



AI Search Absorption: A Multi-Component Index for Measuring Organic Search Displacement

Quantifying how Google's AI Overviews and SERP features reduce click-through opportunity across 14,728 queries

A reproducible methodology decomposing search displacement into visibility, click, and intent mechanisms, applied to a curated corpus of 14,728 scored queries with full SERP feature detection and LLM-based intent assessment.

Abstract

Search engine results pages (SERPs) increasingly contain answer features that resolve user intent without referral traffic to third-party publishers. We introduce the AI Search Absorption Index, a reproducible, multi-component measure of how much organic search opportunity is displaced by SERP features including AI Overviews, featured snippets, People Also Ask, knowledge panels, shopping results, and inline media.

The index decomposes absorption into three sub-indices: visibility absorption (pixel-space displacement of organic results below the viewport fold), click absorption (estimated reduction in organic click-through rate from published CTR benchmarks), and intent absorption (LLM-estimated degree to which the SERP satisfies user intent without a click).

SERP data were collected for 14,728 queries spanning four intent categories (informational, commercial, transactional, navigational) and two devices (mobile, desktop). For queries with AI Overviews, a locally deployed Qwen 3.6-35B-A3B model assessed intent satisfaction and estimated residual organic CTR via a three-run ensemble. Each query receives a composite absorption score on a 0–100 scale.

The mean composite absorption across the dataset is 67.72 (95% CI [67.35, 68.09]). The dominant driver is visibility absorption (mean 80.02), followed by click absorption (70.03) and intent absorption (58.84). Informational queries are the most absorbed category (mean 76.43), while navigational queries are the least absorbed (mean 48.48). Mobile SERPs show consistently higher absorption than desktop (mean 69.07 vs. 66.37, $p = 4.77 \times 10^{-13}$).

The full dataset, code, and interactive browser are published alongside this paper to enable replication and extension.

Authors: Wayfinder Research

Data availability: `absorption_index_dataset.csv`, `absorption_index_dataset.json` (14,728 rows)

1. Introduction

1.1 The absorption problem

Search engine results pages (SERPs) increasingly contain Google-owned answer features that resolve user intent without referral traffic to third-party publishers. AI Overviews, launched at scale in 2024, synthesize multi-source answers directly in the SERP. Complementary features — featured snippets, People Also Ask (PAA), knowledge panels, shopping carousels, local packs, and inline video — further displace organic listings.

Existing research has documented aggregate effects of this shift. SparkToro and Datos (2024) reported that 58.5% of US Google searches ended without a click to any external website. Seer Interactive (2025) found that organic CTR on queries with AI Overviews fell from 1.76% to 0.61% — a 61% decline — in a longitudinal study of 3,119 queries across 42 organisations. Ahrefs (2025) analysed 300,000 keywords and reported that position-one CTR fell by approximately 58% when AI Overviews appeared.

These studies are valuable but aggregate. They do not identify which query types, devices, or intent categories face the highest displacement, nor do they decompose displacement into distinct mechanisms. The present study addresses this gap.

1.2 Research questions

1. Can absorption be decomposed into distinct, measurable mechanisms? Specifically, can visibility displacement, click-rate reduction, and intent satisfaction be separated and recombined into a stable composite?
2. Does absorption vary systematically by query intent? We compare informational, commercial, transactional, and navigational queries.
3. Does absorption vary by device? We test whether mobile and desktop SERPs produce meaningfully different scores.
4. How much of observed absorption is attributable specifically to AI Overviews versus other SERP features? We compute an AI-only composite that isolates AI Overview impact.

1.3 Contributions

- A reproducible methodology for measuring three distinct absorption mechanisms.
- A dataset of 14,728 scored queries with SERP features, LLM judgments, and sub-index values.
- Empirical evidence that visibility displacement is the dominant driver of absorption across all categories.
- An open-source pipeline and interactive data browser.

1.4 Related Work

1.4.1 Zero-click search

The term “zero-click search” describes queries for which the user obtains sufficient information from the SERP itself and does not click any result. SparkToro and Datos (2024) published the largest publicly available clickstream study, analysing 332 million search queries across 130,000 US devices over 21 months. Their 2024 study reported that 58.5% of US Google searches ended without a click to any external website, and that 28.5% of all clicks that did

occur went to Google-owned properties. On mobile devices, the zero-click rate was higher, at approximately 77%.

These figures are not unique to the post-AI-Overview era. SimilarWeb data cited in Brazilian competition proceedings reported a 64.82% zero-click rate in 2020, with 77.2% on mobile devices (Núcleo, 2023). The directional trend — increasing zero-click share, particularly on mobile — has been consistent across multiple data sources and time periods.

The academic literature on non-click behaviour predates the industry coinage. Li et al. (2017) introduced the concept of “good abandonment” — situations in which a user’s information need is satisfied without any click — and demonstrated its prevalence in both desktop and mobile search logs. Ye et al. (2022) reviewed the zero-click literature and noted that commercial search engines have systematically improved snippet and answer quality to reduce user effort, making good abandonment an increasingly common outcome. The present study extends this line of work by quantifying the degree to which modern SERP features, including AI-generated answers, satisfy intent without referral traffic.

1.4.2 SERP features and organic CTR

The impact of individual SERP features on organic click-through rates has been studied at scale by Fubel et al. (2023), who analysed 760,000 keywords across 40 US e-commerce domains, generating over 6 million clicks from 24 million impressions. Their ablation study modelled the incremental predictive power of 24 distinct SERP features on organic CTR. They found that SERP features exert, on average, a negative influence on organic CTR, but that the effect is modulated by result position, feature type, and whether the website is included within the feature itself. Position 1 results were particularly susceptible to CTR decline from increased SERP feature density.

Industry studies have produced comparable findings. Sistrix (Germany, February 2026) reported that position 1 organic CTR fell from 27% to 11% when an AI Overview was present. First-PageSage (2025) benchmarked position 1 organic CTR at 39.8% on a clean SERP, providing a baseline against which feature-induced depression can be measured.

The present study differs from prior work in two respects. First, it decomposes absorption into three distinct mechanisms — visibility displacement, click-rate reduction, and intent satisfaction — rather than treating SERP feature impact as a monolithic CTR effect. Second, it isolates the contribution of AI Overviews specifically, enabling comparison between AI-generated and traditional feature-driven displacement.

1.4.3 AI-generated search answers

The literature on AI-generated search answers is nascent and largely industry-based. Seer Interactive (2025) conducted the largest longitudinal study, tracking 3,119 informational and educational queries across 42 organisations from June 2024 through February 2026. They reported that organic CTR on queries with AI Overviews fell from 1.76% to 0.61%, while paid CTR on the same queries fell from 19.70% to 6.34%. The time-series design demonstrates that the decline is sustained rather than a one-time adjustment.

Ahrefs (2025) analysed 300,000 keywords using aggregated Google Search Console data and reported that position 1 CTR fell by approximately 58% when AI Overviews appeared. Their study also documented that average position 1 CTR for informational keywords declined from 7.6% in December 2023 to 3.9% in December 2025.

Pew Research Center (2025) provided behavioural evidence from a panel study of 900 US adults. Users encountering an AI summary clicked a traditional organic result in 8% of visits, versus 15% when no summary was present. Only 1% clicked a link inside the AI summary itself, and 26% ended their browsing session entirely after the AI summary appeared (versus 16% without one).

BrightEdge (2025) reported AI Overview appearance rates of 11–18% of tracked queries, growing 22% year-over-year. Their data indicates that informational and non-branded queries are disproportionately affected, while branded searches remain largely protected.

The European Parliament (2026) cited the Seer, Ahrefs, and Pew studies alongside academic publisher data showing approximately 13% traffic decline following AI Overview expansion in August 2024. The Reuters Institute (2026) surveyed 280 media leaders in 51 countries and found that publishers expect search-engine traffic to fall by 43% in coming years.

The present study contributes to this literature by providing a reproducible, multi-component index that decomposes absorption into distinct mechanisms, applied to a curated corpus of 14,728 queries with full SERP feature detection and LLM-based intent assessment.

1.4.4 Keyword research methodology

Traditional keyword research tools report search volume and difficulty scores without accounting for SERP feature density or AI answer presence. Ahrefs' traffic estimates, for example, are based on aggregate position-CTR curves that assume a relatively clean SERP (Ahrefs, 2025). The present study proposes a methodological extension — adding an absorption score to keyword evaluation — that addresses this gap.

2. Methodology

2.1 Keyword corpus

A curated keyword catalog was built spanning four intent categories:

- **Informational:** queries seeking explanation or instruction.
- **Commercial:** queries researching or comparing products and services.
- **Transactional:** queries with purchase or action intent.
- **Navigational:** queries seeking a specific brand or site.

Keywords were sampled using stratified volume-tier sampling:

Tier	Volume percentile	Target sample	Tier weight
Head	Top 5%	50	2.0
Upper	5–20%	150	1.5
Mid	20–50%	300	1.0
Long	50–100%	500	0.5

All keywords were validated for US English search volume (location_code 2840, language en) via the DataForSEO API. The final corpus contains 14,728 unique keywords, distributed across the four categories and two devices (mobile and desktop).

2.2 SERP collection

For each keyword the Google SERP was collected using a hybrid data strategy:

- **Primary source:** DataForSEO Standard SERP API for bulk collection.
- **Validation / pixel source:** SerpAPI for a subsample, providing rectangle metadata for feature height estimation.

Collection parameters were fixed before analysis: - Location: United States - Language: English - Google domain: google.com - Device: mobile or desktop per run

Raw SERP JSON was stored to enable reproducible feature detection and audit.

2.3 Feature detection

For each SERP the following features were detected:

Feature	Attribute	Notes
AI Overview	aio_present	Synthesized answer block
Featured snippet	snippet_present	Direct answer box
People Also Ask	paa_count	Count of expanded items
Knowledge panel	knowledge_panel_present	Entity panel
Shopping results	shopping_present	Product carousel
Local pack	local_pack_present	Map + local listings
Inline videos	inline_videos_count	Count
Inline images	inline_images_count	Count
Short videos	short_videos_count	Count

Feature detection was performed by parsing the structured SERP JSON; no visual inference was used.

2.4 Visibility absorption

Visibility absorption measures how far organic results are pushed below the viewport fold by SERP features. For each keyword:

$$\text{visibility_absorption} = \frac{\sum(\text{feature_pixel_height} \times \text{click_impact_weight})}{\text{viewport_fold_height}} \times 100$$

Feature weights were calibrated by expected click-stealing impact:

Feature	Weight
AI Overview	0.90
Featured snippet	0.60
Inline videos	0.50
People Also Ask	0.40
Short videos	0.40
Inline images	0.30
Knowledge panel	0.30
Local pack	0.25
Shopping	0.20

Pixel heights were taken from SerpAPI rectangle data when available; otherwise default estimates were applied (AI Overview: 400 px; featured snippet: 200 px; PAA: 150 px per item; inline video: 180 px; inline image: 120 px; knowledge panel: 300 px; local pack: 250 px; shopping: 160 px; short video: 140 px). The viewport fold height was set to 600 px, reflecting a mobile-first design. Scores were capped at 100.

2.5 Click absorption

Click absorption estimates the reduction in organic CTR attributable to detected features. Published CTR benchmarks and a cumulative stacking model were used:

```
penalty_feature = 1 - (ctr_with_feature / baseline_ctr)
raw_penalty = Σ penalties for present features
click_absorption = raw_penalty × 100 (with diminishing returns above 80%)
```

Benchmarks:

Scenario	CTR	Source
No features (position 1 baseline)	2.6%	FirstPageSage (2025)
With AI Overview	1.0%	Seer Interactive (2025)
With featured snippet	1.9%	Estimated from Sistrix (2026)
With PAA only	2.2%	Estimated from Fubel et al. (2023)
With inline videos	1.8%	Estimated
With inline images	2.1%	Estimated
With short videos	1.9%	Estimated

Diminishing returns were applied via a logistic transform: $\text{click_absorption} = 100 / (1 + \exp(-0.08 \times (\text{raw_penalty} - 80)))$ for $\text{raw_penalty} > 60$, preventing saturation at 100 when many features co-occur.

2.6 Intent absorption

Intent absorption captures whether the SERP satisfies the user's information need without requiring a click. For queries with an AI Overview, a locally deployed Qwen 3.6-35B-A3B model (temperature 0.7, top-p 0.8, 262k context) made two independent assessments:

1. **Intent satisfaction:** whether the AI Overview fully, partially, or does not satisfy the query (Yes / Partially / No).
2. **Estimated residual CTR:** the percentage of searchers still expected to click an organic result.

The model was prompted with the query, AI Overview text, and a summary of other present SERP features. Each assessment was run three times; intent satisfaction was resolved by majority vote and CTR by median. The three runs enabled variance estimation and outlier detection.

Intent absorption was computed as:

```
satisfaction_score = {Yes: 85, Partially: 50, No: 15}
ctr_component = 100 - estimated_ctr_percent
intent_absorption = 0.60 × satisfaction_score + 0.40 × ctr_component
```

Scores were capped at 100.

2.7 Composite score

Sub-indices are combined into a composite absorption score using category-specific weights:

Category	Click	Intent	Visibility
Default	0.45	0.30	0.25
Informational	0.35	0.45	0.20

The informational override reflects that intent satisfaction is the primary mechanism for how-to and explanatory content, where AI Overviews directly answer the query.

Both unweighted mean scores (intuitive average across queries) and volume-weighted scores (market-impact view) are reported. Headline results use unweighted means unless otherwise stated.

2.8 AI-only composite

To isolate the effect of AI Overviews specifically, an AI-only composite was computed using only AI Overview contributions: - AI-only visibility: pixel contribution from AI Overview only; - AI-only click: CTR penalty from AI Overview only; - AI-only intent: same intent absorption, but only when an AI Overview is present (otherwise 0).

This allows comparison of total absorption versus absorption attributable solely to AI Overviews.

2.9 Validation and quality control

- **LLM self-consistency:** three independent runs per query, with variance flagged for review. The median absolute deviation between runs was 4.2 points for intent absorption and 6.8 points for estimated CTR.
- **Spot-check protocol:** a stratified sample of 100 LLM judgments is proposed for review against human SEO expert ratings (target: Cohen's $\kappa \geq 0.60$ or Krippendorff's $\alpha \geq 0.61$). This validation has not yet been conducted.
- **Heuristic baseline comparison:** the LLM-based index was compared against a binary heuristic (AIO present = 100, absent = 0). The LLM-based composite outperformed the heuristic by Spearman $\rho = +0.34$ against external CTR benchmarks (Seer Interactive, 2025; Ahrefs, 2025).

2.10 Metrics

- Mean and distribution of each sub-index and composite score;
- Category-level and device-level aggregates;
- AIO prevalence (% of SERPs with an AI Overview);
- Primary driver per category (the sub-index with the highest mean);
- Top-N and bottom-N queries by composite score;

- 95% bootstrap confidence intervals for headline means (10,000 resamples, seed 42).

3. Limitations

Geographic scope. All data is for US English Google SERPs. Results may not transfer to other locales, languages, or search engines.

Temporal scope. SERPs were collected in June 2026. Google changes feature prevalence and layout continuously; the dataset is a snapshot, not a time series.

CTR benchmarks. Click absorption relies on published industry benchmarks rather than first-party click data. Actual CTR varies by brand, vertical, position, and result type.

LLM judgment. Intent absorption depends on a local LLM's reading of query satisfaction. While ensemble aggregation reduces variance, model-specific biases may remain. Human validation is pending.

Pixel estimates. Visibility absorption uses SerpAPI rectangle data where available and default estimates otherwise. The 600 px fold height is mobile-first; desktop folds differ.

Category taxonomy. Intent classification is based on the catalog labels. Real-world queries can be ambiguous or multi-intent.

No answer-level evaluation. The study measures displacement of organic opportunity, not whether the AI-generated answers are accurate, safe, or useful.

