



The Architectural Visualization AI Visibility Index

Which visualization tools AI answer engines actually name, and why

126 prompts · 6 buyer personas · 4 answer engines · 504 responses

AEO Wrangler B2B Software Index · September 2026

Executive summary

We asked ChatGPT, Gemini, Google AI Mode, and Microsoft Copilot 126 buying questions about architectural visualization software, spanning six buyer situations from concept sketching to studio production pipelines. No prompt named a vendor or a product. Every name in this report comes from what the engines said. None came from what we asked.

Across 504 responses, the engines named 247 distinct software products, 101 of them visualization tools from 89 different vendors, and cited 3,282 URLs from 455 domains.

Nine findings stand out.

Four vendors own the answer. Chaos is named in 120 of 126 prompts and scores 77.0 on equal channel weight. Epic Games follows at 64.1, D5 at 46.2, and Lumion at 28.0. The fifth-placed vendor, Blender Foundation, scores 8.7. Together the four account for 66.6 percent of all visualization vendor mentions. The field of 89 vendors is real, but in AI answers it is a four-vendor market with a long tail.

New companies make up a third of the field and a tenth of the answers. 34 of the 101 visualization products named come from independent companies that launched them in 2022 or later, almost all of them AI image generators. Together they earn 9.5 percent of visualization product mentions, and only six are named often enough to clear the inclusion floor. The one place they break through is the concept stage.

The new entrants have taken the concept stage. At least one independent AI-era newcomer appears in 19 of the 21 concept exploration prompts. In the other five personas combined, new entrants appear in 18 of 105 prompts. LookX ranks second in concept exploration behind Chaos, ahead of D5, Epic, and Lumion. Once a design model exists, the established vendors hold nearly every answer.

The most visible AI-era product belongs to the market leader. Veras, built by EvolveLAB and acquired by Chaos in February 2025, is named in 36 prompts, more than any independent new entrant. Established vendors launched or bought nine of the AI-era products in the dataset, and those nine earn 6.8 percent of mentions against 9.5 percent for the 34 independent newcomers.

The engines are not describing the same market. Lumion is named in 53.2 percent of ChatGPT prompts and 12.7 percent of Gemini prompts. Unreal Engine shows the reverse, at 20.6 percent on ChatGPT and 46.8 percent on Gemini. A buyer's shortlist depends heavily on which assistant they happen to open.

Legacy products are carried by memory. Products released before 2015 are backed by a citation to their vendor's own site in 44.6 percent of mentions, against 58.7 to 60.9 percent for newer products. On Google's two engines the gap is wider, 18.9 percent against 32.5 to 34.1 percent. Twinmotion and Unreal Engine are named on Gemini and AI Mode on reputation alone: their vendor's site supports 7 and 6 percent of those mentions. Long-established vendors can do little AEO work and still rank, because the engines remember them.

ChatGPT and Copilot read vendor websites. Google's engines read everyone else. 97.7 percent of ChatGPT's citations and 94.0 percent of Copilot's point to vendor-owned domains. For AI Mode the figure is 46.4 percent, and more than a third of its citations are "best software" and comparison articles. When ChatGPT names a visualization

vendor, it cites that vendor's own site 75.3 percent of the time. Gemini does so 13.3 percent of the time.

Vendors are writing the lists that Google's engines read. Maxon's two "best architectural visualization software" articles were cited 53 times, mostly by AI Mode, in answers that usually did not name Maxon at all. The AI rendering newcomers draw 67 of their 135 AI Mode citations from listicles they published themselves.

For buyers, ChatGPT is the engine to trust and Gemini the one to check. We compared every price claim for Enscape, Twinmotion, Lumion, V-Ray, Corona, and Veras with current vendor pricing pages. Every verifiable claim from ChatGPT (35 of 35) and Copilot (33 of 33) matched the published list price. Gemini matched in 12 of 19, gave no sources for a third of its answers, and cites a vendor's own site for only 13.3 percent of the vendors it names. Gemini and AI Mode each told a small practice that Twinmotion is free below \$100,000 in revenue. The real threshold is \$1 million. Copilot is as accurate as ChatGPT but far narrower: 37 of its answers named one visualization product or none.

1. Why this index exists, and what it does not measure

Every vendor in this category now has a marketing team asking the same question: when a buyer asks which rendering tool to use, does our name come up? Until now the answers have been anecdotal, based on a prompt someone tried once and a screenshot.

This index replaces that with a fixed, disclosed, repeatable instrument.

What it measures is visibility: whether a product is named, how often, by which engine, and to which buyer. It is a share of voice study.

It does not measure quality, fit, or market position. Being named often is no evidence that a product is good, and being absent is no evidence that a product is bad or losing customers. Read every low placement in this report as a statement about AI visibility and nothing else.

This edition adds one thing the index has not covered before. Section 10 looks at the same data from the buyer's side and asks a simpler question: when these engines give specific facts, such as prices and license terms, are the facts right?

Editorial independence. AEO Wrangler has taken no vendor money in this category, has no paid placements in this index, and offers none. No vendor saw this report or its methodology before publication.

2. Method

Category boundary. Software whose primary job is producing architectural imagery or experiences from design intent: offline render engines, real-time engines, AI image generators built for architecture, general-purpose AI image generators, and immersive and interactive platforms. A product is excluded from the core category if its primary job is authoring the building model itself. That test is what separates Enscape, which we count as core, from Revit, Archicad, SketchUp, and Rhino, which we count as design authoring even though each has rendering built in.

Core and extended tables. The core ranking contains the 101 visualization products. The extended ranking in section 5 contains every software product named, 247 in all, with the adjacency labeled: design authoring, modeling and DCC, asset libraries, render management, coordination and model review, and ten other adjacent types. Hardware such as GPUs and headsets, IT infrastructure such as identity and cloud services, and image models from the companies behind the engines themselves are excluded from both tables.

Vendor consolidation. Vendor figures combine a vendor's products by union, so a prompt counts once however many of the vendor's products it names. A prompt naming both Enscape and V-Ray counts once for Chaos. Acquired products roll up to the current owner: Veras (EvolveLAB, acquired February 2025) and Anima (AXYZ, acquired July 2023) to Chaos, and IrisVR Prospect and The Wild (acquired March 2022) to Autodesk. Per-product figures for the multi-product vendors appear in Appendix C.

The Unreal rule. Twinmotion is built on Unreal Engine, and answers describing Twinmotion often mention Unreal as its underlying technology ("powered by Unreal Engine 5", "Lumen and Nanite via Unreal"). Those mentions describe Twinmotion and are not counted as Unreal Engine mentions. Unreal named as a separate product, an export target, or part of a pipeline still counts.

Personas. Six, chosen because each carries a materially different vendor set.

Code	Persona	Buyer
CE	Concept Exploration	Architect or designer at concept stage, before a developed model exists
DI	In-Model Design Iteration	Architect or design engineer inside an authoring session, on a model that changes daily
CP	Client and Stakeholder Presentation	Project architect or principal preparing for a client meeting, design review, or public consultation
DM	Developer and Real Estate Marketing	A developer's marketing lead, or the practice serving one
IX	Immersive and Interactive	Firm or developer needing a walkthrough, VR design review, sales gallery, or configurator
SP	Studio Production Pipeline	Owner or technical director of a visualization studio, or the visualization team inside a large practice

Prompts. 21 per persona, 126 in total, in three categories.

High Buyer Intent (54 prompts) phrases an active purchase decision. Example: "I'm a one-person architecture practice and I want to start pitching concepts with rendered imagery instead of hand sketches. What software should I buy?"

Category With Intent (54 prompts) phrases a requirement that only a product can satisfy. Example: "I work alone as an architect and need architectural concept imaging software with a free or low-cost tier that permits commercial use of the output. What meets that?"

Open Question (18 prompts) phrases the problem as how or what approach. Example: "How do architects generate concept imagery before a detailed model exists?" These are counted with the other two categories, because most of them named products.

Each persona's prompts are split across three firm size bands (solo or freelance, studio or practice, large or enterprise), 36 prompts per band, with 18 open questions left size-neutral. 45 prompts describe the modeling environment generically (a BIM model under revision, a massing model, an IFC file from another party, a polygonal modeling package, a parametric surface model). No prompt names a vendor, product, or brand.

