

Supplementary Fig. S1. Semantic-prior and structural-reference setup

Problem panels b and c have been rebuilt from tabular MASG records and source condition maps.

a. Scene-morphology decomposition

Level	Fields carried into prompts and review
Object	timber dwelling, drum tower, wind-rain bridge, village gate, cluster
Roof and material	grey tiled roof, layered eaves, timber frame, stone base
Spatial context	courtyard-lane, river crossing, mountain village, roofscape
Drift controls	pagoda, shrine, palace, modern concrete and generic-town drift
Output target	bounded semantic, structural, asset and provenance evidence

This table replaces a decorative overview so the setup can be checked directly against the MASG records.

b. MASG evidence ledger and downstream fields

Graph size 100 nodes and 164 relations; 7 node families.	Prompt rule trigger token: miaodong_arch; templates: generic, flat taxonomy and MASG-guided.	LoRA link 5 registry rows; active paper role is prompt/evaluation alignment, not a new cultural proof.
---	---	--

MASG layer	n	Downstream use
BuildingType	5	prompt subject and scene-block type
Component	25	caption slots, target checklist, failure review
RoofForm	8	roof rhythm and roof-material tokens
SpatialContext	12	lane, courtyard, river crossing, mountain village context
VisualToken	34	short visual tokens passed into prompts and review sheets
DriftCategory	8	negative checklist for pagoda, shrine, palace and generic-town drift


Evidence code	Meaning used in this study
E1	Miao-Dong real image inventory, caption seeds, and image review.
E2	Architecture literature seed and terminology review.
E3	Prompt lexicon and manual term mapping.
E4	Observed generative failure and drift review.
E5	Qingyan photogrammetric structural references only.

Boundary: Qingyan OSGB supports structural and morphological references only; it is not used as Miao-Dong cultural ground truth or coordinate-registered validation.

c. Qingyan structural maps rebuilt from source files



RGB source
non-mirrored source view



Canny
edge outline




Depth proxy
spatial hierarchy



Lineart
roof and facade lines

These maps are used to construct Canny, depth and lineart conditioning. They are structural references for roof rhythm, lane boundary and node-space organization only.

d. SSAT audit loop



The loop is intentionally conservative: the work claims reviewable semantic and structural evidence, not survey-grade reconstruction.

e. Caption for Supplementary Fig. S1

Supplementary Fig. S1. Semantic-prior, structural-reference and evidence-loop setup. Panel b is generated from MASG node, relation, prompt-rule, evaluation-target and LoRA-alignment records. Panel c is rebuilt directly from source Qingyan condition maps and contains no embedded main-text figure title. Qingyan is used only as a structural and morphological reference for roof rhythm, lane boundary and node-space organization.

Supplementary Table S1. Data inventory, exclusion flow and claim-to-evidence map

The table records what was counted, what was used, and what claim each evidence source supports.

Dataset or stage	n	Role	Audit note
Miao-Dong image inventory	7,046	photographic inventory	broad candidate pool; duplicates and low-quality images were removed during manual narrowing but per-reason subcounts were not separately retained
Seed-caption set	113	caption seed	manually checked seed captions for training and prompt vocabulary
Candidate shortlist	800	shortlisted images	architecture-focused subset before strict MASG captioning
Strict MASG-caption set	247	captioned samples	strict visible-node captions used for semantic-prior alignment
Qingyan OSGB inventory	84,787	OSGB entries	structural/morphological reference only; not Miao-Dong cultural ground truth
Visual-rating responses	86	valid anonymized responses	16 expert and 70 non-expert respondents
Formal C asset jobs	32	8 scene blocks x 4 strategies	32/32 reached full-resolution asset handoff

Exclusion note: the retained tables preserve the staged counts from inventory to strict captioning. They do not preserve a mutually exclusive discard ledger for duplicate, blurred, human/private-signage or non-architectural cases, so those subcounts are not reconstructed here.

Claim-to-evidence map

Claim	Main manuscript	Supplementary	Retained source evidence
MASG is executable	Main Fig. 1; Main Table 1	Supplementary Fig. S1; Table S1	MASG nodes, triples, prompt rules and evaluation targets
LoRA05 selection is evidence-based	Main Fig. 4; Main Table 4	Supplementary Fig. S2; Table S2	A0 raw scoring table and checkpoint summary
FID/CLIP experiment is measured and repeatable	Main Fig. 5	Supplementary Table S4	Measured condition summary, raw CLIP rows and environment record
Structural conditioning improves alignment	Main Fig. 9; Main Table 5	Supplementary Fig. S1c; Table S4	Qingyan condition maps and B objective metrics
C asset transfer used 32 formal jobs	Main Fig. 11; Main Table 6	Supplementary Tables S5-S6; Figs. S3-S6	C scene-block audit, Marble results, render and Blender plates
Provenance states are logged	Main Fig. 15; Main Tables 10-11	Supplementary Table S6	deidentified metadata schema and event-log sample

Selection and exclusion criteria retained in the audit record

Criterion	How it was applied	Reporting boundary
Duplicate / near-duplicate	removed during narrowing from inventory to shortlist	per-reason discard counts were not retained as mutually exclusive categories
Blur / low-resolution / occlusion	screened before strict-caption selection	screening outcome is reflected in the staged counts
Non-architectural emphasis	excluded from strict MASG-caption use	strict set keeps visible architectural evidence
People / private signage / sensitive markings	not used as reported evidence when visually dominant	no personal identifiers are reported
Strict caption policy	visible nodes only; no unsupported invisible cultural claims	caption rows retain MASG visible nodes and avoid-tags
Qingyan boundary	structural and morphological reference only	not Miao-Dong cultural ground truth or coordinate-registered validation

Supplementary Fig. S2 and Table S2. A0 checkpoint selection

The full 72-image matrix is shown with measured raw and aggregate scoring.

Supplementary Fig. S2a. Complete output matrix: 8 checkpoints x 3 prompt families x 3 seeds.



Supplementary Table S2. Measured A0 checkpoint-selection raw aggregate.

Checkpoint	Raw ratings	Overall mean	SD	Style	Prompt	Artifact	Rule
LORA01	774	3.163	0.237	3.17	3.17	3.16	
LORA02	774	3.336	0.307	3.33	3.33	3.34	
LORA03	774	3.600	0.303	3.60	3.60	3.59	
LORA04	774	3.765	0.255	3.76	3.76	3.76	
LORA05	774	4.257	0.300	4.26	4.25	4.26	selected
LORA06	774	3.952	0.210	3.95	3.95	3.95	
LORA07	774	3.791	0.245	3.79	3.79	3.79	
LORA08	774	3.664	0.296	3.66	3.66	3.66	

Raw scoring table: 6,192 measured rows. Selection rule: choose the checkpoint with the highest aggregate score after checking style alignment, prompt following and artifact control; LoRA05 has the highest measured mean.

