



Designing for Disclosed Incompleteness: Data Governance Through a Cultural Constraint Engine for Heritage Image Generation

BIN ZHU

YIHAN KE

**Author affiliations can be found in the back matter of this article*

DISCUSSION PAPER

u[ubiquity press

ABSTRACT

Generative image and video systems can produce persuasive representations of intangible cultural heritage while silently introducing errors in culturally constitutive features. This discussion paper presents the Cultural Constraint Engine (CCE), a knowledge-base prototype developed around Foshan lion dance, and advances disclosed incompleteness as a data-governance principle for heritage-oriented generation. Six design requirements—retrievability, translatability, excludability, control recommendation, failure-frequency warning, and traceability—were informed by the frequency structure of 500 generation failures documented during an AI-assisted documentary production. CCE separates 51 cultural-meaning records from 495 production-specific rendering configurations, preventing implementation scale from being mistaken for cultural coverage. Record-level provenance and coding status are exposed through a live audit that reports source attachment, configuration coverage, semiotic classification, and unverified content. Independent coding of 47 records achieved 89.4% agreement; the reported Cohen's κ was .83. Clustered disagreements identified unresolved interpretive boundaries in decorative-pattern classification. A single-shot walkthrough demonstrates an integrated workflow from symbol retrieval and constraint construction to an attributable JSON package and pending generation record. The evidence establishes the architecture, internal relationships, and operation of the prototype for one case; it does not demonstrate improved cultural accuracy, usability, comprehensive coverage, community validation, or cross-domain transferability. The paper argues that heritage-generation tools should make evidential limits visible rather than conceal uncertainty behind plausible defaults. Linking cultural categories and semiotic codes with provenance and rendering options supports comparative interpretation and teaching while keeping evidential limits visible.

CORRESPONDING AUTHOR:

Yihan Ke

Faculty of Fine-Applied
Arts and Cultural Science,
Mahasarakham University,
Mahasarakham, Thailand
yihan.k@msu.ac.th

KEYWORDS:

cultural heritage; generative
artificial intelligence; data
governance; disclosed
incompleteness; knowledge
base

TO CITE THIS ARTICLE:

Zhu, B., & Ke, Y. (2026).
Designing for Disclosed
Incompleteness: Data
Governance Through a
Cultural Constraint Engine for
Heritage Image Generation.
*Journal of Open Humanities
Data*, 12: 141, pp. 1–11. DOI:
[https://doi.org/10.5334/
johd.653](https://doi.org/10.5334/johd.653)

(1) CONTEXT AND MOTIVATION

Generative text-to-image and text-to-video systems have become routine tools for producing visual representations of intangible cultural heritage. Museums, cultural agencies, students, and independent creators increasingly use them to depict practices such as lion dance, paper-cutting, and temple ritual (Foka & Griffin, 2024). Although these systems generate visually coherent and persuasive outputs, visual fluency does not ensure cultural accuracy. A generated image may confidently assign a lion head the wrong colour, social role, or ritual setting without indicating that the representation is uncertain or incorrect.

Foshan lion dance, a Cantonese performance tradition included in China's national intangible cultural heritage list (Qiu et al., 2024; Xie, 2020), makes this problem especially visible. Its lion heads are not merely decorative but encode culturally specific roles and temperaments. For example, golden-yellow and white-bearded, red-and-black and long-bearded, and black-white-cyan and short-bearded forms are conventionally associated with Liu Bei, Guan Gong, and Zhang Fei, respectively (Zhan & Wan Abdullah Thani, 2024a, 2024b). Performance settings and ritual procedures are likewise meaningful: architectural context, eye-dotting (dian jing), and plucking the greens (cai qing) carry conventions whose sequence and etiquette are not interchangeable (Xia, 2020). These relationships between colour, form, character, setting, and procedure may be readily recognised by knowledgeable practitioners, yet generative systems do not independently determine whether such combinations are culturally appropriate.

This study arose from the production of an AI-assisted documentary short on Foshan lion dance. Field observations, expert interviews, production spreadsheets, dissertation materials, and iterative corrections of failed generations yielded a substantial body of cultural and technical knowledge. However, that knowledge remained fragmented and difficult to use within generative workflows. It was not consistently searchable by cultural symbol, did not clearly distinguish constitutive features from optional rendering choices, and left no durable record of which claims had been source-verified or applied during a generation attempt.

