ARTEMIS

An AI Framework for Semantic Processing and Natural Language Reasoning in Reactive Heritage Digital Twins

The ARTEMIS Approach to Integrating and Exploring Cultural Heritage Data

13h October 2026,
Austrian Academy of Sciences, Vienna

Miriana Somenzi

ARIADNE RI

Aida Himmiche

ARIADNE RI

ARTEMIS

Applying Reactive Twins to Enhance Monument Information Systems

ARTEMIS is a European research project developing **Reactive Heritage Digital Twins (RHDTs)** to support the understanding, monitoring and preservation of cultural heritage.

Its aims**:**

- **Connect heterogeneous heritage information:** historical records, scientific analyses, 3D models and environmental data.
- **Develop Reactive Heritage Digital Twins**: integrating dynamic information and simulations to support monitoring and conservation.
- **Support heritage professionals**: providing interoperable tools for research, knowledge sharing and informed decision-making.



https://artemis-twin.eu/

ARTEMIS

Teaching Heritage Digital Twins to Understand Us

Reactive Heritage Digital Twins promise **rich, interconnected knowledge**, but only if they can transform thousands of heterogeneous heritage documents into **structured semantic representations** and make the results accessible to everyone.

BUILDING THE KNOWLEDGE → **MAKING IT ACCESSIBLE**

Two complementary AI components. One continuous semantic workflow.



ARTEMIS

The Challenge

Heritage knowledge is difficult to build and, once built, it is difficult to access

BUILDING THE KNOWLEDGE

Heritage information is scattered across reports, PDFs, images and different languages. Transforming it into structured semantic knowledge still depends on experts annotation that cannot scale to growing document collections

MAKING IT ACCESSIBLE

Knowledge Graphs provide rich semantic connections but interacting with them usually requires SPARQL and ontology expertise, placing their contents beyond the reach of many heritage practitioners

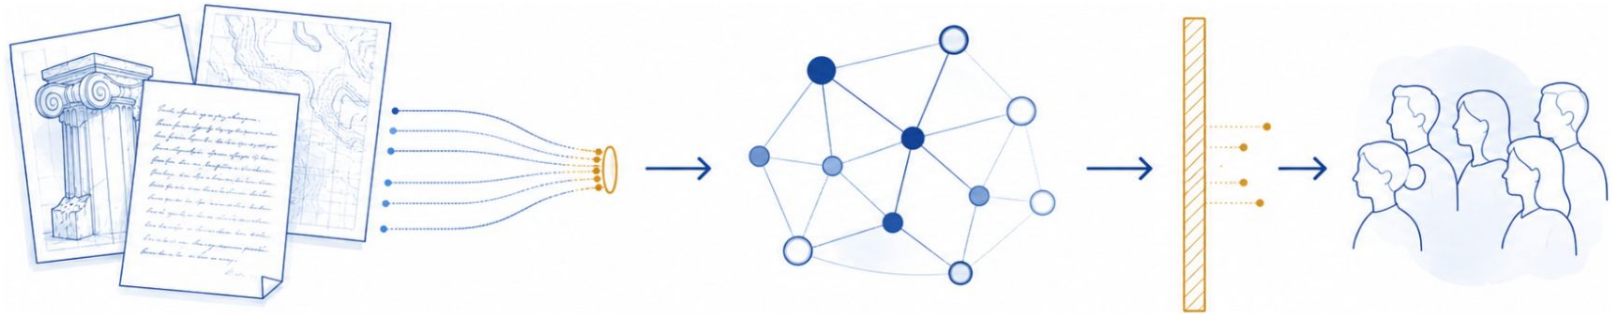

ARTEMIS

Two Complementary AI Contributions

01

BUILDING THE KNOWLEDGE

AI Semantic Annotation Pipeline

- Extract concepts from heterogeneous heritage documents
- Link them to domain ontologies
- Validate the annotations and generate the RDF

02

MAKING IT ACCESSIBLE

Artemisia Reasoning Engine

- Accept questions in natural language
- Retrieve and reason over the Knowledge Graph
- Return evidence-backed answers with provenance

ARTEMIS

What is a Graph?

A graph is a *mathematical representation* of a system where **entities** are represented as **nodes** and **relationships** are represented as **edges** (connections, links between nodes).

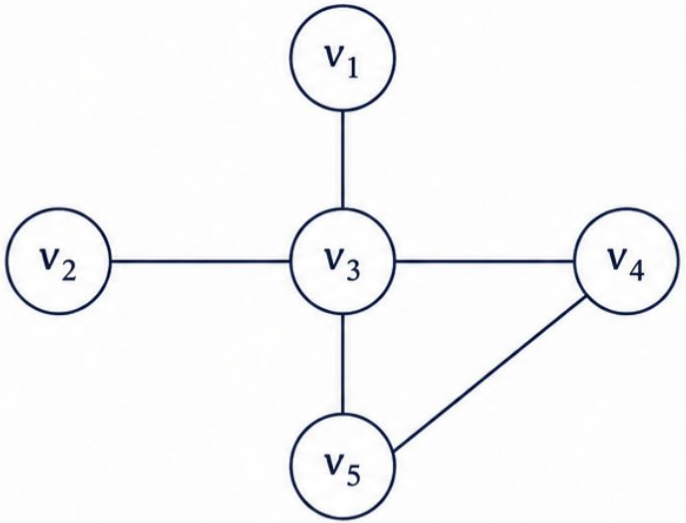

FORMAL DEFINITION

A graph $G = (V, E)$ where:

$V = \{ v_1, v_2, \ldots, v_n \}$
set of vertices (nodes)

$E \subseteq V \times V$
set of edges (relationships)

Each edge $e = (u, v) \in E$ connects two vertices $u, v \in V$.

ARTEMIS

Heritage is naturally a Graph



ARTEMIS

Ontologies: a shared semantic vocabulary

Ontologies are structured, formal representations of concepts in a domain and the relationships among them.

They provide a common vocabulary so that different datasets can interoperate.

They are expressed through a conceptual model containing

- **classes**, such as objects, actors, places, and events
- **relationships** that exist between those classes

They form a **directed graph** that makes knowledge explicit and interpretable

This structure allows both humans and machines to navigate complex heritage data coherently

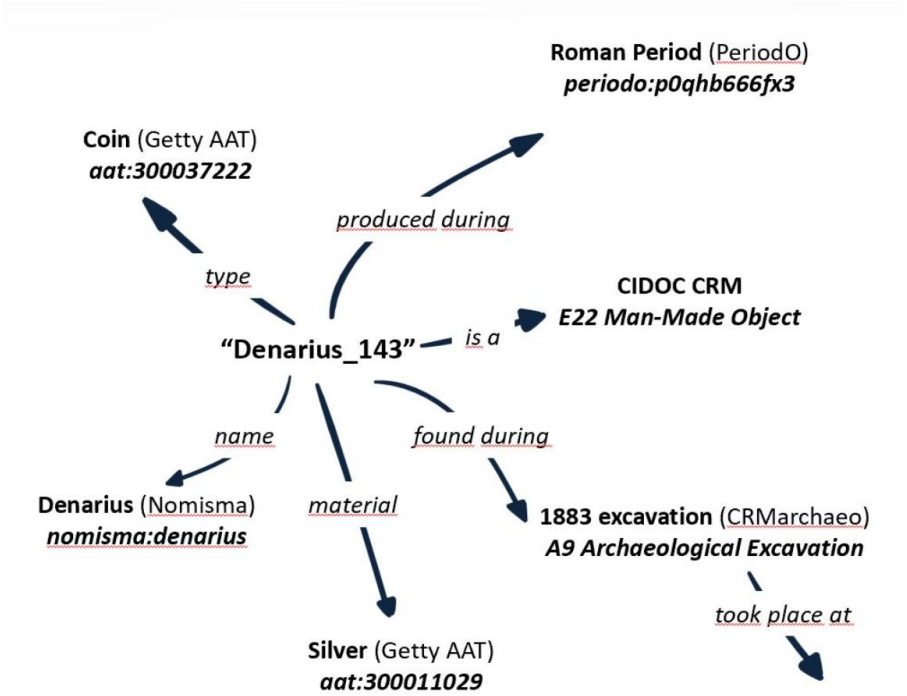

One resource, several assertions

An RDF triple connects a subject to an object through a named predicate.

SUBJECT	PREDICATE	OBJECT
Portrait of Leo X	consists of	wood (plant material)
Portrait of Leo X	consists of	oil paint
Portrait of Leo X	has type	paintings (visual works)
Portrait of Leo X	has current location	Galleria degli Uffizi

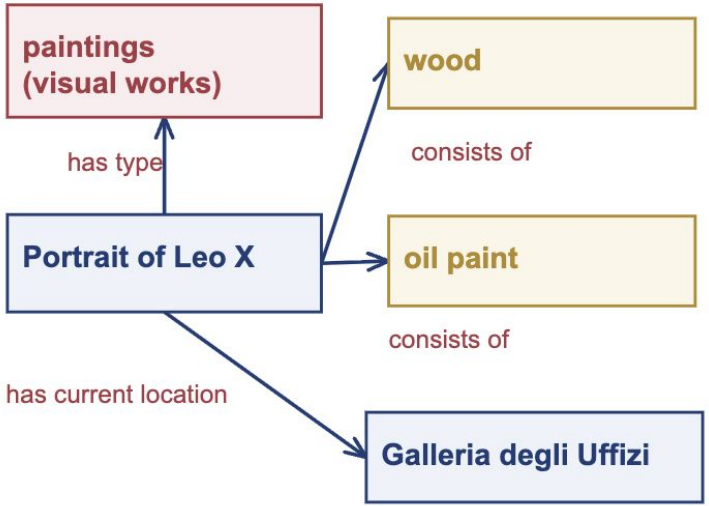

ARTEMIS

From Ontologies to Knowledge Graphs

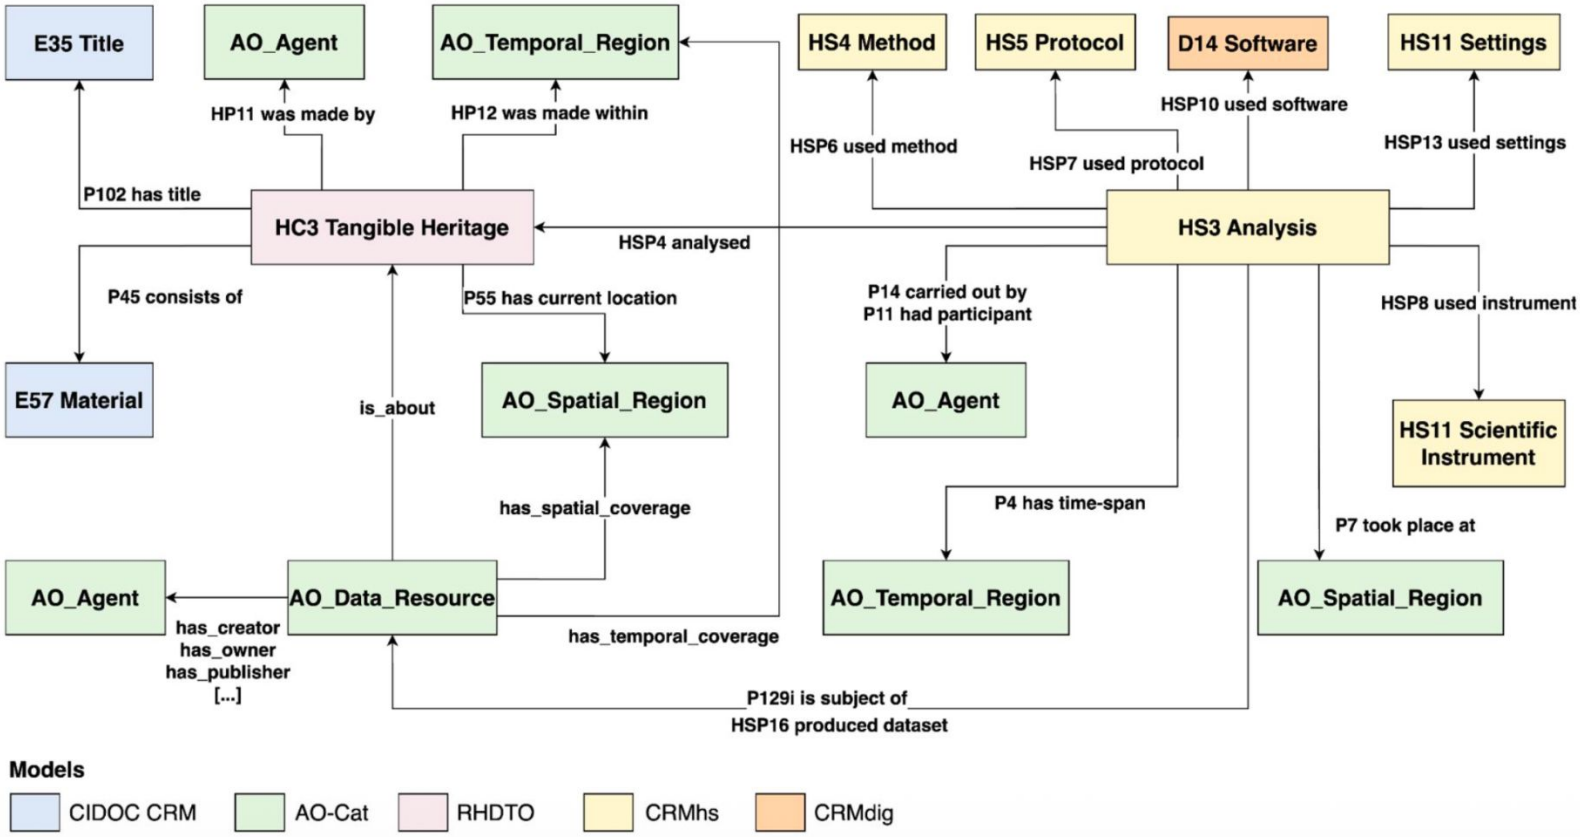

ARTEMIS

From Ontologies to Knowledge Graphs



ARTEMIS

The Knowledge Extraction Bottleneck

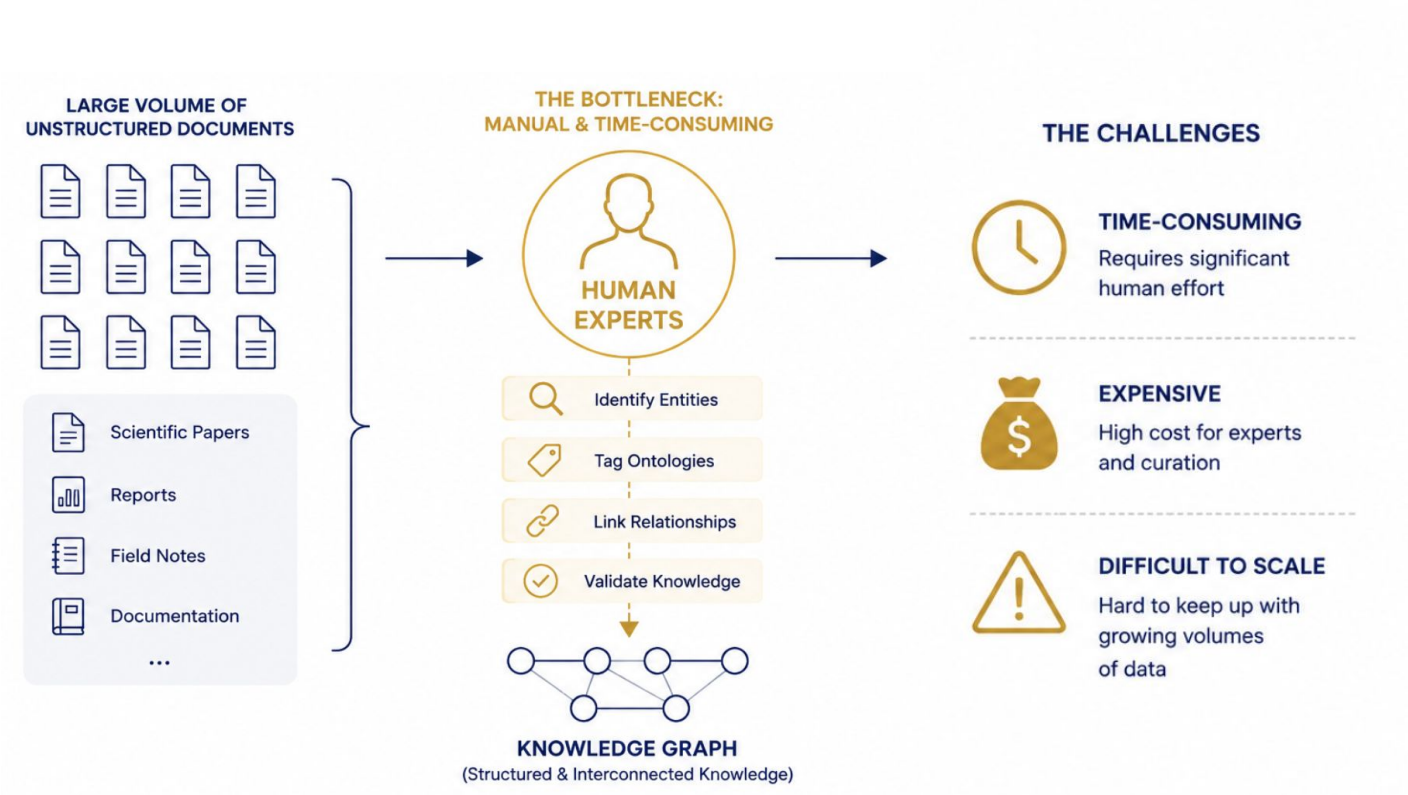

ARTEMIS

The Complexity Challenge

