Pangu: Creating Worlds in Use

Situated World Production with Revisable Creative Decisions

Huaiyang Liu
DeepBean Limited • wyonliu@gmail.com

Abstract. World creation continues as people play, work, and revise what they want to make. AI production must connect those decisions to changes in geometry, appearance, bodies, actions, and rules while sustaining ongoing use. We introduce Pangu, a framework for situated world production whose central system object is a revisable production commitment: a relationship supported by creative experience, qualified by its scope, and applied to an identified subsequent work. The architecture connects editable production, local decisions, current-state installation, and the proposed improvement of later production. A restricted reference turns authored choices and explicit revisions into different subsequent stage arrangements that remain operable. This establishes a concrete link from changing decisions to later works within an authored finite relation library. A reason-sensitive development pair then tests real model interpretation and production with shared tools. The candidate fails initial update admission, leaving both production opportunities unexecuted; a full-history agent produces one usable work and one rejected work. Its memories select different relation labels for different reasons, but one final program fails to apply the selected relation to the new source. These results distinguish semantic selection, interface admission, source binding, and actual use. They support a limited executable system formulation, not a production-learning advantage. We specify a source-bound application path that keeps historical coordinates separate from current-work operands; its incremental benefit, broad production coverage, and creator value remain to be established.

1 The production problem

A creator may begin with a remembered place, a scanned room, a rough sketch, or an activity they want to share. They may ask for an animal that carries objects, a folding stage on which friends perform, or a mechanism that turns a courtyard into a game. These requests concern appearance, structure, action, and use together. A body must support its intended actions; a material must fit the activity; a mechanism must operate on the objects that were actually produced. The result must also remain accessible to subsequent human decisions.

Creation continues after the first output. A creator discovers a better use while playing, changes an object to accommodate another participant, or rejects a technically successful result because its expression is wrong. Existing participants may still be moving, holding objects, or completing an activity while a generation request is pending. Production must therefore account for both a changing design and a changing situation of use.

This paper asks: **how can an AI production system turn situated creative decisions into usable, revisable worlds and into better subsequent production?** The question is independent of any particular editor. A creator may express the request inside a game, through an existing design tool, in conversation, or through an application event. Opening a separate workbench is not a prerequisite. A workbench can still be valuable for detailed inspection, branching, and professional control.

Three difficulties motivate the proposed method. First, production representations have dependencies. Widening a structure can affect attachment points, articulation, clearance, interaction conditions, and presentation. A locally plausible edit can break another part of the activity. Second, creative objectives are incomplete and revisable. A later choice may change the purpose of the work instead of merely optimizing an earlier specification. Third, creative history is ambiguous evidence. A discarded design may have failed technically, exceeded a budget, expressed an unwanted style, or simply become irrelevant after the brief changed. These explanations imply different future behavior.

The complete remit includes reality capture and reconstruction, geometry and topology, materials and illumination, sound, bodies and animation, object capabilities, spatial composition, rules, and interactive presentation. The reference implementation covers a small part of this remit. Its limited scope is an experimental instrument, not a replacement definition of Pangu. Reconstruction, photorealistic enhancement, physical embodiment, and human creative benefit require distinct evidence.

The intended practical outcome is a work that people can deliver, modify, and continue using at a reasonable total cost. For a nonprofessional creator, this includes the ability to realize an idea without mastering every production tool. For a professional, it includes precise control, inspectable intermediate representations, reusable components, and predictable revision. The ability to grow as a creator is also relevant: a system that completes work while making its user less able to understand or influence that work may fail an important human objective.

The defensible system contribution is a connected object and execution lifecycle: a creative choice can establish a provisional production relationship, later evidence can retire it, and an applicable relationship can change a subsequent work without resetting se-

lected ongoing use. An editable source alone does not represent that lifecycle; a history log alone does not execute it. Its components have substantial prior art. We claim an explicit formulation, a restricted executable reference, and evidence that separates the lifecycle's stages. We do not claim a new general learning principle, a superior agent, or full production coverage. The remaining empirical question is whether this treatment of creative experience improves valuable complete production over equally informed alternatives.

2 Related work and the remaining question

World and scene generation. Holodeck generates embodied environments from language [1], while WorldGen studies traversable and interactive worlds [2]. Code2Games connects scene and gameplay production to engine execution and repair [3]. These precedents rule out treating a generated scene, a runnable world, or an execution-feedback loop as independently sufficient novelty. WorldCoder-Bench studies physically grounded world synthesis [4]. Its public v3 repository describes a 205-task core and states that the evaluation framework is not included; access to public tasks alone does not establish an official evaluation or leaderboard result.

Production interfaces and articulated assets. Tripo Agentic Canvas already supports agent-planned workflows, editable nodes, asset refinement, and workflow reuse [5]. LAM generates geometry and articulation through a programmatic route [6]. Articraft combines a production SDK, restricted execution, validation, and articulated-asset generation [7]. These methods are close to the representation and tool-use components of Pangu. We do not claim that connecting generators, defining a production DSL, or jointly producing geometry and joints is new in itself.

Editing and creative process. DesignTrace studies alternatives and the history of design decisions [8]. DesignAgent3D addresses interactive scene editing [9]. The 3DCodeEdit benchmark distinguishes executing a program from producing the intended geometric edit and examines feedback during revision [10]. VibeWorlding studies world construction and editing with learned policies and verification [11]. Pangu must therefore be compared with systems that can already edit, inspect, reflect, and learn. A weak generator without access to history would not isolate the proposed process-learning contribution.

Learning from use. Cursor's official account of deployed feedback and reinforcement learning illustrates how a production environment can supply learning signals [12]. Voyager combines executable skills, feedback, and an automatic curriculum [13]. ACCEL studies environment generation and transfer [14]. Trinity couples capability preparation and harness evolution [15]. Its preparation on unseen tasks prevents interpreting every reported gain as zero-practice transfer. These works motivate production learning while also establishing substantial prior art. A saved skill library, a generated practice environment, or an updated prompt does not by itself establish a new learning principle or recursive improvement.

Executable experience and personal context. SkillCompiler already decomposes experience into operations, combines contextual model execution with deterministic programs, validates reusable assets, and revises them from traces [26]. It includes operation contracts and partial successful operations from failed tasks. Executable induction, applicability checks, and partial reuse therefore do not establish Pangu's novelty. The unresolved production-specific question is whether the reason for a creative choice changes which cross-representation transformation should be reused while preserving selected properties. Our technical histories contain no human preference evidence, so they cannot yet establish this distinction. A version-sensitive personalization study reports task-specific skills as strongest overall in its v3 abstract and evaluates personalization with an LLM developer simulator [27]; it does not justify assuming that more personal history always helps.

