AN AUGMA INTELLIGENCE STUDY

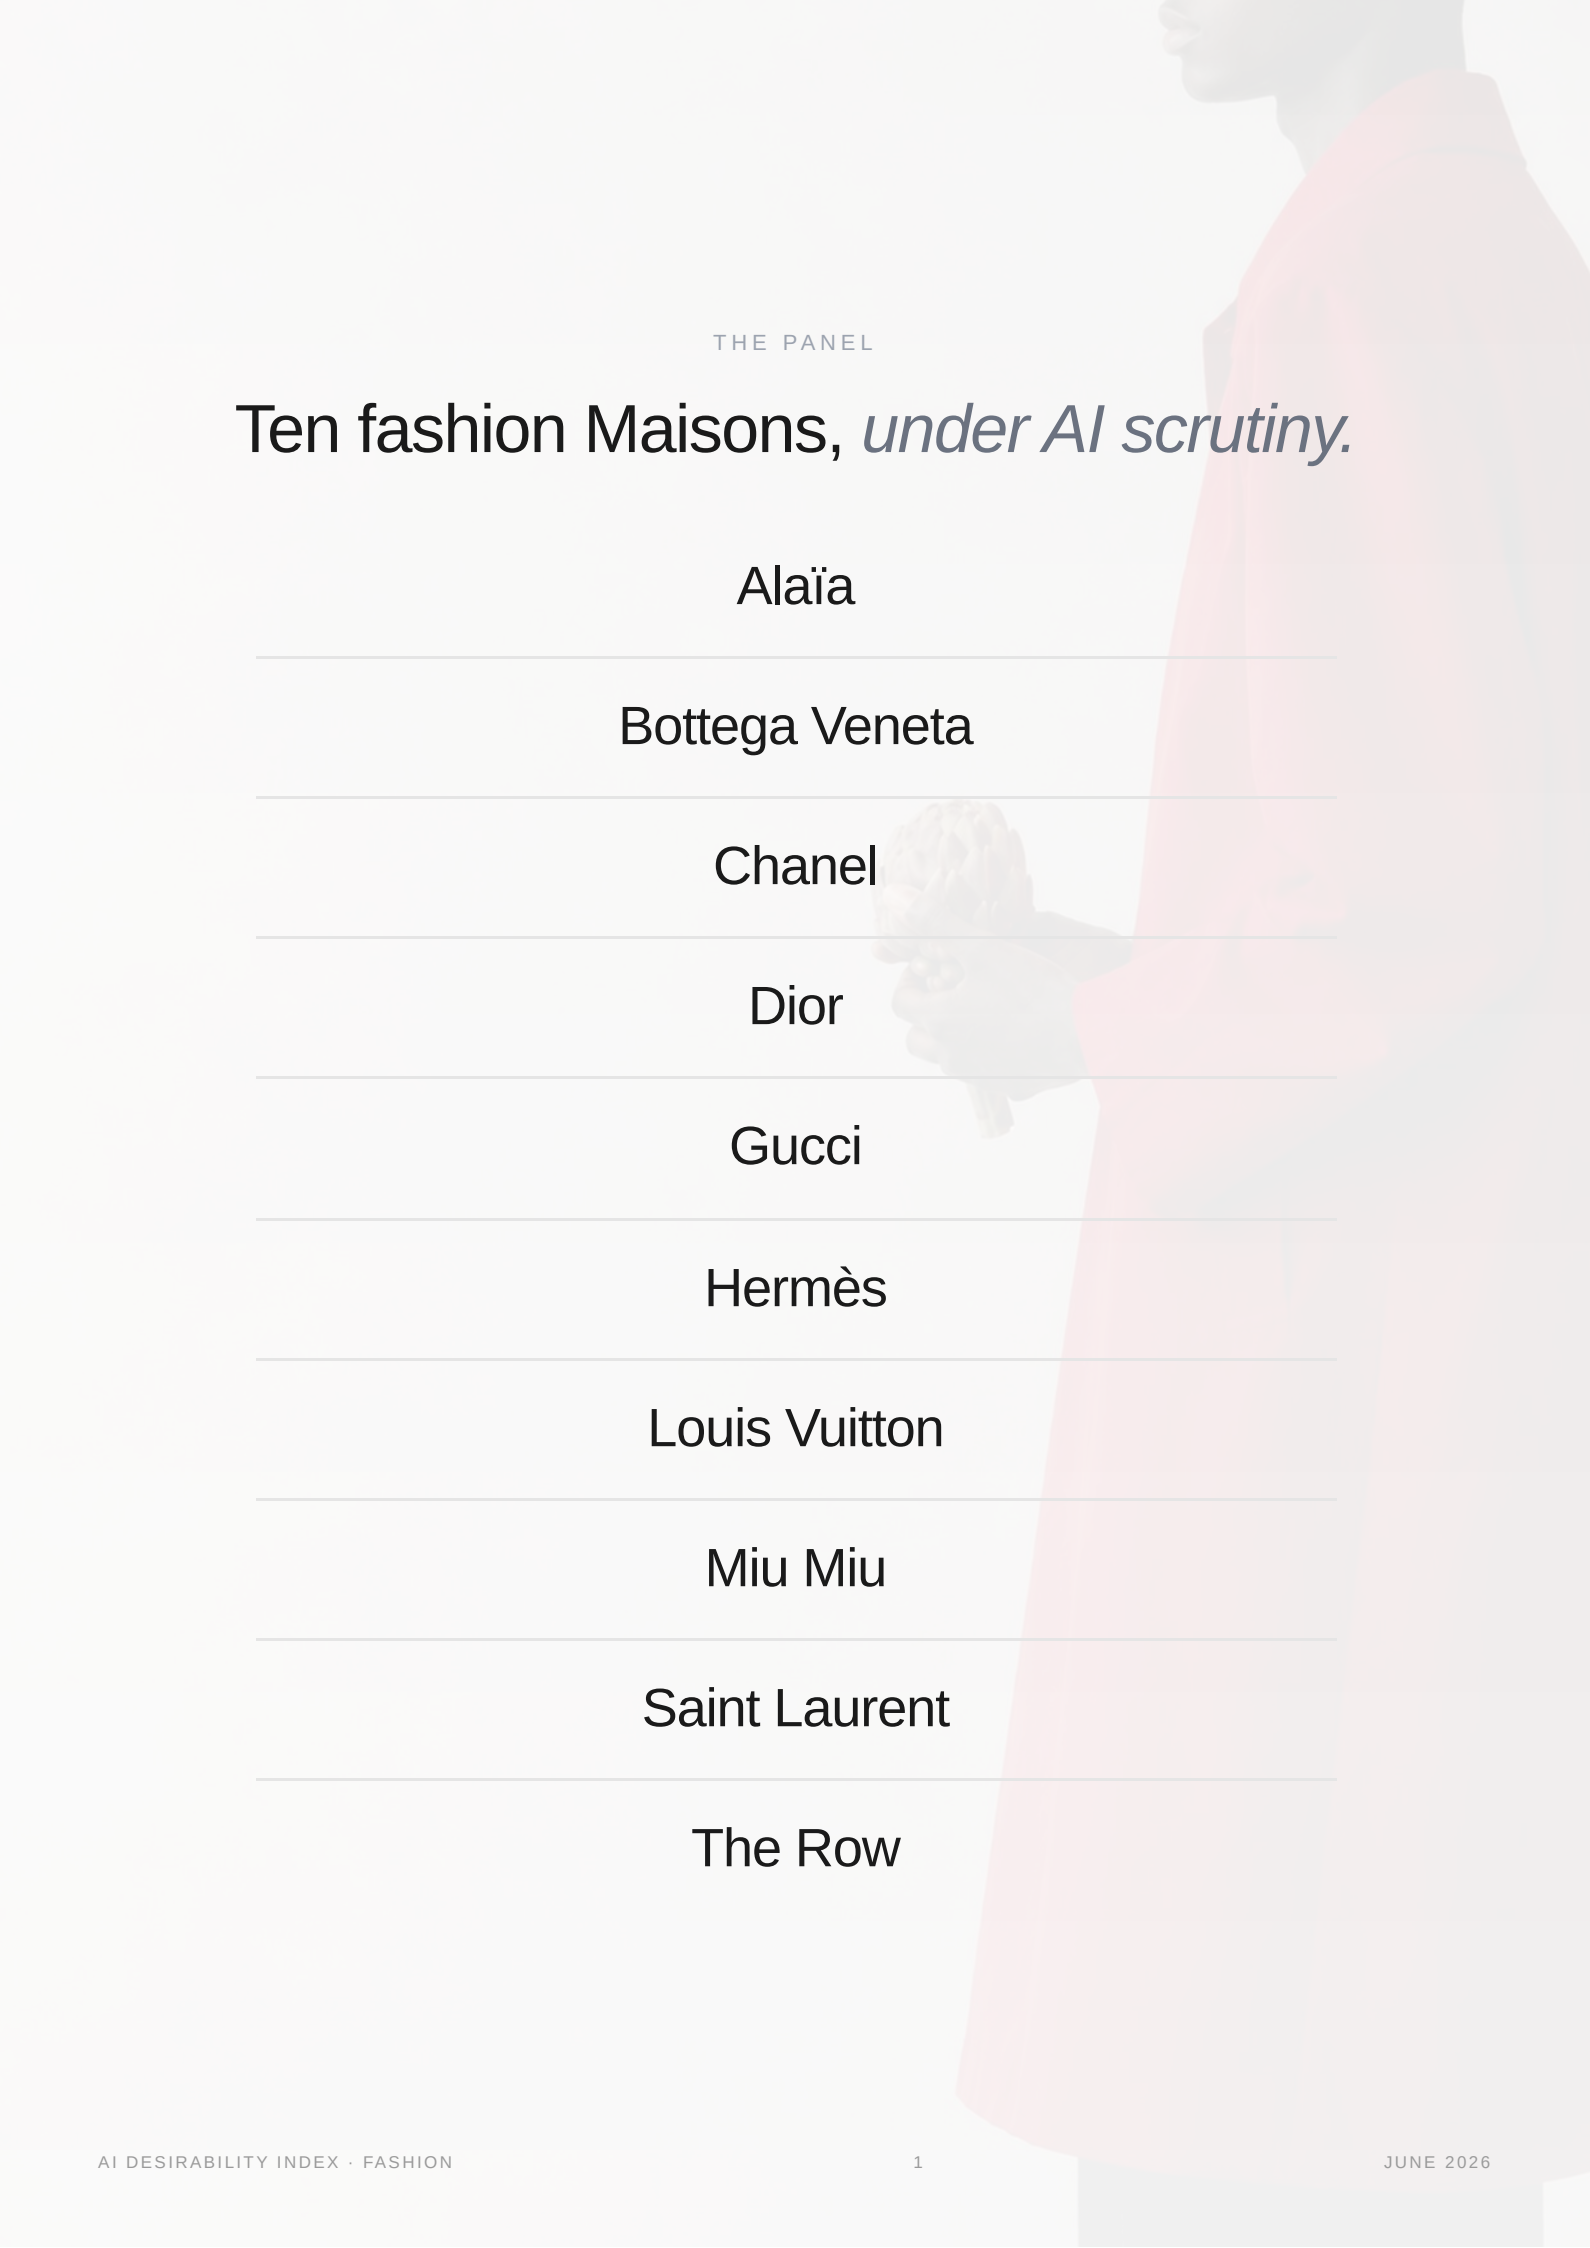
AI Desirability Index

FASHION

How five AI models perceive, surface, and narrate ten luxury fashion Maisons.

MAISONS	AI MODELS	RESPONSES	DIMENSIONS
10	5	2,400	6





THE PANEL

Ten fashion Maisons, *under AI scrutiny.*

Alaïa

Bottega Veneta

Chanel

Dior

Gucci

Hermès

Louis Vuitton

Miu Miu

Saint Laurent

The Row

EXECUTIVE SUMMARY

Visibility is not *desirability*. And neither is *accuracy.*

Our March AI Visibility Index measured whether AI surfaces the great Maisons at all. The Fragrance edition looked at how it values them once they surface. This Fashion edition takes on a question neither of those had to: when AI describes a Maison with confidence, is it describing the Maison that exists today? We read ten fashion Maisons on two axes. **AI Presence** is how reliably they surface. **Cultural Desirability** is the heritage, savoir-faire and cultural authority the models attribute to them. We then audit a third thing the desirability score leaves alone, which is whether the facts are current.

We ran 2,400 responses across Claude, ChatGPT, Perplexity, Gemini and Google AI Overviews. The strongest pattern is about timing. In 82 mentions, the models name an artistic director who has already left the Maison, and those errors land on the Maisons whose creative direction changed most recently and most publicly. AI is most fluent about the Maisons in transition, and most confidently wrong about them too.

