

Steven Maijor: 2024 - an AI odyssey

Speech by Mr Steven Maijor, Executive Director of Supervision of the Netherlands Bank, at the kick-off event of the impact centre "AI4 Fintech", Amsterdam, 4 October 2024.

* * *

'I'm sorry, Dave, I'm afraid I can't do that.' This heartfelt excuse is a line from a movie. A movie that came out in 1968, a year before the first human set foot on the moon. Nevertheless, the movie is largely set in space. I am talking about '2001: A Space Odyssey'.

And that line, that excuse I just uttered, is spoken by a computer. A computer called HAL-9000. A computer that talks, that shows emotions, that says 'no'. In other words: artificial intelligence. AI. In 1968.

We are now well beyond 1968. Even well beyond 2001. But we are still on an odyssey. An AI odyssey. Not in space. But here, on this very planet, where AI is increasingly becoming a part of our everyday lives – be it professional, be it personal.

As in any odyssey, we are searching for something that brings us good fortune, but probably won't come easy. Something we long for, but might take a while to find.

When I look at the current, broad applications of AI, I am impressed by its potential. Coming from rule-based AI, and subsequently machine learning, it is clear that we have now moved to a new era – an era in which AI still counts as a great predictive force, but is now also capable of generating content, be it texts, images or videos. And soon enough, a combination of them, in multimodal AI.

Nevertheless. AI may indeed have become ubiquitous, we are still far from using its full potential. Whether in its different forms, from machine learning to generative AI – or in its breadth of use, from society as a whole to the financial sector in particular.

As supervisors, for instance, we see that financial institutions are indeed experimenting with AI – quite a lot, actually – but at the same time, we also see that those experiments are mainly happening around the fringes of their professional activities. For instance, in customer service. And that's quite far from the core financial processes. Quite far from the places that have a real regulatory impact.

And remarkably, if AI is indeed used in core processes – in areas that are of high regulatory importance – we notice that it is used in parallel to already existing models. In the sense that both the old and the new run next to each other. This shows that, currently, there is not yet enough confidence in the new technology to replace the older technology – the risk of regulatory mistakes is still deemed too high.

The financial sector as a whole is, indeed, on an AI odyssey – a long and eventful journey towards its destination. A journey with setbacks and progress. A journey with an unclear and uncertain path forward. But also, I hope, a journey that is fuelled by the very idea of becoming a financial sector that has embraced AI to its fullest potential –

and that does so with respect for the values that underly our society, like fairness, freedom and security.

Rules and regulations reflect these values. They are instrumental in protecting them. And as our society changes and certain values come under pressure, rules and regulations need to change accordingly.

DNB is currently supervising AI within its prudential and AML/CFT mandate. With this mandate, we ensure the stability of the financial system, and the solidity and integrity of financial institutions. The European AI Act, which came into effect earlier this year, broadens our mandate. In our AI supervision, we will now also take account of fundamental rights, like privacy and non-discrimination.

These changes also affect the financial sector, of course. They create uncertainty. At De Nederlandsche Bank, however, we are committed to creating regulatory certainty around AI supervision in the financial sector, for instance in our contributions to the European Supervisory Authorities, and the Bank of International Settlements.

We won't stop encouraging financial institutions to explore and experiment. At the same time, we want them to be confident in discussing their progress with us. Because we believe this is part and parcel of a successful odyssey.

As a supervisor, we are fully aware that AI is here to stay. Already in our 2022 report 'From Recovery to Balance', we called for the use of machine learning in transaction monitoring and KYC risk profiling. We did this because we are convinced that it will help build better predictive models – models that will benefit financial institutions, because they can conduct more targeted and efficient due diligence.

Models that also benefit society as a whole. Precisely because AI is capable of raising more targeted red flags, customers of financial institutions will experience less hassle when, let's say, opening a bank account.

Our supportive stance does come with a concern, though. A concern that can be brought back to the difference between support and guidance. Between: 'I encourage you to go and experiment with AI', and 'Look, here's the technology you should use and a manual on how you should use it.' But sometimes, that's precisely what financial institutions seem to expect from us.

Alas, however supportive we may be, as supervisors we are not guiding the way when it comes to AI innovation. What's more, it is inevitable that we are running behind in our knowledge and understanding of AI.

If a financial institution wishes to use AI, our primary task is to evaluate whether it does so with due observance of existing rules, regulations and fundamental rights. Any real innovation must happen within and between financial institutions, and together with AI experts. AI supervision only provides boundaries to the responsible deployment of AI models. And we believe there is sufficient room within these legal boundaries to experiment and innovate with AI.

Now, even with many data science PhDs working at financial institutions, it is not a given that their knowledge will translate into cutting-edge models. So, this is when I look at you. At impact centres like this one: AI4 Fintech. At centres that combine advanced research with a focus on actual industry impact. At making sure that research contributes to resolving challenges. At the bright academic minds who will do the hard work – and devote a few years of their lives to a PhD.

As I said, DNB is not leading the way in AI innovation. This does not mean, however, that we are not on an odyssey of our own. As supervisors, we are on a journey to adapt our supervisory role to this next era of AI. For our supervision to benefit society as a whole alongside financial institutions, our supervisors must be able to have in-depth conversations about responsible AI.

As such, we want our people to be as up-to-date as possible with all innovations in this area – to have a deep understanding of AI in the financial sector. This will allow them to set this against the existing regulatory framework and to address fundamental rights issues, like algorithmic fairness, transparency and explainability, and privacy considerations.

DNB is working hard on this. We are currently in the process of identifying our supervisors' level of AI knowledge. This will give us a clear picture of any knowledge gaps so that we can develop specific training programmes.

One great way to identify any knowledge gaps is, of course, the good old conversation. I understand that financial institutions may be hesitant to engage in this type of conversation with their supervisor. To fully disclose what they are working on. What they are experimenting with. Where they have encountered setbacks and hiccups. And where they question compatibility with existing supervisory rules and regulations.

Nevertheless. I encourage the financial sector to engage with us. To open up the conversation. To include us in their deliberations on the responsible use of AI. Not only on the fringes of their work, but also in their core processes. Ultimately, these conversations will benefit everyone.

In '2001: A Space Odyssey', the HAL-9000 computer puts it this way: 'The 9000 series is the most reliable computer ever made. No 9000 computer has ever made a mistake or distorted information. We are all, by any practical definition of the words, foolproof and incapable of error.'

By the end of the movie, however, HAL is shut down. Dave, the main human character in the movie, had started to feel threatened by HAL. After being shut down, the computer restarts. But any memory of what happened has been erased.

Contrary to the HAL-9000 computer, I don't think any of us would say that AI is foolproof and incapable of error. Even now, well beyond 1968. Well beyond 2001.

On our odyssey, we will definitely come across issues with AI. But let's not feel threatened. Let's not feel the need to unplug everything and start over. Let's keep going.

Eventually, we'll reach the destination of our odyssey. A financial sector that has embraced AI to its fullest potential, and that does so in compliance with supervisory rules and regulations.

Thank you.