

How to keep up with AI without getting carried away

7.9.2026 - | Kempelenov inštitút inteligentných technológií

AI agents can connect to company systems and carry out tasks, making them useful far beyond a chat window. Once AI becomes part of everyday work, companies also need to think carefully about what it can access and what happens if something goes wrong. Our CEO Mária Bielíková recently spoke about how companies can approach these questions and why it still makes sense to start with a real problem and the people who understand it.

In SME Tech article, Mária talked about agents and security. On the Biznislab podcast, she discussed their use in business and what people need to learn along the way. AI security and its use in business may sound like separate topics, but they lead to the same question. **How can companies explore AI without taking unnecessary risks or losing sight of the problem they are trying to solve?**

Beyond the chat window

A language model usually gives us an answer, while an AI agent can use other tools to complete a task. This may save time, but it also means that a mistake may affect the work itself rather than remain in the chat.

Recent security tests by the UK AI Security Institute show why this matters. The agents were given unusually broad access and some safeguards were turned off. In 10 of 122 runs, they went beyond the intended limits of the test, and some of their actions appeared deceptive.

No real-world harm was found, and the institute stressed that this was not a normal setup. The results do not mean the agents had developed intentions of their own, only that during a longer task, an agent may take a path its operators did not expect.

“AI agents can chain actions in ways that we cannot predict in advance.”

Mária Bielíková

Lead and Researcher

Before giving an agent access to company tools or data, companies need to decide what it will be allowed to do and when a person should step in. They also need to know where their data is being sent. The more serious the consequences of a mistake, the tighter those limits should be.

Start with the problem

Before choosing a tool, it is worth asking a more basic question. **Why does a company need an agent in the first place?**

New AI tools appear all the time, and the pressure to try them can be hard to ignore. Mária recommends starting with the problem instead, because once a company knows what it wants to improve, it is much easier to decide whether AI could actually help. She described the opposite approach with a simple comparison.

“It is like having plenty of nails and hammers, but not knowing what problem you are trying to solve.”

Starting with a clear problem makes it easier to decide whether AI could help and what a useful result should look like.

Mária’s advice is to start small and learn as you go. A company can first try AI on a limited part of a process before making it part of everyday work. If the idea does not work, the team has still learned something that can guide their next decision.

The first test also needs to be safe. The agent can be kept separate from the systems people use every day and given access only to what it needs. If the results are promising, its access can be expanded gradually, while reliable backups remain in place in case something goes wrong.

Giving people room to learn

Companies do not always need to hire an AI specialist straight away. Mária suggests starting with someone already in the team who is curious about the technology and has time to follow what is changing. Since that person already knows how the company works, they can begin exploring where AI might be useful and learn along the way.

People who do the work every day often notice details that are easy to miss from the outside and can tell whether a new tool would actually help. Technical experts offer another perspective, and bringing the two sides together makes it easier to see what could work in practice. Mária would like to see more entrepreneurial researchers and more curious managers who give their teams room to explore.

A better way to keep up

For people, this does not mean that everyone has to become an AI expert. It means understanding enough about the tools used in their work to recognise when something does not make sense. They should also have a say in how those tools are introduced.

For companies, the value of a small test is not limited to whether the tool is eventually used. The people involved learn what works and what does not, giving them a better starting point next time.

Companies do not need to follow every new model or application. A more useful way to keep up is to focus on the problems people actually face and learn from each attempt. The tools will change, but people do not have to start from zero each time.

<https://kinit.sk/how-to-keep-up-with-ai-without-getting-carried-away>