Engines and dates. ChatGPT (GPT-5.6 Sol, High), Gemini (3.6 Flash Extended), Google AI Mode, and Microsoft Copilot (Think Deeper, temporary chat). Responses were collected between September 12 and September 21, 2026. Every prompt was run in a fresh incognito or temporary session on each engine, so no answer could be influenced by a prior prompt, account history, or personalization. Each response was captured in full with its source list.

Extraction. Products were identified with a curated lexicon of alias patterns, built by reading all 504 responses, with no predefined vendor list. The lexicon was locked before scoring. Sources were extracted as URLs, deduplicated within each answer, reduced to registrable domains, and classified by type.

3. The weighting decision

The four engines produce different volumes of evidence. Copilot returned 9.5 citations per prompt, AI Mode 8.7, ChatGPT 6.1, and Gemini 1.8. Gemini gave no sources at all for 42 of 126 prompts.

That asymmetry does not distort the mention counts, because each engine answered the same 126 prompts and named a similar number of products: 10.5 per answer for Gemini, 9.7 for AI Mode, 9.6 for ChatGPT, and 8.0 for Copilot. Source data is less balanced, so every source figure in this report is expressed as a share of that engine's own citations.

The ranking uses **equal channel weight**: each vendor's coverage rate within each engine, averaged across the four. Every engine contributes one quarter, however much it says.

We publish no usage-weighted score in this edition. Reliable usage shares for these four engines do not exist in a form we could defend. AI Mode sits inside Google Search and Copilot inside Windows and Microsoft 365, so their usage cannot be separated cleanly from the products they are embedded in. Every table reports per-engine rates beside the composite, so readers who prefer a different weighting can recompute the ranking themselves. **Union coverage**, the number of prompts where a vendor was named on any engine, appears as a reference column.

Inclusion floor. A vendor or product must be named in at least 8 of the 126 prompts (6 percent) to be ranked. Below that, the difference between two appearances and one is noise, so those entries are reported as an unordered set in Appendix A. Ties are broken by union coverage, then persona breadth, then alphabetically.

4. Core visualization ranking

4.1 By vendor

The primary leaderboard. The equal weight score sets the rank: it is the average of the four engine coverage rates. Each engine column shows the percentage of the 126 prompts in which the vendor was named on that engine.

Rank	Vendor	Equal weight	ChatGPT	Gemini	AI Mode	Copilot	Union (of 126)	Personas
1	Chaos	77.0	83.3%	82.5%	81.7%	60.3%	120	6
2	Epic Games	64.1	64.3%	73.8%	63.5%	54.8%	105	6
3	D5 (Dimension 5)	46.2	62.7%	42.1%	42.1%	38.1%	89	6
4	Lumion (Act-3D)	28.0	53.2%	12.7%	22.2%	23.8%	73	6
5	Blender Foundation	8.7	4.0%	8.7%	15.1%	7.1%	26	5
6	Autodesk	8.7	7.9%	7.9%	8.7%	10.3%	20	5
7	LookX	8.3	6.3%	10.3%	8.7%	7.9%	22	4

Rank	Vendor	Equal weight	ChatGPT	Gemini	AI Mode	Copilot	Union (of 126)	Personas
8	Maxon	5.8	4.8%	7.1%	5.6%	5.6%	17	4
9	Arkio	5.4	7.1%	7.1%	2.4%	4.8%	11	1
10	Midjourney	5.2	4.8%	7.1%	7.1%	1.6%	16	3
11	Shapspark	5.0	5.6%	7.1%	3.2%	4.0%	9	2
12	Unity Technologies	4.8	1.6%	3.2%	8.7%	5.6%	16	5
13	OTOY	4.0	4.0%	7.9%	3.2%	0.8%	15	3
14	Gendo	3.6	1.6%	4.0%	7.1%	1.6%	11	3
15	Stability AI	3.2	1.6%	4.8%	4.8%	1.6%	9	1
16	Krea	3.0	2.4%	4.8%	4.0%	0.8%	13	4
17	SENTIO	3.0	3.2%	2.4%	3.2%	3.2%	8	1
18	ArchiVinci	2.2	4.0%	0.0%	4.0%	0.8%	10	3
19	Vizcom	2.0	0.8%	6.3%	0.0%	0.8%	8	2

Below this line, 70 further vendors were named in fewer than eight prompts.

Reading the table. The structure is four vendors, then a cliff. Chaos is named on at least 60 percent of prompts on every engine, and its lead rests on breadth: five products above the inclusion floor, covering concept work (Veras), design iteration and presentation (Enscape), and production rendering (V-Ray, Corona, and Vantage). Epic reaches second through two products, Twinmotion and Unreal Engine, that between them cover almost every situation a buyer can describe. D5, a product first released in 2020, has reached third on the strength of a single renderer. Lumion is fourth, but its position depends heavily on ChatGPT, where it is named 53.2 percent of the time, four times its Gemini rate.

The gap between fourth and fifth is 19.3 points. No vendor below fourth reaches 16 percent coverage on any engine.

4.2 By product

The same measure applied to individual products. Chaos places five products in the top ten and Epic two.

Rank	Product	Vendor	Equal weight	ChatGPT	Gemini	AI Mode	Copilot	Union
1	Twinmotion	Epic Games	50.4	59.5%	52.4%	49.2%	40.5%	91
2	D5 Render	D5	46.0	62.7%	42.1%	42.1%	37.3%	89
3	Enscape	Chaos	46.0	57.1%	44.4%	43.7%	38.9%	85
4	V-Ray	Chaos	43.3	42.9%	52.4%	50.8%	27.0%	85
5	Unreal Engine	Epic Games	31.5	20.6%	46.8%	36.5%	22.2%	71
6	Lumion	Lumion	27.8	53.2%	12.7%	22.2%	23.0%	73

Rank	Product	Vendor	Equal weight	ChatGPT	Gemini	AI Mode	Copilot	Union
7	Corona	Chaos	26.4	26.2%	33.3%	31.0%	15.1%	55
8	Veras	Chaos	13.3	11.1%	17.5%	13.5%	11.1%	36
9	Chaos Vantage	Chaos	9.7	7.9%	18.3%	7.9%	4.8%	30
10	Blender Cycles / Eevee	Blender Foundation	8.7	4.0%	8.7%	15.1%	7.1%	26
11	LookX AI	LookX	8.3	6.3%	10.3%	8.7%	7.9%	22
12	Redshift	Maxon	5.8	4.8%	7.1%	5.6%	5.6%	17
13	Arkio	Arkio	5.4	7.1%	7.1%	2.4%	4.8%	11
14	Midjourney	Midjourney	5.2	4.8%	7.1%	7.1%	1.6%	16
15	Shapspark	Shapspark	5.0	5.6%	7.1%	3.2%	4.0%	9
16	Unity	Unity Technologies	4.8	1.6%	3.2%	8.7%	5.6%	16
17	Autodesk Workshop XR	Autodesk	4.6	6.3%	3.2%	4.0%	4.8%	8
18	Octane	OTOY	4.0	4.0%	7.9%	3.2%	0.8%	15
19	Gendo	Gendo	3.6	1.6%	4.0%	7.1%	1.6%	11
20	Stable Diffusion	Stability AI	3.2	1.6%	4.8%	4.8%	1.6%	9
21	Krea	Krea	3.0	2.4%	4.8%	4.0%	0.8%	13
22	SENTIO VR	SENTIO	3.0	3.2%	2.4%	3.2%	3.2%	8
23	D5 Lite	D5	2.8	7.9%	0.8%	0.8%	1.6%	12
24	Arnold	Autodesk	2.8	0.8%	3.2%	2.4%	4.8%	11
25	Lumion View	Lumion	2.8	7.9%	0.0%	0.8%	2.4%	11
26	ArchiVinci	ArchiVinci	2.2	4.0%	0.0%	4.0%	0.8%	10
27	Vizcom	Vizcom	2.0	0.8%	6.3%	0.0%	0.8%	8

D5 Render and Enscape tie at 46.0 and are separated by union coverage, 89 prompts against 85. The top four products sit within 7.1 points of each other, and each has a different shape: Twinmotion and Enscape are the most consistent across engines, D5 Render is strongest on ChatGPT, and V-Ray is strongest on the two Google engines.

