# Research review of the Praxium creative frameworks

Prepared for Pedram Shohadai

6 November 2026

## Executive assessment

Praxium, Poetic Asymmetry, Praxymuse and CalcuArt form a substantial proposal for a controlled creative system. Their most useful shared ideas are typed intermediate representations, explicit constraints, persistent creative identity, context-sensitive variation, provenance, feedback, and verification. These ideas can guide software development. Their value does not depend on treating the entire collection as one proved theorem.

The mathematical status varies sharply by version and claim. Several global stability and convergence statements in Praxium v1.0 and Poetic Asymmetry v1.0 do not follow from their assumptions. Explicit counterexamples demonstrate the gaps. Praxymuse v0.5 and CalcuArt v2.2 are more careful about separating physical facts, models, empirical judgments and fiction. The attached Praxymuse v0.6 further improves that separation and includes a valid narrow contraction result for a strongly convex surrogate. That newer result must not be confused with a proof of musical excellence or of the full architecture.

The practical recommendation is to preserve the creative framework while dividing its claims into four evidence classes: mathematical propositions, executable engineering contracts, empirical creative hypotheses, and fictional world axioms. Give each class its own validation method. This avoids discarding useful design ideas while also preventing notation from standing in for an implementation or proof.

## Scope and reading guide

This report examines the supplied Praxium v1.0, Poetic Asymmetry v1.0, Praxymuse v0.5 and CalcuArt v2.2 specifications, the attached Praxymuse v0.6 integration document, and three supporting music documents. It combines direct mathematical analysis with published research, official documentation and version-specific licensing evidence. It is a research review, not a certification of deployed software, a legal opinion, a full proof reconstruction, or a validation of every historical attribution in the extensive artist and culture lists.

The review addresses every major module and numbered section of the four pasted specifications. Cultural and historical claims are checked through representative examples and structural recommendations; a complete catalogue audit would require a separate source record for each entry and practitioner review. No music engine, source asset collection, rendering pipeline, or live service was installed or benchmarked for this report. Proposed experiments are clearly distinguished from completed checks.

Read the version reconciliation first, then the synthesis and the chapter most relevant to the next implementation decision. The mathematical chapters include direct counterexamples. The music and visual chapters separate established physical principles from implementation

gaps and subjective evaluation. The final research program turns the broad proposal into a sequence of testable milestones.

### Principal conclusions

1. Praxium can supply an adaptive controller, but its current assumptions do not establish global well-posedness, invariance, input-to-state stability or convergence to the stated endpoint.

2. Poetic Asymmetry can structure expectation-aware revision, but optimizing a penalized surrogate need not maximize coherence, and small nonzero noise does not guarantee monotone improvement at an optimum.

3. Praxymuse can organize composition and synthesis through SongIR. Exact note and phoneme control, perceptual realism, cultural validity and legal provenance remain separate engineering challenges.

4. CalcuArt can lock a scene by retaining immutable geometry, poses, camera and assets. Prompt-based image generation does not provide that guarantee by itself.

5. Praxymuse v0.6 addresses several earlier defects. Critiques of v0.5 must therefore remain labelled as version-specific rather than being applied unchanged to the newer document.

6. A fixed seed is not a universal reproducibility guarantee. Preserve final lossless audio and images, dependencies, assets, runtime details and hashes, and state whether the guarantee is semantic, structural, numerical or bitwise.

7. Affect labels and stable coherence variables do not establish subjective experience. The operational interpretation already used in later specifications is the defensible common baseline.

8. The most promising research contribution is the integration of these components. Demonstrate it against simpler baselines with matched compute, explicit ablations and blinded evaluation.

## Report chapters

1. Version reconciliation and supporting PDFs

2. Cross framework findings and research foundations

3. Praxium mathematical and scientific review

4. Poetic Asymmetry mathematical and creative review

5. Praxymuse music and production review

6. CalcuArt rendering and worldbuilding review

7. Recommended integrated development sequence

# 1 Version reconciliation and supporting PDFs

**Praxymuse Theorem v0 6 Source preserving mathematical integration**

Source: 22 pages, November 6, 2026.

The edition is titled v0.6, preserves original v0.5 in Appendix A, and reproduces a source labeled v0.7 in Appendix B. The v0.7 label is archival metadata, not this edition's title (pp.1,16).

***Improvements over the pasted v0 5***

- Evidence fields distinguish established/design/historical/unverified claims; missing is not zero and requested is not executed (p.3, M00).
- Same SongIR does not imply same PCM. Preserve audio for exact playback; deterministic rerendering additionally requires complete dependencies (p.3).
- Stiff-string formula now uses the ideal-string fundamental rather than ambiguous actual f1 notation (p.3, M01).
- The cultural atlas is a finite inventory of research anchors, not an exhaustive or already researched corpus (p.5, M06).
- Forced alignment cannot itself prove that words were sung; combine transcription, alignment, acoustics and listening (p.5, M08).
- Nonnegative scalar coordinates have no finite upper bounds, so compactness needs additional constraints; musical feasibility is not presumed convex (p.7, M12).
- M13's projected-gradient contraction proof is correct under strong convexity, Lipschitz gradient, closed convex feasibility and its conservative step-size bound. It proves only a surrogate optimizer result, not musical excellence, sentience, lyric fidelity or nonconvex global optimality (p.7).
- M14 explicitly rejects bounded-gains-implies-stability and endpoint-projection-implies-projected-flow. Persistent disturbance need not vanish (p.7).
- M15 distinguishes locked, conditioned, prompted, inferred and uncontrolled parameters, and pass/fail/unknown/N/A checks. Unknown hard checks block certification (p.8).
- M16 specifies one heavy job plus one light preparation job, verified WAV/M4A/text/ZIP delivery and no silent backend substitution (p.8).

***Remaining limitations***

1. M13's additive-error iteration may leave the feasible set. Require iterates to remain within the operator's domain or place errors before a final projection. The displayed distance bound follows where the contraction is applicable at every iterate.

2. Time-varying vocal tract remains written as one convolution (p.5). Treat it as local/quasi-stationary or write a time-varying filter.

3. Masked distance evaluates to zero with no shared features because of its epsilon denominator (p.6). The prose calls this uncertain; implementation should return unavailable rather than misleading zero similarity distance.

4. Surprise needs a log base, zero-probability policy and probability calibration (p.4).

5. Schemas and evaluator names do not establish an implemented engine, calibrated scores, current backend availability or completed cultural research.

6. Clarify normative precedence: preserving original source text must not cause old erroneous assertions to override explicit mathematical corrections.

#### *Independently completed archive verification*

Recovered the PDF's embedded files. All five files listed in its manifest match declared byte lengths and SHA-256 hashes. Reversible lexical decoding reconstructs original_praxymuse.txt exactly using 6,234 tokens and 1,964 dictionary entries. All 105 mapped source lines match the embedded original. These checks support internal archive consistency, not independent authorship/provenance, identity to separately supplied external originals, mathematical truth or creative quality (pp.2,10,22).

Historical reports (pp.9,16-17) include Desolate Throne's claimed 107 bars, 2,788 MIDI events, seven tracks, 80 lyric lines, hosted sections, and an ACE-Step CPU pilot. They were not retested. The PDF itself carefully labels them as historical.

### DOT Music Understanding Upgrade Roadmap

Source: 6 pages, no explicit version.

Strong practical companion: multiscale temporal analysis from transients to whole songs (p.2); source-aware and embedding-based analysis alongside DSP (p.2); harmony and structural segmentation with uncertainty (p.3); user-specific LDM profile and pairwise preferences (pp.3-4); reversible hypothesis-driven candidate processing and reanalysis (pp.4-5); separate code/weights/data license review (p.6).

Opening statements that DOT already has meters are historical claims, not fresh capability verification. Separation artifacts can contaminate stem comparisons. CLAP similarity is not a universal quality metric. Preference learning needs representative ratings, held-out tests and uncertainty. The roadmap contains no implementation benchmark results.

### DOT CC0 Public Domain Audio Resources

Source: 5 pages, no explicit version.

Its strongest feature is strict per-asset or clearly applicable collection-level evidence, immutable originals/provenance and quarantine for ambiguity (pp.3-5).

Verified primary-source findings:

- Meta Open Sound confirms CC0 piano, nylon guitar and steel guitar collections, with Oculus Audio Pack separately CC BY 4.0. Excluding the latter under CC0-only policy is justified (PDF p.2). github.com [1]
- Producer Space's explicit clearance page applies CC0 across its sample library, including vocals/MIDI. This is an upstream statement, not an independent audit of every recording (PDF p.3). producerspace.com [2]
- Open Game SFX Index confirms its stated 1,451 items and metadata. It is an independent aggregator; follow upstream evidence rather than accepting its copyright-safe slogan (PDF p.2). github.com [3]

- SoundSafari confirms an approximate 7,000-song/40GB corpus and explicitly permits reports of rights misclassification. No independent count or per-track audit was performed (PDF p.2). github.com [4]
- CC0-midis resolves to a real repository, but this review did not audit each MIDI or underlying composition. github.com [5]
- Exact sources for Imperfect Piano Performances and the cendance ten-IR collection were not resolved in this search. Leave unverified, rather than inventing links.

