




Interdisciplinary Uses of AI in the Humanities:

Developing an Epistemically Grounded LLM-System for Archival Exploration

Sebastian Oliver Eck,* Ashley Duraiswamy, Puyu Wang, Yishun Lu - funded by CrossLinks Award 2026, St Cross College
ARCA - AI for Research, Collections, and Archives, University of Oxford, UK



The Problem

Large-scale digitisation of archives has created vast machine-readable archives that no individual scholar can realistically analyse. This ‘archival deluge’ presents a methodological challenge:

How can historians engage with an ever-expanding evidential record while maintaining the principles of historical source criticism?

Why Existing AI Tools for Archival Exploration Fall Short

Conventional AI systems are designed to optimise fluent responses rather than historically defensible claims. They often obscure provenance, flatten ambiguity, reconcile contradictory evidence without justification, and provide little transparency regarding how conclusions were reached.

➡ These behaviours conflict with established scholarly practice.

Our Central Claim

AI systems cannot be reliably and thereby safely adopted in the humanities unless the epistemic requirements of scholarship are encoded directly into the system architecture.

We propose a **scholarly AI system**: a source-grounded architecture that translates those requirements — an **Epistemic Constitution** — into retrieval and synthesis itself.

Conventional AI Systems	Our Scholarly AI System
Optimises for <i>plausible answers</i>	Optimises for <i>historical rigour</i>
<i>Hides</i> uncertainty	<i>Preserves</i> uncertainty
Opaque provenance	Sentence-level traceability
Stateless	Cumulative research record