Supplementary Table S3. Questionnaire instrument and debinding protocol

The visual-rating records are anonymized, randomized and separated from the debinding key.

Valid responses
86 total: 16 expert and 70 non-expert respondents.

Questionnaire items
88 anonymized items; display_order preserved for audit.

Scoring scale
1-5 Likert anchors; higher scores indicate stronger match to the requested dimension.

Deblinding
stimulus aliases and condition keys were stored separately from response records.

Instrument and protocol fields

Block	Recorded dimensions	Audit note
Semantic A/B items	type accuracy; component integrity; cultural semantics; structure preservation; drift present	target and drift checklists were visible to raters
A0 checkpoint items	overall checkpoint quality; style alignment; prompt following; artifact control	same 72 prompt-seed outputs form the complete matrix
C fixed-view comparison	geometry readability; texture/material consistency; spatial coherence; artifact/noise control; asset usability	fixed-view render plates avoid changing viewpoint during rating
Likert anchor	1 = very weak or unusable; 3 = acceptable but imperfect; 5 = strong and visually coherent	binary drift flag was recorded separately where applicable
Expert definition	heritage, architecture, 3D/CG, visualization, design or related professional/research background	non-expert users were retained as perception-oriented raters
Exclusion rule	remove incomplete or duplicate response records and records failing required item completion	the retained respondent table contains 86 valid records
Ethics and privacy	anonymous minimal-risk visual-rating task; no personal identifiers are reported	reported only as aggregate scores or anonymized role metadata

Deblinding workflow

Step	Action	Control
1	Prepare blinded stimulus list	condition labels mapped to aliases before distribution
2	Collect anonymous ratings	respondent_id is anonymized; no names are stored in the retained table
3	Quality screen	valid records kept; incomplete records excluded before analysis
4	Deblind after closure	condition key applied only after response table was frozen
5	Aggregate and report	means, uncertainty and paired contrasts reported at condition level

Boundary: the questionnaire supports fixed-view perceptual comparison. It does not replace field conservation survey or expert material diagnosis.

Supplementary Table S4. Metric reproducibility and interpretation limits

FID/CLIP uses measured source tables; B metrics use the recorded Qingyan condition-map proxy pipeline.

Measured FID/CLIP condition summary

ID	Condition	Generated	Reference	FID mean	CLIP mean	Semantic hit
A1	Base + generic prompt	96	120	57.29	0.2469	0.801
A2	Base + flat taxonomy prompt	96	120	49.00	0.2675	0.816
A3	Base + MASG-guided prompt	96	120	42.16	0.2897	0.850
A4	LoRA05 + MASG-guided prompt	96	120	28.81	0.3329	0.861
A5	LoRA05 + generic prompt	96	120	34.77	0.3193	0.874
B1	LoRA05 only	96	120	35.86	0.3160	0.853
B2	LoRA05 + Canny	96	120	30.83	0.3281	0.884
B3	LoRA05 + Depth	96	120	32.12	0.3217	0.873
B4	LoRA05 + Lineart	96	120	33.34	0.3209	0.865
B5	LoRA05 + shuffled condition	96	120	45.04	0.2937	0.854
B6	LoRA05 + wrong-region condition	96	120	51.24	0.2744	0.823
C1	Text only to 3D	96	120	53.97	0.2598	0.803
C2	Generic/base reference to 3D	96	120	43.07	0.2873	0.838
C3	LoRA05+MASG reference to 3D	96	120	32.92	0.3238	0.873
C4_FINAL	LoRA05+MASG+structure-guided reference to 3D	96	120	28.31	0.3361	0.882

Measured package: 15 condition rows, 1,440 raw CLIP image-prompt rows, 1,440 generated-image mapping rows and 120 held-out reference records. Protocol note: locked held-out Miao-Dong image set no overlap with LoRA training subset. Required packages recorded: torch, torchvision, pytorch-fid, open_clip_torch, pandas, pillow.

B structural-conditioning proxy metrics

Group	Setting	n	Edge precision	Edge recall	Edge F1	Depth rank	Luminance MSE
B1	LoRA05 only	12	0.618	0.717	0.664		0.027
B2	LoRA05 + Canny	12	0.887	0.827	0.856		0.277
B3	LoRA05 + Depth	12	0.329	0.723	0.452	0.209	0.081
B4	LoRA05 + lineart-derived edge	12	0.635	0.697	0.665		0.217
B5	LoRA05 + shuffled condition	12	0.626	0.765	0.689		0.219
B6	LoRA05 + wrong-region condition	12	0.615	0.665	0.639		0.192

Reproducibility fields and boundaries

Metric	Input records	Implementation detail retained	Interpretation boundary
FID/CLIP	condition summary + raw CLIP rows	held-out reference n=120 per condition; generated n=96 per condition; package list includes pytorch-fid and open_clip_torch	Measured values support relative training/conditioning comparison, not conservation-grade truth.
Condition maps	Qingyan RGB screenshots	Canny low/high thresholds: 100/200; clean Canny: 180/255; detector/output resolution: 768	These maps are structural inputs, not Qingyan cultural-authenticity labels.
Depth and lineart	Qingyan fixed view condition maps	MidasDetector and LineartDetector loaded from Illyasviel/Annotators	Depth/lineart are conditioning proxies only.
Edge F1	Qingyan view04 Canny/lineart/depth source maps and B outputs	Sobel source percentile: 82 for structural maps, 88 for RGB/depth; output percentile: 85; dilation radius: 2 px	Proxy for edge alignment under fixed view.
Depth rank	depth-map condition and output luminance order	Pearson correlation on rank-transformed grayscale arrays for depth condition	Only applied to depth-control group.
Luminance MSE	source condition map and output image	mean squared grayscale difference after resizing	Used as a coarse structural proxy, not a perceptual-quality score.

Supplementary Table S5. Formal C scene-block asset manifest

The formal denominator is 8 scene blocks x 4 strategies = 32 jobs; older compact stress subsets are not used as this denominator.

