

THE RISK AND REWARD OF RETRIEVABILITY

Why global brands risk becoming invisible if AI Agents and algorithms do not retrieve and recommend them.

With thanks to all our roundtable attendees and contributors, who provided an 'inside-out' view of how AI adoption and retrievability practices are being applied within organisations.



INTRODUCTION

For decades, marketing has made brands visible and resonant to people. The lines have been drawn, and roles and remits have been clear. Today, brands face a critical second audience, the AI intermediary.

These super intelligent machine gatekeepers shape the information and brands people encounter through search and discovery every day—from the algorithms that govern social media, YouTube and news feeds to the responses generated by AI Agents such as ChatGPT, Claude, Perplexity and Gemini.

The consideration journey has been deeply affected. Research shows that what was once a multi-touchpoint journey spanning several channels and significant periods of time can now be condensed into a relatively short conversation with an AI Agent that does much of the ‘thinking’ for the users.

This shift does not render classical brand salience obsolete. Instead, it highlights the need for evolution: brand building for humans that captures attention and emotion and brand building for machines that can be parsed, retrieved, and corroborated.

Becoming Retrievable to machine audiences is as much a technical challenge as it is a cultural one. Every area of an organisation and its marketing operations need to come together to enable this evolution.



THE COMMERCIAL UPSIDE OF RETRIEVABILITY IS IMMENSE

YOU CAN BECOME THE DEFAULT CHOICE.



PwC's 2026 AI Performance Study found that 20% of organisations surveyed captured nearly three-quarters of all AI-driven business impact, exposing a widening divide between leaders and laggards. This is in large part because Agents do not compile a shortlist of ten links; they produce a handful of recommendations and reasons to choose. This has two main consequences:

- 01

DECISION OFFLOADING

Consumers delegate all or part of the evaluation phase of decision making, reaching purchase decisions with fewer manual touchpoints. Furthermore, traffic originating from conversational agents displays significantly higher intent: (analysis from Ahrefs showed that while AI visitors represented just 0.5% of overall traffic, they generated over 12% of commercial signups—converting at over twenty times the rate of organic search) meaning that the machine recommendations are trusted and acted upon.
- 02

WINNER-TAKES-MOST VISIBILITY


Brands engineered for retrievability are consistently retrieved and recommended, while digitally fragmented competitors are entirely left out of the picture, shrinking the customer's consideration set and to fewer options.

That sounds compelling! But we must remember that an AI agent cannot recommend what it cannot see and verify. **The brand will only be surfaced if there is sufficient brand equity and credible information available to find.**

As System1 and Effie's *The Creative Dividend* showed, creative excellence and media scale drive over 60% of campaign business impact, directly accelerating profit growth. We know this works for our human audiences, but why does this matter if Agents don't 'read' ads in the same way?

Agents crawl, digest and cross-reference information at a scale no human could. While they may not 'watch' the ads, they continuously ingest the cultural wake those campaigns leave behind: high-authority editorial coverage, community conversation, structured reviews and cross-platform sentiment.

Building machine equity alongside human equity is no longer a secondary technical optimisation task for back-office SEO teams. It is a core strategic discipline that must be architected into brand comms planning from day one.



The commercial downside of not being retrievable is immense

YOU BECOME INVISIBLE

(and engineered out of the choice set entirely).

A brand missing from a relevant AI answer may lose an opportunity to be considered. Appearing in the answer is no guarantee of success, either.

What if the model serves up an old version of your pricing, hallucinates a service you killed off two years ago, or fixates on a random bad review from 2021? Being retrieved is only half the battle-accuracy and context are the rest. Marketing teams now have to know which questions their buyers are actually asking, what answers engines spit out, and whose sources those systems rely on to shape responses.

Add to that the fact that people aren't clicking through to websites like they used to, and the penalties for being overlooked start stacking up fast:

01

ZERO-CLICK

SparkToro and Datos showed that 58.5% of Google searches (US & EU) ended without a click. Pew backed this up, finding that people click traditional links only 8% of the time when an AI summary pops up, compared to 15% when it doesn't. Gartner makes a bold prediction that traditional search volume will drop 25% as conversational bots take over, claiming that a traffic-first mindset is short-sighted. These collectively echo the need to be retrieved and recommended.

02

THE
"COST OF DULL" HITS
RETRIEVABILITY HARD

The Extraordinary Cost of Dull highlighted that boring, forgettable ads cost brands an extra £13.29 billion media spend just to match the impact of creative, distinctive work. Well, that penalty now extends to the machine audience. Bland brands don't spark cultural conversation, influencer commentary or earned media coverage, giving AI only crumbs to ingest. And remember, an AI agent cannot recommend what it cannot see and verify. Campaigns that fail to generate dense, authoritative evidence across the web give Agents zero reason to recommend it over a competitor with a loud, verified footprint.

