CONFERENCE WORKSHOP 2026

Working Sustainably with AI in Digital Heritage Research: A Case-Based Walkthrough



Larissa Terranova

Anthropologist & Digital Humanities Researcher

Centro Studi Americanistici "Circolo Amerindiano"

🔗 linkedin.com/in/larissa-terranova/

Workshop Outline (90 Minutes)

- **Part 1:** Context Setting & The Sustainability Dimensions
- **Part 2:** System Walkthrough: Ingestion, Retrieval, & Evaluation
- **Part 3:** Critical Reflection: Domain Scenarios
- **Part 4:** Interactive Discussion: Trust & Skill Development
- **Part 5:** Open Challenges & Collaborative Futures

Learning Outcomes

Gain insight into AI research workflows and practical decisions.

Apply critical frameworks for evaluation.

Recognize irreplaceable human decision points.

CONTEXT SETTING

Why Sustainability Matters

Beyond efficiency metrics to responsible AI integration.

Three Dimensions of Sustainability

Research

Supporting ongoing scholarly work without creating fragile technical dependencies or unmanaged tech debt.

Interpretive

Ensuring automation enhances rather than erodes the human expertise and cognitive authority of the researcher.

Ethical

Addressing environmental costs of LLMs and maintaining strict algorithmic accountability in digitization.

SYSTEM WALKTHROUGH

The Heritage Lens Architecture

A step-by-step analysis of a Multimodal RAG Prototype.

PART 2

The Heritage Lens approach relies on friction as a feature

“Augmentation, not automation.”



To stress-test this balance, we apply the system to a highly complex, 3,000-year-old ethnoarchaeological reality: The Olmec Colossal Heads.

NotebookLM

System Architecture Overview

Phase	Mechanism	Scholarly Function
Ingestion	PyMuPDF + LangChain	Multimodal source extraction.
Embedding	CLIP (Vision Transformer)	Cross-modal semantic bridging.
Storage	ChromaDB	Vectorized archival retrieval.
Evaluation	LLM-as-a-Judge	Automated 'Agent' grading.

Heritage Lens: Friction as a Feature

Standard Automation	Augmentation (Heritage Lens)
Efficiency & Speed	Deepened Human Expertise
Frictionless UX	Friction as a Feature
Passive Reviewer	Active Interpreter
AI guesses the single answer	AI surfaces the conflict
Immediate outputs	Deliberate human checkpoints



NotebookLM

Three critical decision points requiring human authority

	Decision Point	AI Reality	The Human Imperative
Row 1 (Ingestion / Curation)	Authority	Garbage in, garbage out.	Human curation establishes baseline scholarly validity before AI sees data.
Row 2 (Query / Retrieval)	Gaps	Semantic search fails when terminology varies across decades.	Catching context collapse. Recognizing absence is a deeply human skill.
Row 3 (Output / Verification)	Trust	The system synthesizes from retrieved sources.	Final meaning-making belongs to the researcher. You are the author.