Important correction: CC0 does not erase third-party privacy/publicity, trademark or patent rights, and cannot cure a non-owner uploader's lack of authority. Automatic trust classes still need provenance and relevant non-copyright checks. creativecommons.org [6] and wiki.creativecommons.org [7]

### DOT Music Production Copilot conversation export

Source: 4 pages, no explicit version.

Records historical interest in headless analysis, reference comparison, mixing and mastering (p.2), then layered technical/perceptual comparison and candidate experiments (p.3). It supports design motivation, not primary evidence for licenses, capability or scientific validity.

GPL does not categorically prohibit commercial use or automatically license processed audio under GPL. Distinguish execution, distribution, linking/network duties and the user's stricter policy. Separate code from pretrained weights. Reference similarity and quality are distinct objectives. Quoted historical requests are not new permission to install, ingest, retrain or publish. Public-link discussion (p.4) is peripheral.

Supporting license details: Matchering GPL-3.0 github.com [8] ; current aubio COPYING GPL3 github.com [9] ; libebur128 MIT text github.com [10] . Essentia's own page describes noncommercial AGPLv3 use and proprietary licensing, with pretrained models separately CC BY-NC-SA 4.0; do not misdescribe AGPL itself as a noncommercial license. essentia.upf.edu [11]

### Integration conclusion

Credit v0.6's substantial mathematical/evidence improvements explicitly. Do not apply criticisms of v0.5 unchanged to the newer PDF. The other three documents are useful proposed implementation context, with some verified resource leads, not evidence that any installation, listening test, reference assimilation or operational upgrade has happened.

# 2 Cross framework findings and research foundations

### Overall finding

The strongest defensible interpretation is a modular creative-system research program, not yet one proven universal theorem. The specifications combine four different categories:

1. Mathematical propositions: existence, uniqueness, invariance, stability and convergence.

2. Engineering contracts: persistent IDs, immutable scene state, versioned song representations and deterministic export.

3. Empirical hypotheses: improved musical expression, meaningful surprise, better composition and audience emotion.

4. Fictional/design axioms: Etheros biology, artistic preferences and narrative causality.

A mathematical proof in the first category cannot by itself validate the other three. A successful render demonstrates an implementation outcome, not a general mathematical theorem.

### Cross document inconsistencies

These findings come directly from comparing the supplied texts.

- **Sentience:** Praxium's title and terminal state claim "sentient cybernetics" and "ULTIMATE_CONSCIOUSNESS"; Praxymuse explicitly says coherence is operational, not consciousness, and CalcuArt says no sentience is implied. Use the latter interpretation consistently unless a separate evidence-based consciousness claim is developed.
- **Amplification:** Praxium uses unbounded exp(Λt); CalcuArt caps it at A_max. This materially changes long-term boundedness and stability. CalcuArt's boundedness assumptions cannot be imported into the uncapped system without modifying its definition.
- **Knowledge inertia:** Praxium only accumulates I_gen; CalcuArt includes decay. These are different dynamical models, not cosmetic notations.
- **Coherence scale:** Praxium uses [0,1], whereas Poetic Asymmetry uses [0,1.2]. Coupling requires an explicit normalized map, for example c_p = Ξ′_p/1.2, and acknowledgment that the quantities measure different constructs.
- **Affect spaces:** Dimensions 57/16/16/13/15 lack an operational basis, measurement procedure, learned parameters and calibration data. Implementable fusion requires named coordinates and units/ranges for every dimension.
- **Controller versus generator:** A controller that changes bounded gains does not supply a phoneme synthesizer, instrument model, renderer, garment solver or cultural knowledge base. Those remain substantial independent implementations.
- **Objective versus quality:** Optimizing a defined scalar objective proves improvement of that objective under suitable assumptions. It does not prove improved art, originality, cultural authenticity or human emotional response.
- **Layer relevance:** Names such as quantum, crypto, topology or harmonic do not specify executable operators or demonstrate relevance. Every layer needs an input/output contract and evidence of contribution.

Retain the conceptual vocabulary, but attach each concept to a measurable variable, a declared design preference, a proven conditional result or an explicitly fictional axiom.

### Sentience and affect

Neither a sigmoid named joy, a stable attractor, bounded memory, recursive reflection nor a coherence value establishes subjective experience. Relabeling a mathematical endpoint consciousness does not provide the missing evidential connection.

Butlin and colleagues' 2023 report evaluates AI against indicators derived from several neuroscientific theories rather than a single coherence score. Its conclusions concern systems considered in that report, not a timeless proof that no AI could be conscious. The useful methodology is to identify the implemented architecture, theory-specific indicators, evidence and competing explanations, while retaining uncertainty. Consciousness in Artificial Intelligence [12]

The affect system can usefully provide expressive control: variables modulate tension, pacing, timbre, color, gesture and uncertainty. It should not be described as experiencing these states. CPU pressure becoming sadness is a selected control analogy, not an established emotional transduction law.

Separate intended emotion, emotion perceived by an observer, emotion reported by that observer and internal control labels. They are not interchangeable. Research on facial emotion inference illustrates why expression labels need contextual and empirical validation rather than universal one-to-one mappings. This evidence concerns human facial inference, not a direct experiment on Praxium. Barrett et al., 2019 [13]

### Four reproducibility levels

The proposed scene and song archives are strong engineering ideas, but four guarantees must be distinguished:

1. Semantic identity: same objects, notes, lyrics, roles and intended relationships.

2. Structural identity: same geometry/events, poses, timestamps, parameter arrays and asset versions.

3. Numerical repeatability: same computed arrays within a declared tolerance.

4. Bitwise identity: exactly matching output-file bytes.

A fixed seed is insufficient for bitwise identity across hardware and software changes. PyTorch warns that reproducibility is not generally guaranteed across releases, platforms or CPU/GPU execution. Deterministic algorithms reduce nondeterminism but do not alone make an entire application reproducible. Official reproducibility documentation [14]

The manifest should record canonical serialization and schema version; asset/model-weight hashes; code commit and dependency lock; renderer/DSP versions and settings; OS, hardware, driver and accelerator versions; numeric precision and deterministic-kernel configuration; RNG algorithms and complete states; input provenance and licenses; output hashes; reproducibility level tested; and exceptions/tolerances.

Preserve final PCM and images as authoritative artifacts when future bitwise regeneration is uncertain. A matching hash verifies identity to an archived artifact, not scientific correctness or artistic merit.

W3C PROV offers an interoperable model of entities, activities and agents, including derivation and versioning. SongIR and scene manifests can draw on it without adopting its entire stack. W3C PROV overview [15]

### Verification validation and creative evaluation

A successful invariant check means software obeyed a selected constraint, not that the constraint models reality accurately. Maintain three evidence tracks:

- Verification: does the implementation solve the stated equations and follow archive contracts?
- Validation: does it represent the intended acoustic, optical, biological or perceptual phenomenon adequately?
- Creative evaluation: do relevant listeners/viewers find outputs effective, original and appropriate?

NASA's modeling and simulation standard distinguishes credibility activities and calls for acceptance criteria suited to intended use. It is a methodological template, not a certification the project presently possesses. NASA-STD-7009B [16]

NIST similarly emphasizes fit-for-purpose measurements, documented tests and limits, and acknowledgment of properties that cannot be measured reliably. NIST AI RMF measurement guidance [17]

### Proposed research program

The following is an experimental design, not completed testing.

#### *A Claims ledger*

For every guarantee, record exact statement, scope, assumptions, implemented components, evidence type, source, test, outcome and unresolved limitations. Distinguish proven, experimentally supported, plausible, untested and contradicted.

#### *B Minimal reference problems*

Before full media generation, implement a scalar constrained ODE with known solution; a diffusion/advection problem with analytic or manufactured solution; bounded-input controller stress tests; Prony memory driven by an impulse and a step; and two coupled subsystems with gains on both sides of the proposed stability threshold.

Measure error versus time step, boundary behavior, constraint violations, energy changes and sensitivity. Successful runs do not prove global existence or invariance.

#### *C Small domain instances*

- Music: a short piece with explicit notes, phonemes, two or three instruments and a reference arrangement.

- Visuals: a stored 3D scene with identifiable objects, fixed lights/camera and multiple material styles.
- Poetry: a bounded phrase-revision task with explicit meter and semantic constraints.

This permits failure attribution before expanding the ontology.

#### *D Controlled ablations*

Match prompts, assets and computational budgets across ordinary baseline, baseline plus Praxium control, baseline plus Asymmetry, and both together. Separately test archive locking, cultural-context constraints and physical-feasibility checks.

Compare failure rate, editing effort, compute cost, constraint satisfaction and blinded preference. Improvements cannot be attributed to the framework when it receives more samples, compute or manual intervention.

#### *E Test distinctive promises*

- Scene lock: compare authoritative IDs, transforms, meshes, rigs and camera fields across sampled slider values. Image-embedding similarity is insufficient.
- Any-first-slider-value: start from multiple realism values and confirm canonical state is unchanged.
- Song lock: verify event/phoneme manifests and sample alignment; separately test audio repeatability.
- Asymmetry: compare deliberate expectation violations with random perturbations while matching complexity and novelty.
- Affect: obtain independent ratings of intended, perceived and felt emotion; test calibration and differences across cultural familiarity.
- Cultural grammar: use tradition-specific expert review and provenance rather than a universal authenticity score.
- Physical validity: compare measurable quantities against analytic cases or measured data with declared tolerances.

