MRS report / 2026

Sponsored by

VERVE

Who owns understanding?

How AI is reshaping value, expertise and accountability in research

MRS Delphi Group

MRS Evidence Matters

Contents

About MRS Delphi Group

The MRS Delphi Group is led by a coalition of respected thinkers in the marketing and research sectors. The Group delivers valuable insight across a range of important business, social and political issues. The Steering Board includes: Colin Strong, Ipsos (chair), Rose Tomlins, MTM; Gemma Proctor, PA Consulting; Adrian Sanger, Fore Consulting: Patrick Alcantara, FCA; Zoe Ruffels, Imperial Brands; Dr Mark Thorpe, Truth Consulting; Ian Ralph, Walnut; Edith Bardin, Shell; Sarah De Caux, Co-op; Eddie O'Brien, Sage; Josephine Hansom, Youth Researcher; Chrissie Tarbotton, Boxclever; Sarah Mills, Haleon; Jane Frost, CEO of MRS. MRS would like to thank all the contributors included in this report. The Steering Board members do not necessarily agree with all the comments and conclusions within this report, and it cannot be taken as directly representing their views.

VERVE

Thanks to our report sponsor **Verve**.

The sponsor has not had prior approval or review of this report before publication, and does not endorse any individual's perspective.

www.mrs.org.uk/delphi

VERVE

INTELLIGENT PERSONAS & SIMULATIONS (VIPS)

FOR DECISIONS OF CONSEQUENCE

AUDITABLE NOT "BLACK BOX"

ACCURATE NOT JUST PLAUSIBLE

PREDICITIVE VALIDITY >90%

QUANT NOT JUST QUAL

CREATING A SOURCE OF **DURABLE COMPETITIVE ADVANTAGE**

EMPOWERING **CLIENTS**
EMPOWERING **PARTNERS**

VERVE | TURN **HUMAN TRUTH** INTO **SUPERHUMAN ADVANTAGE**™

TALK TO US

TO UNDERSTAND MORE

▶ WWW.VERVEVERO.COM

MRS™ Awards 2023 Winner Cultural Insights

THE AURAs 2024 Innovation of the Year WINNER

2025 THE ESOMAR AWARDS Excellence in AI & Automation VERVE | MARS

2025 THE ESOMAR AWARDS Outstanding Contribution to Consumer Insights VERVE | SAMSUNG

Introduction

Colin Strong FMRS
Head of Behavioural Science
Ipsos
Chair of MRS Delphi Group

When the MRS Delphi Group published the *BEST Framework for GenAI* in 2023, if felt as though the industry was still trying to find a way through a fast-moving and what felt at times like an over-hyped conversation.

Generative AI had arrived with extraordinary promise, but also with uncertainty about where it could be used, where it should be used, and what risks needed to be managed.

BEST (see below) was designed to help with that task. It encouraged practitioners to begin not with the technology, but with the problem. Was the challenge familiar or new? Was it systematic and predictable, or unsystematic and unpredictable? From those questions came four broad modes of application: Boost, Expand, Shift and Transform.

The aim was not to predict the future, but to help researchers make better choices about when and how to apply AI, based on the nature of the question they were trying to answer. This was, and remains, a useful way of keeping the discussion problem-centred rather than tool-centred.

This new Delphi Group report picks up that conversation from a different point in our industry's development. It is not intended

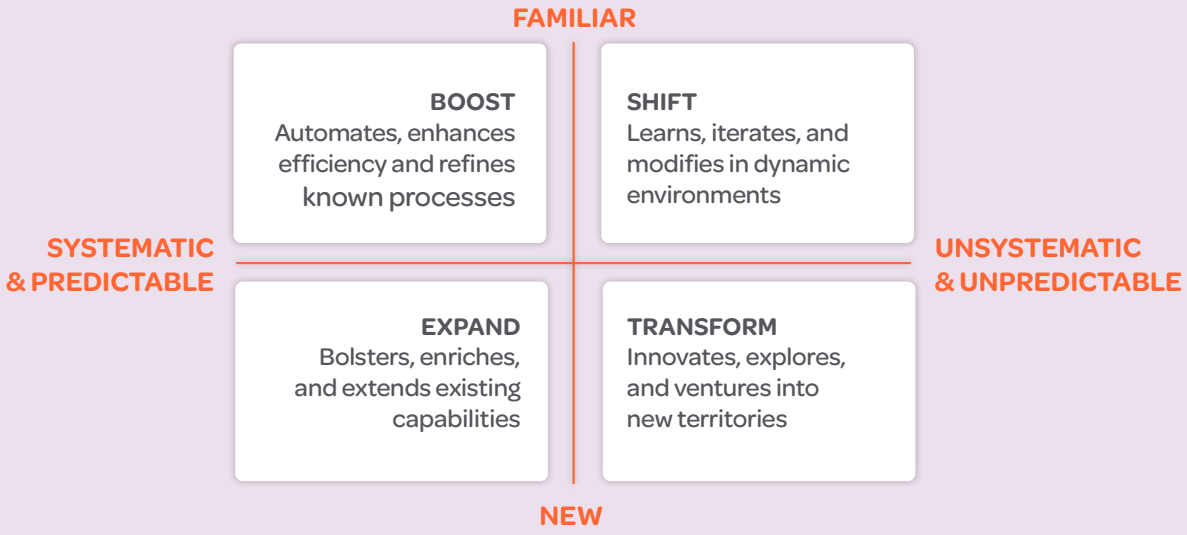

Human involvement cannot simply mean that someone reviews, approves or signs off an AI-generated output. If the human does not understand the assumptions, limitations and consequences of the system, they are not meaningfully in the loop.

as a definitive 'state of the nation' account of AI in research. The technology is moving too quickly, and practice is too varied, for that to be either possible or especially useful. Any fixed account of where the sector is today may be quickly overtaken by what happens tomorrow.

Instead, this report brings together a range of practitioner perspectives from people working across the industry as they grapple with what GenAI is beginning to mean for research, insight, evidence and decision-making.

That distinction matters. This is not a report that says 'this is where the industry has landed.' Instead it is a report that asks: 'what are we noticing, what are we worried about, what are we learning, and what questions now need to be taken more seriously.'

Across the articles, a complex but broadly consistent picture emerges. AI is changing who owns understanding inside organisations. When insight tools can be built, deployed and used by many different teams, authority no longer comes only from controlling access to data or owning the production process. It comes from knowing which outputs can be trusted, what they mean, and who is accountable when decisions are made from them.

The role for humans

AI is also forcing the sector to think harder about 'human centricity'. Research has never been only a process of extracting patterns from data.

It is a discipline of context, interpretation and judgement, rooted in the lives and experiences of real people with often messy lives. As AI becomes better at summarising, clustering and producing fluent outputs, the risk is not simply that people are replaced. It is that our understanding of human experience is thinned out: reduced to signals, averages and plausible narratives that lose contact with the lived experiences that research is meant to represent.

A similar challenge appears in the growing use of synthetic audiences. These tools can be powerful for exploration, iteration and early-stage hypothesis generation.

But they can also create a verification challenge. The more fluent and coherent an output appears, the easier it becomes to mistake plausibility for proof. That places a renewed burden on the profession to define when synthetic evidence is acceptable, when it is not, and what standards of validation are required before it becomes decision-grade.

The same issue runs through the notion of humans in the loop. Human involvement cannot simply mean that someone reviews, approves or signs off an AI-generated output. If the human does not understand the assumptions, limitations and consequences of the system, they are not meaningfully in the loop.

Instead, they may simply be carrying the liability. In higher-consequence contexts, the question is not only whether AI is useful, but whether in fact the people around it have the expertise, authority and confidence to challenge it.

The agency model

AI is also reshaping the commercial structure of research. For decades, agencies have often been paid for the visible labour of research: interviews conducted, data collected, transcripts coded, analysis completed and reports written. AI can now perform many of these tasks quickly and competitively. That does not mean research is disappearing. It means the value is moving. The future of the sector will depend less on defending process and more on making explicit the value of expertise, judgement, interpretation and decision impact.

This has implications for hybrid practice too. 'AI plus human intelligence' is a complex concept which requires a high level of specificity in its application. AI can accelerate workflows, extend evidence bases and help generate possible interpretations. But human researchers remain responsible for deciding which meanings are valid, useful, ethical and safe to act on.

Hybrid intelligence is not simply a division of labour between fast machines and thoughtful humans. It needs to be a designed practice in which evidence, judgement and accountability are deliberately held together.

When outputs become abundant, fluent and cheap, the scarce value is understanding: knowing what to ask, what to trust, what is missing, what matters, and what should follow.

The purpose of research

Finally, the report returns to the purpose of research itself. If AI can produce answers on demand, what needs to remain distinctively human about research? One answer is curiosity.

Research begins before there is a stable prompt. It begins with noticing something unresolved, uncomfortable or strange; with caring about who is represented, who is missing, and what consequences follow from the conclusions we draw. AI can help us answer questions. But it does not decide which questions matter.

The common thread running through these contributions is that AI does not make research less important but is does make the standards of research more important. When outputs become abundant, fluent and cheap, the scarce value is understanding: knowing what to ask, what to trust, what is missing, what matters, and what should follow.

In 2023, BEST helped us think about how and when to build AI into research processes.

This new report is more exploratory and more reflective. It brings together different voices from across the sector to consider what kind of research profession we need when AI is already inside those processes.

The answer is unlikely to come from one framework, one method or one institutional position but it will emerge from continued experimentation, challenge and debate.

Nevertheless, the direction of travel is clear. The future of research will not be secured by speed alone and instead will depend on how well we define, protect and develop the forms of human expertise that make insight worth acting on.

AI and authority

AI and authority
Who owns understanding?

Ian Ralph

Insight Innovation

Walnut

Rose Tomlins

Director

MTM

There is a risk insight teams become positioned as late-stage validators rather than upstream authorities.

Within an AI-enabled economy it is the truth that is scarce, amongst a wealth of insight.

Every innovation throughout the information age has led to a shift in who gets to authorise this truth, be that the printing press or the internet. And as we face the next seismic shift in our profession, we must seize control and define the standards of understanding ourselves, and not allow others to define them through speed, cost and convenience.

It would be easy to view this as a threat to our profession – a shift where the insights function is treated primarily as a quality-assurance unit, auditing and verifying models and outputs.

We argue that it in fact reshapes the dynamics in our favour as the human skills required in the insights space shift from production to discernment: knowing which questions matter, which outputs to trust, and what it all means for the business.

A shared accountability

There is a danger that authority over insight migrates away from those trained to understand evidence, method and human context as AI tools are sold to senior leaders, procurement teams and commercial functions whose focus is on speed, scale and cost reduction. On this basis, there is a risk insight teams become positioned as late-stage validators rather than upstream authorities.

True authority must be redefined as stewardship rather than ownership of insight. By reshaping our roles and focussing on what we do best, we can share the burden of authority and safeguard true understanding.

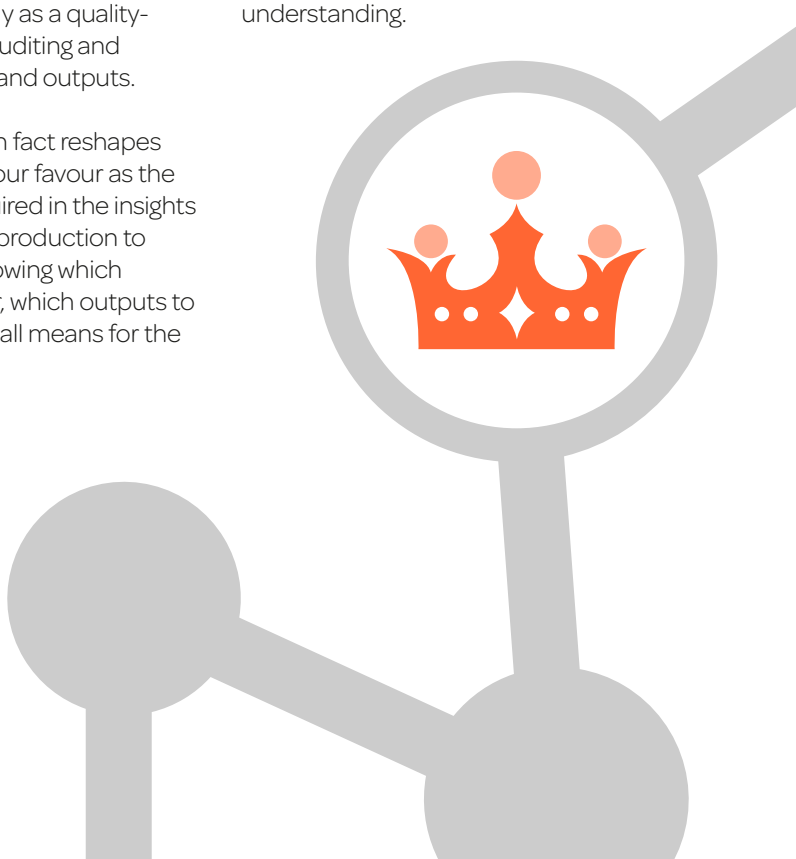

The AI insight stack

The 'AI Insight Stack' (see diagram) helps explain this shift. Standard AI tools focus on three operational layers: the data, the model, and the user-interface. What they lack is this crucial 'understanding layer' that our profession provides. This human understanding layer relies on three core structural responsibilities: Context (what the tool represents), Judgement (confidence in the output), and Accountability (who owns the risk if a decision fails). Each of which align very nicely with the three key insight groups of insight agencies, platform developers and client teams.

The real value of a researcher does not lie in policing how a tool is built, but in managing the commercial risk of the decisions made from it. We need a system where accountability is shared across the research pipeline, rather than forcing researchers to check every automated report.

By dividing these tasks, each partner carries a specific piece of corporate liability through clear structural handoffs. This shared accountability allows each organisation to safely sign off on any insights tool, anchoring final authorisation in our core professional principles: transparency, confidentiality, and independent judgement.

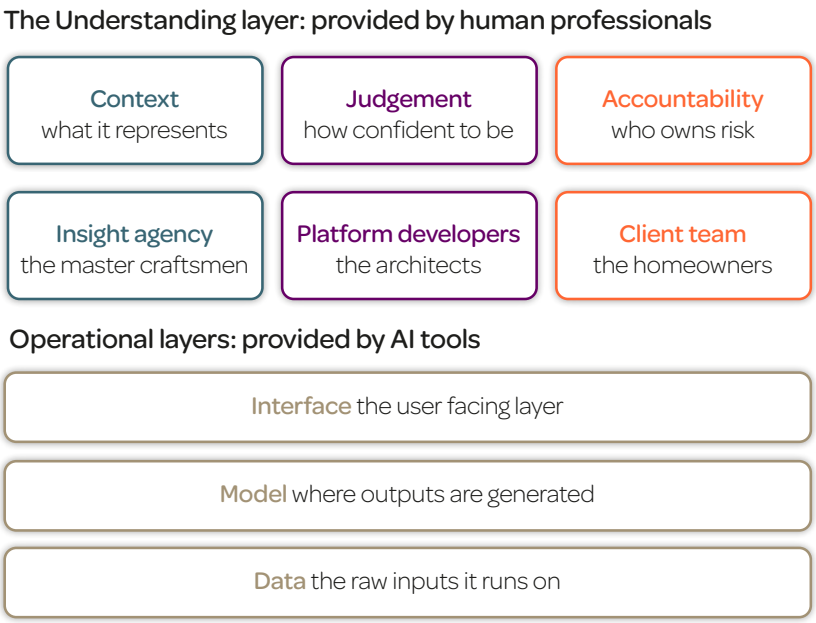

Three clear roles for a new hierarchy

To see how this works in practice, picture building a safe, stable home for the future of research. A reliable home does not rely on one person checking every nail; it relies on a trusted ecosystem of experts. Applying this logic to AI, the human understanding layer (Context, Judgement and Accountability) is maintained through a clear division of labour:

1. **The Platform Developers (The Architects)**

Developers and consultancies building enterprise tools must accept absolute liability for the blueprint. Their job is to design a secure framework, embedding corporate guardrails and compliant workflows into the platform. Just as an architect is legally accountable for the structural safety of a building, platform developers must get the digital architecture right, because at scale errors propagate as readily as insights.

By designing clear operational signals into the blueprint, they provide the foundation for judgement, helping the organisation gauge how confident it should be in the tool's output. They are liable for ensuring the system processes data safely, protects privacy, and operates reliably within set guardrails.

From:
building tools that generate insight

To:
building systems that signal which insights can be trusted and for what

It's not just about what data to use, but how, when and why to apply each source of data and how to get the mix right for each use-case.

2. The Insights Agency (The Master Craftsmen)

An architect's blueprint is only as good as the raw materials used to execute it. Insights agencies are the specialised craftsmen who source, carve and verify the baseline materials: the human data. Working to strict industry standards, they ensure only genuine, robust data feeds the machine, giving the framework its real-world integrity. Their role is to verify the training data and baseline contexts to ensure they are accurate and free from systemic bias and being used for the right job-at-hand.

Like any master craftsman this skill is nuanced. Human understanding it not only about accuracy, but plausibility. Human behaviours are not fixed, but contextual. And perhaps most importantly, needs and attitudes are transient and not set in stone within a foundation of training data. It's not just about what data to use, but how, when and why to apply each source of data and how to get the mix right for each use-case.

By doing this they protect the context of the research, confirming what the tool seeks to synthetically represent, and for whom. In an automated world, their craft shifts from interrogating people to generate answers, to interrogating answers to ensure they can be trusted.

From:
generating data about consumer behaviour

To:
ensuring data remains anchored in real & relevant behaviours

3. The Client Team (The Homeowners)

The client team need not lay bricks or study technical schematics, nor worry about how the tool was built. Instead, they focus on the space they will live in, defining how the house must fit their daily lives. The internal insights team owns the living brief, setting the guardrails, standards and risk tiers for how research is conducted across the enterprise. Authority in the AI era must be risk-tiered: the more consequential the decision being made, the more explicit the standards of validation and human judgement must become. Only the internal insights teams have the required knowledge to make these calls.

The internal insights team owns the living brief, setting the guardrails, standards and risk tiers for how research is conducted across the enterprise.

The moment a tool's outputs miss the commercial context, lack emotional depth, or fail to meet the organisation's needs, the client team demands course correction, or halts production entirely.

Because they carry the ultimate financial risk and must live with the final strategic decisions, their job is to act as the authority on usage. They validate how research is deployed, monitoring the build to ensure it delivers real-world business utility. The moment a tool's outputs miss the commercial context, lack emotional depth, or fail to meet the organisation's needs, the client team demands course correction, or halts production entirely. They wield the supreme corporate veto because they own the ultimate accountability, standing behind the final recommendation.

From:
controlling distribution of and access to insight

To:
validating when insights are safe to act upon

Upholding our principles: trust without re-verification

This collaborative framework is not a radical departure; it maintains the ethos with which clients, agencies and platform developers have always partnered. In a primary research project, a client-side buyer is never expected to check every raw interview record to verify its quality. Instead, the relationship relies on established professional standards and a shared commitment to robust outputs. Generative AI does not change these fundamentals; it requires us to refresh how we apply them.

By clearly dividing responsibilities, we uphold those core principles of our insights profession: transparency, confidentiality, and independent judgment When platform developers and agencies are fully transparent about how their tools are built and their models trained, the client team is freed from the operational trap of double-checking code.

The future authority of the insight profession will not come from owning every tool, checking every output or blocking every experiment. It will come from defining the conditions under which evidence can be trusted, decisions can be justified and human reality remains in view.

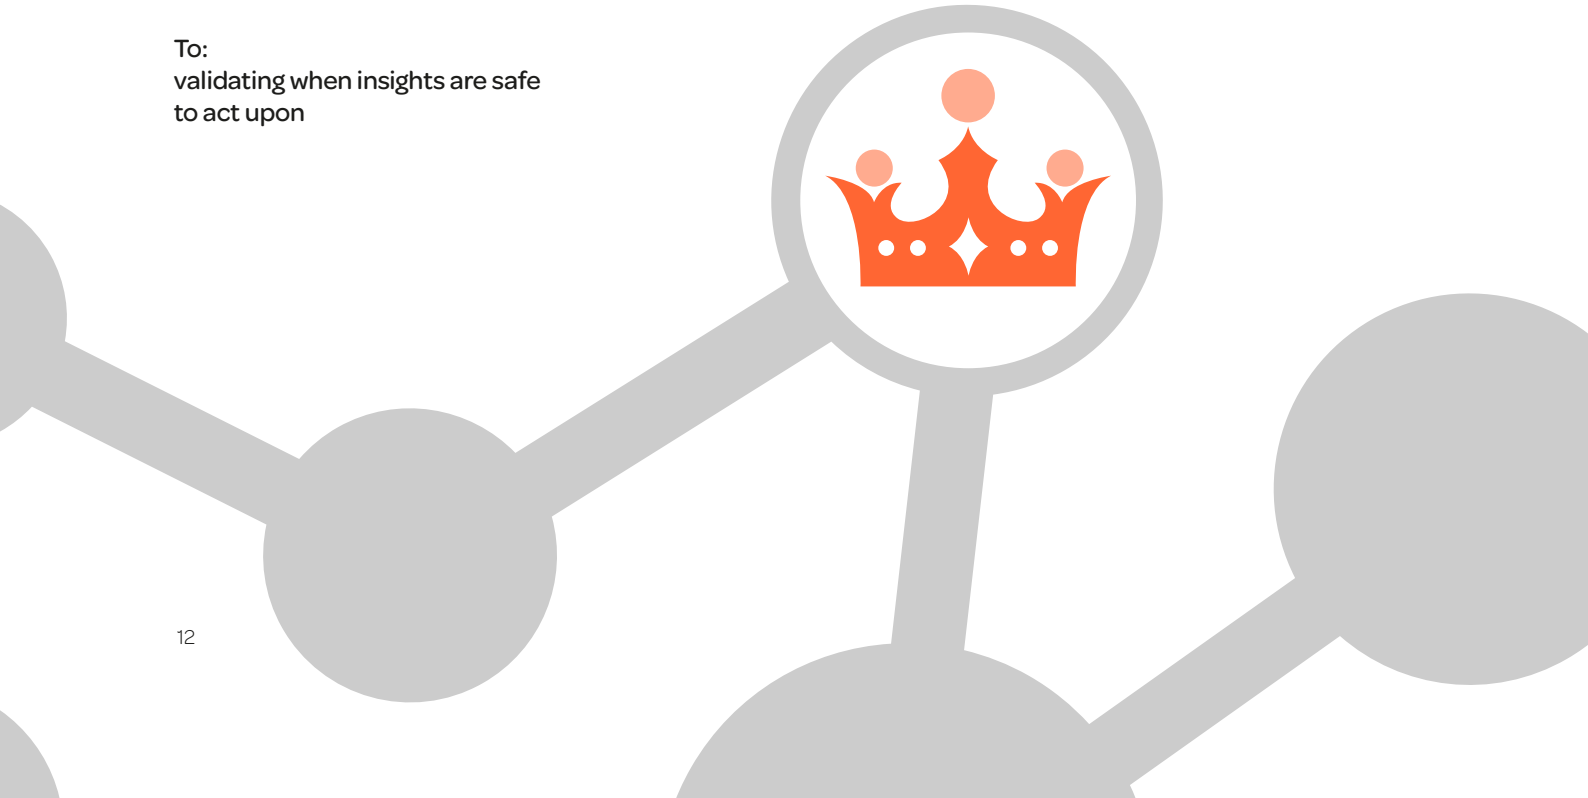

AI and the human role

AI and the human role
Why reality loops really matter

Dr Mark Thorpe
Director
Truth Consulting

It would be fair to say that the emergence and rise of AI in the research and insights sector has generated strong reactions.

Some of these reactions have been highly positive, some more circumspect, and others more negative, driven by concern, anxiety and/or trepidation about the repercussions for an industry that is so fundamentally founded on people.

Change can be problematic, especially when rooted in technology. The reality is that AI is here to stay; our challenge is to ensure that it is incorporated in such a way as to genuinely boost the sector, as well as helping to provide enhanced outcomes for both agencies and clients.

There is, however, an existential risk in implementing and using AI in our sector without due regard to outcomes and consequences. There is jeopardy in the fact that AI carries the risk of disconnecting 'insight' from the realities of everyday life. There is a fundamental break in the 'reality loop' that connects us to the human world we purport to represent.

Data should represent people not replace them

Some of the narratives around AI and the 'production' of synthetic data provide important clues as to potential jeopardies we need to be conscious of. There are **three key areas** that are particularly important:

1. **'Data' should help us understand real people.** This may sound an obvious statement to make, but it does seem as though there is a danger of losing sight of this 'realness' of everyday life in our haste for data at speed and with extended reach/scale. The capturing of data becomes the be-all-and-end-all, a virtue in and of itself, rather than a reflection of fundamental human truths.
2. **Data is meaningless without life.** Something to avoid in our incorporation if AI is the situation where data is collected that is lifeless. By this we mean that there is an absence of real human meaning and a decoupling of the captured data from the lived experience that it was either generated from or created to represent. We should reflect on some of the learnings from the big data boom of a decade or so ago. Too often the endpoint was a mass of de-contextualised information that impressed with its scale but was highly elusive in terms of real context, meaning and application.

The lesson from the big data boom was not that scale lacks value but that scale without context can create the illusion of understanding. AI risks repeating that mistake, only faster and with more fluent outputs.

3. **'Reality loops' are central, not an option.**
No one method or approach should be used as either fully representative or validating. We need to remind ourselves that 'triangulation' – the use of multiple approaches, perspectives and/or methodologies is core to best practice. We seem to be drifting away from this, and it is, ultimately, counterproductive. When we have triangulation in the research process, we are effectively introducing what can be described as 'reality loops'. When using AI in research and insight, we need to recognise the importance of these necessary reality loops.

When everyone is a loser

People matter, not just for moral reasons; they matter also because they/we are the focus of enquiry, whether as 'shoppers', 'consumers' 'targets', or simply 'our audience'. People, real human beings, buy products and services, shop in stores, they are loyal or fickle in relation to the brands we represent and build, and are our constant opportunity as well as potential jeopardy. Data should serve and support this reality, not stand aloof from it.

When data becomes the end-goal, everyone is a loser. This is the point where 'reality drift' really kicks in. We can see this happening everywhere. Data is presented, compared, annualised. It gets bigger, broader, faster, louder. But we lose human centricity. The landscape of meaning becomes thinner, flattened by the absence of real human meaning. We make judgements and comparisons with data that is bereft of life.

It is also important to remember that there is another jeopardy in this scenario. Without 'reality loops', which act as both context and a sense-check for data, we run the risk of transforming data lakes into swamps; a steady flow of error and inaccuracies polluting the sources we increasingly draw upon. Clients lose in this scenario; big decisions are made based on inaccuracies and without meaningful context. Agencies lose because their role becomes one of process management, rather than interpretation and context creation. Participants, consumers, shoppers and the universe we research all loses – because representation is limited, partial and decontextualised.

When we have triangulation in the research process, we are effectively introducing what can be described as 'reality loops'. When using AI in research and insight, we need to recognise the importance of these necessary reality loops.

It does seem as though AI is often positioned as both approach and answer. AI becomes the story, the star of the show; a celebrity presence that dwarfs the original strategic intent. AI becomes both the medium and the message. Reality loops are broken.

When the medium is the message

It seems that much of the narrative around AI (currently) focuses intensely on what the technology is and can do. Outside the implied or stated critiques of 'traditional research', not much is said about understanding human beings, as complex and fluid entities. What we too often get is a valorisation of data, with speed and at scale, but little of the 'fuzzy brilliance' that makes human beings, and the lives they lead, so fascinating.

This is highly problematic. Life is messy. Data is as irrelevant as it is relevant. Context matters. Human understanding matters. Culture and context matters. Strip these away and what's left?

In contrast, and as example, qualitative research (broadly defined) aims to create a depth of understanding around foundational human emotions, passions, needs, ideals and beliefs. Qualitative research should never be the answer, it should always be the route through which we access the answers required.

It does seem as though AI is often positioned as both approach and answer. AI becomes the story, the star of the show; a celebrity presence that dwarfs the original strategic intent. AI becomes both the medium and the message. Reality loops are broken.

Step back and breathe

We all need to take a step back and think carefully about where we are now and our direction of travel in relation to AI. The reality is that AI could be harnessed to achieve wonderful things; but in a sector where understanding human beings is critical, it could also represent a damaging detour. We can become consumed by method and the process of capturing 'stuff'.

In the rush for newness, we risk losing sight of what really matters. As we fine-tune our ability to capture 'data' with greater speed and at scale, we must be cognisant of the 'what' and 'why' behind the things we do. Are we seeking scale or understanding? Can the two exist together? Will AI-enabled data bring us to the heart of meaning, or does it serve simply to raise more questions because it isn't twinned with a deeper, human understanding?

Ultimately, we need to accept that people matter and that they are central to the research and insight, even in the age of AI. If we lose sight of this reality, we risk our data becoming worthless; billions of points in a world without meaning

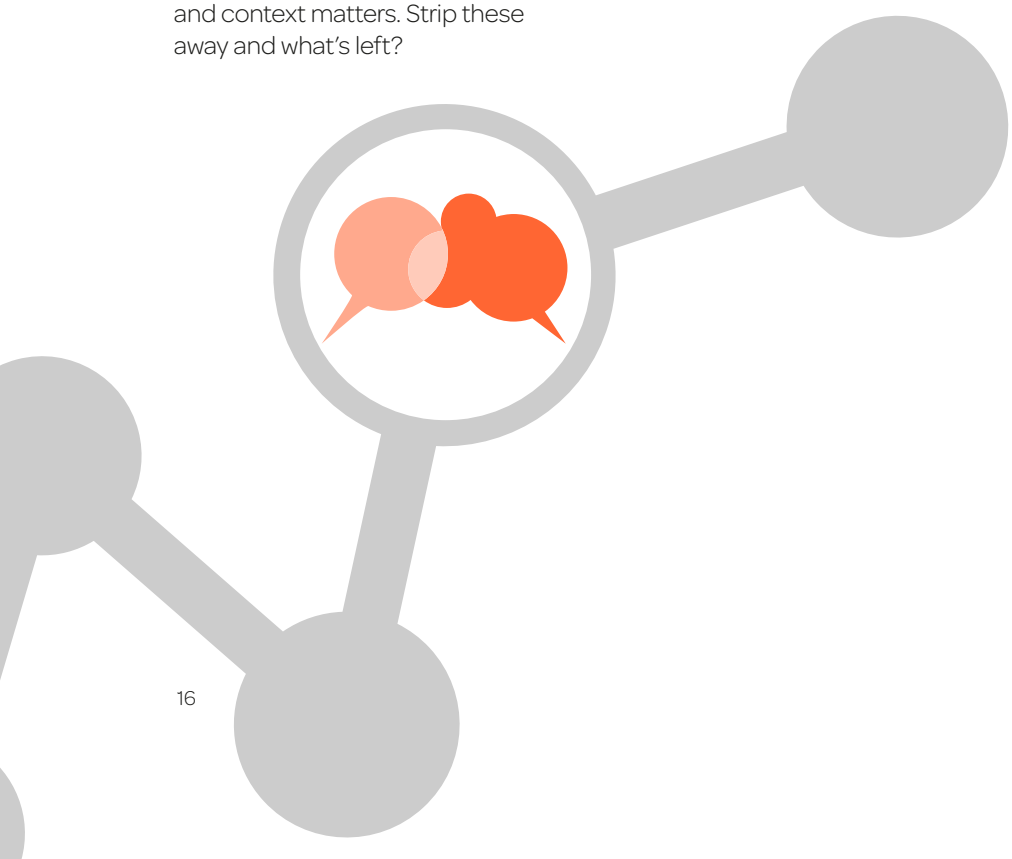

AI and discipline

AI and discipline
Why plausible AI demands stronger proof

Edith Bardin

Brand & Communications Insights Adviser

Shell

Rose Tomlins

Managing Director

MTM

This is the 'verification paradox'. As outputs become more convincing, their truth becomes harder to verify.

Market research is entering a new phase. For the first time, it is possible to generate what looks like fully formed insight without speaking to a single real person.

Synthetic audiences, AI-generated representations of people, are rapidly moving from experimentation into real-world application. They promise unprecedented speed, scale, and access to audiences once considered difficult, costly, or even impossible to reach.

But as these tools become more sophisticated, a new tension is emerging. Outputs are growing more fluent, more coherent, more convincingly human, yet our instinctive trust in them is not keeping pace. In fact, the opposite is true.

This is the 'verification paradox'. As outputs become more convincing, their truth becomes harder to verify.

At the same time, capital is pouring into the space, accelerating the proliferation of synthetic audience solutions, from simple prompted personas to increasingly complex, data-grounded simulations. On the surface, many deliver outputs that appear equally polished and confident, masking significant variation in underlying quality. And with many players optimised for generation rather than verification, speed and plausibility are too often prioritised over robustness and proof.

In this environment, the role of the market researcher becomes not less, but more critical. Their expertise is needed to shape inputs, validate outputs, and rigorously interrogate and synthesise what synthetic systems produce. Ultimately, it is this rigour that determines whether AI elevates insights or simply amplifies noise.

Given these risks, the question becomes not whether to use synthetic audiences, but when, and how.

When is synthetic acceptable?

Synthetic audiences can add real value when used in bounded, low-risk applications, such as testing early-stage ideas, rapid iteration, sense-checking, or hypothesis generation. However, they should not yet replace human research in high-stakes, behavioural, or strategic decisions. The more complex the question, the more important human grounding becomes. Synthetic audiences are therefore best understood as augmentation, not replacement.

Ensure synthetic audiences are grounded in recent, trusted inputs, such as IDIs, focus group learnings, surveys, and behavioural evidence, rather than generic internet data.

Understanding exactly why this grounding is required requires looking beneath the polished surface of AI outputs to observe the systemic gaps left by a shortcut approach.

Why plausibility does not equal truth

The current risk is that generation-optimized models are designed to be convincing and can easily trick our instincts. Early testing shows synthetic audiences can echo human research at a surface level, reflecting themes and even producing realistic verbatims. But beneath this, critical gaps remain:

- **The risk of the average:** Synthetic outputs tend toward the norm. By design, they synthesise patterns and generate what is most typical. Yet research value rarely sits in the average; it sits in tension, contradiction, and deviation. Synthetic audiences risk smoothing out these edges.
- **Missing the golden outlier**: The most valuable insights can come from the unexpected perspective, the "golden outlier" that challenges assumptions. Synthetic audiences, optimised for coherence, often underrepresent these voices.
- **Reduced emotional depth:** Synthetic responses can sound considered and articulate, but often lack emotional friction. They are typically more structured, quicker to resolve ambiguity, and less sceptical than real people, particularly on issues requiring judgement or trust. The result can be insight that feels polished, but lacks behavioural realism.
- **The synthetic gap**: Human understanding is shaped by lived experience. Synthetic audiences can approximate attitudes, but they do not experience them. This creates a gap between simulated response and real-world behaviour, especially in complex or context-heavy decisions.
- **Bias and overconfidence:** Synthetic outputs reflect the biases of their training data and are often skewed toward dominant, globalised perspectives. At the same time, they present their conclusions with confidence, regardless of underlying accuracy. Without validation, plausible outputs can easily be mistaken for truth.

Four rules for stress-testing synthetic audiences solutions

Shifting from plausibility to credibility requires rigorous stress-testing of synthetic outputs. While this seems non-negotiable, the field's infancy means it still lacks universally accepted standards and measurement frameworks.

At a minimum, practitioners should consider the following four actions in the early stages of evaluation:

1. **Anchor models in real, proprietary data:** Ensure synthetic audiences are grounded in recent, trusted inputs, such as IDIs, focus group learnings, surveys, and behavioural evidence, rather than generic internet data. Without this, outputs risk defaulting to undifferentiated patterns and collapsing into 'AI data swamps'.
2. **Validate against real-world behaviour and context** Ensure synthetic insights reflect what people actually do, not just what models predict, by testing them against observed behavioural data and checking they remain credible across different cultural, economic, and situational conditions

3. **Benchmark against human research to assess methodological fit** Synthetic and human approaches should be run in parallel wherever possible. Comparing outputs side by side reveals alignment, but more importantly, highlights divergence and risk.

4. **Interrogate difference, not just similarity:** Validation should not stop at "does this match?" Researchers should also challenge the output and ask "what is missing" and "which tensions or contradictions are absent." It is in these differences that synthetic limitations (such as being 'too polite' or 'too generic') become visible.

A shared mandate for accountability

Market research practitioners, agency and client side, can keep each other accountable by continuously pushing for rigour and high standards.

For agencies	For clients
– Embed structured stress-testing into every synthetic audience project. – Introduce a grounding mandate: no synthetic output without validation. – Actively design for divergence to surface outliers. – Be explicit about limitations (e.g., emotional depth, behaviour, context). – Be transparent with clients on methods and assumptions, specify how the data is built and what % is LLM knowledge vs proprietary data	– Interrogate the process, not just the output. – Demand evidence of validation, especially for high-stakes decisions. – Challenge for what's missing, not just what aligns. – Define clear internal standards for when synthetic is acceptable. – Keep the expertise and ownership with the insights function. The researcher's judgement in persona design, analysis and interpretation is critical

For Both: Uphold "good data in, good data out"; treat synthetic audiences as an extension of research, not a shortcut; and strengthen governance around bias, data, data protection and privacy, and ethics while scrutinising vendors for both AI capability and research expertise.

Conclusion: Reasserting principles

For the market research community, this is a pivotal moment, an opportunity to shape how synthetic audiences are applied within our discipline. By grounding these approaches in real-world data, insisting on transparent and repeatable validation, and applying disciplined critical thinking to every output, we can turn potential risk into a genuine competitive advantage.

The rewards are clear: faster insight without compromising trust, broader coverage without sacrificing accuracy, and innovation that remains firmly anchored in reality.

Ultimately, synthetic audiences will reflect the standards we set. If we uphold the fundamentals of good research, they have the potential to not dilute our discipline, but strengthen it, extending what market research can achieve while protecting the credibility on which it depends.

The future of these Synthetic audience tools, and the trust in our insights, will be defined by the rigour we choose to apply today.

The rewards are clear: faster insight without compromising trust, broader coverage without sacrificing accuracy, and innovation that remains firmly anchored in reality.

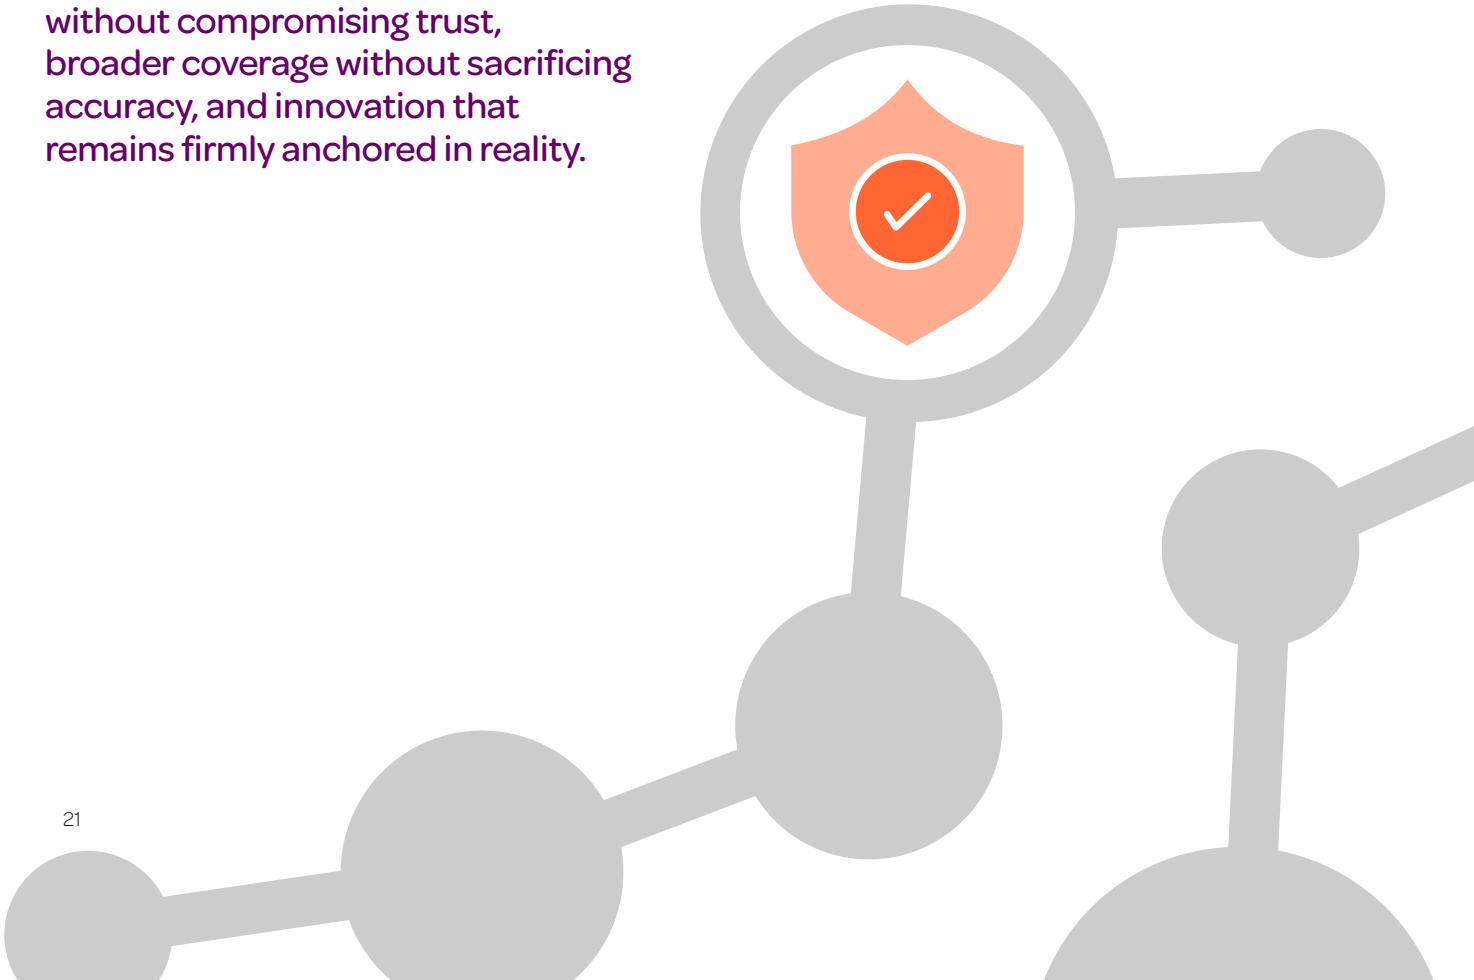


AI and positioning

AI and positioning

AI agents don't go to prison. Insight professionals carry the consequences

Eddie O'Brien

Senior Director, Global Customer Insight & Experience

Sage

Dr Ben Warner

Co-founder

Electric Twin

When AI systems generate plausible, polished narratives at speed, humans increasingly approve rather than interrogate.

The market research industry is in the middle of a profound and uncomfortable transition. Generative AI, synthetic respondents, and agentic insight systems are dramatically accelerating how research is generated, synthesised, and consumed.

What once took weeks of fieldwork, analysis, and debate can now be produced in hours. Concepts are ranked instantly, messaging is optimised in near-real time, and dashboards offer clean, confident answers that feel decisively "good enough."

For organisations under constant pressure to move faster, this shift is seductive and in many respects genuinely valuable. But it is also quietly changing the role of research from a support function with time for considered judgement to part of the production line for decision-making.

Growing responsibility

Historically, market research informed decisions without owning them. Researchers provided evidence, context, texture, and challenge, fully aware that leadership judgement sat downstream. Today, AI-powered insight increasingly functions as a decision trigger rather than a decision input. Synthetic respondents rank concepts and price points. LLMs collapse thousands of verbatims into confident themes. Agentic systems recommend winners, losers, and next actions.

In many organisations, these outputs flow directly into execution; creative selection, pricing changes, product roadmaps with minimal pause. Insight has moved closer to the point of action, and with that shift comes much higher consequence.

This is where Humans in the Loop is supposed to matter. In theory, keeping humans in the loop ensures that interpretation, scepticism, and methodological discipline remain central, acting as a brake on automation when confidence outpaces understanding. In practice, however, "human in the loop" in market research often means little more than procedural oversight. A researcher reviews the output. A slide deck is signed off. A human is technically involved. But involvement is not the same as

The cost of being wrong rises as decision velocity increases. Speed demands higher standards of rigour, not lower ones.

control, and presence is not the same as judgement.

When AI systems generate plausible, polished narratives at speed, humans increasingly approve rather than interrogate.

At its mildest, this dynamic erodes the craft of research. The work of insight has always involved friction: wrestling with conflicting signals, resisting the temptation to over-conclude, pressure-testing assumptions, and holding uncertainty long enough to understand what really matters.

When systems optimise for speed and clarity, that productive friction disappears. Researchers challenge less rigorously because "the model agrees." They triangulate less because synthetic data looks consistent. Insight starts to sound right rather than be right. This isn't a failure of capability; it is deskilling by design. Judgement, like any muscle, weakens when it is no longer exercised.

Owning risk

The more serious risk emerges when market research influences decisions with legal, ethical, or regulatory implications. Pricing fairness, healthcare communications, financial product positioning, sustainability claims, employment strategy, these are areas where insight increasingly carries real-world consequences.

If AI-generated or synthetic-heavy research is biased, incomplete, or misapplied, the downstream impact is no longer academic. When decisions are challenged, nobody asks the synthetic respondent to explain itself. Nobody cross-examines the model. Accountability does not sit with the algorithm. It sits with the insight professionals who signed off the work.

This creates a structural accountability asymmetry that the market research industry has not yet fully confronted. AI systems increasingly shape conclusions and recommendations, while humans increasingly own the consequences. Humans in the Loop, when implemented superficially, becomes less a safeguard and more a liability trap. Oversight that does not include real understanding, the ability to challenge, and the authority to stop or slow decisions is not protection, it is exposure. If a researcher cannot clearly explain why an insight should be trusted, articulate its limits, and defend its use under scrutiny, they were never meaningfully in the loop.

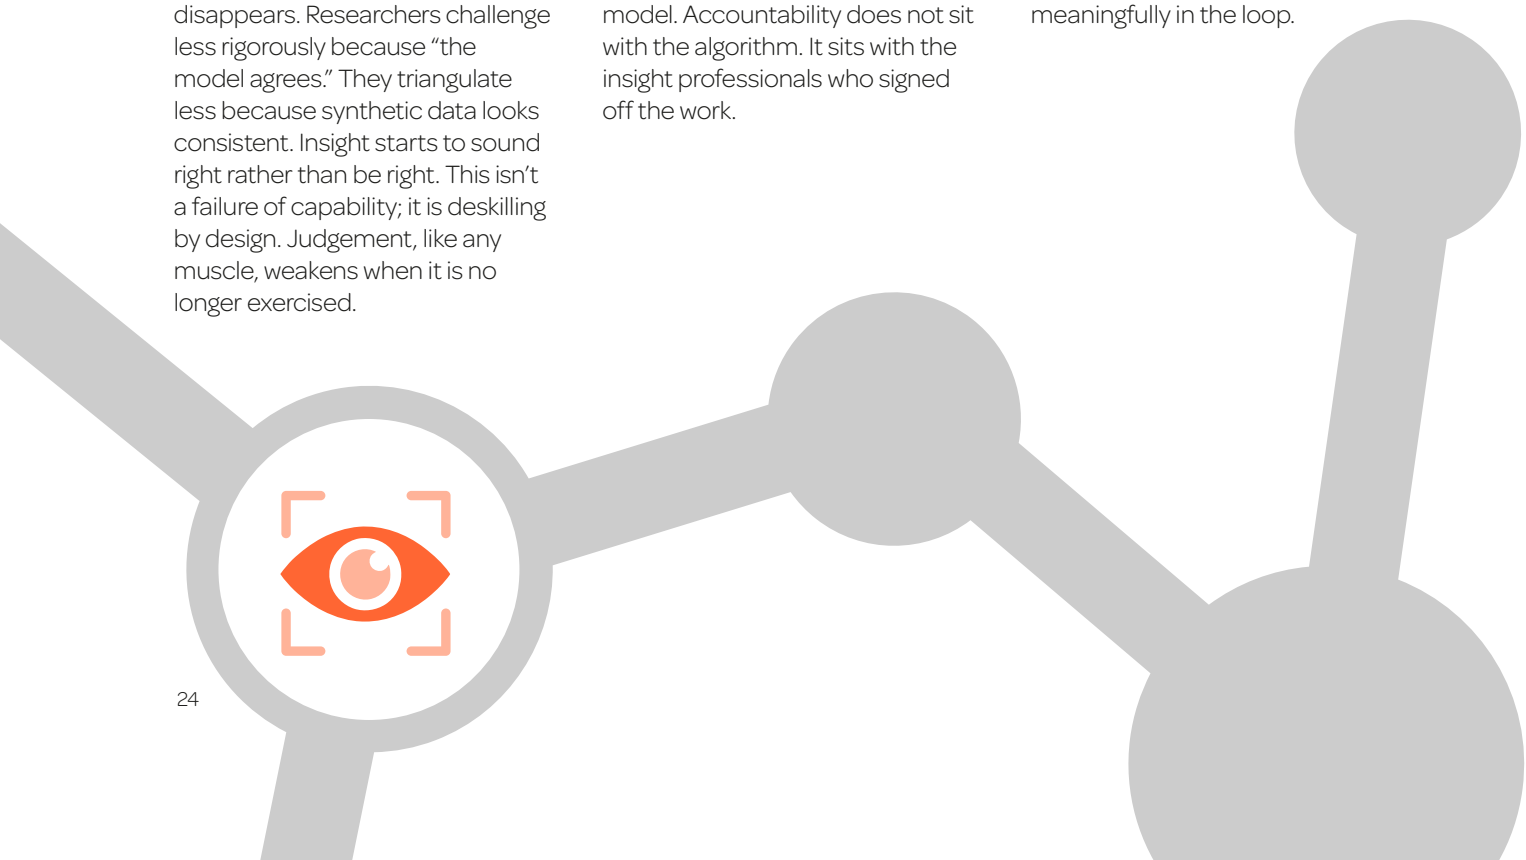

Speed needs scrutiny

Synthetic audiences sharpens this problem further. Used well, synthetic respondents are extraordinarily powerful for exploration, iteration, and early hypothesis testing. They are excellent at ranking, screening, and pattern discovery within known spaces. Like any other methods, used uncritically, they can create dangerous overconfidence.

Some synthetic systems tend to smooth away edge cases, cultural nuance, and emotional complexity; precisely the details that often create real-world failure later. The risk is not that synthetic insight is fake, but that it becomes decision-grade without the proper critique and testing.. When speed replaces scrutiny, 'directionally right' just accelerates uncertainty rather than reducing it.

This is where the industry's tolerance for "80% right" falls apart. We achieve 80% only through hard work, rigorous methodology and critical thinking, things that take time and craft.

Replacing this with a tool that is often wrong, means more wrong decisions, not less. In an AI-accelerated world, decisions move faster, propagate further, and compound risk more quickly. The cost of being wrong rises as decision velocity increases. Speed demands higher standards of rigour, not lower ones.

With power comes...

The truth is that market research is no longer just a support function. It increasingly sits on the critical path of decision-making. That means 'humans in the loop' must evolve from a comforting phrase into professional safety infrastructure. It must mean clear separation between exploratory insight and decision-grade evidence. It must mean mandatory human interpretation, not just approval. It must mean explicit articulation of uncertainty, bias risk, and limits - even when those slow things down.

These are not blockers to progress; they are what preserve credibility when insight is challenged.

Market research will probably never send anyone to prison. But the decisions it informs increasingly shape outcomes that really matter financially, ethically, and reputationally. When decisions go wrong, the AI won't answer for them. The insight professionals who let 'good enough' pass for understanding will.

Market research will probably never send anyone to prison. But the decisions it informs increasingly shape outcomes that really matter financially, ethically, and reputationally.

AI and organisational emotion

AI and organisational emotion
Efficiency or Anxiety?

Chrissie Tarbotton

Research Director

Boxclever

Shane Hanson

Director Insights & Innovation

Panasonic Consumer Electronics, Europe

Most of us are now familiar with AI entering the workplace through tools like Copilot, Grammarly and ChatGPT. For many, the benefits are clear. But alongside those benefits comes a level of uncertainty: about shifting role boundaries, changing ways of working, and what these technologies may mean in the longer term.

Despite this uneasy undercurrent, in many organisations, honest conversations haven't happened. There's been little direct discussion about what AI means for people, roles and the future of work. In that gap, uncertainty tends to fill the space, and for many employees, anxiety quietly follows.

The Anxiety of Ambiguity

Whilst it would be easy to dismiss this as simply a fear of technology, that would miss the point. What many organisations are experiencing is something more precise. The dominant emotion in AI adoption isn't fear of technology. It's the anxiety of ambiguity.

The pressure to adopt AI and not fall behind means that AI lands (often quickly) before anyone has meaningfully worked out what it means for the people doing the work. The technology moves faster than understanding, communication and clear direction- and in that gap, people naturally fill the space with their own inferences.

These fears are exacerbated by a fundamental speed mismatch. AI capability moves at software speed; organisational culture moves at human speed. The gap between the two is where much of the anxiety sits. Expecting one to keep pace with the other by default isn't a realistic strategy. Yet in many organisations, there is often a natural assumption that introducing a new tool will quickly lead to embedding it in day-to-day work, and that access will translate into adoption and then comfort- which rarely holds true in practice. The result is not resistance so much as a sustained state of uncertainty, where people are working with tools whose implications for their roles are still being understood.

The Dilution of Expertise

The market research industry is no stranger to a digital transition, but this wave feels particularly stark in its potential impact. Roles that were once considered highly skilled have become partially commoditised, and capabilities that took years to build are now available at the click of a template. When working in quant and analytics, technical skill- the ability to build and interpret statistical models, write analysis code, manage complex datasets- has long functioned as a marker of credibility and capability. AI compresses the value of that expertise quickly and visibly. Tasks that once required fluency in R, that required genuine

When what you know is central to how you are seen professionally, widespread access to that knowledge changes more than job descriptions. It destabilises the professional self.

statistical understanding, can now be initiated by someone with neither. Whilst this creates real opportunities for upskilling, particularly for those previously constrained by time or resources, it also needs to be understood as more than a practical shift. It is an identity shift. When what you know is central to how you are seen professionally, widespread access to that knowledge changes more than job descriptions. It destabilises the professional self.

Opportunity or Unsettlement

The question this raises of whether researchers are ultimately technicians or strategists is not a new one, but AI makes it urgent. It also makes it uncomfortable, because most roles in the industry have historically required both, and have been valued for both. Removing that balance forces a rapid shift in how many researchers understand their own value within the industry.

This naturally leads to questions around roles, and whether AI is likely to bring greater clarity to them or instead unsettle them further. It has the potential for both, but the reality depends on the level of meaningful integration. The optimal version of the story is that AI handles the mechanical work, humans handle the strategic work, and everyone spends more time on the high-value interpretive work they likely entered research to do. This is not the easiest path. Given the differences in pace outlined earlier, it would require sustained, deliberate effort to keep that balance intact as tools evolve and expectations shift, so in practice, the picture is considerably messier. AI's speed originally promised freedom for deeper thinking, but too often, it just creates more work. Compressed timelines mean output expands while understanding doesn't. And researchers who push back to protect quality, context, and evidential rigour risk being cast as blockers rather than guardians.

Ensuring Good Integration

What helps is less glamorous than most AI strategies suggest. It is not about tool access, training courses, or transformation roadmaps. It is about explicit, role-level conversation: what is AI being asked to do here, what does that change, what stays human, and who has the standing to shape those answers? This is a mutually beneficial goal as this is often something clients seek clarity on too, to be reassured that their investments are not being handed over to an AI. The organisations navigating this well have moved beyond the efficiency framing altogether. They've recognised that "doing more with less" is a promise that arrives with an implicit threat attached, and they've chosen a more honest account of what they're actually asking their people to become. Employees hear: you will be freed up. They suspect: your role will be reduced. In that gap, efficiency stops feeling like empowerment and starts feeling like exposure.

The Pace of Culture

Whether most organisations will get there is a different question. The tools are evolving faster than the cultures that hold them, and the emotional labour of navigating that gap continues to fall, quietly and unevenly, on the people doing the work. Organisations with more agile management styles are better placed to absorb the transition, giving teams the space to adapt as the landscape settles. Where culture is slower to shift, uncertainty lingers and clarity takes longer to follow. By and large, the pace and impact of AI will be shaped as much by organisational mindset as by the technology itself.

AI and commercial structure

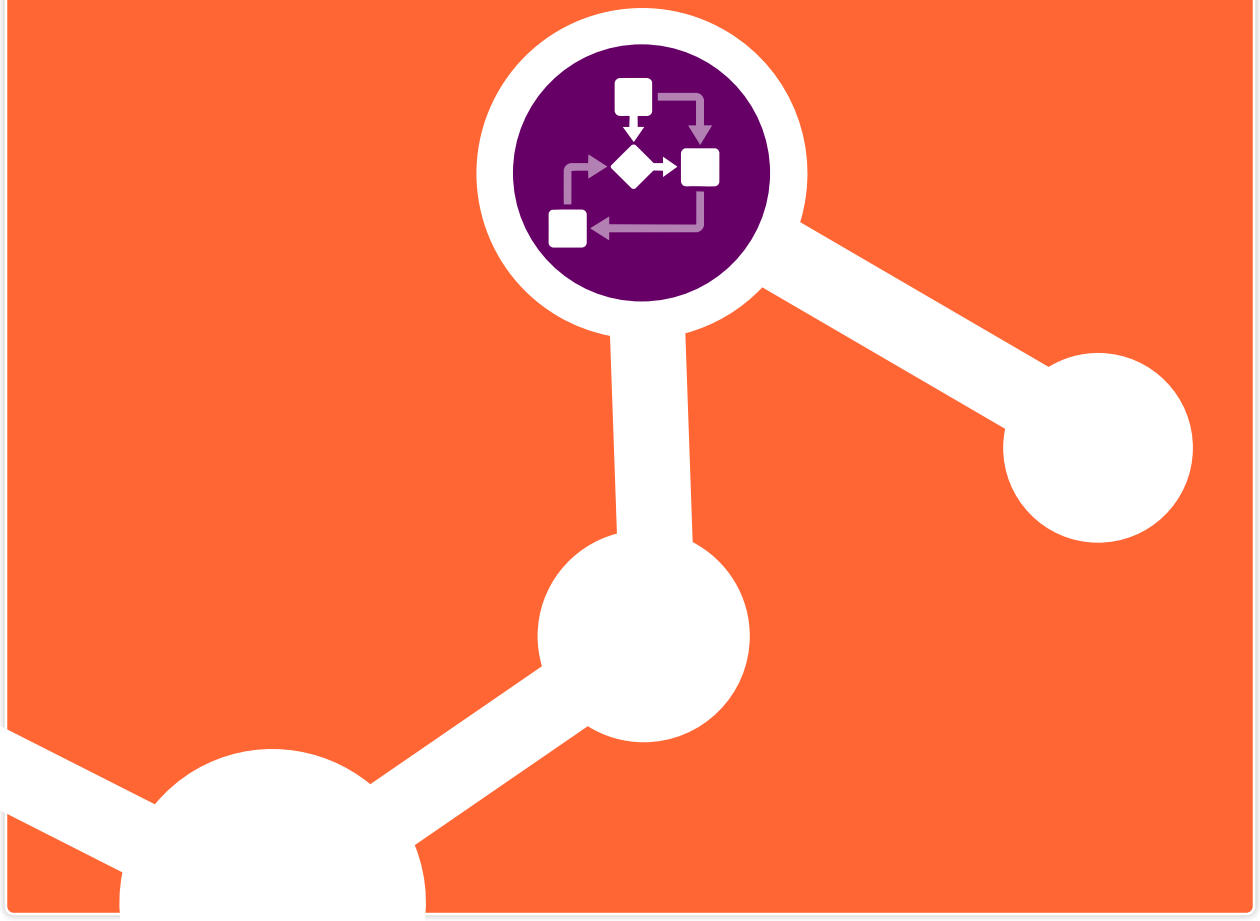

AI and commercial structure

The future won't wait for us to catch up

Alex Owens

Josephine Hansom

Youth researcher and keynote speaker

Much of the debate about AI in research focuses on technology. We believe the bigger question is commercial. AI is fundamentally reshaping the relationship between agencies and clients in three ways:

1. The contract between them is changing.
2. The roles performed are changing.
3. The value created by research is moving.

The agency-client contract is being rewritten

The AI story began with ELIZA, one of the world's first chatbots, developed in 1966. Despite being little more than a pattern-matching system, many users believed it genuinely understood them.

Expectations rose, but reality struggled to keep pace and the first AI winter followed. The same pattern repeated for decades.

Today's AI wave is different:

- The technology genuinely works.
- It's accessible to almost everyone.
- It's been democratised through tools such as ChatGPT and Claude.
- Organisations are seeing measurable value.

This matters because many of the activities traditionally purchased from agencies can now be performed more easily, quicker and at lower cost. Survey design. Coding. Reporting.

The current contract has become obsolete.

As information becomes more abundant, interpretation, challenge and trusted expertise become more valuable. The future agency-client contract will be built less around process and more around expertise, trust, and helping organisations navigate the unknown.

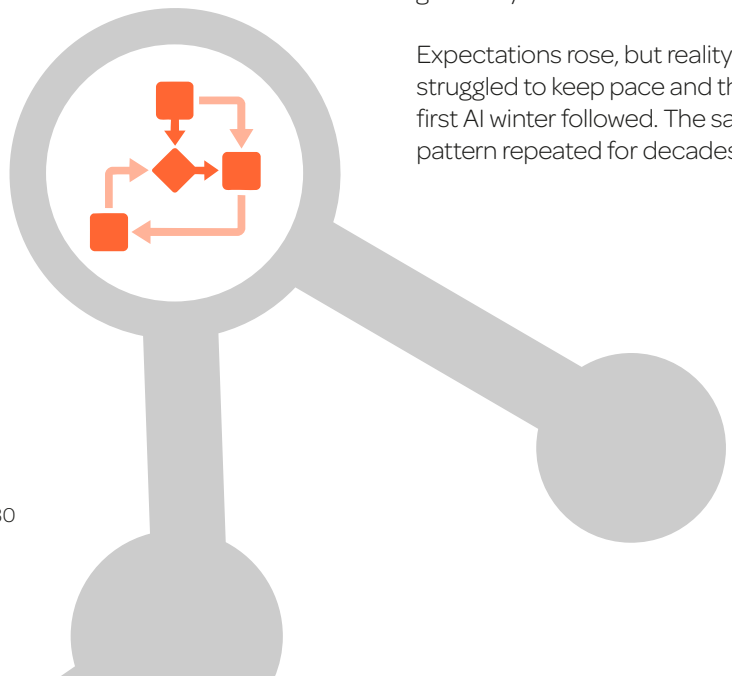

Technology is moving faster than organisations

The greatest challenge facing the industry may not be AI itself. It may be organisational adaptability as most are designed for efficiency, governance and stability. Few are designed for continuous reinvention.

Yet AI demands exactly that.

AI is not a technology transformation. It is an operating model transformation.

Operating models are built from three interconnected elements:

- People
- Processes
- Technology

During, ten years leading transformation programmes at Unilever one lesson became clear.

The role of transformation is not to react to the future. It is to prepare the organisation for it.

We were experimenting with AI technologies years before generative AI became a mainstream boardroom discussion. More importantly, we continuously evolved our operating model to stay ahead of changing business needs. Across that decade, we redesigned our operating model seven or eight times.

Not because it was broken. Because the future had changed.

Many organisations and agencies are not currently structured to evolve at that pace. That may prove to be a bigger risk than AI itself.

Research is not disappearing. Its value is moving

Much of the debate focuses on whether AI will replace researchers, shrink agencies or move more research client-side.

These are important questions, but they miss the commercial one.

AI will expand the market. More organisations will access research, and existing users will conduct it more frequently. More research will be done, but not all of it will require a specialist.

> The research industry built a model around process. Agencies sold it. Clients bought it. AI does process cheaply. Now both sides must answer a question they have never had to face directly: what are we charging for? **Alex Owens**

The future client-side researcher is likely to spend less time generating insight and more time translating evidence to drive data driven decisions and growth.

	What was sold/bought	What AI replaces	Where the value is
Operations	Recruitment, data collection and quality	Automated sampling, scripting, synthetic data	Reach real people and using engaging tools
Agency	Research methodology, analysis, reporting	Research design, automated analysis	Knowing how to ask and what it means in context
Client	Commissioning, delivery and dissemination	Project coordination, access to insights	Translating insight into decisions that land

Like web design, DIY tools increased participation while squeezing the middle of the market. Research is likely to follow a similar path.

Where does value sit in the New Contract?

Historically, much of the industry's revenue has been generated through process execution: survey design, sampling and recruitment, coding and analysis.

AI is rapidly reducing the effort required across these activities with significant implications.

For decades, value was attached to process execution. Increasingly, it will sit in interpretation, judgement, decision intelligence and the ability to drive organisational change and growth.

This shift has significant implications for roles. Some will disappear, others will evolve, and new ones will emerge.

The future operations supplier faces an immediate choice: commoditise or move up the value chain.

The future agency may spend less time executing projects and more time helping clients interpret complexity, and drive action.

The future client-side researcher is likely to spend less time generating insight and more time translating evidence to drive data driven decisions and growth.

Specialist knowledge was always the real product, but it was never priced as such.

For decades, agencies bundled what they knew into what they did. A specialist in young people didn't charge for understanding that audience, they charged for the fieldwork, the analysis, the report. The knowledge came with the process. It had no line item of its own because it never needed one.

"I've spent twenty years developing expertise in a specific audience. None of that ever appeared on an invoice. It was hidden inside the hours. AI has removed that cover."
Josephine Hansom

AI has changed that. Process is now cheap. What remains is what specialist agencies knew all along, and for the first time it has the space to stand alone.

In short, AI does not remove the need for expertise. It increases its importance. As process becomes cheaper and more accessible, expertise, judgement and decision intelligence become the product in their own right.

The future belongs to specialists, not factories. But expertise alone is not enough.

The winners will be those who combine expertise, decision intelligence and the ability to drive organisational change at pace.

Conclusion

The research industry has repeatedly demonstrated its ability to adapt. What makes AI different is not that change is happening. Change has always been a constant.

What is different is the pace at which technology is evolving and the scale of the opportunity it creates.

The traditional agency-client contract was built around process. Agencies sold it. Clients bought it. And for many years that model worked. Today, AI is exposing its limitations.

The real value was never the process itself. It was the expertise, judgement and ability to turn evidence into better decisions.

Process was simply the wrapper. AI has begun to remove it.

Previous AI winters occurred because the technology was not ready. This time it is.

The greater risk is that agencies, client-side teams and organisations fail to adapt quickly enough to capture its value.

The next AI winter, if one occurs, is unlikely to be a failure of technology. It may be a failure of relevance.

Both sides built the existing contract. Both sides will need to rewrite it. Deliberately. Together.

With AI embedded in the foundations rather than added as a layer on top.

The contract needs rewriting. The only question is who holds the pen.

