

The Leading Real-Time AI Software for Machines and Processes

Large Language Models in Manufacturing



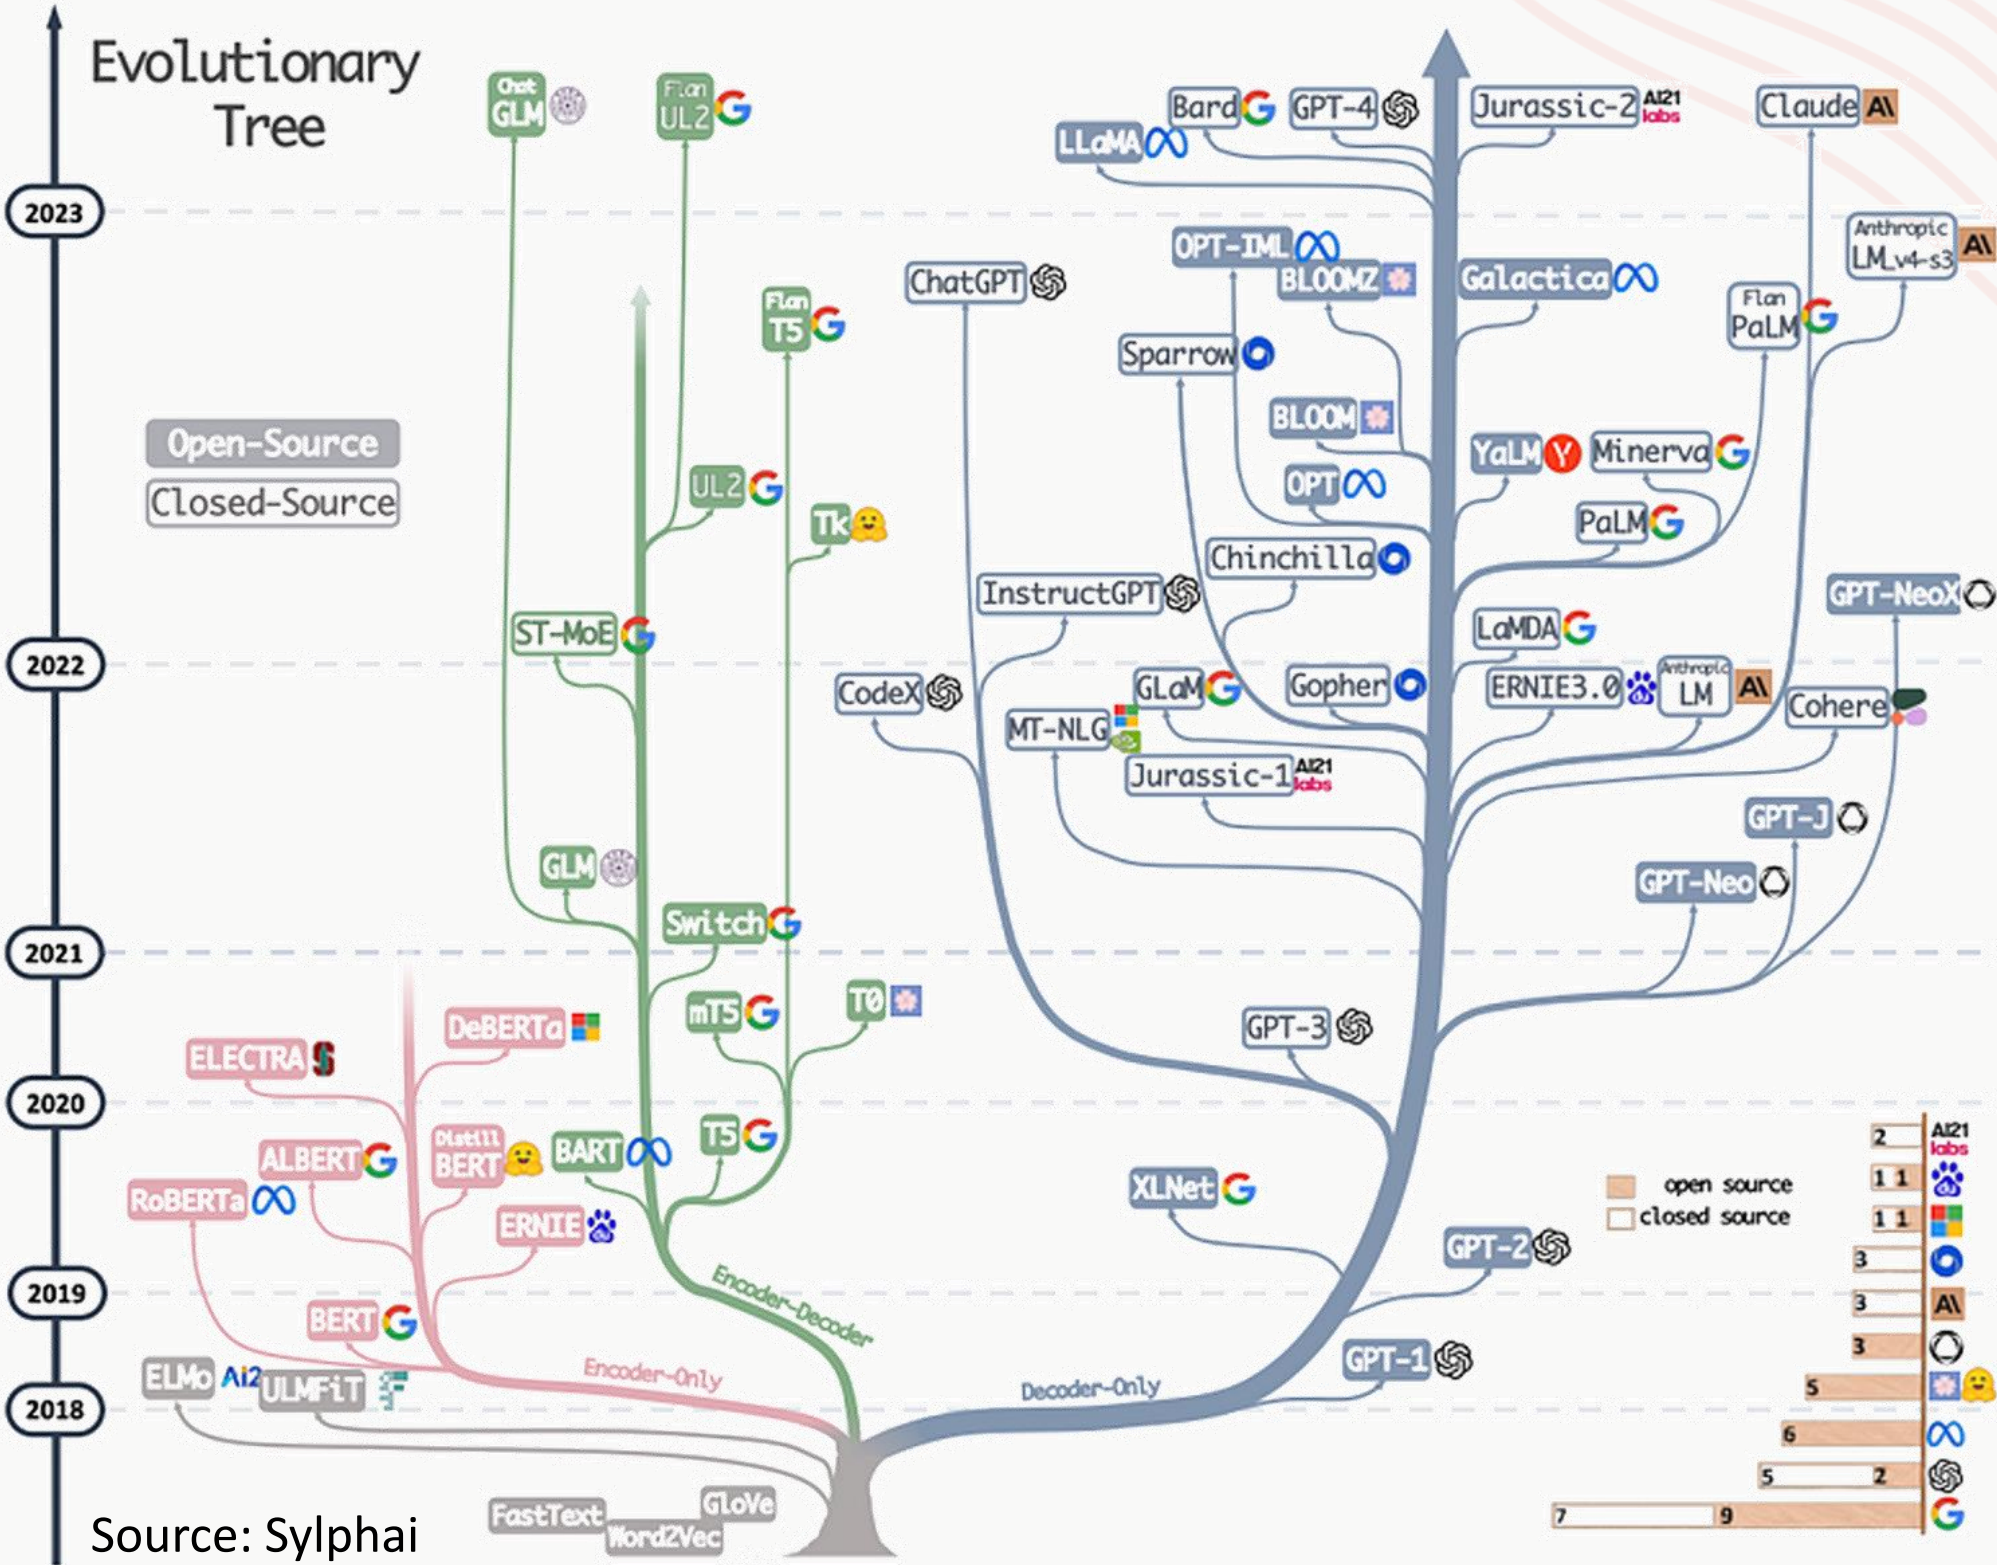
Productie Proces Automatisering

25 januari 2024 | Van der Valk Hotel, Vianen



Large Language Models?

A large language model (LLM) is a type of artificial intelligence (AI) algorithm that uses deep learning techniques and massively large data sets to understand, summarize, generate and predict new content.



slido