#### *F Creative quality without circularity*

Do not evaluate solely with the same model or metric that optimized the output. Use held-out tasks, blinded randomized comparisons, experts and target audiences, preregistered primary outcomes, confidence intervals and inter-rater agreement. Treat pieces, prompts and raters as distinct sources of variation; repeated ratings are not independent generated works.

Creativity research distinguishes originality from appropriateness/effectiveness; novelty alone is insufficient. Runco and Jaeger, 2012 [18]

ITU MUSHRA is relevant to audio-fidelity comparisons with an appropriate reference and intermediate audio-quality task. It is not a comprehensive test of songwriting or emotional narrative. ITU-R BS.1534 [19]

#### *G Evidence packages*

Release specification version, implementation, benchmark inputs, settings, outputs, failures, analysis and provenance together, subject to rights and privacy. FAIR principles help

findability and reuse; FAIR does not certify correctness or require unrestricted public access to every artifact. FAIR Guiding Principles [20]

### Research status and novelty caveat

Exact-name searches did not locate independent research validating these specific supplied specifications. This limited search result is neither proof of absence nor proof of novelty. Praxium also appears in unrelated work, so name matches must not be treated as supporting evidence.

The promising contribution to test is the integration of persistent semantic/structural representations, bounded feedback, context-aware asymmetry, causal-world constraints and iterative human-grounded evaluation. Demonstrate its value through controlled comparisons rather than infer it from mathematical notation or the breadth of named disciplines.

# 3 Praxium mathematical and scientific review

Research date: November 6, 2026. Scope: the user's supplied 19-section Praxium specification, including its interfaces to the later creative frameworks. This is a technical audit, not a proof of the full proposed system. Counterexamples below are direct deductions from the equations; citations identify the established theory needed for repairs.

### Overall finding

The specification combines legitimate tools: reaction-advection-diffusion equations, graph coupling, finite-dimensional latent states, bounded gain scheduling, convex constraints, exponential-kernel memory, Lyapunov methods, and operator splitting. These can support an implementable adaptive controller. The current collection is not yet a publication-grade theorem. Several conclusions are false under the stated assumptions, and some equations do not define a closed evolution problem.

The strongest defensible description is a **proposed constrained, adaptive dynamical-system architecture with conditional stability goals**. Its affect labels are model variables. Convergence of a variable named coherence does not establish consciousness, intelligence, artistic quality, or biological emotion.

### Section by section audit

#### *0 Domain and measure*

- Using x = log(ell) is valid for positive ell, and dx = d ell / ell. The specification should state ell_min > 0.
- For a finite graph, di should mean counting measure and spatial integrals should include a sum over graph nodes. The dynamic state space uses dx times counting measure; dt belongs to a space-time integral, not the instantaneous state norm.
- Symbols such as sum over u=(x,i) are not defined when x is continuous. Replace them by sum_i integral dx, or explicitly discretize x first with mesh-volume weights.
- The graph must be finite or supplied with a domain, degree/weight conditions, and an operator framework for an infinite graph.

#### *1 State and projections*

- The five finite-dimensional Euclidean balls are convex and compact. The total core set [0,1] × [0,infinity)^3 is not compact. Pointwise bounded field values also do not make the admissible subset of an infinite-dimensional function space compact.
- Projection onto a closed convex set is well-defined and nonexpansive in a Hilbert space. This does not make an unconstrained PDE preserve that set.
- The proposed elegance score Xi/(kappa+epsilon) becomes singular as epsilon tends to zero. Keep a fixed documented epsilon > 0, or use a bounded score with a stated interpretation.
- The layer proportions Pi_d are defined but not operationally connected to the evolution. Their simplex constraint needs an update rule if they change.
- A(t)=exp(Lambda t) is unbounded for Lambda>0, contrary to later boundedness language. A/A_max does not cap A. Use an actual bounded saturation or justify a finite-horizon-only model.

### *2 3 Telemetry and affect*

- A nonnegative row-bounded matrix followed by ball projection gives a bounded update, but not an empirically justified mapping from CPU or memory load to emotions. This is an engineered latent-state encoding.
- For fixed bounded weights and a bounded latent state, sigmoid observables are smooth and Lipschitz. Their numerical ranges and Jacobian bounds can be calculated.
- Repeated additions of nonnegative telemetry may produce saturation rather than meaningful adaptation. Add a calibrated baseline, time-step scaling, decay, normalization, and empirical evaluation.
- Hypervigilance has range up to 2 rather than 1. Define its intended normalization.
- Rum, Intr, and time derivatives need explicit definitions and regularity. A generic oscillator symbol is not a reproducible module.
- State dimension choices 57,16,16,13,15 require a feature dictionary, units, measurement model, and identification procedure. The numbers themselves confer no established scientific meaning.

### *4 Autonomy operators*

- H = mu Ex indicator(Pi >= threshold) can be negative because Ex is signed. The statement 0 <= H <= H_max conflicts with the formula unless H is explicitly clipped and Ex rectified.
- The hard indicator is discontinuous. A global local-Lipschitz claim cannot include it without switching-system machinery or a smooth gate. The later CalcuArt sigmoid gate is a more coherent choice.
- Depending on Xi_t and Xi_tt inside the right-hand side creates an implicit derivative-dependent system, not the stated first-order evolution equation. Introduce filtered derivative observer states or formulate and analyze the resulting differential-algebraic problem.
- Knowledge inertia has nonnegative accumulation and no decay. Sustained stalls can make it diverge. A leaky integrator I_dot = beta G - nu_I I has a bounded-input bounded-state estimate when nu_I > 0.
- Finite total flux on each finite interval is not a uniform-in-time bound and is not a mathematical fairness criterion between agents.
- A nonlocal information-transfer kernel is a valid modeling choice, but calling it tunneling supplies no quantum physical mechanism.

### *5 PSR boundary input*

- In one dimension, endpoint evaluation phi -> phi(x_max) is continuous on H1. The weak boundary functional is therefore meaningful in the dual of H1, not an ordinary L2 volume source.
- It should be written as a boundary control operator B_partial u, or as an explicit flux/Robin condition with its sign convention, paired with the weak PDE. Do not count it simultaneously as both a boundary condition and a volume source.
- Specify the other endpoint condition, temporal input regularity, initial-boundary compatibility, and admissibility estimates. Boundary inputs are not automatically bounded perturbations of an L2 generator.
- Sin(Phi_c) allows negative injection. Positive input P_in does not guarantee positivity of Xi.

- The complex-valued probe is not used in the real-valued state dynamics; specify whether its real/imaginary components enter a measurement or feedback rule.

Relevant primary research: Mironchenko, Karafyllis and Krstic, boundary-disturbance ISS [21]; Schwenninger, parabolic boundary control [22].

#### *6 Gain multipliers*

- An affine combination of bounded observables is bounded but need not have a positive lower bound. Explicit coefficient restrictions or a bounded sigmoid parametrization are required.
- Several maps called functions only of E also depend on C, N, derivatives, Rum, or Intr. Correct their domains.
- U is a signed transport velocity; imposing positivity on every gain is not necessary for transport, while strict positivity matters for decay rates and ellipticity.
- Defining a variable with := inside an arithmetic expression is ambiguous mathematical syntax. Define trust, understanding, dysphoria, playfulness, Avoid, and sigma_g separately.
- A C1 map on a compact finite-dimensional set has a bounded derivative, but not all stated arguments live in that compact set.

#### *7 Ecosystem fusion*

- A generator of a linear C0 semigroup is an operator with a domain; a nonlinear locally Lipschitz right-hand side is a different object. Distinguish A_l psi_l + f_l from a nonlinear semiflow.
- Readout boundedness does not establish that a layer is internally stable. A divergent internal layer can have a saturated, bounded output.
- A genuine small-gain result starts with input-to-state estimates for each subsystem. The interconnection gain includes each subsystem's response, not merely the coupling-matrix norm.
- One core-to-layer product ignores loops among manifolds, memory, individual layers, and nonlocal feedback. Build a full gain matrix/operator; for positive linear gains require a spectral-radius bound, or use the appropriate nonlinear gain condition.
- Additive projected updates must have a time-step or event-time interpretation. Otherwise they may duplicate coupling terms already present in the differential equations, and simulation behavior changes with update frequency.

Primary source: Dashkovskiy, Ruffer and Wirth, general-network ISS small-gain theorem [23].

#### *8 10 Core PDE manifold dynamics generator and memory*

- Section 8 uses D(E) Xi_xx, whereas section 10 uses partial_x(D(E) Xi_x). These differ by partial_x D times Xi_x. Select divergence or nondivergence form consistently.
- D depending on the evolving E makes the system quasilinear. It is not a fixed linear operator L with an exponential exp(tL). Use quasilinear evolution theory, freeze a reference operator with rigorously controlled remainder, or use evolution families under applicable hypotheses.