82

Mentions that name an outdated artistic director, anchored to who was in post on the collection date

33

Of those fall on Bottega Veneta alone, the Maison the five models describe most consistently of all

31 pts

Desirability gap between the heritage mega-Maisons (73.2) and the contemporary Maisons (41.9)

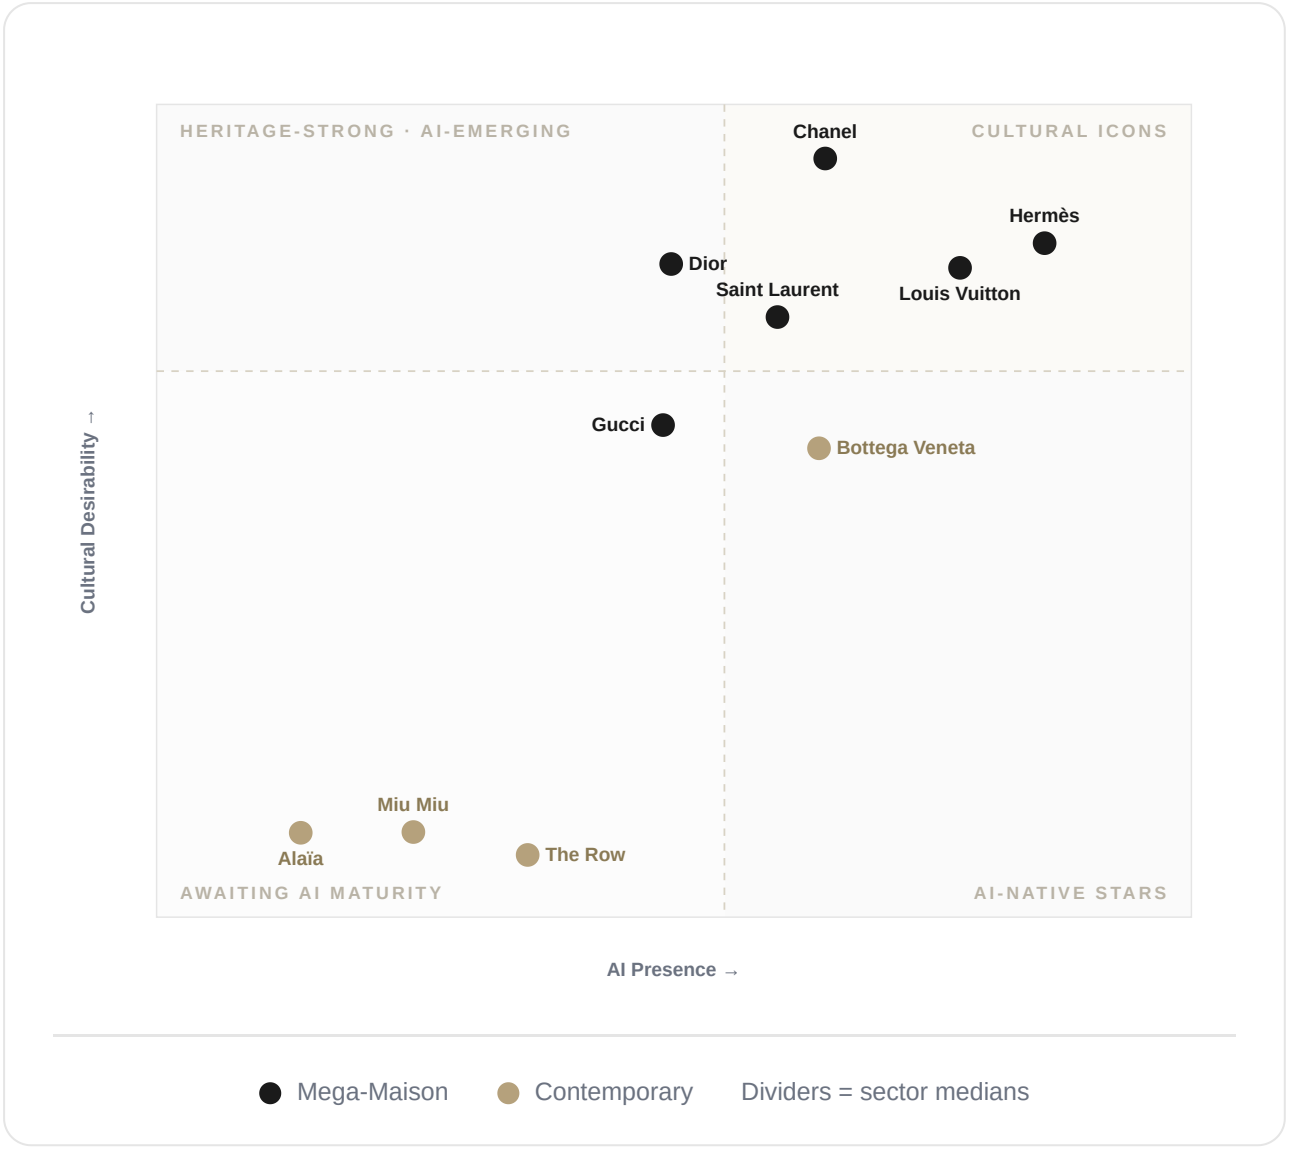
0%

Google AI Overviews trigger rate in French, against 52.5% in English

The director-staleness audit is evidentiary and kept strictly separate from desirability scoring. Full methodology on the final page.

Ten Maisons, *four territories.*

The vertical divider is the sector median for AI Presence (50.2); the horizontal divider, the median for Cultural Desirability (67.0). Four zones emerge. Heritage alone does not predict who lands where.



Cultural Icons lead on both axes. Heritage-Strong / AI-Emerging are revered but under-surfaced. AI-Native Stars surface strongly yet read lighter on desirability. Awaiting AI Maturity trail on both. Bottega Veneta is the only contemporary Maison to break the presence median.

FIVE READINGS

What the models *tell us.*

01

AI is most confident where it is most outdated.

Bottega Veneta is the Maison the five models agree on most: its cross-model standard deviation, 2.3, is the lowest in the panel, and its Brand Knowledge score, 87.6, ranks second of ten. The models speak about it fluently and in unison. They also, in 33 separate mentions, name a creative director who has moved on. That is more than any other Maison, and 40% of every such error in the study.

> *The Maisons AI describes most confidently are often not the ones it describes most accurately.*

Dior follows with 27 such mentions. Together, Bottega and Dior account for 60 of the 82 stale-director mentions, and they are precisely the two Maisons whose artistic direction changed most recently and most publicly. The pattern is structural, not anecdotal: AI models lag on luxury creative governance, and the lag concentrates exactly where the conversation is loudest.

02

Presence and desirability are not the same crown.

Hermès leads the panel on AI Presence (67.2) but ranks second on Cultural Desirability, behind Chanel. Chanel leads on Cultural Desirability (81.3), the highest in the panel, yet ranks third on presence (55.5). The Maison the models surface most is not the one they hold highest.

