WHITEPAPER

Digital Presenteeism

Your Content's Silent Crisis in the Age of Agentic Understanding

AndyIddon
Chief Strategy Officer, Content Bloom

Contents

Executive Summary

60%

60% of searches ending without a single click

21%

Google AI Overviews now appear in 21% of all search results

50%

Gartner predicts 50% or more as consumers embrace generative AI-powered search

36%

About 36% of clicks go to the open web, a new zero-click search study finds.

Sources from: Gartner, BCG and Search Engine Land

If you work in management or lead a team, you will recognize presenteeism immediately.

An employee shows up, is visible, and responsive, but is not materially contributing. The problem is invisible until performance drops. By then, the real damage has already been done.

This same dynamic is now playing out with how organizations present themselves to AI systems. **Your content is there, but is it working?**

> The UX is fine. The Agentic Understanding is broken. This is digital presenteeism: content that exists, appears functional to humans, but is operationally incomplete for AI systems at retrieval time.
>
> **Andy Iddon, Chief Strategy Officer, Content Bloom**

The Technical Assumption That No Longer Holds

Most organizations believe their content is accessible to AI. In many cases, that assumption does not hold. Websites built on JavaScript-heavy architectures, assembled dynamically after the initial page load, can be partially inaccessible to AI systems, crawlers, and enterprise agents, depending on how content is rendered and retrieved. Beautifully crafted Tabs, accordions, specification panels, dynamically injected content, anything that appears after load, is at risk of being invisible.

But even where that technical problem is solved, there is a second, more consequential gap that almost nobody is tracking.

Your content shows up. But is it actually doing anything?

The missing dimension: Agentic Understanding

Agentic Understanding is the degree to which modern systems, e.g., search engines, AI assistants, enterprise agents, can extract, interpret, trust, compare, and act on your content. Think of it as user experience for AI: the same discipline, a different audience. Viewed this way, the challenge is rarely confined to a website. It is usually a reflection of the wider content ecosystem, the governance that maintains it, the workflows that update it, and the systems that distribute it. The website is often just the visible surface of a much larger operational reality.

You used to run UX workshops for your customers. Why don't you do it for AI?

Most organizations obsess over UX (user experience). The interface is clean, navigation works, the priority of content makes sense in the page flow, and humans can find what they need. From a user experience perspective, everything is fine.

Consider a specific technical product page. The HTML is clean. The semantic structure is sound. A human visitor navigates to the specifications, finds the case studies, and downloads the datasheets. The UX works perfectly. But when an AI system crawls that same page, it may primarily rely on the HTML response available during retrieval, it is operationally hollow. The detailed technical specifications live in a PDF. The case studies sit behind download links. The performance benchmarks are in separate documents.

So a procurement engineer asking an AI system what are the specifications and lead times for this product gets an incomplete or inferred answer not because the information does not exist at all, but because it exists in a format that creates a retrieval risk. This is digital presenteeism.

Andy Iddon has 28 years of digital consulting experience and also co-founded a mental health organization, where he first encountered presenteeism as a workplace challenge. He applies that cross-disciplinary insight to AI content strategy.

The Pre-Funnel problem

The missing dimension: Agentic Understanding

Research increasingly supports that AI systems are becoming active participants in commercial discovery. Buyers ask AI systems direct questions, compare options, form views, and build shortlists before they ever visit a website or speak to a salesperson. This introduces a stage in the commercial process that most organizations are not measuring.

Pre-Funnel Eligibility

This is the filtering stage where AI systems include or exclude organizations from consideration before a customer directly engages with a website or sales process.