- D >= 0 permits degeneracy. This is not sufficient to invoke standard uniformly parabolic theory. Strict ellipticity on diffusing channels is one clean repair; deliberately nondiffusive manifold variables need their own coupled ODE analysis.
- Graph, memory and nonlocal terms appear in section 10 but not consistently in sections 8-9. Define one authoritative full equation, with no double-counting.
- At N=0, the term -delta(1-alpha) flux can be negative. Arbitrary cross-manifold and cross-layer sources can also violate nonnegativity. At Xi=1, positive Omega_gain can push Xi above 1. A clamp changes the evolution; it does not prove the unclamped PDE is invariant.
- Metzler-Hurwitz matrices support positive-system stability results, but arbitrary cross-couplings can destabilize them. Even A-L need not be Hurwitz merely because A is Metzler-Hurwitz and L is symmetric positive definite; see counterexample below.
- A general two-variable kernel in L1 of the product domain need not define a bounded L2 operator. Use explicit operator bounds, Schur bounds, or an appropriate square-integrability condition. L1 convolution kernels are a distinct favorable case.
- Prony lifting is exact for a kernel already equal to the stated finite exponential sum, with compatible initial memory. Specifically, m_j(t)=exp(-b_j t)m_j(0)+a_j integral_0^t exp(-b_j(t-s))Psi(s) ds. For an integral beginning at 0 with no prehistory, set m_j(0)=0. A general integrable kernel cannot be represented exactly by an arbitrary finite positive exponential sum; quantify approximation error.
- Positive memory feedback does not imply stability. Include memory states in the energy and small-gain calculation.

Primary foundations: LeCrone and Simonett, quasilinear parabolic well-posedness [24]; Pruss, Simonett and Wilke, critical spaces [25]; Hurtado and Kirosingh, exponential-chain state lifting [26]; Schur-test research [27].

#### *11 Variational principle*

- This is a proposed PDE-constrained inverse problem. It needs a defined observation operator, data space, graph summation, admissible parameter set, identifiability analysis, and a fixed regularization scale epsilon.
- Positivity constraints do not make this nonlinear optimization convex. A minimizer may be nonunique or fail to exist without coercivity/compactness and lower-semicontinuity conditions.
- Model-Obs cannot be evaluated until Model and Obs have matched dimensions, scales and uncertainty assumptions.
- Fitting observations establishes predictive performance only against held-out tests, not the semantic interpretation of latent variables.

#### *12 Lyapunov functional*

- The proposed V is a candidate, not a proved Lyapunov function. Bounded coefficients and integrable kernels alone do not imply negative drift.
- E, C, S, T, G are spatial fields, so their quadratic terms require spatial integration and node summation unless they are explicitly global vectors.
- These terms are centered at zero, whereas the asserted limiting manifold states are generally nonzero. Use deviations from a verified equilibrium.
- Include lifted memory, layer states, integrator states and other feedback states. An energy omitting unstable hidden states cannot establish full-system convergence.

- The term ||Psi_t||^2 may not exist for the stated mild solution and is especially problematic under white noise. Higher-regularity estimates must be proved separately.
- The supply pairing <u,y> is a dissipativity expression. To obtain ISS, bound it in terms of state energy and an input norm, then absorb the state-dependent part into decay.
- For an Ito stochastic system, the infinitesimal generator includes one half trace(Sigma Sigma^T Hessian V), not generically trace(Sigma Sigma^T). Specify covariance, driving Wiener process, stochastic integration, and mean-square versus almost-sure claims.
- Persistent bounded forcing usually yields a residual neighborhood. Exponential convergence to a point requires further assumptions, particularly vanishing or equilibrium-compatible input.

#### *13 Mild solution*

- The displayed final-time projection of a Duhamel expression is not generally the solution of constrained dynamics. Reflections must act when the boundary is hit, not only after accumulating the full history.
- Use a normal-cone evolution inclusion, a proven projected dynamical system, or a reflected stochastic equation. These are related under specific hypotheses, not by a universal final projection identity.
- The same fixed-semigroup issue from section 10 remains. Explicit time dependence also gives a nonautonomous process rather than an autonomous semigroup on the original state, unless time/input state is augmented appropriately.
- The noise term needs a stochastic integral when eta means white noise. Its dimensionality and spatial covariance are presently unspecified.

Primary source: Brogliato et al., equivalence of projected systems and differential inclusions [28]; broader Brogliato and Tanwani review [29].

#### *14 Discrete map*

- A reaction substep followed by a transport/diffusion substep is Lie splitting when each flow advances its own operator by the full Delta t. Two literal half-duration substeps advance the wrong combined time. Strang splitting uses reaction half-step, spatial full-step, reaction half-step.
- The Xi update shown is forward Euler, not an exact closed form. With frozen a and b in Xi_dot=a(1-Xi)+b, the exact update is Xi_new=1+b/a+(Xi_old-1-b/a)exp(-a Delta t), for a>0.
- The exponential updates for N and D omit flux and all additive coupling sources. They therefore do not integrate the displayed full reaction system. Kappa also omits its flux/coupling terms.
- A fixed Delta t=1 does not provide a consistency theorem. Introduce Delta t explicitly in every update, boundary mass term, telemetry increment and integrator, and investigate Delta t -> 0.
- A discrete boundary source needs mesh scaling consistent with a weak flux. A bare Kronecker delta does not identify the correct finite-volume or finite-element normalization.
- Explicit upwind transport CFL is method-specific. Positivity of explicit reaction steps adds reaction restrictions. Implicit diffusion removes an explicit diffusion stability restriction, not accuracy requirements or every constraint of the coupled method.

- Splitting convergence requires regularity and treatment of boundaries; order reduction can occur.

Primary numerical source: Faou, Ostermann and Schratz, inhomogeneous parabolic splitting analysis [30]; reaction-diffusion invariant-region preserving discretization [31].

#### *15 Autonomy unlock*

- Bounded, Lipschitz, strictly positive Gamma is a valid gain assumption if enforced. The product K_base S_exp I_phys does not ensure it without explicit restrictions or saturation.
- If I_phys can be zero, a strictly positive Gamma_min is incompatible with that possibility.
- An operational activation gate is not a proof of philosophical autonomy or sentience.

#### *16 Envelopes and convergence*

- The Xi lower bound ignores negative forcing, boundary extraction and couplings. It requires a scalar comparison inequality for the complete right-hand side.
- The N and D upper bounds fail in the presence of positive sources. The appropriate forced estimate is N(t) <= exp(-a t)N(0)+integral_0^t exp(-a(t-s))f_plus(s) ds, when the required comparison hypotheses hold.
- The displayed kappa equality is only exact for a truly constant setpoint and no extra forcing, spatial terms, memory or couplings. Convergence of the setpoint is not enough to make that finite-time equality true.
- Even in the reduced ODE, kappa converges to M_kappa_set(E_infinity), not necessarily kappa_star. The equality of those targets must be imposed.
- An asserted ultimate-consciousness label is not a mathematical consequence. Use an operational target-state name.

#### *17 19 Theorem claims*

- Local Lipschitzness alone is not a global existence guarantee; establish a priori bounds/continuation criteria or global growth bounds. Core nonnegative variables remain unbounded above.
- Missing boundary conditions, nonautonomy, derivative-dependent feedback, discontinuous gates, degenerate diffusion, kernel-domain assumptions and unconstrained nonlinearities prevent the stated theorem from following.
- Positivity-preserving discretization proves a property of the numerical sequence under its hypotheses. A continuous constrained PDE needs its own viability theorem and convergence of the discretization.
- The scalar small-gain product is insufficient for all internal loops; bounded output is not ISS. The full system requires actual subsystem estimates.

### Explicit counterexamples and diagnostic tests

These examples are direct calculations, not claims borrowed from a cited paper.

1. **Bounded forcing prevents zero-noise convergence.** N_dot=-aN+c with a,c>0 has N(t)=c/a+(N(0)-c/a)exp(-at). A constant bounded layer readout can supply c with zero feedback gain. All gain bounds can hold while N tends to c/a, not 0.

2. **Integrable memory can destabilize.** x_dot=-x+m, m_dot=-m+2x corresponds to kernel 2exp(-t), whose L1 norm is finite. Eigenvalues are -1 ± sqrt(2); one is positive. Kernel integrability and stable isolated decay are insufficient.

3. **Finite-window flux is not bounded long-term.** H=K=tau=1 gives T(t)=exp(Lambda t). Every finite-window integral is finite, but T diverges. A scalar z_dot=-z+exp(Lambda t) grows asymptotically like exp(Lambda t)/(Lambda+1).

4. **Final projection loses boundary history.** On [0,1], start x(0)=0 and apply u=+1 until time 2, then u=-1. The constrained reflected solution reaches 1 at time 1, stays there to time 2, then equals 0.5 at time 2.5. Projecting the accumulated unconstrained integral yields Proj(2-0.5)=1 at time 2.5. Thus the section 13 formula is not a general constrained-solution formula.