Solution: Building an epistemically grounded LLM-system for archival exploration

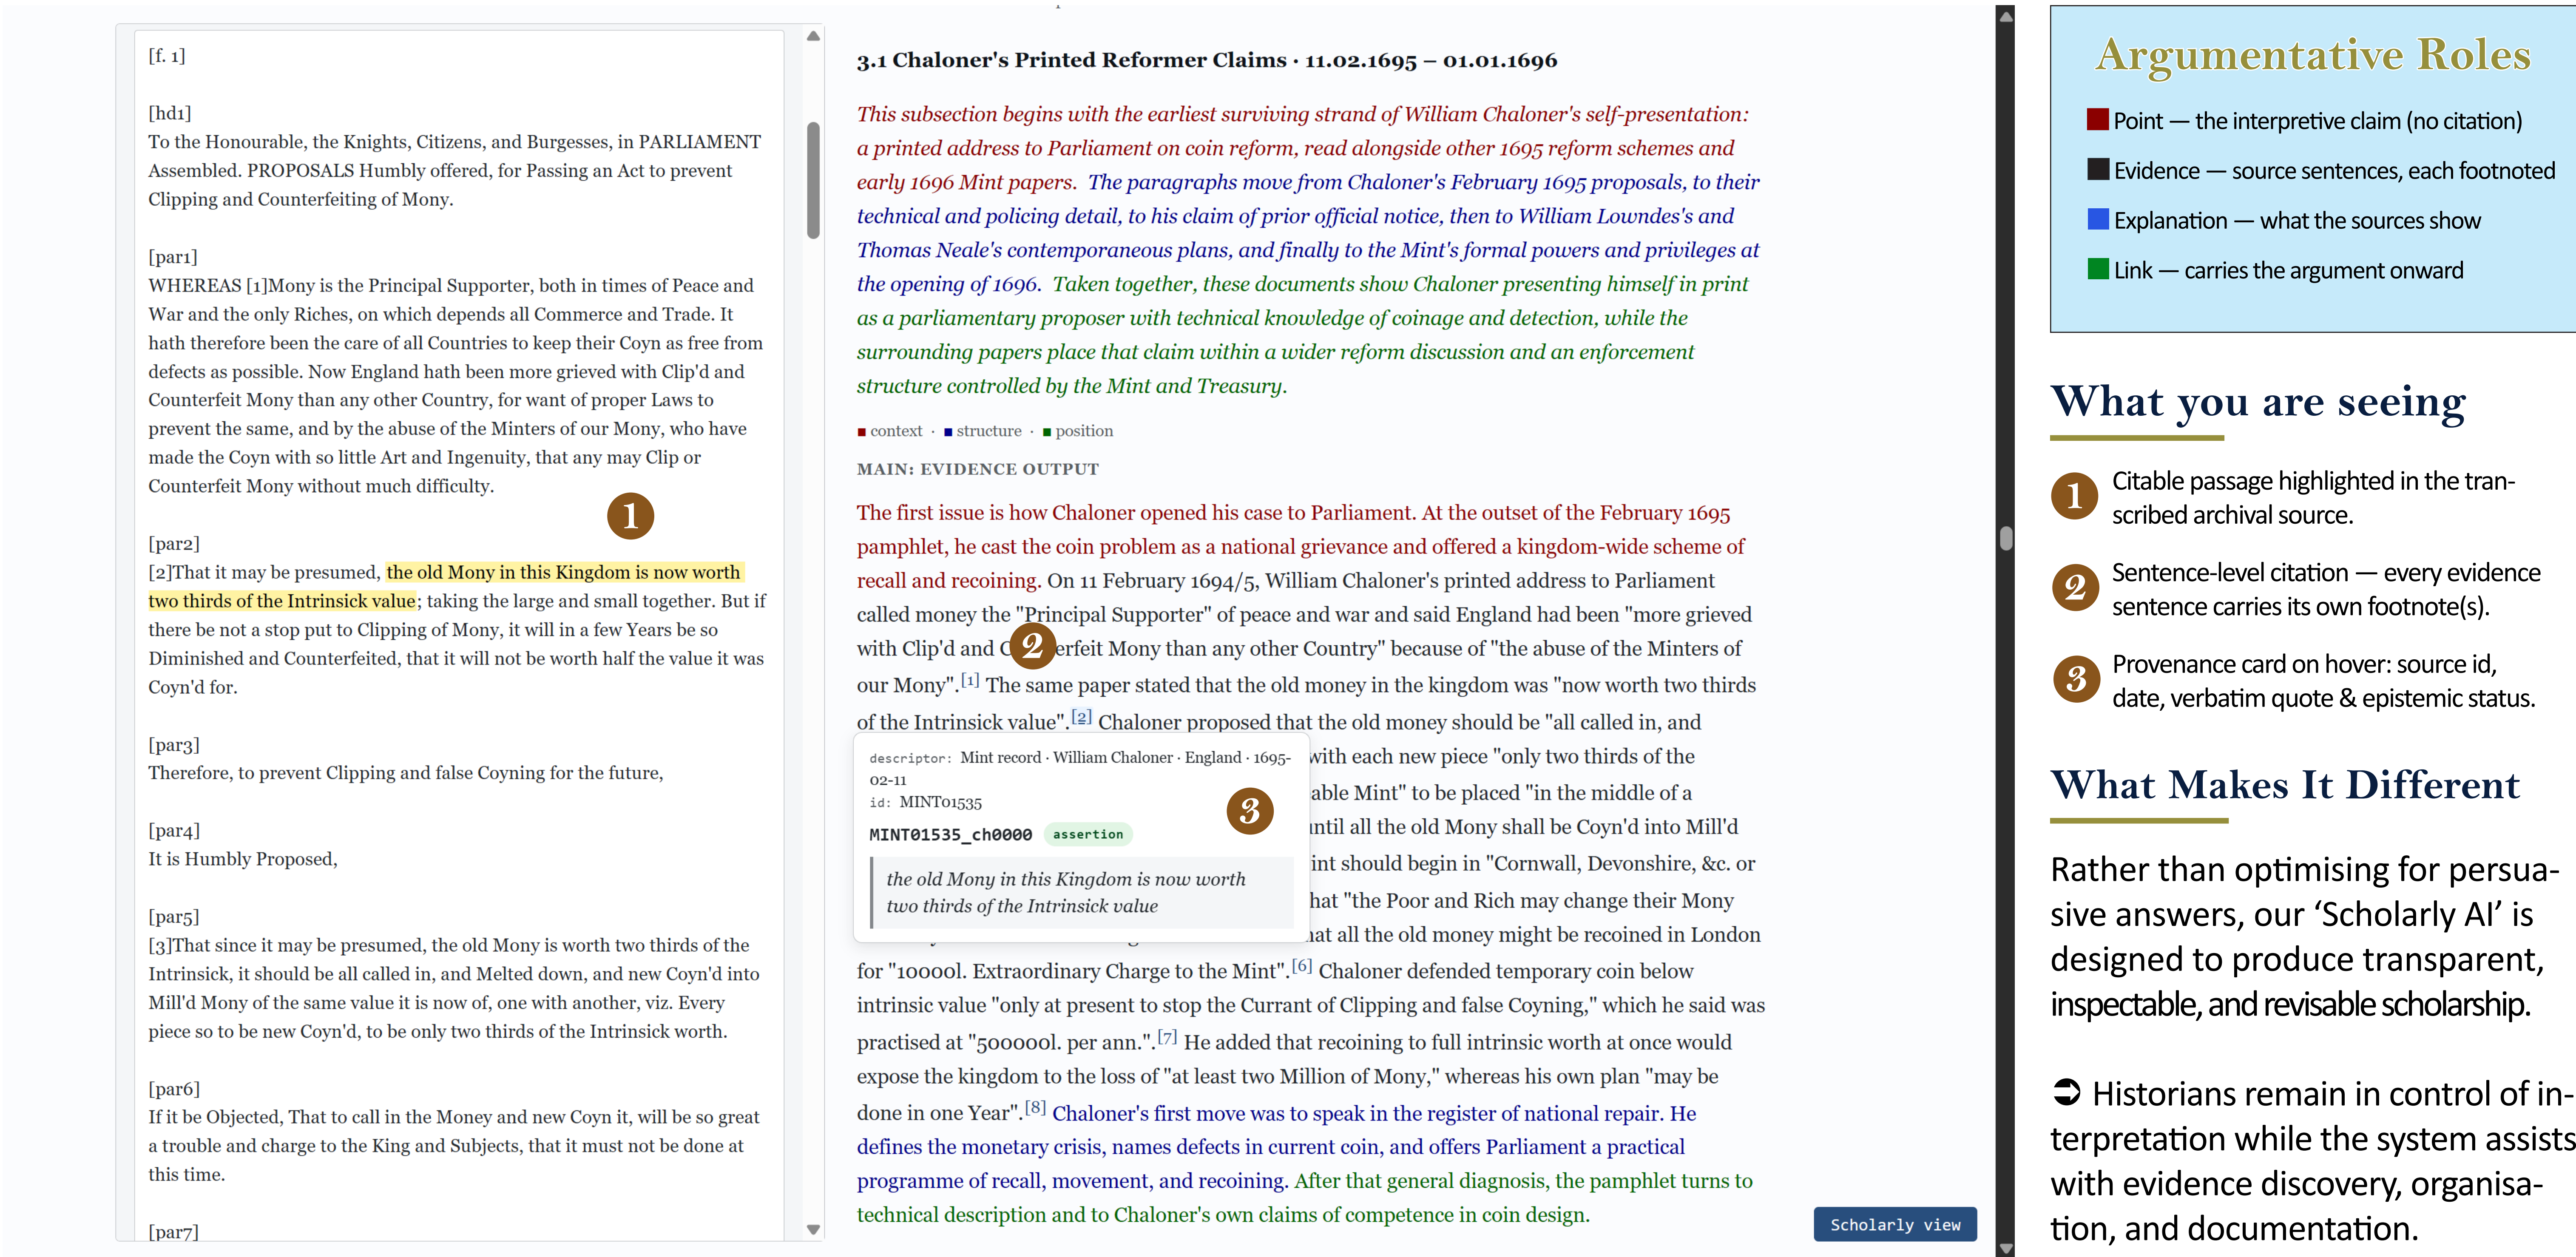


Figure 1. The Evidence Report for one subsection of the Chaloner Prosecution run (Royal Mint papers, 1692–1700). Left: transcribed source with a highlighted citable passage. Right: source-grounded synthesis in PEEL structure, where every claim is traceable to a specific retrieved document.

Our Epistemic Constitution

- I: Traceability — Every claim cites a specific source

II: Bounded agency — Synthesis only from retrieved evidence

III: Language fidelity — Source language reproduced verbatim

IV: Ambiguity preserved — Never silently resolved
- V: Contradiction shown — Conflicting sources kept side by side

VI: Accumulation — Every independent attestation stays visible

VII: Contextual continuity — Argument context carried forward

VIII: Revision & versioning — Cumulative, revisable research record

Collaborate With Us!

We are seeking collaborations with historians, archivists, and digital humanists who would like to evaluate the system using their own collections and help shape the future of trustworthy AI for historical research.

➡ Send us your own datasets and archival sources for processing, and watch our tool sort through your material!

What it produces — one research question

261	1,688	85,789	\$12.70
primary sources synthesised	sentence-level citations	words of source-grounded prose	computation cost

From a single research question the system assembled 261 primary sources into 5 sections / 27 subsections with 1,688 footnoted citations. Whole-run compute: ≈ \$12.70.

Our Workflow

Applying Large Language Models (LLMs) to Historical Source Corpora and Involving Domain Experts in the Process:

