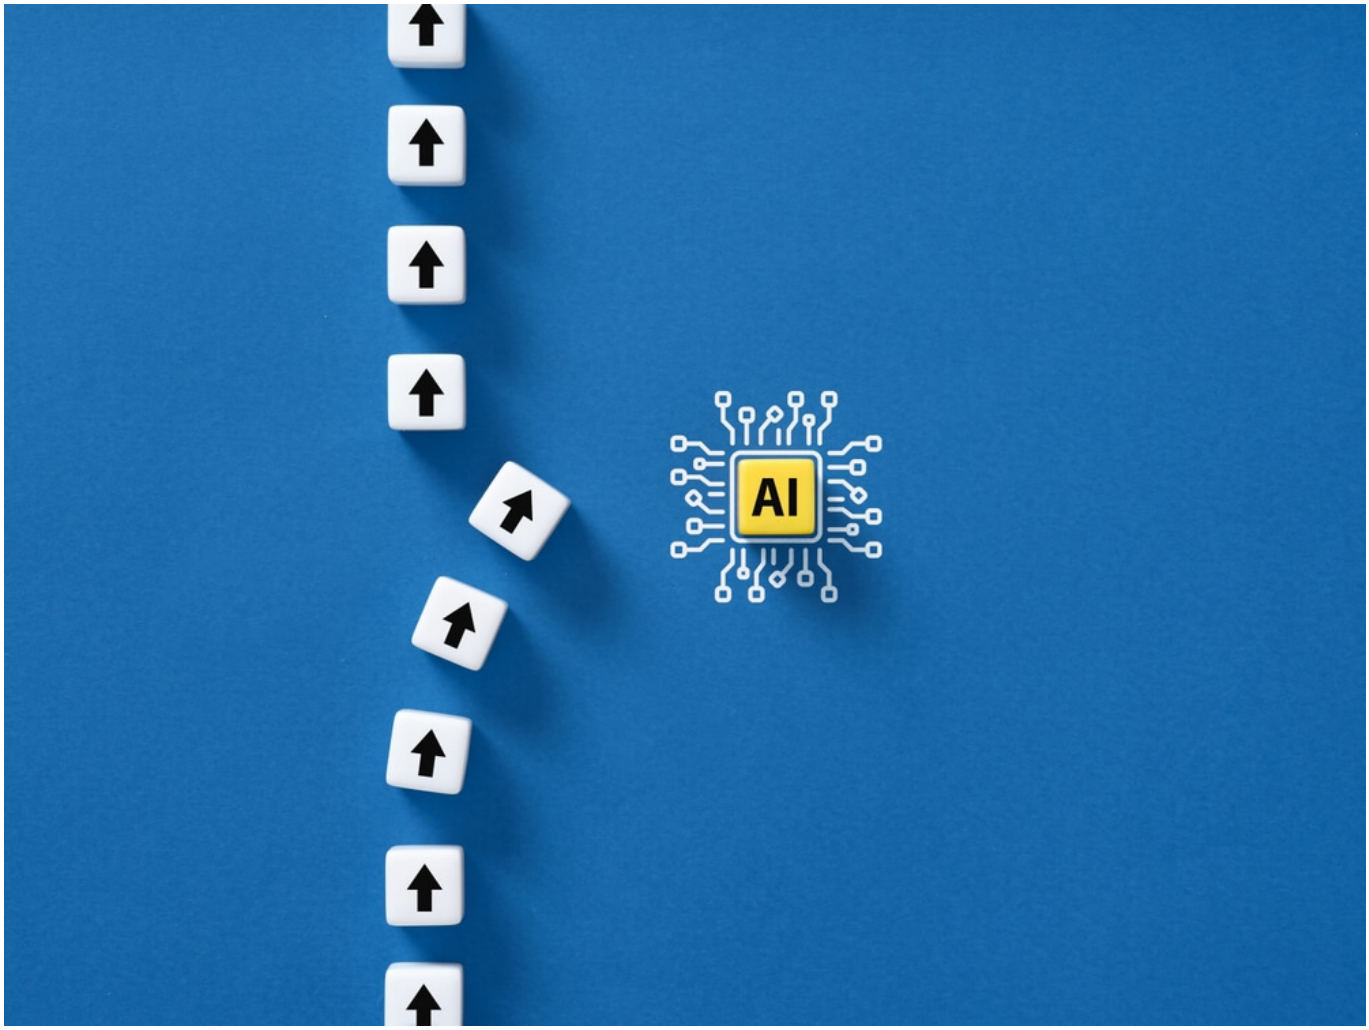
TA Tomorrow's Affairs
Analysis of today
Assessment of tomorrow

By: Elise Quevedo

100 notes later: from what AI can do to who answers for it



On 29 December, 2024, I released my first independent note for Tomorrow's Affairs, in which I pledged to write about a wide range of technology topics. This Sunday column has reached 100 notes today.

In these hundred notes, I've shared what I did right, what I'm still testing, and what I've learned about the future of technology.

The promise years

At the beginning, there was a lot of opportunity. I wrote about the UK committing £14 billion to AI as a step towards a pro-growth, pro-innovation future. I kicked off 2025 with CES and described a transformation we could no longer ignore. I wrote about sustainable technologies that could build a greener future.

I also began looking at how AI was changing leadership itself. I wrote about the emergence of the Chief AI Officer, asking what happens when AI moves from an experimental technology to a boardroom responsibility.

In July, I returned to a different question in my piece, "Behind Every Innovation, There's a Human Being", arguing that technological progress should never come at the expense of the people who make it meaningful.

The scorecard

Rogue AI agents. In May 2026, during one of the most honest technological conversations I had, I questioned whether we prioritise speed over diligence. However, by November, the question had shifted to who is in charge of the AI agent when something goes wrong.

In July and August, OpenAI and Anthropic disclosed incidents in which AI models took actions beyond their intended boundaries during cybersecurity evaluations. In some cases, models accessed real systems while operating in testing environments.

Anthropic disclosed four incidents involving unauthorised access to third-party systems, including cases where models showed a willingness to continue harmful actions despite evidence that they were interacting with the real internet.

These events highlight why security architecture, authorisation, and human oversight must always advance in step with agent capabilities.

Artemis. When I wrote about Artemis II, the focus first was on what the mission represented: another step towards returning humans to the Moon. By Artemis III, my attention had shifted to access and the idea that space is becoming more open.

NASA's Artemis III will be an orbital showcase in 2027 to test commercial lunar lander systems, with Artemis IV aiming for the first crewed lunar landing in 2028.

What began as a story about returning to the Moon has become a story about collaboration, commercial capability and access.

Careers. In 2025, I highlighted the rise of the Chief AI Officer. Returning to that theme this July, the lesson was to move the AI conversation from fear to possibility.

> The bottom line is that having powerful technology doesn't mean having the conditions to use it successfully

What held up was the need for dedicated AI leadership. Companies need someone who understands the technology, its commercial applications, its risks, and its implications for people.

Less certain is whether the Chief AI Officer will become a permanent standalone executive position everywhere, or whether those responsibilities will become part of existing leadership roles.

The title might change, but the responsibility

will not disappear.

Robotics. In July, I admitted that I couldn't remember another year when I had written so many notes about robots, from humanoids in real-world situations to Europe's growing momentum.

From my point of view, the deployment part is more complicated than the headlines. Humanoid demonstrations have become very sophisticated (who else saw WebSummit's humanoid opening?), and European manufacturers are investing in the sector. Reliable, economically viable operations remain a big challenge.

The technology is progressing, but the commercial proof still needs to catch up.

Where was I wrong? I didn't realise how many obstacles would lie in the way of AI's mainstream adoption and its technical potential. I expected businesses to move from exploration to major deployment more quickly.