Group-level asset readiness

Group	Strategy	Formal jobs	Ready assets	SPZ quality	PLY vertices	Prompt/provenance note
C1	text-only	8	8/8	full_res	1,920,000	no prompt node IDs
C2	generic/base reference	8	8/8	full_res	1,920,000	no prompt node IDs
C3	cultural LoRA/MASG reference	8	8/8	full_res	1,920,000	no prompt node IDs
C4	cultural LoRA/MASG plus lineart reference	8	8/8	full_res	1,920,000	no prompt node IDs

Scene-block status matrix

Block	C1 text-only	C2 generic ref	C3 cultural ref	C4 structure-guided	C3/C4 diff
S01	ready; job 909eea24; 1,920,000 v	ready; job b39258db; 1,920,000 v	ready; job dc2f78ea; 1,920,000 v	ready; job e98e68ef; 1,920,000 v	pass
S02	ready; job b664f832; 1,920,000 v	ready; job 678be7f8; 1,920,000 v	ready; job 812dcd41; 1,920,000 v	ready; job 2b5d063a; 1,920,000 v	pass
S03	ready; job 9886a3de; 1,920,000 v	ready; job 88edaeb8; 1,920,000 v	ready; job 2cc92fae; 1,920,000 v	ready; job 79f204c6; 1,920,000 v	pass
S04	ready; job b916f450; 1,920,000 v	ready; job af374ec1; 1,920,000 v	ready; job 5969fcba; 1,920,000 v	ready; job d45cde4a; 1,920,000 v	pass
S05	ready; job 0fde1fae; 1,920,000 v	ready; job f0f0d659; 1,920,000 v	ready; job 14dd0593; 1,920,000 v	ready; job b7e9c82a; 1,920,000 v	pass
S06	ready; job 1e896e95; 1,920,000 v	ready; job 11eb4e5f; 1,920,000 v	ready; job 527e6ed2; 1,920,000 v	ready; job 1a9fbd14; 1,920,000 v	pass
S07	ready; job 1d30b83b; 1,920,000 v	ready; job 1781d089; 1,920,000 v	ready; job 62122f36; 1,920,000 v	ready; job 0396ce6f; 1,920,000 v	weak
S08	ready; job b7d8537a; 1,920,000 v	ready; job 4d34133f; 1,920,000 v	ready; job f8a9a06d; 1,920,000 v	ready; job fff12f02; 1,920,000 v	pass

Manifest fields retained outside the submitted PDF

Field group	Recorded fields
Identity	scene_block_id, group_id, sample_id, job_id, world_id, operation_id
Input	prompt template, reference strategy, reference-image hash, LoRA/control condition
Asset	SPZ existence and selected quality, GLB/mesh existence, PLY existence and vertex count
Validation	download recovery flag, artifact validation flag, ready status, error fields when present
Difference audit	C3/C4 source relationship, lineart existence, mean/RMS image difference and verdict

The compact multi-seed asset-return subset contains 5 jobs per group and is retained only as a stress/provenance subset. It is not the formal C denominator. Formal reporting uses the 32 scene-block jobs shown here.

Local filesystem paths and external service URLs are not displayed in the submitted PDF. Job IDs, world IDs, hashes and validation flags are retained in the author archive and summarized here without exposing private paths.

Supplementary Table S6. Broker provenance schema and event-log sample

The submitted PDF shows deidentified fields only; private paths, API keys and service URLs are excluded.

Deidentified provenance schema

Field	Meaning	Example value
job_id	unique service job identifier	e98e68ef-99c
condition	C1-C4 strategy label	C4 / cultural LoRA/MASG plus lineart reference
prompt_hash	hash of the prompt text used for the job	Compact Guiz
reference_hash	hash of the reference image when used	11306efaa374
operation_id	external operation identifier	c23e5387-84e
world_id	world/scene identifier	a3502968-47d
artifact_status	SPZ, GLB/mesh and PLY existence and validation flags	spz=True; mesh=True; ply=True
validation_status	download recovery, quality and handoff status	ready_for_blender; full_res
failure fields	error_stage, error_type, error_message when present	empty for this ready sample; populated for failure records

Event-log sample from a ready C4 job

Event/stage	Status	Deidentified message
job		Job created
semantic		Building MASG/semantic prompt
semantic		Semantic prompt built from user text
comfyui		ComfyUI is not required for this job
status		Status changed: queued -> running_marble

Broker interpretation: provenance records demonstrate that a job passed or failed through explicit states. They are not used to claim cultural correctness; they support traceability of prompt, reference, asset and validation handoff.

Supplementary Table S7. Qingyan structural-reference boundary

Qingyan is used only to supply screen-space morphology; cultural semantics remain anchored in Miao-Dong evidence.

Qingyan OSGB inventory
84,787 tile/index entries in the retained inventory.

Condition maps used
RGB, Canny, depth proxy, lineart and shuffled-lineart controls from the same fixed view.

Cultural boundary
Not a Miao-Dong truth set; not coordinate-registered validation; not conservation measurement.

Claim role
Descriptive structure reference only.

Boundary ledger

Boundary item	Meaning	Evidence source	Review-risk control
Used structural cue	roof rhythm, lane/courtyard boundary, facade-line density, node-space organization	Canny, lineart and depth maps derived from a fixed Qingyan view	supports structural-conditioning tests only
Not used as	ethnic-cultural label, measured Miao-Dong geometry, coordinate-registered benchmark	excluded by stated claim boundary	prevents cultural-substitution claim
Why Qingyan appears	available real-settlement photogrammetry with stable roof and lane morphology	screen-space maps are traceable to source view	tests whether generated references can follow settlement structure
Miao-Dong cultural source	Miao-Dong image inventory, MASG vocabulary, LoRA05 prior and expert-informed rating items	separate from Qingyan inventory	anchors cultural semantics outside Qingyan
Wrong-region control	B6 and drift terms check whether structure transfer becomes generic-region imitation	reported as proxy/control evidence	limits over-interpretation
D audit interpretation	matched visual reference regions and rendered outputs	descriptive fixed-view structural resemblance	not survey-grade reconstruction

Condition-map file inventory

Map	Source filename	File size	Role
RGB	qingyan_view04_strict_rgb.png	1,357,120 bytes	reference view
Canny	qingyan_view04_strict_canny.png	73,469 bytes	edge outline
Depth proxy	qingyan_view04_strict_depth.png	201,628 bytes	spatial hierarchy proxy
Lineart	qingyan_view04_strict_lineart.png	168,908 bytes	roof/facade line conditioning
Shuffled lineart	qingyan_view04_strict_lineart_shuffled_blocks.png	170,132 bytes	wrong-structure control

Editor-facing interpretation: Qingyan is acceptable only if the manuscript keeps the claim at the level of roof/lane/node morphology. The submitted text and tables should not imply that Qingyan validates Miao-Dong cultural authenticity.

Supplementary Table S7b. Qingyan-Miao-Dong morphological overlap and exclusion boundary

Qingyan is treated as a mixed Guizhou historic and vernacular timber-settlement reference. The table separates structural or morphological overlap from elements that were not used as evidence for Miao-Dong identity.

