



AI believes everything. You don't.

How to critically evaluate AI output and apply human judgment where it counts

A whitepaper by Laurens Vandebergh, newITera
Business Development Manager AI and Project Manager
2026

Executive Summary

AI produces output with impressive ease and great confidence. But ease and confidence are no guarantee of quality. Anyone who blindly adopts AI output is not only taking over the value, but also the errors, the assumptions, and the blind spots.

Discernment is the third core competency within AI Fluency and is about the ability to critically evaluate AI output. Not with distrust, but with the professional judgment that belongs to responsible use. Discernment is the counterpart to Description: where Description is about what you instruct AI to do, Discernment is about what you get back.

This whitepaper explains how Discernment works at three levels, which signals to recognize when AI goes wrong, and how to give effective feedback to structurally improve the quality of AI collaboration.

Discernment is your quality control system for AI collaboration. Not the last line of defense, but an active and continuous part of how you work with AI.

The problem: AI always sounds certain

One of the most deceptive characteristics of large language models (LLMs) is their tone. AI presents incorrect information with the same fluent confidence as correct information. There is no hesitant pause, no footnote, no warning. The output reads convincingly, even when it is wrong.

This makes Discernment harder than it appears. With a human colleague, you recognize doubt through tone, body language, or phrases like “I believe...” or “as far as I know...”. AI does not give those signals. The responsibility to judge lies entirely with you.

Three situations where Discernment is critical

Do you recognize these moments?

- AI presents specific figures, names, or sources. These are the most sensitive spots for hallucination. The more specific the claim, the greater the risk.
- AI gives advice or recommendations. The model does not know what your organization truly needs. It extrapolates from patterns, not from context.
- AI output goes directly to a client, partner, or management. At that point there is no room for correction after the fact.

In all these situations, Discernment is not a luxury but a baseline requirement. Not because AI is poor, but because responsibility for the end result always lies with the person, not the model.

What is Discernment within AI Fluency?

Discernment is the third competency within the 4D framework of AI Fluency. It is the mirror competency of Description: where Description is about giving clear instructions, Discernment is about making clear judgments. Together they form a continuous feedback loop that steadily improves the quality of AI collaboration.

Discernment works at three levels:

PRODUCT DISCERNMENT	PROCESS DISCERNMENT	PERFORMANCE DISCERNMENT
---------------------	---------------------	-------------------------

<i>Evaluate the quality of the AI output itself</i>	<i>Evaluate how AI arrived at the output</i>	<i>Evaluate how AI behaved during the collaboration</i>
Is it correct? Does it fit? Does it add value?	Is the reasoning logical? Are the steps transparent?	Is it communicating at the right level and in the right tone?

Product Discernment: evaluate what you receive

The most direct form of Discernment is evaluating the output itself. Is the information factually correct? Does the structure fit the goal? Is the tone appropriate for the audience? Does it genuinely add something to what you already knew?

Product Discernment requires domain expertise. You can only assess the quality of a legal opinion, a financial analysis, or a technical report if you know the field. This is precisely why the best AI users are first experts in their own domain.

Process Discernment: evaluate how AI reasoned

Sometimes the output seems acceptable at first glance, but the reasoning underneath is flawed. AI can reach a correct conclusion via faulty reasoning, or get caught in a circular argument without this being immediately visible in the output.

Process Discernment means: look not only at what AI says, but also at how it got there. Ask AI to explain its reasoning. Check whether assumptions were made explicit. Recognize when the model is answering the wrong question.

Performance Discernment: evaluate how AI behaves

The third level concerns the interaction itself. Is AI communicating at the right level for your audience? Does it respond appropriately to corrections? Is the collaboration efficient, or does it take increasing effort to keep the conversation on track?

Performance Discernment gives you the signals to course-correct in time: by adjusting your instructions, refreshing the context, or restarting the conversation from a sharper foundation.

Discernment in practice: a working checklist

The checklist below gives the key questions to ask yourself at each level when evaluating AI output, along with concrete actions when in doubt.

PRODUCT DISCERNMENT	PROCESS DISCERNMENT	IF IN DOUBT
• Is the information factually correct? • Does the output fit the goal and audience? • Is the structure clear and coherent? • Does it genuinely add value?	• Is the reasoning logically structured? • Were steps skipped or assumptions made? • Is the order of arguments correct? • Did AI focus on the right problem?	• Verify factual claims independently • Ask AI to explain its reasoning • Restart with sharper instructions

From evaluation to improvement: giving effective feedback to AI

Discernment does not end with identifying a problem. What follows is at least as important: how do you give AI feedback that actually leads to better output? Vague comments produce vague improvements. Effective feedback is specific, names the problem, and gives direction.

