

uninovis

DATA FOR L.I.F.E.
EUROPEAN UNIVERSITY




ARTIFICIAL INTELLIGENCE FOR EUROPEAN UNIVERSITY ALLIANCES: RESEARCH, TEACHING AND ADMINISTRATION

Málaga, June 15-19, 2026

Day 2. 10:00-11:30

Session 1. Theory



UNIVERSITÉ
**SORBONNE
PARIS NORD**

KAUNO
KOLEGIJA

Tampere University
of Applied Sciences

thws
Technical University
of Applied Sciences
Wittenberg-Schweinfurt


UNIVERSIDAD
DE MÁLAGA

V:
Università
degli Studi
della Campania
Luigi Vanvitelli

UNIVERSITY
OF TIRANA

THE HAGUE
UNIVERSITY OF
APPLIED SCIENCES


Co-funded by
the European Union



Before, we start, I have another question for you

- **Which names are appropriate for AI agents?**
 - A robot-like name that does not have proper syllables such as R2-D2 / C-3PO (Star Wars)
 - A human sounding name that can be real (e.g. Watson (IBM) or made-up (e.g. Alexa, Cortana...)
 - I do not like that AI Agents have names







Agenda for this session


- Review of day 1
- RAG
 - The documents database
 - The database system
- Comparing one-shot with RAG
- Hands-on





1. Review of Day 1





Review of Day 1

- The 4-step model:

Perception → Reasoning → Action → Production

- The oneshot agent:

Perception > Action (LLM call) > Production → (No reasoning)


- The system prompt structure:

identity -- rules -- strict

- Measuring reliability (Rabanser et al, 2026)

Consistency	Same answer each time?
Robustness	Handles unusual inputs?
Predictability	Knows when it will fail?
Safety	How severe are failures?








2. RAG Agents



What is RAG?

- Retrieval-Augmented Generation (RAG) is a technique that allows to complement the identity section of the system prompts with relevant data from a documents **database**
- The LLM will receive a different context information for each query





What do you need to create a RAG agent?

1) A documents database (e.g. a list of research papers, a list of documents describing your University, etc.)

2) A database system that will process the documents for you (we will not manipulate it directly)





RAG Agents: The documents database


It's a collection of text documents that the agent treats as its sole source of knowledge. They are stored in a folder typically named **data** or **docs**.

Q: Which documents should you include in the folder?

Consideration	Why it matters
Scope	Focus on your area of interest. Irrelevant documents add noise and reduce retrieval quality
Granularity	Smaller, focused documents retrieve better — a chapter per file is better than the whole book in one file.
Format	Clean text retrieves better than scanned PDFs. Tables and figures are often lost in extraction. Markdown is the most reliable format.
Redundancy	Avoid. If two documents say the same thing differently, the agent may retrieve conflicting passages and confuse the LLM.
Currency	Outdated documents produce outdated answers. The agent doesn't know which version is newer.




Important: never add confidential data!



Practical tips to create a RAG database

- **Name files descriptively** — the filename appears in source citations
academic_calendar_2026.md ✓
doc3.pdf ✗
- **Add section headers to your documents** — chunking algorithms use headers to split intelligently -
- **Test retrieval, not just answers** — a wrong answer might mean the right document was retrieved but the wrong chunk, or the right chunk wasn't retrieved at all (eg.
- **Start small** — 2-3 documents, test thoroughly, then add more






RAG Agents: The database system

There are two variants, each one with its strengths and weaknesses:

Variant	How it searches	Strengths	Weaknesses
Vector-based	Embeddings + semantic similarity	Synonyms and paraphrases	Requires embedding model; heavier
Vector-less	Keyword matching (BM25)	Lightweight; no embedding model; fast	Misses semantic similarity