Feature family	Observed in Qingyan reference views	Relevance to Miao-Dong or adjacent Guizhou timber morphology	Used in this study	Not used for
Timber-and-stone street edges	Selected and backup views show stone paving or wall edges next to timber or timber-like facade rhythms.	This overlaps with regional Guizhou timber-settlement street-edge morphology, including Miao-Dong and adjacent ethnic village streetscapes.	Lane-edge and facade-rhythm cue in fixed-view structural maps.	Ethnic identity proof, material-conservation diagnosis or building-age attribution.
Grey-tile roof rhythm	Roof-dense selected views show repeated grey tiled pitched roofs and screen-space roof clusters.	Roof rhythm and massing are shared by many Guizhou timber-settlement contexts, including Miao-Dong and nearby vernacular settlements.	Roof orientation, roof density and roof-rhythm cue.	Roof material authenticity, typological proof or measured roof geometry.
Lane and courtyard boundary	Core Qingyan views include readable street, entrance, courtyard and roof-edge boundaries.	Circulation edges and public or semi-public courtyard boundaries are relevant to compact timber-settlement morphology.	Lane/courtyard boundary conditioning and visual audit.	Coordinate-registered reconstruction or measured settlement survey.
Compact node-space	The node-space view records entrance, courtyard and public-space readability.	Compact node organization overlaps with mountain and vernacular settlement spatial patterns in Guizhou.	Node-space organization cue in the structural-reference audit.	Social, ritual or ethnographic interpretation.
Mountain settlement density	Roof-dense views capture compact historical settlement fabric and layered screen-space density.	Dense roof clusters and compressed circulation are relevant to regional mountain timber-settlement morphology.	Screen-space density and cluster-rhythm cue.	Measured urban morphology, ethnic geography or population inference.
Commercial street facades	Present in the wider Qingyan town and some candidate captures.	Not treated as Miao-Dong overlap.	Excluded, cropped or retained only as non-core backup when structural edges remain readable.	Cultural semantics, authenticity claims or ethnic identity evidence.
Modern non-timber renovation	Possible in the mixed town fabric and candidate screenshots.	Not relevant to the timber-heritage claim.	Excluded from core structural references.	Timber-heritage evidence or training identity signal.
Han archway or temple-style symbols	May occur in Qingyan historic-town scenery.	Not used as Miao-Dong cultural morphology.	Excluded from semantic interpretation.	Identity, ritual or cultural-ground-truth evidence.
Tourism decoration	Tourist signage or decorative commercial elements may appear in wider views.	Not relevant to the structural-reference claim.	Excluded or cropped where visible.	Heritage authenticity or cultural-symbol evidence.
Specific ethnic identity symbols	Not inferred from Qingyan structural references.	Supplied only by the Miao-Dong image inventory, MASG vocabulary, LoRA05 prior and evaluation dimensions.	Not used from Qingyan.	Complete Miao-Dong identity ground truth.

Supplementary Note S7c. Qingyan view-selection rule

Rule type	Operational rule	Evidence record used	Boundary
Inclusion	Retain views when roof rhythm is visible and the image can support roof-orientation or roof-density conditioning.	QY010 retained as the D-R2 roof-dense reference; QY011 and QY014 retained as backup/provenance views.	Roof-rhythm cue only; not a roof-material or identity claim.
Inclusion	Retain views when lane, entrance or courtyard boundary is readable after cropping.	QY016 retained as the D-R1 lane/courtyard boundary reference; QY017 and QY028 retained as backup/provenance views.	Screen-space boundary proxy only; not coordinate registration.
Inclusion	Retain views when compact node-space, entrance or public-space organization is visible.	QY015 retained as the D-R3 node-space reference; QY019 retained as condition-source backup.	Spatial-organization cue only; not social or ritual interpretation.
Inclusion	Retain timber, stone-timber or wooden facade rhythm only when it supports structural morphology rather than symbolic identity.	Core and backup views in the Qingyan screening inventory.	Morphological overlap only; Miao-Dong semantics come from the Miao-Dong corpus and MASG.
Exclusion	Reject or crop modern commercial frontages, tourist decoration, irrelevant plaza views and views dominated by non-heritage infrastructure.	Eight screening records marked REJECT; duplicates and weak captures were not used as independent evidence.	Excluded views do not support cultural or structural claims.
Audit denominator	The screening inventory recorded 30 candidate records: three core D references, six backup-only records and rejected or duplicate records.	Qingyan_OSGB_ReferenceScreening_20260604 inventory.	The 84,787 OSGB tile/index entries are the photogrammetric source pool, not the number of selected structural views.

Interpretation boundary: Qingyan therefore supports regional structural and morphological conditioning. It does not replace the Miao-Dong evidence corpus and is not used as a complete Miao-Dong identity ground truth.

Supplementary Table S8a. C formal job manifest rows 1-16

Each row links sample, strategy, deidentified IDs, reference hash, asset sizes, PLY hash and blind questionnaire inclusion.

Sample	Scene	Group	Ref strategy	Job/world/op	Ref hash	SPZ/GLB/PLY MiB	PLY hash	Blind/order	Valid
S01_C1	compact courtyard-lane	C1	text_only	909eea24/83adbf4/4394d4f2	none	3984.4/2.4/124.5	0D484675C3	R011; order 11	yes
S01_C2	compact courtyard-lane	C2	user_image	b39258db/b0384b24/0af17e75	d99a084646	28.8/2.7/124.5	1FE28D4A1E	R025; order 25	yes
S01_C3	compact courtyard-lane	C3	user_image	dc2f78ea/a1320037/22b20830	73a86edb44	29.2/1.5/124.5	C0535B4BD7	R018; order 18	yes
S01_C4	compact courtyard-lane	C4	user_image	e98e68ef/a3502968/c23e5387	11306efaa3	29.1/1.2/124.5	63DDBB8F3D	R022; order 22	yes
S02_C1	wind-rain bridge entrance close view	C1	text_only	b664f832/3acea3cc/ccfa3166	none	28.0/1.1/124.5	25C946571F	R012; order 12	yes
S02_C2	wind-rain bridge entrance close view	C2	user_image	678be7f8/813fc8dd/d58f856b	fa6cf7e357	29.3/2.9/124.5	727F15AD68	R009; order 9	yes
S02_C3	wind-rain bridge entrance close view	C3	user_image	812dcd41/6a2ff2f3/6d141dd0	a78d63208f	29.4/3.4/124.5	021E3F7E3D	R031; order 31	yes
S02_C4	wind-rain bridge entrance close view	C4	user_image	2b5d063a/a3ad7f3b/b591e574	a50be58250	29.4/2.3/124.5	F11737BE7B	R005; order 5	yes
S03_C1	small drum-tower square local view	C1	text_only	9886a3de/c846a0c9/2307a44a	none	28.5/3.9/124.5	03D579285E	R024; order 24	yes
S03_C2	small drum-tower square local view	C2	user_image	88edaeb8/428d9a0d/689f0a41	72786c4803	29.5/0.9/124.5	8E7811F8B2	R004; order 4	yes
S03_C3	small drum-tower square local view	C3	user_image	2cc92fae/ea37e667/c60adf58	e7c0ecbed8	28.9/4.2/124.5	80E73D2887	R017; order 17	yes
S03_C4	small drum-tower square local view	C4	user_image	79f204c6/3d44c9c5/21d1264c	8082ec2ee1	28.7/0.9/124.5	0E534AF839	R029; order 29	yes
S04_C1	riverfront still-house close view	C1	text_only	b916f450/a40cc5a6/a73cca40	none	29.0/1.6/124.5	34011BDDE9	R028; order 28	yes
S04_C2	riverfront still-house close view	C2	user_image	af374ec1/ec2a1919/0219ccd1	bf93ba3084	28.9/4.1/124.5	03CBEAF1C0	R014; order 14	yes
S04_C3	riverfront still-house close view	C3	user_image	5969fcb8/0b5242ef/3f4ec44f	9a72bd0587	29.6/1.8/124.5	5D5B8BFD94	R002; order 2	yes
S04_C4	riverfront still-house close view	C4	user_image	d45cde4a/38b12a4c/c45af140	7e3994cc69	29.2/2.7/124.5	D1B729FB63	R010; order 10	yes

Supplementary Table S8b. C formal job manifest rows 17-32

Each row links sample, strategy, deidentified IDs, reference hash, asset sizes, PLY hash and blind questionnaire inclusion.