IF YOU SEE THIS	THE PROBLEM IS	SAY THIS
<i>Output is factually incorrect</i>	Product Discernment: hallucination or knowledge boundary	"This is incorrect: [fact]. Correct this and indicate where you sourced it."
<i>Answer misses the point of my question</i>	Process Discernment: AI solved the wrong problem	"My actual question was X, not Y. Try again with that focus."
<i>Tone or level is off</i>	Performance Discernment: wrong behavioral register	"The tone is too [formal/informal]. Rewrite for [audience]."
<i>AI keeps circling the same answer</i>	Process Discernment: circular reasoning or stuck context	"Start fresh: forget the previous answer and approach this from [different angle]."

The pattern is consistent: effective feedback names the problem concretely, explains why it is a problem, and gives AI a clear new direction. The more precise the feedback, the greater the chance of a usable follow-up.

What does Discernment require from an organization?

Individual Discernment is a personal skill. But organizations that invest in it structurally build something more durable: a culture of critical thinking around AI use.

Domain expertise remains irreplaceable

Discernment is inseparable from subject matter expertise. Someone who is not an accountant cannot adequately evaluate an AI-generated financial statement. Someone who is not a lawyer cannot screen an AI-generated contract for legal risks. AI amplifies the productivity of experts. It does not replace the expertise.

Review processes as a standard

Organizations that take AI seriously build review steps into AI-supported work that leaves the organization or influences decisions. Not as bureaucratic control, but as a professional standard, comparable to how an editor checks texts or a controller verifies figures.

Discernment as a shared language

When a team knows the three levels of Discernment and uses the same terminology, feedback on AI work becomes more concrete and efficient. "This is a Product Discernment issue: the facts are wrong" is a more precise signal than "this doesn't seem quite right." Shared language accelerates learning.

Discernment and Diligence go hand in hand

Discernment is the skill of recognizing errors. Diligence is the responsibility to act on them. Someone who spots a mistake in AI output but sends it anyway has Discernment but not Diligence. Both competencies are needed for responsible AI use.

How do you develop Discernment in your organization?

Discernment is a skill you develop through practice, reflection, and sharing experiences. A phased approach works best:

- **Step 1** — Make it visible: discuss in your team which AI output was wrong or incomplete over the past month. Not to assign blame, but to recognize patterns. What types of errors come up most often? Are they Product, Process, or Performance issues?

- **Step 2** — Build the habit: choose one type of AI task your team performs regularly and make the Discernment checklist part of your workflow. Small rituals are more effective than sweeping policy changes.
- **Step 3** — Learn from feedback: document which feedback formulations work best for your specific AI use. Build a shared library of effective corrections and rephrasing. That is knowledge that adds value for the whole team.

An AI Readiness Assessment maps out how strongly your organization currently scores on Discernment and the other AI Fluency competencies, and which targeted steps will have the most impact.

Conclusion

AI believes everything. It has no doubt, no intuition, and no sense of what does or does not fit in your specific context. That is not a weakness of the technology. It is a characteristic that calls for an equally strong counterpart: human judgment.

Discernment is that judgment in structured form. It turns AI collaboration from blind trust into a considered partnership, where you determine what makes it to the finish line.

Organizations that take Discernment seriously do not just get better results from AI. They also build the trust needed to embed AI structurally in their work processes.

How sharp is the Discernment in your organization?

newITera's AI Readiness Assessment maps out how your people, processes, and governance score across all four AI Fluency competencies, including Discernment.

Contact us at laurens.vandebergh@newitera.nl or visit www.newitera.nl

About newITera & Laurens Vandebergh

newITera is a Dutch IT consultancy based in Beuningen, Gelderland. Founded in 2011, initially as a specialized SAP firm, newITera has grown into a broad digital transformation partner. Today, newITera combines deep SAP expertise with services in data analytics, AI, cloud strategy, and low-code platforms. With a team of more than a hundred consultants, newITera serves leading enterprises and multinationals across their full digital transformation journey. What sets newITera apart is its combination of early adoption and deep specialization, as an independent partner that values human connection just as highly as technical excellence.

Laurens Vandebergh is Business Development Manager AI and Project Manager at newITera. With a broad background in technology and organizational change, he guides organizations in defining their AI strategy, conducting an AI Readiness Assessment, and realizing concrete use cases.

Contact: laurens.vandebergh@newitera.nl | www.newitera.nl