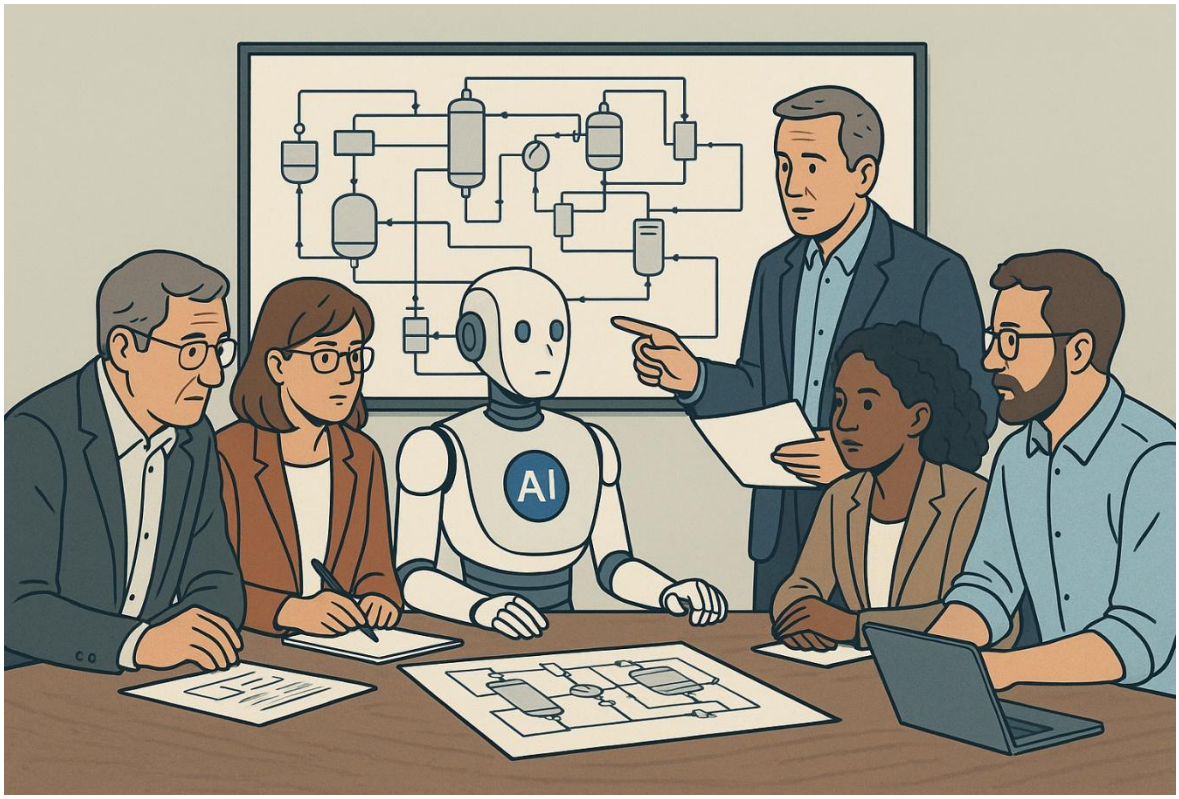

AI – The Wizard of Haz?

David Slater

There is a quiet shift underway in process safety. Artificial intelligence systems are beginning to appear not as peripheral search tools or novelty gadgets, but as capable assistants in hazard studies. They can read process descriptions, scan P&IDs, interpret operating procedures, and retrieve relevant incidents faster than any human in the room. This naturally leads to a provocative question:

If AI can already perform the structured elements of a HAZOP, then what exactly are the people there to do?

It is a variation of what we could call the *inverse Turing test*. In education, once AI could generate essays that passed university assessment criteria, it exposed the weakness not of students—but of the assessments themselves. If a machine can successfully complete the task, the task may not be measuring human understanding at all.

Applied to process safety, the thought is uncomfortable because it exposes something we rarely articulate: a significant portion of HAZOP effort has traditionally been spent on **information handling**, not thinking. The worksheets, formatting, cross-references, incident recall, guideword prompts and transcription are essential—but they are not where insight lives.

And yet, despite everything AI can now automate, HAZOP teams still matter. Their value lies not in reproducing what is already written down, but in what happens *between* the lines: curiosity, challenge, doubt, judgement and shared interpretation. These are the qualities no model can genuinely replicate.

