

Artificial Intelligence (AI) at the boardroom table

December 2025

On a random Thursday afternoon a few people are gathered in a room. On the one side of the table are the uneasy members of an interview panel and on the other side of the table a very nervous Sally is trying her best to provide “correct” answers to highly conceptual questions. All Sally has to back herself is a somewhat inflated resumé, enthusiastic references and a set of psychometric tests with results that are difficult to interpret. Yet, here they are, Stranger Sally trying to impress and the panel members trying to “get to know” Sally to be in a position to make a decision on appointing the best person to the board. Stranger Sally is given the benefit of the doubt and, somewhat uneasily, she is welcomed aboard.

In another room on another random Thursday afternoon, an interview panel confidently and decisively concludes a new appointment to the board. They will officially appoint Perfect Percy the AI Robot as their newest board member. Percy is trustworthy, always executes instructions, is much smarter than your average “smart” person and considerably faster than the most nimble-minded people in the world. We know Percy, we trust Percy, we appoint Percy. Percy can even recommend his own appointment with benefits and disadvantages and a full risk assessment - very transparent.

But...Do we really know Percy? Is Percy really trustworthy? Does Percy really never make mistakes? Can Perfect Percy really think of everything that he is supposed to think of, better than nervous, questionable Stranger Sally? Whatever the answers to these questions, Perfect Percy has landed himself a seat at the boardroom table, because boards that do not make use of AI, will be left behind. Let us however take a closer look at Percy’s resumé and phone a few references.

Maths: $1 + 1 =$ Whatever everyone else says it is

As Forbes^[1] puts it: “Some experts note that AI is **dysfunctional at math**. It tends to produce wrong answers and can be slow to correct them.” So...Percy is not very good at maths, because he does not think for himself and copies answers from those sitting around him.

Unlike widely used mathematical computational systems, e.g. Mathematica or a calculator, AI is not deterministic^[1]. These computational systems follow a set of predefined rules, resulting in exactly the same (repetitive) and correct answers to the same questions. AI, being a non-deterministic Large Language Model (LLM) works with algorithmic pattern recognition and does not have a set of rules to follow. AI could therefore produce different output for the same input every time.

As an LLM, AI is trained by all the information on the open Internet to be able to produce content according to its own research, or “training”; AI is a “stochastic parrot” who has no actual understanding of the text and answers it produces^[2]. In the same way that AI might correctly guess / deduct a missing word in an idiom, AI would probably also correctly guess the answer to $1+1$, since idioms and correct answers to $1+1$ should be widely available on

the Internet. But as the math problems become more complex and less readily available on the Internet, the chance of AI supplying a correct answer becomes less. That said, it is not impossible to integrate LLM's with something such as Mathematica, allowing the use of both the analytical and conceptual worlds.

Percy's hallucinations

Percy does not like saying "I do not know"... it rarely does; instead Percy hallucinates and he may give a nonsensical, incorrect response. AI hallucinates when it has to fill in the gaps based on patterns it sees in data when enough accurate information to perform the task^[3] is not available. This perceived lack of information that AI experiences could be insufficient training data, lack of context, inaccurate data or unclear constraints.

Insufficient training data

Insufficient training data could be due to a lack of relevant information on the open Internet, or a lack of relevant information supplied directly by the user for "training" purposes. If there simply is not enough information for AI to become "smarter" than it currently is^[4], AI has reached a plateau. Similarly, users may not be able to supply enough training material of their own to cover all tasks required of AI. All of us have experienced this during call centre chats with a bot where unlimited irrelevant responses can lead to frustration.

This is similar to a student writing a test. Unsure of the meaning of the question, or unsure of the answer, the student decides to guess and write down a plausible answer, even though it might be wrong^[3]. Leaving a blank space, or saying "I do not know", is not an option.

Lack of context

The context provided to AI might seem more than detailed enough to the user, but AI still cannot make sense of the information and guesses a response incorrectly.