5. Extended ranking including adjacent products

The same measure applied to every software product named, whatever its category. This table answers a different question: when a buyer asks an AI assistant about visualization, what are they actually shown?

Rank	Product	Category	Equal weight	ChatGPT	Gemini	AI Mode	Copilot	Union
1	Revit	Design authoring	64.3	65.9%	60.3%	69.8%	61.1%	109
2	Rhino	Design authoring	58.9	59.5%	60.3%	63.5%	52.4%	106
3	SketchUp	Design authoring	51.4	57.1%	48.4%	50.8%	49.2%	98
4	Twinmotion	Core	50.4	59.5%	52.4%	49.2%	40.5%	91
5	Archicad	Design authoring	46.6	50.8%	46.0%	44.4%	45.2%	88
6	D5 Render	Core	46.0	62.7%	42.1%	42.1%	37.3%	89
7	Enscape	Core	46.0	57.1%	44.4%	43.7%	38.9%	85
8	V-Ray	Core	43.3	42.9%	52.4%	50.8%	27.0%	85
9	3ds Max	Modeling / DCC	38.9	34.9%	42.9%	45.2%	32.5%	73
10	Unreal Engine	Core	31.5	20.6%	46.8%	36.5%	22.2%	71
11	Lumion	Core	27.8	53.2%	12.7%	22.2%	23.0%	73
12	Corona	Core	26.4	26.2%	33.3%	31.0%	15.1%	55
13	Vectorworks	Design authoring	25.6	38.9%	12.7%	13.5%	37.3%	67
14	Blender	Modeling / DCC	25.0	17.5%	27.0%	35.7%	19.8%	63
15	Datasmith	Connector / live link	18.3	13.5%	27.0%	19.8%	12.7%	48
16	Cinema 4D	Modeling / DCC	17.5	16.7%	20.6%	17.5%	15.1%	43
17	Veras	Core	13.3	11.1%	17.5%	13.5%	11.1%	36
18	Autodesk Forma	Generative design	10.3	14.3%	8.7%	2.4%	15.9%	25
19	Maya	Modeling / DCC	9.9	5.6%	12.7%	11.9%	9.5%	28
20	Chaos Vantage	Core	9.7	7.9%	18.3%	7.9%	4.8%	30

The design tool sets the frame. Revit is the most-named product in a visualization study. The prompts never mention Revit, but buyers describe a BIM model or a design workflow, and the engines answer by naming the tools on either side of the render. Design authoring products account for 27.6 percent of all product mentions, core visualization products for 43.6 percent (34.5 AEC-built and 9.1 general-purpose), and modeling and DCC tools for 10.1 percent. The remaining 18.7 percent spreads across thirteen adjacent categories led by coordination and model review, live-link connectors, and asset libraries.

For a visualization vendor, that means the answer is usually organized around the host application. "Which renderer works inside Revit" is the frame the engines reach for, and a vendor whose integrations are poorly documented is competing from outside that frame.

Which authoring tool the engines assume. Thirteen prompts describe "a BIM model under active revision" without naming the software. Revit was named in response to all 13, on all four engines for 12 of them and on three engines for the other. Archicad was named in response to 13 as well, but less consistently, and Vectorworks in 10. When a buyer says BIM without saying which, the engines assume Revit.

6. Per-engine leaderboards

The four engines agree on who the top four vendors are and in what order. Below the top four, and in how strongly they favor each of the four, they disagree considerably.

6.1 ChatGPT (GPT-5.6 Sol, High)

Rank	Vendor	Coverage	Own-domain citations
1	Chaos	83.3%	194
2	Epic Games	64.3%	92
3	D5 (Dimension 5)	62.7%	124
4	Lumion (Act-3D)	53.2%	62
5	Autodesk	7.9%	69
6	Arkio	7.1%	13
7	LookX	6.3%	6
8	Shapspark	5.6%	9
9	Maxon	4.8%	5
10	Midjourney	4.8%	5

ChatGPT is the most concentrated engine. The top four vendors take 75.3 percent of its visualization vendor mentions, the highest share on any engine, and 97.7 percent of its citations point to vendor-owned sites. ChatGPT reads vendor documentation, pricing, and licensing pages closely, and drew only six of its 766 citations from listicle pages. It is also where Lumion and D5 are strongest: D5 is named in 62.7 percent of ChatGPT prompts and Lumion in 53.2 percent, their best figures on any engine.

(Autodesk's own-domain citations here are mostly to Revit, 3ds Max, and licensing pages. Its visualization products are named in 10 ChatGPT prompts.)

6.2 Gemini (3.6 Flash Extended)

Rank	Vendor	Coverage	Own-domain citations
1	Chaos	82.5%	19
2	Epic Games	73.8%	3
3	D5 (Dimension 5)	42.1%	8
4	Lumion (Act-3D)	12.7%	3
5	LookX	10.3%	0

Rank	Vendor	Coverage	Own-domain citations
6	Blender Foundation	8.7%	0
7	Autodesk	7.9%	4
8	OTOY	7.9%	0
9	Arkio	7.1%	0
10	Maxon	7.1%	1

Gemini names more products per answer than any other engine, 10.5 on average, and shows the least evidence for them: 1.8 citations per prompt, with 42 of 126 answers carrying no sources at all. Epic reaches its best figure on any engine here, 73.8 percent, on the strength of Unreal Engine, which Gemini names in 46.8 percent of prompts. Lumion's 12.7 percent is its worst figure on any engine.

Because Gemini cites so little, its rankings are the closest available view of what the model absorbed in training. The low citation support rates reported for Gemini in section 9 reflect that low volume.

6.3 Google AI Mode

Rank	Vendor	Coverage	Own-domain citations
1	Chaos	81.7%	84
2	Epic Games	63.5%	13
3	D5 (Dimension 5)	42.1%	28
4	Lumion (Act-3D)	22.2%	13
5	Blender Foundation	15.1%	4
6	Autodesk	8.7%	13
7	LookX	8.7%	0
8	Unity Technologies	8.7%	7
9	Gendo	7.1%	14
10	Midjourney	7.1%	1

AI Mode is the retrieval engine of this study and the most open to influence. It cited 1,097 URLs from 300 domains, more than twice as many domains as any other engine, and only 46.4 percent of its citations were vendor-owned. YouTube supplied 10.1 percent of its citations, visualization studios 14.1 percent, and other software companies 16.3 percent. AI Mode is also where the AI rendering newcomers perform best: Gendo reaches ninth, and new entrants appear in 30 AI Mode prompts against 13 on ChatGPT.

6.4 Microsoft Copilot

Rank	Vendor	Coverage	Own-domain citations
1	Chaos	60.3%	258
2	Epic Games	54.8%	180
3	D5 (Dimension 5)	38.1%	89
4	Lumion (Act-3D)	23.8%	54
5	Autodesk	10.3%	147
6	LookX	7.9%	15
7	Blender Foundation	7.1%	29
8	Maxon	5.6%	19
9	Unity Technologies	5.6%	11
10	Arkio	4.8%	12

Copilot cited the most URLs of any engine, 1,191, and 94.0 percent of them were vendor-owned. It names fewer products per answer than the others, 8.0, and 23 of its answers named no visualization product at all. When Copilot does name a visualization vendor, it cites that vendor's own site 89.6 percent of the time, the highest support rate in the study.

7. Persona leaderboards

Twenty-one prompts per persona, ranked by equal channel weight within the persona. Each table shows the top six; lower places are discussed in the text where they matter. Vendors whose persona rank differs from their overall rank by four or more places are marked with an asterisk.

Concept Exploration (CE)

Rank	Vendor	Score	Overall rank
1	Chaos	70.2	1
2	LookX*	44.0	7
3	D5 (Dimension 5)	38.1	3
4	Epic Games	27.4	2
5	Midjourney*	26.2	10
6	Stability AI*	19.0	15

The one persona where the leaderboard is rewritten. LookX, a product launched in 2023, ranks second, ahead of every established vendor except Chaos, and Chaos holds first here largely through Veras. Just below the table, Gendo and Adobe (through Firefly) score 17.9 in seventh and eighth. Lumion, fourth overall, falls to ninth at 16.7, and Krea is tenth. Six of the top ten are AI image vendors. This is the persona where a product that needs no 3D

model can win outright, and the engines treat it that way.

