Journal of Innovation

OMG Object Management Group
An EDM Association Community

Beyond AI Agents: Creating the Conditions for Breakthrough Intelligence

Authors:

Kyle Hatchard BEng (Hons) CMC PGCE DEP
STEM-TEC LTD, Director
kyle@stem-tec.com
https://stem-tec.inloop.studio/

Joshua Dickson MSc., ADAP (acc.), EMDR (acc.)
EMDR Consultant, Clinical Director Resurface Behavioural Health Ltd
josh@resurfaceuk.com

George Bissell BSc (Hons)
Inloop.studio
inloop.studio

CONTENTS

Beyond AI Agents: Creating the Conditions for Breakthrough Intelligence

Abstract:
How much of an organization's wisdom lives in quiet corners, unspoken and unwritten? Research suggests that a huge portion – perhaps 80% of critical knowledge – is tacit, existing only in employees' minds and habits (Sugarwork, 2025). Such knowledge is *embodied* in people's experiences and even their physical intuition. Traditional innovation methods struggle to tap this deep well. Meetings and workshops often trigger fear and performance games that keep real insights locked away. But a new approach is emerging. By combining human-centered design with **agentic AI** – AI that acts as an autonomous collaborator – organizations can create psychologically safe, **human-aware** discovery processes to unlock tacit wisdom. In this article, we explore the four key stages of innovation as practiced by Inloop.studio – **Discover, Design, Decide,** and **Validate** (*see Figure 1*) – and show how each phase can harness embodied cognition and AI's capabilities to drive breakthrough innovation. The tone is academically rigorous yet warm and human-centered, inviting engineers and innovators alike to explore how technology and psychology can work together to transform the way we solve problems.

1.1 INTRODUCTION

1.2 DISCOVER: WHERE IS OUR HIDDEN KNOWLEDGE?

What keeps organizations from accessing their own hard-won wisdom? A core issue is **tacit knowledge** – the know-how people carry in their bodies and unconscious minds. Organizations can't always access their own best wisdom because that wisdom resides not in documents, but in embodied, emotional, and intuitive experiences — the kind that Neuroscientist Antonio Damasio's somatic markers describe. Damasio famously described how **"somatic markers"** – gut feelings and bodily sensations – guide decision-making in ways we often can't articulate (Damasio, 1994). In fact, the most valuable institutional wisdom isn't in databases or reports; it lives in the muscle memory of veteran engineers, in the intuitive "feel" a manager has for what works and what doesn't. When someone says, *"something about this strategy doesn't feel right,"* they are accessing years of embodied experience that traditional analysis might overlook. The challenge is that **extracting these buried insights is a difficult task**. Conventional

brainstorming sessions and group workshops often *fail to unearth* tacit wisdom – not for lack of effort, but due to the limitations of human neurobiology.

Under the spotlight of a group meeting, people's brains sense **social threat**. If we fear looking incompetent or being judged, the limbic system engages in a defensive state. Just as tacit knowledge resides in the body and is guided by somatic markers, the amygdala's activation under social threat blocks access to that embodied intelligence. When people feel unsafe or judged, their nervous systems prioritize survival over reflection — so instead of tapping into deep insight, they perform. Psychologist Matthew Lieberman notes that social pain lights up the same neural circuits as physical pain. In a conference room full of peers and bosses, this threat response is almost inevitable. **Amygdala activation** prepares us for fight-or-flight, not for subtle insight. As a result, traditional discovery methods are *"neurologically optimized for social performance rather than authentic knowledge sharing,"* as one analysis put it. The loudest voices dominate not because they hold the best ideas, but because they're least sensitive to social risk. Meanwhile, genuine experts often remain quiet, especially in cultures where admitting uncertainty is perceived as unsafe. Harvard's **Project Aristotle (2016)**, which studied hundreds of teams at Google, found that **psychological safety** was the **number one predictor of team performance (Insights, 2025)**.

Yet paradoxically, it's often *most lacking* among senior leadership teams – precisely where new ideas and frank knowledge-sharing are most needed. Without psychological safety, curiosity is replaced by caution, and innovation stalls before it starts.