Existing work addresses parts of this problem without fully resolving it. Heritage informatics provides rigorous structures for representing cultural knowledge, but such structures are generally designed for documentation, retrieval, or inference (Doerr, 2003; Hou & Kenderdine, 2024; Sartini et al., 2023) rather than prompt-based generation. Conversely, controllable-generation research offers increasingly precise ways to govern visual form but does not ordinarily represent whether the controlled form is culturally appropriate. The unresolved issue is therefore both technical and epistemic: how can curated heritage knowledge be translated into actionable generation constraints while making the knowledge base's coverage, verification status, and uncertainty inspectable?

To address this gap, the paper presents the Cultural Constraint Engine (CCE) as a design case study organised around the principle of disclosed incompleteness. Rather than concealing missing knowledge through plausible defaults, CCE makes unverified, uncoded, and unrated content visible to users. The study derives six design requirements from 500 documented generation failures, implements a two-layer knowledge base comprising 51 cultural-meaning records and 495 rendering configurations, and demonstrates an auditable workflow culminating in an exportable JSON constraint package. These materials establish the system's design rationale and integrated operation, but they do not demonstrate improved cultural accuracy, usability, comprehensive coverage, or transferability beyond the single heritage case.

(2) RELATED WORK

Heritage information systems formally represent cultural knowledge through ontologies such as CIDOC-CRM, supporting interoperable documentation, semantic querying, and inference (Doerr, 2003; Hou & Kenderdine, 2024; Sartini et al., 2023). Whereas these systems prioritise retrieval and reasoning, generative systems consume prompts and limited control signals. CCE adopts typed fields and record-level provenance but prioritises operational use in generation workflows rather than ontological completeness.

Prompt libraries organise reusable descriptors for style, composition, viewpoint, and visual quality (Pektaş & Sağlam, 2025), while diffusion models, structural conditioning, and adapter-based fine-tuning increasingly control generated form. Heritage applications include paper-cutting, historical façades, and Chinese landscape painting (Dai et al., 2024; Kuang et al., 2025; Peng et al., 2024), alongside work evaluating artefact fidelity and digital visualisation as evidence (Stejskal, 2025). These methods govern rendering without determining whether a form, gesture, or setting is culturally appropriate. CCE therefore separates constitutive constraints from optional rendering choices.

Creativity-support research evaluates tools through exploration, expressiveness, immersion, and human–AI agency (Shneiderman, 2007). CCE supports creative production but is not presented as a usability-tested creativity-support system. Its rationale instead draws on human–computer interaction research concerning system-state visibility, intelligibility, calibrated trust, explanation, and meaningful human control (Bellotti & Edwards, 2001; Lee & See, 2004; Miller, 2019; Shneiderman, 2020). This literature suggests that systems should communicate evidential limits rather than encourage confidence exceeding their reliability. CCE applies this principle through provenance indicators, explicit unrated states, and continuously updated coverage audits.

CCE is not implemented as conventional retrieval-augmented generation. It retrieves culturally typed records through explicit symbol selection, preserves standing negative constraints, and exposes source and verification states. This comparison locates the system architecture without implying that comparable governance mechanisms could not be incorporated into other retrieval-based designs.

Finally, research on hallucination, representational bias, sequence consistency, and the semiotics of generated imagery demonstrates that plausible outputs may remain unsupported or culturally distorting (Foka & Griffin, 2024; Ji et al., 2023; Putland et al., 2025). CCE accordingly treats generated images as consequential representations rather than neutral prompt translations (Leone, 2024; Peirce, 1931–1958; Short, 2007). Its contribution is not a new theory of transparency or cultural authenticity, but a heritage-specific design proposition: generation constraints should remain linked to inspectable evidence, while absent or contested knowledge remains visibly incomplete.

(3) DATASET DESCRIPTION

The deposited dataset documents the construction and governance of the Cultural Constraint Engine (CCE), developed for heritage-informed image generation using Foshan lion dance as a bounded case. The Version 5 public file visual-dictionary-data.js contains 51 cultural-meaning records in SYMBOL_LAYER and 495 rendering configurations in VARIANT_LAYER, matching the corresponding full working totals of 51 and 495. These are counts of structured records, not comprehensive cultural coverage. Public release of these records does not imply release of all underlying field materials or coding worksheets. Supporting application, provenance, audit, and pipeline files accompany the records; materials whose community-knowledge permissions remain unresolved are excluded.

REPOSITORY LOCATION

The dataset is openly available through Figshare (Zhu & Ke, 2026) at <https://doi.org/10.6084/m9.figshare.33137087>.

REPOSITORY NAME

Figshare.

