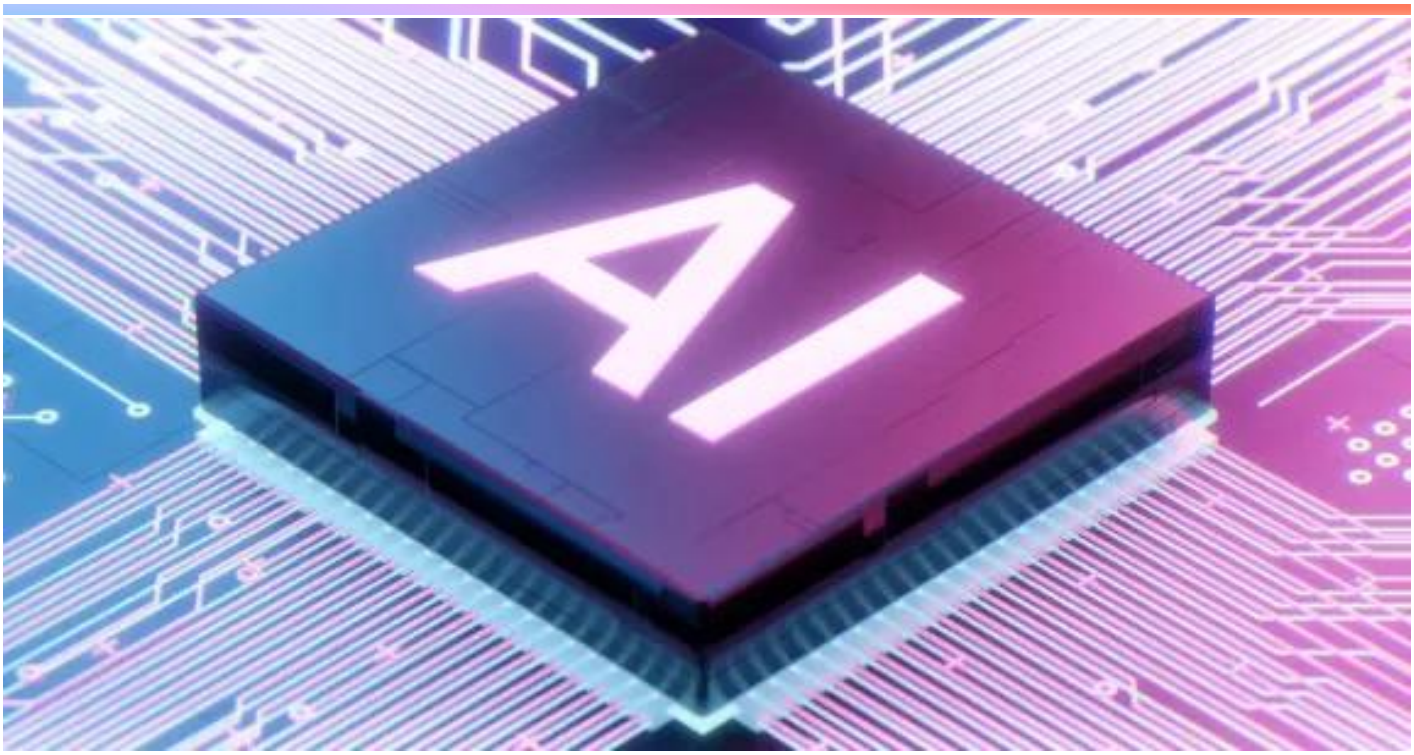


Generative AI makes forecasting accessible

19 March 2026, 10:54

Mihail Mihaylov
Wim Codenie

This article is part of a three-part series on Generative AI for time series forecasting. In the introductory article, we explored why forecasting is difficult in industrial environments. In this second article, we show how Generative AI makes forecasting far more accessible. Continue to the next article to see how these models can be turned into working tools in minutes.

Time series forecasting becomes practical for every industrial company

Manufacturing companies sit on a mountain of data. Sensors monitor machines, energy systems track consumption and production software records operations every minute. The data keeps growing. Yet reliable forecasts remain rare. Building forecasting models has traditionally required specialised expertise and long development cycles. Generative AI is now lowering this barrier and making forecasting far more accessible for industrial companies.

The forecasting paradox in industry

Many industrial teams would benefit from predicting what happens next. They want to anticipate machine maintenance, estimate future energy demand and plan production capacity more accurately. These tasks sound simple. Their execution is not.

Building forecasting models has traditionally required specialised knowledge and time. Engineers must select algorithms, design input features, train models and validate results before a model becomes reliable.