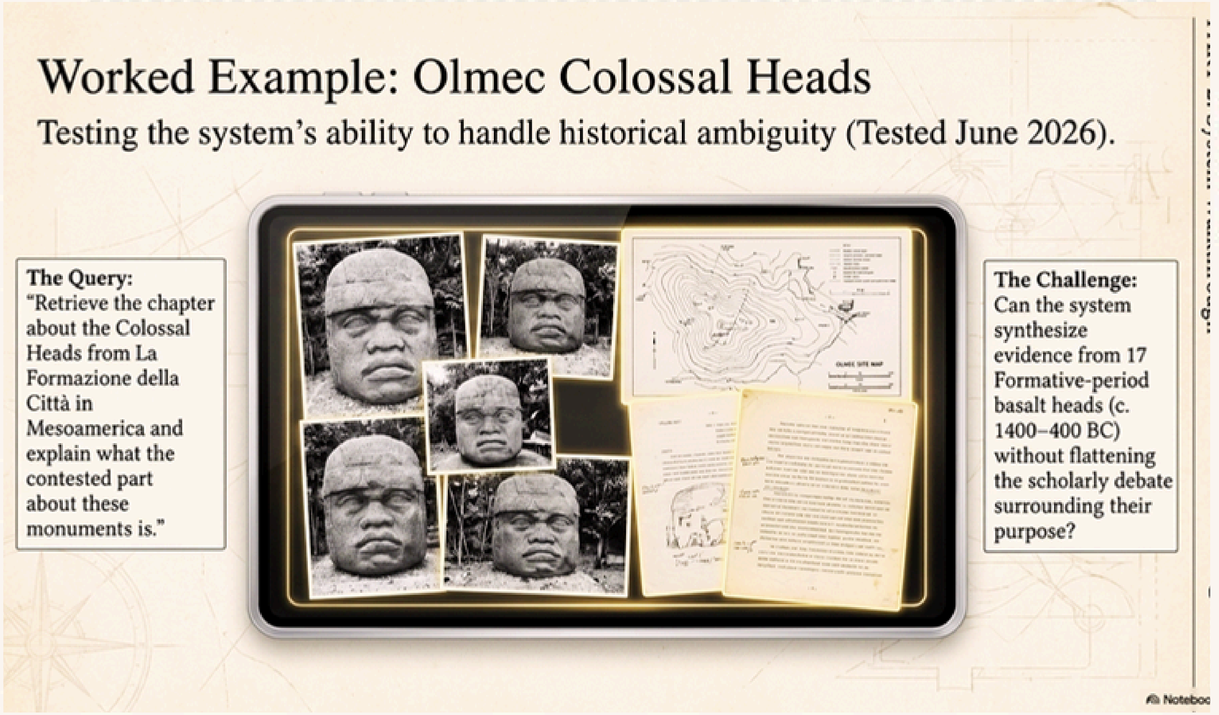
NotebookLM

Case Study: Olmec Colossal Heads

Query: "Retrieve the chapter on Colossal Heads from 'La Formazione della Città' and explain contested parts."

- 17 Formative-period basalt heads.
- Surfaced multiple scholarly interpretations: recarved thrones, ritual mutilation.
- System Evaluation: 5/5.