OBJECT NAME

Dataset for “Designing for Disclosed Incompleteness: A Cultural Constraint Engine for Heritage-Informed Image Generation”. The deposited object consists of the permission-cleared symbol-layer and configuration-layer records, accompanied by application files, audit and provenance files, and pipeline documentation.

JavaScript data and application files, JSON provenance and risk files, a CSV conformance audit, HTML, Python scripts, and Markdown documentation; Figshare Version 5.

CREATION DATES

2025-07-01 to 2026-05-30.

DATASET CREATORS

Bin Zhu, Mahasarakham University: conceptualisation; investigation; field-based knowledge collection; failure-log compilation; dataset design; data curation; software development; and preparation of the symbol and configuration records. Yihan Ke, Mahasarakham University: supervision; validation; independent cross-coding; quality review; project administration; and review of the dataset documentation.

LANGUAGE

English is used for variable names, field labels, coding categories, technical documentation, and generation-oriented constraint descriptions. Chinese terms and romanised Cantonese or Mandarin expressions are retained where required to identify culturally specific symbols, roles, rituals, architectural features, or inscriptions. English explanations are provided to support interpretation and reuse.

LICENSE

Creative Commons Attribution 4.0 International (CC BY 4.0). Users may share and adapt the deposited materials provided that appropriate attribution is given, a link to the licence is supplied, and any changes are indicated. The licence applies only to the openly deposited, permission-cleared materials and does not remove ethical responsibilities relating to community-held cultural knowledge.

PUBLICATION DATE

The dataset was first published on 2026-08-01; Version 5 was released on 2026-08-06.

(4) METHOD

This study used a single-case design to examine the construction and operation of the Cultural Constraint Engine (CCE), a software implementation of the Culturally Constrained Iconographic Generation Framework. The case was bound to materials assembled during production of an AI-assisted documentary short about Foshan lion dance. The study evaluates whether the proposed data-governance architecture is constructible and internally coherent; it does not test usability, comparative effectiveness, or reductions in cultural inaccuracy.

The first empirical material comprised 500 classified generation-deviation entries recorded during documentary production. Each entry represented an attempt judged unusable and assigned one primary outcome category: cultural-accuracy, physical or motion-dynamics, or general visual-quality failure. Usable outputs were not entered, so 500 is not the total number of generation requests and cannot serve as a denominator for an overall failure rate. The 251 cultural-accuracy entries were further classified into five mutually exclusive sub-patterns. Because the log originated as a production record, entries were accumulated according to production needs rather than sampled. Classification was retrospective and single-coded, while dates, complete prompts, and per-attempt model versions were not consistently retained. The resulting counts describe the composition of one failure log, not failure probabilities or population estimates.

The cultural-symbol layer was compiled from published literature, the first author's field observations in Foshan, and semi-structured interviews with ten informants representing lion dance practice, coaching, lion-head making, heritage research, folklore, documentary

and animation production, and digital media. Records were selected according to production requirements rather than from a defined population, producing 51 cultural-meaning records. A configuration layer was generated for the 15 symbols rendered during production by combining documented symbols with predefined visual variants, producing 495 configurations. Each record stored structured provenance, coding, verification, and status fields. Source attachment indicates documented provenance, not community ratification. The fieldwork string is inherited from one dataset-level campaign and should not be read as a separate observation log for every symbol.

All 51 symbol records were assigned a dominant Peircean semiotic aspect—icon, index, or symbol (Peirce, 1931–1958; Short, 2007)—by the first author. The second author independently coded 47 records after four examples used for training were excluded. Dominant-aspect agreement was 89.4% (42/47); the externally computed Cohen’s κ was reported as .83. The five disagreements were retained and analysed rather than reconciled. The coefficient is reported as a numerical value rather than mapped onto conventional verbal bands (Landis & Koch, 1977). Reliability estimates apply only to the dominant-aspect field. The dataset metadata records that 19 initial codes were tool-proposed and confirmed record by record by the first coder on 28 July 2026, before the blind second pass on 30 July; 18 fall within the 47-record comparison. The public files do not contain the per-record second-coder ratings needed to recompute agreement, so the reported statistics remain dependent on the external coding worksheets.

CCE was implemented as a client-side single-page application using HTML, CSS, and vanilla JavaScript. Static JavaScript files contain the data and version metadata; audit measures are recalculated whenever the application loads. A single-shot walkthrough used the symbol with the highest documented failure count and the spatial context with the second-highest count. It demonstrated integrated operation and JSON export but was not an evaluation. The Version 5 repository file listing includes prototype application source files (index.html, app.js, and traceability.js), data files, and supporting audit scripts. These deposited files accompany the prototype description; their presence does not establish independent replication or effectiveness.