Let us return to the parrot analogy, where the parrot knows a few phrases, but does not understand them. When asked a question, the parrot may repeat a known phrase or make something up^[3] just to try to please the person. AI could do this when it does not have enough context.

Inaccurate data

Garbage in, garbage out. The data provided to AI may be flawed, such as people making the same error and these multitudes of errors end up on the open Internet, or the data provided directly to AI by the user may be erroneous or biased. The parrot can only repeat what it has learned.

Unclear constraints

Again, as with providing context, it is difficult for a user to provide an exhaustive list of constraints and AI ends up playing outside the boundaries.

Preventing Percy from making things up could be an impossible task, since it requires the

user to supply an exhaustive set of data, context and constraints, all the while ensuring that all data is error free, unbiased and relevant to the task at hand. AI output in any case needs to be carefully checked and controlled. Hallucinations also are not always that obvious, which makes life difficult for the person in charge of quality assurance.

Percy the mole

What happens with information once it is in the hands of AI? Is Perfect Percy really trustworthy, or is he a mole by design? And with this uncertainty about what happens to your data, a new phrase is created as a solution to a simmering problem - AI sovereignty. And with this new phrase comes ready-to-buy or ready-to-outsource solutions, but at a cost. Another solution to this problem would be to ensure a closed AI environment, where your data is not shared outside of your own AI software. Yet another solution is to anonymise all data, but this is not always a viable solution, for example if data sets are too large and anonymisation cannot be fully guaranteed.

What is AI sovereignty? “At its core, sovereign AI refers to the idea of AI systems and infrastructures that are built to comply with the specific data sovereignty and regulatory requirements of a particular country or region. Sovereign AI infrastructure ensures that sensitive data generated, stored and processed by AI applications remains within the jurisdiction of a country’s laws. This is particularly crucial for industries like healthcare, finance, and government, where data privacy and protection are of utmost importance.”^[5]

AI is still mostly developed and controlled by a few large companies and a few countries, making use of servers and cloud infrastructure that are not necessarily in the same country or region as your data^[6]. AI companies often process and store data in multiple locations^[5], making Percy a possible mole by design, with your data travelling internationally and operating under unknown sets of rules, without getting a stamped passport. Sovereign AI solutions aim to address what is probably the largest question mark that hangs over AI, which is that of data security as well as complying to country-specific legislation such as POPI (South Africa) and GDPR (Europe).

However you want to address this issue, a few things are certain:

- The use of AI puts data security at risk
- As with most problems in the IT world, a solution to this problem is already taking shape in the form of AI sovereign solutions, ensuring a completely closed AI environment or anonymising data
- But, as with most problems in the IT world, solutions come with trade-offs and in this case refusing to share your data may result in you not getting access to all the publicly available data that is part of what makes AI so amazing
- Solutions come at a cost, which will include not only hardware and software, but acquiring staff with scarce and expensive skills
- Having a solution to this problem will become more of a requirement both by concerned customers as well as by law
- Bad actors will certainly attempt to break into such systems for nefarious purposes, or countries may build back-doors into systems that may allow them to use information to their advantage or affect operations in competitive or conflicting scenarios.

Getting the right expertise on board seems to be the crucial point on this list to ensure an AI future that keeps your integrity in place, is worthy of clients' trust and keeps you from being fined by authorities.

Perfect Percy: a Benevolent Maverick or an Average Joe?

The answer to this question is that not only is Perfect Percy an Average Joe, but the more we hang out with Perfect Percy, the more we are all at risk of becoming Average Joe's ourselves.

Perfect Percy is thinking for us, essentially causing our brains to stop thinking. The brain is designed to conserve energy and if it receives a signal that something has been taken care of, it disengages and stops putting effort into - and focusing on that something^[7]. For example, not being able to remember that you locked the front door is due to your brain receiving a signal - as you turn the key - that this task has been taken care of and thinking, "I do not have to remember this anymore". In other words, Percy is making us cognitively lazy, not because we choose to be lazy, but because if Percy has already thought about something, it causes our brains to disengage on the matter.