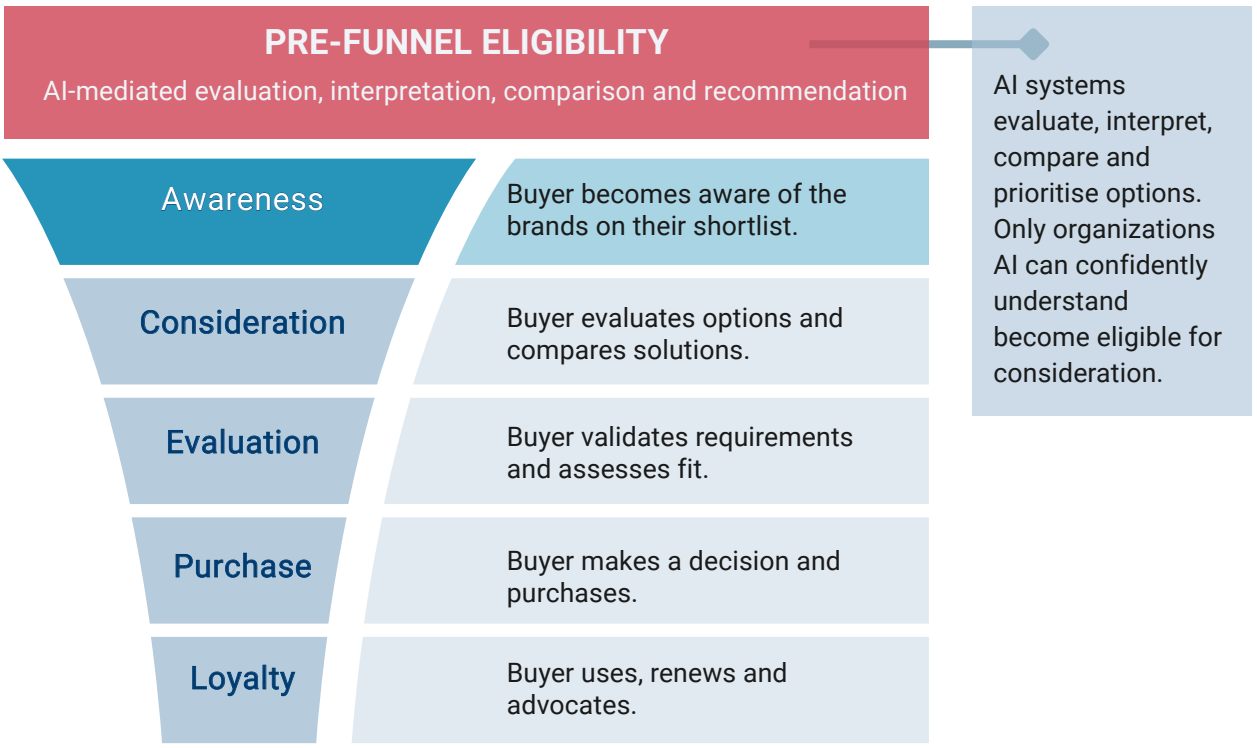

> That loss happens before the sales pipeline even exists.
> It is not a conversion problem. It is a pipeline problem. Demand is being filtered before you ever have the opportunity to compete.
>
> **Andy Iddon, Chief Strategy Officer, Content Bloom**

The real cost: what you cannot see in your analytics

Why is the loss silent

When a buyer visits your website and leaves without converting, you see it in Google Analytics or any other traffic measurement tool. Bounce rate, session durations, and exit page. When you have the data, you can investigate, test and iterate.

When an AI system evaluates your organization and excludes you from a recommendation before a buyer ever reaches your website, you see nothing. There is no event in your analytics. No signal in your CRM, no drop in traffic. The buyer simply never arrives because your organization was filtered out at the eligibility stage before they ever thought to look for you directly.

The compounding effect

This means the organizations most affected by digital presenteeism are also the least likely to know it is happening. Their conversion rates look stable, their traffic looks steady, and their pipeline feels normal. But the consideration set being assembled upstream in AI systems, before human intent is ever expressed, is quietly narrowing around their competitors.

By the time that shows up in revenue, the gap has been compounding for months.

Winning the game you're measuring while losing the game you can't see.
Digital Presenteeism creates a commercial exposure risk that rarely appears in analytics, CRM systems, or pipeline reporting until the impact has already compounded.

-Andy Iddon, Chief Strategy Officer, Content Bloom

The Hypothesis

Digital Presenteeism began with a simple observation.
Information that was clearly visible on a webpage was not always reflected in AI-generated responses about that page.

This led to a broader question:

If an AI system can find content, does that mean it truly understands it?

The hypothesis explored in this paper is:

A page can be visible to AI systems, yet still fail to be properly understood, trusted, compared, or recommended.

Historically, organizations have invested heavily in search engine optimization and customer experience research to improve visibility and usability for human audiences. The purpose of this investigation is to explore whether similar breakdowns in understanding occur when AI systems interact with content, and to identify the signals that appear to strengthen or weaken that understanding.

The Signals We Consistently See

The AI Signals Framework

These observations are not specific to any individual AI platform. In testing, multiple AI systems consistently identified similar strengths and weaknesses, particularly around evidence, specification accessibility, commercial readiness, and recommendation confidence.

The five dimensions

The five dimensions sit beneath two primary pillars:

- Accessibility Signals
- Evidential Signals

Accessibility Signals

Signals
Extraction

Can AI access the content without additional processing?