Note: while casual site visits drop, buyer qualification increases. When an AI agent recommends a brand, the downstream user converts at higher velocity because the machine pre-filtered the choice.

03


COMPRESSED
CONSIDERATION SETS

McKinsey's research on the consumer decision journey showed that brands in the initial consideration set can be up to three times more likely to be purchased than brands outside it. AI agents compress that consideration window into one or two recommendations. If your brand isn't legible and verified across the web for machines to retrieve, it risks being excluded from the decision.

In summary

Classic marketing has not stopped working. It's that all the great work you're doing to win over human hearts is just getting blocked at the door because machines can't read the signals.

This is where evolution is needed.



OVERCOMING THE
OPERATIONAL BARRIERS
VIEWS FROM THE INSIDE

To understand how organisations can navigate this evolution while facing into real-world headwinds, we held a series of roundtables to explore the topic. Five key operational tensions, strategic implications and practical solutions emerged:

01

BUDGETS & RESOURCE
ALLOCATION

02

PLANNING
ORGANISATIONAL SILOS

03

CHANNEL DIVERSIFICATION AND
GENERATIVE SEARCH OPTIMISATION (GEO)

04

CREATIVE: ESCAPING THE
SAMENESS TRAP

05

WAYS OF WORKING
& MEASUREMENT

01	BUDGETS & RESOURCE ALLOCATION
02	PLANNING ORGANISATIONAL SILOS
03	CHANNEL DIVERSIFICATION AND GENERATIVE SEARCH OPTIMISATION (GEO)
04	CREATIVE: ESCAPING THE SAMENESS TRAP
05	WAYS OF WORKING & MEASUREMENT



Marketing leaders report growing pressure on how they deploy budgets and tightening purse strings frequently lead to cuts in innovation. In larger organisations, different departments may address the challenge in different ways, which can cause fragmentation and misalignment.

Strategic implication

Budgets should be structured to deliver impact among human and machine audiences, which requires ring-fencing dedicated investment for innovation.

Practical application

Successful teams apply a 70:20:10 investment model, 70% in established core channels, 20% in tested opportunities ready to scale, and 10% dedicated to emerging retrieval work like structured-data improvements and conversational-discovery pilots.

“We have to prove that innovation directly impacts core business, balancing long-term brand equity building with short term results. We do this by working collaboratively cross-team to develop strategic initiatives and use a shared language of success to demonstrate impact”

Lisa Hale, Head of Broadcast, Consumer PR, Social Media & Brand Activation, Specsavers

01	BUDGETS & RESOURCE ALLOCATION
02	PLANNING ORGANISATIONAL SILOS
03	CHANNEL DIVERSIFICATION AND GENERATIVE SEARCH OPTIMISATION (GEO)
04	CREATIVE: ESCAPING THE SAMENESS TRAP
05	WAYS OF WORKING & MEASUREMENT



Over the past few years, organisations have tended to view AI through an operational or cost-management lens-looking at automation and workflow efficiency-rather than AI as a critical audience and distribution channel.

Strategic implication

Paid, earned, shared, and owned media must come together to form integrated communications planning discipline that treats campaigns and always-on activity as a single ecosystem, rather than as a linear(unconnected) organisational framework for messaging.

Practical application

Many organisations described cross-functional task forces set up to address retrievability. These may be independent teams working across departments or hybrid councils made up of members from different functions. Their work ranges from shared quarterly planning to internal upskilling sessions.

"AI brand performance is a live discussion between our teams as we explore the best tools and metrics to invest in. There's a real appetite to collaborate across digital marketing and communications teams, and it's a healthy challenge for us to tackle together"

Helen Smith Strategic PR Lead, Slimming World.

01	BUDGETS & RESOURCE ALLOCATION
02	PLANNING ORGANISATIONAL SILOS
03	CHANNEL DIVERSIFICATION AND GENERATIVE SEARCH OPTIMISATION (GEO)
04	CREATIVE: ESCAPING THE SAMENESS TRAP
05	WAYS OF WORKING & MEASUREMENT



Traditional planning has relied on a few core channels to do the heavy lifting. However, generative search has accelerated the need to diversify the channel mix.

GEO is on everyone’s lips. Every channel has a role to play, and it’s increasingly important to get the mix right. A recent study of 2.37 million queries found that Google AI Overviews appeared on 25.8% of searches - meaning roughly one in four search queries are served with an AI retrieved and recommended answer that bypasses traditional paid and organic links. Conversational AI answers can also bypass paid and organic links, leaving brands vulnerable.