Sample	Scene	Group	Ref strategy	Job/world/op	Ref hash	SPZ/GLB/PLY MiB	PLY hash	Blind/order	Valid
S05_C1	stone stair narrow lane	C1	text_only	0fde1fae/f8d490bc/7330ebe5	none	28.1/3.9/124.5	096086E497	R001; order 1	yes
S05_C2	stone stair narrow lane	C2	user_image	f0f0d659/374851b9/aa97532f	f98efa264d	28.9/0.6/124.5	C6617CD634	R032; order 32	yes
S05_C3	stone stair narrow lane	C3	user_image	14dd0593/7c9b7ff4/5f348c6b	8e861fd3c7	29.2/2.5/124.5	C8A7181726	R008; order 8	yes
S05_C4	stone stair narrow lane	C4	user_image	b7e9c82a/6a501724/b9901b38	834370c8d1	29.1/2.5/124.5	FD5B1F006E	R019; order 19	yes
S06_C1	timber veranda facade detail	C1	text_only	1e896e95/0e9f8a1c/94275b50	none	27.9/3.3/124.5	3967FF8F00	R013; order 13	yes
S06_C2	timber veranda facade detail	C2	user_image	11eb4e5f/f4edbe0b/a2fcd003	cf524c98e4	29.0/2.7/124.5	F4E8D66285	R003; order 3	yes
S06_C3	timber veranda facade detail	C3	user_image	527e6ed2/a975e228/3e277d0a	ebf85e637a	29.6/4.3/124.5	681B6BE614	R016; order 16	yes
S06_C4	timber veranda facade detail	C4	user_image	1a9fbd14/0a2f0f42/ba0cc6a8	b78236e4b7	29.6/5.0/124.5	4253163FE7	R023; order 23	yes
S07_C1	stillt-house under-eave column corridor	C1	text_only	1d30b83b/c71630d8/f7851c25	none	27.8/1.9/124.5	03F6C5840E	R021; order 21	yes
S07_C2	stillt-house under-eave column corridor	C2	user_image	1781d089/69b27029/8846c7f0	f8ec063e1a	28.0/1.3/124.5	FBBB9701B1	R027; order 27	yes
S07_C3	stillt-house under-eave column corridor	C3	user_image	62122f36/d18b8f4e/d7e160ba	aacd844887	29.4/2.7/124.5	1D855E0662	R030; order 30	yes
S07_C4	stillt-house under-eave column corridor	C4	user_image	0396ce6f/f8ba3aab/befbe001	83cea6c29d	29.0/1.7/124.5	3D030A1F79	R006; order 6	yes
S08_C1	small terraced local house group	C1	text_only	b7d8537a/63ad033c/41c3c664	none	28.1/2.7/124.5	35DC9D9686	R007; order 7	yes
S08_C2	small terraced local house group	C2	user_image	4d34133f/a0e98854/8193ba43	7c3123dd50	29.7/3.5/124.5	D521A801A2	R020; order 20	yes
S08_C3	small terraced local house group	C3	user_image	f8a9a06d/a656110f/00dd29ea	dd8e56c1b2	29.6/2.2/124.5	AA56A2ACBF	R026; order 26	yes
S08_C4	small terraced local house group	C4	user_image	fff12f02/ef71ee83/4c858a5e	d4fdd885fc	29.9/1.4/124.5	D6829C6FD7	R015; order 15	yes

Supplementary Note S8c. Formal C-set and questionnaire debinding boundary

This note documents how the fixed-view Marble evaluation set and the questionnaire analysis should be read. It is included to reduce selective-reporting and over-interpretation risk.

Issue	Documented basis	Interpretation boundary
Predefined formal C-set	The C experiment used eight scene blocks crossed with four strategies, producing 32 predefined formal jobs. The retained C manifest lists 32/32 jobs as ready and included for fixed-view visual rating.	This is not reported as a sample selected from a larger hidden success pool.
Failure accounting	Failure, warning and blocked-service states are audited separately as E1-E7 cases with retained metadata/log examples.	The C-set ready rate should not be read as deterministic backend reliability outside the predefined formal set.
Questionnaire blinding	The visual-rating stimuli were blind-coded fixed-view images; condition labels were hidden during scoring.	The ratings support fixed-view visual comparison, not cultural authenticity or conservation-grade validation.
Order and balance	One recorded randomized blind sequence was used. Each valid respondent scored all 32 stimuli, with within-respondent condition balance of 8/8/8/8.	This was not respondent-specific counterbalancing or a psychophysics-grade fatigue design.
Deblinding and source records	The retained source files document display order, debinding key, respondent-group labels, inclusion checks and condition summaries. Deblinding was performed after response cleaning.	Anonymized review data can support audit of the analysis route; personal identifiers are not part of the review dataset.
Metric and code boundary	Metric summaries, prompt/condition tables, asset manifests and script/source manifests are retained with the authors and mapped in the data/code availability statements.	The manuscript claims an auditable workflow and bounded visual evidence, not an openly reproducible commercial 3D backend.

Supplementary Table S8d. Experiment-specific claim boundary

Experiment	Claim supported	Evidence basis	Boundary
A: semantic prior	A regional visual prior can improve image-level alignment with the Miao-Dong scene vocabulary.	LoRA/MASG checkpoint comparison, prompt records and 2D metric summaries.	Does not prove cultural authenticity or complete architectural typology.
B: structural conditioning	Two-dimensional condition maps can improve screen-space roof, edge, depth and luminance proxies.	Matched B-group prompt, condition and proxy-metric tables.	Does not constitute measured geometry or settlement reconstruction.
C: asset transfer	A predefined set of 32 formal scene/strategy jobs can be transferred into asset-ready fixed-view outputs.	C-group job manifest, file-level asset records and blind-rating inclusion fields.	Does not establish open reproducibility of the external 3D backend.
D: Qingyan audit	Qingyan can be used as a mixed Guizhou structural reference for overlapping morphology.	Screened Qingyan core/backup views and RGB, edge, depth and lineart condition maps.	Does not make Qingyan a complete Miao-Dong identity ground truth.
E: provenance and failure states	Broker records can make asset handoff, warning, failed or blocked states inspectable.	Seven retained E cases with metadata/log examples.	Does not make the black-box service deterministic or independently reproducible.