In-Model Design Iteration (DI)

Rank	Vendor	Score	Overall rank
1	Chaos	86.9	1
2	Epic Games	75.0	2
3	D5 (Dimension 5)	71.4	3
4	Lumion (Act-3D)	46.4	4
5	Blender Foundation	11.9	5
6	Maxon	9.5	8

The most conventional persona. D5 is named in all 21 prompts, the only vendor to reach full coverage in any persona apart from Chaos. Live-link and plugin integration with the authoring tool dominate the answers. AI newcomers appear in 7 of the 21 prompts, and none ranks in the top six.

Client and Stakeholder Presentation (CP)

Rank	Vendor	Score	Overall rank
1	Epic Games	82.1	2
2	Chaos	81.0	1
3	D5 (Dimension 5)	67.9	3
4	Lumion (Act-3D)	51.2	4
5	Blender Foundation	3.6	5
6	Midjourney*	3.6	10

One of two personas where Chaos does not rank first. Twinmotion's free license for smaller firms and its presentation features carry Epic to the top. Lumion's 51.2 is its best persona score. The drop from fourth place to fifth, 47.6 points, is the steepest in the study.

Developer and Real Estate Marketing (DM)

Rank	Vendor	Score	Overall rank
1	Chaos	88.1	1
2	Epic Games	73.8	2
3	D5 (Dimension 5)	61.9	3
4	Lumion (Act-3D)	31.0	4
5	Blender Foundation	21.4	5

Rank	Vendor	Score	Overall rank
6	OTOY*	14.3	13

Chaos's highest score anywhere, driven by V-Ray and Corona. This is the persona where image quality is a revenue instrument, and the engines answer with offline render engines. OTOY's Octane appears in 9 of the 21 prompts here, and 9 of its 15 appearances overall come from this persona. Maxon matches its 14.3 in seventh, through Redshift.

Immersive and Interactive (IX)

Rank	Vendor	Score	Overall rank
1	Epic Games	69.0	2
2	Chaos	61.9	1
3	Autodesk	35.7	6
4	Arkio*	32.1	9
5	Shapspark*	28.6	11
6	D5 (Dimension 5)	20.2	3

The most distinct vendor set of the six. Unity (19.0) and SENTIO (17.9) follow in seventh and eighth. Arkio, SENTIO, and Resolve appear here and almost nowhere else: every one of Arkio's 11 appearances comes from this persona. Autodesk reaches third through Workshop XR and the products it acquired with The Wild. D5 and Lumion, strong in most other personas, fall away.

Studio Production Pipeline (SP)

Rank	Vendor	Score	Overall rank
1	Chaos	73.8	1
2	Epic Games	57.1	2
3	D5 (Dimension 5)	17.9	3
4	Blender Foundation	14.3	5
5	Maxon	9.5	8
6	Autodesk	8.3	6

The persona where the core table matters least. Studio pipeline answers are dominated by adjacent tools, render management (Deadline, Pulze, Royal Render), asset management (Connector, Anchorpoint, Perforce), and render farms, which is why scores below second place are so low. Lumion falls to eighth, with a score of 6.0.

Firm size. The top four do not change order in any of the three size bands. The only movement is below them: OTOY rises in the solo band, Unity in the studio and enterprise bands, and Adobe in the enterprise band. Firm size changes the price the engines quote far more than the names they give.

Summary of persona effects. Concept exploration is a different market from the other five, and immersive is a different market again. In the four personas in between, the top four vendors take almost every answer. For a vendor outside the top four, the persona view shows exactly where a foothold exists (concept for AI generators, immersive for VR platforms, production for renderers such as Octane) and where none does.

8. Established vendors and new entrants

The engines named 101 visualization products. That is more than most people in the industry could list, and it raises an obvious question: is the field being reshaped by newer, faster companies, or do the established vendors still own it?

We looked up the first public release year of every product in the core table and grouped them into four cohorts. Products launched in 2022 or later are split by who makes them, because several established vendors launched or acquired AI products in that period.

Cohort	Products	Above floor	ChatGPT	Gemini	AI Mode	Copilot	All engines
Established (released before 2015)	32	11	50.7%	53.2%	56.0%	53.2%	53.2%
Second wave (2015 to 2021)	26	6	33.6%	30.8%	26.1%	32.1%	30.6%
AI era, established vendor (2022 or later)	9	4	9.0%	5.5%	5.0%	7.9%	6.8%
AI era, new entrant (2022 or later)	34	6	6.7%	10.6%	12.9%	6.8%	9.5%

Percentages are each cohort's share of visualization product mentions on that engine.

The established products still carry half the answer. Twinmotion (first released 2005), V-Ray (2002), Lumion (2010), Unreal Engine (1998), and Corona (2009) are five of the seven most-named products. The two second-wave products in the top seven are Enscape (2015), now owned by Chaos, and D5 Render (2020). D5 is the one company in the top four that did not exist a decade ago, and its rise to third is the largest competitive shift in this category that the data can see.

The new entrants are many, and concentrated. 34 independent products launched since 2022 were named, a third of the core table, and almost all of them are AI image generators. Six clear the inclusion floor: LookX AI (22 prompts), Midjourney (16), Krea (13), Gendo (11), ArchiVinci (10), and Stable Diffusion (9). Their visibility sits almost entirely in one persona:

Persona	Prompts naming at least one new entrant (of 21)
Concept Exploration	19
In-Model Design Iteration	7
Client and Stakeholder Presentation	6
Developer and Real Estate Marketing	3

Persona	Prompts naming at least one new entrant (of 21)
Immersive and Interactive	1
Studio Production Pipeline	1

The disruption is real and it has a location. At the front of the design process, before a model exists, AI generators are the default answer. Once a building model exists, the engines send buyers back to the established renderers.

New entrants do better on the engines that read the live web. Independent AI-era products appear in 30 AI Mode prompts and 26 Gemini prompts, against 13 on ChatGPT and 13 on Copilot. That is the signature of content-driven visibility: products too new or too small to be well represented in training data, found through retrieval. Section 9 shows how much of that retrieval runs through listicles the newcomers published themselves.

The established vendors are buying and building their way into the AI era. Nine AI-era products come from established vendors: Veras (acquired by Chaos with EvolveLAB), D5 Lite, Lumion View, Autodesk Workshop XR, Adobe Firefly, Chaos Envision, Graphisoft AI Visualizer, SketchUp Diffusion, and Siemens Immersive Explorer. Veras alone is named in 36 prompts, more than any independent newcomer, and it is the main reason Chaos leads the concept exploration persona. The pattern in the data is incumbents absorbing the new category faster than new companies can build visibility in the old one.

The engines are shopping on price. One more signal in the answers bears on competitive pressure. The engines described Twinmotion as free in 57 answers, quoted D5 Pro at \$360 a year beside Enscape Solo at \$574.80 and Lumion Pro at \$1,149, and recommended \$20 to \$29 monthly AI subscriptions for concept work. ChatGPT told one solo architect that Twinmotion's free license for businesses under \$1 million in revenue is "an almost absurd advantage". Whatever a vendor's list price, the assistants frame it against a free or low-cost alternative, and they do so unprompted.

9. Source analysis: what is actually driving visibility

Coverage alone cannot distinguish a product the engines remember from a product the engines can find. That requires reading the mention data and the citation data together.

9.1 How the engines cite

Engine	Answers with sources (of 126)	Citations	Per prompt	Unique domains	Top 10 domains share	Vendor-owned share
ChatGPT	114	766	6.1	87	74.3%	97.7%
Gemini	84	228	1.8	113	34.2%	61.0%
AI Mode	117	1,097	8.7	300	42.8%	46.4%
Copilot	126	1,191	9.5	135	66.2%	94.0%

Engine	Answers with sources (of 126)	Citations	Per prompt	Unique domains	Top 10 domains share	Vendor-owned share
All engines	126	3,282	6.5	455	52.4%	76.6%

A citation is a unique URL within one answer. ChatGPT and Copilot concentrate on a small set of vendor domains. AI Mode spreads across 300 domains, most of them cited once.

9.2 The citation support rate

For every visualization vendor mention, we checked whether the same engine cited that vendor's own domain in the same answer.

