Europe, February 2025

Time Series Foundation Models (TFMs)

Introduction deck

LUNAR

Executive summary

- Foundation models (FMs) or base models, are task-agnostic, pre-trained, large-scale neural networks that can be adapted to numerous downstream tasks.
- Within FMs, large language models (LLMs) trained on terabytes of internet-based textual datasets like CommonCrawl have become popular in recent years due to product examples like ChatGPT. As LLMs have attracted hundreds of billions of EUR, and FMs like GPT-3/GPT-4 have led to a myriad of derivative models tailored to specific downstream tasks (e.g., specific assistants). These recent investments in LLMs have also shed light on the unit economics in an AI-world: CAPEX required to train SOTA (state of the art) models increases~2,5x/annum[1], compute required to train SOTA models is growing at ~4,6x per year[1], rapid commoditization of models occurs due to open-source competition, different geographic regions tend to build their FMs[2], and rapid declines in token costs for existing models occur at -86%/annum[3].
- New categories of FMs will occur in the coming years, typically with alternative use cases. The main examples include geospatial foundation models (GFMs), time series foundation models (TFMs), and material foundation models (MFMs). An alternative way to segment foundation models is to focus on the data modality being used, the main categories using that logic are language foundation models, vision foundation models and multimodal foundation models – the latter category using multimodal input data and typically allowing multimodal generalizations (e.g. text to image, image to text).
- Time Series Foundation Models (TFMs, or TSFMs) are foundation models specifically focused on time-sensitive data. TFMs are typically trained on cross-domain datasets (e.g., energy, stock, commodity, sentiment data), and allow better cross-domain generalizations as output. It is well conceivable that TFMs will also be more accurate than traditional models due to less overfitting and better long-term predictions. TFMs may have a significant impact on all markets where time-sensitive predictions are important, such as energy markets, financial markets, climate (risk) forecasting, supply chain predictions, and other areas (e.g., political sentiment predictions).
- Lastly, we dive into the AI economics making these use cases viable and highlight three main trends relevant to understanding AI economics: i) the increase in training costs for new state-of-the-art (SOTA) models, ii) the increase in the total amount of available models, and iii) the decrease in token costs over time. First, investment for developing foundation models is significant, with the first ~1bln EUR models in sight – while at the same time cheaper variants keep popping up (e.g., DeepSeek, Llama). Second, in the last four years open-source model repositories have grown very rapidly in with a 70X increase in the number of models available. Language-based models are still the dominant category, but this can change when more high-quality and affordable satellite data becomes available at scale. Finally, given the intense competition from open-source models, API token costs tend to fall rapidly, as was shown in recent years for GTP-3 and GPT-4 where token costs fell ~90% per year. These trends support the premise that geospatial models can have a wide set of economically viable use cases in the coming years, as the willingness to invest is present and usage costs are declining rapidly.

[1] Source: Epoch.ai; [2] ChatGPT in the US, Mistral in Europe, DeepSeek in China – with largely similar capabilities, [3] source: deeplearning.ai.

Foundation models are multi-purpose models...