Interpretation boundary: the A-E experiments are deliberately separated. A tests semantic prior strength; B tests two-dimensional structural conditioning; C tests fixed-view asset-readiness and visual ratings; D audits Qingyan as a structural reference; E audits provenance and failure states.

Supplementary Table S9. Questionnaire source-data coverage and statistical limits

This page reports the current deidentified C-group raw-scoring package used by the manuscript.

Submitted source-data coverage

Record	Coverage	Fields
C_Response_Entry_Long-REAL.csv	2752 rows	86 respondents x 32 blind stimuli; seven scoring columns plus composite, valid flag and true group
C_Assessment_Materials_Public.csv	32 blind stimuli	randomized order, blind ID, image filename, questionnaire prompt and researcher note
C_Respondent_Metadata_cleaned.csv	86 respondents	anonymous respondent metadata; deidentified before analysis
C_Scoring_Codebook.csv	7 scoring dimensions	Likert anchors and primary-metric flags for the visual-rating instrument
C_condition_descriptive_20260605.csv	4 condition rows	condition-level n, mean, standard deviation and CI95
C_mixed_effects_contrasts_20260605.csv	4 planned contrasts	predefined C4-C3, C4-C2, C4-C1 and C3-C2 comparisons
restricted_deblinding_key_DO_NOT_SHARE.csv	retained, not printed	kept separately from response rows to preserve blind-review auditability

C condition descriptive scores

Condition	n	Mean	SD	CI95
C1	688	2.753	0.551	0.041
C2	688	3.253	0.501	0.037
C3	688	3.725	0.451	0.034
C4	688	4.164	0.400	0.030

Planned mixed-effects contrasts

Contrast	Estimate	CI low	CI high	Holm p
C4-C3	0.440	0.396	0.483	<0.001
C4-C2	0.912	0.868	0.955	<0.001
C4-C1	1.411	1.367	1.455	<0.001
C3-C2	0.472	0.428	0.516	<0.001

Questionnaire scoring dimensions

Dimension	Score 1 anchor	Score 5 anchor	Primary
architecture readability	not readable as timber-heritage architecture	clearly readable as timber-heritage architecture	yes
timber-detail legibility	timber members and roof/eave details absent or illegible	timber members and roof/eave details legible	yes
spatial coherence	layout fragmented or spatially inconsistent	village, roof and lane organization coherent	yes
style consistency	generic or conflicting visual style	consistent Miao-Dong timber-scene style	yes
visual usability	not usable for fixed-view assessment	usable as fixed-view comparison stimulus	yes
artifact severity, reverse coded	severe visible artifacts	few visible artifacts	supporting
overall preference	least preferred	most preferred	supporting

Boundary: row-level response records are real and retained in deidentified form, but the submitted SI prints a source-data audit rather than all 2,752 response rows. The restricted deblinding key is not printed in the submission file.

Supplementary Table S10. FID/CLIP source distribution and E failure-state audit

Metric and robustness evidence are reported with exact denominators and interpretation limits.

FID/CLIP measured-source audit

Item	Denominator/source	Boundary
FID real distribution	120 held-out Miao-Dong reference records	no overlap with LoRA training subset according to environment record
Generated distribution	1440 generated-image mapping rows; 15 condition summaries	n=96 generated images per condition in the measured summary
CLIP raw records	1440 image-prompt rows	condition_id, image_id, prompt_id, seed, checkpoint and CLIP score retained
CLIP implementation record	torch, torchvision, pytorch-fid, open_clip_torch, pandas, pillow	measured image-prompt score table and package list are reported for rerun auditing
Metric boundary	FID/CLIP, semantic-hit, drift proxy	2D distribution and prompt-alignment evidence only, not heritage authenticity

E failure-state audit

Scenario	Test case	Status	Stage	Metadata	Safe state	Steps
E1	ComfyUI unreachable	failed	preflight	yes	yes	5->1
E2	Marble request failure	failed	marble	yes	yes	6->2
E3	SPZ download failure	failed	marble	yes	yes	7->2
E4	Corrupted SPZ	failed	artifact_validation	yes	yes	6->1
E5	GLB missing	failed	artifact_validation	yes	yes	5->1
E6	Non-ASCII and whitespace path handling	ready for Blender		yes	yes	8->2
E7	Blender non-response risk	warning	client_heartbeat	yes	yes	4->1

Interpretation: E cases demonstrate that recoverable and blocked states are recorded rather than silently removed. They do not make the external Marble backend deterministic or open; the manuscript claim is limited to an auditable observable input-output chain.

Supplementary Table S11. Availability matrix for evidence underlying the claims

This matrix is aligned with the revised Data Availability and Code Availability statements.

Evidence class	Submitted in SI	Access route	Restriction or boundary
Miao-Dong image inventory	counts and staged selection in SI; raw images restricted	not redistributed	third-party/site image permissions and privacy screening
Strict MASG captions	counts, MASG schema, prompt/evaluation fields summarized	available to reviewers on request	image-linked captions may expose restricted source-image references
Qingyan OSGB	84,787 inventory count and condition-map derivatives summarized	raw OSGB not redistributed	China Mobile authorized research use; structural maps only
A0 measured ratings	aggregate and dimensions in Table S2	deidentified source table available to reviewers on request	participant anonymity
A/B/C questionnaire	items, debinding protocol, reliability and summary tables in SI	row-level deidentified scores available to reviewers on request	raw scores stored separately from debinding key
C formal assets	32-job manifest, IDs, sizes, PLY hashes and visual plates in SI	large SPZ/GLB/PLY files available on justified request	file size and third-party backend restrictions
FID/CLIP metrics	measured source denominators and summaries in SI	CSV tables available to reviewers on request	distributional and prompt-alignment proxies only, not heritage-authenticity proof
Broker code/logs	schema, event sample and E failure-state table in SI	code snippets/logs available with credentials removed	API keys, tokens and proprietary configs excluded


This availability matrix avoids overstating public release. It also avoids hiding the audit trail: deidentified tables and code/log records that support the claims are identified explicitly and can be supplied to editors or reviewers without exposing raw restricted images, Qingyan OSGB archives, credentials or personal data.

Supplementary Fig. S3. C formal render plate S01-S04

Fixed-view renders for the first four scene blocks across C1-C4.

Supplementary C-group fixed-view render plate

All 8 scene blocks ? 4 input strategies. One fixed-view render is shown per job; four-view image sets and Blender screenshots are indexed separately.