Engine	Vendor mentions	Backed by own-domain citation	Support rate
Copilot	318	285	89.6%
ChatGPT	441	332	75.3%
AI Mode	427	164	38.4%
Gemini	444	59	13.3%
All engines	1,630	840	51.5%

By vendor, for the ranked vendors:

Vendor	Mention prompts (all engines)	ChatGPT	Gemini	AI Mode	Copilot	All engines
Chaos	388	86.7%	15.4%	53.4%	93.4%	60.1%
Epic Games	323	65.4%	3.2%	8.8%	79.7%	36.5%
D5 (Dimension 5)	233	78.5%	13.2%	43.4%	95.8%	59.2%
Lumion (Act-3D)	141	62.7%	18.8%	21.4%	90.0%	55.3%
Blender Foundation	44	40.0%	0.0%	10.5%	88.9%	27.3%
Autodesk	44	100.0%	30.0%	45.5%	92.3%	68.2%
LookX	42	62.5%	0.0%	0.0%	100.0%	35.7%
Maxon	29	50.0%	0.0%	100.0%	85.7%	55.2%
Arkio	27	100.0%	0.0%	0.0%	100.0%	55.6%
Midjourney	26	66.7%	11.1%	11.1%	100.0%	30.8%
Shapspark	25	71.4%	44.4%	50.0%	100.0%	64.0%
Unity Technologies	24	50.0%	0.0%	36.4%	71.4%	41.7%

Vendor	Mention prompts (all engines)	ChatGPT	Gemini	AI Mode	Copilot	All engines
OTOY	20	40.0%	0.0%	0.0%	0.0%	10.0%
Gendo	18	100.0%	40.0%	88.9%	100.0%	77.8%
Stability AI	16	0.0%	0.0%	0.0%	50.0%	6.2%
Krea	15	100.0%	0.0%	20.0%	100.0%	33.3%
SENTIO	15	100.0%	0.0%	75.0%	100.0%	73.3%
ArchiVinci	11	100.0%	n/a	60.0%	100.0%	81.8%
Vizcom	10	100.0%	37.5%	n/a	100.0%	50.0%

Mention prompts count one per prompt per engine, so the maximum is 504. "n/a" means the vendor was not named on that engine.

Two mechanisms. A vendor can be visible because the model learned it or because content exists and gets retrieved. ArchiVinci, Gendo, SENTIO, and Autodesk sit at the top of the support table: when they are named, their own sites are usually cited, which means they are being found. Stability AI, OTOY, and Blender Foundation sit at the bottom: they are named without supporting citations, which means they are being remembered. Among the big four, Epic has the lowest support rate by a wide margin.

Epic is the clearest recall case among the leaders. Epic is named in 93 Gemini prompts and 80 AI Mode prompts, and its own domains are cited in support of 3 and 7 of them. On Google's engines, Twinmotion and Unreal are visible almost entirely on reputation. That is a strong position today and an exposed one over time, because training data refreshes on a cycle no vendor controls, while retrieval can be influenced every week.

Chaos is strong on both. Chaos combines the highest coverage with support rates above 85 percent on ChatGPT and Copilot and 53.4 percent on AI Mode. It is both remembered and findable, and that combination explains the size of its lead.

9.3 Memory or retrieval: why legacy products hold their place

The support rate can be read product by product, and with launch years beside it, it answers a question every vendor in this category should care about: is a product being named because the engine remembers it, or because the engine found it?

Product	Launched	Mentions	Support, all engines	ChatGPT and Copilot	Gemini and AI Mode (mentions)	Reading
Lumion View	2025	14	100%	100%	100% (1)	Retrieval-driven
Autodesk Workshop XR	2023	23	87%	100%	67% (9)	Retrieval-driven
D5 Lite	2026	14	86%	100%	0% (2)	Retrieval-driven
ArchiVinci	2024	11	82%	100%	60% (5)	Retrieval-driven

Product	Launched	Mentions	Support, all engines	ChatGPT and Copilot	Gemini and AI Mode (mentions)	Reading
Gendo	2024	18	78%	100%	71% (14)	Retrieval-driven
SENTIO VR	2017	15	73%	100%	43% (7)	Retrieval-driven
Veras	2022	67	69%	100%	46% (39)	Retrieval-driven
Enscape	2015	232	66%	89%	41% (111)	Retrieval-driven
Shapspark	2013	25	64%	83%	46% (13)	Mixed
D5 Render	2020	232	59%	85%	28% (106)	Mixed
Arkio	2021	27	56%	100%	0% (12)	Mixed
Lumion	2010	140	55%	71%	20% (44)	Mixed
Redshift	2014	29	55%	69%	44% (16)	Mixed
V-Ray	2002	218	54%	88%	32% (130)	Mixed
Arnold	1998	14	50%	86%	14% (7)	Mixed
Corona	2009	133	50%	83%	30% (81)	Mixed
Vizcom	2021	10	50%	100%	38% (8)	Mixed
Chaos Vantage	2020	49	47%	88%	27% (33)	Leaning recall
Twinmotion	2005	254	43%	79%	7% (128)	Leaning recall
Unity	2005	24	42%	67%	27% (15)	Leaning recall
LookX AI	2023	42	36%	83%	0% (24)	Third-party retrieval
Krea	2022	15	33%	100%	9% (11)	Third-party retrieval
Midjourney	2022	26	31%	75%	11% (18)	Third-party retrieval
Blender Cycles / Eevee	2011	44	27%	71%	7% (30)	Recall
Unreal Engine	1998	159	25%	63%	6% (105)	Recall
Octane	2010	20	10%	33%	0% (14)	Recall
Stable Diffusion	2022	16	6%	25%	0% (12)	Recall

All 27 ranked products. Mentions count one per prompt per engine; the Gemini and AI Mode column shows the number of mentions on those two engines in brackets. Support means the same answer cited a domain owned by the product's vendor. Readings: 65 percent or more is retrieval-driven, 50 to 64 mixed, 35 to 49 leaning recall, below 35 recall. LookX, Midjourney, and Krea are labeled third-party retrieval: they are new products with little training history, named on Gemini and AI Mode almost without citations to their own sites, which points to the comparison lists those engines read (sections 9.5 and 9.6).

The column that separates the two mechanisms is the last numeric one. ChatGPT and Copilot cite the vendor's own site for almost every product they name, old or new, so their support rates are high across the board. Gemini and AI Mode cite far less and far more widely, so a product that is named there without a supporting citation is being

named mostly from memory.

By cohort, weighted by mentions:

Cohort	Ranked products	Support, all engines	Gemini and AI Mode
Released before 2015	11	44.6%	18.9%
Released 2015 to 2021	6	60.9%	32.5%
Released 2022 or later	10	58.7%	34.1%

The oldest products are the ones the engines remember. On Google's engines, the products released before 2015 are backed by their vendor's own site in fewer than one mention in five. Twinmotion (7 percent), Unreal Engine (6 percent), Blender (7 percent), and Octane (0 percent) are named there almost entirely from what the models learned in training. The answers with no sources at all make the same point. Gemini gave 42 answers without citing anything, and products released before 2015 took 57.2 percent of the visualization mentions in those answers, against 51.4 percent in its sourced answers. Unreal Engine alone appears in 26 of the 42.

This helps explain why legacy vendors rank well without much AEO effort. A product that has been discussed in forums, tutorials, documentation, and reviews for fifteen or twenty years is embedded in the training data of every model, and the engines reach for it without needing to find anything new. The newer products earn their place the other way, by being found: Gendo, ArchiVinci, Workshop XR, and Veras are named with a supporting citation in most of their mentions.

Memory is a strong position and a decaying one. Training data refreshes on a cycle no vendor controls. Retrieval indexes update continuously, and every week of published content from a newer competitor adds to what the engines can find. A legacy product whose visibility rests on recall has no defense against a competitor that systematically out-publishes it in the retrieval layer, and the retrieval engines are exactly where the AI newcomers are concentrating their content. The vendors that hold both positions, such as Chaos with Enscape and V-Ray, are the ones best placed to keep their rank.

9.4 Where the engines disagree

The gap between a product's coverage on the two Google engines and its coverage on ChatGPT, for ranked products with the largest gaps in each direction.