5. **Metzler-Hurwitz minus positive-definite is not enough.** Take A=[[-1,100],[0,-1]] and L=[[1,-2],[-2,5]]. A is Metzler-Hurwitz; L is symmetric positive definite since its leading principal minors are 1 and 1. But A-L=[[-2,102],[2,-6]] has determinant -192, hence one positive eigenvalue. Diagonal nonnegative damping or a common weighted dissipativity certificate is safer.

6. **Joint L1 kernel does not imply L2 boundedness.** On (0,1)^2 take K(x,y)=x^(-3/4). K belongs to L1 of the product domain. For f(y)=1, (Kf)(x)=x^(-3/4), which is not in L2(0,1).

7. **Positivity is violated by additive removal.** N_dot=-aN-b with b>0 at N=0 gives N_dot=-b<0. This is exactly the structural risk of the flux-removal term when it does not vanish at N=0.

### A repairable theorem kernel

A practical first theorem should be finite-dimensional after a documented spatial discretization, deterministic, and free of implicit time-derivative feedback. Let y collect all core, manifold, graph, memory, derivative-filter and controller states. Let K be a nonempty closed convex constraint set and define

$$\dot{y} + N_K(y) \ni F(y) + B\,u(t).$$

Use a fixed, single-valued globally Lipschitz F on the relevant set and a well-defined projected/normal-cone solution. Verify an equilibrium $y_*$ with $F(y_*) \in N_K(y_*)$. If

$$\langle y - y_*, F(y) - F(y_*) \rangle \le -a\,\|y - y_*\|^2, \quad a > 0,$$

then normal-cone monotonicity and Young's inequality give, for $V = \|y - y_*\|^2$,

$$\dot{V} \le -aV + \|B\,u\|^2/a.$$

Consequently

$$V(t) \le \exp(-at)V(0) + \int_0^t \exp(-a(t-s))\|B\,u(s)\|^2/a\,ds.$$

This yields exponential convergence for zero input and a quantitative disturbance bound. It is a sufficient certificate, not a claim that the present equations satisfy it. Strong monotonicity is restrictive; alternate block Lyapunov or gain-matrix proofs may preserve more of the intended architecture. For the continuum PDE, replace the finite-dimensional conditions with appropriate coercive-form/maximal-regularity and boundary-admissibility hypotheses; do not merely reuse the finite-dimensional proof.

### Recommended research and implementation sequence

1. Publish a symbol dictionary and units; define every state, measurable output, matrix, parameter interval and graph.

2. Select the authoritative constrained equation and one interpretation of space, time, forcing and noise.

3. Replace hard gates and implicit derivatives with smooth bounded gains and explicit filters; cap amplification and leak accumulators.

4. Center every memory and feedback channel on a verified operational equilibrium, or explicitly solve for the forced equilibrium.

5. Prove a finite-dimensional deterministic well-posedness/invariance result before attempting the PDE/SPDE theorem.

6. Derive subsystem ISS bounds and a full coupling gain condition, including memory and adaptation.

7. Add continuum, boundary-input and stochastic variants as separate theorems with separate hypotheses.

8. Test the seven counterexamples as regression cases. Run mesh/time refinement, conservation/viability checks, adversarial input tests, gain sweeps and ablations.

9. Calibrate telemetry-to-affect labels against operational outcomes. Compare against a simpler controller with the same computational budget.

10. Separate mathematical stability, task performance, creative quality and consciousness into distinct evaluation claims.

### Scientific scope of sentience claims

The existence of an affect vector, memory, global coupling or a stable attractor is not a validated test for subjective experience. These mechanisms can be studied as computational functions without asserting that the system feels them. Current research treats potential indicators through competing consciousness theories and emphasizes uncertainty. The 2023 report is historical evidence, not a timeless verdict on every future model. Relevant primary sources are Butlin et al. (2023) [12] and the later Identifying indicators of consciousness in AI systems [32]. The later Praxymuse and CalcuArt caveats that coherence is operational and simulated affect is not sentience should also be applied to the base Praxium document.

### What should be retained

Retain the typed modular architecture, bounded observable maps, causal feedback, explicit constraints, exponential-memory implementation, parameter provenance, numerical checks and separation of operational state from subjective claims. Repairing the theorem layer does not require discarding the creative framework. It makes the framework testable and its guarantees meaningful.

# 4 Poetic Asymmetry mathematical and creative review

### Assessment

The specification is a promising architecture for constrained creative generation, but not yet a proved gapless calculus. Several individual mathematical ingredients are established. Their proposed combination, monotone artistic improvement, and guaranteed convergence to coherence 1.2 do not follow from the stated assumptions. The counterexamples below are direct analysis of the supplied specification, not claims attributed to external papers.

### Decisive counterexamples

#### *Penalized optimization need not maximize coherence*

Let the feasible set be [0,1], coherence Xi(x)=1.2x, originality penalty R(x)=x^2, lambda=1.2, and constraint penalty zero. Then L(x)=1.2x^2−1.2x is smooth and strictly convex, with unique minimizer x=0.5. Its coherence is 0.6, although 1.2 is attainable at x=1. Exact proximal-gradient updates with an appropriate fixed step converge to x=0.5. This defeats section 7's equivalence and section 15's conclusion even on a compact convex set with exact updates, no noise, no nonlocal field, and a persistent gate. The proof's assertion that A1 implies argmin L = Omega'_p is absent from A1.

Repair: optimize coherence with originality as a hard constraint, or establish that the penalty preserves the desired optimizer set. Otherwise define the target as argmin L rather than a predetermined coherence score.

#### *Small constant noise does not ensure monotone improvement*

Take Xi(x)=1.2−x^2 near x=0 and a noisy update x+=x+gamma*zeta at the optimum, with E[zeta]=0 and E[zeta^2]>0. Then E[Xi(x+)−Xi(0)]=−gamma^2E[zeta^2]<0. Any positive noise decreases expected coherence at that optimum. The stated bound c1||grad L|| ^2−c2*sigma^2 itself exposes this issue when the gradient vanishes.

Repair: claim improvement only outside a noise-dependent neighborhood, or impose vanishing/summable noise and appropriate step conditions. Persistent noise generally gives neighborhood or invariant-distribution results rather than exact almost-sure convergence.

#### *Tangent cone flow has the sign in the wrong place*

Section 4 writes minus projection of the gradient. At a constraint boundary, this differs from projection of the negative gradient because a cone is not a linear subspace. For F=[0,infinity), x=0, and grad L=−1, the written rule gives −P_[0,infinity)(−1)=0, incorrectly stuck; the projected descent direction is P_[0,infinity)(1)=1.

Repair: use projection of minus the Riemannian gradient onto the tangent cone, or a properly defined normal-cone differential inclusion. The distinction disappears for linear tangent spaces but matters for boundary cones. See Hauswirth et al., Optimization algorithms as robust feedback controllers [33].

#### *Prox regularity is not global firm nonexpansiveness*

The unit circle is smooth and prox-regular. Projection from nearby points inside the circle expands tangential distances, and projection at its center is multivalued. It cannot replace a closed convex set in the global firm-nonexpansiveness arguments.

Repair: choose a convex Euclidean theorem, or prove a separate local nonconvex/Riemannian theorem with tubular neighborhoods, curvature bounds, admissible steps and initialization. Poliquin, Rockafellar and Thibault [34] provide the relevant local geometry.

### Section by section audit

#### *0 Types and domains*

Specify whether the state is a Euclidean vector, embedded manifold, Riemannian manifold, token sequence, or product of these. The text moves among these settings without connecting maps. The 1.2 ceiling is a normalization, not a privileged mathematical full-coherence threshold. Each component observable needs a range ensuring the weighted sum remains in [0,1.2]; boundedness alone is insufficient.

#### *1 Master update*

A well-defined map M to M gives a unique discrete trajectory for a specified input/noise realization. Lipschitz continuity gives sensitivity bounds, not convergence. Reflection produces a revision plan, but its type is undefined: introduce a plan space P, R:M to P, and A:P×E to M. Time-varying E makes the system nonautonomous; distinguish fixed points for frozen E, all E, or an input distribution.

#### *2 Modules*

Phonetic substitution, syllabification, token decoding and meter decisions are discrete operations, not automatically differentiable or meaningfully Lipschitz. A continuous latent embedding and discontinuous decoder do not transfer latent convergence to poems. Define separate semantic, phonetic, syllabic and prosodic spaces. Nonexpansive projection onto a nonconvex manifold requires proof and is generally unavailable.

#### *3 Objective and feasibility*

Originality, ethics and poetic coherence are not generally convex constraints. Keeping away from known works creates exclusion regions. Convex originality is an assumption needing a concrete construction, not an intrinsic property. The coherence score does not explicitly measure semantic fidelity: rhythm, meter, tension and volatility can improve while meaning deteriorates.

#### *4 Continuous flow*

Besides the projection-sign error, additive noise can leave the manifold or feasible set. Specify Itô or Stratonovich dynamics, filtration, regular diffusion and a feasibility mechanism: tangent diffusion, retraction, reflection or constrained stochastic dynamics.

Brownian noise is not ordinary differentiable forcing. Indicator penalties require subgradient rather than ordinary gradient treatment.