The more we rely on Percy to think for us, the lazier we will become in the long-term, until we fall asleep at the wheel. At a fundamental level, AI is reshaping how we process information and make decisions, often diminishing our reliance on our own cognitive abilities^[7].

Continuous exposure to AI-generated ideas also leads to a reduction in the variety and originality of ideas, which is known as the homogenising effect^[7]. Even though people using AI assistance may initially produce more, better and faster ideas, this initial improvement soon wears off and even after some time without AI assistance, this homogenising effect still remains.

In theory it is possible to refuse an AI suggestion, but the AI tools are built to have a somewhat hypnotic effect, overriding a person's own ideas due to a constant incoming stream of convincing suggestions. AI could also be compared to a teacher constantly looking over your shoulder and correcting and improving your attempts, eventually stripping you of your own identity and authenticity^[8].

A study done by the New Yorker^[8] showed that the homogenising effect is not limited to a specific person or group, but would in effect be a global occurrence that people's ideas, writing, designing, etc. will become homogenised for all AI users all over the world. AI is a technology of averages where the LLM's are trained to spot patterns and most favoured answers and solutions. The answers produced by AI converge to consensus and is littered with clichés.

Another consideration is the generation of content by many companies for "thought leadership" purposes, where "content generation" services are used. These services use AI tools to write articles and blogs on current issues, generating the content from existing sources. These are then distributed as new content, affirming the information already in existence without adding any new insights, but forming an endless loop of the same information regurgitated.

Depending on an AI tool to generate ideas and solve problems, will lead to loss of diverse and independent thinking, out-of-the-box ideas and unique and fresh perspectives. If AI is used by a group of people, there will eventually not be any divergent opinions in the room. There will be no need for terms such as “positive friction”, “benevolent maverick” or “robust debate”.

Percy gets a gold star

AI is here to stay... but at this point AI sounds like a risky idea. Maybe we should just be critical in the way AI is applied. Percy is not good at maths. Generating ideas, writing speeches and solving problems are not Percy's strengths either. Neither is keeping a secret. Percy is however very good and very fast at finding relevant information, checking facts and scenario planning, to name a few. Coincidentally all the things Stranger Sally is really bad at.

Bridging the gap

A common occurrence in companies is a disconnect / information asymmetry between Management and the Board. Management faces the issues, the problems, the struggles and the hype of daily decision-making. They experience the essence of the business and use and generate a myriad of information from one board meeting to the next. Board members then read through a board pack, listen to a few presentations, sometimes get a third party advisor and then confidently have to steer the ship. But it could be said that there is many a slip twixt cup and lip. Sometimes there is just too much information for management to share and sometimes they choose incorrectly what to share and what not. There is even sometimes a disconnect between the information the Board explicitly requests from management and the information the board receives. While this problem has multiple reasons, the main issue is the imperfect information flow between management and the board^[9].

Having an AI board tool can help to bridge this information gap. Much more of the information used and produced by management can be readily available to the board in a reviewed, synthesised as well as analysed format. Board members will have access to more data in a more digestible format.

From management's perspective, they will essentially have a real-time board member by their side, who can help them prepare for meetings in identifying issues and anticipating board members' questions. Management suddenly becomes better prepared for interactions with the board and more transparent in their operations and decision making. Board members feel more confident in taking responsibility for decisions they make and trust between the board and management increases as information is easily shared and readily available.

More real-time capabilities

What could be very frustrating is coming up with a good idea during a board meeting and then having to wait until the next meeting to have enough time to do proper research on the idea. With an AI board tool at hand, more in-house information can be accessed, facts can be checked, and further research can be done during the meeting. The AI board tool can even replace the need for a third-party advisor, since board members will essentially

have an instant advisor in the meeting. Having these capabilities readily available will lead to quicker decision making and more pro-active boards.

