



GenAI in practice: from scepticism to strategic applications

19 May 2025, 10:50

Nick Boucart

Reflecting on the first experiments with generative AI

Generative AI has undergone remarkable development in recent years. The introduction of powerful language models such as ChatGPT has led to a wave of experiments, insights and applications. Yet the central question remains: how can we use this technology in a way that is relevant, robust and valuable?

At Sirris, we guide businesses in the process of making technology tangible. **Our approach is pragmatic:** we start from real needs and challenges, not from hype. In an earlier article, [‘Four months of experimenting with GenAI’](#), we shared insights from a series of internal projects. One of the people driving this exploratory programme is Nick Boucart, senior technology advisor at Sirris.

As an expert in software, SaaS models and digital strategies, he combined his critical eye and technical expertise in a series of concrete experiments. What follows is his **personal account of how he learnt about and came to apply generative AI in practice.**

GenAI: from scepticism to targeted experiments

It must have been sometime in the autumn of 2023. The mainstream introduction of ChatGPT was almost a year in the past. At that time, I personally had no serious plans to start using it. My first impression was of an interesting technology, but also lots of hype, unreliable output and **limited added value** for my daily work. I had tried prompting, but the results had rarely been convincing, so I tended to prefer a good old Google search.

However, my interest started to grow, partly due to conversations with colleagues who had actually gained valuable results from ChatGPT through intensive practice. That got me thinking. **Maybe I** an actual prototype



In October, I decided to get started systematically. I developed **a first GenAI prototype**: a chatbot that could generate answers based on the entire content of the Sirris website. It worked simply but effectively.

- I used Python code to scrape our website,
- stored the data as embeddings in a vector database,
- and then linked it to a generative language model via LangChain.

This produced an RAG (retrieval augmented generation) architecture that formulated answers based on our own content.

The **results were mixed**. The chatbot performed well on questions for which the answer was explicitly given on our website. In other cases, though, for example in response to a question about what you can make with salmon and figs, the system came up with a recipe. Obviously, we don't have a culinary section on our site. This was a clear illustration of the limitations of language models without content control.

Four insights from practice

After this first exercise, **I continued to experiment**. I adjusted prompts, re-indexed the content and tested variants in the configuration. Gradually, I gained a better understanding of what works and what doesn't. Some key lessons:

1. Ensure that the technology starts from a clear need

Using a chatbot without having a clear problem to solve is rarely worthwhile. Generative AI only has an impact when it meets a specific need. Otherwise, it's just a pointless technological gimmick.

2. Generative AI is a Swiss army knife, not a magic wand

Generative AI isn't a magic wand, but a flexible tool. The challenge lies in defining the application correctly. Jobs-to-be-done are a useful framework in this context.

3. Long prompts don't make short work

Language models predict the most likely next word, based on enormous quantities of data. This makes them powerful, but also potentially error-prone. Anyone who wants reliable output would do well to integrate their own data using techniques such as RAG (retrieval augmented generation), tool calling and, recently, MCP (model context protocol).

4. Users are unpredictable. And that's putting it mildly

User interactions aren't always predictable. Some users will ask innocent but unusual questions, while others will deliberately try to test the system. It's important to think about what is and isn't permissible, and to adjust the system accordingly.

So what's happening now?

Since then, I've **developed dozens of prototypes** and spoken to numerous companies about using GenAI. The core challenges from my first experiment remain just as relevant today, but

enthusiasm among a wider public has really taken off since the launch of ChatGPT. For companies, this presents both opportunities and challenges. How do you deal with this technology as a software company? How do you ensure that your applications remain robust, reliable and secure, even with GenAI components?

At Sirris, we use the motto '**made real**', and it's equally applicable to generative AI. We help companies to look beyond the buzzwords and **implement technology in a way that's realistic and sustainable**.

It's for this reason that the [Coock+ project LISA](#) was launched on 1 May, an initiative of Sirris and KULEuven DistriNet. In this project we're **looking at how to use generative AI applications responsibly in software development**.

Are you interested in joining the LISA user group?

[Contact our expert Nick Boucart](#)

Like to know more?

You can also read the article '[Four months of experimenting with GenAI](#)', in which we summarise the wider lessons from our internal projects at Sirris.

Are you interested in getting started with GenAI yourself?

[Have a look at more practical examples and get in touch](#)

Authors



Nick Boucart