reliable estimates and introducing sequential sampling for higher precision. Our study exceeds this threshold by 28× per prompt-persona combination.

<https://graphite.io/five-percent/demystifying-randomness-in-ai>

## 8.2 Rudzki, Tomek (2026). AI Overviews is the most undertracked AI search: 500,000 prompts show why

Peec AI Blog. Analysis of 500,000 prompts shows that AI Overviews appears in 87 % of Google searches, including 88.5 % of bottom-of-funnel queries, and reaches 2.5 billion monthly users, making it the largest AI search surface in the world. Referenced in Finding 4

to contextualize the strategic importance of AI Overview as the most persona-sensitive platform in our study.

<https://peec.ai/blog/ai-overviews-is-the-most-undertracked-ai-search-500-000-prompts-show-why>

## **9. Data access**

The full underlying data, including all 24 metric tables (Visibility, Share of Voice, Sentiment, and Average Position, each across 5 models plus an all-models aggregate), all delta calculations (absolute change vs. P0), are available in the accompanying files.

Methodology questions and data inquiries can be directed to [yannick.reihs@radyant.io](mailto:yannick.reihs@radyant.io).

## **10. About Radyant**

Radyant is an organic growth agency specializing in AI search optimization (AEO) and SEO. We work with brands including Allianz Suisse, Heyflow, osapiens and Stackfield, helping them build visibility in both traditional search and LLMs.

Our research is motivated by practical client work. The questions in this study are questions we ask every day for the brands we represent.

[radyant.io](https://radyant.io)