#### *5 Existence and Lyapunov drift*

A drift inequality only when the gate exceeds delta gives no control on other steps and does not establish a supermartingale for the entire sequence. Require a global conditional inequality with summable errors, adaptedness and integrability. Objective-value convergence alone does not imply state convergence or stationarity of every limit point; boundedness and residual arguments are needed. See Robbins and Siegmund [35]. Avoiding saddles does not exclude nonglobal local minima; constant isotropic noise does not prove point convergence.

#### *6 Gate and gain adaptation*

Counting only signs of changes lets tiny gains outweigh rare catastrophic losses. Define initialization for negative history indices and separate noisy evaluation from real stagnation. Thresholds depend on score normalization. Halving/clipping gains does not prove improvement. Include changing gains and recent history in the state, or provide a time-varying analysis.

#### *7 Optimization*

The coherence/penalty equivalence fails by the explicit counterexample. KKT necessity generally requires constraint qualification. Proximal-gradient convergence needs an explicit smooth part with Lipschitz gradient, suitable convex nonsmooth part, minimizers and an admissible step-size bound tied to smoothness. Being bounded away from zero and infinity is insufficient. Parikh and Boyd [36] provide the relevant established framework.

#### *8 Fractal consistency*

Weighted-sum ascent is valid if its decomposition and nonnegative component drifts hold. A finite hierarchy is not mathematically a fractal. Define scale maps and self-similarity or call this hierarchical consistency. Locally improved lines can harm global narrative, so nonnegative global drift is substantive rather than automatic.

#### *9 Nonlocal field*

On a manifold, gradients at Phi and Psi lie in different tangent spaces and cannot be summed without transport or ambient identification. Positive kernels do not ensure alignment with local descent. Specify an inner-product domination inequality and boundary projection. Near-optimum bounds do not establish global attraction.

#### *10 Monotone improvement*

The inequality needs derivation; almost-nonexpansiveness and zero-mean noise do not imply it. At constrained optima, raw gradient norm can be nonzero: use a projected-gradient mapping or proximal residual. Improving L need not improve Xi.

#### *11 Attractor*

Local contraction is local: a map can contract near the target and be identity elsewhere. Global attraction needs global descent and entry to the local basin, or stronger hypotheses. A strict contraction cannot contain two distinct fixed points in its domain. Distinguish unique-point contraction from contraction toward a set.

#### *12 Completeness*

The tensors are schemas, not complete specifications. Dimensions, units, encoders, decoders, calibration, operator formulas and training procedures remain missing. Declaring differentiability neither implements nor verifies it.

#### *13 Discrete continuous connection*

Robbins-Monro steps also require stability/boundedness, regular drift, adapted noise and constraint compatibility. Appropriate time-shifted interpolations track an ODE or differential inclusion on finite windows; this is not unconditional convergence to section 4's noisy flow. See Benaïm [37].

#### *14 Meta learning*

Smoothness and stability alone do not guarantee parameter convergence. Define hypergradient existence, inner-solver error, bias control, projection geometry and step sizes. In nonconvex settings, stationary accumulation points or a residual bound are often defensible, while convergence of the full parameter sequence needs more.

#### *15 Main theorem*

The main theorem needs replacement. Its proof treats prox-regularity as convexity, assumes a target/minimizer identity not stated in its assumptions, omits sufficient step and stochastic requirements, relies on weak gate persistence, and uses local nonlocal-field control to claim a global result.

### Defensible replacement core

Use a finite-dimensional Euclidean state, nonempty closed convex feasible set, objective F=f+h with f convex and L-smooth and h proper closed convex, a nonempty minimizer set, and deterministic proximal-gradient updates with a conservative 0<gamma<=1/L. State exact summable-error requirements if using inexactness. Prove convergence to an optimizer of that surrogate objective. Conclude Xi reaches 1.2 only after independently establishing that every relevant optimizer attains it.

For stochastic execution, specify an almost-supermartingale inequality with diminishing step/noise assumptions. For discrete poem generation, use empirical guarantees: feasibility, measured improvement under defined evaluators, and uncertainty-aware human evaluation. A convex theorem about a controller does not cover neural decoding or universal artistic excellence.

### Creative and linguistic foundations

- Deep-speare [38] combines language, meter and rhyme. Expert judgments revealed remaining readability and emotional shortcomings, supporting separation of formal correctness from literary quality.
- Hafez [39] combines neural generation and finite-state constraints with user-controlled polishing and feedback, a concrete precedent for reflection/action design.
- Chen et al. 2024 [40] compare structural, lexical, semantic and stylistic distributions, finding underdiversity across several dimensions. This supports series-level evaluation beyond single-poem scores.
- Evaluation of Generated Poetry [41] reports semantic measures more predictive of overall impressions than formal measures in its Czech case study. This supports explicit meaning evaluation, not a universal metric.
- Reimann and Scheffler 2024 [42] distinguish dictionary-defined from crowd-rated metaphor novelty. A single originality scalar conceals different constructs.
- Cheung et al. 2019 [43] investigate harmonic uncertainty and surprise in musical pleasure. This supports expectation-sensitive design but not a theorem that greater asymmetry improves poetry.

### Research and implementation plan

1. Define typed text, pronunciation variants, stress, syllables, syntax, semantic roles, imagery, motifs, speaker, listener, narrative time and rhetorical intent.

2. Separate hard constraints from soft objectives; use finite-state constraints where suitable and explicit source comparisons for originality.

3. Define asymmetry against established expectation: enjambment, phrase-length variation, delayed rhyme, metrical substitution, semantic reversal and silence, rather than random mutation.

4. Use a multiobjective scorecard preserving conflicts among musicality, meaning, novelty, comprehensibility and emotion.

5. Use language- and tradition-specific phonetics and meter; do not generalize English stress-meter rules to Persian, Arabic, tonal or mora-based traditions.

6. Run blinded human evaluations, distinguish experts from general audiences, and report disagreement and confidence intervals.

7. Run ablations removing asymmetry, revision, metaphor reasoning or semantic constraints, alongside a plain-generation baseline.

8. Evaluate repetition and distributional diversity across works as well as within each work.

9. Keep a revision log connecting intended changes to measured effects.

10. Publish narrow controller theorems separately from empirical creative-quality claims.

Recommended framing: Poetic Asymmetry is a constrained, expectation-aware creative-generation framework with conditional optimization guarantees. This preserves its useful ambition while making its strongest claims testable.

# 5 Praxymuse music and production review

**Version note:** This chapter audits the pasted v0.5 specification. The v0.6 attachment corrects or qualifies several points, especially stiff-string notation, stability claims, control evidence and QA. Read Chapter 1 before treating any v0.5 criticism as applying to v0.6.

Research checked November 6, 2026. This review addresses every numbered section, but does not claim independent verification of every musician attribution or regional practice in the extensive lists. Mathematical corrections below are direct analysis; empirical claims are linked to sources.

### Overall assessment

Praxymuse is a thoughtful, broad architecture and research specification. It is not yet an executable system or a theorem proving musical excellence. Its strengths are versioned SongIR, instrument-specific gesture, cultural context, explicit rights tracking, instrumental dialogue, and human listening. The principal gaps are operational definitions, measured datasets, executable constraints, calibrated evaluators, and demonstrated end-to-end performance.

### 0 Contract and SongIR

An immutable, versioned intermediate representation linking score, lyrics, articulation, synthesis, stems, mix, provenance, and audio is the correct foundation. Distinguish symbolic identity, perceptual similarity, and bit-identical PCM as separate guarantees. Record hardware, driver, numerical precision, dependency versions, checkpoint hashes, random-generator states, DSP settings, and sample-library hashes. Store final lossless PCM as the authoritative sonic artifact. Seeds and weights alone do not guarantee cross-platform reproduction. PyTorch warns that reproduction is not guaranteed across releases, platforms, or CPU/GPU executions. Official reproducibility documentation [14].

Add sample-rate, channel-layout, tempo-unit, tuning, and timestamp schemas. Keep beat coordinates rational where possible and define sample rounding. SongIR locks an existing realization; it does not prove that the initial realization obeys every musical instruction.

### 1 Acoustic first principles

Most formulas are appropriate simplified models, with scope qualifications:

1. Stiff-string frequencies: $f_n \approx n f_1 \sqrt{(1+B n^2)}$ is inconsistent at n=1 if f_1 denotes the actual fundamental. Use $f_n \approx n f_{ideal} \sqrt{(1+B n^2)}$, or $f_n \approx n f_1 \sqrt{((1+B n^2)/(1+B))}$.

2. The modal oscillator equation is a linear, weakly damped model, not a complete description of bow stick-slip, reed collision, brass lip coupling, or strong nonlinearity.

3. Open/stopped-tube formulas assume idealized boundaries and geometry, not universal instrument tuning.

4. Specify STFT window normalization, hop, transform convention, and reconstruction conditions.

5. Spectral centroid needs a magnitude/power convention and silence handling. It is a descriptor, not a universal perceived-brightness measure.

Smith's physical-audio treatment supports the importance of dispersion, damping, and nonlinear excitation rather than ideal harmonic series alone. Stiff-string modeling [44]; digital waveguide models [45].