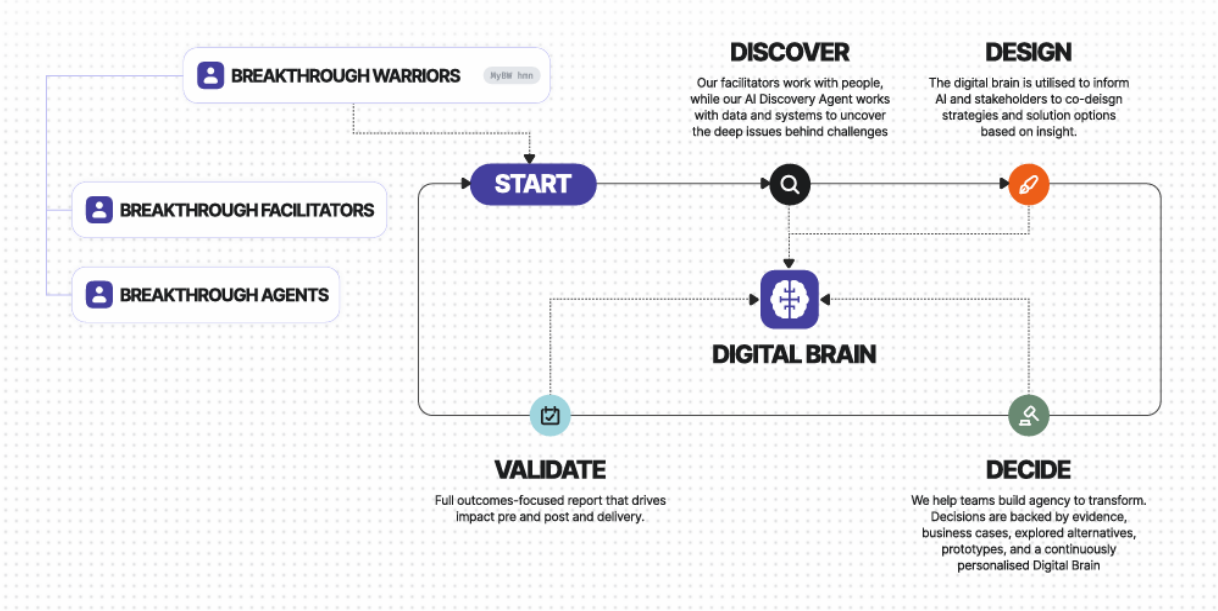

Figure 1 Inloop. Studio's process flow of Discover, Design, Decide, Validate, and Engineering Innovation Breakthrough in a week. Source Inloop.studio (2025)

2 Keystone Law PLC – A case example:

Context: Keystone Law, a distributed law firm with over 400 lawyers, sought to understand how AI could meaningfully contribute to its operations. Maurice Tunney, Director of IT, confessed, "I don't know what the problem is that I'm trying to solve." Fiona Servers, Director of Community, was equally unsure and sceptical of AI.

Embodied dynamics: The firm's culture mirrored the broader legal sector, characterized by a highly professional, regulated, and conservative approach to adopting new technology. This caution created uncertainty about where AI could safely add value without undermining trust or confidentiality. This uncertainty acts as a barrier to clarity in direction, increasing the "trough of disillusionment" entrapment described in Gartner's Hype Cycle (Fenn & Raskino, 2008), delaying the organization's movement toward the "slope of enlightenment" and eventual productivity gains

AI mediation: On the first day, **168 stakeholders were interviewed independently and simultaneously.** The agentic AI synthesized responses, surfacing recurring emotional language and subtle patterns of uncertainty that would have been lost in group discussion. AI provided a psychologically safe environment that allowed time for thoughtful reflections on inefficiencies, opportunities, and potential areas for automation. (relate to the article in some form - reference the references)

Embodied shift: The tone of conversation shifted from apprehension to agency. Fiona, once sceptical, became an active builder, prototyping an AI-assisted matching tool she called *Lawyers Logic*. Maurice gained clarity when collaborating with AI, clarifying each challenge step by step.

Outcome: The firm emerged with a clear roadmap toward an AI-first operating model, including initiatives in compliance, conflict checking, and intelligent client-lawyer matching. The process did more than generate ideas; it reframed AI from an external threat to an internal capability.

Agentic AI offers an alternative path to discovery that cleverly sidesteps our social wiring. Imagine each employee can interact one-on-one with an AI partner – a sort of intelligent sounding board – *instead of* voicing ideas in a crowded boardroom. This shifts the context from public performance to private reflection. Suddenly, the nervous system perceives safety: there's no one to judge "silly" questions or half-baked ideas. Individuals can take their time to think and even listen to their gut feelings without peer pressure. In fact, interacting privately with an AI can help people maintain a calmer physiological state – the parasympathetic "rest-and-digest" mode, linked to openness and insight. This defensive response is not merely psychological discomfort - it reflects measurable changes in nervous system function. A parasympathetically dominant state enables access to the very somatic markers and intuitive cues that underpin tacit knowledge. In simple terms, when people feel safe, their nervous systems literally make it easier to think deeply and sense what's true. Polyvagal theory (Porges, 2011) explains how social threat triggers either sympathetic activation (fight-or-flight) or dorsal vagal shutdown (freeze response), both of which inhibit access to higher-order cognitive functions. When teams operate in these defensive states, they cannot access the ventral vagal state characterized by social engagement, curiosity, and cognitive flexibility - the very conditions necessary for innovation.

Heart Rate Variability (HRV) provides objective evidence of these shifts. Higher HRV correlates with enhanced prefrontal cortex function, improved emotional regulation, and better creative problem-solving (Thayer et al., 2009). Traditional group workshops systematically reduce HRV through social evaluation pressure, whilst private AI-mediated reflection maintains the physiological conditions associated with optimal cognitive performance. This is not speculative - clinical research demonstrates that creating psychologically safe conditions produces measurable outcomes.

The implications for organizational knowledge extraction are profound. When Damasio (1994) described 'somatic markers', he was identifying how institutional wisdom is encoded not just cognitively but throughout the nervous system. An experienced professional's 'gut feeling' represents years of pattern recognition stored as embodied knowledge - accessible only when their nervous system perceives safety. AI-mediated discovery creates this safety not through reassurance but through structural design: removing the neurobiological triggers that activate threat responses.

Heart rate variability (HRV), a measure of nervous system relaxation, tends to be higher when we feel safe and unhurried. Higher HRV is associated with better emotional regulation and

creative problem-solving. By conducting discovery through AI dialogs on personal devices (for example, over the course of a day or two during the "Discover" phase), organizations enable team members to share knowledge while maintaining an optimal mental state. There's no rush to blurt out an answer before the meeting moves on; instead, each person can engage when they feel ready and **"centered,**" leading to more thoughtful input.