Product	ChatGPT	Gemini	AI Mode	Google average minus ChatGPT
Lumion	53.2%	12.7%	22.2%	-35.7
D5 Render	62.7%	42.1%	42.1%	-20.6
Enscape	57.1%	44.4%	43.7%	-13.1
Twinmotion	59.5%	52.4%	49.2%	-8.7
Corona	26.2%	33.3%	31.0%	+6.0

Product	ChatGPT	Gemini	AI Mode	Google average minus ChatGPT
Blender Cycles / Eevee	4.0%	8.7%	15.1%	+7.9
V-Ray	42.9%	52.4%	50.8%	+8.7
Unreal Engine	20.6%	46.8%	36.5%	+21.0

Lumion's gap is the largest in the study in either direction. It is fourth on ChatGPT at 53.2 percent and falls to 12.7 percent on Gemini, with only three own-domain citations on that engine. ChatGPT reads Lumion's own site and names Lumion. Gemini, answering mostly from training, rarely does. AI Mode, retrieving from third-party lists, lands in between.

The mirror image is Unreal Engine, which the Google engines treat as a mainstream architectural visualization answer and ChatGPT treats as a specialist one. V-Ray and Corona lean the same way. On this evidence, Google's engines favor the production-grade engines and ChatGPT favors the tools that are easiest for an architect to adopt.

9.5 What kinds of pages get cited

Page type	ChatGPT	Gemini	AI Mode	Copilot	All engines
Listicle or comparison	0.8%	38.2%	35.0%	0.5%	14.7%
Product or other page	38.0%	29.4%	36.4%	41.9%	38.3%
Documentation or support	42.7%	8.8%	6.2%	33.2%	24.7%
Pricing or licensing	10.6%	0.9%	0.8%	11.8%	7.1%
Blog or article	5.2%	15.8%	19.8%	1.4%	9.4%
Homepage	2.7%	7.0%	1.8%	11.1%	5.8%

The engines split cleanly in two. ChatGPT and Copilot cite documentation, product, and pricing pages, and almost never a "best of" list. Gemini and AI Mode draw more than a third of their citations from listicles and comparison articles.

The listicles are not neutral. Of the 483 listicle citations in the dataset, 222 point to pages published by software vendors. The most-cited listicles:

Page	Publisher type	Citations	Prompts
maxon.net: best architectural visualization software	Vendor (Maxon)	33	32
maverickframe.com: best architectural rendering software	Visualization studio	30	29
rendershop.ai: best architectural visualization software	Vendor (Rendershop)	29	26
maxon.net: best architectural rendering software	Vendor (Maxon)	20	20
meltplan.com: best 3D rendering software for architects in 2026	Software company	16	16

Page	Publisher type	Citations	Prompts
d5render.com: best architectural visualization software	Vendor (D5)	13	13
maverickframe.com: best AI rendering tools	Visualization studio	13	13
gendo.ai: top 10 AI rendering software for architects 2026	Vendor (Gendo)	11	11
moldaspace.com: best architectural rendering software	Vendor (MoldaSpace)	10	10
visioplan.io: AI rendering tools for architects compared	Software company	9	9

Maxon's two articles were cited 53 times, and in only 19 of those citations did the answer name Maxon. Maxon is supplying the reference list that AI Mode uses to answer questions about everyone else. Rendershop, a product first released in February 2026, has a list cited in 26 prompts, while Rendershop itself is named in one.

9.6 Incumbents and challengers in the citation layer

Group	Engine	Share of visualization vendor mentions	Share of visualization vendor-owned citations	Own-domain listicle citations
Big four (Chaos, Epic, D5, Lumion)	ChatGPT	75.3%	69.8%	3
	Gemini	59.9%	32.0%	7
	AI Mode	61.8%	33.4%	23
	Copilot	70.1%	62.1%	1
AEC-built AI rendering challengers	ChatGPT	4.8%	4.4%	0
	Gemini	8.1%	39.8%	23
	AI Mode	11.0%	32.7%	67
	Copilot	6.6%	3.4%	0

On Gemini and AI Mode, the AI rendering challengers hold a far larger share of the vendor-owned citation layer (32.7 to 39.8 percent) than of the mentions (8.1 to 11.0 percent). Half of their AI Mode citations, 67 of 135, come from their own listicles. They have found the one channel where publishing moves the needle quickly, and they are using it. The conversion into mentions is still modest, but the engines are reading them.

9.7 Who owns the citation layer

Domain	Type	Prompts	Citations	Engines
chaos.com	Vendor	109	536	All four
d5render.com	Vendor	75	247	All four
autodesk.com	Vendor	68	222	All four
twinmotion.com	Vendor	56	132	All four

Domain	Type	Prompts	Citations	Engines
youtube.com	Video	56	112	Gemini, AI Mode
maxon.net	Vendor	52	87	All four
maverickframe.com	Visualization studio	52	76	Gemini, AI Mode
lumion.com	Vendor	51	132	All four
epicgames.com	Vendor	49	135	All four
rendershop.ai	Vendor	26	30	Gemini, AI Mode
myarchitectai.com	Vendor	24	30	Gemini, AI Mode
superrendersfarm.com	Render farm	22	38	Gemini, AI Mode
reddit.com	Community	22	25	ChatGPT, AI Mode
meltplan.com	Software company	21	21	Gemini, AI Mode
moldaspace.com	Vendor	20	24	AI Mode

Three observations for anyone trying to influence these systems.

Trade media has left the most influential slot in the category open. Trade media and publications account for 1.5 percent of all citations, and review, directory, and analyst sites for 0.8 percent. Architizer, the most-cited trade title, appears in 11 prompts, all on AI Mode. AEC Magazine appears in 5. Yet the kind of page trade media has always owned, the independent comparison of tools, is exactly what Gemini and AI Mode read most: listicles and comparisons are more than a third of their citations. Of the 483 listicle citations in the study, trade publications wrote 21. Software vendors wrote 222, visualization studios 104, and other software companies 94. The engines are reading buying guides, and the industry's independent voices are hardly writing any of them. A trade title that published current, structured, well-sourced comparisons of visualization tools, with prices, license terms, and host application support stated plainly, would be competing against vendor-written lists for the same citations, with a credibility advantage the vendors cannot claim. The opening is on Google's engines. ChatGPT and Copilot cite almost nothing but vendor sites, and trade media would have to earn its way in there the slow way, through training data.

YouTube is a major source on AI Mode and nowhere else. AI Mode cited YouTube 111 times across 55 prompts, 10.1 percent of its citations. ChatGPT and Copilot never cited it. For a category that sells visually, video tutorials and walkthroughs are a real retrieval surface on Google.

Visualization studios are cited as experts. Studios such as Maverickframe account for 14.1 percent of AI Mode's citations. Their "best software" guides reach buyers at the moment of choice, which makes studio relationships a visibility channel as well as a sales one.

10. The buyer's view: which engine to trust

Everything above treats the engines as a channel vendors want to be seen in. Buyers see the same answers from the other side, and for them the questions are simpler. When an assistant states a fact, is it right? And when it gives a shortlist, is the shortlist complete?

10.1 Price and license accuracy

Prices and license terms are the easiest facts to check, and the ones a buyer is most likely to act on. We extracted every price and eligibility claim the engines made for six products with publicly listed prices (Enscape, Twinmotion, Lumion, V-Ray, Corona, and Veras) and compared each with the vendor's own pricing page on September 21, 2026. A claim counts as matching if it is within 5 percent of a current list price, or states the correct eligibility threshold.

Engine	Answers quoting any price (of 126)	Price claims checked	Match current list price	Off by more than 5%	Wrong eligibility threshold	Could not verify
ChatGPT	51	38	35	0	0	3
Copilot	44	33	33	0	0	0
Gemini	36	26	12	6	1	7
AI Mode	13	4	3	0	1	0

A checked claim is one product priced in one answer. "Could not verify" covers monthly-billing prices and license types the vendor pages did not display.

ChatGPT and Copilot quote the price list. They give exact figures to the cent (ChatGPT: "\$574.80/year (\$47.90/month billed annually)" for Enscape Solo), distinguish named from floating licenses, and multiply correctly for team seats. Every verifiable claim from either engine matched the vendor's current page. This fits what section 9 showed: both engines cite pricing and licensing pages directly.