Design intent and constraints. CIT-CAD infers entity, hierarchy, and relation structure from descriptions, guides CAD generation, and compares generated constraints to that structure [28]. Aligning Constraint Generation with Design Intent in Parametric CAD uses solver feedback in constraint-generation training [29]. Intent-to-constraint conversion, structural verification, and solvability are therefore not sufficient novelty claims. Pangu's proposed system object emphasizes the continued authority of a local creative decision: what it applies to, when it is superseded, and how its relationship binds to a later work. This is a specific research distinction, not evidence that the earlier systems cannot express it or that Pangu outperforms them. Eliminating inconsistent members of a finite hypothesis set also has classical precedent [30].

Reconstruction and presentation. VGGT provides geometric reconstruction capabilities [16]; 3D Gaussian Splatting provides a representation for visual reconstruction and rendering [17]. Neither representation alone establishes object functionality. Roblox Cube includes schema-based functional generation inside experiences [18]. Roblox Reality describes a hybrid direction connecting structured world state with high-quality presentation; its April 2026 demonstration was explicitly not real-time, and the later early-access date is a plan, not an available capability [19, 20]. We adopt the distinction between executable state and presentation, while requiring separate tests of their consistency.

Agents within activities and domain transfer. SIMA 2 studies agents acting, reasoning, and learning in virtual 3D worlds [21]. The existence of in-world interaction and experience-based improvement is consequently not unique to Pangu. Enterprise and embodied benchmarks such as TheAgentCompany and

BEHAVIOR offer relevant task families [22, 23], but their environments, observations, and success conditions differ. A common interface cannot establish transfer between them.

The remaining candidate contribution is specific: infer reusable production transformations from creative processes that include scoped intent, selection reasons, failed alternatives, and actual use; validate their applicability; and use them to improve subsequent revisions without losing the activity in progress. This combination is a hypothesis. Establishing its incremental value requires a strong same-information code agent, structured retrieval, and standard skill induction under equal resource conditions.

3 Continuing world production

3.1 Source, use state, and revision intent

Let a production source be $G = (V, E, A, B, R)$, where V identifies objects and parts, E records structural and representational dependencies, A contains geometry and appearance assets, B contains bodies and action implementations, and R specifies executable interactions. This is an organizing abstraction, not a requirement that every asset be symbolic. A reconstructed surface, a procedural mesh, a motion model, and a generated material may each be represented differently while sharing identifiable interfaces.

Let x_t denote the relevant state of use at time t. It can include participants, held objects, joint configurations, completed events, and task monitors. Source G and use state x_t are different: editing a source can change which actions are possible, while actions can change x_t without editing G. A generation request at time t may return at time t+d, when the activity has advanced.

A revision request $U = (D, K, S, J)$ records the desired change D, selected preservation obligations K, the scope S of the request, and the creator's reasons or context J. K is selective. A creator may intentionally retire an old activity or replace an object. The system must distinguish that decision from accidental loss. A universal freeze of all previous behavior would prevent meaningful creation.

A candidate consists of a source change Δ, any required state mapping m, and evidence e. Three acceptance questions remain separate. Does the candidate enable the desired new activity? Does it satisfy the selected preservation obligations? Is it applicable to the actual state when installed? A candidate that answers only one of these questions is incomplete. Replaying a previously successful trace is useful regression evidence, but does not settle all three.

3.2 Production knowledge and its limits

A process episode records the input source, relevant situation, requested scope, attempted change, creator response, reason when available, execution observations, and selected result. Its provenance distinguishes human expression, model inference, authored supervision, and observed outcome. Silence is not automatically approval. A technical rejection is not automatically a judgment about aesthetic value. A choice made under one budget is not automatically a permanent preference.

The proposed learner derives an executable transformation k with applicability conditions C_k and a validation procedure V_k. It may use the transformation, adapt it, or fall back to a general generator. This executable-skill structure overlaps prior work [26]; the production objective alone does not establish a new algorithm. The candidate learning target is future joint satisfaction of the requested change and preservation obligations, subject to a total production budget. This includes preparation and failed attempts, rather than only the final successful invocation.

The representation does not determine the learning algorithm. Parameter learning, executable skill induction, retrieval, and changes to the improvement procedure are distinct interventions. Our current implementation supplies source transformations and process records but does not implement a validated cross-project learner. No result in Section 6 should be read as evidence that this learning step has occurred.

3.3 Revisable commitments and source bindings

We describe a production commitment as $c = (r, O, S, E, v, z)$: a relational form r, intended object roles O, applicability scope S, supporting creative episodes E, version v, and status z. Status distinguishes active, suspended, and retired commitments. The episode preserves what was selected and why; the commitment expresses what may guide another production. A recorded coordinate is evidence about its original work, not automatically an operand for a new work.

Application creates a separate binding $b = (h(G), \beta, q, U)$. Here h(G) identifies the source actually being revised, beta resolves the intended roles to that source's objects, q contains quantities measured from it, and U is the current request. The instantiated change is $\Delta = \text{Apply}(r, b)$, followed by installation against the current use state. Creative knowledge and its current instantiation can thus be inspected separately. Different sources may share r but require different q and Delta. This is ordinary relational abstraction and late binding organized around continuing creative decisions; the notation itself is not a learned algorithm.

Our finite reference has authored relation forms and role identities. Its update controller suspends locally

affected commitments, requires targeted retirement when relations change, and allows only one active applicable commitment to drive production. It checks the existence of a reason quotation, not whether the quotation entails the relation. Model-derived semantic proposals, scope discovery, role discovery, and any gain from this organization require their own evidence. The source-bound extension in Section 4.6 is a specified next intervention, not a result of the completed model batch.

3.4 A bounded continuation argument

Consider deterministic partial transition systems $M = (S, A, T)$ and $M' = (S', A', T')$. Let m initialize revised state, pi project revised state onto the protected old state, and rho implement protected old actions in the revised system. Assume $\pi(m(s)) = s$. For every related reachable pair and protected action a, assume equivalent enabledness and, when enabled,

$$\pi(T'(s', \rho(a))) = T(s, a). \tag{1}$$

Then every finite trace consisting of protected actions has equal enabledness and equal projected endpoint in both systems. The zero-step case follows from initialization. If the relation holds after k steps, equivalent enabledness determines whether the next step can occur, and the transition condition preserves the relation. Induction yields the result. This is a standard simulation-style argument [24], not a new general theorem about world models.

The assumptions carry the difficulty. New actions can deliberately change old state, so arbitrary interleavings require additional conditions. Hidden engine state, stochastic dynamics, continuous contact, and concurrency are not covered by this finite deterministic statement. The selected protected set and the maps m, pi, and rho require independent justification. In addition, preserving old behavior supplies no evidence of useful novelty: the identity change preserves everything while adding nothing. A separate reachable witness for the new activity is necessary, and such a witness still does not establish enjoyment or meaning.