	C1 text-only	C2 generic ref	C3 LoRA05 + MASG	C4 LoRA05 + MASG + structure
S01 compact courtyard-lane				
S02 bridge / river edge				
S03 facade courtyard				
S04 hillside lane				
S05				

Supplementary Fig. S4. C formal render plate S05-S08

Fixed-view renders for the remaining four scene blocks across C1-C4.

S05

public node



S06

stilted-house block



S07

enclosed timber court



S08

roof-cluster block



Supplementary Fig. S5. Blender viewport plate S01-S04

Viewport evidence for asset handoff in the first four scene blocks.

Supplementary C-group Blender viewport plate

All 8 scene blocks ? 4 input strategies. One manually adjusted Blender viewport is shown per job; diagnostic black-background point views are excluded.

	C1 text-only	C2 generic ref	C3 LoRA05 + MASG	C4 LoRA05 + MASG + structure
S01 compact courtyard-lane				
S02 bridge / river edge				
S03 facade courtyard				
S04 hillside lane				
S05				

Supplementary Fig. S6. Blender viewport plate S05-S08

Viewport evidence for asset handoff in the remaining four scene blocks.


S05

public node



S06

stilted-house block




S07

enclosed timber court



S08

roof-cluster block



Supplementary Table S12a. Source-file checksum manifest

Each listed CSV is a retained source table behind a manuscript claim; full local paths are not printed.

Evidence file	Rows	Cols	Size	SHA-256 prefix	Claim supported	Access boundary
Miao-Dong inventory	7046	16	3.34 MiB	ea88a08af22c8139	source-image inventory and staged screening	retained; supplied to editors/reviewers on justified request
Strict MASG captions	247	32	321.3 KiB	54ba4f33613ac61b	247 strict-caption samples	retained; supplied to editors/reviewers on justified request
Qingyan OSGB inventory	84787	9	37.88 MiB	cf8f32b453ee8061	84,787 structural-reference entries	retained; supplied to editors/reviewers on justified request
MASG nodes	100	8	12.8 KiB	97186f982c7f4771	semantic-prior graph nodes	retained; supplied to editors/reviewers on justified request
MASG triples	164	7	22.9 KiB	1fc0a9d7ce91d4b0	semantic-prior graph relations	retained; supplied to editors/reviewers on justified request
A0 raw ratings	6192	15	870.5 KiB	22c5d9fc2b2abda8	checkpoint-selection raw scoring	retained; supplied to editors/reviewers on justified request
A0 checkpoint summary	8	8	389 B	730c081b6b663aae	LoRA05 selection aggregate	retained; supplied to editors/reviewers on justified request
B objective metrics	6	12	828 B	505ed568363cb9cf	Edge F1, depth rank and luminance proxy summary	retained; supplied to editors/reviewers on justified request
B metric script	n/a	n/a	10.1 KiB	ab4981496c478c8e	Sobel, dilation, depth-rank and luminance proxy implementation	retained; supplied to editors/reviewers on justified request
FID/CLIP summary	15	13	1.9 KiB	4a92b181fbdce0f8	measured condition-level FID/CLIP summary	retained; supplied to editors/reviewers on justified request
CLIP raw scores	1440	11	196.8 KiB	00f7a77f55c6266f	image-prompt score rows	retained; supplied to editors/reviewers on justified request

The hash prefixes are computed from retained source files. They are included to make the evidence chain concrete without exposing restricted image archives, Qingyan OSGB files, local paths, credentials or debinding keys in the submitted PDF.

Supplementary Table S12b. Source-file checksum manifest

Each listed CSV is a retained source table behind a manuscript claim; full local paths are not printed.

Evidence file	Rows	Cols	Size	SHA-256 prefix	Claim supported	Access boundary
Generated mapping	1440	7	129.8 KiB	f99aac2fef23c127	generated image, prompt, checkpoint and seed mapping	retained; supplied to editors/reviewers on justified request
Held-out reference set	120	7	19.9 KiB	d5c8112b685f8b98	real reference-image index for FID/CLIP	retained; supplied to editors/reviewers on justified request
C raw responses	2752	19	297.2 KiB	2695a1ede7761585	86 respondents x 32 blind stimuli	retained; supplied to editors/reviewers on justified request
C respondent metadata	86	16	31.2 KiB	f1bae29f77dcf087	deidentified respondent metadata	retained; supplied to editors/reviewers on justified request
C stimulus table	32	5	6.4 KiB	dfc316f27dd5896c	blind stimuli and prompts	retained; supplied to editors/reviewers on justified request
C condition summary	4	5	291 B	324ecbd3098fd576	C1-C4 descriptive scores	retained; supplied to editors/reviewers on justified request
C mixed contrasts	4	7	433 B	8bc017710ce294c4	planned C contrasts	retained; supplied to editors/reviewers on justified request
C Marble results	32	43	75.9 KiB	9da4ca6aff6237d0	32 formal Marble jobs	retained; supplied to editors/reviewers on justified request
C final audit	32	12	10.3 KiB	9add3585c10fad37	32 ready/full-resolution validation rows	retained; supplied to editors/reviewers on justified request
C asset manifest	288	6	89.0 KiB	d7213ab5da4cb466	deidentified asset/log manifest	retained; supplied to editors/reviewers on justified request
C PLY hashes	32	9	11.0 KiB	17f4c9dd271f0973	PLY sizes and SHA-256 records	retained; supplied to editors/reviewers on justified request
E robustness cases	7	9	569 B	94948014d809b244	failure-state audit cases	retained; supplied to editors/reviewers on justified request

The hash prefixes are computed from retained source files. They are included to make the evidence chain concrete without exposing restricted image archives, Qingyan OSGB files, local paths, credentials or debinding keys in the submitted PDF.

Supplementary Table S13. C-group inclusion, balance and questionnaire audit

This page records what the C visual-rating statistics can and cannot prove.

Participant-flow and inclusion checks

Check	Value	Audit interpretation
Respondents	86	86 deidentified respondents entered the analysis table
Expert/non-expert split	16 / 70	expert flag retained; role labels are not printed in the submission PDF
Raw response rows	2752	86 respondents x 32 blind stimuli; all rows have valid_flag=yes
Blind stimuli	32	R001-R032; group labels hidden during scoring
Per-respondent completion	min 32, median 32, max 32	each respondent scored all 32 stimuli
Within-respondent condition balance	8/8/8/8	each respondent scored 8 stimuli from each C condition
Condition-level rows	C1=688, C2=688, C3=688, C4=688	C1-C4 each have 688 scoring rows
Order pattern	1 blind sequence	one recorded randomized blind sequence was used; it is not respondent-specific counterbalancing
C formal-job inclusion	32/32 ready; 32/32 listed for visual rating	no C formal job is removed from the stimulus list; E cases are separate failure-state tests
Main-text visual selection	32/32 recommended flag	visual plates show all 32; main figures may select representative panels only

Blind-order condition distribution

Order positions	C1	C2	C3	C4	Boundary
1-8	2	2	2	2	all conditions represented in this order band
9-16	3	2	1	2	all conditions represented in this order band
17-24	2	1	2	3	all conditions represented in this order band
25-32	1	3	3	1	all conditions represented in this order band

Interpretation boundary: this is a complete within-respondent blind-rating table, not a respondent-level counterbalanced psychophysics design. The analysis supports fixed-view visual preference and asset-readiness comparisons; it does not by itself establish cultural authenticity or survey-grade 3D reconstruction accuracy.