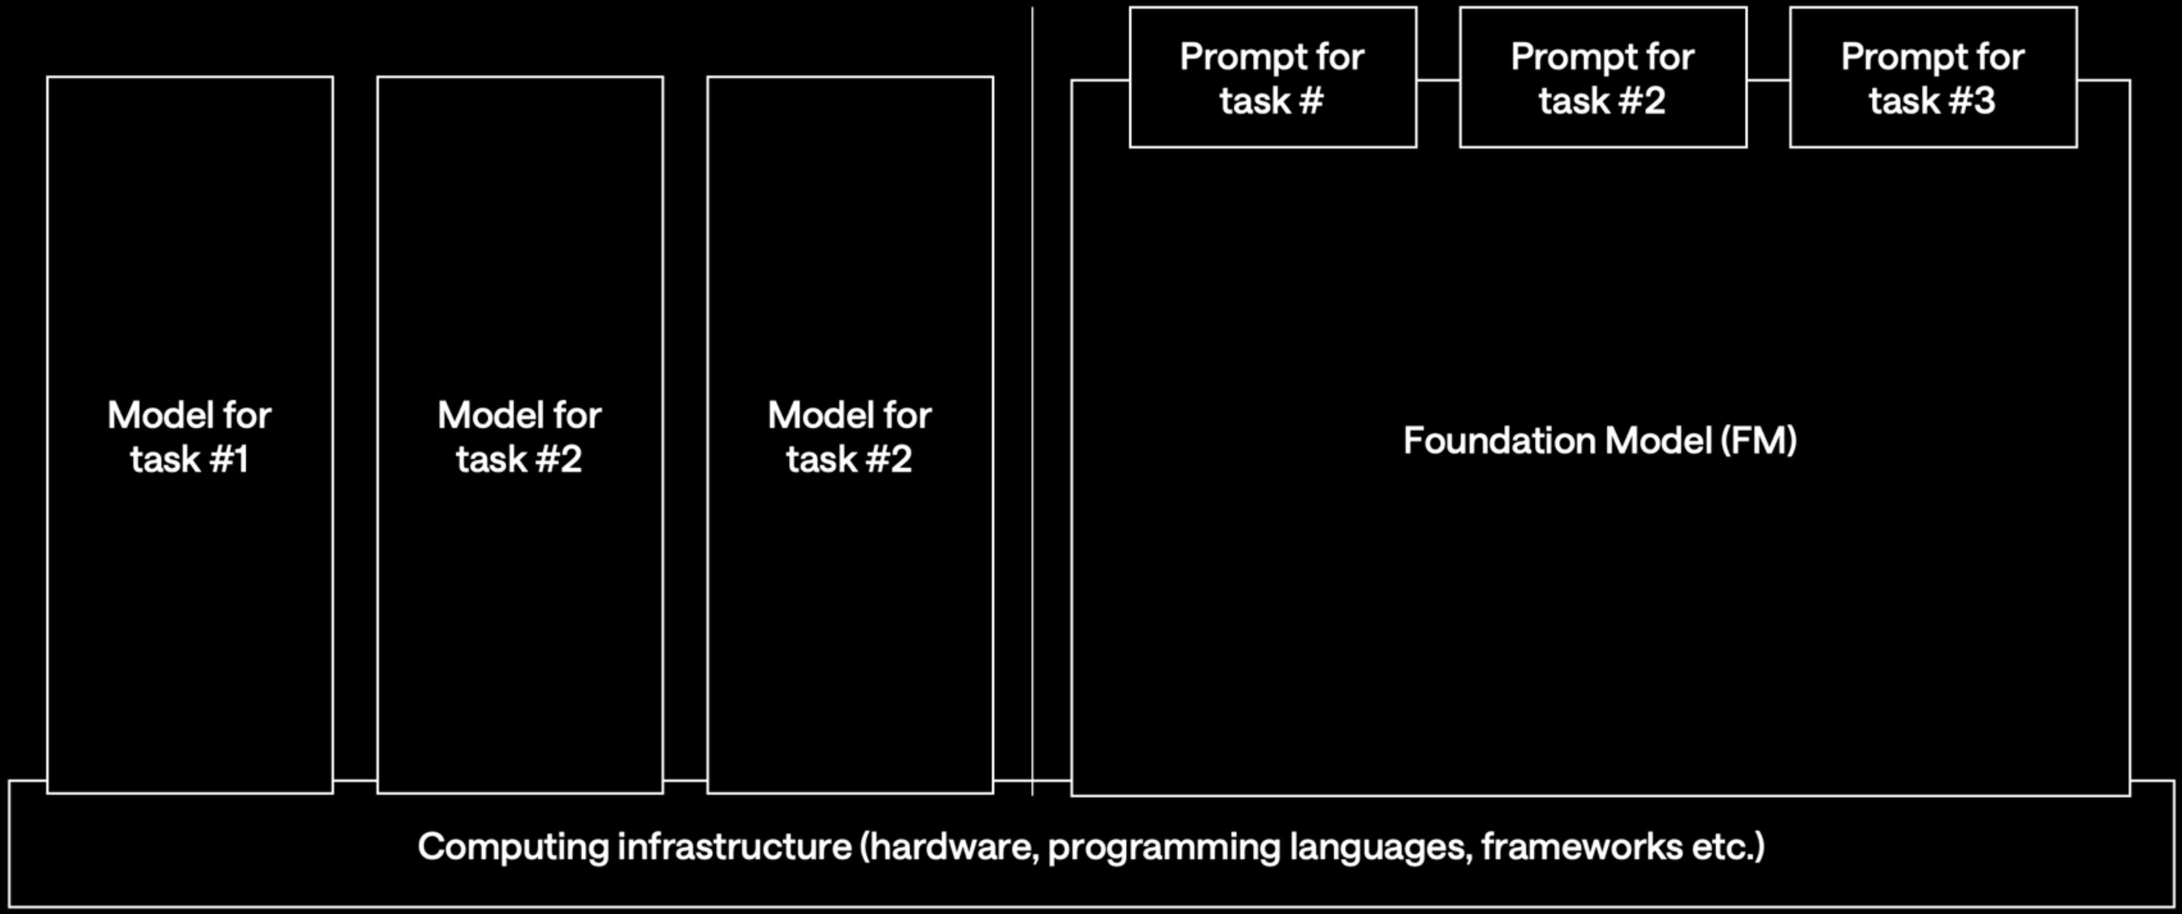

Source: University of Amsterdam (2024): Foundation models are platform models.

...where TFMs are trained on cross-domain temporal data

TFMs are trained on cross-domain temporal datasets[1] ...