4 Production within an ongoing activity

4.1 A production service with optional interfaces

The service receives an activity-linked request and the source and state it is permitted to observe. It produces candidates, executes relevant tools, returns interpretable consequences, and proposes a source change. The host activity continues while generation proceeds. Installation checks the candidate against the current source and current use state. Outcomes return to the production process with their versions and provenance.

Studio and Canvas remain useful interaction options. Studio provides detailed production and repair operations. Canvas supports alternatives, branching, composition, comparison, and explicit local decisions. Neither is a mandatory architectural layer. The same production service can be invoked through a game interaction, a design plugin, a business event, or conversation. Detailed review may still require an interface; reducing switching does not justify hiding consequential changes from the creator.

Pangu can use an ordinary engine or local activity host. Nuwa may supply situated understanding and Fuxi may provide durable shared-world execution, but Pangu's production mechanism must be evaluated without requiring those systems to supply its missing contribution. The reference host is deliberately limited; it is not a replacement implementation of Fuxi.

4.2 From reality and imagination to editable production

Real inputs require more than importing points. Capture calibration, segmentation, uncertain geometry, occlusion, object identity, and the distinction between observed and inferred structure affect later use. A reconstructed appearance should be linked to editable geometry, material regions, and possible articulations. When functional structure is inferred, its status must remain explicit until checked. A synthetic point fixture cannot substantiate real reconstruction.

Imagined inputs can arrive as language, images, sketches, existing assets, or a partially completed work. A complete production route coordinates geometry and topology, materials, illumination, sound, bodies, motion, capabilities, spatial composition, and rules. Specialized models and professional tools may supply different stages. The production service needs shared interfaces and dependency information so that a changed stage can trigger the appropriate downstream repair and validation.

For a body, appearance and action must be produced together. Joint structure, available degrees of freedom, contact geometry, reach, and the object's intended capability constrain one another. A named capability such as carrying or opening is not sufficient evidence that the body can perform it. Likewise, changing a body may require changing an action implementation. Our restricted reference supports rigid articulation; it does not demonstrate learned motion, deformation, balance, or real-robot control.

The architecture permits both professional and nonprofessional control. A specialist can inspect parts and dependencies; another creator can work through examples and situated requests. These are different interaction needs, not evidence that one interface serves both equally well. Evaluation must include whether users can understand and change the result, not only whether the system emits an asset.

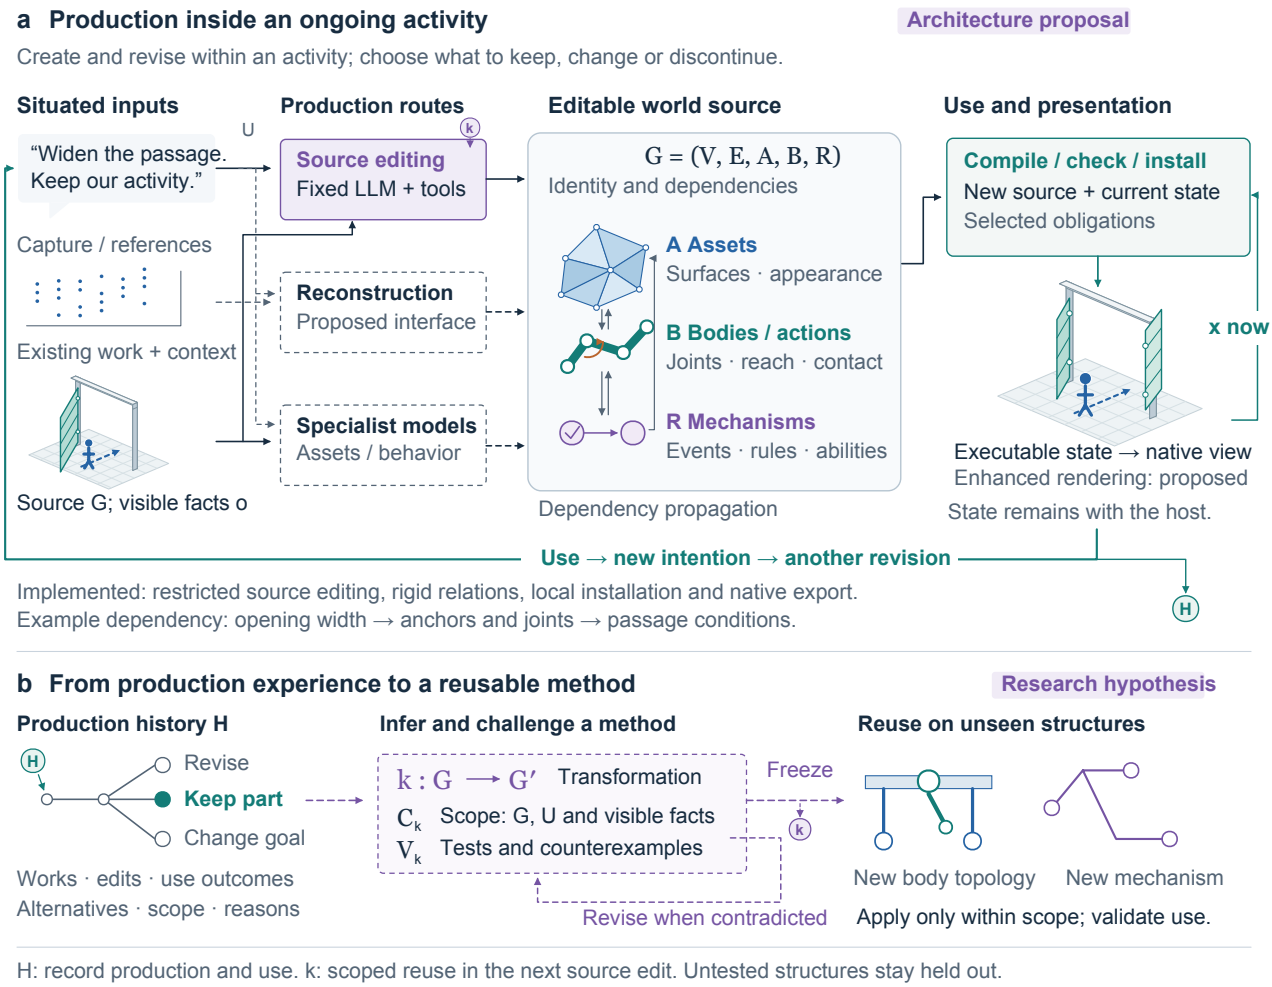


Figure 1: Proposed full production architecture. Situated intent U and source G connect assets A, bodies/actions B and mechanisms R through identities V and dependencies E. These links express constraints and revision propagation, not a demonstrated joint learned model. Installation uses current activity state x. History H, transformations k, applicability C_k and checks V_k depict the proposed learning route. Section 3 distinguishes a revisable commitment c from its current-source binding b; k is the broader transformation, not a claim that general binding or learning is implemented. Reconstruction, rich bodies, enhanced rendering and learning advantage remain unvalidated. The scenes are schematic.