(5) RESULTS AND DISCUSSION

(5.1) FAILURE PATTERNS AND DESIGN REQUIREMENTS

Of the 500 documented failure entries, 251 (50.2% of logged failures) were classified as cultural-accuracy failures, 130 as physical or motion-dynamics failures, and 119 as general visual-quality failures. Five sub-patterns accounted for the cultural-accuracy category: Liu Bei lion morphology or role-colour deviation (78 instances), spatial or architectural misplacement (61), generalisation of historical or memory-bearing symbols into non-specific “oriental” decoration (53), percussion or sound-cue mismatch (31), and incorrect textual inscriptions (28).

These patterns informed six requirements: retrievability, translatability, excludability, control recommendation, failure-frequency warning, and traceability. Their evidential bases were not equivalent. Translatability, excludability, and warning were directly associated with recorded high-frequency patterns. Retrievability was a logical precondition rather than an independently counted failure. Control recommendation was motivated by an unquantified subset of cases in which textual correction did not prevent recurrence. Traceability responded to inscription and sound-cue errors, although the proposition that incomplete records caused recurrence remains an interpretation of the production history. The failure log therefore provides an empirical anchor for the design, not proof that the resulting requirements improve generation outcomes.

The combined 114 instances of spatial misplacement and historical-symbol generalisation showed why cultural knowledge must be translated from narrative descriptions into concrete visual constraints. The 78 role-identity failures additionally showed the need for explicit exclusions and advance warnings. However, the recorded values are counts rather than rates because the number of attempts per symbol or scene was unavailable. CCE consequently labels entries by observed failure frequency—high (≥ 60), medium (25–59),

low (1–24), or unrated—rather than estimating risk. “Unrated” denotes absent evidence, not safety. [Table 1](#) summarises the failure patterns and corresponding design requirements.

REQUIREMENT	OPERATIONAL PURPOSE	ASSOCIATED PATTERN	DOCUMENTED INSTANCES
R1 Retrievable	Make cultural records searchable	Precondition for R2–R6	Not independently counted
R2 Translatable	Convert cultural meaning into visual constraints	Spatial/context errors; symbol generalisation	114 (61 + 53)
R3 Excludable	Express known-wrong renderings as prohibitions	Role-colour and morphology deviation	78
R4 Control-recommending	Recommend visual control when text is insufficient	Unquantified subset of morphology/context errors	Not isolated
R5 Failure-frequency warning	Display observed recurrence before generation	Role-colour and morphology deviation	78
R6 Traceable	Record knowledge, settings, and outcomes	Inscription and sound-cue errors	59 (28 + 31)

Table 1 Failure patterns and corresponding design requirements.

(5.2) KNOWLEDGE-BASE STRUCTURE AND AUDIT RESULTS

The implemented knowledge base separates cultural meaning from rendering variation. Its symbol layer contains 51 records describing meaning, constitutive visual features, historical context, provenance, and coding status. Its configuration layer contains 495 rendering records, generated from 15 symbols through 33 configuration slots per symbol. These comprise three base styles and combinations of six palettes with five ornament sets. This distinction corrects an earlier single-layer representation in which 495 rows could have been mistaken for 495 independently documented cultural facts.

Only 15 of the 51 symbols have configurations; the remaining 36 (70.6%) contain cultural knowledge without production-specific rendering variants. Configuration coverage is therefore 29.4%. The figure of 495 represents implementation scale, not breadth of cultural documentation. Separating the layers also distinguishes features that must remain culturally stable from stylistic variables that may legitimately change.

The live audit reported non-empty bibliographic source lists for 24 of 51 records, while 50 carry a legacy provenance/verification flag, first-pass dominant-aspect coding for all 51, and independent blind coding for 47. Agreement on the dominant Peircean aspect was 89.4%, with an externally reported Cohen’s κ of .83; the public files do not permit independent recomputation. The five disagreements were clustered within Form/Pattern records. The first coder emphasised resemblance and classified these motifs as icons, whereas the second emphasised meanings assigned through homophony or cosmology and classified them as symbols. This pattern identifies a boundary in the operational definitions rather than simply random coding error. The secondary-aspect field was not independently assessed. Here, category denotes a cultural domain, such as Field, Sound, Form, or Ritual, whereas semiotic aspect denotes the dominant icon, index, or symbol relation. These fields offer complementary descriptions; neither is a measure of evidential authority.