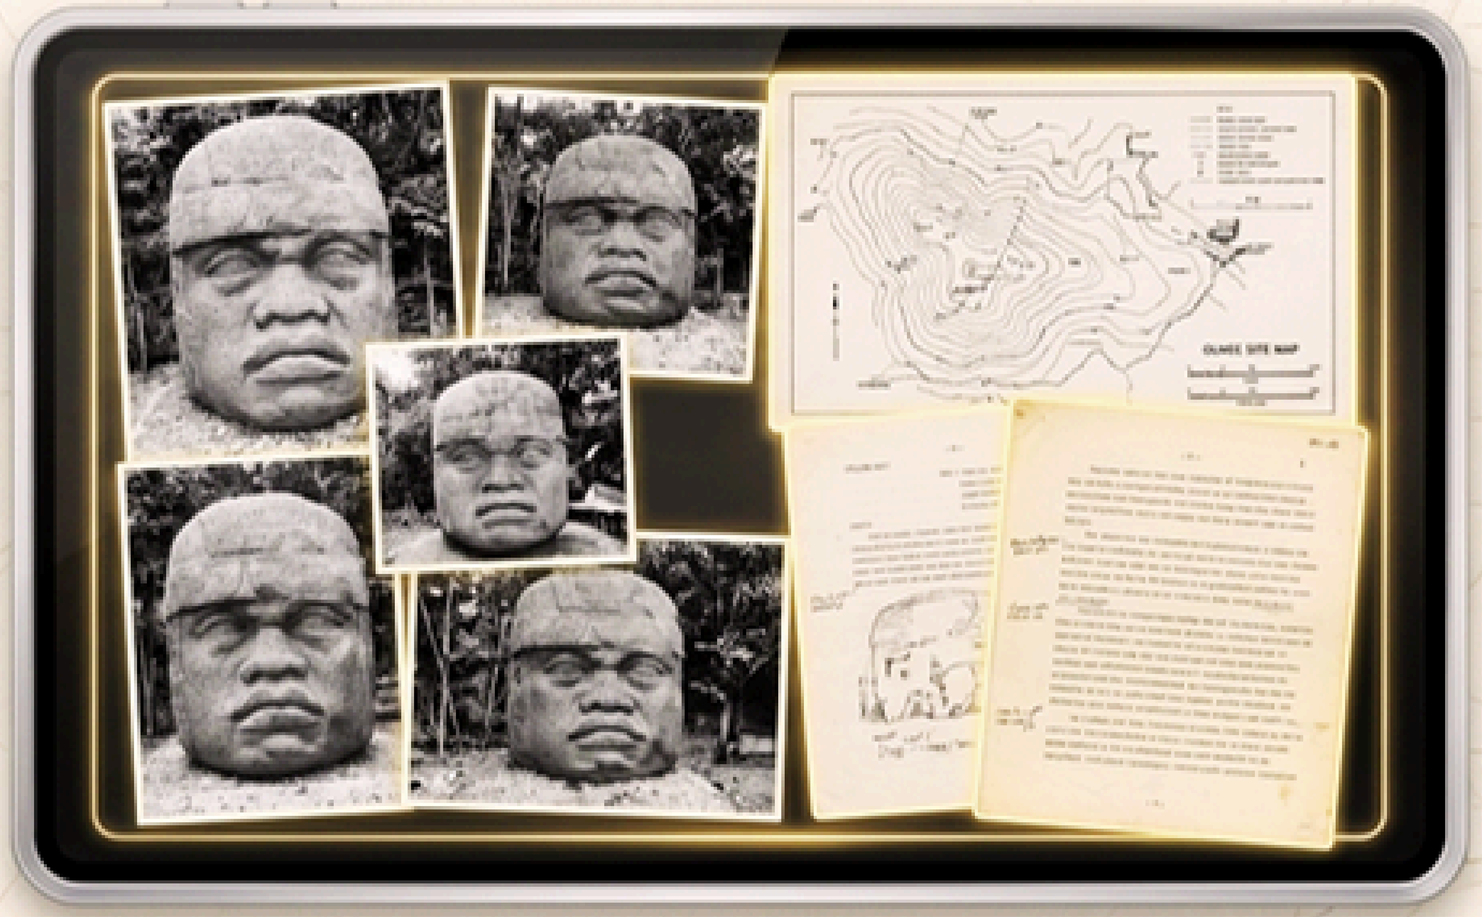
Insight: The system surfaced interpretive contest, not a flattened consensus.



Worked Example: Olmec Colossal Heads

Testing the system's ability to handle historical ambiguity (Tested June 2026).

The Query:
"Retrieve the chapter about the Colossal Heads from La Formazione della Città in Mesoamerica and explain what the contested part about these monuments is."



The Challenge:
Can the system synthesize evidence from 17 Formative-period basalt heads (c. 1400–400 BC) without flattening the scholarly debate surrounding their purpose?

CRITICAL REFLECTION

Domain Context Scenarios

Mapping sustainability to specific heritage challenges.

Facilitated Discussion: Navigating the Trade-offs



Effort Redistribution

Where in your specific workflow does AI genuinely save time versus creating entirely new categories of verification labor?



Skill Development

How do we train emerging scholars to evaluate AI synthesis if they haven't manually done the foundational, frictional retrieval work themselves?



Institutional Contexts

How does the threshold for acceptable AI automation differ between a GLAM public portal and an academic peer-reviewed publication?



Sustainability Metrics

How do we measure "sustainable" AI integration when efficiency and speed are no longer the primary goals?

NotebookLM

Facilitated Reflection Exercise

A. Indigenous Knowledge

Safeguarding sacred material from algorithmic scraping. Who has the authority to decide what gets ingested?

B. Museum Engagement

Balancing accessibility with accuracy in public-facing bots. Where does the conservator become irreplaceable?

C. Dissertation Research

How does AI assistance change expertise development? When does efficiency work against scholarly growth?