4. Results

4.1 Aggregate absorption

Across all 14,728 queries:

Metric	Mean	95% CI
Composite absorption	67.72	[67.35, 68.09]
AI-only composite	41.65	[41.17, 42.13]
Visibility absorption	80.02	[79.54, 80.49]
Click absorption	70.03	[69.52, 70.54]
Intent absorption	58.84	[58.51, 59.18]

The primary driver across the full dataset is visibility absorption.

4.2 Absorption by intent category

Category	Queries	Composite	AI-only	Visibility	Click	Intent	Primary driver
Informational	6,034	76.43	52.17	92.30	84.25	63.30	Visibility absorption
Commercial	4,434	72.09	46.55	84.03	74.40	58.68	Visibility absorption
Transactional	2,842	52.00	24.22	59.52	46.87	53.44	Visibility absorption
Navigational	1,418	48.48	16.01	56.28	42.31	51.22	Visibility absorption

Informational and commercial queries cluster in High-to-Critical absorption. Transactional and navigational queries fall into Moderate-to-High absorption, with navigational queries the least absorbed.

4.3 Absorption by device

Device	Queries	Composite	Visibility	Click	Intent
Mobile	7,374	69.07	81.57	72.07	59.19
Desktop	7,354	66.37	78.46	67.99	58.49

Mobile scores are consistently higher across all categories and sub-indices (mobile vs. desktop composite: $\Delta = 2.69$, $p = 4.77 \times 10^{-13}$, Welch's t-test), consistent with smaller viewports amplifying feature displacement.

4.4 AI Overview prevalence

Category	Queries	AIO Count	AIO Prevalence
Informational	6,034	5,054	83.8%
Commercial	4,434	3,379	76.2%
Transactional	2,842	1,162	40.9%
Navigational	1,418	374	26.4%

AI Overviews appear most frequently in informational and commercial queries. Prevalence is similar across devices within each category (difference < 1 percentage point).

4.5 Distribution of absorption

The distribution of composite scores is right-skewed: a long tail of queries scores above 75 (Critical), while branded navigational queries form a left mode near zero (Figure 1). Sub-index histograms show that visibility and click absorption cluster at high values, while intent absorption is more dispersed (Figure 2).




Figure 1: Distribution of composite absorption scores by intent category.




Figure 2: Distributions of visibility, click, and intent absorption scores across all queries.




Figure 3: Boxplot of composite absorption scores by intent category.





Figure 4: Boxplot of composite absorption scores by device.



Figure

5: Pearson correlation heatmap of sub-indices and composite score.

4.6 Highest- and lowest-absorption queries

Highest absorption (composite > 95):

Keyword	Composite	Category	Device
what is hackrf one	96.7	commercial	mobile
how to poach eggs	96.3	commercial	desktop
best fuel injector cleaner	95.9	commercial	mobile
xbox series s specs	95.9	commercial	mobile
easiest way to make candles	95.9	commercial	mobile

These are informational or commercial how-to/what-is queries that AI Overviews answer directly.

Lowest absorption (composite < 5):

Keyword	Composite	Category	Device
indeed login	3.7	navigational	desktop
capital one login	3.7	navigational	mobile
ziprecruiter login	3.7	navigational	mobile
schwab login	4.5	navigational	mobile
american airlines flight status	4.5	navigational	mobile

These are branded navigational queries where the user explicitly seeks the brand's site.

4.7 AI-only versus total absorption

The AI-only composite (41.65) is substantially lower than the total composite (67.72). The difference (26.07 points) indicates that approximately 38% of measured absorption is attributable to non-AIO SERP features such as featured snippets, PAA, and inline media.

Category-level AI-only composites confirm the pattern: informational queries show the largest AIO-specific effect (52.17), while navigational queries show the smallest (16.01). The gap between total and AI-only composite is largest in categories where non-AIO features are common.

4.8 Sub-index correlations

	Visibility	Click	Intent	Composite	AI-only
Visibility	1.000	0.921	0.199	0.905	0.665
Click	0.921	1.000	0.241	0.934	0.852
Intent	0.199	0.241	1.000	0.528	0.375
Composite	0.905	0.934	0.528	1.000	0.819
AI-only	0.665	0.852	0.375	0.819	1.000

Visibility and click absorption are very highly correlated ($r = 0.921$), reflecting that SERP features both displace organic results and reduce CTR. Intent absorption is more weakly correlated with the other two sub-indices ($r \approx 0.20$ – 0.24), supporting the decision to measure it as a distinct mechanism.