Heritage knowledge is a rich multiplex network

How to deal with all this complexity?

AI! ….. AI?

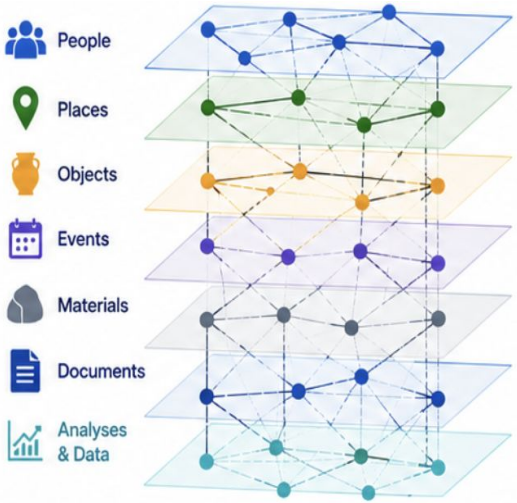



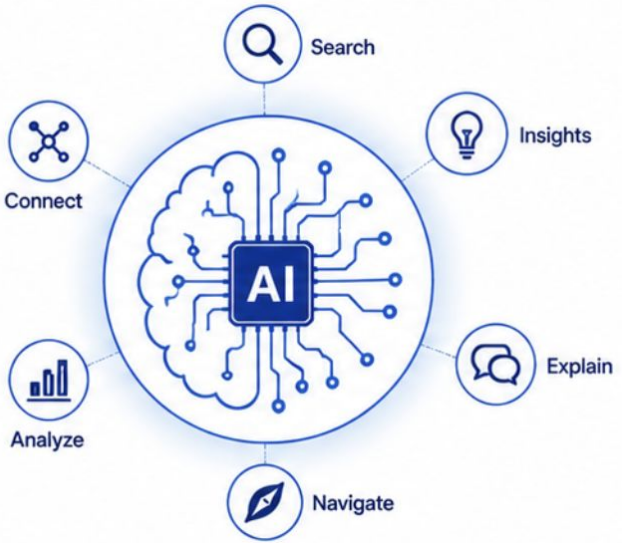

ARTEMIS

What if there was an **AI-based component** that could help you read your scientific documents and extract structured knowledge?

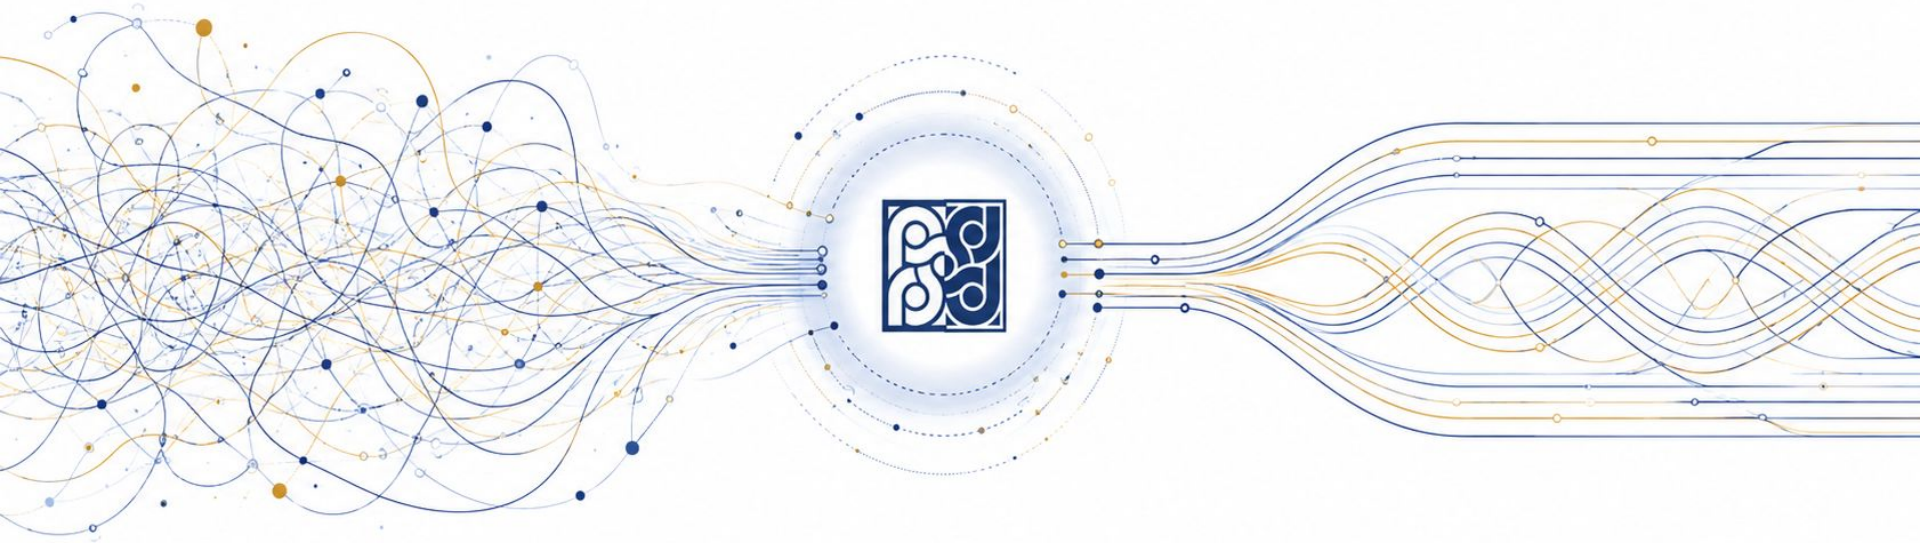

ARTEMIS

Artemisia Oracle: the pipeline

The ARTEMIS AI-assisted pipeline designed to transform unstructured cultural heritage documentation into structured, ontology-aligned knowledge