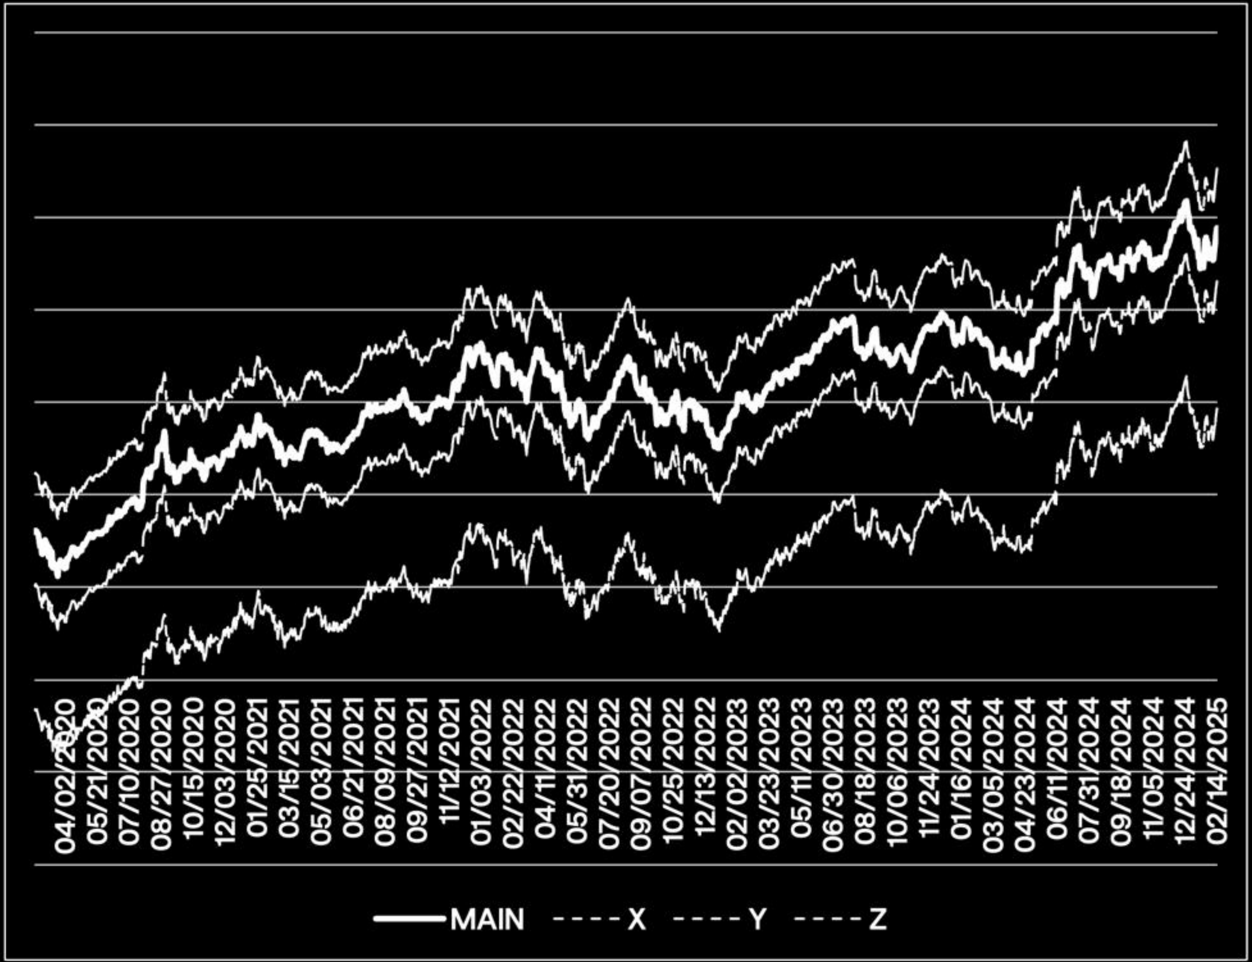

...allowing a new paradigm of time series modeling.

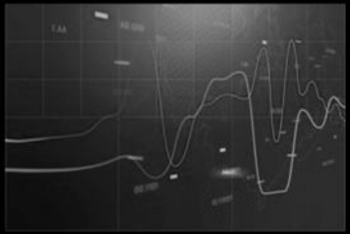

Real-time forecasting

Cross-domain generalization[2]

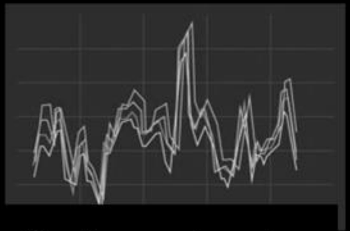

Continuous learning

Multimodal capacity

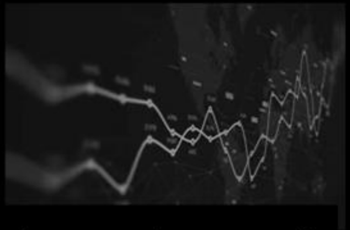

Improved accuracy[3]

Fast deployment

[1] Dummy dataset based on AAPL stock prices; [2] TFMs are typically trained on datasets from different industries (e.g., finance, energy, climate) and should thus be able to generalize better cross-domains; [3] To be assessed per use case, but the i) diversity of input data, ii) transformer infrastructure (vs ARIMA, RNNs) should theoretically allow for less overfitting and better performance on long time horizons.