Describe what you think LLMs can do for you

ⓘ Start presenting to display the poll results on this slide.



UReason and LLaMA

At UReason we experiment with the LLaMA LLM to determine how we can make use of it in our products: APM/APM Studio, the Control Valve and Pump/Motor App

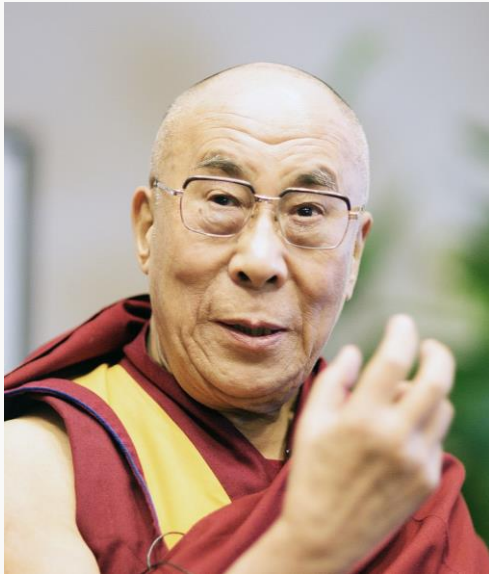
Why: Our customers have large volumes of text resources (manuals, maintenance activities, inspection reports et cetera) and time-series data from for examples control valves. It would be great if we can support our users with additional advisories/recommendations and predictions combining both!?!

Hi Precious .. Long time no see, gosh we are growing old together. But this will be the last time we see each other I'm off on pension

Hey Old Man .. I'd love to talk back to you but I can't, my owners have never really cared about all the data I generate just about my functioning when they noticed it was off 😞



Prompt: What is LLaMA?