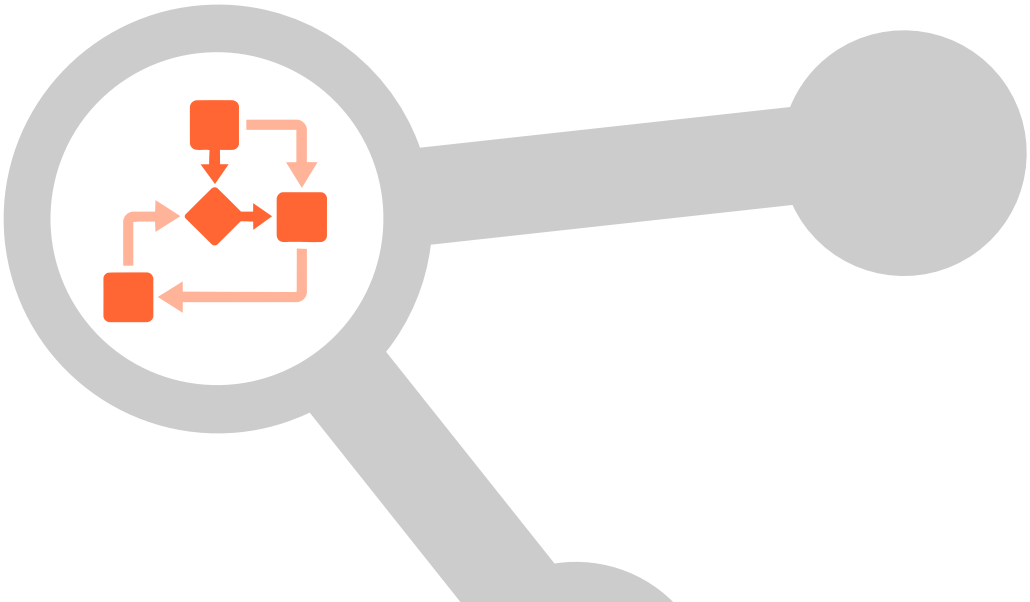

AI and hybrid practice

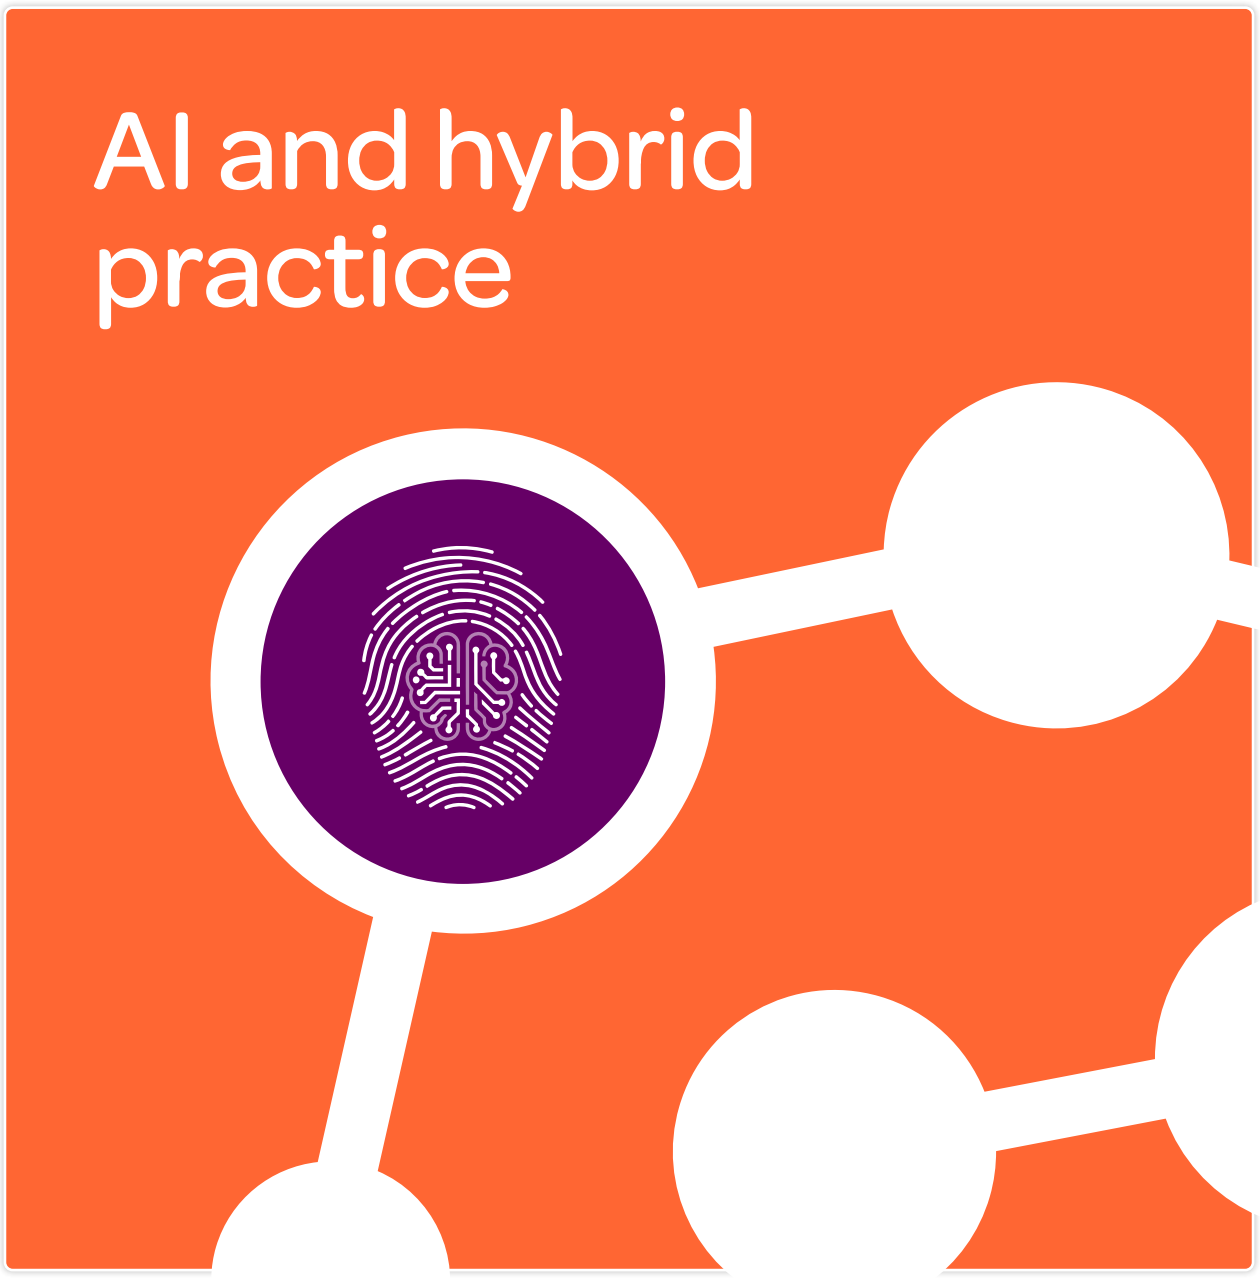

AI and hybrid practice
The human edge

Adrian Sanger

Co-Founder

Fore Consulting

Sarah De Caux

Lead Analytics & Insights Manager

Co-op

In our own daily work, whether on the agency side or managing brands, we talk constantly about combining AI with human intelligence. It's a great line for a pitch slide or a board deck. In this chapter we set out to explain what that actually means in practice.

Turns out that's not so easy. AI is so embedded in the research process and being without it now feels like cutting the WiFi. It's indispensable when it works and we grow more fluent and more dependent with each passing day.

On faster and cheaper, AI wins comfortably. Sets you thinking, if machines can process data faster and cheaper than any human team then what is the distinctive contribution of researchers?

In adopting more AI in more parts of our work and more often, have we outsourced our knowledge and experience for the sake of getting the job off the list. It's a question of accountability. Who actually owns understanding? Is it the system that processes the information or the person who decides what that information means?

Humans served by AI

We tend to paper over this question of accountability by wrapping it in the comfortable language of 'AI plus human intelligence'. It's become our standard industry positioning because it sounds balanced. It promises augmentation instead of replacement.

In practice, AI has already settled into parts of the research workflow, most of them operational and increasingly interpretive.

Seen this way, the distinction is more nuanced than it appears. Rather than a clean split between processing and interpretation, AI's role is increasingly to generate

Stage	AI contribution	Human contribution
Data processing	Cleaning transcripts, tagging themes, structuring datasets	Judging relevance, filtering noise
Pattern recognition	Detecting correlations, clustering responses, surfacing anomalies	Asking why patterns exist, applying meaning to them
Output generation	Drafting summaries, charts, first pass reports	Constructing narrative and implications

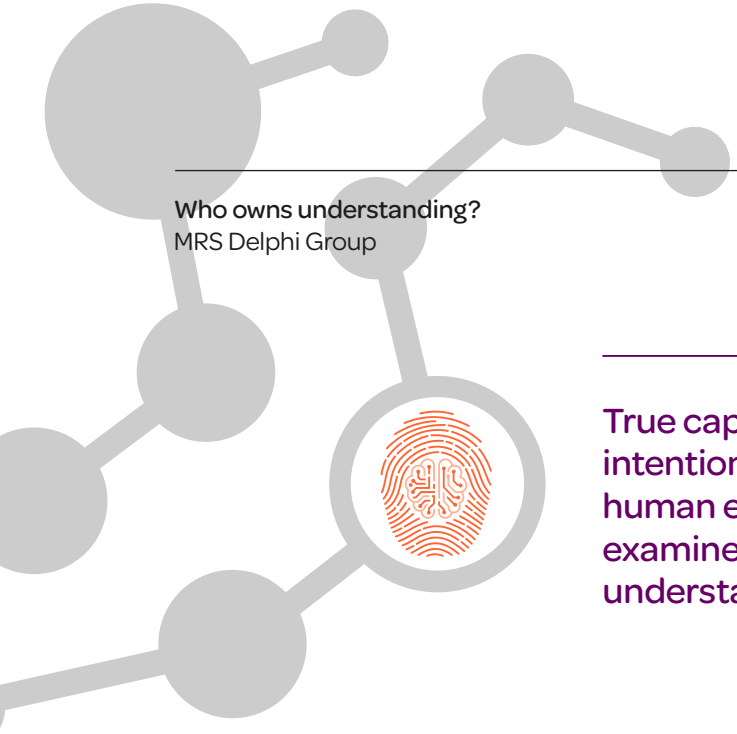

True capability evolution means we must intentionally design new ways to build human expertise. You cannot cross-examine an AI output if you do not deeply understand the data it is summarizing.

and organise possible meanings. It can cluster themes and draft implications at warp-speed. The uniquely human task is deciding which of those meanings are actually valid, useful and safe to act on

To our minds, the value shifts entirely from producing analysis to producing meaning. Anyone or anything can generate a pattern. The real human task is deciding whether that pattern actually matters and attributing meaning to it.

Consider a pack design test. The AI platform instantly aggregates eye-tracking data and flags a massive success: a bold new logo is drawing huge visual attention.

The machine registers a winning signal. An experienced researcher applies experience. That high attention score isn't an index of desire. Consumers are staring because they are confused. The human exercised the judgement to take other information into account and urge caution

Hybrid intelligence

We believe the future of consumer insight lies in hybrid practice rather than pure artificial intelligence.

Artificial intelligence replaces individual tasks, but hybrid intelligence entirely reshapes human roles. When you look at how this plays out on live projects, three patterns emerge

1. **The efficiency engine**
   The first role for AI is straightforward.

   Much of research contains structured, repeatable tasks that historically required substantial human labour. AI handles these quickly and reliably.

   This should reduce cost and shorten timelines. Efficiency alone does not produce insight. Still, it accelerates the pipeline through which insight may emerge.

   Automation improves the mechanics of research but alone it does not resolve interpretation.

2. **The collaborative colleague**
   A second model appears in more exploratory workflows where AI becomes less a tool and more a partner in analysis.

   Used well, this dynamic accelerates thinking and enables rapid exploration and progression of ideas. At the same time Gen AI notoriously struggles to maintain a consistent analytical lens across complex datasets. If a project involves comparing multiple consumer sub-segments, the machine frequently drifts, applying different evaluative criteria to different groups.

   This is why the ping-pong dynamic requires high experience. The researcher must constantly audit the machine's consistency, acting as the stabiliser across the data.

3. **The Evidence amplifier**
   The most interesting role for AI emerges when it expands the evidence base of insight work.

   Many modern datasets are too large or complex for traditional analysis alone. AI systems can detect patterns that would otherwise remain invisible.

   Use cases in this area include identifying emerging consumer trends from social media chatter, spotting anomalies across millions of transactions, uncovering hidden correlations in long-term customer engagement and mapping complex pathways

in multi-channel journeys that humans will not easily track.

But expanded evidence does not automatically produce understanding. Data does not explain itself. Insight emerges when patterns are interpreted and placed within context. More evidence is only an advantage if you have the experience to cross-examine it.

The Human edge

If machines dominate speed and scale, the human contribution becomes clearer:

1. Researchers act as **narrative anchors**, weaving together multiple sources of evidence and deciding which signals deserve attention.
2. They also function as **cultural translators**. Consumer behaviour sits within wider social and economic contexts. AI can detect statistical relationships. Humans explain their significance.
3. Finally, there is **ethical responsibility.** Research involves judgement about data sources, representation and interpretation. AI processes information. Responsibility for the conclusions drawn from that information still rests with people.

Large organisations are fragmented. Data sits across departments, platforms and programmes of work that rarely connect cleanly. Humans navigate these gaps.

Why human intelligence cannot be abdicated

There is also a structural reason human interpretation remains necessary. AI systems analyse what is available to them but don't recognise what is missing.