Early evidence for this AI-mediated discovery is compelling. **Harvard Business School researchers** found that when Procter & Gamble employees used AI assistance to generate ideas, they produced significantly more high-quality, novel solutions (HBS, 2025). The AI acted as a *"cybernetic teammate,"* giving individuals much of the benefit of a skilled collaborator (HBS, 2025) Participants who brainstormed with a GPT-4 based assistant were **three times more likely** to have an idea rated in the top 10% of all ideas, compared to people working without AI (HBS, 2025) In other words, the technology helped surface *breakthrough insights* that might otherwise have stayed buried. Interestingly, employees using the AI also reported feeling **more excitement and less anxiety** during ideation, relative to those working solo without AI (HBS,2025). This suggests that a well-designed AI partner doesn't just produce more ideas – it **changes the emotional climate** of discovery for the better. It creates a **psychologically safe "container"** where people share more honestly. As a result, quiet experts who might never speak up in a meeting often contribute rich insights via AI. Their knowledge gets captured without the usual social filters. One internal study observed that these private AI sessions *"invert"* the usual dynamic – authentic expertise rises to the top, while bluster and performative confidence fall by the wayside. In essence, **the Discover phase becomes a human+AI dialogue**, mining the tacit wisdom of individuals in a nonthreatening way. The organization gains a much fuller picture of the problem space – including those subtle warnings and creative hunches that people *feel* but seldom voice. It lays a foundation of insight into which truly novel solutions can be built.

2.1 DESIGN: CAN AI HELP US THINK OUTSIDE THE BOX?

Once we've gathered raw insights, the next challenge is turning them into innovative solutions. **Design** is the phase of idea generation, prototyping, and creative iteration. It's where human creativity must flourish. But creativity, too, has an embodied, non-linear nature. We've all experienced that *spark* when a problem suddenly "clicks," often after a period of frustration or wandering thought. Psychologist Mihály **Csikszentmihályi (1990)** famously described the state of **flow** – a feeling of energized focus and losing track of time – as the engine of peak creative performance. Flow occurs when we strike an optimal balance: **clear goals, immediate feedback, and a challenge level that stretches but doesn't overwhelm us**. Traditional design meetings rarely hit this sweet spot. Goals can be ambiguous, feedback is slow or muddled by politics, and a group of diverse people can't all be challenged at their personal optimum level. The result is often brainstorming that feels like slogging through mud – or design sessions that gravitate to safe, familiar ideas.

Agentic AI has the potential to **turbocharge the Design phase** by creating conditions that are much closer to a *creative flow*. How? As a participant in the process, an AI can provide instant

feedback and adapt to everyone. For example, when a team member proposes a concept, an AI design assistant might quickly sketch a mockup or simulate the idea, providing a concrete visualization within minutes. This is immediate, judgment-free feedback that something *works*, partially works, or fails – and why. The AI can also prompt with questions that nudge the person slightly beyond their current thinking ("Have you considered an approach that uses X?"), effectively adjusting the **challenge level** on the fly. In a human-only workshop, a facilitator can't possibly tailor the discussion in real time to each person's skill level, but an AI can. Teams at companies like General Motors have already used generative AI tools to explore *hundreds of design permutations* and identify improvements that a human might miss. In one case, GM and Autodesk's generative design system developed a seat bracket that was 40% lighter yet 20% stronger than the original – a design that no engineer had conceived before (Quinnox, 2025). This demonstrates how AI can transcend conventional boundaries, suggesting truly novel solutions while humans provide the goals and constraints.

Perhaps most importantly, AI helps **sustain the creative momentum** that humans often lose. In a typical design sprint, there's a day when uncertainty peaks – the concept sketches look wobbly, confidence dips. It turns out this *Day 3 uncertainty* is not a flaw but a **feature** of breakthrough innovation. Neurologically, our brains need to venture into the unknown and feel a bit lost before reorganizing ideas into a new pattern (Dietrich, 2004; Siegel, 1999). It's during this uncomfortable phase that many teams lose their nerve and fall back to safe ideas. This Day 3 discomfort is not incidental - it reflects the neurological architecture of a genuine breakthrough. Cognitive neuroscience reveals that authentic insight requires temporary destabilization of existing mental models (Dietrich, 2004). The brain must venture into uncertainty before reorganizing information into novel patterns. This parallels what trauma therapy refers to as the 'window of tolerance' (Siegel, 1999)—the zone of arousal where learning and integration occur. Too little challenge leads to disengagement, while too many challenges trigger a defensive shutdown. Effective innovation processes must maintain this optimal zone.

Traditional design workshops rarely achieve this balance because they cannot adapt to individual nervous system states. One person's energizing challenge is another's overwhelming threat. However, AI-mediated design can adjust dynamically - sensing when to push exploration and when to consolidate, when to offer scaffolding and when to step back. This responsive adaptation maintains participants within their individual windows of tolerance, enabling sustained creative engagement.

The 250% problem-solving enhancement observed in flow states (Chi & Snyder, 2012) occurs precisely because flow involves transient hypo-frontality - temporary reduction in prefrontal cortex activity that paradoxically enhances pattern recognition (Dietrich, 2004). The inner critic quiets, defensive monitoring reduces, and associative thinking accelerates. An AI partnership can facilitate these conditions by handling evaluative functions externally, allowing human cognition

to remain in generative mode. This is the neurobiological basis for why participants describe AI-assisted design as 'mind-blowing' - they are accessing cognitive states rarely achieved in conventional business contexts.

But a supportive AI partner can encourage *perseverance through the dip*. By reliably generating *workable ideas or insights at each iteration, AI assures human designers that progress is being made*. Team members stay engaged rather than hitting a wall. Studies on human-AI co-creation at P&G have shown that teams using AI achieve **higher "boundary-breaking" creativity**, mixing technical and commercial ideas more freely than teams without AI (HBS, 2025). They also developed concepts more quickly – the AI-assisted teams reduced concept development time by ~13%, and even individuals using AI worked 16% faster on ideas (HBS, 2025). Speed alone isn't everything, but it matters: rapid prototyping allows more cycles of learning within the same timeframe.