The Chair: From Gatekeeper to Catalyst

Traditionally, the Chair is responsible for managing process, pace and fairness. In experienced and confident teams, this can feel like orchestration. But in more hierarchical cultures—where questioning senior engineers or contractors may feel disrespectful—the Chair's role becomes much more subtle and essential. They must help the team voice concerns they would otherwise keep to themselves.

Historically, Chairs achieved this through questioning, tone, silence—and, at key moments, **war stories**. A well-chosen incident from experience or industry memory can shift the room from polite observation to genuine exploration.

Now, AI expands that capability. With access to global incident data, regulatory findings, technical literature and near misses, AI can act as a library of lived experience—one far larger than any single practitioner could accumulate alone. Used well, it strengthens rather than displaces the Chair's influence: the experience remains human-led, but the examples become richer, broader and more relevant.

The Secretary: From Scribe to Sense-Maker

In every HAZOP, the Secretary ensures continuity and accuracy. They track actions, record discussions, maintain structure and prevent insights from being lost.

AI can now handle much of this administrative load with perfect recall. It can format consistently, flag contradictions and maintain traceability—quietly doing what humans often struggle with when hours of intense discussion accumulate into fatigue.

This doesn't remove the need for the Secretary. Instead, it shifts their role: from capturing words to ensuring that what is captured reflects the *intent* of the conversation. In other words, from typing to interpreting.

The Team: Interpretation, Not Enumeration

With AI now capable of presenting the "obvious" hazards and failure modes, the human contribution becomes clearer. Teams are not there to identify every deviation—they are there to **interpret** those deviations in context.

For example, AI may generate a plausible risk: reverse nitrogen inflow pressurising a pump header. But the team knows whether operators routinely rely on "feel" rather than instrumentation, whether winter startup conditions behave differently, or whether maintenance shortcuts change how equipment actually performs. These contextual insights turn abstract hazard into credible scenario.

This is where the "Wizard of Haz" metaphor becomes useful. When AI is first introduced in polished form—confident, seemingly omniscient—the team's guard goes up. The first hallucination becomes the character-unmasking moment, like Dorothy pulling back the curtain to find an ordinary man working levers.

But when AI is introduced transparently—limitations first, capability second—the experience changes. It becomes less an oracle and more a collaborator. A tool that helps the team start deeper, not faster.

Why the Framing Matters

Prospect theory offers a useful insight here: people are more motivated to avoid loss than to pursue a gain. If AI is framed as efficiency—faster meetings, reduced effort—adoption will be cautious.

If instead the message is:

Using AI in this way lowers the chance of missing a major accident hazard that could harm people or the business, because it brings global memory, consistent logic and fatigue-free recall into the room,

then resistance softens. Because the greatest fear in hazard analysis is not doing more work than necessary—it is discovering too late that something critical was overlooked.

So, Is AI the Wizard of Haz?

In a way, yes—but not in the sense of an all-knowing magician.

AI is the quiet intelligence behind the curtain: tireless, consistent, deeply knowledgeable and surprisingly useful. But the human team remains essential. Only they can apply judgement, weigh trade-offs, read silence, notice discomfort—and stand behind the conclusions.

No AI system will ever sit in a regulatory hearing or a boardroom and say,
"We believed this configuration was safe."

Only people do that.

So rather than resisting AI or surrendering to it, the opportunity is to treat it as a **co-intelligence**: a tool that prepares the ground, broadens awareness and sharpens discussion—while humans remain firmly responsible for the decisions that matter.

Used well, AI will not replace the HAZOP team.
It will make them better.

References & Further Reading

Keddie, A. (2025). *PS Blog 101 – Meet Your Wizard of Haz*.
Process Safety Matters. *(The article that sparked this discussion and named the "Wizard of Haz" moment.)*

Slater, D. (2025). *The Inverse Turing Test – Why AI Forces Us to Rethink What Learning and Assessment Really Mean*.
ResearchGate.
https://www.researchgate.net/publication/397601199_The_Inverse_Turing_Test_-_Why_AI_Forces_Us_to_Rethink_What_Learning_and_Assessment_Really_Mean

Slater, D. (2025). *A Functional Systems Approach to Hazard and Operability Analysis*.
ResearchGate.
https://www.researchgate.net/publication/397676889_A_Functional_Systems_Approach_to_Hazard_and_Operability_Analysis

Slater, D. (2024). *FRAMOP – Extending HAZOP with FRAM*.
ResearchGate. https://www.researchgate.net/publication/394379279_FRAMOP_-_Extending_HAZOP_with_FRAM