LUNAR

TFMs may well be a significant industry in the coming years

Data modality: Language data | Vision data | Multimodal data

Model type: Geospatial (GFMs) | Time Series (TFMs) | Urban FMs (UFMs) | Material (MFMs)

Sector: Energy | Finance | Climate | Supply Chain | Other

Sector	Energy	Finance	Climate	Supply Chain	Other
Value estimate	**1-5B EUR/year[1].**	**5-10B EUR/year[2].**	**5-25B EUR/year[3].**	**1-10B EUR/year.**	**1-50B EUR/year.**
	Energy prices, energy demand, wind speeds, grid load forecasting, renewable energy generation predictions etc.	Stock prices, commodity prices, bond yields, credit risk, GDP, inflation, unemployment, central bank rates, market sentiment, portfolio optimisation.	Daily max/min temperature prediction, GHG concentrations, rainfall, droughts, surface temperatures, sea temperatures, pollution levels, long-term forecasting etc.	Demand forecasting, inventory optimisation, consumer sentiment, inflation prediction, global economic signals, geopolitical risks, transport disruptions etc.	Geopolitical risk assessments, social sentiment analysis, sanctions impact analysis, economic stability forecasting, fraud detection etc.

[1] Smart grids (management) to reach ~50B USD by 2030, the value shown assumes a 2-10% increase in efficiency using better models; [2] Algorithmic trading market to reach ~30B USD by 2030, a large part of which will like be done with AI; [3] Wide range of estimates – depending on the severity of climate disasters. Note that the California wildfires alone cost ~250B USD.

While training costs for foundation models is generally increasing...