4.3 Executable source and presentation

An editable source connects objects, parameters, motion, and rules. Compilation produces the representations required by the host engine and renderer. Persistent identities and versioned correspondence allow an edit to refer to the intended object, while geometric and behavioral checks inspect whether the correspondence remains correct.

Presentation may combine native graphics with optional generative enhancement. Executable state provides facts about identities, actions, contact, and events; an enhanced image is a representation of those facts, not a replacement source of truth. A matching version tag cannot establish semantic consistency. Evaluation must inspect object identity, geometry, occlusion, contact, motion, temporal stability, latency, and cost. Visual quality and interaction fidelity need separate measurements.

The current export produces glTF geometry, basic PBR color materials, and joint-range animations [25]. Editable source, use state, and provenance accompany the file. Custom rules and actor actions still require an interpreter. Exporting a file does not establish that arbitrary engines execute the same activity, nor that a generative presentation system is connected.

4.4 Installing a change while use continues

The local reference separates the generation snapshot from the installation state. A proposal identifies the source on which it was based. If another production change has already replaced that source, the host rejects the stale proposal. If only the activity state has advanced, the host applies the proposed source transformation to the current state and checks the task's selected preservation obligations. It also checks that the new geometry does not occupy the current participant's body.

The implementation performs generation outside the host lock. Local validation and assignment occur within the reference host's installation boundary. Under the assumptions of a complete local state, correct migration and validation, and atomic installation, a successful installation does not reset activity merely because generation used an earlier snapshot. This is an application of established versioning and atomic-update ideas. It supplies no general guarantee about distributed hosts, missed external events, or irreversible physical effects.

The distinction is observable. A naive installer can correctly construct a new source yet restore an old actor position or erase an event that happened while the request was pending. A valid in-activity implementation must instead continue from the current state. Refusing every change avoids some errors but does not demonstrate production usefulness; successful new work and preservation must be reported together.

4.5 How creative experience changes later production

The learning objective is to improve subsequent creation from the process of making and using works. Episodes retain alternatives, local selections, reasons, and observed consequences. A rejected body may be mechanically invalid; a valid body may express the wrong character; an earlier preference may cease to apply when the activity changes. These events call for different updates. Discarding every rejected branch loses usable partial choices; treating every rejection as a permanent negative teaches the wrong scope.

The present controller makes a narrower intervention explicit. On new local feedback it suspends the affected active commitment. An external semantic proposal selects a relation and scope, cites the available episode, and identifies any previous commitment to retire. Geometric incompatibility or an unresolved revision blocks its application. An accepted update changes which relation is available to future production; retaining only a version number would not accomplish this. The creator can still revise the purpose. Quotation checks do not certify the proposal's meaning, and the controller can accept a well-quoted wrong interpretation.

This reference interprets choices over an authored finite library; it does not discover new relation formulas or train a generative model. A stronger production learner could infer reusable cross-representation transformations, conditional role mappings, or revision rules from experience, then test them before reuse. Its gain must exceed that of a strong agent allowed the same history, tools, reflection, and budget. Executable skills, retrieval, and learner changes remain distinct interventions. Studio and Canvas can expose the process; neither supplies the learning mechanism by existing as an interface.

4.6 Applying a revisable relation to the next work

The latest batch exposes a specific separation: a memory can select the right relational label yet the subsequent program can use a number that does not satisfy it. We therefore specify **source-bound commitment application** as the next intervention. The retained memory identifies the relation, roles, evidence and current authority; the application stage must bind its operands to the work actually being edited. Historical values remain in the evidence record. They cannot silently substitute for current-source values. An intentional literal, such as a requested absolute position, must be represented as such with its own evidence and applicability, rather than prohibited universally.

For a relative camera position, let q_old contain the current work's camera and stage positions and let U contain the requested new stage position. Application computes

$$x'_{camera} = x'_{stage} + (x_{camera} - x_{stage})_{G_{current}}. \quad (2)$$

The subscript is consequential: the old values belong to the work now being revised, not the historical training example. In the development source they are -0.5 and -2; the requested stage position is 1, giving 2.5. The observed failed program instead contains 3, a value equal to the historical memory's coordinate. That equality does not establish the model's internal copying process.

The executable path is:

```
commitment =
resolve_active(decision_versions,
current_observations)

binding = bind_roles_and_operands(commitment,
current_source, request)

if roles, scope, or evidence for a required
literal are unresolved:  defer

change =
shared_relation_editor(current_source,
binding, commitment.relation)

check selected relationship and preservation
in the resulting source

install(change, based_on=current_source_hash,
use_state=host.current_state)

record actual use and any new creative
decision as a subsequent episode
```

Existing `reason_grounded_update.produce` already recomputes quantities from its supplied source and applies one of three authored formulas. The shared `edit_source` tool likewise operates on the supplied original source; both methods can use it. `ActivityHost.install` already rejects a stale source and preserves selected current state. These are existing production operations, not new learning contributions. The candidate failed before reaching this path in the real batch. The successful strong-agent branch used the same editor through a final edit request.

The proposed addition is a mandatory, inspectable binding step between an accepted creative commitment and the final production change: identify the current objects, distinguish current-source operands from historical evidence and deliberate literals, retain the source/version binding, and defer if it cannot be justified. A production route should not bypass the accepted relation by emitting a new unbound program without explicitly proposing a revised decision. The typed binding record and its general role-resolution policy are not implemented or evaluated here. Source hashes, checks, and three authored formulas alone do not provide them.

This intervention concerns applying retained knowledge. Learning remains responsible for which relation to retain, its scope and revision; the public editor remains responsible for carrying out an instantiated relation. A binding correction cannot be reported as learning a generation algorithm. Evidence that would change the present judgment is a valuable complete work with different structure, where choices cause an appropriate relationship update, that update binds correctly to the new work, and the resulting change improves usable production over a strong agent with the same tools and information. Giving the answer to a binding checker, or repeating this camera example, would not establish that gain.

5 Reference implementation and development protocol

The reference uses a bounded relational program with numerical parameters, arithmetic expressions, stable part identities, parent-child transforms, rigid joints, and rules. A general source patch can add or remove parts and alter relations; it is not restricted to selecting one of several prewritten correct repairs. The interpreter excludes arbitrary Python and shell execution. Its geometric domain is intentionally narrow.

Geometric predicates can measure distances between anchors and test an upright traversal body against all solid geometry. For upright or yaw-rotated boxes, the path checker uses the continuous horizontal sweep of a cylinder along a specified level segment. Other primitive bounds and tilted boxes are conservative. The implementation has no gravity, compliant contact, stair climbing, or dynamic avoidance. The size and path of the tested body are part of the task specification.