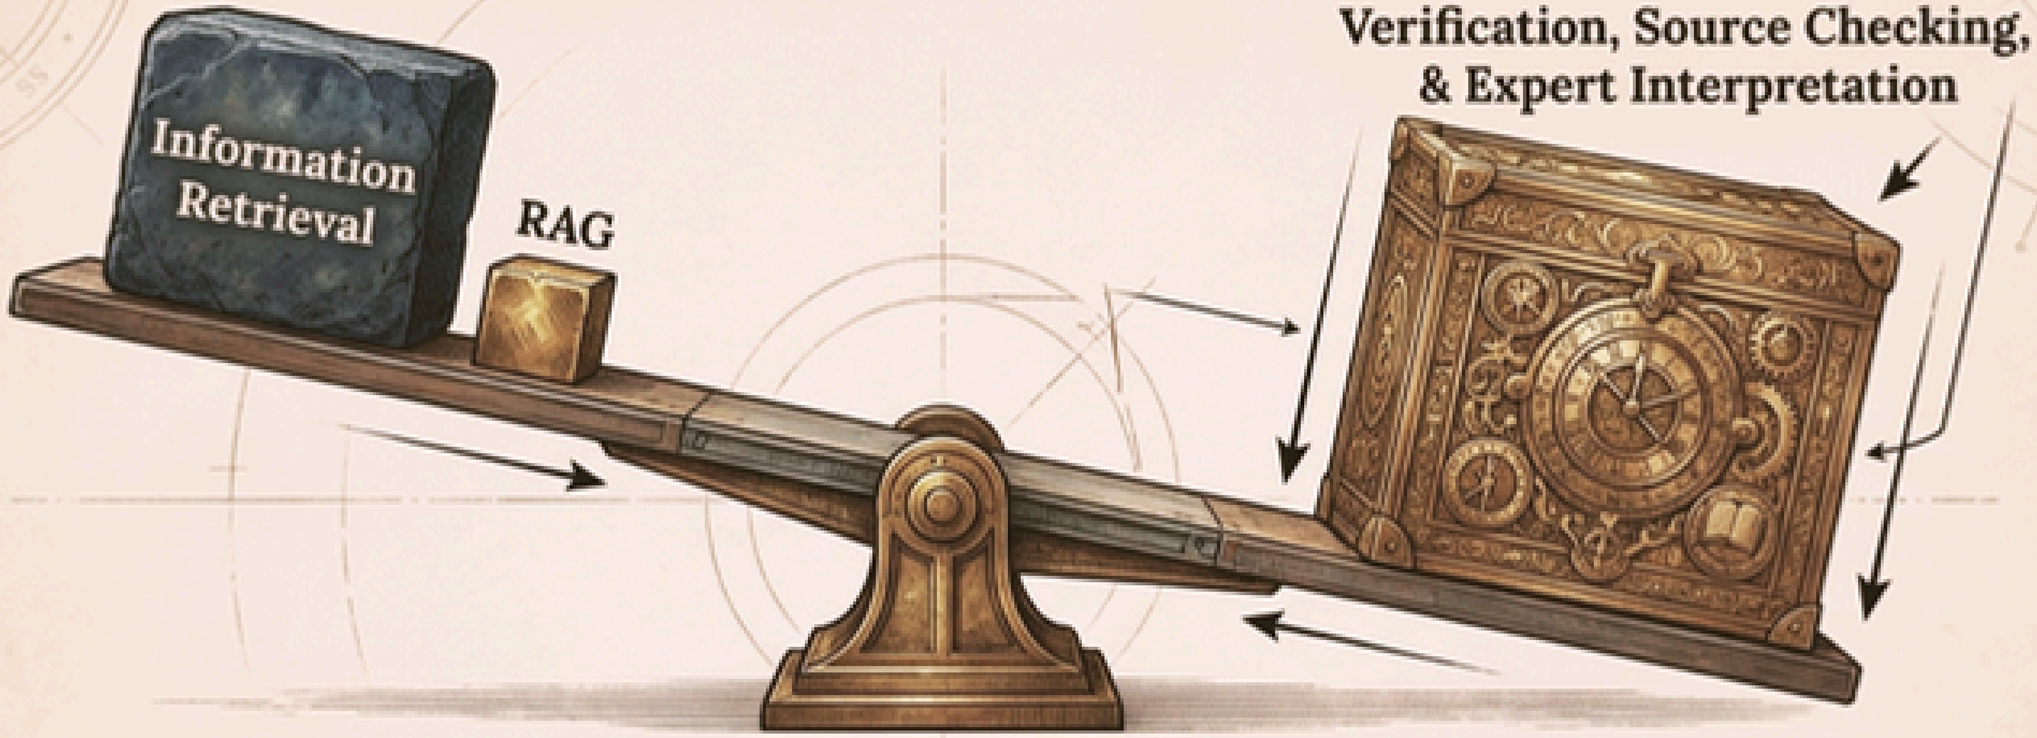
INTERACTIVE DISCUSSION

Friction & Authority

Where do we draw the line between helpful and dangerous?