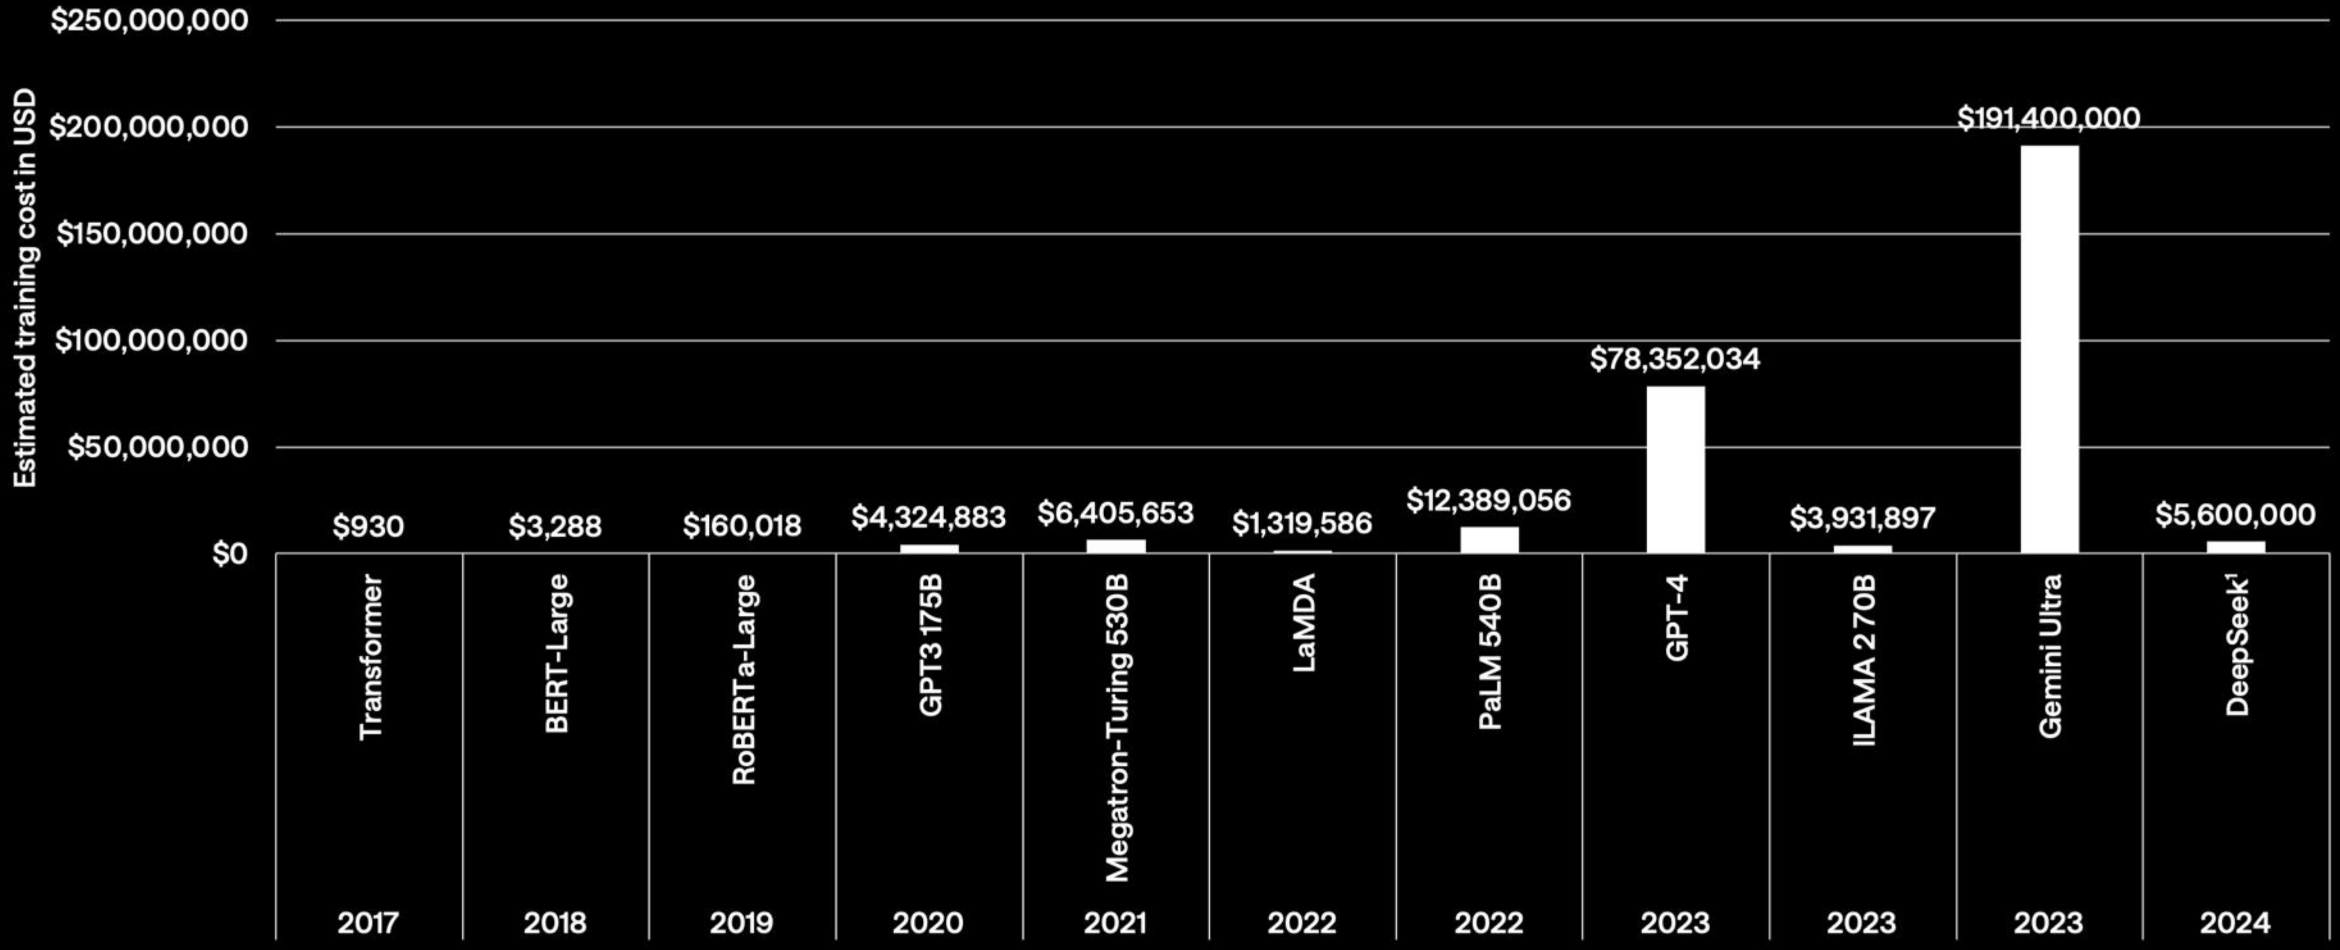

Source: epoch.ai; [1] Note the discussion around DeepSeek's true training costs, with estimates from 5,6mln USD to 1,6bln USD.

...the competition between models is fierce...

Open-source models grew ~70X in the last four years...

