LLMs: generating innovative and effective collaborations in AI

Prof. Damien Ernst, Prof. Gilles Loupe and Lize Pirenne

LIÈGE université

NRB AI &c Data Xperience

Introduction to AI

- Artificial Intelligence
  - This refers to machines able to learn, reason and adapt.
  - Used in multiple industries and domains such as healthcare, finance, entertainment, autonomous vehicles.

- AI research communities are organised around several sub-communities, for example:
  - Natural Language Processing (NLP)
  - Computer Vision (CV)
  - Intelligent Robotics (that benefits from progresses in NLP and CV)
  - Reinforcement learning
  - Deep Learning

The rise of LLMs in AI

- A large language model (LLM) is a computational model capable of language generation or other natural language processing tasks.

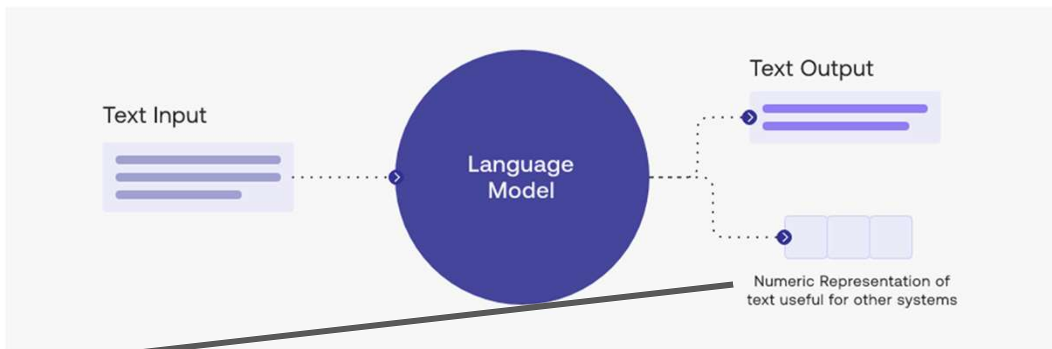

https://www.bureauworks.com/blog/what-is-large-language-models-llm

Inside the LLM, everything is represented as embeddings (vectors of fixed size). This is a very useful representation that can be shared with numerous other AIs.

- LLMs were initially introduced in the AI NLP community. Today they are being effectively applied accross the board, such as in reinforcement learning.

What is Reinforcement Learning?

- Reinforcement Learning (RL) is a powerful AI paradigm where an agent interacts with its environment to maximise rewards it accumulates.

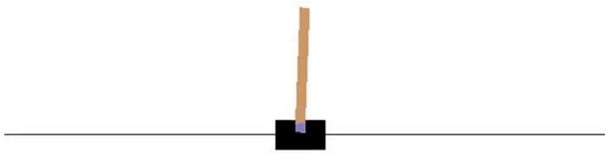

- The key problem in RL: the design of the correct reward function. For simple tasks, a simple reward function that can be hand-crafted can work well, e.g., reward = cos θ for the inverted pendulum swinging on a cart.

- What about designing the right reward function for the pen spinning problem?

Eureka: LLM for crafting the reward function in RL

- Let the LLM figure out what a good reward is by trial and error.
  - Start with a human-designed reward: Reward at time *t* = distance(pen_position(*t*), ideal_position(*t*))

Pen position → Average

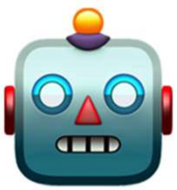

Add angular velocity

→ A little better

Add rotation reward

→ Much better!

Entire process is automatic!

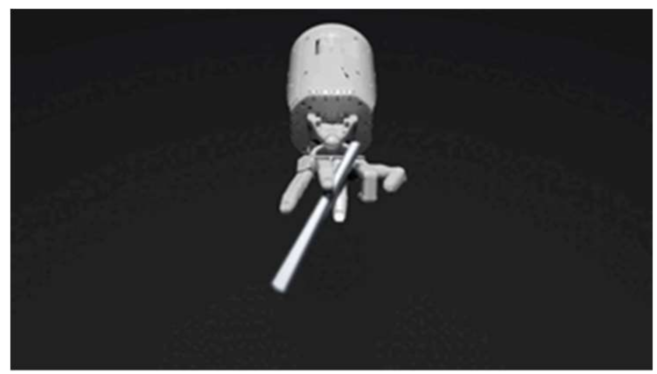

[Eureka: Human-Level Reward Design via Coding Large Language Models]