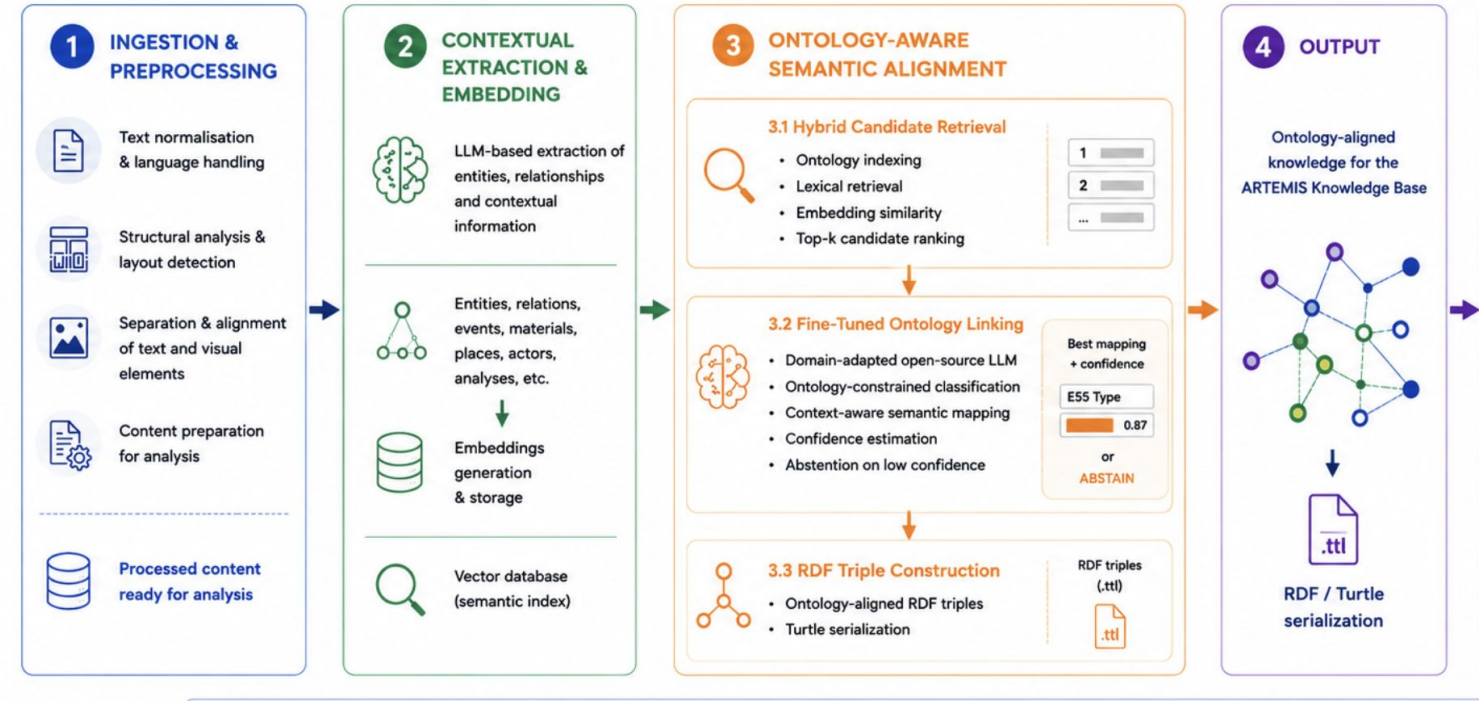

ARTEMIS

Artemisia Oracle: the backend components



ARTEMIS

Beyond Entity Extraction: Why Context Matters

Extracting relevant terms and entities is not enough: their meaning depends on the **context** in which they appear

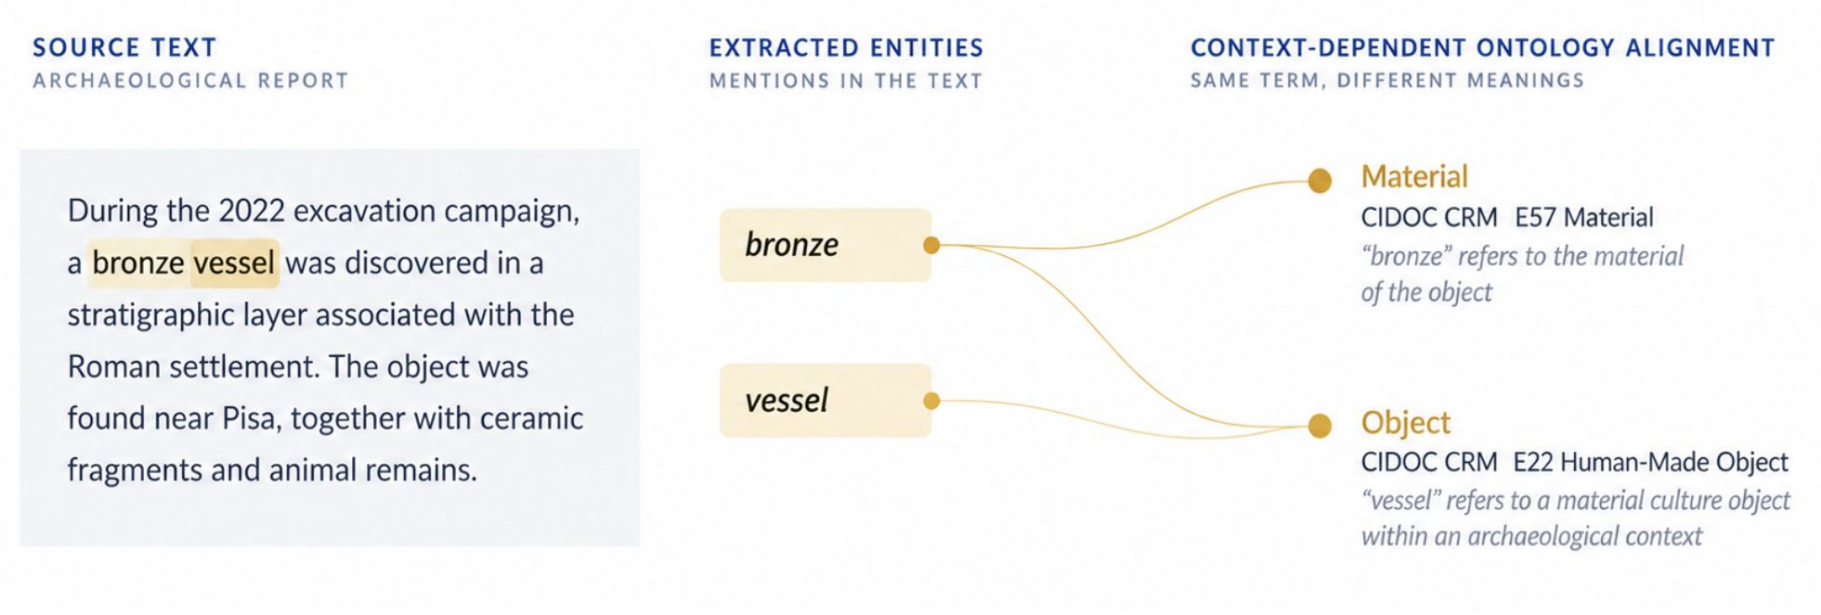

ARTEMIS

Embeddings

Word embeddings are **dense numerical representations** of text that capture semantic meaning

- Each word/sentence is converted into a vector of numbers:

  bronze -> [0.82, 0.71, 0.15, −0.24]
  copper -> [0.79, 0.75, 0.12, −0.20]
  vessel -> [−0.31, 0.12, 0.89, 0.65]

- Words with **similar meaning** have **similar vectors**, so they're close in a semantic space

- This allows us to compare and retrieve information based on meaning, not lexical match

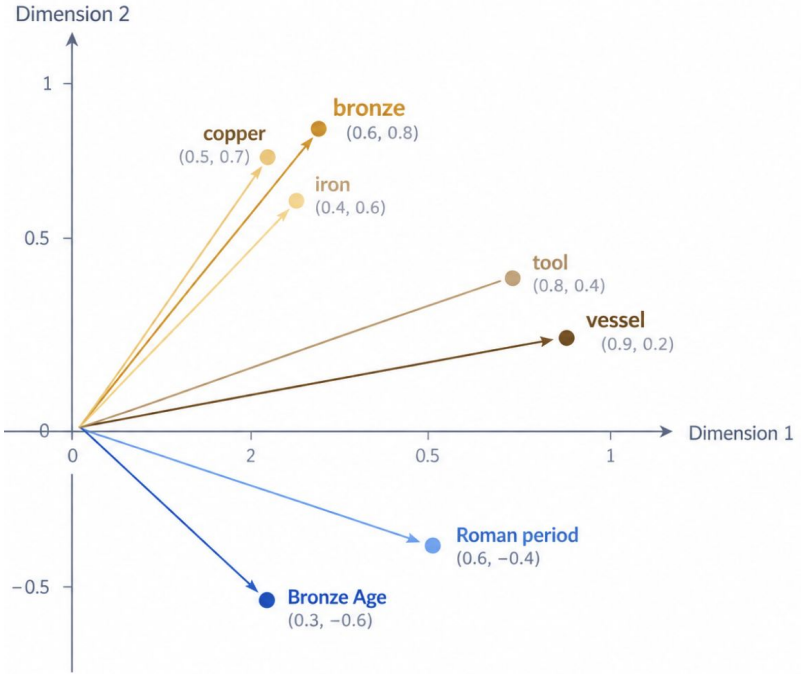

ARTEMIS

Context Aware Retrieval

Artemisia Oracle uses **SentenceTransformers** and **FAISS** to retrieve relevant contextual information for extracted entities, supporting their subsequent alignment with domain ontologies