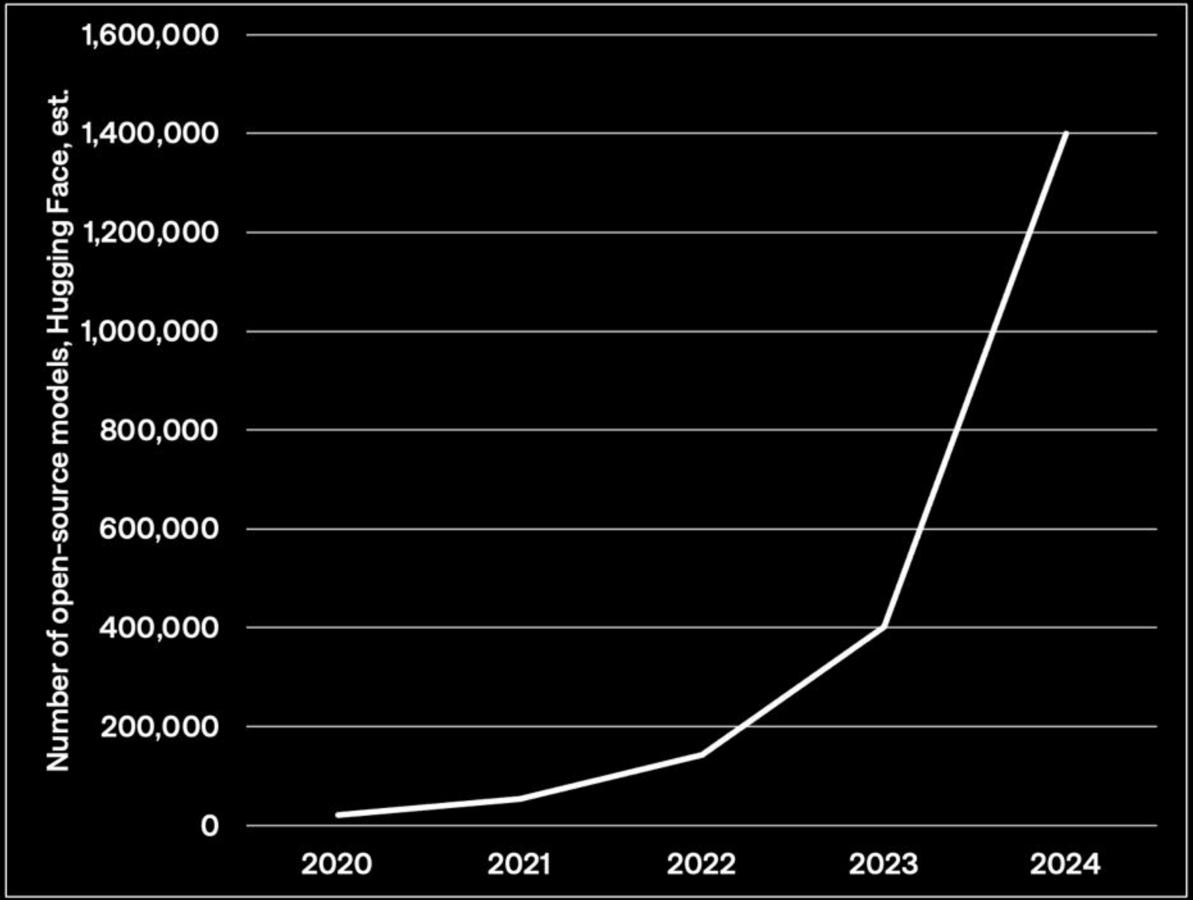

...and more players are releasing large (over 10^23 FLOP) models.