Crucially, **quality** didn't suffer – it improved. When AI was involved, even employees with less product design experience contributed ideas on par with those of seasoned experts (HBS, 2025). In other words, AI support helped democratize design skills across a broader group. This is an emotional win as well as a technical one: people often feel energized and empowered when they can make meaningful contributions to creative work. Indeed, in that P&G experiment, participants working with AI reported higher enthusiasm and lower frustration (HBS, 2025). It's as if having an AI "co-pilot" frees them to imagine boldly, because the AI handles some of the grunt work and provides a safety net.

3 Digital Isle of Man – From Stuck Design to Defined Product

Context: Digital Isle of Man, a government agency responsible for developing products that strengthen the island's digital economy, had spent over a year and a half working on the *Data Assets Foundation*. Despite significant effort, the team struggled to define and finalize the product requirements.

Embodied dynamic: The extended design period had created fatigue and uncertainty. The team's expertise and ideas were dispersed across meetings and documents, yet consensus on the product's core purpose and specifications remained constant. Stakeholders were spread across jurisdictions, and centralised coordination was a challenge. Extended innovation cycles often lead to cognitive fatigue and fragmentation of insight, particularly in distributed teams where knowledge is scattered across asynchronous documents and meetings. Research shows that AI-mediated discovery can consolidate dispersed inputs and re-align stakeholders on core purpose, even in complex, cross-jurisdictional environments (Mollick & Kifayat, 2023; Edmondson, 1999).

AI mediation: Over the course of one week, the AI conducted interviews with all parties involved in the project, synthesizing their perspectives into a single, coherent structure. Acting as an autonomous collaborator, it consolidated areas of agreement, clarified overlaps, and generated

a complete Software Requirements Document (SRD) that the team had been unable to produce after months of iteration.

Embodied shift: The experience replaced frustration with relief and renewed momentum. Seeing their collective input transformed into a concrete, organized specification gave the team a sense of progress and shared ownership.

Outcome: In five days, the project moved from stagnation to clarity. The AI delivered a well-defined product specification aligned with stakeholder intent, allowing the agency to proceed confidently into development. This demonstrated how AI can serve as a design catalyst, translating accumulated expertise into actionable direction.

4 Emirates New Development Bank – Designing a Future Built on Trust

Context: Emirates NDB's leadership team worked with Inloop in collaboration with INSEAD over three days to explore how the bank wanted to be seen in the public eye in 2035.

Embodied dynamic: The team had been engaged in strategic discussions about innovation but needed a focused space to define what "trust" would mean for the next generation of banking.

AI mediation: With Human+AI collaboration at the forefront, each leader reflected on what makes a brand trusted. The AI summarized these insights into shared principles and facilitated a structured exercise to deconstruct how existing banks build or lose trust. The group then used these principles to co-design what Emirates NBD will look like in 2035.

Embodied shift: The AI sessions helped create psychological safety, enabling open discussion and rapid alignment around trust as a design foundation.

Outcome: By the end of the three days, the leadership team had a defined vision for *NBD 2035*, a concept grounded in trusted values and supported by a clear set of brand design principles.

One participant described the AI as a **"collaborative thought partner, rather than a search engine"**, and that mindset shift was key to deeper creativity *(HBS, 2025)*. This underscores an important point: **we must reconceptualize AI as a teammate, not just a tool** (HBS, 2025). In design mode, which means engaging with the AI almost like you would a colleague – bouncing ideas off it, asking it to critique a sketch or stress-test an assumption. The AI might generate 10 variations of a concept, whereas a human would usually stop at 2, thereby expanding the design space to explore. It can simulate edge cases or run virtual stress tests on a design, as AI prototypes now do in industries from software to aerospace (Quinnox, 2025). All of this provides human designers with rich feedback to consider when forming their intuition. We still rely on human judgment to decide what "feels right" or aligns with user needs – those somatic markers again – but AI ensures we're not limited by what we can manually imagine or analyze.

Finally, embodied cognition reminds us that **design is a physical act as much as a mental one**. Great ideas often emerge from building and trying tangible prototypes. Generative AI accelerates

this *"think by doing"* loop. With modern AI tools, an engineer can generate 3D prints of a dozen component variations overnight to see which one *feels* most robust in hand. A UX designer can use AI to instantly create interactive app mockups and then physically observe users interacting with them the next day. By shortening the gap between idea and tangible experiment, AI allows teams to **listen to their embodied reactions** – that subtle sense that "this version is awkward, but that one flows nicely." In short, the Design phase becomes a dance between human intuition and AI augmentation. When done right, the outcome is a *flow state* for the team: clear goals, continuous feedback, and just the right amount of challenge to keep everyone fully engaged. Research shows that achieving flow can **boost problem-solving performance by over 200%**. This aligns with anecdotal reports of "mind-blowing" design sessions where teams accomplish in hours what used to take weeks. In these moments, the boundary between the human and the technology teammate fades – it just feels like a very creative team at work. And that's exactly the point. **Human-AI co-design** can unlock levels of innovation that neither could alone, by blending our embodied creative genius with the AI's tireless idea generation.

4.1 DECIDE: SHOULD WE TRUST THE GUT OR THE DATA?

After discovery and design, teams often face a tricky question: **Which idea or strategy do we move forward with?** In the Decide phase, the goal is to evaluate options, make informed decisions, and allocate resources effectively. This is where analysis traditionally weighs heavily – market research, cost-benefit calculations, risk assessments. Yet here too, the *embodied side of knowledge* plays a role. Seasoned leaders often rely on intuition or "gut sense" when making big calls, especially under uncertainty. Rather than dismissing this as irrational, science suggests those instincts are real intelligence at work: people with high **interoceptive awareness** (sensitivity to internal bodily signals) tend to make *better decisions under uncertainty*, presumably because they can subconsciously integrate complex cues (Mehling *et al.*, 2012). Understanding interoception requires recognizing that bodily sensations are not mere 'feelings' separate from cognition; they represent rapid, unconscious computations that integrate complex environmental cues (Craig, 2002). Research using the Multidimensional Assessment of Interoceptive Awareness demonstrates that individuals with higher interoceptive awareness make demonstrably better decisions under uncertainty (Mehling et al., 2012). A leader's 'gut sense' about a strategy reflects pattern matching across thousands of previous experiences, processed below conscious awareness, and signaled through somatic channels.