For many companies, the effort is difficult to justify. This creates a paradox. The organisations that would gain the most value from forecasting are often the least equipped to build it. Generative AI is beginning to change that situation.

What makes traditional forecasting so expensive?

To understand the impact of Generative AI, it helps to look at how forecasting has **traditionally** worked. Developing a forecasting model involves **several technical steps**. Each stage requires expertise, experimentation and iteration.

Engineers must first select a suitable algorithm. Options range from classical statistical models to modern machine learning approaches such as gradient boosting.

The next step is feature engineering. Raw industrial data rarely works directly as model input. Engineers must transform measurements into meaningful signals that a model can interpret. After that comes model training and parameter tuning. Models are evaluated, adjusted and tested again until performance becomes acceptable.

This process often **takes weeks or months**. It also requires experienced machine learning engineers who understand both modelling techniques and the industrial context. As a result, forecasting remains concentrated in organisations with **dedicated data science teams**. Smaller companies often face difficult choices. They rely on basic regression models, outsource forecasting tasks or operate without predictive insight.

A new approach: time series foundation models

At Sirris, we explored a different approach based on time series foundation models. These models represent a new category of Generative AI designed specifically for **forecasting tasks**.

The concept is straightforward. A **pre-trained model** is downloaded from an open-source repository. Historical time series data is provided as input. The model then **produces forecasts immediately**. No model training is required. Feature engineering becomes largely unnecessary and extensive tuning is often optional.

This approach resembles how large language models generate text. A language model can produce coherent sentences without being trained on a specific topic because it has learned general linguistic patterns during large-scale pre-training.

Time series foundation models follow the same principle. During pre-training they learn common patterns in time series data, including trends, seasonal behaviour and recurring cycles. Once trained, the model can **analyse completely new datasets** right out of the box.

The practical implication is significant. Forecasting tasks that previously required specialists and weeks of development can now be **implemented in hours**.



© Unsplash

Testing the approach in practice

At Sirris, we evaluated this approach by reproducing benchmark experiments from scientific literature. We compared the results with well-tuned traditional forecasting models.

The foundation models performed on par with industry-grade forecasters and in several cases achieved stronger results, despite requiring no training. These findings show that Generative AI can deliver reliable forecasting accuracy while significantly reducing the effort required to forecast data.

Industry feedback

Industry feedback confirmed the practical value of this approach. A large European energy company explained that they manage hundreds of internal forecasting projects. Many teams rely on simple regression models because more advanced solutions are too expensive to develop repeatedly. The company sees time series foundation models as a potential **default baseline** because they are **easier to deploy** and deliver **stronger predictions**. Another important advantage is deployment.

The models run fully on-premise, even on a standard laptop, which means sensitive industrial data never leaves the organisation.

Understanding the limitations

Generative AI forecasting models also have limitations. The most important challenge concerns **explainability**. Traditional machine learning models can often reveal which features influenced a prediction. Foundation models behave more like black boxes. They provide accurate forecasts but offer limited insight into their reasoning.

For many operational applications this is acceptable. However, some decisions require deeper explanation. In those cases companies may combine Generative AI forecasts with more interpretable machine learning models.

What this means for industrial companies

The main impact of Generative AI is economic rather than purely technical. Forecasting models no longer require large investments in expertise and development time. When the cost of deploying a forecasting baseline becomes low, many new use cases become viable. Industrial companies can forecast dozens or even hundreds of data streams across their operations.

Patterns appear earlier and new optimization opportunities emerge. This shift also enables new digital services such as demand forecasting, price forecasting and predictive maintenance.

From insight to action

Many companies see the potential of Generative AI forecasting but do not know where to start. Sirris supports organisations through workshops and hands-on sessions that explain the technology and its practical applications. For companies ready to go further, Sirris also offers project-based support to implement Generative AI forecasting solutions in real industrial environments.

Explore the full series

This article is part of a three-part series on Generative AI for time series forecasting:

Intro: Your operations generate data. Are you forecasting with it?

Article 1: Generative AI makes forecasting accessible

Article 2: From prompt to prototype in 7 minutes

Discover it yourself in 3 hours

Join the Sirris masterclass on GenAI timeseries forecasting. No coding required. No prior ML experience needed.

You will leave with a practical step-by-step approach that you can apply directly to your own data. Almost free for Flemish companies thanks to VLAIO subsidy.

Register for the masterclass

Authors



Mihail Mihaylov



Wim Codenie