

From Prompt Logs to Traceable Governance


A Provenance Application Profile for GenAI-Assisted Multilingual Thesaurus Localization

Core question

How can an AI-assisted KOS decision remain traceable
after the meeting, the prompt, and the review are over?

Shu-Jiun (Sophy) Chen
Institute of History and Philology, Academia Sinica


NKOS Workshop at DCMI2026
3 August 2026, Seoul



Images: National Palace Museum; The Met; UNESCO Memory of the World


Why this matters for multilingual KOS

Multilingual thesaurus localization depends on several forms of professional judgment.



**GenAI can supply useful candidates quickly.
The professional question becomes: what must be recorded for later review?**

What prompt logs are not enough



C02 motifs: why a PT/NPT decision needs provenance

One English concept, two legitimate Chinese scholarly senses

AAT concept boundary

distinct or separable design elements
usually decorative
distinguished from patterns

Chinese scholarly senses

紋飾 wen-shi · visible decorative form
母題 mu-ti · recurring or symbolic sense

Editorial decision

Earlier Chinese PT

母題

mu-ti



PT

紋飾

wen-shi



NPT retained

母題

mu-ti



Qianlong tea bowl with floral scroll holding shou character, National Palace Museum, Taipei

Governance question: Why did 紋飾 become the PT while 母題 remained an NPT?

Final KOS record preserves the outcome; provenance preserves the reasoning path.

L13 Paiban: historical connection is not equivalence



Paiban (拍板), The Metropolitan Museum of Art, public domain

Paiban / 拍板

AI surfaced candidate relations; expert reconfirmation assigned governed statuses.

Textual cue / AI candidate	Expert reconfirmation	Governed KOS status
chongdu 舂牘 historical name / access point	can be included	NPT controlled historical access point
du 牘 “adapted from” / predecessor relation	not interchangeable	historical predecessor / related concept not NPT

Governance question: Why did two historically related candidates receive different KOS statuses?

Vocabulary record preserves statuses; provenance preserves textual cue, review, rationale, and version change.

From two cases to three provenance requirements



C02 motifs
**Why did the PT/NPT
status change?**
concept boundary +
editorial rationale



L13 Paiban
**Why did candidate
relations diverge?**
textual cue + expert
reconfirmation + version change



**What must remain
visible after the final
KOS record?**



**Reconstruct the
conditions**
model · prompt · task
context

**Link evidence, criteria,
and roles**
sources · textual cues ·
editor / expert

**Explain decisions and
changes**
KOS status · rationale ·
version path

These requirements are translated into competency questions and AP record elements.

From provenance requirements to 18 competency questions

CQs specify what the provenance record and the AP must be able to answer.

Inputs

3 provenance requirements

conditions · evidence/roles · decisions

2 KOS workflows

AAT localization · local vocabulary

Workflow stages

task · evidence · AI · review · revision

Role-based questions

analytical perspectives

Derivation

Identify questions at
Task · Turn · Result

Merge overlaps
Separate case-specific
questions
Group by necessity

CQ architecture

Core-L1 · 10 CQs

minimum provenance closure

CQ1–3
Task

CQ4–5
Conditions

CQ6–7
Evidence

CQ8–10
Decision

Level-2 · 8 CQs

task-activated governance modules

CQ11–12
Quality

CQ13–15
Risk

CQ16–18
Version

The CQs drive AP record elements, executable queries, and minimum checks.

Prov-AI AP: a CQ-driven KOS application profile built on PROV-O

Reuse → specialize → constrain → validate

Prov-AI reuses PROV-O, then specifies a minimum accountable record for AI-assisted KOS decisions.

Standards reused

PROV-O backbone

Entity · Activity · Agent
used · generated · derived
associated with

Profile method

DCMI application profile style
selected terms + local constraints
human-readable release package

Supporting standards

SKOS · DCMI · SHACL
optional L2: DQV · Web Annotation

Reused operational structure

WorkItem
prov:Entity

TaskRun
prov:Activity

TurnRun
prov:Activity

Record roles around each AI turn
prompt · source/evidence · model response · decision note

Result
governed
candidate
or final output

Artifact
Prompt, source,
response, or
decision note

Agent → Association → Role
editor · expert · model · system

What KOS-specific

A minimum accountable KOS decision record

- **KOS concept + task context**
what concept is being reviewed, and for which task
- **Evidence + model response**
what evidence was available, and what AI produced
- **Editorial decision + responsibility**
what was decided, why, and by whom
- **CQ-driven queries and checks**
Can the record answer the required questions?

Prov-AI's contribution lies in the KOS-specific links and minimum requirements, not in the run units themselves.

How Prov-AI records the C02 motifs decision

One record connects the evidence, responsibility, decision, and rationale.

EARLIER RECORD 母題 = PT · 紋飾 = NPT → REVIEWED DECISION 紋飾 = PT · 母題 = NPT

CQ6 EVIDENCE

What supported the change?

- AAT scope note
- *motifs/patterns* boundary
- Three Chinese art-history sources

PROV-AI RECORD SourceEvidence

CQ9 RESPONSIBILITY

Who did what?

- AI expanded the candidate set.
- Editors interpreted the evidence and made the final decision.

PROV-AI RECORD Agent +
Association + Role

CQ10 DECISION + RATIONALE


What was decided—and why?

紋飾 → PT · visible decorative form
母題 → NPT · iconographic access term

PROV-AI RECORD Result +
DecisionNote

One record shows what supported the decision, who was responsible, and what was decided—and why.

Why is 紋飾 the PT—and 母題 the NPT?



The query returns the decision path—not only the final labels.

conditions → evidence and responsibility → decision and rationale

Release checks: did Prov-AI behave as intended?

3

internal cases
C02 · L13 · L08

30/30

3 cases × 10 core questions
Returned the expected answers

34/34

Pre-defined record checks
Correctly accepted, rejected, or flagged
5 accepted · 26 rejected · 3 flagged



L08: additional internal check case

Boundary: minimum release checks—not external validation.

Image: Suizhou Bianzhong of Marquis Yi of Zeng, UNESCO Memory of the World / Hubei Provincial Museum

From prompt logs to traceable governance

Prov-AI keeps the decision path visible.



A KOS term is more than a label. Provenance preserves why candidates were accepted, redirected, or rejected.

Next: Level-2 checks and broader testing

Thank you