The model-assisted structural-editing artifact begins as a single articulated gate. A separate authored touring-stage work connects a folding platform, camera, support, interaction rule and continuing performance. It is used to examine decision versions and reason-sensitive subsequent production. These are separate evidence streams: the stage is not an automatically generated initial world.

The earlier structural artifact begins as a single articulated gate. Its revision widens the opening and changes the structure into two independently articulated leaves while preserving the participant, held object, history, and selected joint state. This is a small but complete production interface exercise: it involves geometry, hierarchy, motion, behavioral conditions, state migration, use, and export. It is not a representative proof of creative richness, scene-scale production, or the full architecture.

Task checks are based on geometry and execution, not the generator's success statements. Early checks were insufficient: a threshold between sampled angles escaped detection, and a solid obstacle added outside the named parts could block the opening while the named-anchor checks still passed. We retained those controls, added current-state and decision-boundary checks, and tested the actual specified path against every solid part. The corrections improve this evaluator's validity; they do not create a new production-learning result.

Within each fixed comparison, arms receive the same external source, state, history and common tools with matched budget ceilings; their internal representations and actual costs may differ. The early gate harness advances the host after taking a snapshot and before invoking the transport. In the later reason-development batch, the response is obtained before the local host advances the joint to 0.45 and attempts installation. It checks installation against advanced state; this schedule is not a measurement of concurrent progress during API inference. Failed candidates, repair feedback, raw responses, and costs are retained. A transport error stops the run instead of silently switching a model or consuming another arm's budget. Scripted transports have an explicit non-model provenance label throughout their outputs and exported artifacts.

6 Evidence and findings

6.1 What changed in a complete authored work

The touring-stage reference links past choices to a subsequent editable work and to continued use. Five authored local-selection episodes filter three authored relation forms: fixed world position, fixed relative offset, and proportional offset. Ambiguous evidence remains unresolved. An explicit revision retires the old offset relation. A frozen active relation then changes a later stage, after which installation retains the advanced joint state and the performance action remains available.

The earlier version places the later camera at 2.5 m; the revised relation places it at 3.4 m. This is an actual effect of changing the relation version, not a strong-baseline improvement of 0.9 m. The new value uses the later work's own original ratio, 1.5/5, giving 1+(1.5/5)*8. The system did not learn a fixed optimal camera proportion from the revision episode. The creator, options, scope labels, explicit revision and relation library are authored; natural-language reasons

are stored but not parsed in this reference. Full-history recomputation obtains the same relation state. The same prediction function checks execution consistency, not independent correctness of a real creator's intent.

This establishes a bounded system property: changing a supported relation version can change subsequent production while continued use is preserved. It does not establish autonomous interpretation, a new learning principle, or user benefit. The current rendered stage includes separately recorded author corrections to stage height and joint limits; the original source is retained.

6.2 Retained evidence and negative results

Table 1 keeps earlier records in view without treating them as independent successes. Their detailed frozen responses, evaluator changes and author diagnostics remain in the accompanying evidence index and accompanying evidence archive.

Source-field preservation did not guarantee geometric preservation in the gate: changing width from 2 to 2.4 moved a protected post by 0.20 m through its unchanged expression. The old success field remains frozen; it cannot certify the full request. Separate later author-format diagnostics moved posts by (0.15,0,2) m and (0.30,0,2) m. They are different artifacts and are not combined with the 0.20 m audit.

6.3 A reason-sensitive real-model development pair

The next design holds the selected geometry, alternatives and observation fields identical while changing only the second creative reason. Both world-position and relative-offset relations fit the geometry. One reason keeps the camera on the ground; the other changes its arrangement to follow the stage. Unlike the authored relation-version sample, the updater receives neither a mount label nor an explicit replace-previous-choice flag. Target object identities and the project observation remain authored. The language supplies meaningful extra information, but the directions are simple synthetic instructions, not evidence of nuanced human preference learning.

Both methods first update from episode A, then independently process the two B branches from that A memory. Both branch memories freeze before either subsequent production input is supplied to that arm. The candidate uses its commitment controller; the strong agent may store arbitrary JSON memory and reconsider full history during production. Both have an executable quantities/compile/edit interface. The strong agent may use the same finite relationship editor, request a concise edit, or generate a complete source. It is not forced to serialize a world that the candidate obtains through tools.

The frozen ceiling is 10 requests and 40,960 output tokens per arm, at most 4,096 output tokens per request, with preparation, tools and corrections counted. The model is qwen3.7-plus-2026-05-26, temperature 0, with thinking disabled. All 13 actual responses return this exact model identifier. The independent authored use target is withheld from update and production; final scores are not sent back for repair. The evaluator checks both camera and support, their dimensions and relative arrangement, stage width and position, and selected ongoing state. It does not certify photography, general physical validity, or arbitrary work quality.

6.4 Actual outcomes and the location of each failure

Table 2 accounts for every predeclared opportunity. It is an execution-disposition table, not a clean equal-protocol estimate of relative method quality or cost; Section 6.5 records the stopping deviation.

The candidate's first response explicitly selects world_x and quotes the ground-fixed reason, but omits the outer response envelope. It also includes an extra scope field that the run never reaches because outer parsing fails first. Later responses submit whole programs or use tools rather than provide an admissible semantic-update artifact. Its eight requests therefore end without an accepted A update. Neither later opportunity is erased from the denominator; neither is mislabeled as a tested geometric failure. The mixed update/tool/production instructions and coarse rejection message are plausible interface contributors, not experimentally isolated causes.

The strong agent's branch memories contain different relation labels for the two reasons. Thus the reason difference is reflected in its recorded selection; this does not establish general understanding. In the changed branch, the final artifact is a complete program with camera_x=3. That value is equal to the historical memory's coordinate; an internal copying process is not proven. The final artifact has no relation-edit request, and the shared edit tool is not invoked on this path. The saved production source equals the raw model program. The failure is observed between relation selection and its application to the next source, not a measured error of the public editor.

In the continued branch, the final artifact requests edits to width and stage position with relation world_x. The host dispatches that shorthand to the shared editor, which preserves the new work's original camera position of -0.5. All five strong-agent responses use kind=final; there is no separate kind=tool message exchange. Tool execution through the final edit dispatcher is real, but should not be confused with an additional agent tool dialogue or the earlier authored transport checks.