The audit further showed that the semiotic aspect was completely determined by category for Field and Sound, varied within Form and Ritual, and could not be evaluated in four categories containing too few records. These results do not establish independence between category and aspect. More importantly, attached provenance must not be equated with cultural validation. The legacy “verified” flag on 50 records denotes recorded provenance rather than independent cultural validation; only 24 contain explicit bibliographic references. No record documents the stronger states of cross-source corroboration, practitioner review, or community validation. Making those distinctions visible prevents a field observation, an informant statement, and

a community-ratified account from being presented with equivalent authority. For users, “source-attached” means that a claim can be traced to an identified source, not that it has been independently confirmed. Cross-source corroboration would require supporting evidence from additional sources, practitioner review would require documented assessment by relevant practitioners, and community validation would require a documented community review process with its scope made explicit. These states should not be inferred from coding agreements or treated as universal certification of cultural authenticity. [Table 2](#) summarises the knowledge-base coverage and coding status.

METRIC	VALUE	INTERPRETIVE NOTE
Symbol-layer records	51	Units of cultural meaning
Rendering configurations	495	15 symbols × 33 slots; implementation scale
Records with bibliographic references	24/51	Non-empty source lists; not community validation
Blind-coded records	47/51	Four training examples excluded
Dominant-aspect agreement	89.4%; $\kappa = .83$	Reported externally; paired ratings not deposited
Symbols with configurations	15/51 (29.4%)	Production-driven coverage
Symbols without configurations	36/51 (70.6%)	Cultural records without rendering variants

Table 2 Knowledge-base coverage and coding status.

(5.3) SYSTEM OPERATION AND WORKED DEMONSTRATION

CCE implements nine functions across four interface areas: the Visual Dictionary, Constraint Engine, Dataset Audit, and Generation Log. The Visual Dictionary retrieves cultural and rendering records with their provenance. The Constraint Engine constructs positive prompts, supplies typed negative constraints, recommends control methods, displays failure-frequency evidence, and exports a structured package. The Dataset Audit recalculates coverage and coding measures from the current records without modifying them. The Generation Log registers runs and returns completed outcomes to the relevant symbol record. [Figure 1](#) summarises the CCE architecture and auditable workflow.

The negative-constraint component organises recurrent errors into six classes: morphological or species confusion; role-identity drift; spatial or architectural drift; ritual or procedural error; ornamental or chromatic drift; and stylistic or semantic distortion. Record-specific exclusions are added separately. This taxonomy improves inspectability over a flat list, but its classes are not equally supported by the log. In particular, ritual entries such as eye-dotting and plucking the greens lacked linked failures and were consequently changed from judgement-based risk labels to “unrated”. The ritual-error class therefore remains a design precaution rather than a log-derived finding.

The worked demonstration combined the Liu Bei lion, the most frequent cultural-failure pattern, with the ancestral-hall setting, the second most frequent. Retrieval returned record FLD-S-030 with source status and provenance. The system constructed positive image and video prompts, applied shared and symbol-specific exclusions, recommended reference imagery, LoRA, and Canny-edge control, and assigned a high failure-frequency flag based on 78 documented instances across five linked records.

Export produced a JSON constraint package containing the symbol record, prompts, negative-constraint classes, settings, provenance, coding and source status, control recommendation, engine and dictionary versions, dataset fingerprint, and a unique run identifier. Model, parameter, outcome, and reviewer fields remained visibly pending until supplied by the operator. This establishes that the nine functions operate as an integrated pipeline for the demonstrated input. It does not validate their performance across symbols or show that the resulting prompt produces a more culturally accurate image.




Figure 1 Cultural Constraint Engine architecture and auditable workflow. Solid arrows show operational flow and read-only audit inputs; the dashed arrow shows reviewed performance-history write-back.

(5.4) INTERPRETATION AND LIMITATIONS

The design case demonstrates that disclosed incompleteness can be implemented within the CCE prototype. CCE does not determine cultural authenticity; it exposes the evidence needed for human judgement, including what is constitutive, what is sourced, what has previously failed, and what remains unknown. In this respect, the audit is not an auxiliary administrative feature but a principal interface alongside generation controls. This design applies transparency and calibrated-trust principles to heritage-oriented generative workflows (Bellotti & Edwards, 2001; Lee & See, 2004; Miller, 2019; Shneiderman, 2020).

The evidence nevertheless remains limited. No controlled usability or effectiveness study was conducted, and the single walkthrough cannot demonstrate reduced cultural error or superiority to careful manual practice. The architecture was developed from one production and one heritage tradition. Its two-layer separation, explicit status vocabulary, and live audit may be transferable, but this has not been tested. The failure taxonomy, frequency bands, and 33-slot expansion are case-specific and should not be reused for other traditions without reconstruction from relevant evidence. The potentially transferable principles concern how evidence is organised and disclosed, rather than which cultural combinations are permissible. For example, separating cultural claims from rendering choices and retaining provenance and unresolved states may guide another project, whereas lion-role colour conventions, ritual sequences, and local architectural constraints require tradition-specific evidence and review.

Independent coding and source attachment also do not amount to community validation. The sources and informants were selected within the authors' project, and no independent panel of lineage holders or performers ratified the records. Moreover, a knowledge base may inadvertently stabilise a living tradition as a fixed authorised account (Smith, 2006; UNESCO, 2003), despite evidence that heritage authenticity is relational and context-dependent (Zhu et al., 2025). The 51 records should therefore be treated as dated, revisable representations rather than a definitive standard.

Finally, the browser-local run register is installation-specific rather than collaborative, and outcome write-back updates performance history without altering cultural content. Cultural revision remains a human editorial responsibility. Accordingly, the study establishes the existence, structure, and internal operation of CCE, while leaving usability, comparative performance, comprehensive coverage, community validation, and cross-domain transferability as matters for future research.

(6) IMPLICATIONS/APPLICATIONS

Cultural heritage generation systems should treat the disclosure of knowledge limitations as a core function. CCE implements this principle by distinguishing cultural meaning, rendering variation, provenance, coding status, observed failure frequency, and unresolved gaps. Performance controls should therefore be paired with mechanisms showing what a system knows, what supports that knowledge, and where human judgement remains necessary, especially when persuasive outputs may circulate as authoritative representations of living traditions.

Heritage institutions, museums, cultural agencies, and creative practitioners could adapt the two-layer architecture to organise production knowledge. A cultural-meaning layer can record constitutive features and their provenance, while a separate configuration layer can define permissible stylistic variation for exhibitions, educational media, documentary production, or public communication. The six requirements derived in this case—retrievability, translatability, excludability, control recommendation, failure-frequency warning, and traceability—provide questions that other projects can use when designing their own systems. However, the Foshan-specific failure counts, negative-constraint classes, and rendering configurations should not be transferred unchanged. Each heritage practice requires its own source base, community participation, failure evidence, and rules governing legitimate variation.

The dataset and exported JSON structure also have research and technical applications. Researchers may use the records to examine the separation between cultural claims and rendering instructions, evaluate provenance and status vocabularies, compare alternative semiotic coding schemes, or develop governance-oriented retrieval and prompting systems. The structured package could serve as an interchange format linking heritage databases, prompt-generation interfaces, image-generation tools, and evaluation records. The live-audit approach may also inform other small, expert-curated datasets by replacing static claims of completeness with measures recalculated from the current data. Reuse must nevertheless preserve attribution, verification status, and explicit “unrated” states; missing evidence should not be recorded as low risk or cultural permissibility.

For an illustrative reuse exercise, a researcher can retrieve the publicly deposited Liu Bei Lion record FLD-S-030, compare its golden-yellow face and white-beard descriptors with its exclusions, and inspect its references, informant codes, and dominant Icon and secondary Symbol codes. The record has no linked configuration-layer variants (`variantCount = 0`), but its image and video prompts demonstrate how cultural descriptions are translated into rendering instructions. Its high frequency band refers to 78 failures across a five-record pattern, not 78 failures for this symbol alone. This exercise supports comparison and teaching without establishing cultural validity or generation effectiveness.

For education, the dataset can support teaching about digital heritage documentation, prompt design, data provenance, semiotic classification, and the responsible use of generative AI. Students can inspect how narrative cultural knowledge is translated into positive descriptors, negative exclusions, and control recommendations, while also identifying where such translation simplifies or stabilises contested cultural meanings. The disagreement between coders over decorative motifs provides a useful case for discussing interpretive plurality rather than treating classification as mechanically objective.

These applications remain prospective. Before operational adoption, CCE requires practitioner and community review, controlled usability testing, cross-case replication, and evaluation of whether exported constraints affect generated outputs. The system should support revision and contestation rather than become a fixed authority over authenticity. Disclosed incompleteness is therefore a baseline governance commitment: it cannot guarantee cultural accuracy, but it keeps uncertainty visible.

AI DECLARATION

Generative text-to-image and text-to-video systems were used in the production context from which the generation-failure records were derived. Their outputs were treated as production artefacts, not as sources of cultural authority. Generative AI was not used to conduct interviews, establish cultural ground truth, calculate agreement, or validate records. OpenAI’s ChatGPT assisted with language editing, condensation, structural organisation, and alignment with the Discussion Paper template. The authors reviewed and verified all AI-assisted text and remain responsible for the manuscript’s accuracy, originality, citations, interpretations, and conclusions.

DATA ACCESSIBILITY STATEMENT

The permission-cleared data and supporting files, including pipeline documentation are available through Figshare (Zhu & Ke, 2026) at <https://doi.org/10.6084/m9.figshare.33137087>. Material for which community-knowledge permissions remain unresolved is not included in the public deposit.

The underlying doctoral research was reviewed by the Human Research Ethics Committee of Thepsatri Rajabhat University and determined to be exempt from full ethical review (TRU-EC 29/2026). Informants provided informed consent and were anonymised as E01–E10. Only permission-cleared records were included in the public dataset.

ACKNOWLEDGEMENTS

This research project was financially supported by Mahasarakham University. The authors gratefully acknowledge the lion dance practitioners, coaches, lion-head makers, heritage researchers, folklore specialists, and media professionals who contributed contextual knowledge to the wider doctoral project. The authors also thank Mahasarakham University and Thepsatri Rajabhat University for institutional and administrative support.

FUNDING INFORMATION

This research project was financially supported by Mahasarakham University (Grant No. Res-In70020069).

COMPETING INTERESTS

The authors have no competing interests to declare.

AUTHOR CONTRIBUTIONS

Bin Zhu: Conceptualisation, Data curation, Formal analysis, Investigation, Methodology, Software, Visualisation, and Writing – original draft.

Yihan Ke: Project administration, Supervision, Validation, and Writing – review & editing.

AUTHOR AFFILIATIONS

Bin Zhu

Faculty of Fine-Applied Arts and Cultural Science, Mahasarakham University, Mahasarakham, Thailand

Yihan Ke

Faculty of Fine-Applied Arts and Cultural Science, Mahasarakham University, Mahasarakham, Thailand

REFERENCES

- Bellotti, V., & Edwards, K. (2001). Intelligibility and accountability: Human considerations in context-aware systems. *Human-Computer Interaction*, 16(2–4), 193–212. https://doi.org/10.1207/S15327051HCI16234_05
- Dai, M., Feng, Y., Wang, R., & Jung, J. (2024). Enhancing the digital inheritance and development of Chinese intangible cultural heritage paper-cutting through Stable Diffusion LoRA models. *Applied Sciences*, 14(23), Article 11032. <https://doi.org/10.3390/app142311032>
- Doerr, M. (2003). The CIDOC Conceptual Reference Module: An ontological approach to semantic interoperability of metadata. *AI Magazine*, 24(3), 75–92. <https://doi.org/10.1609/aimag.v24i3.1720>
- Foka, A., & Griffin, G. (2024). AI, cultural heritage, and bias: Some key queries that arise from the use of GenAI. *Heritage*, 7(11), 6125–6136. <https://doi.org/10.3390/heritage7110287>
- Hou, Y., & Kenderdine, S. (2024). Ontology-based knowledge representation for traditional martial arts. *Digital Scholarship in the Humanities*, 39, 575–592. <https://doi.org/10.1093/llc/fqae005>
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55, Article 248. <https://doi.org/10.1145/3571730>
- Kuang, Z., Zhang, J., Li, Y., & Fukuda, T. (2025). Preserving architectural heritage in urban renewal: A Stable Diffusion model framework for automated historical facade generation. *npj Heritage Science*, 13, Article 256. <https://doi.org/10.1038/s40494-025-01826-4>
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174. <https://doi.org/10.2307/2529310>

- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. <https://doi.org/10.1518/hfes.46.1.50.30392>
- Leone, M. (2024). Semiotics of the black box: On the rhetorics of algorithmic images. *Visual Communication*, 23, 426–451. <https://doi.org/10.1177/14703572241247120>
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38. <https://doi.org/10.1016/j.artint.2018.07.007>
- Peirce, C. S. (1931–1958). *Collected papers of Charles Sanders Peirce* (C. Hartshorne, P. Weiss, & A. W. Burks, Eds.; Vols. 1–8). Cambridge, MA: Harvard University Press.
- Pektaş, Ş. T., & Sağlam, B. (2025). Semiotics-based prompt engineering for architectural text-to-image generation processes. *Estoa*, 14(28), 121–135. <https://doi.org/10.18537/est.v014.n028.a09>
- Peng, X., Hu, Q., Fan, F., Xie, P., Zhang, Y., & Cao, R. (2024). FHS-Adapter: Fine-grained hierarchical semantic adapter for Chinese landscape paintings generation. *Heritage Science*, 12, Article 265. <https://doi.org/10.1186/s40494-024-01370-7>
- Putland, E., Chikodzore-Paterson, C., & Brookes, G. (2025). Artificial intelligence and visual discourse: A multimodal critical discourse analysis of AI-generated images of “dementia.” *Social Semiotics*, 35, 228–253. <https://doi.org/10.1080/10350330.2023.2290555>
- Qiu, L., Abdul Rahman, A. R., & Bin Dolah, M. S. (2024). Exploring the evolution and cultural significance of Guangdong lion dance in Chinese traditional culture. *International Journal of Academic Research in Business and Social Sciences*, 14(2), 2131–2143. <https://doi.org/10.6007/IJARBS/v14-i2/20905>
- Sartini, B., Baroncini, S., van Erp, M., Tomasi, F., & Gangemi, A. (2023). ICON: An ontology for comprehensive artistic interpretations. *ACM Journal on Computing and Cultural Heritage*, 16, Article 51. <https://doi.org/10.1145/3594724>
- Shneiderman, B. (2007). Creativity support tools: Accelerating discovery and innovation. *Communications of the ACM*, 50(12), 20–32. <https://doi.org/10.1145/1323688.1323689>
- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe and trustworthy. *International Journal of Human–Computer Interaction*, 36(6), 495–504. <https://doi.org/10.1080/10447318.2020.1741118>
- Short, T. L. (2007). *Peirce’s theory of signs*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511498350>
- Smith, L. (2006). *Uses of heritage*. London: Routledge. <https://doi.org/10.4324/9780203602263>
- Stejskal, J. (2025). Art-historical empiricism and digital visualization of cultural heritage. *Synthese*, 205, Article 132. <https://doi.org/10.1007/s11229-025-04959-y>
- UNESCO. (2003). *Convention for the safeguarding of the intangible cultural heritage*. Paris: UNESCO. Retrieved from <https://ich.unesco.org/en/convention> (last accessed: 6 September 2026).
- Xia, Z. W. (2020). 醒狮表演的礼节研究. *北京舞蹈学院学报*, (1), 77–82. <https://doi.org/10.3969/j.issn.1008-2018.2020.01.011>. Access via <https://d.wanfangdata.com.cn/periodical/CiBQZXJpb2RyY2FsQ0hJU29scjkyMDI2MTAwOTA4MDAzMRIRYmp3ZHh5eGlyMDIwMDEwMTEaCGpnaHd0amd5>
- Xie, Z. Y. (2020). 非遗传承主体存续的文化社会基础——对佛山醒狮习俗的历史考察. *湖北民族大学学报 (哲学社会科学版)*, 38(1), 133–142. <https://doi.org/10.13501/j.cnki.42-1328/c.2020.01.015>
- Zhan, S., & Wan Abdullah Thani, W. J. I. (2024a). A study on the colour elements and symbols of Foshan awakening lion head. *International Journal of Art & Design*, 8(2), 32–41. <https://doi.org/10.24191/ijad.v8i2.1960>
- Zhan, S., & Wan Abdullah Thani, W. J. I. (2024b). Conveying spirit through form: An exploration on the shape symbols of Foshan awakening lion head. *Art and Society*, 3(1), 89–100. <https://doi.org/10.56397/AS.2024.02.08>
- Zhu, B., & Ke, Y. (2026). *Dataset for “Designing for disclosed incompleteness: A cultural constraint engine for heritage-informed image generation”* (Version 5) [Data set]. Figshare. <https://doi.org/10.6084/m9.figshare.33137087>
- Zhu, P., Chi, X., Ryu, H. B., & Han, H. (2025). Experience economy and authenticity in the heritage tourism sector: A multiple-dimensional approach. *Acta Psychologica*, 257, Article 105118. <https://doi.org/10.1016/j.actpsy.2025.105118>

TO CITE THIS ARTICLE:

Zhu, B., & Ke, Y. (2026). Designing for Disclosed Incompleteness: Data Governance Through a Cultural Constraint Engine for Heritage Image Generation. *Journal of Open Humanities Data*, 12: 141, pp. 1–11. DOI: <https://doi.org/10.5334/johd.653>

Submitted: 06 September 2026

Accepted: 02 October 2026

Published: 27 October 2026

COPYRIGHT:

© 2026 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <https://creativecommons.org/licenses/by/4.0/>.

Journal of Open Humanities Data is a peer-reviewed open access journal published by Ubiquity Press.