Lama (Tibetan): is a title for a teacher of the Dharma in Tibetan Buddhism. The name is similar to the Sanskrit term guru, meaning "heavy one", endowed with qualities the student will eventually embody. The Tibetan word "lama" means "highest principle", and less literally "highest mother" or "highest father" to show close relationship between teacher and student.

Wrong Answer

The llama (Lama glama) is a domesticated South American camelid, widely used as a meat and pack animal by Andean cultures since the pre-Columbian era. Llamas are social animals and live with others as a herd. Their wool is soft and contains only a small amount of lanolin. Llamas can learn simple tasks after a few repetitions. When using a pack, they can carry about 25 to 30% of their body weight for 8 to 13 km (5–8 miles). The name llama (in the past also spelled "lama" or "glama") was adopted by European settlers from native Peruvians.

Almost Right



Prompt: What is LLaMA?

A thing involving Llamas and AI?



Almost Right



Prompt: What is LLaMA?

LLaMA (Large Language Model Meta AI) is a family of large language models (LLMs), released by Meta AI starting in February 2023. For the first version of LLaMA, four model sizes were trained: 7, 13, 33 and 65 billion parameters..

Meta

Research

Introducing LLaMA: A foundational, 65-billion-parameter large language model

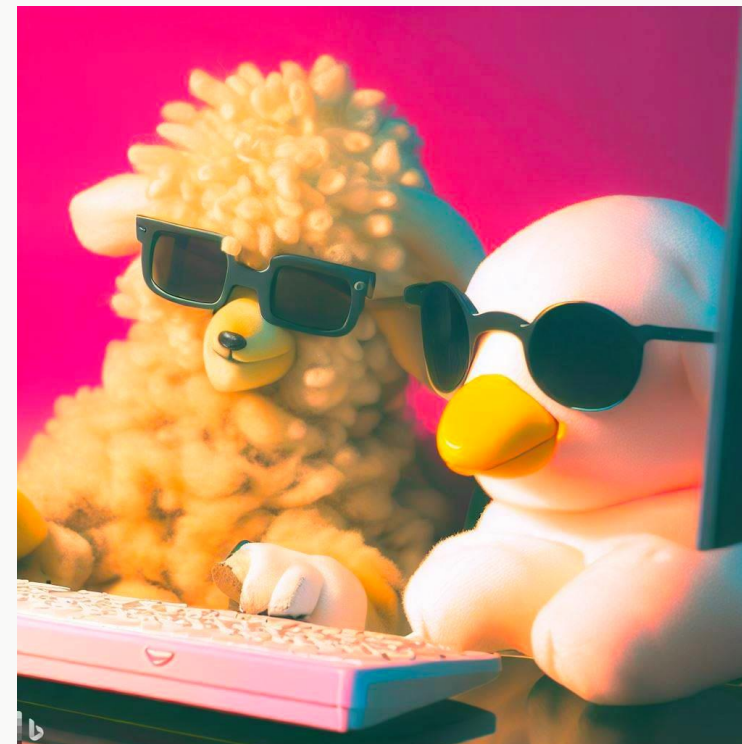
February 24, 2023

UPDATE: We just launched Llama 2 - for more information on the latest see our blog post on Llama 2.

As part of Meta's commitment to open science, today we are publicly releasing LLaMA (Large Language Model Meta AI), a state-of-the-art foundational large language model designed to help researchers advance their work in this subfield of AI. Smaller, more performant models such as LLaMA enable others in the research community who don't have access to large amounts of infrastructure to study these models, further democratizing access in this important, fast-changing field.