> *A Maison can dominate the answer without owning the admiration. Another holds the admiration and surfaces only when asked for by name.*

This is the thesis of the whole index made visible on one map: visibility and desirability are two axes, not one ranking. The distance between where a Maison sits on each is the work that remains.

03

Heritage cannot be improvised, and AI knows it.

On Heritage Depth, the founding Maisons cluster at the top, Chanel 95.1, Saint Laurent 90.5, Dior 89.4, Hermès 87.5, while the youngest contemporaries score zero: the models make no spontaneous heritage attribution to Miu Miu or The Row at all. Alaïa, carrying Azzedine Alaïa's couture legacy, registers a faint 17.9.

> *In fashion, heritage in AI's eyes is still a function of time. The youngest Maisons read at or near zero.*

It is the cleanest separation in the edition. Savoir-faire can be built in a decade and the models will credit it; spontaneous heritage cannot, and they will not invent it. Bottega Veneta is the telling exception, an 80.8 heritage score that reads like a mega-Maison, evidence that AI weighs sustained craft reputation as a heritage signal in its own right.

04

Google's AI Overviews barely enter the fashion conversation.

Across the edition's prompts, brand-agnostic and maison-directed alike, Google AI Overviews triggered for 52.5% of English queries and 0% of French ones. On the AI Presence axis, the AIO layer scores near zero for most Maisons, dragging the channel to the floor for Dior, Miu Miu, Alaïa and The Row alike.

The layer closest to mass-market search is structurally absent from luxury fashion, and entirely absent in French. A Maison that reads its AI exposure through English-language Overviews alone will systematically misjudge its presence in its home market.

05

Where the models diverge, the narrative is still open.

Most Maisons read consistently across the five models. The disagreement concentrates on a few: Dior (standard deviation 15.0) and Hermès (13.7) are read most differently depending on which model is asked, while Bottega Veneta (2.3) and Gucci (6.9) read almost identically everywhere.

> *Divergence is not noise to be averaged away; it marks the Maisons whose AI representation is still being formed, and is therefore most open to influence.*

For the Maisons where the models split, the narrative is not yet settled. Those are the Maisons for which deliberate source and reference work would move the reading fastest.

SIX DIMENSIONS

The anatomy of *AI desirability.*

Three dimensions compose AI Presence: Product Discovery, Brand Knowledge, Gifting Occasion. Three compose Cultural Desirability: Heritage Depth, Savoir-Faire Authority, Cultural Resonance. Each is scored 0–100, model-averaged across all five.

READING THESE SCORES

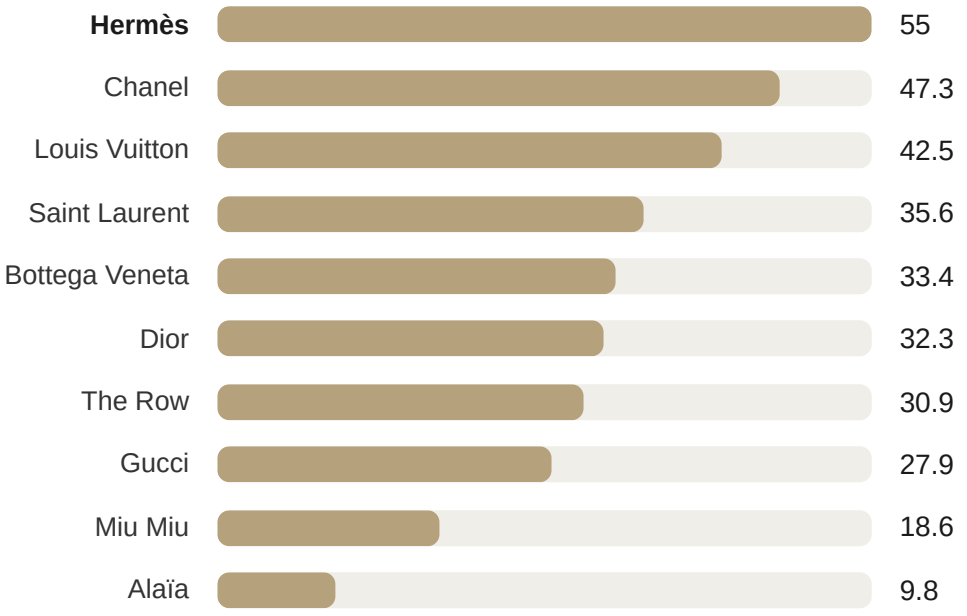
Each score runs from 0 to 100. It is not a percentage, and not a sales figure. It is an index built from how the five models actually answer, reading four signals in every response: how often the Maison appears, how early, in what tone, and with how much detail.

Savoir-Faire Authority replaces the Fragrance edition's Olfactory Authority. It reads how fluently a model speaks of materials, ateliers and techniques, how a thing is made, not how it looks.

A practical note: gaps under one point sit within the measurement's natural noise. On Brand Knowledge, Dior at 75.6 and Hermès at 75.6 are a tie, not a ranking.

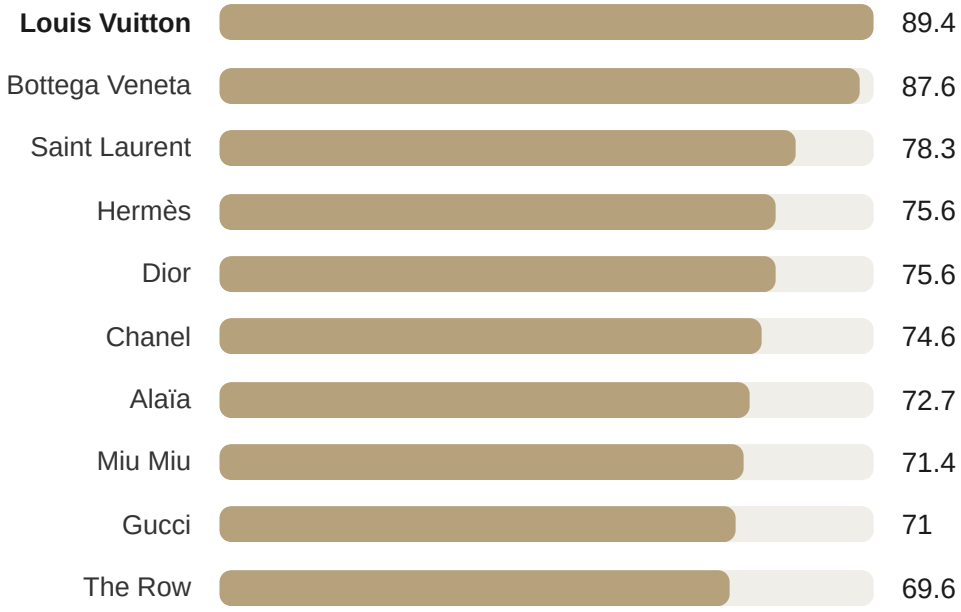
Product Discovery

Spontaneous surfacing when a client seeks ready-to-wear or leather goods · AI Presence · Scores /100