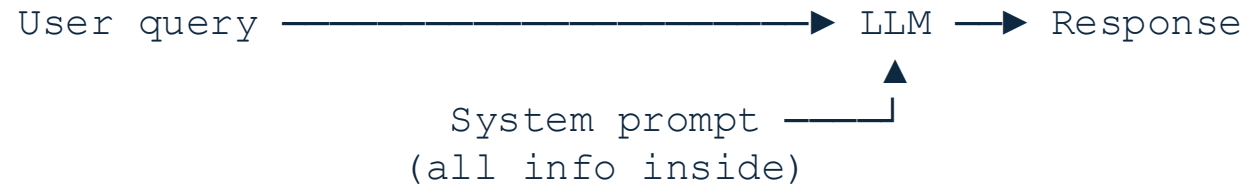


3. One-shot versus RAG

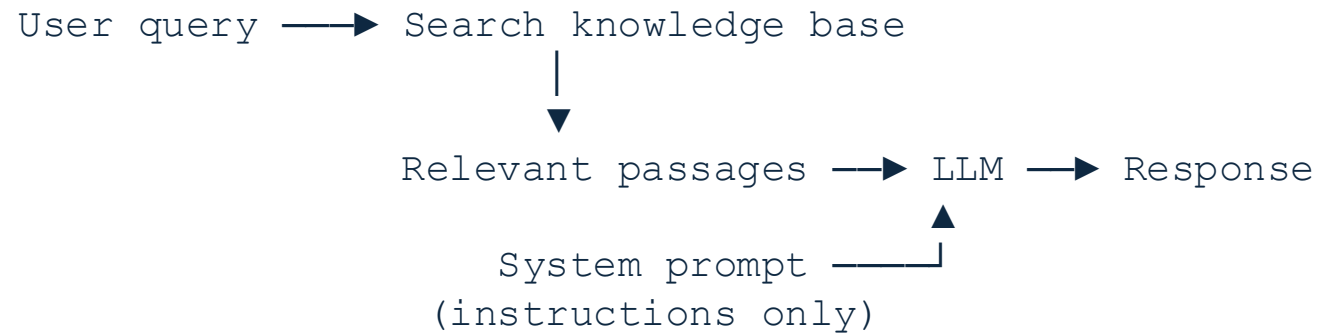


One-shot vs RAG — Flow Diagram

ONESHOT AGENT




RAG AGENT



Knowledge base: handbook.pdf, paper1.pdf, notes.md...





Example: what the LLM sees

1) THE USER'S QUERY

How many faculties does UMA have?

2) THE RETRIEVED CONTEXT

[Source: university_overview.md]

The University of Málaga (UMA) was founded in 1972. It currently has 18 faculties and schools, distributed across two main campuses: Teatinos and El Ejido.


[Source: campus_facilities.md]

The Teatinos campus hosts 14 of the 18 faculties, including Engineering, Science, and Economics. El Ejido houses Fine Arts, Medicine, and Architecture.

3) THE SYSTEM PROMPT:

You are UMA Assistant, specialized in the University of Málaga. Answer ONLY from the provided context.





The expected answer


The LLM responds:

"According to university_overview.md, UMA has 18 faculties and schools. Of these, 14 are on the Teatinos campus and the remaining 4 on the El Ejido campus (campus_facilities.md)."

Note:

- It cites both source documents
- It combines information from two different chunks





One-shot vs RAG — Reasoning

One-shot

- The prompt contains instructions AND knowledge
- The query goes straight to the LLM


RAG

- The prompt contains only instructions, knowledge lives in documents
- The agent now reasons before answering

Important implication:

The quality of the answer now depends on the quality of the search, not just on the LLM.
RAG introduces a new point of failure.





Live demo: Comparing oneshot/RAG


We have prepared two agents:

- One with oneshot architecture: day2_oneshot
- One with RAG architecture: day2_RAG

Both share the same information and the same rules in the system prompt.

Let's ask the same question to both agents side by side, using queries that expose the difference.





Exercise:


<https://uninovis.widening.eu/bipia>

Section: Session Materials: Day 2 exercise

What you will need for RAG agent:

- 1) BIP_info.md
- 2) step_by_step_guide.md
- 3) Terminal_Guide.md
- 4) Tommi_Lite_Course_Book.md






Oneshot vs RAG — Summary


	Oneshot	RAG
Knowledge source	LLM's general training data	Your documents
Course-specific questions	No	Yes
Risk of hallucination	High — invents confidently	Lower — grounded in passages
Can hallucinate?	Yes — always	Yes — when KB doesn't cover it
Needs a knowledge base	No	Yes
Complexity	Minimal	Medium


Key insight: RAG reduces hallucination but does not eliminate it.





Takeaways





What you should have learnt

1. How a RAG agent works in terms of the 4-step model
(the Reasoning step now searches a knowledge base)
2. Criteria used to create a documents database
3. The key differences between oneshot and RAG agents






uninovis

DATA FOR L.I.F.E.
EUROPEAN UNIVERSITY

END OF SESSION 1



UNIVERSITÉ
**SORBONNE
PARIS NORD**

KAUNO
KOLEGJA

Tampere University
of Applied Sciences

thws
Technical University
of Applied Sciences
Wittenburg-Schweinfurt

UNIVERSIDAD
DE MÁLAGA

V: Università
degli Studi
della Campania
Luigi Vanvitelli

UNIVERSITY
OF TIRANA

THE HAGUE
UNIVERSITY OF
APPLIED SCIENCES

Co-funded by
the European Union