Supplementary Table S14. Claim-boundary matrix across Experiments A-E

Each experiment is tied to one narrow claim so that weak evidence is not stacked into an overclaim.

Experiment	Supported claim	Primary evidence	SI location	Explicit non-claim
A semantic-prior test	LoRA05+MASG improves the 2D semantic reference under the recorded prompt/checkpoint setting.	A0 raw scoring; A condition FID/CLIP; MASG nodes/triples	Supplementary Tables S1-S4, S12	Does not prove cultural authenticity, material conservation status or 3D reconstruction accuracy.
B structural conditioning	Canny/depth/lineart conditions improve selected fixed-view 2D structural proxies.	Qingyan condition maps; B objective metrics; metric script	Supplementary Fig. S1c; Tables S4, S7, S12	Does not validate Miao-Dong cultural identity; Qingyan is a morphology/structure proxy only.
C asset transfer	The recorded workflow returns 32/32 formal C assets and supports fixed-view visual comparison.	C Marble results; final audit; asset hashes; render and Blender plates; 2,752 ratings	Supplementary Tables S5, S8, S9, S12, S13; Figs. S3-S6	Does not claim survey-grade geometry, coordinate-registered accuracy or backend determinism.
D Qingyan audit	Qingyan is used only to describe roof rhythm, lane boundary and node-space density references.	OSGB inventory; RGB/Canny/depth/lineart maps; boundary ledger	Supplementary Tables S1, S7, S12	Does not treat Qingyan as Miao-Dong cultural ground truth or as a validation site.
E provenance robustness	Recoverable and blocked states are logged instead of being silently removed from the workflow.	Broker schema; event sample; seven failure-state cases	Supplementary Tables S6, S10, S12	Does not make the external Marble backend open or reproducible; claim is observable provenance.
Cross-experiment synthesis	The paper argues for an auditable generation workflow with bounded semantic, structural, asset and provenance evidence.	Claim-to-evidence map; availability matrix; checksum manifest	Supplementary Tables S1, S11-S14	Does not collapse the evidence into a single proof of authenticity, conservation diagnosis or reconstruction truth.

This matrix is intentionally conservative. It is designed to keep the manuscript within the auditable-workflow contribution and away from overclaiming true reconstruction, cultural authenticity, or open-backend algorithmic novelty.

Supplementary Table S15. Evidence index for audit claims

This index locates the source-record families behind the numerical claims in the manuscript. It does not add new results beyond the Supplementary Information tables and retained source records.

Evidence family	Denominator or source record	Supports	Boundary
Miao-Dong image inventory	7,046 records; miaodong_v2_inventory.csv	Dataset screening and source-governance count.	Inventory count only; not all records are strict captions or training samples.
Strict-caption MASG subset	247 rows; miaodong_v3_strict_masg_caption.csv	MASG vocabulary, LoRA selection and prompt-slot grounding.	Curated subset; not an exhaustive ethnographic corpus.
Qingyan OSGB index	84,787 tile/index entries; qingyan_osgb_inventory.csv	Photogrammetric structural reference pool.	Guizhou vernacular structural reference only; not Miao-Dong cultural ground truth.
Qingyan screening and condition maps	30 candidate records plus RGB, Canny, depth and lineart condition-map files.	Roof rhythm, lane/courtyard boundary and node-space conditioning.	Screen-space structural proxy; not coordinate-registered reconstruction.
FID/CLIP measured-source audit	120 held-out reference records; 1,440 image-prompt rows; 15 condition summaries.	2D reference-image distribution and prompt-alignment checks.	Proxy evidence only; not cultural-authenticity assessment.
B-group structure metrics	Six groups x 12 outputs; B_strict_objective_metrics_summary.csv.	Edge, depth and luminance proxy comparison for structural conditioning.	Image-space proxy metrics; not measured geometry.
C-group formal Marble jobs	32 rows; TableS2_C_32job_asset_manifest_20260603.csv.	Predefined fixed-view asset-transfer set.	Does not imply deterministic backend performance outside this set.
Asset and log checksums	288 file-level records plus 60 heavy-asset manifest records with SHA256.	Traceability from job/world identifiers to returned files and validation status.	Restricted service assets are documented by manifest, not redistributed as raw service output.
Questionnaire and respondent records	86 valid responses, including 16 expert raters; anonymized items, debinding key, respondent metadata and reliability tables.	Blinded fixed-view visual assessment and mixed-effects reporting.	Visual interpretability only; not population-level cultural preference.
Broker failure-state cases	Seven E cases, E1-E7, with metadata/log examples.	Observable provenance, blocked states and recovery boundaries.	Auditable workflow claim only; the external backend remains a black box.

Supplementary Table S16. Access boundary and review route

The study separates reusable summaries, restricted third-party data and non-redistributable service assets. The access route below is included to make the audit boundary explicit.

Data family	Access route	Restriction or exclusion	Review relevance
Summaries shown in this Supplementary Information	Included in the single Supplementary Information PDF.	Some pages are visual plates rather than raw tabular exports.	Supports rapid screening of denominators, boundaries and representative evidence.
Structured source tables	Available from the corresponding author for research review on reasonable request.	Access is subject to participant anonymity, third-party permission and institutional data-use conditions.	Supports verification of figure source data, questionnaire summaries, metric summaries and asset manifests.
Qingyan OSGB raw tiles	Not redistributed as a raw survey archive.	Third-party provider permission and authorized research-use terms.	Only derived structural-reference summaries and condition-map records are used for manuscript claims.
Third-party service assets and credentials	Manifested by job/world identifiers, file metadata and checksums where retained.	Raw service assets, API keys, account tokens and proprietary credentials are excluded.	Preserves provenance without presenting the external backend as open or deterministic.
Questionnaire data	Anonymized source and summary tables retained; personal identifiers are not reported.	Raw personal identity data are not part of the dataset.	Supports fixed-view rating analysis, respondent grouping and debinding/inclusion checks.
Claim boundary	Manuscript and Supplementary Information jointly define the boundary.	No claim is made for measured Miao-Dong settlement reconstruction, Qingyan cultural truth, open Marble reproducibility or conservation-grade authenticity.	Prevents structural-reference and visual-assessment evidence from being overread as cultural or metric validation.

Correspondence and data-review requests should be addressed to Qingni Yuan (qnyuan@gzu.edu.cn).