Strategic implication

Every channel—paid, earned, shared and owned - now feeds the signal layer that LLMs (Large Language Models) ingest, and the Agents retrieve and recommends. The mix needs to be considered and weighted appropriately, because overreliance on traditional channels can create blind spots in machine retrieval.

Practical application

In our roundtables, marketers described shifting their focus towards higher-impact channels such as YouTube and communities such as Reddit, alongside a growing emphasis on earned media. Integration with technical teams was a core priority, however, the discussions highlighted the need for greater connectivity to ensure that all activity is aligned and able to deliver maximum impact.

“Organisations are often fundamentally dysfunctional because everyone has a different set of priorities. Getting everyone on the same page and focused on the same goals requires a focused leadership and a real organisational commitment.”

Katie Cope, Marketing Consultant

01	BUDGETS & RESOURCE ALLOCATION
02	PLANNING ORGANISATIONAL SILOS
03	CHANNEL DIVERSIFICATION AND GENERATIVE SEARCH OPTIMISATION (GEO)
04	CREATIVE: ESCAPING THE SAMENESS TRAP
05	WAYS OF WORKING & MEASUREMENT



Generative tools have made content creation cheap and almost limitless contributing to the rise in “AI Slop” and “AI Convergence” (similar messaging across categories and brands due to everyone using the same tools and prompts).

We have already discussed the impact of dull creativity and the need for narratives that have cultural resonance and travel beyond their initial placement. Paid, earned and owned activity all have value, but earned thinking is attracting renewed industry focus as brands look to build the credibility and third-party validation that both people and machines take notice of.

Strategic implication

Creative that stands out and cuts through is the only defence against human and machine invisibility. Culturally resonant narratives that generate conversation and talkability, supported by sharp brand codes, semantic hooks and clear opinions, continue to capture human attention and give LLMs clear anchors to organise and retrieve.

Practical application

Brands should audit their guidelines to embed these principles and ensure that everyone responsible for activation understands how their work will amplify or dilute the brand’s digital footprint. Organisations also need a shared view of what good creative looks like for both human and machine audiences.

“There’s a real awareness around using AI purely to create. If everyone puts the same prompts in, organisations are likely to get the same answers out, flooding the market with undifferentiated campaigns. Human input and distinctive point-of-view have never been more critical.”

*Kacee Fennelly,
Senior Social Media & PR Manager, TalkTalk*

01	BUDGETS & RESOURCE ALLOCATION
02	PLANNING ORGANISATIONAL SILOS
03	CHANNEL DIVERSIFICATION AND GENERATIVE SEARCH OPTIMISATION (GEO)
04	CREATIVE: ESCAPING THE SAMENESS TRAP
05	WAYS OF WORKING & MEASUREMENT



Campaign-only planning leaves natural gaps in brand activity during which secured AI authority begins to decay and signals weaken. Many organisations are therefore considering how campaigns and always-on activity can better work together.

Taking the organisation on this journey, and resolving misaligned taxonomies, language, KPIs and definitions of success, also emerged as a challenge.

Strategic implication

Organisations must establish a single, shared language of success. Marketers need to consider how isolated campaigns integrate with an always-on architecture, evaluating machine retrievability, brand equity and business outcomes as one interconnected system.

Practical application

Organisations should develop shared taxonomies, language and measures of success, clarifying what is tracked centrally and what each channel owns. Machine-retrieval indicators should be linked directly to business KPIs, giving finance and executive leaders a clear reason to invest in retrievability. Paid, earned and owned teams should coordinate through a shared plan to keep brand signals active and reduce algorithmic decay.

"Measurement is currently disjointed across channels. When reporting fails to tie signal presence back to bottom-line impact, protecting innovation investment becomes difficult."

Steph Bridgeman, Managing Director at Experienced Media Analysts and AMEC Board Member since 2021

THE ENGINE ROOM
HOW LLMS, ALGORITHMS,
AND AGENTS COLLABORATE

To generate machine retrieval, marketing teams must stop treating “AI” as a black box. In conversational discovery, three functions interact to decide whether your brand surfaces, gets retrieved and recommended. What you’ll find in practice, is that these three layers often overlap and their architectures vary widely across platforms:

01

THE LLM IS
THE ‘BRAIN’

Large Language Models understand nuance, intent, and semantic context. They determine how queries are parsed and how the final answer is structured and phrased. If your brand positioning is generic or undifferentiated, the model lacks distinctive semantic associations to link your brand naturally to the user’s specific problem.