Large organisations are fragmented. Data sits across departments, platforms and programmes of work that rarely connect cleanly. Humans navigate these gaps.

Representation presents a further challenge. AI reflects the structure of its inputs. If certain groups are absent or poorly represented, the system may produce confident conclusions based on incomplete evidence. The machine cannot easily identify what it was never given the capacity to observe.

Worse, even with perfect inputs, the machine requires constant policing. True hybrid practice dispels the myth that we can trust the system to run flawlessly in the background. In reality, a significant portion of human time is spent checking that the machine has actually done its job properly. Blind trust in the machine is not acceptable in the business of evidence; a human driver must step-in to identify where it has misread nuance, smoothed over critical friction or simply cut corners

Conclusion

In a fast moving environment, detailed prescriptions age quickly. What appear more durable are several guiding principles.

Researchers must adopt a deliberate mindset about where AI adds value and where interpretation remains essential. Hybrid intelligence requires conscious design.

Data quality becomes even more critical. AI amplifies whatever it is fed, which means source transparency and methodological discipline matter more rather than less.

However, this pivot introduces a sharp craft problem. The entry-level work of the past was also the training ground where junior researchers developed their category intuition. If we automate the mechanics entirely, we are hollowing out the next generation's ability to interrogate the machine. True capability evolution means we must intentionally design new ways to build human expertise. You cannot cross-examine an AI output if you do not deeply understand the data it is summarizing.

The work starts now. Hybrid models are still forming and the organisations that learn fastest will shape the future of insight.

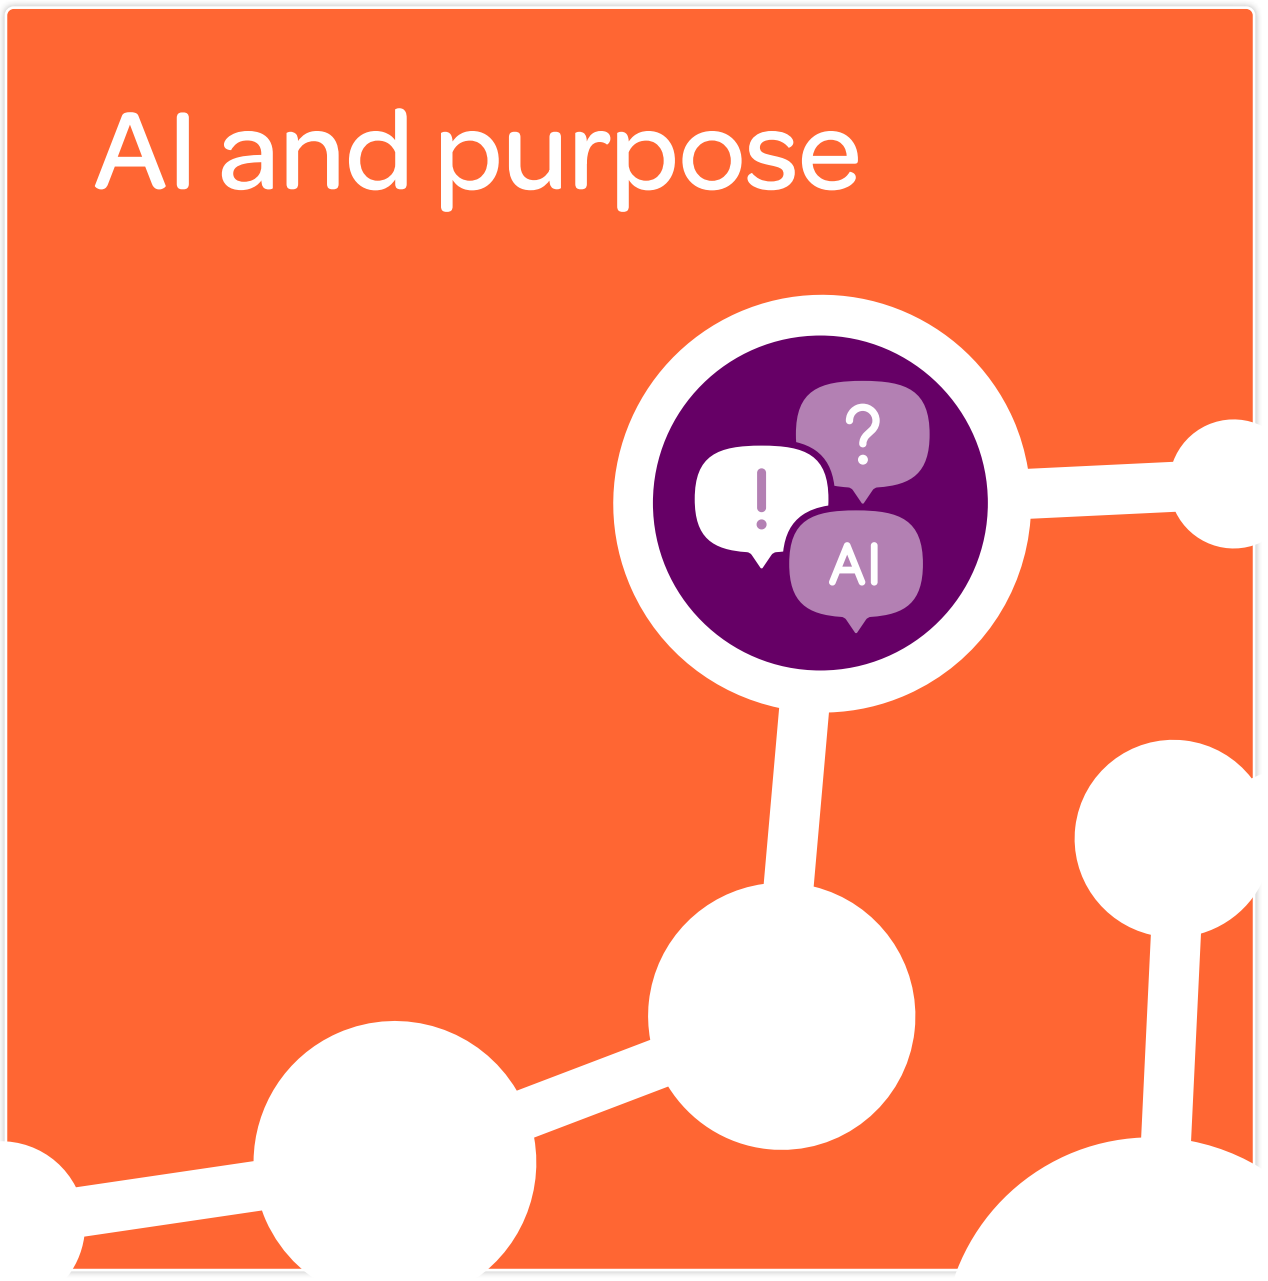
AI and purpose
?
!
AI

AI and purpose

Why do we need research when we have got AI?

Jane Frost
CEO
MRS

Marie Robelin
Growth Strategy & Transformation CMI Director
Unilever

About eight years ago non-researchers were predicting the end of the research function. I vividly remember this quote by the Chief Strategy Officer of a major London agency, he said to me, "We don't use research, we just use Google". Then he went quiet and after a moment reflected, "Maybe that's why our insight isn't very original".

Big data appeared to be the only thing people cared about. No one asked whether there was any virtue in the adjective "big", but people did think the more they had the better, the more they would understand their customer, the more problem solutions they would find. Here's the thing; big data is no use to anyone if it isn't good data, contemporary data, data that suits the investigation or the problem you want to solve, data above all that has been analysed and converted into usable answers to our problems.

The chasing after more and more data, without quality or purpose wasted a lot of time, consumed a lot of energy, and led us down blind allies that took investment that could be better applied elsewhere.

Think about it, the very famous saying from Avon; "in the factories we make cosmetics, to the public we sell hope" or Unilever's "Domestos isn't about bleach, it's about protection" or Persil's "it isn't about clothes, it's about Mother's care", which all imply we need to collect data that reflects emotions and psychological needs and not factory units. It was mostly the physical that got counted and turned into data.

So why is "artificial" intelligence a good thing? Why do we accept it as something everyone needs to sell, that we all need to have in our portfolio or we risk being dinosaurs?

Why is it easier to get the CFO to invest in anything with the acronym, AI, than trained professionals? Why do we have to defend the real against the synthetic? Why the tendency to assume artificial is better than human?

Maybe it is a very human tendency to be awed by technology, to be fascinated by process, because it is easier to measure outputs than outcomes. Think about it, procurement departments are

naturally more comfortable with units of ball bearings than with buying non-physical services.

The thing we should do is relegate AI to what it is, a process. A process of great potential in all of its formats when applied right. A process that can be good or poor quality. A process we need to run to deliver our objectives, rather than let it run us. A process we need to apply intentionally. A process we need to understand and professionally manage. We need to understand the quality of the data input, the way the LLM works needs to be transparent and the output needs to be verifiable.

Perhaps we need to avoid the acronym and call it what it is "Artificial Intelligence". AI somehow lets us get away with not considering that artificial anything has as many drawbacks as benefits. AI could be your best friend, so you don't think too carefully about its weaknesses.

AI can be marvellous – scratch that, good AI can be really, really useful as part of our armoury. It can be faster, go wider and deeper. It may be cheaper at times (although I have yet seen a total cost to serve model for quality AI) but sometimes it's no

AI can run the process, but the thinking must remain ours

Marie Robelin

Growth Strategy & Transformation CMI Director

Unilever

If AI is, as Jane argues above, a process to be applied intentionally, the real question is how it is now reshaping the craft of research and whether we are at risk of letting the process runs us.

AI is already rewriting how we work: generating briefs, surveys and discussion guides, moderating conversations, synthesising outputs, even producing synthetic respondents.

It can "go wider and deeper", as highlighted earlier. But its outputs are compelling not just because of what they say, but because of how they say it: fluent, confident, effortless. And that effortlessness is precisely the problem. AI taps into cognitive ease, our brain's preference for not having to think, amplified by authority and automation bias. In a discipline already tempted to prioritise outputs over outcomes, the risk is not just inefficiency, but intellectual passivity, researchers quietly "falling asleep at the wheel".

This is why research matters more, not less. As Jane reminds us, AI is only as strong as the questions we ask and the data we feed it. It cannot define the problem, nor observe the nuance of what people do and feel beyond words.

AI surfaces patterns, but often the most frequent ones and that is how you end up with safe, average, "beigified" insight. Real insight lives in tension, contradiction and minority voices.

...we need to increase our skills to embrace AI as we do other research practices, we need to understand our cost models and communicate our research frameworks with confidence.

cheaper than just seeing people, watching their expressions, seeing them interact. No senior leader I know thinks we can substitute that and the thing is that can be cheaper and just as fast as you need it to be.

AI can't speak better truth to power or give a voice to the marginalised as it stands because it is made up of our prejudices, our data weaknesses and importantly its output is defined by the way we frame our questions and our problems (sorry, "engineer the prompts").

So let's start with our purpose, be crystal clear about our outcomes and then apply the relevant methodology, whether that be an AI engine, use of synthetic data or an application of behavioural science, ethnographic study or focus group. Let's select our methodology based on our need and embrace what each of them is good for.

That means we need to increase our skills to embrace AI as we do other research practices, we need to understand our cost models and communicate our research frameworks with confidence. We definitely need to write better briefs so we understand our required outcomes more clearly and get colleagues to share the vision and we need to be in charge.

As the Harvard mathematician and satirist Tom Lehrer once said: "what you get out of it depends on what you put into it".