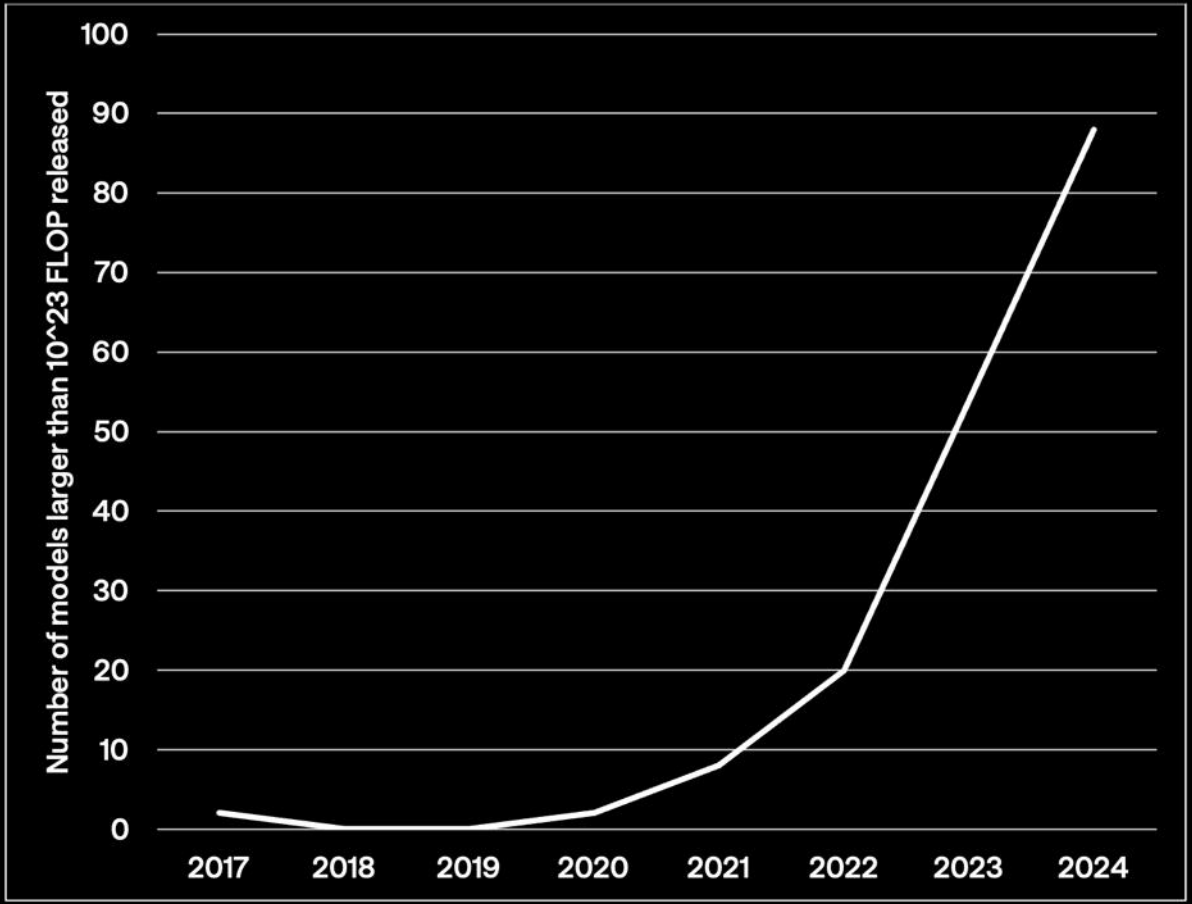

Source: Hugging Face, Epoch.ai.

...and falling token costs will make more use cases economical

GPT-3 token cost fell ~86% in cost per year...