4.9 Pairwise category comparisons

Welch's two-sample t-test (categories have different sample sizes):

Pair	Δ Mean	p-value	Significant at $\alpha=0.05$
Informational vs Commercial	+4.35	1.75×10^{-32}	Yes
Informational vs Transactional	+24.43	$< 10^{-300}$	Yes
Informational vs Navigational	+27.96	2.32×10^{-252}	Yes
Commercial vs Transactional	+20.09	1.34×10^{-259}	Yes
Commercial vs Navigational	+23.61	5.31×10^{-187}	Yes
Transactional vs Navigational	+3.53	1.11×10^{-5}	Yes

All pairwise differences are statistically significant at $\alpha = 0.05$.

5. Discussion

5.1 Visibility is the dominant mechanism

Visibility absorption is the highest sub-index and the primary driver in every category. This finding is consistent with Fubel et al. (2023), who reported that SERP feature density exerts a negative influence on organic CTR, particularly for position 1 results. The high visibility absorption scores (mean 80.02) indicate that organic results are, on average, pushed substantially below the viewport fold.

5.2 Intent category and absorption

Informational queries show the highest composite absorption (76.43), followed by commercial (72.09), transactional (52.00), and navigational (48.48). This ordering is consistent with the AI Overview prevalence data: informational queries have the highest AIO prevalence (83.8%) and navigational the lowest (26.4%). The pattern aligns with Seer Interactive's (2025) finding that informational queries experienced the steepest CTR decline.

5.3 Mobile and desktop differences

Mobile composite absorption (69.07) exceeds desktop (66.37) by 2.69 points ($p = 4.77 \times 10^{-13}$). The difference is consistent across all sub-indices and categories. This finding aligns with SparkToro and Datos (2024), who reported higher zero-click rates on mobile (77%) than

desktop, and with the intuitive expectation that smaller viewports amplify feature displacement.

5.4 AI Overviews and non-AIO features

The AI-only composite (41.65) is 26.07 points below the total composite (67.72), indicating that non-AIO SERP features contribute approximately 38% of measured absorption. Featured snippets, PAA, inline media, knowledge panels, and shopping results collectively displace a significant share of organic opportunity. This finding is consistent with Fubel et al. (2023), who documented negative CTR effects from multiple traditional SERP features.

5.5 Relationship to prior work

The headline composite of 67.72 is directionally consistent with SparkToro and Datos's (2024) 58.5% zero-click estimate and Seer Interactive's (2025) 61% AIO CTR decline. The decomposition into three sub-indices provides a finer-grained view than either aggregate source. The correlation between visibility and click absorption ($r = 0.921$) supports Fubel et al.'s (2023) finding that SERP features both displace results and reduce CTR. The weaker correlation between intent absorption and the other sub-indices ($r \approx 0.20$) suggests that intent satisfaction is a distinct mechanism, not fully captured by position-based or CTR-based measures alone.

6. Validation Protocol

A full human-validation study has not yet been conducted. The protocol below is proposed for future work.

6.1 Sampling design

A stratified sample of 100 (query, SERP) pairs:

Stratum	Count	Selection criteria
High absorption	25	composite_score $\geq$ 75, balanced across categories
Moderate absorption	25	$50 \leq$ composite_score $<$ 75
Low absorption	25	composite_score $<$ 50
AIO present	15	random sample where aio_present = 1
AIO absent	10	random sample where aio_present = 0

6.2 Raters

2–3 independent SEO or search-content experts. Raters see the query and rendered SERP (or structured feature summary + AIO text) but not the LLM's score or reasoning. They provide ratings for intent satisfaction and estimated residual CTR using the same rubric as the LLM.

6.3 Metrics

1. Cohen's κ for categorical intent_satisfied agreement (Yes/Partially/No).
2. Intraclass correlation coefficient (ICC) or Pearson r for estimated_ctr_percent between LLM median and human median.
3. Krippendorff's α for the full multi-rater, multi-category setup.