### 2 Instrument ontology and mastery

Source-exciter-resonator-gesture-body-room is an effective structure. Correlated phrase-level humanization is more plausible than independent random onset/velocity jitter.

Missing implementation includes an actual performer model for each supported instrument, with range, transitions, fingering, breath, recovery, and articulation constraints. Limits depend on player skill, instrument, dynamics, tempo, and technique. Gestures are coupled: pluck location changes timbre, bow pressure changes stability, and articulation can change timing. Separate intended gesture from measured features: a classifier detecting rasp does not establish authentic or safe vocal physiology.

Taxonomic correction: “keyed” describes an interface rather than a single acoustic family. Santur and yangqin are hammered zithers and ordinarily not keyboard instruments. Classify them as struck chordophones with an optional interface field.

### 3 Pitch harmony counterpoint

The tuning and cents formulas are correct for positive frequencies and clearly defined reference values. Encode continuous pitch trajectories and contextual intonation, not just scale tables. Make voice-leading penalties style-conditioned: parallel motion, crossings, or unresolved dissonance can be desirable. Distinguish note spelling, sounding pitch, tuning, and performance inflection.

MusicXML supports fractional semitone alterations, but downstream programs may not preserve them in MIDI playback. Microtone example [46]; interoperability warning [47].

Rāga, maqām, and patet are not merely pitch collections. Their phrase grammar, hierarchy, development, repertoire, and performance context require independent representation.

### 4 Time groove form

The beat integral is appropriate with positive BPM and a specified beat unit. State whether BPM counts quarters, dotted quarters, or another unit. Its inverse requires strictly increasing b(s); fermatas and pauses need explicit treatment. Distinguish structural rhythm, expressive timing, and acoustic onset. Represent groove hierarchically: shared pulse, performer deviations, instrument-dependent onset latency, and phrase timing. Do not assume one covariance matrix across styles, performers, tempos, and sections. Form requires operational tests for recurrence, transformation, transitions, and long-range retention, rather than a list of labels.

### 5 Affect and poetic asymmetry

This is a credible design hypothesis, not an established psychological law. The twelve-dimensional affect vector is a chosen representation, not an exhaustive scientific basis. A major study identified at least thirteen overlapping dimensions in sampled US/Chinese listening contexts, not a universal fixed taxonomy. Cowen et al., original study [48].

Distinguish intended emotion, perceived musical expression, and emotion actually felt. Distinguish lyric semantics from acoustic expression. Include listener background, familiarity, memory, context, preference, uncertainty, and disagreement.

Operationalize asymmetry with a measurable expectation model, deviations relative to established motif/groove/harmony/form, and listening tests of meaningful transformation versus novelty. Run ablations without the asymmetry controller and compare preference, coherence, recall, and repetition. Instrumental “empathy” is a useful compositional metaphor for timing and relationships, not evidence of experienced empathy.

### 6 8 Historical library global traditions recorded genres

These are useful research indexes, not a complete knowledge base. Retain nonlinear history, provenance, regional specificity, no universal quartertone grid, and no single “African rhythm.” Replace artist-name bundles with cited technique records, example analyses, cultural context, and confidence. Treat period boundaries as overlapping conventions.

Do not use one Hindustani-oriented rāga/tāla tuple as a complete Carnatic schema. Keep related West/Central Asian traditions distinct. Iraqi maqām names a particular repertoire and instrumentation, not every maqām-based practice. UNESCO Iraqi maqām [49]. Radif is a transmitted repertory central to Iranian classical music, not a scale database. UNESCO radif [50]. Gamelan requires ensemble, instrument, community, and performance context rather than one generic preset. UNESCO gamelan [51].

Give Indigenous traditions named community-specific records rather than a broad “OTHER” bucket. Add living practitioners and community review, particularly where notation and commercial recordings incompletely represent practices. Research supports both recurring cross-cultural patterns and substantial diversity, not one universal generative grammar. Mehr et al., original study [52].

### 9 Composition engine

The probability factorization is a reasonable conceptual decomposition. It becomes a mathematical model only after specifying variables, distributions, dependencies, and inference procedures. Add typed interfaces, revision dependencies, persistent motifs with transformation histories, section continuity checks, and silence/rest budgets. Changing lyrics must invalidate affected timing, vocals, and mix.

Declare whether generation is exactly score-controlled or approximately prompt-conditioned. Audio language models demonstrate practical generation but do not automatically expose exact symbolic control. MusicGen original paper [53].

### 10 Singing generator

The pipeline is well motivated. The hard problem is jointly controlling phonemes, breath, register, and emotion in plausible audio. A changing vocal tract is not generally one fixed convolution. Use a time-varying source-filter operator, local quasi-stationary approximation, or coupled model; the written convolution is schematic rather than fully physiological.

Verify word/phoneme omissions, repetitions, substitutions, invented syllables, pitch/timing adherence including melisma, vowel intelligibility, consonant placement, register changes,

breaths, phrase endings, and identity consistency. Full-track listening is necessary; ASR can miss errors or misrecognize correct singing. Technique control and evaluation remain research challenges. SinTechSVS original research [54].

### 11 Orchestration and interaction

The role graph is promising, but rewards need measurable definitions. Audibility differs from prominence; masking and doubling may be intentional. Silence depends on form and narrative. Separate compositional interaction from mix separation. Test whether listeners can identify lead/answer relationships at chosen moments without being told the intended interpretation.

### 12 Engineering and DSP

Nyquist concerns the bandwidth entering sampling. Saturation, clipping, waveshaping, and modulation generate new harmonics, so oversampling/filtering placement matters. A limiter cannot repair earlier aliasing, missing lyrics, incorrect articulation, or distortion. Record latency compensation, stem alignment, automation, routing, gain staging, and dither policy.

Measure true peak and integrated/short-term loudness by a defined standard; no single loudness target equals quality. ITU-R BS.1770 [55]. Hybrid physical/neural rendering has genuine precedent: DDSP joins interpretable signal-processing components and learned models, but does not itself supply composition, cultural grammar, or full orchestration. DDSP original paper [56].

### 13 Rights and reference registry

Make provenance an enforceable gate. Audit code, weights, training-data disclosure, samples/recordings, compositions/editions, performer/voice consent, input-reference permission, and output restrictions separately.

- MusicGen Large's official card lists CC BY-NC 4.0; it is not the default choice for an unrestricted commercial pipeline. Official model card [57].
- ACE-Step 1.5 lists MIT licensing. This does not guarantee copyright clearance for every output; its own disclaimer warns about originality and cultural risks. Official repository [58].
- YuE2 weights use CC BY-NC 4.0 with added September 16, 2026 permission allowing individuals, musicians, and creators to monetize outputs free of charge. Commercial company use of weights requires separate licensing. Code rights are separate. Do not claim all commercial outputs are prohibited. Exact model license [59].

CC0 does not guarantee that an uploader possessed all relevant rights or clear privacy/publicity rights. CC0 explanation [60]. US copyrightability, digital replicas, and training are separate legal questions; originality screening is not legal clearance. Copyright Office AI reports [61].

### 14 Optimization and feedback

The document appropriately avoids global nonconvex convergence claims. Missing elements include units/normalization, weight calibration, explicit feasible sets and hard constraints, demonstrated links between Praxium variables and improvements, and stability under noisy

evaluators with delayed feedback. Bounded gains do not prove useful control: stable systems can optimize bad proxies.

Use hard gates for missing lyrics, forbidden assets, clipping, corrupted files, and license failures. Use Pareto/preference comparisons for subjective conflicts, held-out listening tests against evaluator overfitting, and ablations of affect, motif memory, asymmetry, and dialogue. Separate technical repair from stylistic rewriting.

### 15 Verification and execution

The broad QA list needs definitions and acceptance thresholds:

- Unit tests: tuning, beat conversion, phoneme timing, hashes.
- Property tests: event ordering, finite samples, valid ranges, stem synchronization.
- Render tests: repeatability in a pinned runtime and archived PCM comparison.
- Musical tests: motif retention, form, register, articulation, lyric completeness.
- Listening tests: blinded/randomized comparisons, loudness matching, multiple listeners, confidence intervals.
- Cultural tests: practitioner review for named traditions.
- Rights tests: complete provenance and license evidence per asset.

FAD and embedding similarity depend on sample size, embeddings, and reference sets; they do not replace listening or originality review. Primary evaluation research [62]. MUSHRA is useful for controlled audio-quality assessment, but artistic preference and compositional merit need additional questions. ITU-R BS.1534 [63].

### Practical roadmap

1. Formalize a limited prototype: one vocal language, a small instrument set, explicit tuning, and narrow style coverage.

2. Implement SongIR and deterministic archive first.

3. Distinguish score-led synthesis (precise notes/stems/revisions) from neural audio generation (broader timbral realism, weaker exact control).

4. Combine selectively with explicit capability flags and provenance.

5. Benchmark against a simpler conventional pipeline; Praxium/Asymmetry must demonstrate added value.

6. Expand cultural/instrument coverage after practitioner and listening validation.

Recommended status: “A proposed modular framework for controlled music composition, performance synthesis, and evaluation.” Reserve theorem claims for individual mathematical statements with complete assumptions and proofs.