However, corporate environments systematically train professionals to override these signals. The phrase 'I need to sit with this' acknowledges that genuine evaluation requires embodied processing time—yet most business decisions are forced into immediate timeframes. This is where AI-mediated decision support offers unique value. By handling rapid analytical processing externally, AI creates temporal space for human embodied wisdom to emerge. Leaders can

engage with data whilst simultaneously attending to their interoceptive responses, asking 'Does this feel right?' without pressure to articulate why immediately.

This integration of analysis and embodiment parallels therapeutic principles. Just as effective trauma processing requires titration—moving between cognitive understanding and somatic awareness—effective decision-making requires oscillation between data analysis and intuitive sensing. AI partnership enables this rhythm: analyze, reflect, check somatic response, and integrate. Research in human-AI decision collaboration shows this integration produces superior outcomes precisely because it honors both computational and embodied intelligence (Lu & Zhang, 2025). Technology extends our analytical reach, whilst the somatic wisdom provides contextual validation that algorithms cannot access. In other words, a calm stomach or a racing heart can reflect computations our conscious mind hasn't yet sorted out. However, gut feeling *alone* isn't enough – it can be biased by emotion or limited experience. Conversely, a purely analytical, data-driven approach can overlook context or human factors that are not captured in the spreadsheet. The best decisions in innovation combine **analysis and intuition**, head and heart.

Agentic AI is uniquely positioned to help achieve that blend by serving as a kind of **rationality enhancer** that still respects human values. An AI system can sift through vast data, run simulations, and present evidence for each option far faster than any person. It can, for example, forecast outcomes of a new product launch across multiple scenarios or recall how similar past projects fared. This augments our *explicit knowledge*. But importantly, a well-designed AI decision aide will also highlight **uncertainties and trade-offs**, inviting human judgment where it matters. Research on human–AI collaboration in decision-making finds that **people and AI have different strengths**: humans excel at understanding novel situations and ethical nuances, while AI handles complexity and pattern recognition in "big data" domains (Carey School of Business, 2025). When we let each do what it does best, the result can outperform either one alone (Carey School of Business, 2025). For instance, a recent study by Lu and Zhang (2025) showed that in loan approvals, a combination of human officers plus AI predictions led to fewer defaults than either the officers or the algorithm working solo – and it also **mitigated bias** that the AI alone had shown (Carey School of Business, 2025). Humans identified context and fairness issues that the algorithm overlooked, while the AI detected statistical patterns that humans had missed, and together they made the best decisions (Carey School of Business, 2025). The lesson for any innovation team is clear: **use Agentic AI as a collaborative decision consultant, not as a dictator or a crutch**.

Psychologically, getting this collaboration right requires overcoming two opposite pitfalls. One is **over-trust** – blindly following whatever recommendation the AI spits out, which can lead to disaster if the AI's data was incomplete or its objective mis-specified. The other pitfall is **distrust or neglect** – ignoring a valid AI warning because the team is confident in its own opinion. Both tendencies exist. In fact, researchers have found that *experts* are often more resistant to AI advice

because they trust their own experience more (sometimes to a fault), whereas novices might over-rely on AI even when it's wrong (Carey School of Business, 2025). Cultivating the optimal middle ground requires transparency and training. **Explainable AI** is key: when people understand *why* the AI prefers Option A over B, they can better judge whether that reasoning aligns with factors their intuition values. In the loan study, simply providing explanations for the AI's decisions prompted human deciders to think more critically and improved outcomes further (Carey School of Business, 2025). The same principle can be applied in product or strategy decisions. If an AI analysis favors launching a particular feature because "forecasted customer lifetime value is 15% higher," a human might weigh that against a gut sense that the other feature, while less profitable on paper, aligns better with the company's vision (something an algorithm might not quantify). The AI's hard data, combined with human soft insight, leads to a well-rounded decision.

An **agentic AI** can do even more than sit in the background crunching numbers – it can actively guide the decision process to be more balanced. Think of it as a neutral facilitator that keeps the discussion grounded in facts and goals. It might remind the team of their own stated success criteria ("Remember, our priority was to maximize user adoption, not short-term revenue"). It could run real-time "what-if" scenarios during the meeting as new ideas arise ("If we prioritize speed to market and cut testing time by 50%, what's the risk? Here's what the model says…"). This frees the humans to inject the *why* – the strategic intent, the ethical considerations, the brand implications – without getting lost in calculating every scenario by hand. The AI essentially extends our *working memory and analytical reach*, allowing the collective team mind to explore more options simultaneously. Notably, studies of teams using AI have found that they tend to consider a **broader range of factors** and **cross-silo ideas** more frequently (HBS, 2025). By presenting information impartially, AI can reduce the sway of the highest-paid person's opinion or the confirmation bias that sometimes steers decisions. It's harder to ignore inconvenient evidence when an AI puts it plainly on the screen. In this sense, an AI assistant can help enforce intellectual honesty and keep the group's embodied threat responses (such as ego and defensiveness) at bay.