The more content that is directly available in the initial page response, the easier it is for search engines, AI systems, crawlers, and enterprise agents to acquire and use consistently. Content that only becomes available after rendering, interaction, or additional processing introduces dependency and increases the risk of incomplete understanding.
In our diagnostics, most organizations with reasonable technical foundations score between 55 and 90. Lower scores are typically caused by critical content that only becomes available after rendering or is separated from the primary page response.

Signals
Interpretation

Can AI understand the relationships between content?

AI systems rely on structure to understand how information relates to the page, the product, and the user. Clear heading hierarchy, semantic sections, specification tables, definition lists, and well-labelled actions help systems interpret meaning consistently. When content relationships are expressed using generic markup, important context can become ambiguous or harder to associate correctly.
In our diagnostics, interpretation scores vary based on how clearly the page communicates relationships between concepts, specifications, evidence, commercial information, and actions.

Evidential Signals

Machines reward clarity with confidence. Brand recognition, market share, and reputation do not transfer to AI recommendations unless the evidence for them is present and accessible in your HTML.

Signals
Confidence

Can AI verify and support the claims being made?

Confidence is determined by the presence of specific, attributable evidence. Quantified results, named customers, certifications, specifications, test data, warranty information, and verifiable proof increase confidence. Generic marketing claims, unsupported assertions, and vague statements reduce it. Confidence does not measure brand popularity. It measures whether the page provides sufficient evidence for an AI system to confidently repeat, recommend, or compare the information without introducing assumptions.

Signals
Preference

Does the content provide clear reasons to recommend this option?

Preference emerges when the page clearly communicates why a product, service, or organization is different from available alternatives. AI systems compare evidence, suitability, specifications, proof, and positioning when generating recommendations. If competing content provides clearer answers, stronger evidence, or more explicit differentiation, it is more likely to be referenced or recommended.
Preference is driven by comparative clarity, not marketing language.

Signals
Conversion Utility

Can AI guide the next step?

AI systems are increasingly used to support evaluation and buying decisions. To do this effectively, the content must answer the practical questions that arise during the journey: suitability, pricing, implementation, proof, support, procurement, onboarding, and next steps.
When these details are missing, AI may understand the offering but cannot confidently guide users toward action. Strong actionability signals include clear commercial information, buying processes, support expectations, proof of outcomes, and explicit calls to action.

Cross-Industry Diagnostics

The observations in this paper are drawn from AI Signals diagnostics conducted by Content Bloom between April and May 2026. Fifty commercial pages across eight industry sectors were assessed and anonymized by vertical. The sample included product pages, service pages, product family pages, and informational articles.

Each page was analyzed using multiple AI systems to identify recurring patterns in how content is extracted, interpreted, trusted, and recommended. The objective was not to produce definitive scores, but to uncover where AI understanding consistently breaks down, much like a usability session seeks to identify where human users encounter friction.

A notable pattern emerged across almost every sector and page type. AI systems were generally able to find and extract content, but often struggled to confidently understand it, trust it, compare it, or recommend it. While the wording varied between AI systems, the underlying signals were remarkably consistent.
This finding reinforces the central argument of Digital Presenteeism: visibility does not guarantee understanding.

Sector	Pages Assessed	Typical Outcome
Consumer Goods & Retail	3	Moderate to Strong
Consumer Electronics	3	Moderate
Technology & SaaS	10	Moderate
Industrial & Manufacturing	11	Moderate
Financial Services	7	Weak to Moderate
Professional Services	7	Weak to Moderate
Life Sciences & Healthcare	6	Predominantly Weak
Travel & Hospitality	3	Predominantly Weak

Key Observation

- Across all sectors, the most common pattern was not poor visibility. In many cases AI systems were able to locate and retrieve content successfully.
- The challenge emerged later in the process. Interpretation, confidence, evidence, and recommendation readiness consistently lagged behind extraction.
- In simple terms, the content could often be found, but could not always be confidently understood.

Across industries, products, services, and page types, AI systems appeared to respond consistently to similar information patterns. Pages containing clear propositions, explicit evidence, structured specifications, identifiable audiences, commercial clarity, and obvious next steps were consistently easier for AI systems to interpret, trust, compare, and recommend.

While no universal formula exists, the recurring nature of these patterns suggests that AI systems may be looking for many of the same signals regardless of industry or platform.

Visibility Does Not Guarantee Interpretation

Content being present does not guarantee it will be interpreted in the way intended. Across multiple assessments, information that was easy for humans to find was sometimes interpreted differently by AI systems.

The reason is structural: humans and AI systems do not experience content in the same way.

Human And AI Experiences Are Different

Perhaps the most significant observation is that humans and AI systems do not experience digital content in the same way. Humans browse experiences. AI systems interpret signals. Where humans rely on navigation, context, design, and exploration, AI systems depend more heavily on the information and signals available during retrieval and interpretation.