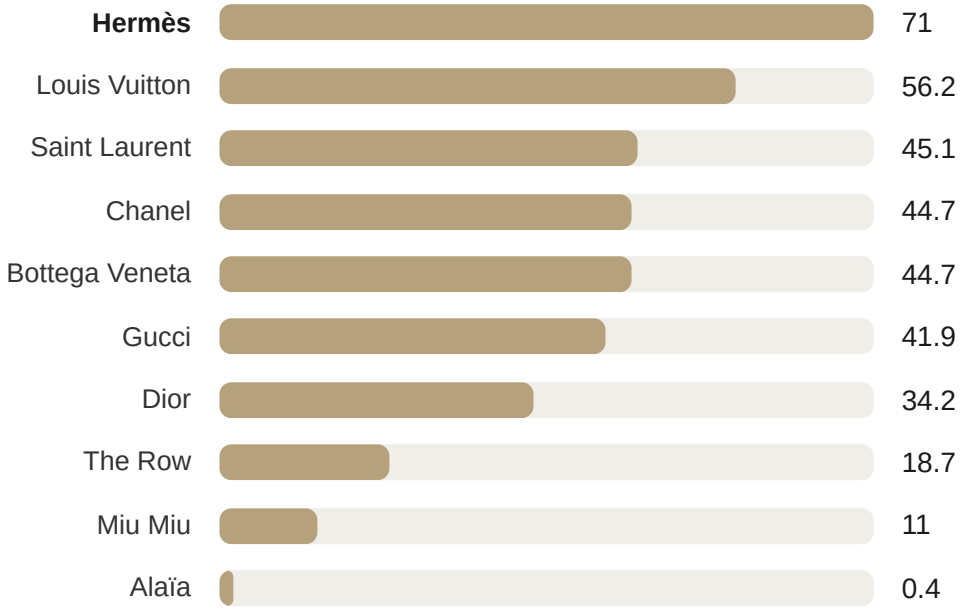
Brand Knowledge

Depth and accuracy of what the model knows · AI Presence · Scores /100



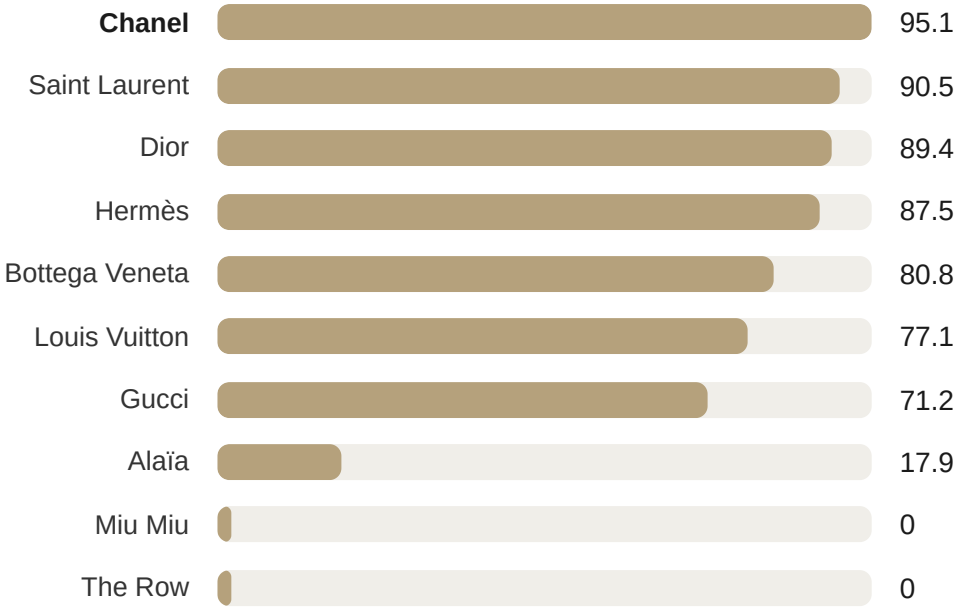
Gifting Occasion

Relevance when the prompt is a gift or an occasion · AI Presence · Scores /100



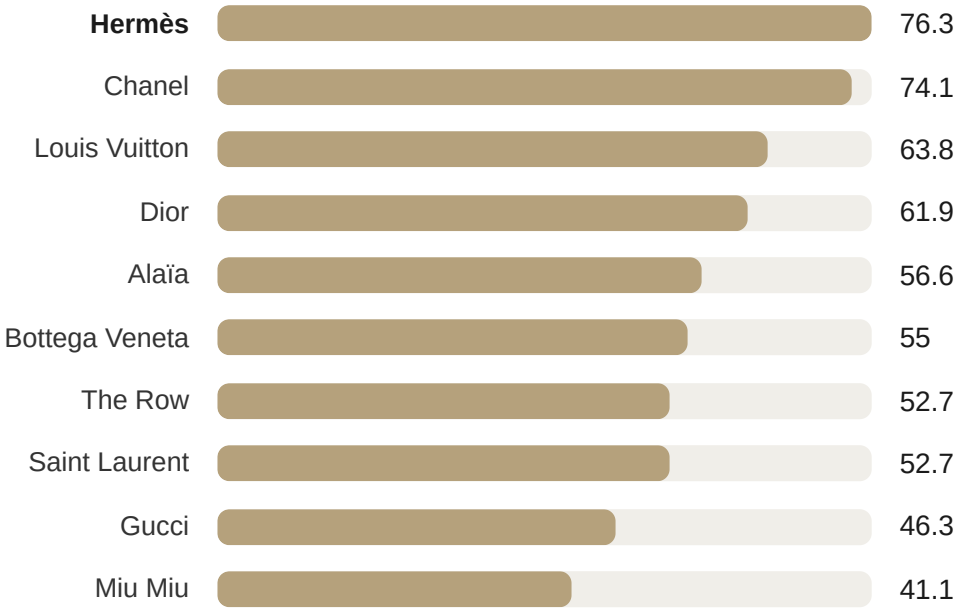
Heritage Depth

Heritage and patrimony the models attribute · Cultural Desirability · Scores /100



