

AI Is Making Knowledge Cheap and Judgement Scarce

And why assessment innovation has never mattered more.

Sylvain Rottier

Co-Founder MindMarqs, Psychologist
July 2026

Abstract: The widespread adoption of generative artificial intelligence is rapidly shifting corporate talent requirements by automating routine knowledge tasks and elevating the strategic value of human judgment. While AI accelerates the velocity of executive workflows, it simultaneously induces cognitive strain, attention fragmentation, and a passive reliance known as "cognitive surrender". Because traditional talent acquisition tools (such as CVs, interviews, and self-report psychometrics) can now be systematically gamed or optimized by AI, organizations require a new assessment paradigm rooted in physiological measurement to isolate genuine leadership skill from luck.

Key Structural Insights

- **The Judgement Premium:** As AI agents handle execution, human workers are increasingly confined to high-stakes verification, data review, and contextual governance.
- **The Attention Deficit:** Knowledge workers have experienced a steep, multi-decade decay in screen focus intervals, dropping from 2.5 minutes in 2004 to just 47 seconds by 2020, compounded by the cognitive cost of "attention residue" during task switching.
- **The Atrophy of Critical Thought:** Confident AI outputs frequently cause human operators to disengage independent reasoning, automating decision loops and causing long-term analytical decay.
- **The Assessment Failure:** Traditional recruitment filters have collapsed under AI-driven coaching, resulting in identical-looking applicant pools and measuring a candidate's prompting skill rather than raw capability.

- **The Observable Solution:** Advanced eye-tracking technology opens the decision-making process by recording involuntary gaze mechanics and micro-behavioral preferences.
 - **Empirical Performance Metrics:** Observable behavioral patterns dictate how evidence accumulates in the brain; analyzing these stable, un-gamable pathways separates structural competence from statistical randomness.
 - **Metacognitive Development:** Beyond selection, providing leaders with empirical data regarding their own visual data-parsing habits serves as the baseline for retraining focus and sharpening organizational judgment.
-

Section 1 : Market Context & The Cognitive Saturation Problem

In 2019, I led the first large-scale deployments of autonomous mobile robots across EMEA, machines cleaning hospitals, airports and offices without human operation. The business case was simple: consistent, tireless, scalable. The cleaning sector became one of the world's first major autonomous mobile robot markets.

But the market did not simply swap humans for machines. At our customers, cleaning staff had to plan, coordinate and interact with systems rather than follow physical instructions. Clients started receiving data and demanding explanations, integrations, governance.

Internally, the capability change required was equally vast. Customer service teams fielded questions that had nothing to do with the machine itself. Simple business cases became transformational implementation plans. Our annual product update became a continuous, fast-evolving service. The work didn't disappear. It changed, transforming roles far beyond the cleaning function of the robot.

Some people thrived. Others struggled. And watching that unfold made one thing clear: when technology shifts, the real impact is not what it technically solves. It is how it changes the requirements placed on the workforce affected by it, not just those working directly with it, but the entire ecosystem of capabilities.

AI is the same shift, at a far greater scale and speed.

Fewer Tasks, Higher Stakes

As I spoke with leaders who implemented AI, a consistent pattern emerged: the technology reduces the value of what we know and increases the importance of how we think.

As AI handles more tasks, humans must provide more inputs to review, more outputs to evaluate, more context to weigh. Fewer tasks. More judgement. Higher

stakes. And the consequences of every judgment are amplified by what an agentic AI workforce can execute on the back of it. Leaders report the same challenge: the hardest problem is not technology. It is maintaining accountability at the speed AI enables.

This is not new. Every major technological transition reshuffles the hierarchy of human capabilities. The printing press made literacy non-negotiable. Electrification elevated those who understood systems over those who understood only machines. AI is doing the same, but at the cognitive level, and across every sector simultaneously.

We are entering a world where judgement matters more than ever.