6.4 Acceptance thresholds

Gate	Threshold	Action if failed
Self-consistency (3-run ensemble)	Krippendorff's $\alpha \geq 0.61$	Revise prompt
Category correlation vs. external benchmarks	Spearman $\rho \geq 0.50$	LLM not adding signal
Human agreement	Cohen's $\kappa \geq 0.60$	Deploy as triage only

7. Conclusions

1. Absorption can be decomposed into visibility, click, and intent mechanisms. The three sub-indices vary independently enough to justify separate measurement: visibility and click absorption are highly correlated ($r = 0.921$), while intent absorption is weakly correlated with both ($r \approx 0.20$).
 2. Informational queries are the most absorbed category (mean 76.43, 95% CI [76.06, 76.81]); navigational queries are the least absorbed (mean 48.48, 95% CI [47.20, 49.78]). All pairwise category differences are statistically significant ($p < 10^{-5}$).
 3. Mobile SERPs are more absorbed than desktop SERPs (mean 69.07 vs. 66.37, $p = 4.77 \times 10^{-13}$).
 4. Non-AIO SERP features contribute approximately 38% of measured absorption. The AI-only composite (41.65) is 26.07 points below the total composite (67.72).
 5. Visibility absorption is the dominant driver across all categories, with a mean of 80.02 (95% CI [79.54, 80.49]).
-

8. Data and Code Availability

- **Full dataset:** `absorption_index_dataset.csv` and `.json` (14,728 rows)
- **Category summary:** `absorption_index_category_summary.csv` and `.json`
- **Interactive browser:** `absorption_index_browser.html`
- **Pipeline code:** `main.py` and `steps/` directory
- **Configuration:** `config.yaml`

The pipeline is designed to be reproducible: fixed seeds, specified prompts, open-source LLM, and stored raw SERP JSON.

9. Future Work

- Longitudinal tracking: monthly re-runs to establish trend lines.
 - Vertical slicing: add health, finance, software, and news verticals.
 - Additional markets: extend beyond US English.
-

- Embedding-model sensitivity: test whether chunker/retriever choices affect visibility scoring.
 - Blind human calibration: expand the LLM spot-check to a larger, fully blind inter-rater study.
 - Answer-level evaluation: measure whether AI-generated answers are accurate and complete, not just displacing.
-

References

1. Ahrefs. (2025). AI Overview impact on organic CTR: Analysis of 300,000 keywords. Ahrefs Blog.
 2. BrightEdge. (2025). AI Overviews at the one-year mark: Prevalence, growth, and vertical impact. BrightEdge Research.
 3. European Parliament. (2026). The impact of AI-generated search answers on digital markets and publisher revenue. Policy Department for Economic, Scientific and Quality of Life Policies, PE 787.211.
 4. FirstPageSage. (2025). Google click-through rates by ranking position in 2026. First-PageSage Benchmark Report.
 5. Fubel, E., Groll, N. M., Gundlach, P., Han, Q., & Kaiser, M. (2023). Beyond rankings: Exploring the impact of SERP features on organic click-through rates. arXiv preprint arXiv:2306.01785.
 6. Li, X., Wang, Y., & de Rijke, M. (2017). Investigating queries and search failures in academic search. Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, 123–132.
 7. Natividad, A., & Fishkin, R. (2025). Zero click marketing. SparkToro.
 8. Pew Research Center. (2025). How Americans interact with AI-powered search results. Pew Research Center Internet & Technology Report.
 9. Reuters Institute. (2026). Journalism, media, and technology trends and predictions 2026. Reuters Institute for the Study of Journalism, University of Oxford.
 10. Seer Interactive. (2025). Organic CTR in the age of AI Overviews: A longitudinal study of 3,119 queries. Seer Interactive Research.
 11. Sistrix. (2026). AI Overviews and organic CTR: Position-level analysis in the German market. Sistrix Data Study.
 12. SparkToro & Datos. (2024). We analysed 332 million queries over 21 months to uncover never-before-published data on how people use Google. SparkToro Blog.
-

13. SparkToro & Datos. (2025). 2024 zero-click search study: For every 1,000 US Google searches, only 374 clicks go to the open web. SparkToro Blog.
 14. Ye, Z., Liu, Y., & Zhang, M. (2022). Understanding zero-click search via brain signals. Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 1234–1243.
-

Appendix A: Prompts

A.1 Intent satisfaction prompt

You are analyzing whether the SERP features fully satisfy a searcher's intent.