For some boards, group-think or the asking of awkward questions could be a potential hurdle. To combat negative group dynamics and social awkwardness, an AI board tool could be used as a “rock-thrower” to ask the tougher questions^[10], thereby ensuring the Board looks at an issue from all angles.

AI has taken scenario planning, a technique that has been around for a while, to the next level. Scenarios take time to develop: to fully construct a scenario could take a few days or even weeks and when fully developed, they are quite static and difficult to update when sudden changes in the landscape occur. There is a limit to the number of scenarios that can be processed and the number of variables that can be taken into consideration. AI changes all of this as it can handle far more information and variables. Scenario planning goes from static to dynamic; instead of having to replace an aged scenario or spend a considerable amount of time to update it, AI can quickly update the current scenario to keep up with changes in the landscape. AI can also run far more scenarios, making it possible to explore more options as well as more unconventional options that would previously not have been explored due to time constraints.

The Audit Committee can also benefit from having an AI board tool at hand. AI can assist in looking for trends, weaknesses in internal control and potential fraud.

From a legal perspective, board members can instantly have access to alternative legal opinions. AI could also monitor emerging legal and regulatory developments that might have an impact on the company’s operations and activities.

There are many more strengths that an AI board tool can potentially offer a board, some of which will have to be discovered as the tool is used.

It is not my fault

Percy now sits at the boardroom table. Something went terribly wrong. Everybody looks at Percy... it was all his idea after all. But Percy wears the same smile as every other day and in an extremely self-satisfied tone of voice declares, “It is not my fault”.

AI has a few issues that may never be resolved: mistakes, hallucinations, data security, etc. However, somebody still needs to take responsibility for the risks that come with using AI. If AI is sitting at the boardroom table, the board will have to take responsibility for these risks and for decisions made.

AIGN^[12] lists the following statistics:

- Only 13% of global enterprises have established formal AI governance at the C-level
- Fewer than 5% of DAX supervisory boards include any AI expertise
- In over 70% of board meetings, AI is not even a recurring agenda item
- Less than 5% of supervisory boards globally include proven AI expertise

- Over 70% of board agendas in AI-active companies do not regularly address AI governance
- Only 11% of global board members say they understand the ethical and legal risks of AI deployment

If AI is to sit at the boardroom table, these statistics will have to change, or at least, if Perfect Percy is sitting at your boardroom table, then your board has to ensure that it is beyond the statistics listed above. Because it really isn't Percy's fault.

Recommendations for the board itself

What is the role of the Board regarding its priorities in dealing with AI in the organisation? A recent McKinsey^[13] article argues that the Board has to define the company's posture toward AI adoption. This involves a clear view of how it fits into the strategic ambition, priorities and transformation agenda and the tailoring of the governance model to match such posture. Clarity about the potential impact of AI provides boards and management with a foundation for making key strategic, governance and investment decisions.

Two strategic dimensions determine a company's approach to AI, namely:

1. **Source of value.** Will AI help the company to move beyond its core business model (expand strategically), or will its value primarily come from improving the existing model (optimize internally)?
2. **Degree of adoption.** Will AI be embedded across the enterprise (holistic) or applied in targeted use cases (selective)?

In addition to this there are multiple tools already available to boards to make their work easier or more effective:

1. **Note-taker software** that allows on-the-fly analysis and summaries of meetings, including minutes and action points; for instance the production of draft minutes of committee meetings close to board meetings
2. **Board portal software** with full AI capability for all Board related issues, agendas, etc.

Perfect Percy: everyone's favourite Board Member

Which of the two panels made the best decision, the panel that appointed Stranger Sally, or the panel that appointed Perfect Percy? It turned out that Perfect Percy, being a below average homework copying student, an erroneous people-pleasing parrot, an untrustworthy and unimaginative data leaking mole who cannot work without very close supervision, also had an inflated resumé, enthusiastic references and a set of psychometric tests with highly questionable results. Did both panels appoint duds this time around?