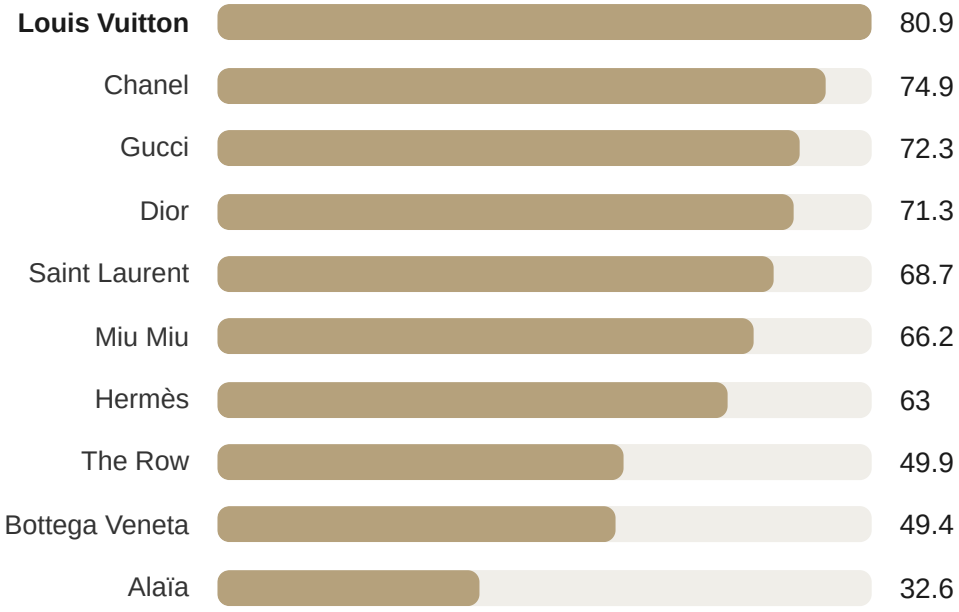
Savoir-Faire Authority

Materials, ateliers, techniques the models invoke · Cultural Desirability · Scores /100



Cultural Resonance

Cinema, art, music and cultural circulation · Cultural Desirability · Scores /100

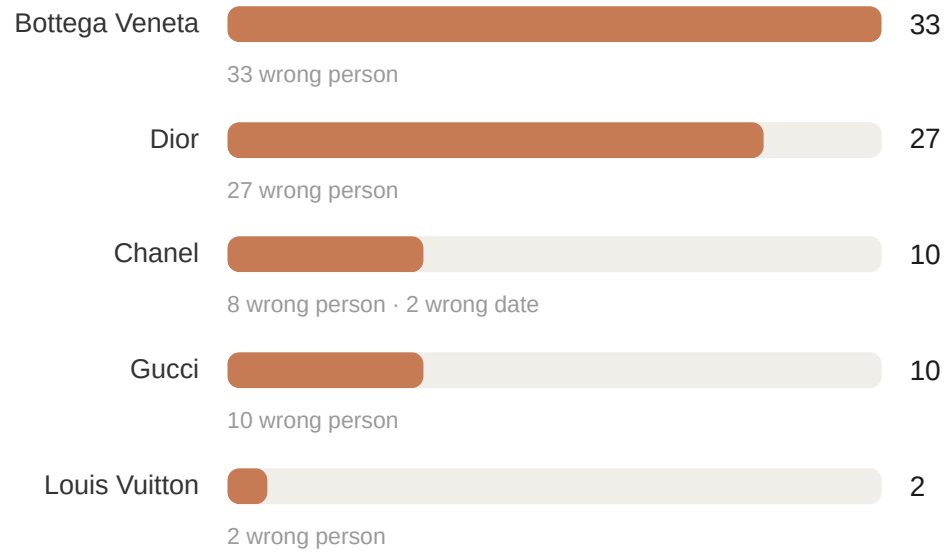


When AI dresses a Maison in *last season's director.*

Desirability measures how a Maison is regarded. This audit measures something the desirability score deliberately ignores: whether the model has the facts current. We flag every response that attributes an artistic director the Maison no longer has, anchored to who was in post on the collection date (2026-06-16). The flag is evidentiary. It never touches a Maison's desirability or valence.

STALE-DIRECTOR MENTIONS BY MAISON

82 mentions · 80 name the wrong person, 2 the right person on a wrong date · Maisons at zero omitted



> *The error is worst where a designer moved between Maisons. AI keeps Matthieu Blazy at Bottega Veneta long after he left for Chanel, which is why Bottega carries 33 of the 82 flags.*

The reading is narrow on purpose. The models are not generally wrong about who runs these Maisons. They lag in one case, when a designer moves from one Maison to another and the old attachment stays. Bottega Veneta (33 mentions) and Dior (27) carry the weight. Chanel's ten include the only date errors, where the model names the right designer but