02

THE
RETRIEVAL SYSTEM

LLMs do not browse the live web alone; external retrieval mechanisms (search indexes, vector databases, and knowledge graphs) fetch the underlying evidence from far and wide. While ranking methods remain largely proprietary and platform-specific, these engines evaluate relevance, fresh consensus, and structured signals across third-party sources. If your digital footprint is thin or contradictory, retrieval engines are less likely to pass your data back to the model as high-confidence context.

03

THE AGENT IS THE
DECISION-MAKER

The agent orchestrates the entire task. When given a complex goal (“Compare the top three enterprise accounting tools with rapid deployment for mid-tier retailers”), the agent formulates a plan: deciding whether to query external web tools, calling specific APIs, synthesising the retrieved evidence, and evaluating trade-offs before delivering a focused recommendation.

THE ENGINE ROOM
HOW LLMS, ALGORITHMS,
AND AGENTS COLLABORATE

When your brand performs across all three fronts, semantic distinction for the LLM, credible supporting evidence for the retrieval system, and clear problem–solution alignment for the agent, you improve its likelihood of being surfaced and recommended. The S.A.V.E. framework is designed to strengthen performance across this interconnected ecosystem.

THE S.A.V.E.
OPTIMISATION
FRAMEWORK

To move systematically from basic visibility to proactive retrievability, apply the S.A.V.E. framework:

STORY
& STRUCTURE

CAN PEOPLE AND
MACHINES FIND
AND UNDERSTAND
THE INFORMATION?

Clear, attention grabbing and culturally resonant brand story backed by dense, corroborating proof points.

Partner with web and SEO leads to ensure high-priority assets, case studies, and schema are cleanly indexed and crawlable by LLMs. Use a topic map to define distinct entity anchors, then partner with technical teams to translate those into machine-crawlable schema and topic maps.

ALIGN

DO OUR SOURCES
AGREE ON THE FACTS?

Audit details of all comms across all channels and in partner ecosystems to eliminate contradiction & build algorithmic confidence. Identify any contradictions or inaccuracies that trigger AI hallucinations or model suppression.

VALIDATE

WHAT INDEPENDENT
EVIDENCE IS AVAILABLE?

Clear, attention grabbing and culturally resonant brand story backed by dense, corroborating proof points.

Partner with web and SEO leads to ensure high-priority assets, case studies, and schema are cleanly indexed and crawlable by LLMs. Use a topic map to define distinct entity anchors, then partner with technical teams to translate those into machine-crawlable schema and topic maps.

EVOLVE


WHAT NEEDS TO BE
CHECKED OR UPDATED?

Continuously monitor branded and category-level conversational prompts across top AI engines.

Refresh outdated assets as market dynamics shift, closing temporal gaps and avoiding signal decay to ensure retrieval and recommendation is continual.


**DEEP DIVE THE OVERLOOKED POWER
OF MULTIMODAL RETRIEVAL (YOUTUBE)**

To generate machine retrieval, marketing teams must stop treating “AI” as a black box. In conversational discovery, three functions interact to decide whether your brand surfaces, gets retrieved and recommended. What you’ll find in practice, is that these three layers often overlap and their architectures vary widely across platforms:

A close-up photograph of a person's hands typing on a white computer keyboard. The lighting is warm and slightly dim, focusing on the keys and the movement of the fingers.

CITATIONS

AI Agents (including Google Gemini and AI Overviews) heavily prioritise video transcripts, automated captions, timestamps, and descriptive metadata to answer complex comparative and instructional queries, particularly in YMYL (Your Money Your Life) categories.

A man with dark hair and glasses, wearing a blue button-down shirt, is looking down at a black smartphone he is holding in his right hand. He is standing in front of a large window with a view of a city street.

**HIGH-TRUST
& CONSENSUS**

Video provides tangible demonstration, creator validation, and unfiltered community conversation via comment sections. These assets are actively indexed by the LLMs behind the AI Agents that are looking for real audience sentiment and real-world efficacy.

A young man with dark hair, wearing a pink t-shirt and blue overalls, is smiling and waving at the camera. He is holding a vlog camera on a tripod. The background shows a balcony with plants and a building.

THE OPPORTUNITY

YouTube has a huge amount of audience and machine clout that optimisations can compound a retrieval footprint across AI Engines, in the new living-room screens, and traditional search results alike.