PART 4

Redefining Our Relationship with Friction



1. Effort Redistribution	2. Skill Development	3. Trustworthiness
Where does AI genuinely reduce effort versus introducing new types of work (e.g., source-checking)?	Where do we draw the line between helpful automation and dangerous deskilling?	Where do you see the greatest friction between automated synthesis and scholarly authority in your own work?

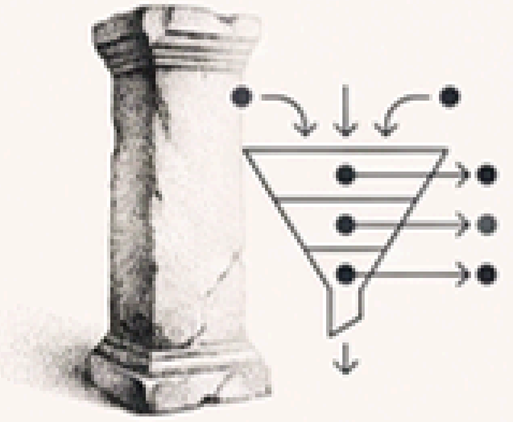
PART 4: Interactive Discussion

NotebookLM

Prompts for Open Discussion

- **Effort Redistribution:** Does AI save time or shift labor to verification?
- **Institutional Contexts:** How do decision points shift for community-led archives?
- **Deskilling Risk:** When does AI automation erode foundational judgment?
- **Sustainability:** Beyond speed, how should we measure 'success'?

Three Pillars for Sustainable AI Practice



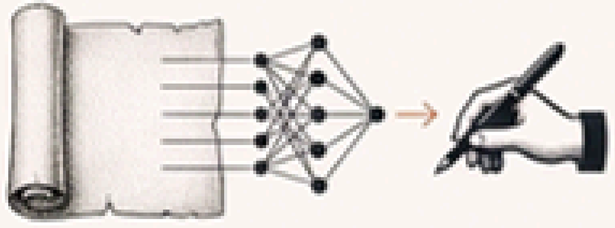
1

Curate with Rigor: The system is only as good as its boundaries. Rigorous human curation before ingestion is the primary defense against algorithmic hallucination.



2

Embrace Friction: Do not design for seamless automation. Design for transparent augmentation that forces the researcher back to the primary source for verification.



3

Protect Expertise: Redistribute computational effort toward retrieval, but ensure that final interpretive authority and historical synthesis never leave human hands.

Reflection: *What is one specific thing you will do differently in your digital practice after today?*