Gemini approximates, and the approximations drift. Gemini's prices are almost always rounded and prefixed with a tilde, and 7 of its 19 verifiable claims missed. It priced Enscape at "~\$500/yr" (list price \$574.80) and Corona at "~\$515 / seat / year" (Solo is \$414, Premium \$562.80). In one answer it gave Enscape as "~\$575/year (~\$85/month)", pairing the annual price with a monthly figure that does not match its \$47.90 monthly equivalent. Several of its prices came from third-party review blogs.

The eligibility errors matter most. A freelance designer asked for concept visuals under \$50 a month. Gemini told them Twinmotion is "100% free for individuals earning under \$100,000/year gross revenue". Asked which tools have a free tier that permits commercial use, AI Mode also gave the threshold as under \$100,000. Epic's published threshold is \$1 million, ten times higher. In the same Gemini answer, the free tier of D5 Render was suggested for commercial pitch work, although D5's Community license is restricted to non-commercial use. A buyer who acted on the first error would pay for software they could have used free. A buyer who acted on the second would be using software outside its license.

AI Mode rarely commits to a price at all. Only 13 of its 126 answers quoted any price, against 51 for ChatGPT. When it did, three of its four checked claims were accurate.

Prices are only one kind of fact, but the pattern matches everything else in this report. The engines that read vendor pages report vendor facts accurately. The engines that read listicles and answer from training report the market's general shape well and its specifics less reliably.

10.2 Completeness and depth

Accuracy is half of a useful answer. The other half is whether the answer gives the buyer enough to decide: a real shortlist, the reasoning behind it, and evidence the buyer can check.

Measure	ChatGPT	Gemini	AI Mode	Copilot
Median answer length (words)	612	374	594	529
Software products named per answer	9.6	10.5	9.7	8.0
Visualization products named per answer	4.5	4.7	4.3	3.0
Answers naming one visualization product or none (of 126)	13	5	9	37
Answers with no sources (of 126)	12	42	9	0
Mentions backed by the vendor's own site	75.3%	13.3%	38.4%	89.6%
Verifiable price claims matching list price	35 of 35	12 of 19	3 of 4	33 of 33

The visualization vendor support rates are the section 9.2 figures. Some single-product answers are legitimate, for example studio pipeline prompts where a render manager is the right answer.

ChatGPT gives the most complete answer. It writes the longest answers, names close to ten products in each, backs three quarters of its vendor mentions with the vendor's own page, and got every checkable price right. A representative example: asked for concept tools under \$50 a month, it recommended one product at \$29, ranked two alternatives with prices, and flagged that the third was above budget on monthly billing.

Copilot is as accurate as ChatGPT and much narrower. Every price it quoted matched, and it cites a vendor source for nine vendor mentions in ten. But 37 of its 126 answers named one visualization product or none, nearly three times ChatGPT's figure. Asked which software publishes interactive walkthroughs from a polygonal model, Copilot answered in 56 words: "The software is Shapspark." For a buyer who wants options, a correct single answer is still an incomplete one.

AI Mode gives a broad view built from other people's lists. It names as many products as ChatGPT and cites more domains than any other engine, but a third of its citations are "best software" articles, many published by vendors, and it quoted a price in only 13 answers. On the same under-\$50 question, AI Mode recommended three tools and described their prices only as "comfortably within your budget", "affordable", and "accessible within your budget", without a single tool price. It is useful for discovering names and weak on the specifics a purchase requires.

Gemini is the engine we would not rely on for a purchase decision. It names the most products and shows the least evidence for them. A third of its answers cite nothing, it cites the named vendor's own site for only 13.3

percent of vendor mentions, and it produced the largest share of pricing errors, including a Twinmotion eligibility threshold ten times too low and a suggestion to use D5's non-commercial tier for paid work. It also returned the one answer in the study that consisted of a board of images with no written recommendation. A quarter of the products Gemini names in a given answer (24.8 percent) are named by no other engine for the same prompt, against 16.6 to 20.5 percent for the others. Some of that is breadth. Without sources, a buyer cannot tell which part.

Our ranking for buyers in this category: ChatGPT first, Copilot second, AI Mode third, Gemini fourth.

ChatGPT combines depth, breadth, evidence, and accuracy. Copilot matches its accuracy with a fraction of its breadth. AI Mode offers breadth with secondhand evidence. Gemini offers breadth with little evidence and the most errors.

10.3 What this means for a buyer

Start with ChatGPT, and check its shortlist against a second engine. The shortlist you get depends on which engine you ask: Lumion appears in half of ChatGPT's answers and one in eight of Gemini's. If you use Copilot, ask it explicitly for alternatives.

Check prices and license terms on the vendor's own page. Any price from any engine should be treated as a starting point, and a price from Gemini as an estimate.

Ask for sources, and look at them. Gemini gave no sources for a third of its answers. When the source is a "best software" list, check who published it: in this study, nearly half of all listicle citations pointed to pages written by software vendors.

Be most careful with license eligibility. Free tiers, revenue thresholds, and commercial-use rules produced the most consequential errors in the dataset, and they are the claims a small practice is most likely to rely on.

11. What this means

For Chaos. The visibility contest is won. Chaos leads on every engine, combines strong recall with strong retrieval on ChatGPT, Copilot, and AI Mode, and is the only vendor to rank first or second in every persona. Veras has extended that lead into the one segment where newcomers were winning. The exposure is narrow: Copilot, where Chaos drops to 60.3 percent, and the client presentation and immersive personas, where Epic leads.

For Epic Games. Second place is secure and rests on two products and a price. Twinmotion is the most-named visualization product in the study, and it is described as free in 57 answers. The weakness is evidence: on Gemini and AI Mode, Epic is cited in support of fewer than one mention in ten. Epic's position on Google's engines is almost pure recall.

For D5. D5 is the one company in the top four founded within the last decade, and it is named in 89 of 126 prompts from a single renderer. It is third on every engine and is named in all 21 design iteration prompts. The engines consistently present it as the value choice against Enscape, which is a strong position until a cheaper option is framed against it the same way.

For Lumion. Lumion's visibility depends on ChatGPT more than any other vendor's. At 53.2 percent on ChatGPT and 12.7 percent on Gemini, it is the vendor most exposed to a buyer's choice of assistant. It is also strongest in the

personas where its product fits best, client presentation and design iteration, and weakest in studio production, then in concept exploration, where AI generators have taken the answer, and immersive work. The fix is in the retrieval layer, where Lumion is cited in support of about one Google mention in five.

For the AI newcomers. The concept stage is the one place where newcomers win answers, with products launched since 2022. The retrieval engines are reading their content, much of it self-published comparison lists. The open question is whether that visibility can move downstream of the concept stage, where the established renderers still hold nearly every answer, or whether acquisition by an incumbent, as with Veras, is the more likely path to scale.

For trade media and independent reviewers. The engines are answering buying questions from comparison articles, and vendors write almost half of them. That is the largest open position in this dataset. A publication that owns the independent buying guide for this category would supply the evidence Gemini and AI Mode are currently taking from vendors and studios, and would do so for every persona at once.

For everyone. The design authoring tool frames the answer, the four engines answer from different evidence, and a vendor measuring its AI visibility on one engine is measuring a quarter of the picture.

12. Limitations

Collection window. Responses were captured between September 12 and 21, 2026, with each engine's prompts run in a window of one to four days. These systems change continuously and a rerun would produce different numbers.

Single run per prompt per engine. No repeat sampling, so variance within an engine is not measured. Some of the long tail is likely noise.

Model version specificity. Results describe GPT-5.6 Sol High, Gemini 3.6 Flash Extended, Google AI Mode, and Copilot Think Deeper as they behaved in September 2026. Other tiers of the same products may behave differently.

Lexicon extraction. Products were matched by curated alias patterns. A product named once under an unusual spelling could be missed. Products named in fewer than five prompts should be treated as directional.

Launch years. Product launch years in section 8 come from vendor sites, press coverage, and public databases. Years for several small AI tools are estimates, and a one-year error would not move any product between the pre-2022 and post-2022 cohorts except at the boundary.

Source classification. Domain types were classified by hand for the most-cited domains and by rule for the rest. Page types were classified from URL patterns.

Citation support is a proxy. A mention accompanied by a citation to the vendor's own domain suggests retrieval-driven visibility. It does not prove causation.