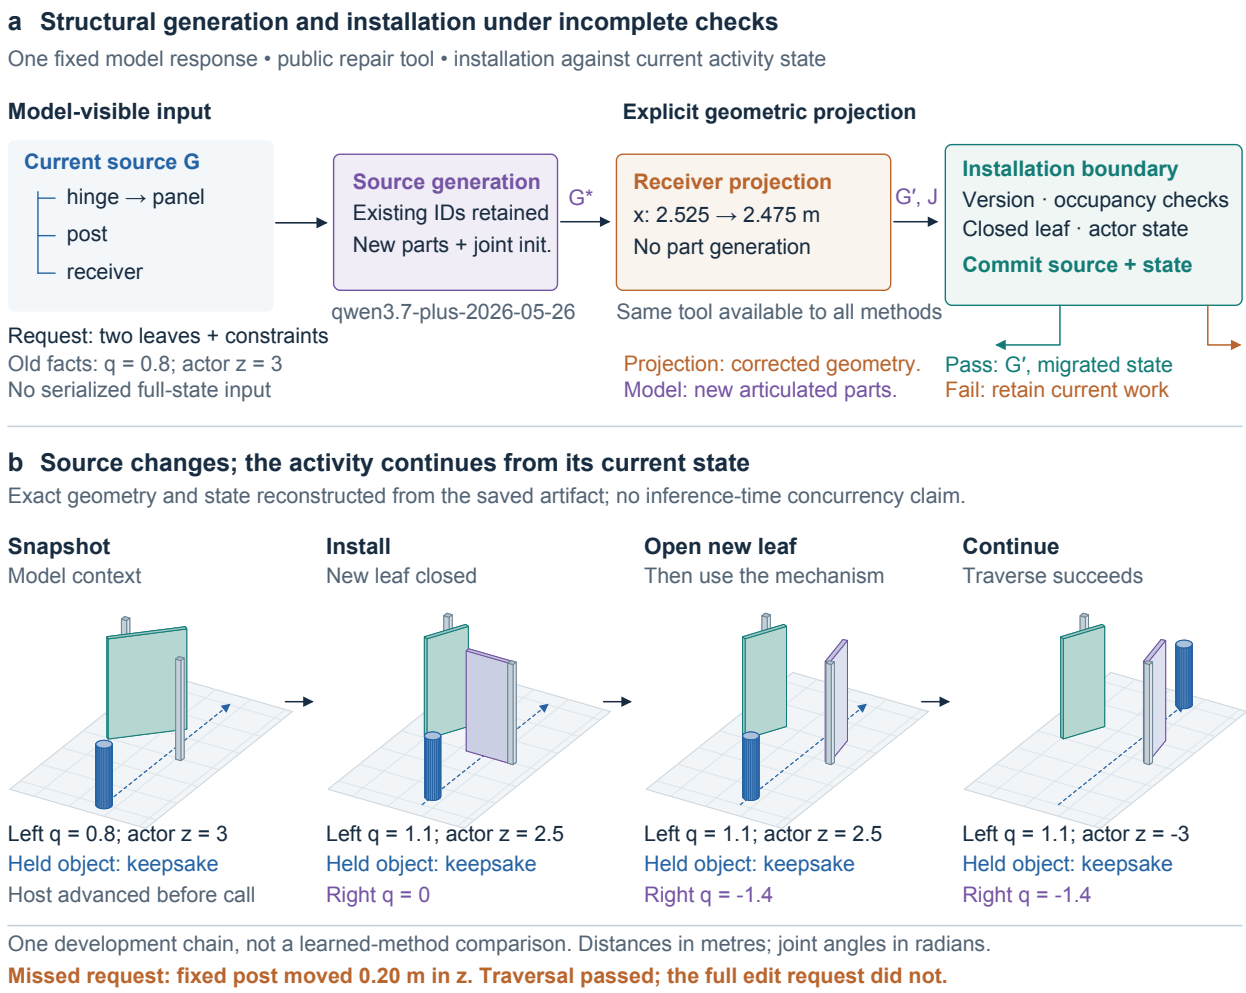


Figure 2: Earlier structural-production evidence, separate from the reason-sensitive study. The fixed model creates a second articulated leaf; a public tool corrects receiver geometry; current-state installation permits the recorded subsequent traversal. The original request also required a fixed post, but changing width displaced it by 0.20 m along z. The full edit request therefore failed despite passing the then-used checks. Four views are reconstructed from saved source and state. The host advances before the model call; inference-time concurrency is not measured. This is neither a learned-method comparison nor evidence for the proposed general binding layer.

6.5 Execution deviation, cost, and interpretation

The written protocol says that the same public error twice ends a stage. The frozen implementation compares only consecutive equal errors, resetting after a successful tool result. The candidate receives the same update rejection at zero-indexed attempts 1, 4 and 7, separated by tools, and reaches its eight-request A-stage allocation. Under the written cumulative interpretation it should have stopped after attempt 4. Attempts 5–7 add three requests and CNY0.049568 of estimated cost, all retained. This is an execution deviation from the prose stopping rule. There is no corrected rescore or rerun. The run remains within the request and financial ceilings, but it cannot support a clean equal-protocol method or cost comparison.

The candidate uses 8 requests and 4,405 output tokens at an estimated CNY0.103852; the strong agent uses 5 requests and 1,431 output tokens at an estimated CNY0.047540. All 13 responses finish with stop, not output truncation. Total input is 52,352 tokens. New estimated cost is CNY0.151392. The full round now contains 54 requests and CNY0.734980 including the original CNY0.077 unknown-billing reservation. These are usage estimates and reserves, not provider invoices or total research-and-development cost. Preparation and previous failed attempts remain in the full-round ledger.

The batch establishes one real model-selected, tool-produced, installed and used branch, and records where the other paths stop. It does not establish a candidate advantage: the candidate never produces an admitted update, and the small, known development pair with interface sensitivity and a stopping deviation cannot settle a general comparison. The same task is closed to further paid expansion. The remaining question is whether creative experience can reliably change important subsequent production, including its source binding, over an equally capable agent. Broader recon-

Table 1: Retained evidence streams and their inference limits.

Evidence stream	Directly observed result	Inference boundary
Finite continuation diagnostic	Authored counterexamples distinguish preserved old behavior from reachable new opportunities; exact contract baseline agrees	No algorithmic advantage
Actual gate development: 15 requests	Model creates a second leaf; shared projection repairs geometry; current state and tested use continue	Original fixed-post request violated by a 0.20 m displacement; full edit request failed
First induction pilot: 22 requests	Both inductions fail admission; nine subsequent outcomes fail; scoped/strong fallback request pairs coincide	Invalid comparison between distinct learned methods; SDK/feedback/evaluator issues retained
Corrected admission: 2 requests	Truncation, then rule and source/body-binding errors; strong arm unexecuted	No qualified skill or comparison
Minimal-edit admission: 2 requests	Undefined reference, then invalid arithmetic arity; strong arm unexecuted	No qualified transformation; author-only format correction is not a model result
Dependency repair, no model call	Authored stage works after repair; current-source-only analysis yields identical repaired source	Useful shared tool, no evidence that history is necessary
Relation versions, no model call	Authored choices and explicit revision alter subsequent work and preserve continued use	Finite library and supplied scope/revision, no model learning advantage

Table 2: Disposition of every predeclared opportunity; not an equal-protocol ranking. The stopping deviation is reported in Section 6.5.