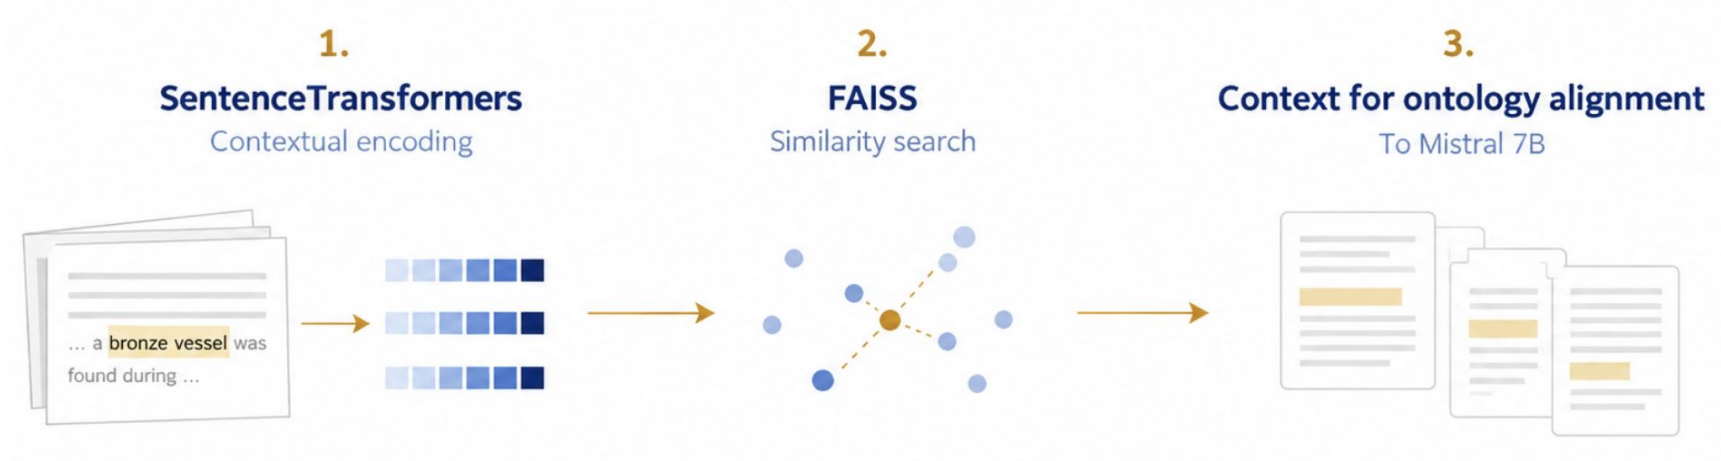

Extracted entities and their surrounding text are encoded as dense vector representations

The vector index identifies semantically similar passages, allowing the system to retrieve relevant contextual information

Retrieved passages are combined with the extracted entities and ontology information

ARTEMIS

Fine-tuning the model

Fine-tuning specializes Mistral 7B for cultural heritage terminology and ontology alignment, while preserving the pre-trained knowledge and leveraging external semantic grounding



MISTRAL 7B Instruct

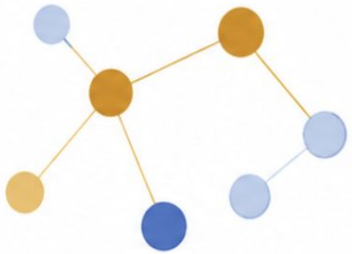

Training data

Annotated cultural heritage examples with validated ontology labels

Supervised QLoRA fine-tuning

Keep the pre trained model frozen and train lightweight adapters on a quantized model

Ontology-aligned predictions

Context-aware classification of cultural heritage entities into formal ontology classes

ARTEMIS

Why fine tuning?

Commercial, proprietary models

General purposes, closed systems

- **Closed and not transparent:** model architecture, training data and update processes are not available
- **Limited customization:** customization is provider dependent and very limited
- **Data and privacy constraints**

Open source/ Open weights models

Domain specific, open and controllable

- **Transparent and reproducible** for the most part: model architecture and weights are available
- **Domain adaptation**
- **More control**: fine tuning and inference can be run locally and data can be used following GDPR and EU AI Act
- **Sustainable and extendable**: adaptable to new domains and tasks without dependence on external providers

AI or not to AI

To trust,

or not to trust?

That is the question.



ARTEMIS

Why do LLMs hallucinate

Large Language Models don’t store **facts,** they store **statistical patterns.**

When they’re generating answers they’re not **retrieving** it from a database, they are **predicting** the most likely next word based on patterns and probabilities observed during their training phase

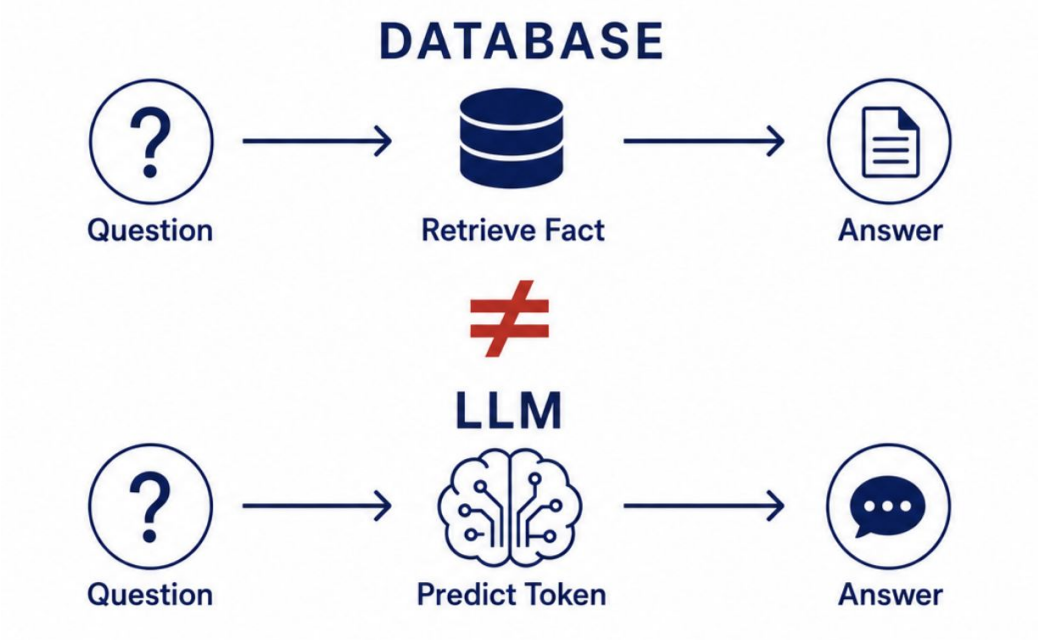

ARTEMIS

Grounding AI with Knowledge Graphs



In **ARTEMIS**, we explore a complementary relationship between open source, fine tuned Large Language Models and Knowledge Graphs: combining the two enables more **reliable** and **explainable** AI systems.