Price checks. Prices were checked against vendor pages on September 21, 2026, up to nine days after the answers were collected. A price change in that window would count against the engine. D5 does not display prices in a form we could capture, so D5 claims were not scored.

No prompt volume data. This study measures how engines answer these questions. It does not measure how often real buyers ask them.

Appendix A: Core products named in fewer than eight prompts

Reported as an unordered set: 3DVista, Aareas Interactive, Adobe Firefly, AI Architectures, Architect Render, ArchLab, Archmaster, ArkoAI, Awen AI, Bentley LumenRT, Bricks & Pixels, Chaos Envision, CloudPano, ComfyUI, DisplaySweet, Fenestra, FLUX, FStorm, Fuzor, Graphisoft AI Visualizer, Gravity Sketch, IdealTwin, Ideogram, IntuiFace, IrisVR Prospect, KeyShot, Krita AI Diffusion, Kubity, Kuula, Leonardo.ai, LuxCoreRender, Marmoset Toolbag, Massing Labs, Maxwell Render, Mindesk, MoldaSpace, MyArchitectAI, NVIDIA Omniverse, Panoee, PromeAI, PromptRender, Rendair AI, Rendershop, Rendervi, Renderwow, ReRender AI, Resolve, Rhino Render, RunDiffusion, Runway, ShapesXR, Siemens Immersive Explorer, SimLab VR, Simplio3D, Sketchfab, SketchUp Diffusion, SmartPixel, Spacely AI, TechViz, The Wild, Thea Render, TouchDesigner, Trezi, Vectary, Ventuz, Verge3D, Vinode, Visoid, Visualizee.ai, VizBase, Vizera, Vrex, Woonmodule, ZeroLight.

Appendix B: Adjacent products by category

Design authoring (20): Revit, Rhino, SketchUp, Archicad, Vectorworks, Allplan, AutoCAD, Autodesk FormIt, Bonsai (BlenderBIM), BricsCAD, Cedreo, Chief Architect, Foyr Neo, Live Home 3D, MicroStation, Planner 5D, RoomLab, Roomle, Snaptrude, Tekla.

Modeling and DCC (8): 3ds Max, Blender, Cinema 4D, Maya, Houdini, CATIA, Creo, SolidWorks.

Asset libraries (18): Chaos Cosmos, Quixel Megascans / Fab, Forest Pack / RailClone, Adobe Substance 3D, Anima (AXYZ), Blendkit, Chaos Scans, GlobePlants, Greyscalegorilla Plus, GrowFX, Humano 3D, Maxtree, PGSkies, Poliigon, Poly Haven, Renderpeople, Scatter5, VizPeople.

Digital asset management and version control (14): Connector, Git / Git LFS, Alienbrain, Anchorpoint, Autodesk Vault, Daminion, Diversion, Flow Production Tracking (ShotGrid), ftrack, iMeshh Asset Manager, KStudio Project Manager, Mudstack, Perforce, Pulse Scene Manager.

Post-production and editing (13): Photoshop, After Effects, CapCut, DaVinci Resolve, Illustrator, InDesign, Krita, Lightroom, Nuke, Premiere Pro, Procreate, SUPIR, Topaz.

Coordination and model review (12): Autodesk Construction Cloud, Navisworks, Revizto, Bentley iTwin, BIMcollab, Dalux, Graphisoft BIMcloud, Novorender, Procore, Solibri, Trimble Connect, VIM.

Render management (11): Deadline, Backburner, CrowdRender, Flamenco, Maxon Team Render, OpenCue, Pulse RenderFlow, Qube!, RenderPal, Royal Render, Tractor.

Generative design and feasibility (9): Autodesk Forma, BuildAmbit, Carve, Dynamo, Finch3D, Hypar, Maket, Rayon, TestFit.

Streaming and sharing (9): Twinmotion Cloud, Pixel Streaming, Arcware, Autodesk Viewer, Furioos, Graphisoft BIMx, Lumion Cloud, PureWeb, Vagon Streams.

File sync and storage (8): Dropbox, Egnyte, Google Drive, LucidLink, OneDrive, Panzura, Resilio Sync, SharePoint.

Cloud rendering and render farms (7): Chaos Cloud, Autodesk Cloud Rendering, Drop&Render, Fox Renderfarm, GarageFarm, RebusFarm, Super Renders Farm.

Connectors and live links (5): Datasmith, Elefront, Pixyz, Speckle, Unity Reflect.

Kiosk and device management (5): ArborXR, KioWare, ManageXR, Meta Quest for Business, SiteKiosk.

Presentation and graphics (5): Canva, Figma / FigJam, Marq, Miro, Prezi.

Review and approval (2): ReviewStudio, Ziflow.

Appendix C: Multi-product vendors, per-product detail

Product	Equal weight	ChatGPT	Gemini	AI Mode	Copilot	Union
Enscape	46.0	57.1%	44.4%	43.7%	38.9%	85
V-Ray	43.3	42.9%	52.4%	50.8%	27.0%	85
Corona	26.4	26.2%	33.3%	31.0%	15.1%	55
Veras	13.3	11.1%	17.5%	13.5%	11.1%	36
Chaos Vantage	9.7	7.9%	18.3%	7.9%	4.8%	30
Chaos Envision	0.6	1.6%	0.0%	0.0%	0.8%	3
Chaos combined	77.0	83.3%	82.5%	81.7%	60.3%	120
Twinmotion	50.4	59.5%	52.4%	49.2%	40.5%	91
Unreal Engine	31.5	20.6%	46.8%	36.5%	22.2%	71
Sketchfab	1.0	0.8%	1.6%	0.8%	0.8%	3
Epic Games combined	64.1	64.3%	73.8%	63.5%	54.8%	105
D5 Render	46.0	62.7%	42.1%	42.1%	37.3%	89
D5 Lite	2.8	7.9%	0.8%	0.8%	1.6%	12
D5 combined	46.2	62.7%	42.1%	42.1%	38.1%	89
Lumion	27.8	53.2%	12.7%	22.2%	23.0%	73
Lumion View	2.8	7.9%	0.0%	0.8%	2.4%	11
Lumion combined	28.0	53.2%	12.7%	22.2%	23.8%	73
Autodesk Workshop XR	4.6	6.3%	3.2%	4.0%	4.8%	8
Arnold	2.8	0.8%	3.2%	2.4%	4.8%	11

Product	Equal weight	ChatGPT	Gemini	AI Mode	Copilot	Union
IrisVR Prospect	2.0	2.4%	2.4%	2.4%	0.8%	6
The Wild	1.2	1.6%	1.6%	0.8%	0.8%	4
Autodesk combined	8.7	7.9%	7.9%	8.7%	10.3%	20

Combined figures are the union of prompts naming any of the vendor's products. They are smaller than the sum of the rows.

Appendix D: Launch years for ranked products

Product	First public release	Cohort
Unreal Engine	1998	Established
Arnold	1998	Established
V-Ray	2002	Established
Twinmotion	2005	Established
Unity	2005	Established
Corona	2009	Established
Lumion	2010	Established
Octane	2010	Established
Blender Cycles / Eevee	2011	Established
Shapespark	2013	Established
Redshift	2014	Established
Enscape	2015	Second wave
SENTIO VR	2017	Second wave
D5 Render	2020	Second wave
Chaos Vantage	2020	Second wave
Arkio	2021	Second wave
Vizcom	2021	Second wave
Veras	2022	AI era, established vendor (acquired by Chaos 2025)

Product	First public release	Cohort
Midjourney	2022	AI era, new entrant
Stable Diffusion	2022	AI era, new entrant
Krea	2022	AI era, new entrant
LookX AI	2023	AI era, new entrant
Autodesk Workshop XR	2023	AI era, established vendor
Gendo	2024	AI era, new entrant
ArchiVinci	2024	AI era, new entrant
Lumion View	2025	AI era, established vendor
D5 Lite	2026	AI era, established vendor

Launch years for all 101 core products are included in the underlying dataset.

Appendix E: Prompt set and underlying data

The complete set of 126 prompts, the six persona definitions, the locked product lexicon, and the full mention and source datasets are available at no charge on request to info@aeowrangler.ai.

AEO Wrangler is an independent consultancy specializing in Answer Engine Optimization. This index is editorially independent and carries no paid placement.