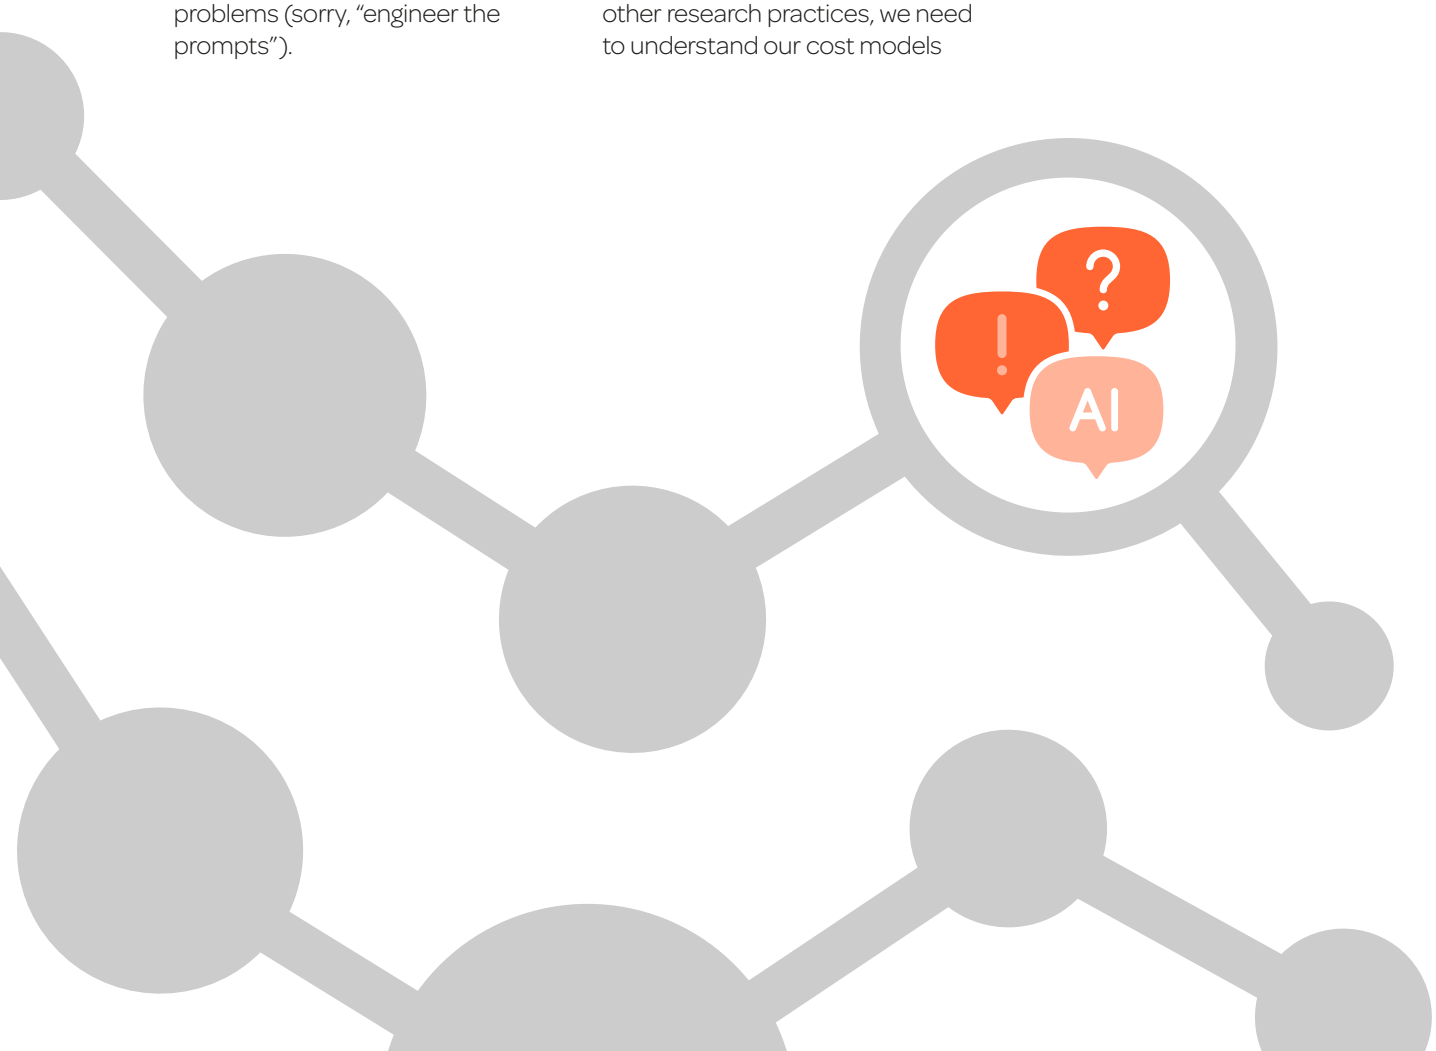

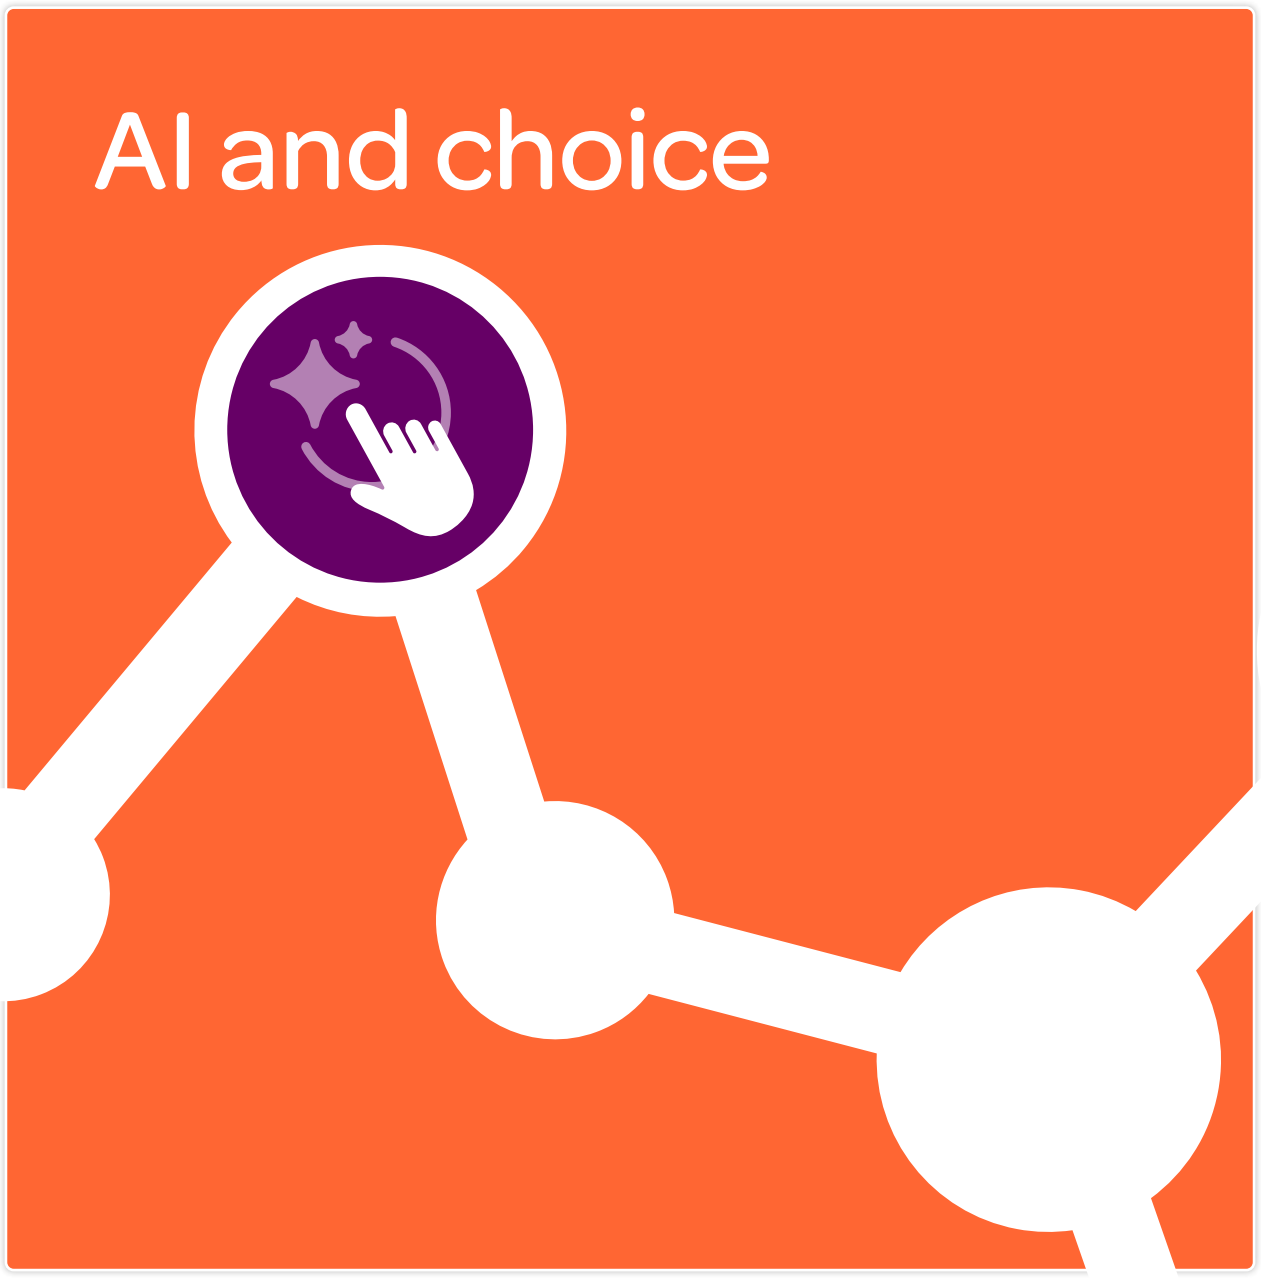
AI and choice

AI and choice
From AI capability to human choice

Dr Zsolt Kiss FMRS
Founder
o-x.ai

When MRS Delphi published the BEST (Boost, Expand, Shift, and Transform) framework in late 2023, it moved the conversation about generative AI beyond hype and provided practitioners with a practical framework for using it to support, extend or change research practice.

Much has changed since 2023. AI adoption is no longer limited to experimentation around isolated tools or use cases. It is becoming embedded in day-to-day research work: helping to design materials, structure analysis, summarise evidence, draft outputs and manage parts of the workflow. As adoption has widened, AI's function within research has also become more varied.

This suggests the framework would benefit from a second dimension: one that captures the role AI plays within the research workflow. BEST gives us a way to identify the kinds of situations where AI can support, extend or transform research practice. A role-based dimension would describe what AI is doing functionally or structurally within that work.

Based on what we are seeing in practice across the industry, including in our own work at o-x.ai, I believe there are four broad roles AI can play in the research process:

1. It can **deliver** parts of the research process, including qualitative coding, summarisation, analysis and first-draft reporting.
2. It can **generate** research inputs, including synthetic data, augmented samples and simulated stimuli.
3. It can **manage** workflows, from elements of project management to work that might previously have sat with junior or mid-level researchers.
4. It can **build** research tools, such as questionnaire scripters, coding aids, analysis utilities and dashboards.

These categories add a workflow layer to BEST. The original framework helps practitioners identify where AI may create value across different kinds of research problems. A role-based view brings the focus closer to the workflow itself: the function AI is performing, the kind of judgement it requires, and the form of oversight researchers need to provide.

This is less about what AI can do, and more about what we choose to keep human because it defines the quality and value of research.

The value is insight, not process

This distinction will become more important as the technology matures. Early debates understandably focused on what AI could and could not do, but capability is a moving target. Model performance is improving, and with the right context, prompting and self-diagnostic systems, hallucination risk can be significantly contained. Agentic systems are making multi-step workflows more realistic. Frameworks built mainly around yesterday's strengths and weaknesses will age quickly.

Much of the discussion has focused on what AI is suited to do: which tasks it can support, where it performs reliably, and where its limitations still matter. As capability improves, that question starts to lose some of its usefulness. The question the industry will soon need to answer is what we want to keep human in research, and what skills we need to develop to do that well. This is less about what AI can do, and more about what we choose to keep human because it defines the quality and value of research.

This brings the discussion back to value. Much of the industry's current use of AI is still focused on micro-productivity: accelerating individual tasks within a research operating model that otherwise remains largely unchanged. Faster transcript turnaround, quicker summaries, cleaner first drafts. These are useful gains, but they do not realise the full potential of AI in our sector.

The bigger opportunity is to rethink the operating model itself. That means asking what the work should become, not just how to do today's work faster. We have historically been paid for much of the visible labour of research: interviews conducted, transcripts coded, analysis completed, reports written. AI can now perform many of these activities quickly and competently, with the requisite supervision. This should not be seen as a threat to the sector. It is an opportunity to shift more of our time and expertise towards the work where research creates most value: providing insight and answering clients' questions, rather than executing process.

The value of research lies in framing the right question, understanding context, interpreting evidence, applying judgement and building trust with clients. If AI reduces the burden of process, it gives researchers more room to focus on those higher-value activities. It also creates an opportunity for the sector to occupy a more advisory role, closer to decision-making, but from a distinctive position, grounded in evidence, methodological rigour and expertise in understanding people.

Building judgement, not just capability

A future in which AI takes on more of the research workflow has direct implications for people, skills and roles. The sector will need to choose which parts of research should remain expert-driven: not because AI cannot assist with them, but because they define the quality, credibility and value of the work.

The issue is not simply keeping a human in the loop, which should be standard, but deciding where human expertise should lead: in framing the problem,

...the sector will need to be more deliberate about how people develop judgement.

making methodological choices, interpreting evidence, challenging weak conclusions and building the client trust that gives research its value.

This has consequences for the skills researchers need, and for how careers develop. Critical thinking, methodological confidence and the ability to challenge fluent, persuasive outputs are basic requirements. They will need to know when an AI-generated answer is plausible but thin, when an interpretation is over-extended, and when the evidence does not support the conclusion.

When it comes to progression, if AI absorbs more of the executional work through which junior researchers have traditionally learnt, the sector will need to be more deliberate about how people develop judgement.

The future researcher may spend less time producing every element of the process by hand, but they will need to be better at directing the work, questioning its outputs and knowing not to outsource the thinking to the machine.

The next phase of AI in research will require frameworks that are useful but provisional. The sector needs enough structure to act with confidence, and enough pragmatism to adapt as the technology, the work and the role of the researcher change.

A three-part series on generative AI

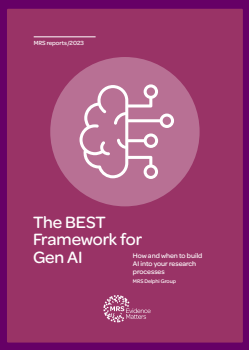

Download all Delphi reports at www.mrs.org.uk/delphi

Private lives?
A look at privacy issues through the lens of the consumer.

The politics of persuasion
What influences voters and how to improve democratic engagement.

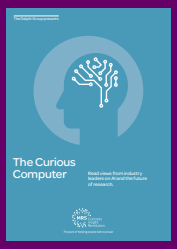

The curious computer
The impact of AI on the research and insight sector.

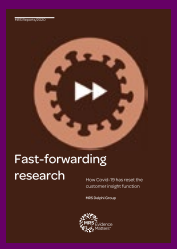

Fast-forwarding research
How Covid-19 has reset the customer insight function.

Rise of the insight alcher
How the data explosion has created a new breed of insights professional.

Towards an insight driven organisation
People, skills and processes.

Prediction and planning in an uncertain world
Helping organisations connect with the future.

Great Expectations
How technology impacts consumer trust.

Managing the multiverse
Understanding a world in transformation.

MRS Evidence Matters