Section 1 : We Are Losing the Capability We Need Most

[Herbert Simon](#) wrote in 1971 that “a wealth of information creates a poverty of attention.” He was making a theoretical argument. How precisely does that land today, fifty years later, as people feel increasingly scattered in a world of constant notifications and AI-generated content.

[Kahneman and Tversky's](#) foundational research (Science, 1974) established the mechanism: under cognitive load, people default to heuristic shortcuts, anchoring on the first number they see, confirming what they already believe, simplifying what is complex. These are not failures of intelligence. They are the predictable responses of a capable system under strain. When attention is fragmented and cognitive resources are taxed, deliberate reasoning cannot fully engage.

[Gloria Mark](#) and colleagues at UC Irvine have tracked focused task time in knowledge workers for over two decades. The decline began well before AI, driven by the rise of the internet and smartphones, and it is steep:



Mark, Attention Span, 2023

Two thirds of sustained attention lost in twenty years. But the problem runs deeper than duration. Each time attention switches, there is a recovery cost, what researchers call attention residue. [Rubinstein, Meyer and Evans](#) established that task-switching produces measurable cognitive penalties: lost time, increased errors, reduced depth of thinking.

The problem is not just that we focus for less time. It is that every switch is a tax on the quality of whatever comes next.

Alongside attention, critical thinking is under measurable strain. A 2025 study by Microsoft and Carnegie Mellon University surveyed 319 knowledge workers and found a direct pattern: the more confident workers are in AI's ability to complete a task, the more they disengaged their own thinking. We are increasingly in a cognitive state resembling what is commonly called doomscrolling, whether through work email, social media or AI prompting: a semi-automatic loop with reduced awareness that crowds out reflection, mental rest and critical thinking.

The Rise of Cognitive Surrender

A more precise term has entered the literature: cognitive surrender. A 2024 systematic review in Smart Learning Environments documented how over-reliance on AI dialogue systems progressively impairs critical thinking and analytical reasoning, not through any failure of the technology, but through the atrophy of capabilities we stop exercising when we hand thinking to a machine.

This is a deeply human reflex, amplified at scale. When a contract comes back from a lawyer, most people do not read it critically, they take it as it is. With AI, that instinct accelerates: the system is always available, always confident, and never visibly wrong. We are not outsourcing thinking occasionally. We are building a reflex of doing so.

AI is raising the bar for human cognition at the exact moment we are losing the ability to clear it.

While traditional psychometric assessments have collapsed under the weight of generative AI scripting, forward-thinking enterprise organizations are shifting from output-based screening to live behavioral observation. [MindMarqs](#) resolves this validation crisis by capturing the raw, un-gamable reality of executive decision-making. By mapping real-time gaze mechanics during simulated market crises, the MindMarqs platform completely bypasses coached application artifacts, isolating true strategic intent from automated or rehearsed responses.

Section 3: The Structural Collapse of Traditional Assessment

If judgement is the critical capability, the obvious question is: how do we identify and develop it in people? This is where things get uncomfortable, because the tools built to answer that question are themselves breaking down.

Most assessment methods were built for a world where a direct link existed between the applicant and their output, reflecting genuine effort and skill. A CV requires you to sit down and write it. An interview required preparation. Self-report instruments were always imperfect, [Crowne and Marlowe](#) documented their limitations as early as 1960, but self reports were a genuine window into who a person was and what they could do.

Generative AI breaks that link. A candidate today can produce a CV optimised to mirror a job description with precision no human editor could match. They can rehearse interview answers with an AI that generates structured, plausible responses to any behavioural question. They can even complete a personality questionnaire by asking a language model how to respond. Several pre-seed investors have told me the same thing: pitches are starting to look alike. The language is polished, the structure impeccable, the differentiation quietly gone.

- **CV Hyper-Optimization:** Language models seamlessly structure resumes to mirror enterprise job descriptions with absolute algorithmic precision, neutralizing standard screening tools (p. 3).
- **Interview Rehearsal Scripting:** Candidates utilize conversational AI engines to generate highly structured, optimized, and plausible responses to traditional behavioral assessment interview questions (p. 3).
- **Psychometric Manipulation:** Large language models can easily parse and answer personality questionnaires to project whatever corporate profile a specific firm desires (p. 3)

This is not about dishonesty. It is structural. AI makes the differences between us less visible. What remains is increasingly a measure of how well someone uses the tool, not of the actual capability the assessment was designed to reveal.

The industry's response has been to lean toward skills-based hiring: case studies, simulations, work samples. This is the right direction. Research by [Schmidt and Hunter \(1998\)](#) established that work sample tests have meaningfully better predictive validity than personality questionnaires, and meta-analyses by [Sackett and colleagues \(2022\)](#) confirm the pattern holds. Assessing what someone can actually do is better than asking them to describe what they think they can do.

But skills-based assessment as we know it has structural limits. It either focuses on a specific task, coding for example, or requires expensive full-day assessments feasible only at the final stages of selection. If we are to assess attention control and judgement at scale, we need tools that are easily deployed, and that AI cannot game.

We are no longer measuring capability. We are measuring how well people use AI.

Section 4: Technical Deep Dive: Eye-Tracking & Cognitive Signals

My background is in psychology, social and cognitive psychology in particular. I have followed the literature for decades while leading multinational businesses through technological change. People often ask what served me more: my economics degree or my psychology studies. I never had to think long about it. Psychology. At the heart of leadership is understanding people, and specifically how individuals and teams make decisions.

Most decision-making is not conscious. Ask someone how they make decisions and they will either say they do not know, or describe a rational process they rarely follow. In practice, what drives an outcome happens largely below the threshold of awareness: how broadly or narrowly we scan information available, which heuristics we deploy, how quickly deliberate reasoning gives way to intuition. These micro-behavioral patterns are stable across individuals, predictive of outcomes, and almost entirely inaccessible through self-report.

Gaze patterns distinguish a person who made a good decision for the right reasons from one who simply got lucky. [Krajovich, Armel and Rangel's](#) landmark 2010 established the mechanism: fixation patterns causally shape how evidence accumulates toward a choice, people integrate approximately three times more evidence for the option they are looking at.

The 2022 systematic review by [Borozan, Cannito and Palumbo](#) in the Journal of Behavioral and Experimental Finance synthesised 51 peer-reviewed studies and reached a conclusion the assessment industry has not yet absorbed: eye-tracking provides access to cognitive processes underlying decisions that no self-report method can reach. [Lahey and Oxley](#), writing in Judgment and Decision Making in 2019, demonstrated that simple eye-movement metrics predict future financial decision performance.

What these patterns reveal, what researchers call micro-behavioral preferences, includes how much information someone genuinely engages with before committing to a choice, how speed trades off against accuracy under pressure, which cognitive biases are most likely to distort their judgement, and whether they can sustain deliberate reasoning when the easier option is to trust intuition.

These are stable, learned behavioral signatures of the cognitive process itself. They cannot be fabricated. A candidate can optimize a CV or rehearse interview answers. They cannot alter the unconscious patterns in their gaze under a real decision task.

This is why eye-gaze measurement is more than an interesting technique, but a scalable way to observe our attention control and cognition directly. What you look at when making a decision is not random, it reflects what you are attending to, what you are processing. Gaze is largely unconscious, but still behavior. It cannot be coached in any systematic way, until you are made aware of it.

How we make decisions does not have to be a black box. We now have the tools to open it.

The financial risk of a bad executive hire in capital markets demands a selection methodology backed by hard neuroscience rather than subjective interviewing. [MindMarqs](#) translates decades of peer-reviewed cognitive literature into scalable, enterprise-ready behavioral signals. The platform's proprietary technology decouples structural problem-

| solving from statistically lucky outcomes, giving HR executive managers the definitive signals needed to protect institutional capital and future-proof their leadership succession.