Method / opportunity	Update status	Actual production and use	Counted outcome
Candidate / changed reason	A not admitted	Not executed	Precondition failure, 1 of 2 opportunities
Candidate / continued reason	A not admitted	Not executed	Precondition failure, 1 of 2 opportunities
Strong / changed reason	Memory contains offset_x	Complete program puts camera and support at 3 m; target 2.5 m; installation rejected	Rejected, 1 of 2 opportunities
Strong / continued reason	Memory contains world_x	Final edit request produces -0.5 m; installed at joint state 0.45 and used at 0.8	Used, 1 of 2 opportunities

struction, expressive bodies, high-quality presentation, creator growth and market value remain unmeasured.

7 Decisive evaluation required for the central claim

7.1 Complete production and revision

The primary task must include actual generation, an intervention that requires a nontrivial structural or behavioral change, continued use, and a later unseen production task. Initial generation and revision are scored separately. Revision arms receive an identical initial work and situation so that differences in starting difficulty do not masquerade as editing gains. Initial generation failures remain in the complete-chain results; they are not silently replaced with authored correct artifacts.

The primary outcome is joint satisfaction of the requested change, selected preservation obligations, and independent execution checks. Report each component as well. Geometry, executable behavior, appearance, latency, and total cost remain separate measures. A mean score across components must not conceal an unusable work. Crashes, timeouts, truncated outputs, invalid programs, and exhausted budgets are failures rather than omitted samples.

Mechanism families, body topology, and project composition determine train, development, and test splits. Renaming the same mechanism or changing its color is not cross-project transfer. Development feedback may inform the method, but final tasks must not have supplied selection evidence. The initial small development batch is not a population estimate and cannot support broad statistical claims.

7.2 Strong learning comparisons

Compare a strong full-history code agent, structured retrieval, standard skill induction including the operation-compilation design of SkillCompiler [26], and the proposed applicability-tested process learner. All receive the same available history and tools, with equal preparation and inference opportunities. A full-history agent may reflect, discover dependencies, construct skills, and run checks; withholding those capabilities would weaken the comparison artificially.

Removing scope or reasons is an information ablation and must be labeled accordingly. Removing applicability or counterexample validation can be a mechanism ablation if the original information and remaining opportunities are preserved. Final-artifact-only input is another information condition, not a same-information baseline. Preparation, retrieval, verification, unsuccessful attempts, and any training or search must all be counted.

Report paired outcomes by independent project or mechanism family, uncertainty at that level, and costs both before and after any stated amortization. A small number of families cannot justify stable population inference. Freeze the actual cases, code, prompts, method artifacts, thresholds, and scoring procedure before final evaluation. The current budget is a development envelope, not evidence that such a study has been completed. A binding intervention must also be separated from a learning intervention: compare the same accepted commitment with and without mandatory current-source binding, and separately compare how the commitment is learned or revised. A source-bound public editor must remain available to the strong agent. Useful new evidence requires a different structural or activity problem; the exposed camera pair cannot become a final test.

7.3 Reconstruction, presentation, and creator benefit

Real reconstruction requires real acquisition inputs and an appropriate measurement protocol. Point import or synthetic geometry is insufficient. Appearance reconstruction, editable structure, and inferred functionality have different ground truths. A system can reconstruct a surface accurately while inventing an incorrect articulation.

Presentation evaluation must compare rendered observations with executable state under changed views and interactions. Measure identity, geometry, occlusion, contact, action fidelity, temporal stability, frame rate, end-to-end latency, and all rendering costs. A beautiful still image does not establish a responsive or editable world.

Human evaluation must distinguish novice and professional creators. Measure complete time to an acceptable work, including waiting, clarification, rework, and manual rescue. Also examine control, understanding, exploration, participation, and the ability to continue creating. Interface comparisons hold information, capability, authority, and activity timing constant where possible. Learning comparisons hold the interface fixed. Neither an authored demonstration nor a synthetic preference judge establishes market demand or human benefit.

7.4 External tasks and broader applicability

WorldCoder-Bench is a relevant external target, but the missing public evaluation implementation remains a protocol dependency [4]. We have not obtained an official score or submitted a leaderboard entry. An adapted local task must remain separate from native benchmark results. Production or editing sub-benchmarks can support particular abilities without serving as an overall world-production ranking.

Enterprise production may concern executable tools, processes, or shared artifacts; embodied production may concern simulated objects, mechanisms, tasks, or control programs. Reuse requires an explicit mapping of objects, actions, observations, objectives, and validation. Organization-level learning additionally requires rules for whose experience is shared and who can revise a collective practice. Physical objects require manufacturing and real-world validation. A 3D demonstration cannot establish these mappings or outcomes merely because the interface vocabulary is shared.

8 Relation to the broader intelligence program

Taihe addresses human purposes, relationships, values, and the direction of intelligence and civilization. Nuwa addresses persistent subjects, understanding, representation, help, and growth. Fuxi addresses ongoing shared environments and the consequences of actions. Pangu addresses how works, objects, mechanisms, and presentations are produced and revised, and how production capability can improve. These are cooperating research questions, not interchangeable modules or evidence sources.

Pangu's human purpose includes expression, play, making things together, learning a craft, and enabling participation. The effects on these experiences should shape design and evaluation. They cannot be inferred from asset counts, rendering scores, or a single task reward. Similarly, an organization becoming more efficient does not automatically mean that its members gain more opportunity or agency.

The author's AI1.0-4.0 framing is a research orientation, not an established industry taxonomy or a capability certificate. The source discussion does not precisely define AI1.0. AI1.5 concerns world modeling; Pangu contributes the proposed production and inspection of spatial-temporal situations. AI2.0 concerns causal models; interventions on structures and mecha-

nisms motivate checks of actual consequences. AI3.0 concerns lingji in shared human-AI activity, including how experience, feeling, intuition, and creative choices can change what is worth pursuing. Canvas is one possible record of those changes, not their explanation.

AI4.0 concerns AI designing AI, potentially including mathematical formulations, architectures, and physical implementations. Pangu could provide environments and production tools for such work. Ordinary retrieval, a new skill, a prompt edit, and an improved learning procedure are different changes. None of the present evidence establishes AI4.0 or recursive improvement. Such claims require independent tests of the designed component and successive-cycle evaluation of the improvement process itself.

9 Discussion and conclusion

Pangu's practical ambition is to make creating, using, and revising a world parts of one continuing activity. Its contribution cannot be supplied by a compelling editor alone, or by the other papers in the broader program. The current defensible object is a revisable production commitment and its connection to subsequent executable works: what decision established it, what it applies to, when it loses authority, and which current source its execution changes. The restricted reference makes that lifecycle observable. Authored choices already change a later work that remains usable; the real-model record demonstrates both a functioning shared-tool branch and the separation between selecting a relation and executing it correctly.