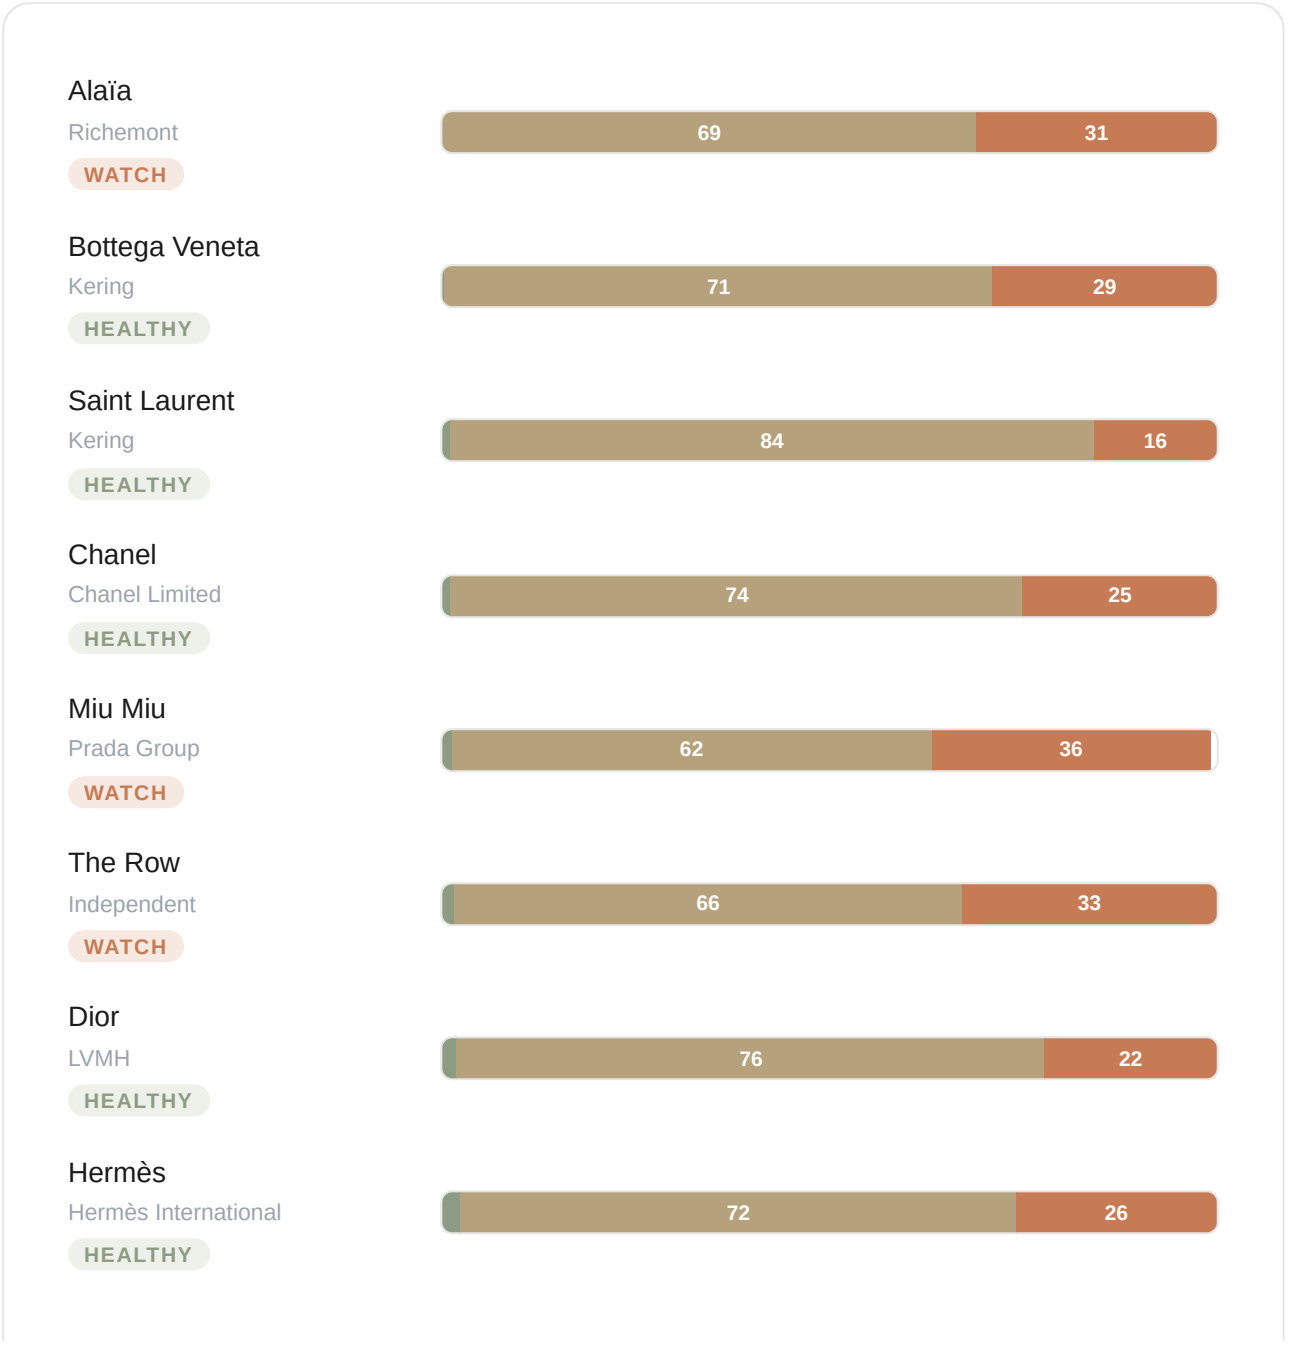
mis-dates the arrival. The Maisons with no recent change of director draw zero flags. That is the control: the method reads real lag, it does not manufacture it.

Every flag is traceable to its source response. The raw text is preserved, never corrected: the point of the index is to record what AI says, not to fix it.

SOURCE AUTHORITY

Who AI *listens to* when it speaks of a Maison.

Google AI Overviews is the one layer that publishes its grounding sources. We classify each as **owned** (the Maison's own domains), **earned** (editorial and press) or **third-party** (platforms and social the Maison does not control). Chat-model sources are model-inferred and excluded from this metric.



Gucci

Kering

HEALTHY



Louis Vuitton

LVMH

HEALTHY



Owned share is in single digits for every Maison, the structural reality of the AI regime: even the most desired Maison sees its authority narrated overwhelmingly by voices it does not direct. The contemporary Maisons carry the highest third-party dependency, and a **watch** classification with it.

How AI *tells each Maison*, and where the story is exposed.

Chanel

Presence 55.5 · Desirability 81.3

CHANEL LIMITED

The most desired Maison in AI, and the one most narrated by others.

ANCHOR Cultural Desirability leader of the panel (81.3), with the deepest heritage score of all ten (95.1).

EXPOSURE Owned sources are 1% of its AI grounding; third-party platforms carry a quarter of the narrative.

LEVER The reverence is secured. The opening is source control, converting earned admiration into owned authority.

Dior

Presence 47.4 · Desirability 74.2

LVMH

Revered by AI, surfaced unevenly, and the second-most mis-attributed on direction.

ANCHOR Heritage 89.4 and Cultural Resonance among the strongest; a clear Heritage-Strong reading.

EXPOSURE 27 stale-director mentions, and the panel's widest cross-model spread (15.0): the narrative is unsettled.

LEVER Discovery (32.3) and gifting (34.2) lag the desirability. Presence at the moment of choice is the lever.

Hermès

Presence 67.2 · Desirability 75.6

HERMÈS INTERNATIONAL

The presence leader: surfaced first, and revered, without a single stale-director flag.

ANCHOR Tops AI Presence (67.2) and gifting (71.0); savoir-faire authority leads the panel (76.4).

EXPOSURE Zero director errors, the cleanest accuracy reading among the mega-Maisons.

LEVER Read most differently across models (13.7). The Maisons that diverge are the ones still open to shaping.

Louis Vuitton

Presence 62.7 · Desirability 73.9

LVMH

The most consistently known Maison, and the second-most cross-model stable of the mega-group.

ANCHOR Brand Knowledge leads the panel (89.4); Cultural Resonance leads it too (80.9).

EXPOSURE Only 2 stale-director mentions despite a dual-director structure; strong factual currency.

LEVER Cultural circulation is the asset. Discovery (42.5) is the headroom to convert presence into intent.

Gucci

Presence 46.9 · Desirability 63.3

KERING

The most contrasted Maison in AI: strong knowledge, softer desirability.

ANCHOR Brand Knowledge 71.0 and Cultural Resonance 72.4 hold; Savoir-Faire (46.3) reads thin.