Still, the final call must resonate at a human level. Deciding to take one innovative concept forward and shelve others can be an emotional moment for a team that poured themselves into those ideas. A human leader's intuition may detect subtle signals – such as excitement, concern, or fatigue – that influence when and how to make decisions. For example, a leader might feel that the team, although leaning toward a particular solution, seems hesitant at its core. Maybe it's worth probing why, or even sleeping on the decision, rather than forcing a quick vote. A purely logical process wouldn't account for this interpersonal sense, but a psychologically savvy one would. The **embodied knowledge** of the team (their comfort level, energy, stress) is data, too. Good decision-making processes incorporate check-ins or "listen to your gut" moments explicitly. Some progressive teams even monitor stress or engagement via wearables and use that as feedback – if collective heart rates are spiking, maybe the discussion has hit a nerve that needs

addressing. In any case, **agentic AI combined with human intuition can yield decisions that are not only analytically sound but also feel right to those who must** implement them. And when people feel ownership and confidence in a decision, they execute it far more effectively.

One more benefit: by involving AI in decisions, organizations can later **audit and learn** from those choices with greater clarity. The AI can log the rationale and data behind a decision, creating a trail that can be reviewed if outcomes go awry. This reduces hindsight bias ("We should have known…") and turns each decision into a learning opportunity. Over time, the AI might even detect patterns – for example, noticing that when teams ignore their initial gut instinct in favor of a shaky data model, the project often fails – and remind us of those lessons next time. In short, the Decide phase becomes a true **human-in-the-loop process**. The AI does the heavy analytic lifting and offers recommendations; humans apply wisdom, ethical judgment, and gut validation. By iterating in dialogue, they converge on a decision that is both *intellectually robust* and *emotionally acceptable*. That's critical for innovation, because major bets usually entail uncertainty and risk – you want your team fully committed to the path chosen, not secretly doubtful or disengaged. The human touch ensures alignment with values and vision, while the AI injects a discipline of evidence. Together, they make decisions that are at once bold and well-grounded.

4.2 VALIDATE: HOW CAN WE FAIL SAFE AND LEARN FAST?

No innovation is complete until it has been thoroughly tested. The Validate phase is all about experimentation, feedback, and learning – essentially asking, *"Did our idea work in the real world, and if not, what do we do next?"* This stage closes the loop of innovation, and it's where theory meets practice. It's also a stage where psychological safety and embodied knowledge remain paramount. Validation often means *failure* on the path to success: prototypes will break, users will reject features, and initial hypotheses will prove wrong. If the team fears blame or embarrassment from these outcomes, they may distort or hide the results, or worse, avoid truly testing the riskiest assumptions. A **psychologically safe validation environment** treats failures as learning opportunities ("fail fast, fail forward" is the mantra). Agentic AI can reinforce this by providing a more neutral lens with results. An AI won't sugarcoat test data, but it also won't assign blame – it just analyzes what happened. This can help teams detach their ego from the experiment's outcome and focus on the facts.

One powerful role for AI in validation is **rapid and automated testing**. Just as AI accelerated design iterations, it can accelerate experimental cycles. For example, AI-driven analytics can monitor how a new feature is used in real-time by thousands of beta users and immediately flag patterns – such as a particular tool crashing for 5% of users on Android devices. In the past, gathering and sifting such feedback might have taken weeks; an AI can do it in minutes and even suggest likely causes. AI can also run *simulations* to validate ideas in virtual environments before touching live customers. Engineers now use digital twin models (highly detailed simulations of systems) to test how a new process might behave under various conditions – reducing the risk when the real rollout happens. The Financial Times reported that using AI for simulations and

stress tests has cut time-to-market by up to **40%** in some R&D projects (Quinnox, 2025). In effect, AI enables teams to "practice in the simulator" and catch issues early. This not only speeds up learning but also builds confidence; the team's collective gut can relax a bit knowing they've tested many angles.

5 Domicilium – Validating Growth Strategy Through Emerging Trends

Context: Domicilium, the internet, and cloud service providers had reached a growth plateau. CEO Phil Adcock asked, "I'm stuck with a growth problem–could AI help?" The team wanted to validate whether their instinct on a new growth strategy could hold up under new market conditions.

Embodied dynamic: Phil and his leadership team brought strong instincts about where growth might come from but lacked evidence to test those intuitions. They sought confirmation that their internal sense of direction aligned with external market trends. Leaders often possess powerful instincts about where opportunities lie but struggle to validate these hunches against evolving market data, dormant internal assets, or adjacent innovations. This gap is especially pronounced in live workshops, where time constraints and cognitive load limit teams' ability to link intuition with real-time evidence (Yao, 2023; Wieland et al., 2022). AI-assisted facilitation bridges this divide by surfacing patterns across internal repositories and external signals during the session itself.

AI mediation: Over several days, the AI partner worked with the team to map patterns in historical growth and compare them with emerging industry signals. This alignment provided real-time validation of Phil's original strategy, connecting the intuitive insights from day one with concrete data and external evidence by the end of the week.

Embodied shift: As the patterns emerged, the team moved from uncertainty to conviction. Seeing their instincts reinforced by objective data created confidence and a sense of focus. The AI became a trusted collaborator, bridging subjective insight and objective validation.

Outcome: On the final day, Phil presented a validated growth strategy supported by AI-driven market evidence and his own experiential understanding. The work identified a £11.5 million opportunity and demonstrated how embodied knowledge, combined with AI mediation, provides a robust foundation for confident growth.