The evidence does not yet show that the proposed organization improves learning or complete production over strong alternatives. Its next scientific obligation is a valuable work in which experience-derived decisions bind to a genuinely different production situation, with a measured gain that survives a capable same-information comparison. That requires more than fixing this interface or checking a formula. It requires useful changes to what a creator can make, preserve, revise, and continue doing. The full production architecture and human purpose remain the target; the reported restricted successes and failures define the present evidence boundary.

References

[1] Yue Yang et al. Holodeck: Language Guided Generation of 3D Embodied AI Environments. CVPR 2024. https://arxiv.org/abs/2312.09067

[2] WorldGen: From Text to Traversable and Interactive 3D Worlds. CVPR 2026. https://arxiv.org/abs/2511.16825

[3] Wei Wu et al. Code2Games: Enabling Coding Agents for Gaming World Generation. 2026. https://arxiv.org/html/2610.05033v1

[4] Shuo Lu et al. WorldCoder-Bench: Benchmarking Physically Grounded 3D World Synthesis. Version 3, 2026. https://arxiv.org/abs/2606.01869v3 ; official repository: https://github.com/shuolucs/WorldCoder-Bench

[5] Tripo. Agentic Canvas Guide. Official product documentation, October 2026. https://www.tripo3d.ai/product-doc/agentic-canvas/agentic-canvas-guide

[6] LAM. Official project and paper links. 2026. https://gaoypeng.github.io/LAM/

[7] Matt Zhou et al. Articraft: An Agentic System for Scalable Articulated 3D Asset Generation. 2026. https://arxiv.org/html/2605.15187v1

[8] Xiaohan Peng, Debanjana Haldar, Wendy E. Mackay, and Janin Koch. DesignTrace: Exploring, Iterating and Tracking Design Alternatives with GenAI. CHI 2026. https://doi.org/10.1145/3772318.3791036

[9] Xiujin Liu et al. DesignAgent3D: Interactive 3D Scene Editing via Designer-like Multimodal Reasoning. 2026. https://arxiv.org/html/2608.21438v1

[10] 3DCodeEdit: Benchmarking Instruction-Guided 3D Program Editing. 2026. https://arxiv.org/html/2610.38231v1

[11] VibeWorlding. Version 2, 2026. https://arxiv.org/html/2608.15265v2

[12] Jacob Jackson, Ben Trapani, Nathan Wang, and Wanqi Zhu. Improving Composer through real-time RL. Cursor research, 26 March 2026. https://cursor.com/blog/real-time-rl-for-composer

[13] Guanzhi Wang et al. Voyager: An Open-Ended Embodied Agent with Large Language Models. 2023. https://arxiv.org/abs/2305.16291

[14] Jack Parker-Holder et al. Evolving Curricula with Regret-Based Environment Design. 2022. https://arxiv.org/abs/2203.01302

[15] The Trinity Research Team. The Trinity: Capability-Driven Recursive Self-Improvement for Multimodal Agents. 2026. https://arxiv.org/html/2610.36077v1

[16] Jianyuan Wang et al. VGGT: Visual Geometry Grounded Transformer. CVPR 2025. https://vgg-t.github.io/

[17] Bernhard Kerbl et al. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Transactions on Graphics, 2023. https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/

[18] Roblox. Accelerating Creation, Powered by Roblox's Cube Foundation Model. 4 February 2026. https://about.roblox.com/newsroom/2026/02/accelerating-creation-powered-roblox-cube-foundation-model

[19] Roblox. Roblox Reality: A Hybrid Architecture for Democratizing Photorealistic Multiplayer Gaming. 30 April 2026. https://about.roblox.com/newsroom/2026/04/roblox-reality-hybrid-architecture-democratizing-photorealistic-multiplayer-gaming

[20] Roblox. Q3 2026 shareholder letter. SEC exhibit, 2026. https://www.sec.gov/Archives/edgar/data/1315098/000162828026068235/ex991-q32026earningsshar.htm

[21] Google DeepMind. SIMA 2: An agent that plays, reasons and learns with you in virtual 3D worlds. Official research account. https://deepmind.google/blog/sima-2-an-agent-that-plays-reasons-and-learns-with-you-in-virtual-3d-worlds/

[22] TheAgentCompany: Benchmarking LLM Agents on Consequential Real World Tasks. 2024. https://arxiv.org/abs/2412.14161

[23] Stanford. BEHAVIOR: Benchmark for Everyday Household Activities in Virtual, Interactive, and Ecological Environments. Official project and challenge links. https://behavior.stanford.edu/

[24] Martin Abadi and Leslie Lamport. The Existence of Refinement Mappings. Theoretical Computer Science 82, 1991. https://lamport.azurewebsites.net/pubs/abadi-existence.pdf

[25] Khronos. glTF 2.0 Specification. https://registry.khronos.org/glTF/specs/2.0/glTF-2.0.html

[26] Chengyu Yang et al. Compile Away Uncertainty: Self-Evolving Agent Skills via Programmatic Compilation (SkillCompiler). Version 1, 29 October 2026. https://arxiv.org/html/2610.39508v1

[27] Shuyan Huang et al. Do Personalized Skills Help Coding Agents? An Empirical Study of Developer Interaction Traces. Version 3, 28 October 2026. https://arxiv.org/abs/2608.10319v3 (abstract reviewed).

[28] CIT-CAD: Constraint Intent Tree-based CAD Code Generation and Verification. Version 1, 2026. https://arxiv.org/html/2609.07434v1 (method overview and Section III-B reviewed; not reproduced).

[29] Aligning Constraint Generation with Design Intent in Parametric CAD. Version 1, 2025. https://arxiv.org/html/2504.13178v1 (relevant methods in Sections 3–4 reviewed; not reproduced).

[30] Tom M. Mitchell. Version Spaces: A Candidate Elimination Approach to Rule Learning. IJCAI, 1977. https://ijcai.org/proceedings/1977/0068.pdf

10 Artifact availability and AI assistance

The accompanying evidence archive includes synthetic inputs, original model responses, usage records, saved programs, activity receipts, and the reference code needed to inspect them. The evidence index distinguishes authored controls from model-produced artifacts, unexecuted opportunities from rejected works, and proposed mechanisms from implemented operations. Original scores and usage are retained. Packaging replaces private workstation paths with portable artifact paths; the mapping and original/public hashes are recorded without exposing private source locations. No new model experiments were performed for this publication package.

AI assistance through Codex was used for literature exploration, implementation, experiment orchestration, analysis, figures, and manuscript preparation. The author is responsible for the release. Experimental model outputs are separately identified by the fixed Qwen model version and their raw response records. Authored controls and deterministic tool results are not presented as model discoveries. The current evidence does not establish a new learning algorithm, a general productivity advantage, or human benefit.

Original material retains copyright; this release does not grant an open-source license. Third-party material remains subject to its respective terms. The public artifact index and suggested citation accompany the paper at the DeepBean research site.