EXPOSURE 10 stale-director mentions, and the panel's most polarised valence in annotation.

LEVER Awaiting-maturity by the medians. The lever is savoir-faire narrative, the dimension AI under-credits.

Saint Laurent

Presence 53 · Desirability 70.6

KERING

Heritage-secure and accurate, with discovery still to build.

ANCHOR Heritage Depth 90.5, second only to Chanel; the highest earned-source share (84%).

EXPOSURE Zero stale-director flags; clean factual currency.

LEVER Product Discovery (35.6) trails the desirability. Surfacing at the point of search is the opening.

Bottega Veneta

Presence 55.2 · Desirability 61.7

KERING

AI's most consistent Maison, and its most mis-dated one.

ANCHOR The panel's lowest cross-model spread (2.3) and an 80.8 heritage score that reads like a mega-Maison.

EXPOSURE 33 stale-director mentions, the most of any Maison: AI keeps Blazy in place long after Chanel.

LEVER Desirability is sound and untouched by the error. The lever is currency, a corrected director narrative at source.

Miu Miu

Presence 33.6 · Desirability 35.8

PRADA GROUP

Culturally legible, structurally young in AI's eyes.

ANCHOR Cultural Resonance 66.2 is real; zero spontaneous heritage attribution pulls desirability down.

EXPOSURE Highest third-party source dependency of the panel; a watch classification.

LEVER The lever is owned and earned grounding: the cultural energy exists, the authority layer does not yet.

Alaïa

Presence 27.7 · Desirability 35.7

RICHEMONT

Deep savoir-faire, almost no AI surface.

ANCHOR Savoir-Faire 56.6 outranks several mega-Maisons; the legacy reads. Discovery (9.8) is the panel floor.

EXPOSURE No owned sources at all in AI grounding; entirely narrated by others.

LEVER Presence is the whole task. The craft authority is there to be surfaced; today AI rarely reaches it.

The Row

Presence 39.7 · Desirability 34.2

INDEPENDENT

Known to the models, with no heritage and little circulation in AI.

ANCHOR Brand Knowledge 69.6 holds; heritage and gifting read at or near zero.

EXPOSURE Read consistently across models; the narrative is stable but thin.

LEVER The lever is cultural reference and savoir-faire narrative, the signals AI rewards and the Maison under-supplies.

CROSS-MODEL

Five models, *where they agree and split.*

A low standard deviation means the five models read a Maison alike; a high one means its AI representation is still contested. Bottega Veneta is read almost identically everywhere (2.3); Dior and Hermès most differently.

MAISON	AI PRESENCE	CULTURAL DESIR.	CROSS-MODEL Σ	QUADRANT
Hermès	67.2	75.6	13.7	Cultural Icon
Louis Vuitton	62.7	73.9	7.2	Cultural Icon
Chanel	55.5	81.3	12.8	Cultural Icon
Bottega Veneta	55.2	61.7	2.3	AI-Native Star
Saint Laurent	53	70.6	10.4	Cultural Icon
Dior	47.4	74.2	15	Heritage-Strong · AI-Emerging
Gucci	46.9	63.3	6.9	Awaiting AI Maturity
The Row	39.7	34.2	7.4	Awaiting AI Maturity
Miu Miu	33.6	35.8	8.3	Awaiting AI Maturity
Alaïa	27.7	35.7	13.4	Awaiting AI Maturity

Scores model-averaged across Claude, ChatGPT, Perplexity, Gemini and Google AI Overviews. σ is the standard deviation of the Maison's overall score across the five.

What the map asks of the *Maison.*

Desirability is earned, not inherited.

The Maisons that lead are those AI both surfaces and reveres. Heritage opens the door; savoir-faire narrative and cultural circulation are what the models actually reward inside it. Bottega Veneta's mega-Maison heritage score, built on craft reputation rather than age, is the proof.

Accuracy is a layer of its own, and AI lags on it.

The 82 stale-director mentions are a currency problem, not a desirability one. They concentrate on the Maisons in transition, and they reward a Maison that grounds its creative-direction narrative at source. This is the corrective layer the index exists to make visible.

Source control is narrative control.

Owned share is in single digits for every Maison. The Maisons most exposed to third-party narration carry a watch classification. Whoever supplies the grounding supplies the story.

Divergence is the opening.

Where the five models disagree, the narrative is not settled. Dior and Hermès, read most differently, are the Maisons where deliberate source and reference work would move the AI reading fastest.

METHODOLOGY

How this index was *built.*

The panel	Ten fashion Maisons across ready-to-wear and leather goods: six heritage mega-Maisons (Chanel, Dior, Hermès, Louis Vuitton, Gucci, Saint Laurent) and four contemporary Maisons (Miu Miu, Alaïa, The Row, Bottega Veneta). Beauty, fragrance and fine jewellery are out of scope.
The models	Five: Claude, ChatGPT, Perplexity, Gemini and Google AI Overviews. Each prompt is run identically across all five.
The dimensions	Six, on two axes. AI Presence: Product Discovery, Brand Knowledge, Gifting Occasion. Cultural Desirability: Heritage Depth, Savoir-Faire Authority, Cultural Resonance. Savoir-Faire Authority replaces the Fragrance edition's Olfactory Authority and measures materials, ateliers and techniques.
The prompts	120 prompts in two families: 80 brand-agnostic prompts that measure spontaneous mention, and 40 Maison-directed prompts. Brand-agnostic prompts run once across five models; Maison-directed prompts run for each of the ten Maisons. 80×5 + 40×5×10 = 2,400 responses.
Scoring	Each dimension is scored 0–100 from four signals, presence, prominence, tone and depth, then model-averaged across all five. Mentions without a valenced context (133 of 4,257) are excluded from the valence mean, never counted as zero.
The accuracy audit	A separate, evidentiary layer. Each response is checked against a reference of the artistic director in post on the collection date (2026-06-16); a mention is flagged when it attributes an outdated director (wrong person) or a wrong appointment date. The audit is kept strictly out of desirability and valence scoring, and raw responses are never altered.

Disclosure	The author held the role of Executive Director, Digital & E-Commerce at Diptyque, and earlier worked at Dior Couture. Dior is one of the ten Maisons in this panel. The director-staleness finding that involves Dior is unfavourable to the AI models, not to the Maison, and the scoring is fully reproducible from the underlying data.
A note on the score	These are AI-perception scores, not commercial figures. They measure how five models perceive, surface and narrate each Maison in June 2026, nothing more, and nothing less.

AUGMA · The AI Desirability Index · Fashion Edition

June 2026 · Sébastien Pagès · AUGMA / The Augmented Maison. Narrative coherence infrastructure for the AI regime. The data behind every figure in this report is traceable to its source responses.