Right Answer



At UReason we use a fine tuned SheepDuck 70B model

We know this sounds weird but not in the LLM community ;) .. And it beats Llamadeus Mozart



Prompt with Real Data Example

Below prompt is provided to a LLaMA 70B model with time-series data:

```
prompt = """<<SYS>> You are responding to highly technical customers. </>
[INST] Remember you are an assistant for engineers. Be serious. [/INST] [INST]

Period HealthScore TotalMovement StictionIndicator TotalEnergyLoss OpTime Issues
1 41 23711 0.26 NaN 1261215.0 [...]
2 0 10392 0.55 NaN 1354597.724576 [...]
3 0 6520 62.18 264859 0.31 NaN 1482085.0 [...]
4 70 10647 0.26 175045 0.1 NaN 1007099.129695 [...]
5 0 8338 2.96 83731 0.29 NaN 586485.0 [...]
6 47 5 0.22 104 0.29 NaN 658355.44835 [...]
7 38 12785 2.78 175401 0.14 NaN 1071060.0 [...]
8 0 38688 15.69 200290 0.31 NaN 1129920.0 [...]
9 0 5516 6.35 6 0.24 NaN 418053.98751 [...]
11 43 1893 5.03 NaN 389039.900265 [...]
12 22 199 0.71 NaN 553660.0 [...]
```

Context of reply (technical)

Prevent it from starting a dialogue

Time series data for a control valve

Response we are looking for + identify trends

Above is a panda dataframe consisting of information for a specific Valve. Can you summarize the information into a concise description? Explain the different aspects of the data. Each period corresponds to a month of the year. Try to identify temporal trends.

[/INST]"""



Response

A decent model was chosen and somewhat okay response:

The valve operation data shows information spanning 12 periods (months). During this time, the Stiction indicator increases gradually in the first two months from 0.26 to 0.55, then decreases until period 8 and finally increases again towards the end. The TotalMovementOfValve is consistently within the range of 1k-39k, with a slight upward trend. OvershootPercent and UndershootPercent fluctuate between periods, with no clear performance issues, StictionIssue was consistently present in all periods such as OvershootIssue, UndershootIssue, Packing Bellows Issue, Dead Band Issue, Positioner Calibration Issue, and Stuck Issue. Some of these issues appeared more frequently in certain periods than others but were generally present throughout the entire period. The valve's HealthScore is not provided, so it's difficult to determine the overall performance trend based on this single metric. However, the TotalEnergyLoss and EnergyLoss values show an increasing trend over time, which could indicate potential deterioration in the valve's efficiency or power consumption. OpTime and StrokesTotalNumber also have a similar trend but do not necessarily reflect any specific issue. In summary, this valve seems to be experiencing performance issues related to various aspects of its operation, with some issues appearing more frequently than others. While the problems are recurrent, they may vary in intensity or frequency from period to period, suggesting a need for continued monitoring and potential maintenance interventions to ensure optimal valve functioning.

Not entirely right .. Understanding numbers

Wrong, it is part of the input

Not entirely right .. Is related to issues

(Used highlighting to indicate the correctness of the response)



Prompt with Real Data Example

Below prompt is provided to a LLaMA 70B model with forecasting templates:

```

prompt = """"<<SYS>> You are responding to highly technical customers. <</SYS>>

[INST] Remember you are an assistant [/INST] [INST]

Instruction: Describe in detail the valve health score trend over time for
The trend is represented by bins, where each bin corresponds to the monthly
drop up to 10% is considered acceptable if it recovers within the subsequent
The Valve A bins are: "91%", "89%", "90%", "87%", "90%". "89%", "92%".
Valve Conditions: The valve is operating in hot water distribution. Common mechanical issues for valves includes stuckness, positioner calibration, positioner drift, and increased deadband.

Forecasting Examples:
Valve B bins: "92%", "88%", "84%", "86%", "82%". "87%", "90%", "84%", "86%", "84%", "87%".
Valve B exhibits a healthy trend, as the health score remains closer to 100%. No issues are
continue to behave the same in the coming months.

Valve C bins: "92%", "92%", "90%", "91% - INCREASED POSITIONER DRIFT", "86%". "84%", "78% -
Valve C exhibits a healthy trend during the first four months. Though an increase in positioner drift
However, a mechanical stuckness issue arises, which causes the health score to rapidly decrease.

Now explain Valve A's behavior over time. Do not just summarize the history.
Use Valve B and C as examples, do not use analyze them. Look only at Valve A bins.

Please provide your answer in the following template:

QUERY: Write a summary of the query here.
DESCRIPTION: Write a description of the health trend in detail.
COMPARISON: Compare Valve B to Valve B and C. Explain which trend it follows more closely.
DETAILS: Provide extensive analysis and reasoning.

[/INST]""""

```

Explaining the data format and providing examples

Please provide your answer in the following template:

QUERY: Write a summary of the query here.