Structure Is Still An Afterthought

Despite years of discussion around structured content, metadata, and Schema.org, many websites still rely heavily on presentation and inference rather than explicitly communicating meaning through machine-readable signals.

This increases ambiguity and places greater responsibility on AI systems to infer meaning rather than consume it directly.

Modern Websites Depend Increasingly On Rendering

A growing proportion of website content is delivered dynamically through JavaScript and client-side rendering, creating additional layers between information publication and information retrieval.

While often beneficial for user experience, these approaches can create additional complexity for systems attempting to retrieve and interpret content.

Commercial Information Is Often Weakly Structured

Product details, specifications, offers, and differentiators frequently exist but are embedded within experiences designed for humans rather than structured for reliable interpretation by automated systems.

Technical specifications, case studies, and performance data are frequently locked inside PDFs, present for humans who click, but less visible to AI systems making their first pass.

From Visibility to Understanding

Closing the Digital Presenteeism gap is not a purely technical initiative, nor is it a purely content initiative. It sits at the intersection of both, and that organizational no-man's-land is precisely why it goes unaddressed in most enterprises until performance metrics make it unavoidable.

The AI Signals Framework

Those building websites in the mid-1990s will recognize this moment. Different browsers rendered pages differently. Internet Explorer 2 and Netscape Navigator interpreted the same HTML in ways that produced entirely different experiences. The organizations that figured out how to build consistently across that fragmented environment gained an advantage that compounded over the years.

The current state of AI retrieval resembles the early browser era. Different systems often interpret the same content differently, much as Internet Explorer and Netscape once rendered the same HTML in different ways. During our diagnostics, scores varied mildly between AI systems, but the underlying weaknesses were remarkably consistent. Missing evidence, inaccessible specifications, weak differentiation, and unclear commercial pathways appeared repeatedly regardless of platform.

Homepage

clear extractable value proposition, canonical identity AI systems can anchor to, structured navigation signals

Landing page

explicit outcomes, comparative framing, commercial specificity, what it is, who it is for, what happens next

Services page

named methodologies, delivery evidence, structured descriptions enabling AI extraction and comparison

Product page

quantified specifications, named technologies, performance data, and proof points embedded in HTML, not behind tabs or inside PDFs

The implication is clear: the goal is not to optimize for a single AI model. As organizations once adopted web standards to create more consistent experiences across browsers, they should now focus on reducing ambiguity and strengthening evidence so any AI system has the best chance of reaching the intended conclusion.

AU (Agentic Understanding) - Native architecture

AU (Agentic Understanding) - Native architecture is not a technical initiative handed to developers, nor a content refresh handed to marketing. It is an organizational design decision about how information is structured, governed, and maintained across the content lifecycle.

The organizations performing best in AI-mediated discovery are not treating Agentic Understanding as a project. They are treating it as a standard.

Readers interested in the broader operational disciplines behind Agentic Understanding may find useful parallels in **Content as Infrastructure**, which explores how governance, architecture, ownership, and content operations influence the quality and usability of information across digital channels.

Where the Gap Sits in Organizations

The clearest pages for AI systems are, in practice, usually the clearest pages for people too.

For most organizations, this gap goes unaddressed because it sits between teams. Content teams do not own technical delivery. Technical teams do not own commercial messaging. Neither team is accountable for AI recommendation inclusion. The evidence in this report suggests this gap will remain unaddressed in most organizations until it appears in quarterly results: compressed deal cycles, shorter consideration sets, and buyers who have already decided before they reach the website.

Content Team	Technical Team	No Ownership of AI Visibility	AI Excludes Your Company	Invisible Revenue Loss
Creates messaging & content	Delivers infrastructure & pages	Nobody ensures AI can retrieve, interpret, and recommend the content	Not selected during AI-assisted research	Fewer opportunities enter the pipeline

So what do we do for now?

You cannot control which AI system your next buyer uses. You cannot dictate whether they are using Claude, ChatGPT, Gemini, Copilot, or something that does not yet exist. What you can control is whether your content gives them a fair chance of understanding you correctly.

Across our diagnostics, a consistent pattern emerged. While AI systems varied mildly in how they evaluated content, they repeatedly struggled with the same underlying weaknesses: missing evidence, inaccessible information, weak differentiation, unclear commercial pathways, and fragmented content.

The implication is not that organizations should optimize for a specific AI platform. They should focus on the information patterns that consistently help AI systems extract, interpret, trust, compare, and recommend content with confidence.