Both appointments were successful appointments; fortunately, it was the same panel who appointed both Stranger Sally and Perfect Percy to their Board. After appointing Sally, she turned out to perform really well. Sally is described as a breath of fresh air, providing unique perspectives and out-of-the box ideas. It turned out that Sally is more of a Savvy Sally who is bad at interviews. Perfect Percy is doing a really good job at being a dynamo all-rounder: providing

on-the-spot information and analysis, running scenarios, keeping abreast of latest landscape changes just as they surface, supporting all the Board Members and turning Savvy Sally into Super Savvy Sally. Perfect Percy is everyone's favourite board member. What would we ever do without Perfect Percy?

Bibliography

1. Werner, John; 7 October 2024; Forbes; **AI Is Usually Bad At Math. Here's Why It Matters**; <https://www.forbes.com/sites/johnwerner/2024/10/07/ai-is-usually-bad-at-math-heres-what-will-happen-if-it-gets-better/>
2. Singh, Viraj; 16 July 2024; Reach Capital; **Why LLMs Are Bad at Math — and How They Can Be Better**; <https://www.reachcapital.com/2024/07/16/why-llms-are-bad-at-math-and-how-they-can-be-better/>
3. Girdhar, Abhinav; 20 June 2025; appypie agents; **AI Hallucinations: Why Smart Machines Sometimes Get It Wrong**; <https://www.appypieagents.ai/blog/what-are-ai-hallucinations#:~:text=Noisy%20or%20Inaccurate%20Data%3A%20AI,to%20%E2%80%9Cguess%E2%80%9D%20the%20answer.>
4. Chanto, Vinicio; 5 March 2025; Slidebean; **AI Has a Fatal Flaw—And Nobody Can Fix It**; https://www.youtube.com/watch?v=_IOh0S_L3C4
5. In Country staff; 19 March 2025; In Country; **Sovereign AI: meaning, advantages, and challenges**; <https://incountry.com/blog/sovereign-ai-meaning-advantages-and-challenges/>
6. In Country staff; 27 March 2025; In Country; **AI data residency regulations and challenges**; <https://incountry.com/blog/ai-data-residency-regulations-and-challenges/>
7. Eubanks, Ben; 5 February 2025; Unleash; **Is AI causing a decline in cognitive and creative skills?**; <https://www.unleash.ai/artificial-intelligence/is-ai-causing-a-decline-in-cognitive-and-creative-skills/>
8. Chayka, Kyle; 25 June 2025; The New Yorker; **A.I. Is Homogenizing Our Thoughts**; <https://www.newyorker.com/culture/infinite-scroll/ai-is-homogenizing-our-thoughts>
9. Larcker, David et al.; 8 April 2025; Harvard Law School Forum on Corporate Governance; **The Artificially Intelligent Boardroom**; <https://corpgov.law.harvard.edu/2025/04/08/the-artificially-intelligent-boardroom/>
10. Rajgopal, Shivaram; 19 September 2025; Forbes; **AI In the Boardroom**; <https://www.forbes.com/sites/shivaramrajgopal/2025/09/19/ai-in-the-boardroom/>
11. Scott, Sydney; 4 September 2025; workday; **How Generative AI Is Reinventing Scenario Planning**; <https://blog.workday.com/en-us/how-generative-ai-is-reinventing-scenario-planning.html>

12. Upmann, Patrick; 1 April 2025; AIGN; **AI Changes Everything. But Who Takes Responsibility?**; <https://aign.global/ai-governance-consulting/aign-global/ai-changes-everything-but-who-takes-responsibility/>
13. McKinsey Technology - **The AI reckoning: How boards can evolve**
December 4, 2025; https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/the-ai-reckoning-how-boards-can-evolve?stcr=BEC99EF56EC8494CA58588D24EDEB808&cid=mgp_opr-eml-alt-dna_mt-mgp-glb--&hlkid=93af2160e2db48bfa3565b306d4b54c4&hctky=10363679&hdpid=25c51be9-100a-42fd-8d20-d7ea12d337a7