LLMs can also use external tools at their own pace

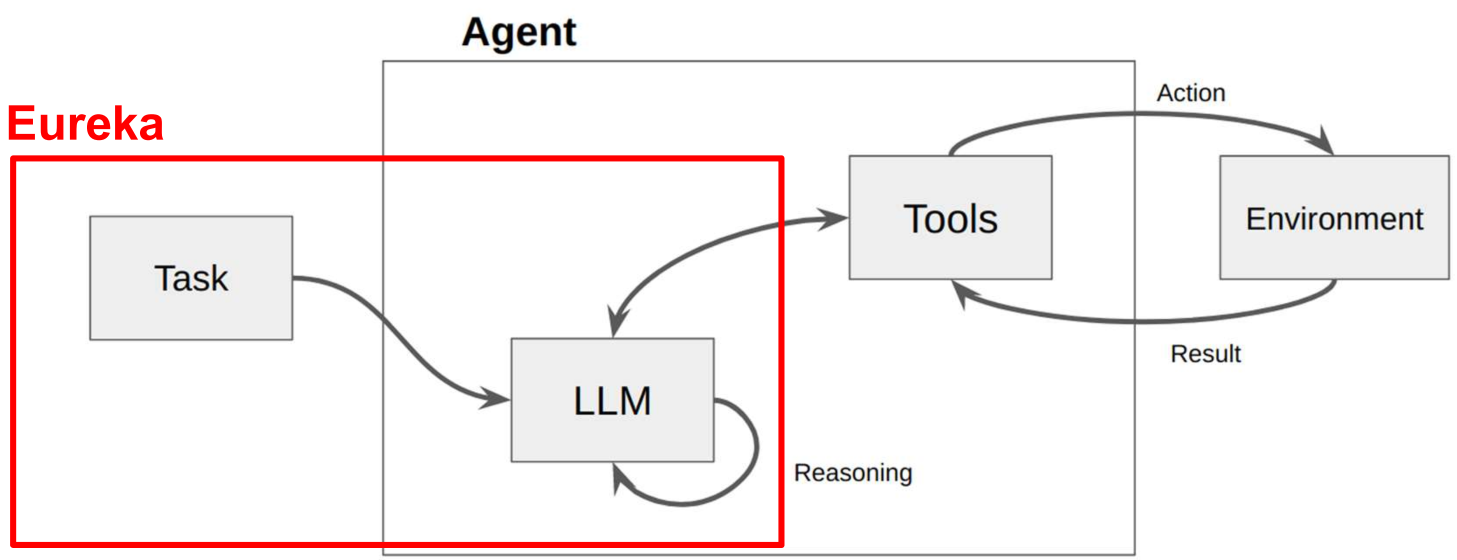

These tools include

- Databases used for memory
- Internet searches for knowledge
- Robot controls for movement

LLMs as agents directly controlling a robot from input goals, without using a reward function

- Let us use everything!
  - Task breakdown
    - High-level initial task
    - Multiple-level atomic tasks
  - Usage of tools
    - Motion (API to robots)
    - Memory (API to databases)

goblin.tools

nvidia ReMEmbR

Usefulness of embeddings for multi-modality and interacting with other AIs

- At the heart of the representation inside an LLM, there are the embeddings. You can embed everything (music, video, etc) the ‘same way’.
- The differents embeddings can be processed by the same LLM to create a powerful representation that can also be useful for other AI (e.g., diffusion models that generate images).

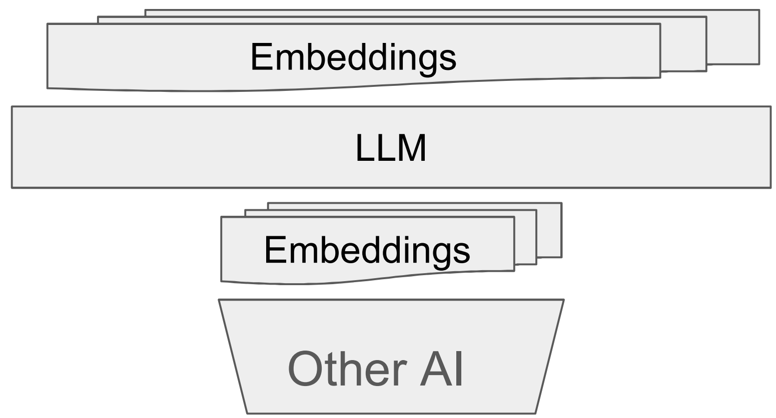

Embeddings are so powerful that LLMs can deal directly without needing to go through any text