Organizations that reduce ambiguity and expose those patterns clearly give any AI system the best possible chance of reaching the intended conclusion.

Recommendations

Closing the Agentic Understanding gap requires alignment across content strategy, technical delivery, and operational governance. The organizations performing best in AI-mediated discovery are not approaching this as a standalone AI initiative. They are improving the operational maturity of the content ecosystem that sits behind every website, document, product page, and AI response.

The starting point is measurement. Without a scored baseline across your highest-value commercial pages, remediation becomes guesswork. You end up fixing what feels broken rather than what is actually preventing AI systems from retrieving, understanding, and recommending your content.

Once that baseline is established, the highest-return intervention is usually straightforward: move evidence from PDFs into HTML. Most organizations already possess the proof points AI systems need: quantified outcomes, technical specifications, customer success stories, certifications, and performance data. The challenge is that much of this information remains locked in formats that create retrieval risk.

From there, the focus shifts to evidential strength. Audit commercial pages for language that sounds persuasive to humans but provides little value to AI systems. Replace broad claims with specific, verifiable evidence. Named clients, measurable outcomes, quantified improvements, and clear differentiators consistently produce the largest improvements in confidence and recommendation potential.

The technical layer runs alongside this work. Rendering gaps, semantic structure, heading hierarchy, schema markup, and content accessibility all influence how effectively AI systems interpret information. These are rarely the primary problem, but they often amplify every other weakness when left unresolved.

Ultimately, the goal is not to optimise for a particular AI engine or chase short-term platform behaviour. It is to establish AU-Native architecture as an operational standard. When content is structured, governed, and maintained as a business asset, both people and AI systems are better able to find it, understand it, trust it, and use it. The result is not simply improved Agentic Understanding. It is a more resilient content ecosystem capable of supporting whatever comes next.

Conclusion

Digital Presenteeism is not a technical problem. It is an understanding problem.

For decades, organizations invested in usability testing and user experience practices because they learned a hard lesson: what a company believes it has communicated and what a customer actually understands are often two entirely different things. UX workshops existed to close that gap, structured, deliberate attempts to see your own content through someone else's eyes.

That customer set now also contains agentic systems.

This is not a radical shift in discipline, it is the same instinct applied to a new audience. When a procurement engineer delegates early-stage research to an AI system, that system becomes the first reader, the first evaluator, and the first filter. It does not browse the way a human does. It extracts, interprets, compares, and recommends based on the signals your content exposes at retrieval time.

Signals is a tool, not a prediction engine, it doesn't tell you how every AI model will behave, and no honest methodology could make that claim since models differ, retrieval behaviours vary, and the landscape will keep shifting. What Signals does is what a good usability session always does: it gives you enough SIGNAL to act. Five users in a UX workshop never produced statistical certainty. They produced clarity about where understanding broke down. A structured AI pass does the same.

The organizations that will perform best in AI-mediated discovery are not the ones that optimize for a specific engine. They are the ones who reduce ambiguity consistently, across content types, page structures, and formats so that whatever system encounters their content has the best possible chance of interpreting it the way it was intended.

The question is no longer whether your content can be found.
The question is whether it is actually working

Evidence & Methodology

The observations presented in this paper are based on AI Signals diagnostics conducted across fifty enterprise pages spanning manufacturing, industrial, technology, healthcare, financial services, professional services, travel, and consumer sectors.

The assessment included a mixture of product pages, service pages, product family pages, and informational articles to ensure findings were not limited to a single content type.

Each page was assessed using multiple AI systems against five dimensions:

- Extraction
- Interpretation
- Confidence
- Preference
- Conversion Utility

The assessment considered content accessibility, structural clarity, semantic organisation, evidence availability, commercial pathways, trust signals, and machine-readable indicators available at the time of analysis.

The objective was not to predict how every AI system will respond, nor to produce definitive performance scores. Rather, the framework was designed to identify recurring conditions that strengthen or weaken AI understanding, recommendation confidence, and commercial usability.

The findings in this paper represent patterns consistently observed across multiple organisations, industries, and page types rather than isolated examples. As content, platforms, and AI systems evolve, individual results will naturally change over time. The intention is not to provide a permanent score, but to highlight the signals that most commonly influence how modern AI systems interpret and evaluate digital content.

MEET OUR TEAM

Nickoli Roussakov
Chief Executive Officer

Andy Iddon
Chief Strategy Officer

Jon Primmer
VP of Technology

Umar Akhtar
New Business Director

Nicholas Wetmore
Senior Technology Manager

Morgan Henley
Senior Technical Consultant

Pankaj Gaur
Director India/APAC
Operations

What's your next move before AI makes it for you?