Of course, no simulation replaces reality entirely. When the innovation finally meets the real world – whether it's a new internal process being piloted, or a product launching to customers – that's when *embodied responses* really come into play. Users will experience a range of visceral reactions, including confusion, delight, frustration, trust, and distrust. Validating an innovation means paying attention to both human responses and hard metrics. Here, Agentic AI can serve as a **sensitive antenna**, picking signals that humans might miss or be slow to quantify. For instance, AI sentiment analysis might comb through customer comments or support chats to detect subtle shifts in sentiment ("Customers are excited about the new feature's speed, but a

few mention it 'feels confusing' to navigate"). It can also correlate physiological data if available – imagine a wearable on test participants showing spikes in heart rate at a specific step in a process, indicating stress. These clues hint at where a design isn't aligning with human expectations or comfort. Incorporating such feedback is part of **embodied validation**: recognizing that success criteria aren't just numerical (revenue up, errors down), but also experiential (people feel good using this, employees feel empowered when operating it).

An enlightened validation approach, therefore, uses **hybrid metrics**. It tracks technical KPIs *and* human factors. A team might set goals not only for improving conversion rates or cycle times, but also for reducing participant stress levels or increasing satisfaction. Measuring psychological impact requires moving beyond satisfaction surveys to physiological validation. Heart rate variability (HRV) tracking during innovation processes provides objective evidence of nervous system regulation; higher HRV indicates a ventral vagal state, which is conducive to learning and creativity. Organizations can establish baseline HRV measurements and track whether AI-mediated processes genuinely reduce threat activation or merely shift its location.

Similarly, competence revelation metrics validate whether hidden expertise is surfacing. In traditional settings, overconfident voices often dominate, while genuine experts remain silent—a dynamic explained by research on the impostor phenomenon (Clance & Imes, 1978). Individuals with authentic expertise often underestimate their abilities precisely because they understand complexity deeply enough to appreciate what they don't know. Adequate validation, therefore, tracks not just who participates but whose contributions prove valuable in implementation. If quieter participants' insights consistently drive successful outcomes, the process is successfully inverting dysfunctional social dynamics.

This validation approach mirrors clinical practice, where outcomes are combined with both objective measures (such as symptom reduction) and subjective experiences (improvement in quality of life). Reductions in anxiety and depression demonstrate that creating optimal conditions for human cognition produces measurable well-being benefits alongside performance gains. Innovation processes should be evaluated in both dimensions: Did we generate better solutions? Did participants experience engagement as energizing rather than depleting? When both metrics improve simultaneously, the approach is genuinely sustainable.

This mirrors practices in fields such as healthcare, where treatment is evaluated based on outcomes and patient-reported experiences. In organizational innovation, we might measure **the quality of participation – did the quiet voices from the Discovery phase contribute to the implementation of the final solution**? Are people openly sharing issues during the pilot, which indicates psychological safety, or hiding them? One could conduct before-and-after surveys of team climate to see if introducing the AI-mediated process has reduced anxiety or increased a sense of empowerment. In fact, early case studies show promising results: in one human-AI discovery program, participants reported a 29% reduction in anxiety and a 27% increase in

positive outlook on challenges after completing the process. At the same time, their measured problem-solving performance improved markedly, consistent with the boost we'd expect from flow states and better cognitive engagement. These are significant validation markers – they tell us not only that we built the thing right, but also that we built the right thing in terms of its human impact.

When an organization treats validation holistically, it truly closes the learning loop. The **Validate phase serves as a launchpad for the next cycle, as** lessons learned are fed back into the Discover phase of the next iteration. Perhaps the AI system logs that a particular line of questioning in Discover led to an insight that proved crucial, so that it will emphasize that next time. Or it notes that users struggled with a feature despite internal enthusiasm, signaling a blind spot in how the team empathized with customers – an insight to address in the next Design phase. Over time, this approach cultivates a *learning organization* with a digital nervous system, where AI and human feedback work together like sensory and motor neurons. Just as our human nervous system continuously monitors and adjusts our body's actions, a digital nervous system in an organization senses what's happening (both in data and in human emotions) and adapts processes accordingly. Microsoft's Bill Gates originally used the "digital nervous system" metaphor to describe real-time information flow in businesses. Still, here we give it new depth: it's not only information flow, but *insight flow* up and down the four stages of innovation.

Through the process of Discover, Design, Decide, and Validate, a common theme emerges, honoring the human element while leveraging the power of AI. Each phase supports embodied cognition by creating the right environment – whether it is private reflection, collaborative flow, balanced judgment, or safe experimentation – and in each, AI serves as an agent that amplifies our strengths and compensates for our weaknesses. It's a far cry from the fear that "AI will replace humans." Instead, humans remain very much *in the loop*, arguably **more** human than ever because we are focusing on what humans uniquely excel at (empathy, intuition, creative thinking, ethical reasoning) and offloading the rest to our AI partners. One consulting CTO aptly said, *"AI is a co-pilot, not a replacement"* (Quinnox, 2025). When AI is integrated with care, teams report feeling *more* empowered and creative, not less (HBS, 2025). This suggests that the right approach to AI in innovation increases **psychological safety** – people see the AI as having their back, not looking over their shoulder.

In summary, the **embodied nature of organizational knowledge** means we must design innovative processes that respect human psychology and tap into the wisdom stored in feelings, not just files. The addition of **agentic AI** – AI that actively participates as a team member – can transform these processes at every stage. It helps individuals safely share what they know in Discover; it accelerates and deepens creative work in Design; it brings data and clarity to decision-making in Decide; and it enables faster, fuller learning in Validate. The result is a virtuous cycle: a system where humans and AI work together to create an environment of curiosity, trust, and high performance. Innovation ceases to be a stressful high-stakes push and becomes more of a *pull* – drawing out the best ideas naturally. People often describe a kind of *emotional warmth* in

such environments: it feels **exciting, meaningful, and fun** to solve problems collaboratively. That human-centric warmth is not incidental; it's the very fabric of how breakthroughs happen. As we stand on the frontier of more **human-aware AI** tools, the opportunity is to build not just faster and smarter organizations, but *healthier and wiser ones* too. The technology may be cutting-edge, but the principle is ancient: when you truly listen – be it to a colleague's heartbeat of anxiety or an AI's data visualization – you learn. And when you know, you grow. The future of innovation belongs to those who can both listen to the people and leverage the machines in one integrated dance of discovery.