ARTEMIS

Keeping the humans in the loop

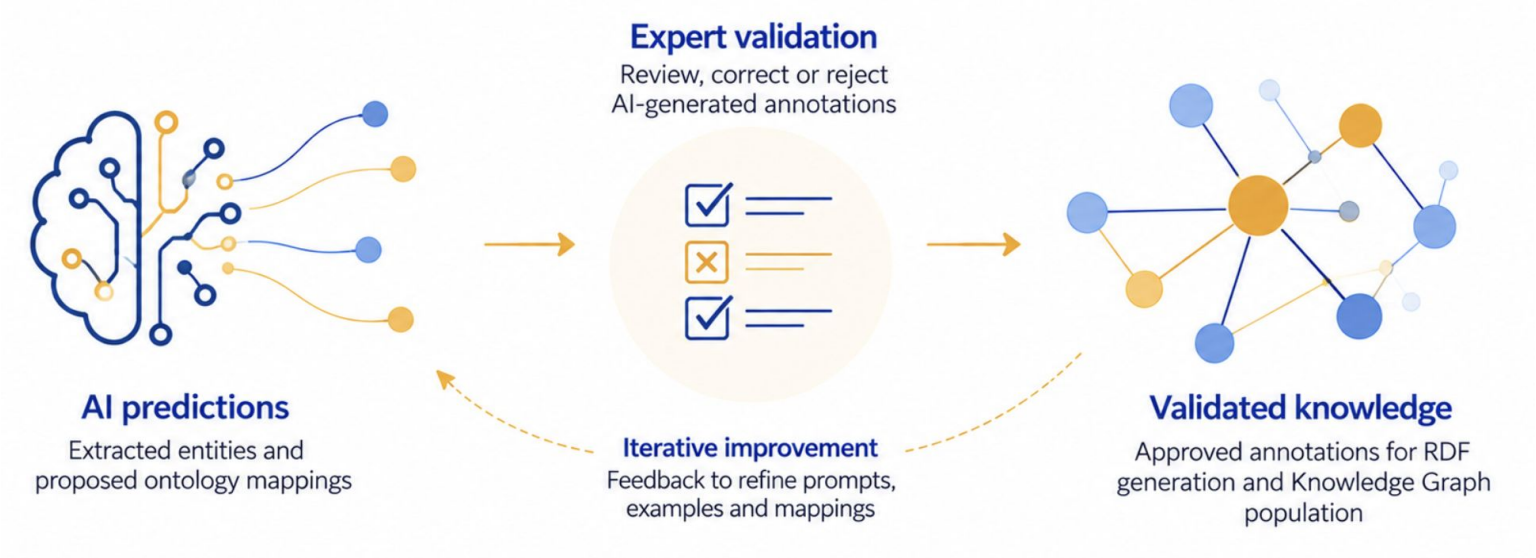

Human oversight is a fundamental part of our approach to responsible AI. Artemisia Oracle proposes ontology classes and properties with a confidence score, but domain experts retain the responsibility of reviewing, correcting and validating these suggestions before they become part of the Knowledge Graph.

ARTEMIS

Natural Language Reasoning Pipelines for KG Querying

ARTEMIS

Two barriers between documentation and an answer

Knowledge acquisition

Reports and datasets contain useful information in heterogeneous formats and languages.

Knowledge access

Structured knowledge requires users to understand identifiers, relations and query syntax.

“Which data resources are linked to the Stonehenge monument?”

A useful answer needs a resolved object and a supported connection.

ARTEMIS

This section will explain how we:

01 Read the graph

Resources, relationships and bounded SPARQL queries

02 Follow the evidence

A question becomes an anchor and a supported connection

03 Explain the answer

Artemisia separates retrieval from language formulation

04 Locate a failure

Evaluate each stage to identify the component to improve

ARTEMIS

One title can refer to two different records

HE Heritage Entity

Stonehenge

A cultural heritage object or site

heritage_entity
HE canonical URI

DR Data Resource

STONEHENGE

A dataset, report or related digital asset

data_resource
DR canonical URI

A shared title does not establish a shared identity.

ARTEMIS

RDF: resources, predicates and literals

Illustrative RDF data

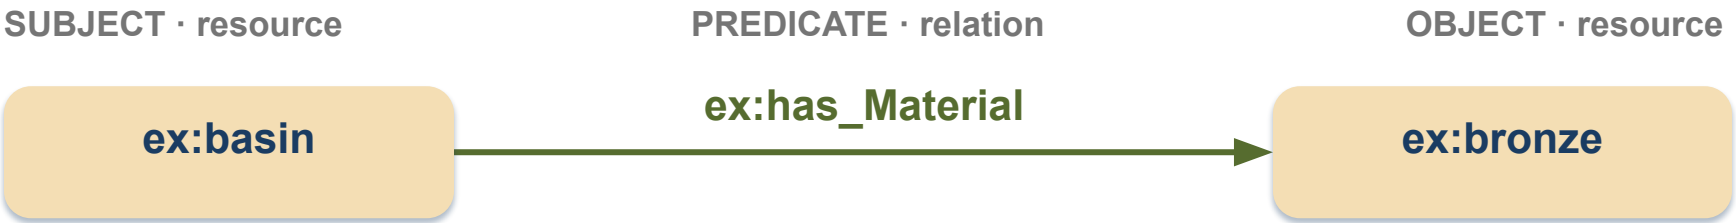

A predicate labels a directed edge. A resource URI identifies a node.

ARTEMIS

SPARQL returns relationships and their source

This query find all data linked to “basin”

```
SELECT ?p ?o ?g WHERE {
  VALUES ?s {
    <https://artemis-kg/entity/basin>
  }

  GRAPH ?g { ?s ?p ?o }
  FILTER(isIRI(?o))
}

LIMIT 30
```

Returned variables

?p	?o	?g
relation	*endpoint*	*source graph*
P45_ consists_of	Bronze	<graph_uri>

GRAPH retains source context.

This implementation also supports incoming edges and predicate filters.

ARTEMIS

Some questions need more than one edge

“Is there a scientific report on Stonehenge?”

N-hop query illustrative path:

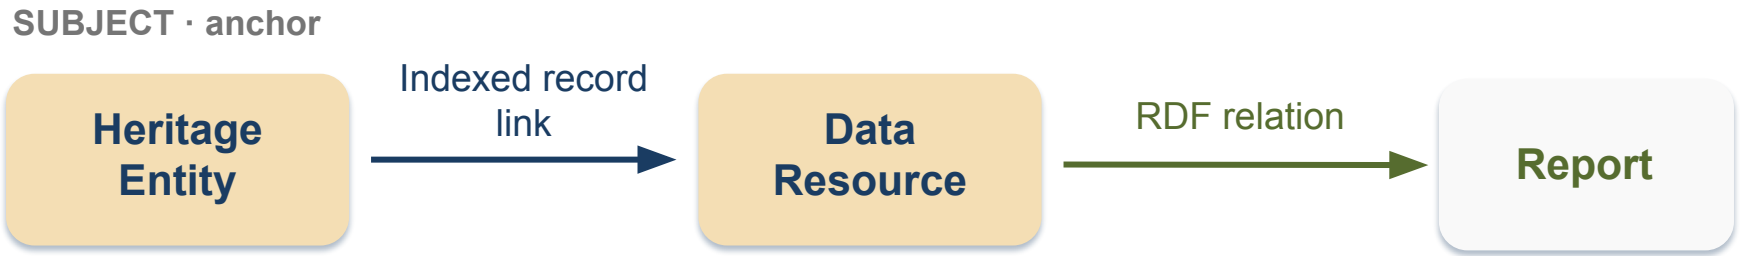

Two hops connect the anchor to the requested endpoint.

Path discovery is bounded to three hops and excludes repeated nodes.

ARTEMIS

Querying difficulty extends beyond syntax

“Is there a scientific report on Stonehenge?”

Question element	Computational decision
Stonehenge	Resolve an HE or DR as the starting record
Scientific report	Determine the requested endpoint type
on/about	Select the supported link operation
available links	Choose evidence associated with the anchor

A valid retrieved connection can still be irrelevant to the question.

ARTEMIS

Artemisia's natural-path pipeline

Deterministic retrieval, followed by evidence-based language generation

User question

PREPARATION
SPARQL templates: RDF paths up to 3 hops

01
Grounding
Resolve anchor
1 Selected record *(if any)*
2 Explicit mention
3 Compatible conversation

02
Planning
Select operation
Determine target type
Set sources and depth

03
Retrieval
Candidate paths
Document links
RDF paths
Hybrid paths
Filter and rank

04
Evidence
Endpoint
Ordered nodes
Typed edges
Provenance

05
Mistral
Ollama
Formulates an answer from the supplied evidence

Grounded answer

ARTEMIS

A request contains more than a sentence

```
{
  "message": "Are there reports on this
monument?",
  "sessionId": "workshop-session",
  "selectedRecords": [
    {"id": "<Stonehenge_ID>",
     "type": "heritage-entity"}
  ],
  "history": []
}
```

Context inputs

Question text

Portal selection
(selected record/s if any)

Session identifier
and recent context

Selection and conversation state can supply a missing explicit name.

ARTEMIS

Grounding follows explicit precedence

1

Selected record

Use a valid portal selection as the authoritative anchor.

2

Current mention

Resolve a new name when selection does not take precedence.