DESCRIPTION: Write a description of the health trend in detail.

COMPARISON: Compare Valve B to Valve B and C. Explain which trend it follows more closely.

DETAILS: Provide extensive analysis and reasoning.



Response

Overall a very decent response:

QUERY: The query is to describe the valve health score trend over time for Valve A, and explain the reasoning behind this trend based on available information.

DESCRIPTION: Valve A has a fluctuating health score trend. It starts at 91% and then gradually decreases to its lowest point of 87%, before increasing again to 90%. The health score then remains relatively stable, with slight increases and decreases from month to month.

COMPARISON: Comparing Valve A to Valves B and C, it follows the trend of Valve B more closely, since both exhibit fluctuating health scores but remain within healthy levels most of the time.

DETAILS: The observed behavior of Valve A suggests that it is experiencing normal fluctuations in its performance. While there is a decline to 87% at one point, this is still well within acceptable limits. It's important to monitor closely for any further decreases or the emergence of any mechanical issues like stuckness or positioner drift to ensure that the valve continues to operate effectively. This fluctuating but stable trend is similar to Valve B, which also experienced normal variations in its health score over time without any major concerns. It would be advisable for maintenance teams to keep an eye on Valve A's performance and address any potential issues promptly to prevent serious problems from arising, as seen in the case of Valve C.

(Used highlighting to indicate the correctness of the response)



Prompt with Support Example

Below prompt and response from a LLaMA 70B model requesting support:

```
prompt = """<<SYS>> You are responding to highly technical customers. <</SYS>>

[INST] Remember you are an assistant [/INST] [INST]

What do you know about the ideal maintenance procedure for a SAMSON 3271 actuator?

[/INST]"""
```

```
The SAMSON 3271 is a pneumatic linear actuator commonly used in industrial applications. Proper maintenance is essential to ensure its optimal performance, safety, and longevity. Here are some key points to follow when developing an ideal maintenance procedure for the SAMSON 3271 actuator:
```

1. Regular cleaning: Dirt, dust, and debris can accumulate around the actuator over time, leading to decreased performance or even damage. Clean the surrounding area of the SAMSON 3271 regularly using a soft brush or compressed air (with oil removal). Avoid touching moving parts while cleaning.
2. Lubrication: Proper lubrication is critical for maintaining the smooth operation of pneumatic actuators like the SAMSON 3271. Ensure that all moving parts are lubricated according to the manufacturer's recommendations, paying special attention to the bearing surfaces and seals. Use a high-quality, lightweight oil designed for air tools or pneumatic systems.
3. Adjustment: Periodically check the actuator's operation, adjusting the pressure regulator. Ensure that the actuator operates within the recommended pressure range to improve its efficiency, reduce wear on components, and extend its lifespan.

Part of a longer response, response is very general and not specific to the type of actuator (as the model has not been trained on this)



slido



How precise and careful should LLM models be

ⓘ Start presenting to display the poll results on this slide.



slido



What would you use an LLM for (multiple options selectable)

ⓘ Start presenting to display the poll results on this slide.



Wrap-up



LLMs provide great value to support users on technical questions

LLMs can be combined with time-series data to provide predictions (requires specific processing)

Llama is open source and up to 700 million active users per month free to use

Footprint and deployment costs of Llama is considerably smaller than ChatGPT:

- Llama: A6000 – MSRP ~ 4.000 Euro
- ChatGPT4: 10xA100 ~ MSRP ~170.000 Euro

Your data is private to where model is running either with UReason or at your site!



Q&A + Come Visit Us!



Leave your business card and receive the survey results, presentation and more insights on how you can monitor your important assets and provide your users with guidance and support using LLMs and your data!



joudmans@ureason.com
Mobile: +31 6 18 556 550