Query: {query}

AI Overview Text: {aio_text}

SERP Features Present: {features_summary}

If an AI Overview is present, focus your analysis on it. If no AI Overview is present, assess whether the OTHER SERP features (featured snippets, answer boxes, videos, images, PAA, etc.) fully satisfy the query.

Think step by step, then respond with ONLY a JSON object:

```
{
  "intent_satisfied": "Yes|Partially|No",
  "confidence": 1-5,
  "reasoning": "...",
  "information_gaps": ["..."],
  "intent_type": "Transactional|Informational|Navigational|Mixed"
}
```

A.2 CTR estimation prompt

You are estimating what percentage of searchers would click an organic result after seeing the SERP.

Query: {query}

AI Overview Text: {aio_text}

SERP Features Present: {features_summary}

Intent Assessment: {intent_assessment}

Use this rubric:

- 0-10%: SERP fully answers the query, no need to click
- 11-30%: Minor information gaps, most users satisfied
- 31-50%: Moderate gaps, significant minority will click
- 51-70%: Major gaps, majority will likely click through
- 71-100%: SERP is inadequate, most users will click

Think step by step, then respond with ONLY a JSON object:

```
{
  "estimated_ctr_percent": 0-100,
  "confidence": 1-5,
  "reasoning": "...",
  "key_driver": "..."
}
```

Appendix B: Bootstrap Confidence Intervals

Bootstrap confidence intervals were computed for all headline means and pairwise category comparisons.

Procedure: 1. For each metric, draw 10,000 resamples with replacement from the full dataset of 14,728 queries. 2. Compute the mean for each resample. 3. Report the 2.5th and 97.5th percentiles of the resampled mean distribution as the 95% confidence interval. 4. Random seed was fixed at 42 for reproducibility.

Software: Python 3.11, NumPy 1.26, SciPy 1.12.

Example (composite absorption): - Observed mean: 67.72 - Bootstrap distribution: 10,000 resampled means - 95% CI: [67.35, 68.09] - Standard error (bootstrap): 0.19

Pairwise comparisons: Welch's two-sample t-test was used for pairwise category comparisons, as categories have different sample sizes and potentially unequal variances. The test statistic and p-value were computed from the observed data; bootstrap CIs for the mean difference were also computed via resampling.

Appendix C: Implementation Notes

The pipeline is implemented in Python across eight steps:

1. `_01_keyword_fetcher.py` — load curated catalog and apply stratified sampling
2. `_02_serp_collector.py` — collect SERPs via DataForSEO and SerpAPI
3. `_03_feature_detector.py` — detect AI/answer features from SERP JSON
4. `_04_visibility_scorer.py` — compute visibility absorption
5. `_05_llm_assessor.py` — LLM intent + CTR assessment with 3-run ensemble
6. `_06_click_calculator.py` — compute click absorption
7. `_07_index_calculator.py` — compute sub-indices and composite scores
8. `_08_report_generator.py` — generate self-contained HTML report

All configuration (weights, benchmarks, prompts, sampling parameters) is centralized in `config.yaml`.

Appendix D: Glossary

- **Absorption:** the degree to which SERP features reduce organic search opportunity for a given query.
 - **AI Overview (AIO):** Google's synthesized answer block generated from multiple sources.
 - **Click absorption:** estimated CTR reduction from detected SERP features.
 - **Intent absorption:** estimated degree to which the SERP satisfies user intent without a click.
 - **Visibility absorption:** estimated pixel-space displacement of organic results below the fold.
 - **Composite score:** weighted combination of the three sub-indices, category-specific.
 - **AI-only composite:** composite score using only AI Overview contributions.
-

Figures

All figures referenced in §4.5 are available in the `plots/` directory accompanying this paper.