3

Conversation state

Reuse a compatible recent referent when no explicit anchor exists.

Ambiguous or stale references require a clearer anchor.

ARTEMIS

Global grounding produces typed candidates

1. Normalize supported case, punctuation...
2. Extract known names, quoted mentions and configured aliases.
3. Search HE and DR indices separately.
4. Deduplicate by record type and canonical identity.

“bronze vessel from Lecce” / Bronze Tripod Basin (HE)

ARTEMIS

The signals behind anchor ranking



The weights are heuristic.

ARTEMIS

Uncertainty changes the response

0.72

Anchor acceptance threshold

0.12

Cross-type ambiguity margin

“What reports are linked to the bronze basin record?”

“I found more than one possible record.
Which one do you mean?”

Clarification preserves the distinction between an HE and a DR.

ARTEMIS

The planner selects a controlled operation

Operation	Example request
Direct record lookup	What is this record’s title?
Explicit document link	Which data resource is linked to this object?
RDF path	Which material is this object made of?
Hybrid path	What report is reached through its linked resource?
Shared-property rule	Do these selected objects share a material?
Unsupported	Outside the available operation coverage (default to a broader reasoning branch)

The plan constraints anchor, endpoint type, evidence source and depth.

SPARQL query execution

This query find all data linked to “basin”

```
SELECT ?p ?o ?g WHERE {
  VALUES ?s {
    <https://artemis-kg/entity/basin>
  }

  GRAPH ?g { ?s ?p ?o }
  FILTER(isIRI(?o))
}

LIMIT 30
```

Returned variables

?p	?o	?g
relation	*endpoint*	*source graph*
P45_ consists_of	Bronze	<graph_uri>

The query execution uses SPARQL templates based on the planner output

ARTEMIS

The evidence package preserves the connection

```
{
  "status": "supported",
  "endpoint": {
    "label": "STONEHENGE",
    "type": "data_resource"
  },
  "path": {
    "nodes": ["Stonehenge", "STONEHENGE"],
    "edges": ["related_data_resource"]
  }
}
```

For each edge

Relation
Direction
Provenance

For each node

Identity and readable label

Full inspectable evidence differs from the compact LLM input.

Mistral formulates the answer from evidence

Generation input

- Question
- Execution status
- Verified endpoint/s
- Reduced path

Generation output →

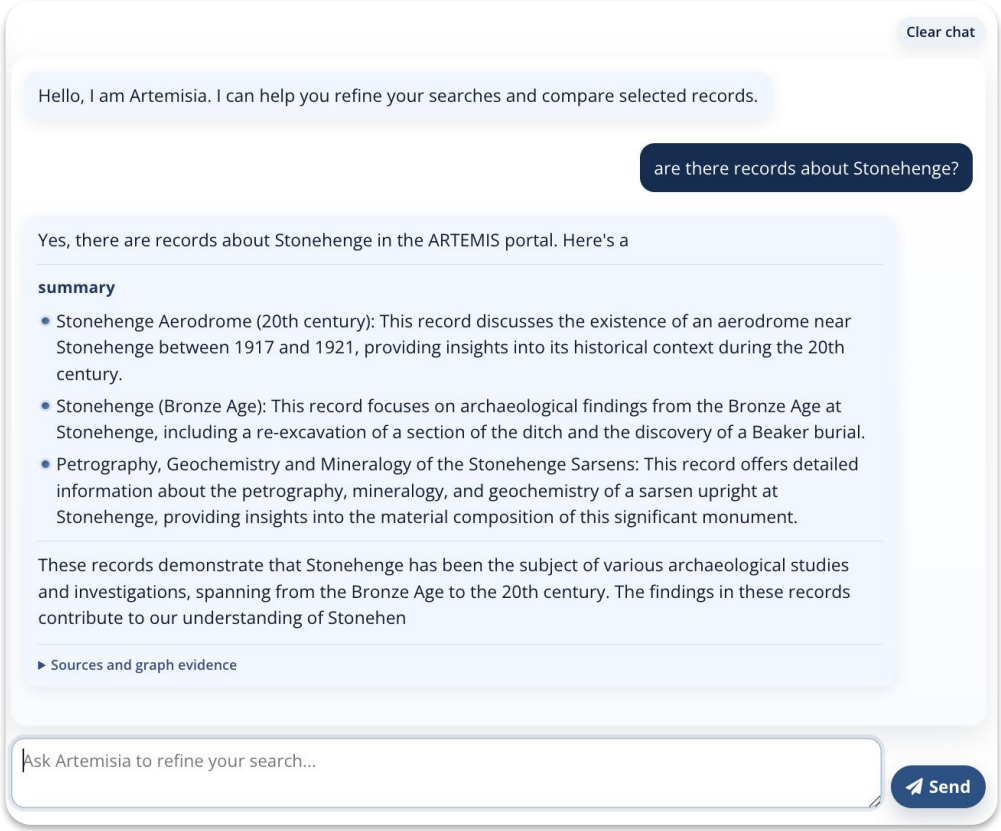

ARTEMIS

A plausible answer can exceed the evidence

Supported by the link

“STONEHENGE is a dataset linked to the Stonehenge monument.”

Enriched response

“This dataset provides a detailed and up-to-date record of the monument.”

A stored link alone establishes neither completeness nor freshness.

No facts are inferred without the proper evidence support, but additional context may be synthesized from the evidence to enrich the response.

ARTEMIS

Evaluation follows the pipeline stages

Anchor accuracy

Did the system select the intended record?

Endpoint accuracy

Did it reach the expected result?

Exact-path accuracy

Did it recover the expected ordered path?

Answer supportedness

Did the response satisfy the evidence-based scorer?

A failure in one stage can propagate into the final answer.

ARTEMIS

From an observed error to a component

Wrong object or dataset	→	Mention extraction and anchor ranking
Correct anchor, wrong endpoint	→	Plan, candidate coverage and path ranking
Missing connection	→	Data coverage, freshness and depth limit
Correct evidence, extra claims	→	Generation input, prompt and output checks

Deterministic execution still depends on rules, data and ranking.

ARTEMIS

The main architectural takeaways

Record	Resolve the record before retrieving its connections.
Evidence	Retain the path before formulating the answer.
Evaluation	Assess grounding, retrieval and wording separately.

ARTEMIS

Current features and ongoing research directions

Established in the implementation

- Collection-specific grounding
- Controlled operations
- Bounded path retrieval
- Inspectable source attribution

Still open

- Broader query coverage
- Stronger endpoint/path selection
- Inventory freshness
- Performance at scale

General-domain accuracy and comparative superiority remain untested.

How do we achieve higher semantic reasoning even at a larger scale?

How can we deal with complex research tasks?

ARTEMIS

Embedding-based retrieval using a query graph

“Which methods were used in the scientific analysis of Stonehenge?”

Interpret the question

Identify the entities and relations.

→

Build a query graph

Typed nodes + relation + unknown node to solve

→

Retrieve candidates

Encode the pattern and Find matching KG pattern.

Original graph evidence determines what Artemisia can answer.

ARTEMIS

The ontology defines the query pattern

“Which methods were used in the scientific analysis of Stonehenge?”

Classify

Stonehenge: Heritage Entity
Requested answer: Method

Link

Resolve Stonehenge to its canonical KG node.

Map the relation

Map other entities and relations to ontology and mark the unknown node.

CONCEPTUAL QUERY PATTERN

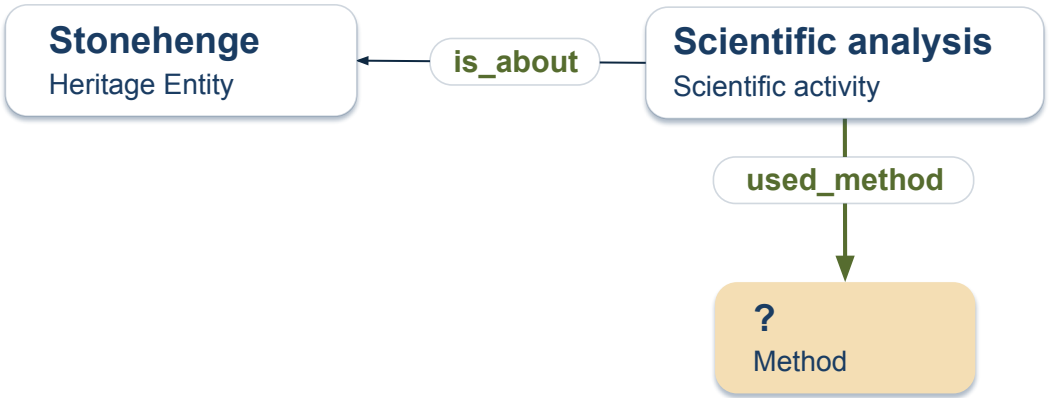

The real ontology determines the predicates, direction and any intermediate nodes.