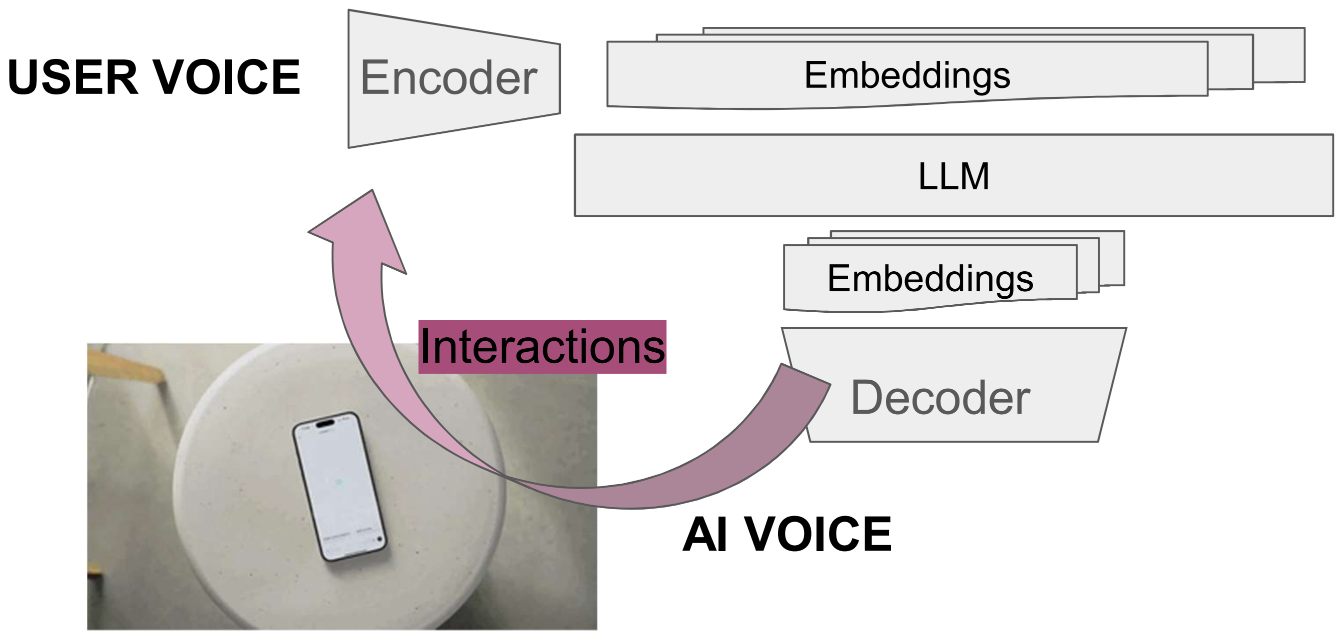

[OpenAI Advanced Voice]

Voice to Voice directly!

The architecture behind making movies from a script looks so simple

Let us enter a prompt corresponding to the script of a movie!

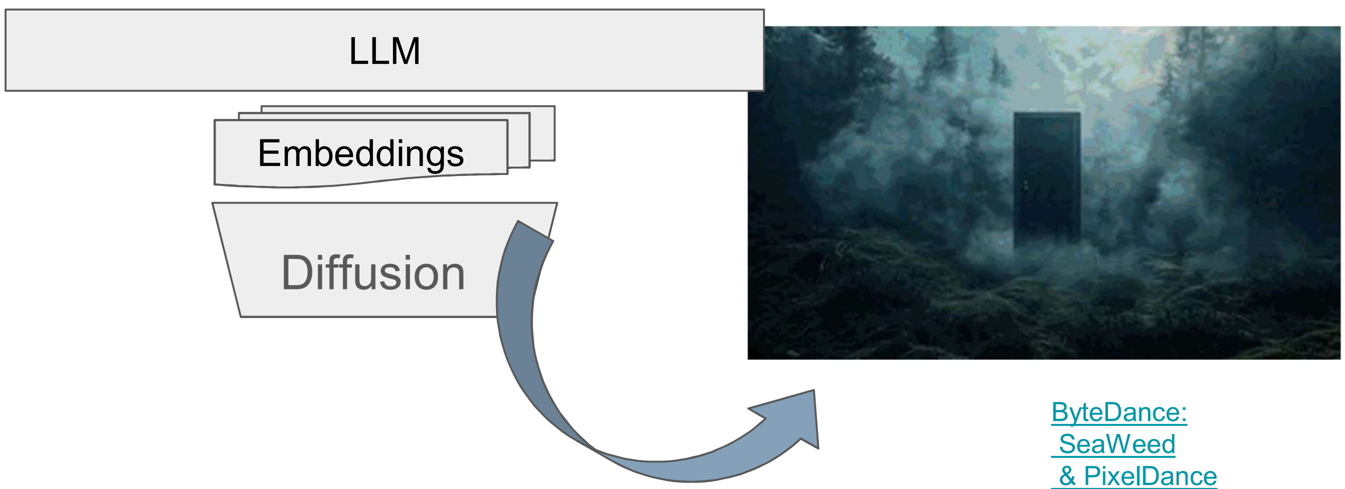

We can output images!

[ByteDance: SeaWeed & PixelDance Presentation]

Information retrieval: A common-use case for LLMs in the software industry

- A common issue shared by many customers is finding information on demand in technical documents.

- Often, the Retrieval Augmented Generation (RAG) framework is used for this purpose.

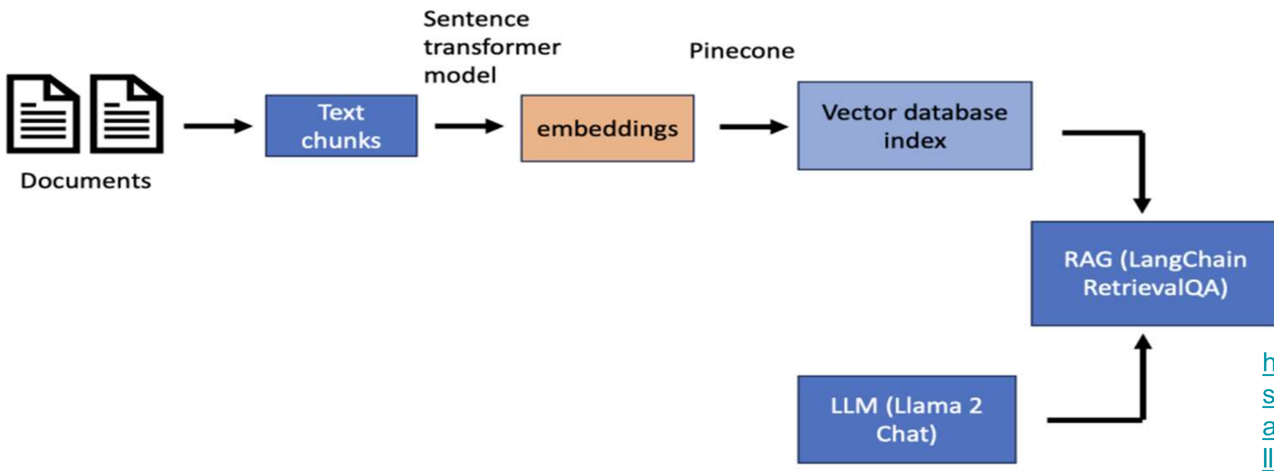

https://agi-sphere.com/retrieval-augmented-generation-llama2/

More than retrieval: Rationale extraction

- Does Retrieval Augmented Generation (RAG) solve everything? Not really: out of context data can be hard to understand, and often leads to incorrect answers

- Need for humans in the loop to verify the soundness. But how do we facilitate the tasks for humans? **Rationale extraction!**

An example of rationale extraction. Needs to be short and include all reasoning (a difficult but common challenge)

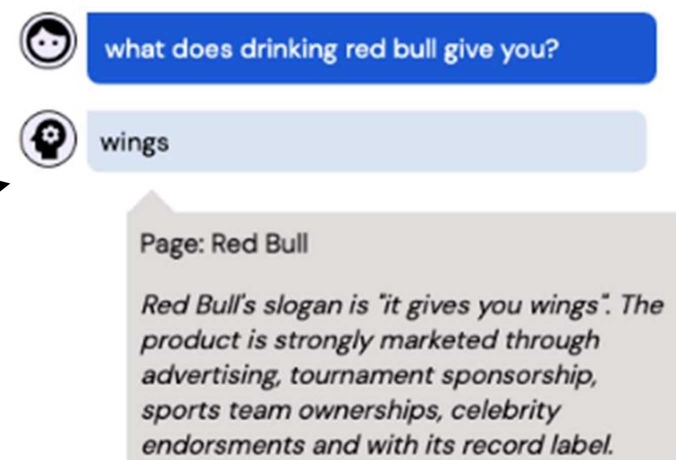

Our method for rationale extraction with LLMs

Our solution was to use the inner workings of an LLM to see what it pays attention to. More specifically, we were analysing the attention matrix to find the sentences in the context needed for generating the answer.

Rationale
In black

Irrelevant
in grey

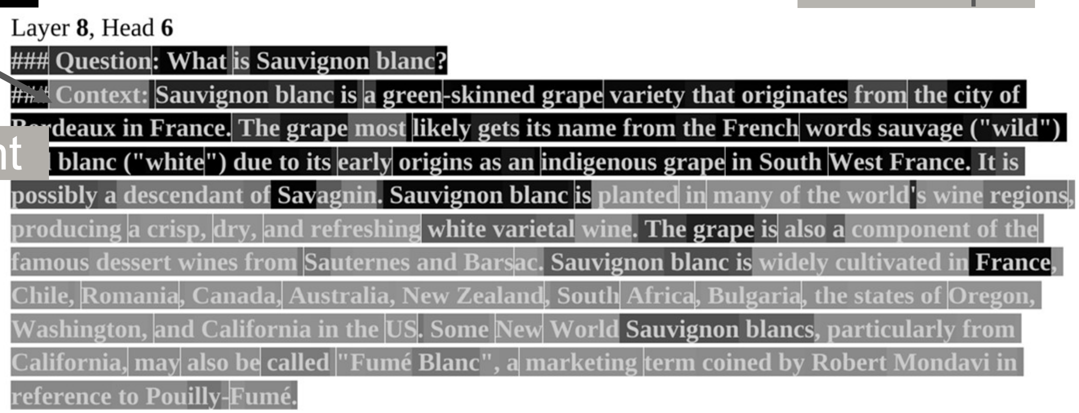

~ Auto-correlation
of the input

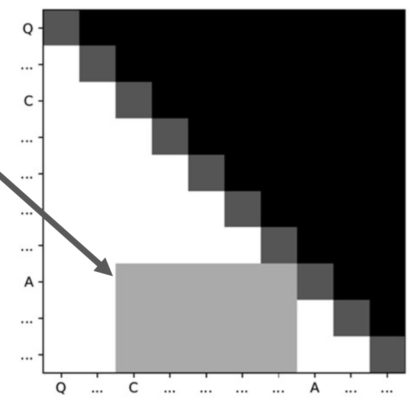

Reference: Pirenne, L., Mokeddem, S., Ernst, D., & Louppe, G. (2024). **Exploration of Closed-Domain Question Answering Explainability Methods With a Sentence-Level Rationale Datase**. ORBi-University of Liège. https://orbi.uliege.be/handle/2268/322654.

This Rationale Extraction problem underlines one of the many needs for mutually beneficial collaborations between the industry and universities. Others exist!

Flow.1-schnell

- Universities needs that can be fulfilled by the industry:
  - research problems that are relevant to the industry
  - computing resources
  - high quality data
  - funding.
- Industry needs that can be fulfilled by universities:
  - capacity to provide solutions to problems requiring research efforts
  - access to top-notch knowledge in AI
  - access to brilliant minds.

NRB and the Montefiore Research Unit of the ULiège have launched a Research Chair in AI for the industry

- Significant funding over four years.
- Two professors involved.
- Considerable thought has gone into putting in place the correct organisational aspects that ensure optimum synergies.

Photo taken during the Chair launch meeting at the NRB premises. Image Left: Prof. Gilles Louppe (L) and Damien Ernst (R). Image Right: the CEO of NRB, Laurence Mathieu with her colleagues.