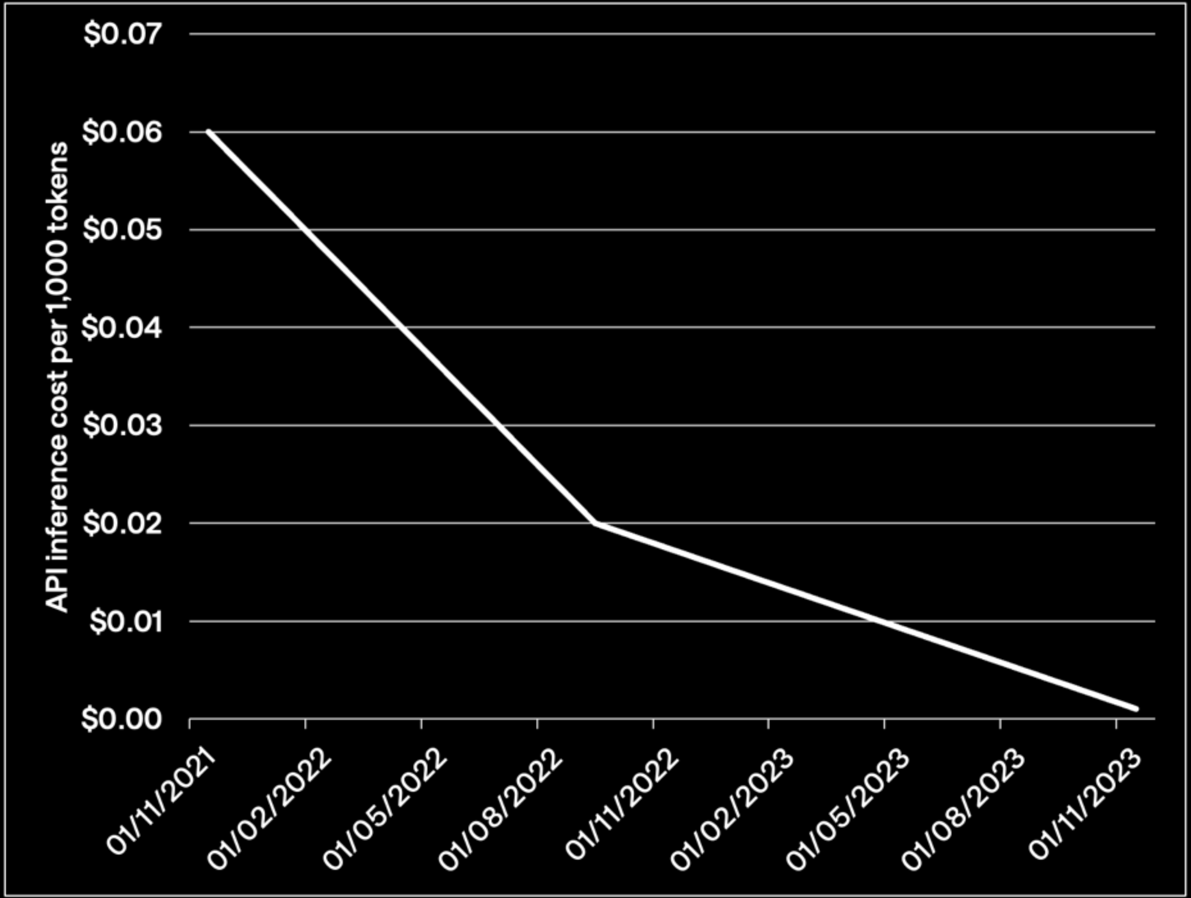

...and GPT-4 tokens showed a similar trend of falling 92% per year.

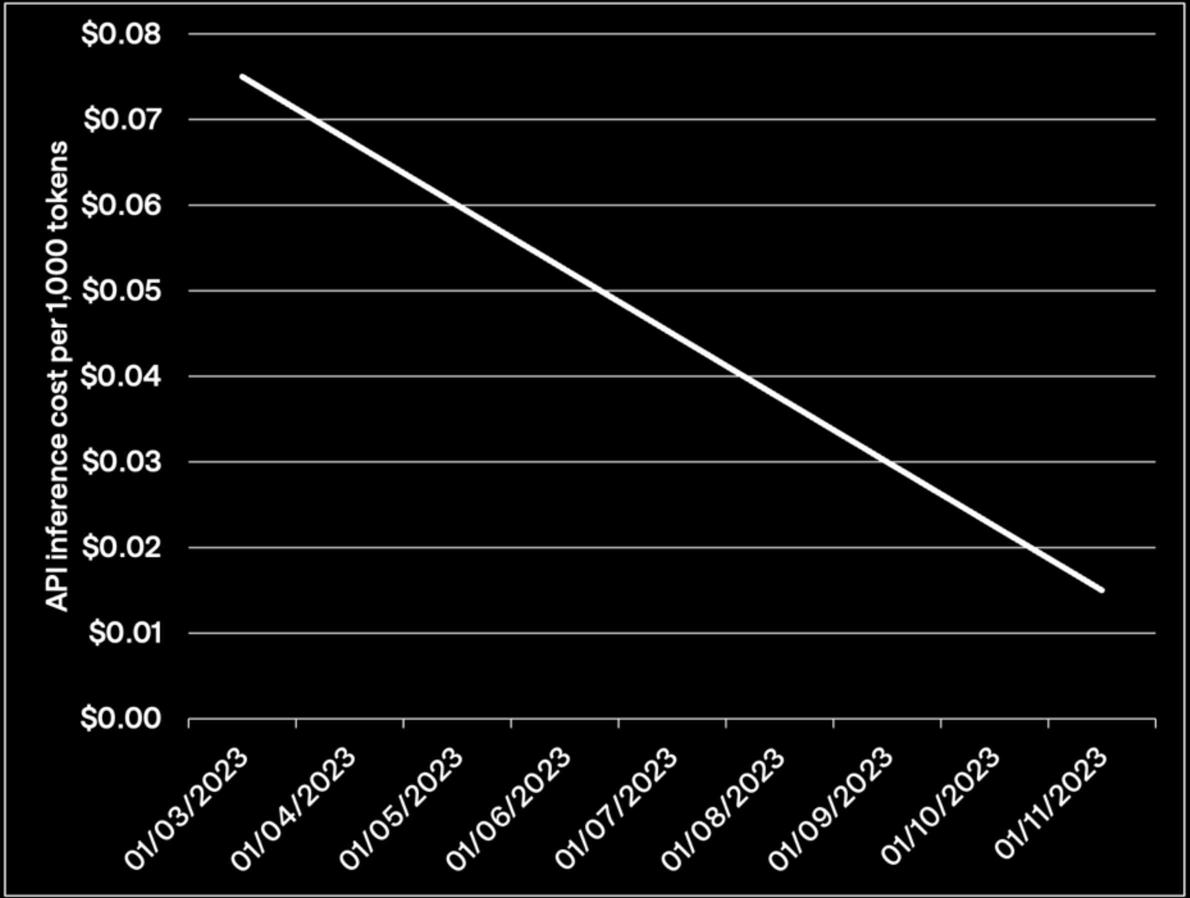

Source: Ark Invest analysis (2024); The Batch (2024).