ARTEMIS

One encoder for query and candidate graphs

OFFLINE: PREPARE KG

KG paths / subgraphs
Prepare candidate graphs

ONLINE: ENCODE EACH QUESTION

Query graph
Anchor, relation, variable

Trained graph encoder

Candidate vectors
Stored in a vector index

Query vector
Compare with the index

Training makes matching query and candidate graphs close in vector space.

ARTEMIS

A similarity match must answer the question

Question: “Which methods were used in the scientific analysis of Stonehenge?”

Similarity search

Returns ranked candidate graphs

Ex: Method node with a verified path to Stonehenge

Include as a supported answer

Ex: Used dataset in scientific activity about Stonehenge

Ranked low: wrong answer type

Ex: Methods used in analysis of prehistoric monuments

No verified subject link: omit

Check the answer type, the anchor and the actual supporting relation.

ARTEMIS

Cross-graph attention reranks the shortlist

Similarity retrieves likely matches. A pairwise model compares them more closely.

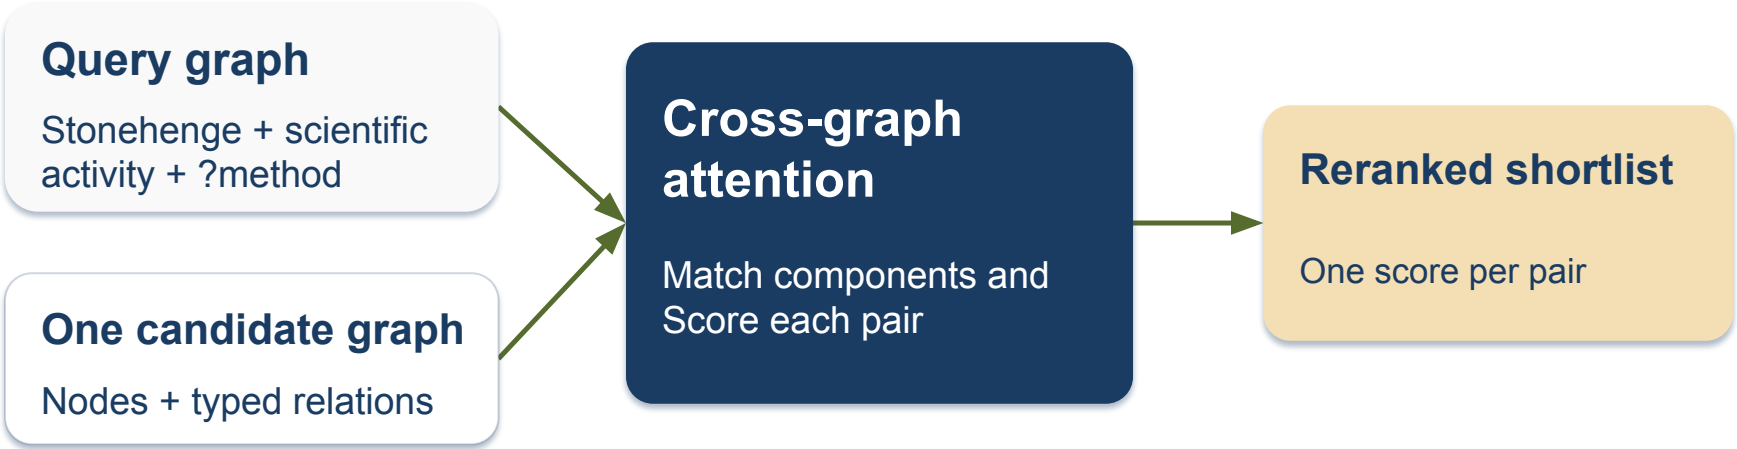

Test whether closer matching improves ranking enough to justify its added latency.

ARTEMIS

Provenance stays attached to each candidate

Offline: store a source reference beside every indexed encoding.

Vector index entry

Candidate vector
Candidate ID
Source graph + KG version

Similarity search returns candidate IDs

Use each ID to recover its original graph evidence.

Recover and validate source evidence

Entity URIs + supporting triples + named graphs
Retain these references in the answer package.

Similarity traces candidate selection. Source triples support the answer.

ARTEMIS

Integration with Artemisia's natural-path branch

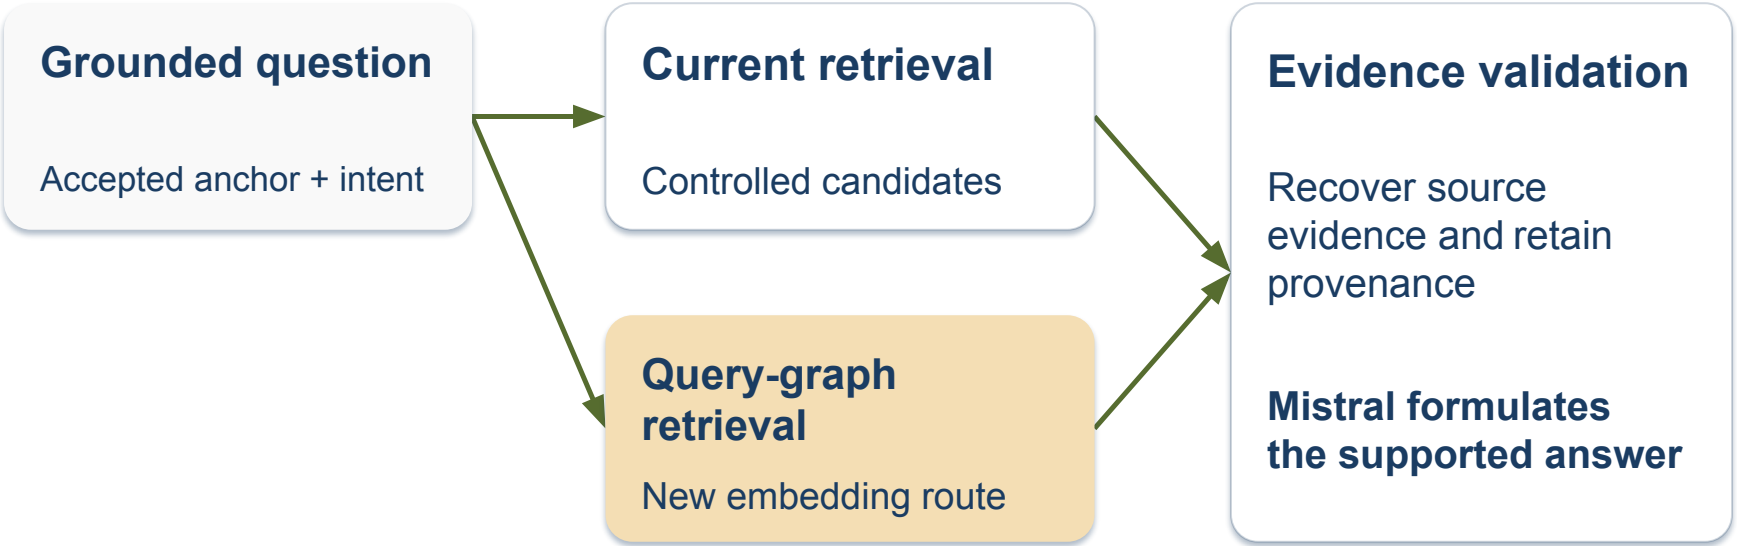

Start by reranking verified paths. Then test wider candidate discovery.

ARTEMIS

Questions and discussion

1. What information is an archeologist most likely to search for?
2. Are there searches a user would perform that artemisia isn't currently facilitating?
3. What evidence should the interface expose to users?

ARTEMIS portal: artemis-portal.ino.cnr.it

ARTEMIS

Thank you!

Any questions?

ARTEMIS is a project funded by the European Commission under Grant Agreement n.101188009.



Funded by
the European Union

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.