References (Harvard style)

- Clance, P.R. & Imes, S.A. (1978). *The imposter phenomenon in high-achieving women: Dynamics and therapeutic intervention*. Psychotherapy: Theory, Research & Practice, 15(3), 241-247.
- Craig, A.D. (2002). *Interoception: the sense of the physiological condition of the body*. Current Opinion in Neurobiology, 13(4), 500-505.
- Csikszentmihályi, M. (1990). *Flow: The Psychology of Optimal Experience*. New York: Harper & Row.
- Damasio, A. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Grosset/Putnam.
- Detert, J.R. & Edmondson, A.C. (2011). *Implicit voice theories: Taken-for-granted rules of self-censorship at work*. Academy of Management Journal, 54(3), 461-488.
- Dietrich, A. (2004). *The cognitive neuroscience of creativity*. Psychonomic Bulletin & Review, 11(6), 1011-1026.
- Edmondson, A.C. (1999). *Psychological safety and learning behavior in work teams*. Administrative Science Quarterly, 44(2), 350-383.
- Fenn, J. and Raskino, M. (2008) *Mastering the Hype Cycle: How to choose the right innovation at the right time.* Boston, MA: Harvard Business Press.
- Herman, J.L. (2015). *Trauma and Recovery* (2015 ed.). New York: Basic Books.
- Kotler, S. & Wheal, J. (2017). *Stealing Fire*. New York: HarperCollins.
- Kruger, J. & Dunning, D. (1999). *Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments*. Journal of Personality and Social Psychology, 77(6), 1121-1134.
- Lieberman, M.D., Eisenberger, N.I., et al. (2007). *Pain and social rejection: social pain and the brain*. In: *The Science of Social Vision*. Oxford University Press.
- Lu, T. & Zhang, Y. (2025). *Collaborative Decision-Making with AI: Field Experiment Evidence*. (Working Paper, Arizona State University).
- Mehling, W.E., et al. (2012). *The Multidimensional Assessment of Interoceptive Awareness (MAIA)*. PLoS ONE, 7(11), e48230.
- Ogden, P. & Minton, K. (2000). *Sensorimotor psychotherapy: One method for processing traumatic memory*. Traumatology, 6(3), 149-173.

- Porges, S.W. (2011). *The Polyvagal Theory: Neurophysiological Foundations of Emotions, Attachment, Communication, and Self-regulation*. New York: Norton.
- Rock, D. (2008). *SCARF: A brain-based model for collaborating with and influencing others*. Neuro Leadership Journal, 1(1), 44-52.
- Shapiro, F. (2001). *Eye Movement Desensitization and Reprocessing: Basic Principles, Protocols, and Procedures* (2nd ed.). New York: Guilford Press.
- Siegel, D.J. (1999). *The Developing Mind*. New York: Guilford Press.
- Siegel, D. J. (2012). Pocket Guide To Interpersonal Neurobiology: *An Integrative Handbook Of The Mind*. WW Norton & Company.
- Thayer, J.F., Hansen, A.L., Saus-Rose, E., & Johnsen, B.H. (2009). *Heart rate variability, prefrontal neural function, and cognitive performance: The neurovisceral integration perspective on self-regulation, adaptation, and health*. Annals of Behavioral Medicine, 37(2), 141-153.
- University of Sydney. Chi, R. P., & Snyder, A. W. (2012). *Brain stimulation enables the solution of an inherently difficult problem.* Neuroscience letters, 515(2), 121-124.
- van der Kolk, B. (2014). *The Body Keeps the Score: Brain, Mind, and Body in the Healing of Trauma*. New York: Viking.

Wieland, B., de Wit, J., & de Rooij, A. (2022). *Electronic Brainstorming with a Chatbot Partner: A Good Idea Due to Increased Productivity and Idea Diversity*. *Frontiers in Artificial Intelligence*, 5, Article 880673. DOI: 10.3389/frai.2022.880673.

Sugarwork. (2025). *Institutional knowledge is a $47 million/year opportunity for US businesses — here's how to approach it.* Available at: *https://www.sugarwork.com/resources/how-enterprises-can-transform-institutional-knowledge-loss-into-value-creationnbsp* (Accessed: September 2025).

Insights. (2025). *Why psychological safety drives successful teams.* Available at: *https://blog.insights.com/en-us/blog/why-psychological-safety-drives-successful-teams* (Accessed: September 2025).

Harvard Business School (HBS) (2025). *When AI joins the team, better ideas surface.* Working Knowledge. Available at: *https://www.library.hbs.edu/working-knowledge/when-ai-joins-the-team-better-ideas-surface* (Accessed: September 2025).

Quinnox. (2025). *AI for rapid prototyping: benefits, use cases & challenges.* Available at: *https://www.quinnox.com/blogs/ai-for-rapid-prototyping/* (Accessed: September 2025).

W. P. Carey School of Business. (2025). *Better together: how AI plus people improves some decision-making.* W. P. Carey News. Available at: *https://news.wpcarey.asu.edu/20250212-better-together-how-ai-plus-people-improves-some-decision-making* (Accessed: September 2025).

ACKNOWLEDGEMENTS

The views expressed in the *OMG Journal of Innovation* are the author's views and do not necessarily represent the views of their respective employers nor those of the Object Management Group® (OMG®), an EDM Association community.