THE EXECUTIVE READINESS CHECKLIST

AUDIT RETRIEVABILITY (SHARE OF MODEL)	Benchmark and baseline where and how your brand is cited in conversational AI answers across category queries, relative to primary competitors.
MODERNISE INFORMATION ARCHITECTURE	Ensure that website schema, machine-readable assets, robots directives and clear proposition copy are accessible to search crawlers and permitted data-ingestion systems.
REBALANCE TOWARD EARNED CONSENSUS	Prioritise third-party authority (reviews, industry analyst recognition, earned media) over self-referential marketing collateral.
ENGAGE WHERE CONVERSATIONS LIVE	Actively participate across industry forums, podcasts, and digital communities where natural language consensus is established.
ESTABLISH SIGNAL CADENCE	Replace isolated campaign spikes with an always-on content and PR calendar to avoid citation decay.
ALIGN BOARD METRICS	Develop an executive-ready retrievability index linked to commercial outcomes, while distinguishing correlation from proven causation.
CLOSE THE CAPABILITY GAP	Train internal marketing, communications and digital teams in generative engine optimisation (GEO) principles.



CONCLUSION

AI HAS AMPLIFIED THE IMPORTANCE AND INFLUENCE OF BRAND BUILDING.

The brands that win in the coming decade will master dual-audience marketing: captivating human emotion while systematically creating trustworthy signals for the algorithms that govern discovery.

Companies that act now can build compounding retrieval authority; those that delay may find that the shortlist has been written without them.

With thanks to all our contributors.

LISA HALE Head of Broadcast, Consumer PR, Social Media & Brand Activation, Specsavers	ROSIE GOGOLEY, Public Relations Lead at Skipton Building Society	STEPH BRIDGEMAN, Managing Director at Experienced Media Analysts and AMEC Board Member since 2021	KATIE COPE, Consultant
KACEE FENNELLY, Senior Social Media & PR Manager at Talk Talk	HELEN SMITH, Strategic Communications PR Lead at Slimming World	GEORGINA PRODHAN, Global Head of Media Relations at Siemens Healthineers	ASHLEY JENKINS, Corporate Communications Manager at Experian
ANDREA KLETTNER, Director Communications, UK & Ireland Avison Young	TONY MAYS, PR Manager EMEA at Interface	ANNE FUELL, Director of Marketing and Development at Girls Day School Trust	

SOURCES

1 | System1 and Effie Worldwide, The Creative Dividend: A Practical Guide to Making Effectiveness Happen in Advertising Today (27 January 2026). System1

2 | Google Search Central, “AI Features and Your Website” (updated 10 December 2025). Google Search Central

3 | PwC, Want ROI from AI? Go for Growth, PwC’s AI Performance Study (13 April 2026). The study surveyed 1,217 companies across 25 sectors; 20% captured 74% of AI-driven returns. PwC

4 | Rand Fishkin, SparkToro and Datos, “2024 Zero-Click Search Study” (1 July 2024). The study reported that 58.5% of US Google searches and 59.7% of EU searches ended without a click; its authors note panel and device-coverage limitations. SparkToro

5 | Athena Chapekis and Anna Lieb, Pew Research Center, “Google Users Are Less Likely to Click on Links When an AI Summary Appears in the Results” (22 July 2025).
Based on browsing data from 900 US adults and 68,879 Google searches in March 2025; traditional-result clicks occurred on 8% of visits with an AI summary versus 15% without one. Pew Research Center

6 | Gartner, “Gartner Predicts Search Engine Volume Will Drop 25% by 2026, Due to AI Chatbots and Other Virtual Agents” (19 February 2024). This is a forecast, not a measured outcome. Gartner

7 | System1, eatbigfish, Adam Morgan, Jon Evans and Peter Field, The Extraordinary Cost of Dull (2024). The report estimated that bringing dull US advertising up to the effectiveness of non-dull advertising would require an additional \$189 billion in media spend; the UK estimate was £13.29 billion. System1

8 | David Court, Dave Elzinga, Susan Mulder and Ole Jørgen Vetvik, “The Consumer Decision Journey,” McKinsey Quarterly (1 June 2009). The research examined purchase decisions from almost 20,000 consumers across five industries and three continents. McKinsey & Company

9 | Patrick Stox, Ahrefs, “Does AI Search Traffic Convert Better Than Traditional Search? For Ahrefs, Yes: 0.5% of Visitors Drove 12.1% of Signups” (16 June 2025).
The reported 23x conversion rate came from Ahrefs’ own website over a 30-day period and should be treated as a single-site, directional finding rather than a universal benchmark. Ahrefs

10 | Trevin Shirey, WebFX, “Where and Why Google’s AI Overviews Appear: A Study of 2.3m Keywords” (30 July 2025). The analysis found AI Overviews on 25.8% of a 2.37-million-query US sample, with substantial variation by query length, intent and sector. WebFX