Section 5: Strategic Business Benefits for High-Pressure Sectors

AI raises the bar for human judgement at the exact moment the underlying capabilities are under pressure. The demand for sustained attention, cognitive inhibition and critical thinking has never been more relevant, precisely as we are asked to input, evaluate and govern an agentic AI workforce. The decline is not inevitable and not universal. Attention, focus and independent thinking are learned behaviours. We can develop them. But first we need to see clearly where we are.

Organisations that fare best through technological change are those that understand what the technology requires of their people, and invest in closing the gap. That starts with measurement.

The cognitive processes that drive judgement, attention control, information processing, resistance to bias, can now be tracked, benchmarked and developed. What was once a black box is opening. For the first time, we have tools that can show an individual or an organisation not just what decisions were made, but how, and where the specific vulnerabilities lie.

Most people sense that their attention is not what it was. Few have seen the data on their own cognitive patterns. That matters, because awareness is the first lever. Research on metacognition consistently shows that people who understand their own decision-making processes make better choices, not because they become smarter, but because they become more deliberate. Once you can see how you are thinking, you can start to change it.

From awareness we can build, strengthening the specific capabilities under strain. The science here is encouraging. Attention is trainable. Critical thinking improves with practice. Cognitive surrender is a reflex, and reflexes can be retrained. Organisations that understand this will not just navigate the AI transition but raise the capabilities of their workforce with it.

We are not short on intelligence. We are short on Judgement when it comes to AI.

And I believe we can close this gap.

References

- Simon, H.A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), Computers, Communications, and the Public Interest. Johns Hopkins Press.*
- Kahneman, D. & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185(4157), 1124–1131.*

- Mark, G. (2023).** *Attention Span: A Groundbreaking Way to Restore Balance, Happiness and Productivity.* Hanover Square Press.
- Mark, G., Gonzalez, V.M. & Harris, J. (2005).** No task left behind? Examining the nature of fragmented work. *Proceedings of CHI 2005*, ACM Press.
- Rubinstein, J.S., Meyer, D.E. & Evans, J.E. (2001).** Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 27(4), 763–797.
- Lee, H., Sarkar, A. et al. (2025).** The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. *Proceedings of CHI 2025*, ACM Press.
- Gerlich, M. (2025).** AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), 6.
- Smart Learning Environments (2024).** Over-reliance on AI dialogue systems and its impact on decision-making and critical thinking: A systematic review.
- Crowne, D.P. & Marlowe, D. (1960).** A new scale of social desirability independent of psychopathology. *Journal of Consulting Psychology*, 24(4), 349–354.
- Schmidt, F.L. & Hunter, J.E. (1998).** The validity and utility of selection methods in personnel psychology. *Psychological Bulletin*, 124(2), 262–274.
- Sackett, P.R., Zhang, C., Berry, C.M. & Lievens, F. (2022).** Revisiting meta-analytic estimates of validity in personnel selection. *Journal of Applied Psychology*, 107(12).
- Flavell, J.H. (1979).** Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, 34(10), 906–911.
- Dunning, D., Johnson, K., Ehrlinger, J. & Kruger, J. (2003).** Why people fail to recognize their own incompetence? *Current Directions in Psychological Science*, 12(3), 83–87.
- Krajovich, I., Armel, C. & Rangel, A. (2010).** Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10), 1292–1298.
- Lahey, J. & Oxley, M. (2019).** Simple eye movement metrics can predict future decision-making performance: The case of financial choices. *Judgment and Decision Making*, 14(3), 223–233.
- Borozan, M., Cannito, L. & Palumbo, R. (2022).** Eye-tracking for the study of financial decision-making: A systematic review of the literature. *Journal of Behavioral and Experimental Finance*, 35, 100702.
- Stanovich, K.E. & West, R.F. (2000).** Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences*, 23(5), 645–665.