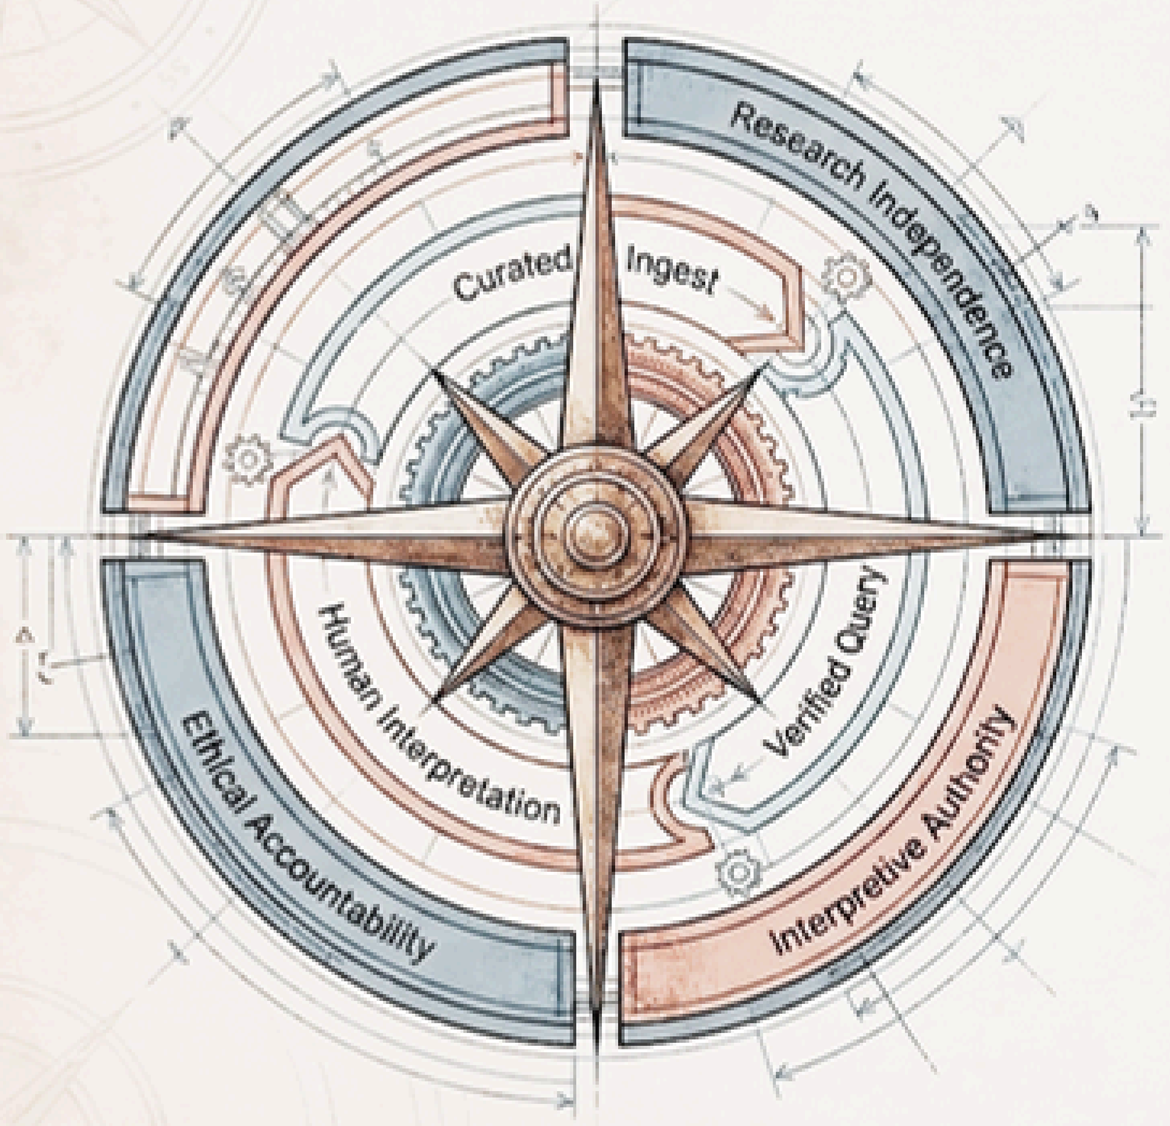
NotebookLM

OPEN CHALLENGES

Collaborative Futures

What remains to be solved in Heritage AI.

Expanding the Blueprint: Unresolved Pathways



Critical questions remain unresolved:

- **Portability:** How adaptable is this RAG architecture and its tollbooths to other specialized, non-Mesoamerican heritage domains?
- **Cultural Sensitivity:** How do decision point structures need to shift dynamically when interfacing with culturally sensitive or restricted materials?
- **Metrics of Success:** How do we empirically measure genuine scholarly enhancement, rather than metrics based purely on output speed?

NotebookLM

What is missing matters as much as what is present

Heritage Lens Agent



Thank you!

Feedback Form



LinkedIn



Scan to share feedback or connect with me

© Larissa Terranova 2026.

Except where otherwise indicated, original content is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International licence (CC BY-NC-ND 4.0). Third-party images and source materials remain subject to their respective copyright and licensing terms.