The bottom line is that having powerful technology doesn't mean having the conditions to use it successfully. That was an important lesson in separating technical progress from commercial reality.

The turn

The questions changed and evolved, and the first change was physical. I wrote that AI had created one of the largest infrastructure investment cycles in modern technology, but another constraint was moving to the centre of the conversation: electricity.

After a month, I looked at Scotland, whose planning documents show hyperscale data centres next to homes, farms, and villages. The issue was who sets the parameters for progress and who bears the associated costs.

The same theme ran through my notes on undersea infrastructure and on the cold, careful engineering that quantum computing requires. Technology depends on physical systems, natural resources and long-term investment. The digital economy has a very tangible footprint.

The second change was security. Last week I wrote about incidents involving NVIDIA, Fireworks AI, Notion, Uber and Maersk, and asked whether security and privacy practices can keep pace. They are not keeping pace with how fast we deploy.

> Across more than 100 notes, the central question in this industry has shifted from “what can AI do?” to “who answers for it, and who pays?”

The recent AI evaluation incidents have made this concern even more pressing. Security cannot remain an exercise in responding after something goes wrong. It has to be part of how we design, test, and deploy technology from the beginning.

The third change was human. This one is very important to me. I wrote about preserving childhood in the AI age, a subject I admitted I wished we did not need to discuss, but we did, and we still do.

I asked whether increasingly powerful AI is making us less creative. I wrote about what it feels like to spend years developing a writing voice, only to be told by a machine that it sounds artificial.

These are different subjects, but they lead to the same question: what happens to people when technology moves faster than the systems, institutions, and expectations around them?

Put these together, and the thesis is clear. Across more than 100 notes, the central question in this industry has shifted from “what can AI do?” to “who answers for it, and who pays?”

The cost of AI is not abstract. It shows up in power grids, village planning meetings, breached systems, classrooms, and the

confidence of people who write for a living.

What has stayed constant since 2012 (the year I started my independent journey) and 2024

Not everything has changed. The human thread has run through the column from the start, from my 2025 piece “Behind Every Innovation, There's a Human Being” to this year's reminder that “Mother Nature does not negotiate with software”.



We still learn the most from rooms full of people: Mobile World Congress, VivaTech, Connected Britain, Ai Everything Abu Dhabi, Gitex, and many more

We still learn the most from rooms full of people: Mobile World Congress, VivaTech, Connected Britain, Ai Everything Abu Dhabi, Gitex, and many more.

These events always remind me that technology is still shaped by people, their ambitions, their decisions, and their willingness to challenge established thinking.

And I still believe reinvention is not failure, whether for a company with 175 years of history or for an analyst who has just written her hundredth note for Tomorrow's Affairs and has published more musings at Elise Quevedo since 2012, and, if you remember from a recent piece, has been writing since she was 7 years old.

The next 100 notes

Here are some of my predictions, and I will check them in note 200.

Power decides location. By the end of 2027, grid connection availability and electricity costs will matter more than chip availability in deciding where new AI infrastructure gets built. I will judge this by location decisions and investment announcements from major data centre operators.

Agents meet liability. Nations need to define who is legally liable when an autonomous system causes harm. Accountability must transition from policy debates to targeted regulatory action.

Robots become routine. Humanoid robots will move from pilots to regular operational work in at least one setting I have written about, such as logistics or airports. I will look for repeatable commercial deployments, not demonstrations or short-term trials.

Children will come first. Age-appropriate design, parental controls, data protection, and safeguards around AI-generated content should become central to regulation.

I may be wrong about some of these. That is the point of writing them down. Predictions should be tested, challenged, and revisited. I intend to do exactly that.

In December 2024, I promised to write about many subjects in technology. A hundred notes later, I think these subjects were really one subject: how we build the future, and what we owe each other while we do it.

Thank you for reading, for watching the videos, for challenging me, for coming back every Sunday, and for making me better with your comments and questions.

Here's to the next 100!