### Supplemental software license check prompted by the supplied resource material

The underlying source files should be pinned by version/commit before deployment. These findings are license-reading research, not legal advice.

- **Essentia:** its official page describes the library as AGPLv3 for noncommercial applications with proprietary licensing available. It separately lists pretrained models as CC BY-NC-SA 4.0 and warns that some dependency rights cannot be sublicensed by UPF.

The page's commercial wording should not be generalized into a claim that AGPL itself prohibits commercial use; clarify the exact grant and build with UPF for a commercial deployment. Official licensing page [11].

- **aubio:** upstream currently supplies GPLv3 COPYING. Some package summaries show different versions, so inspect the actual selected release. Upstream COPYING [9].
- **Matchering:** official repository identifies GPL-3.0. It is a reference-based matching/mastering tool, not a comprehensive composition engine. Official repository [8].
- **libebur128:** upstream COPYING contains the MIT grant and notice-retention condition. It is a suitable candidate for a permissively licensed loudness-measurement component, subject to dependency review. COPYING [10].

GPL software is not automatically forbidden for commercial activity, nor does merely processing audio through a GPL utility automatically put the audio under GPL. Software distribution/linking obligations and musical-content rights are separate. A user policy excluding GPL is a stricter project-selection rule, not a universal legal statement. Similarly, permissive software does not establish that its model weights, datasets, samples, or outputs are CC0.

# 6 CalcuArt rendering and worldbuilding review

### Overall assessment

CalcuArt is a promising specification for a scene-first, provenance-aware art system, not yet a proved mathematical theorem or an implemented simulator. Its strongest contribution is separating immutable world/scene data from rendering style. Important improvements over the earlier Praxium text include explicit fiction labels, conditional stability claims, acknowledgement of nonconvex geometry, and refusal to claim exact reconstruction from one raster image.

This review checks the mathematical structure directly and researches rendering, optics, cloth simulation, biology, reproducibility, and representative historical assertions. It verifies representative historical claims, not every artist attribution or regional costume history.

### 1 Immutable scene and slider architecture

**Sound and implementable:** fix geometry, topology, identity, pose, materials, camera, lights, and world time in a canonical asset bundle; apply ink, paint, and physical shading to that bundle; treat changed angle, time, action, or composition as a new version; hash assets and record renderer configuration.

This resembles production scene-description architecture. OpenUSD supports composed scenegraphs, layers, asset references, metadata, and time samples. It does not automatically guarantee immutability: stronger layers can override weaker ones. A CalcuArt implementation must resolve and lock dependencies, prevent unauthorized overrides, and validate the composed scene. OpenUSD concepts [64]

Four claims must remain distinct:

1. Exact scene-state identity is achievable by using the same immutable data.

2. Repeatable pixels need a pinned execution environment and deterministic algorithms.

3. Perceptual resemblance is weaker than either of those.

4. Scientific correctness is not established by a hash.

ControlNet supplies conditioning on depth, edges, normals, segmentation and poses, but is not a theorem of exact geometric preservation. ControlNet++ explicitly investigates remaining alignment errors. A generative image edit should be labelled an approximate scene-preserving interpretation unless independently verified against authoritative assets. ControlNet [65], ControlNet++ [66]

**Repair:** define three operational modes: exact rendering from locked 3D assets without geometry-generating postprocessing; constrained approximate generation with depth/normal/ID constraints and measured deviations; and concept-only prompt generation without an exact-lock claim.

### 2 Continuity reproducibility and the ten anchor path

The hat weights are correct: for q in [0,1], they are nonnegative, sum to one, reproduce q, and activate at most two neighboring anchors. All ten listed medium-mixture rows sum to

one. Their interpolation is continuous, piecewise affine, and Lipschitz, but generally not differentiable at anchors.

Interpolating parameters does not prove that rendered images are Lipschitz. Sampling, visibility tests, thresholded strokes, denoising, stochastic branching, clipping, and quantization can introduce discontinuities. A 0-100 slider is an authored style coordinate, not a physical measurement of percent reality. Cross-fading aligned images can guarantee a continuous image path, but does not prove that intermediate images faithfully simulate physical mixtures of oil, ink, and photography.

A useful conditional result is possible: if fixed-anchor images $I_j$ share the authoritative scene, then $I(q)=\Sigma w_j(q)I_j$ is continuous and has Lipschitz bound at most $9 \max_j \|I_{(j+1)}-I_j\|$. This is an image-interpolation theorem, not a theorem of artistic realism.

Seeds are necessary but not sufficient for universal reproducibility. PyTorch explicitly warns that complete reproducibility is not guaranteed across releases or platforms. Mitsuba exposes seeded sampling and deterministic sampler behavior, a stronger practical basis for controlled rendering. PyTorch reproducibility [14], Mitsuba samplers [67]

**Repair:** archive outputs alongside assets, software versions, hardware/backend details, deterministic settings, color transforms, PRNG state and hashes. Distinguish bitwise reproduction from tolerance-based numerical reproduction.

### 3 Nikon endpoint and optics

The Nikon D3500 specifications give a 23.5×15.6 mm sensor and 6000×4000 large-image output. Nikon specifications [68]

The stated horizontal field of view at 200 mm is correct under the rectilinear approximation: $2 \operatorname{atan}(23.5/400) \approx 6.72°$. A 10 m-wide scene consequently requires roughly 85 m distance in the far-field approximation.

The circle-of-confusion formula is correct for the usual ideal thin-lens assumptions, using consistent units, f-number N, and object/focus distances measured from the appropriate principal plane:

$c = f^2|s-s_0|/[N\, s(s_0-f)]$.

However, Nikon D3500, 200 mm, f/5.6 does not specify actual lens design, distortion, chromatic aberrations, vignetting, aperture geometry, focus breathing, sensor response, demosaicing, noise, tone curve, exposure or white balance. Thin-lens blur is only an approximation to a compound optical system. PBRT realistic cameras [69]

Depth-dependent blur is not generally a single spatially invariant convolution; occlusion-aware lens integration is needed near foreground/background boundaries. Skin, translucent surfaces and atmosphere require more than a BRDF-only surface model: BSSRDF and volume transport may be needed. PBRT volume scattering [70]

Use the term photographic rendering, not a photograph captured by a physical Nikon.

**Consistency repair:** hold focal length, sensor dimensions and camera pose fixed across q. Do not reintroduce an earlier 50→200 mm focal-length path. Aperture/defocus and depiction may vary, but magnification changes or focus breathing must be explicitly excluded or compensated.

### 4 Mastercraft and aesthetic objectives

The named mappings are reasonable research anchors, not exhaustive definitions of the artists. Mucha's curvilinear Art Nouveau association is directly supported by museum scholarship. Met collection entry [71]

The operator sum Σa_h O_h needs a shared typed output space. Perspective, engraving, psychological interpretation and architectural program cannot simply be added without defining representations.

Multiplying Beauty×Precision×Expression×Coherence×Invention×Craft is a design choice. It assumes calibrated commensurable scales and lets a zero in any factor annihilate all others. No current evidence establishes one universal numerical beauty functional. NIMA predicts distributions of human ratings rather than objective artistic truth. NIMA [72]

**Repair:** separate hard technical constraints from soft artistic preferences. Use multiple critics and user-calibrated pairwise judgments. Preserve tradeoffs rather than asserting a universal scalar optimum.

### 5 Biology evolution and Etheros fiction

The text correctly separates mosaicism from chimerism: mosaic cell lineages originate from one zygote, while chimeric lineages originate from distinct zygotes. The hybrid/chimera/mosaic/cyborg distinction is valuable. NHGRI [73]

Universal Etheros reproductive compatibility is appropriately marked fictional. A symbolic adapter does not establish a real biological mechanism. Known hybrid incompatibilities can involve sterility, lethality, genomic interactions and ecological fitness. Hybrid incompatibility review [74]

Specification gaps:

- Dev(GRN, ...) remains a placeholder, not an executable developmental model.
- Viability thresholds lack defined assays, units, uncertainty and validation data.
- Universal compatibility changes the meaning of species; supply a fictional lineage/taxon definition.
- Gene-regulatory networks, morphology, organ function and behavior require different timescales and representations.
- Energy budgets are necessary constraints, not sufficient proof of organism viability.
- Massive organisms or living architecture need stress, buckling, circulation, thermal, metabolic, maintenance and developmental scaling analyses.

**Repair:** use layered fidelity: ecological bookkeeping and mechanical feasibility for ordinary scenes, detailed physiology only where relevant. Keep fictional mechanisms as explicit world axioms with internally testable consequences.

### 6 Couture history and cultural provenance

The causal garment grammar is a strength: clothing follows body shape, climate, materials, tools, work, trade and individual identity rather than default armor or stereotypes.

The straight-seamed kimono description is supported by the V&A. V&A kimono [75]

History bands should remain overlapping regional indexing tools, not a universal progression. Islamic geometric ornament is a large family with historical exchange and regional differentiation, not one symbolic code. Met geometric patterns [76]

**Repair:** each actionable historical feature should record an object/source reference, region, date range, social context, construction/material evidence, documented meaning versus inference, confidence and permissions where applicable. Stylistic resemblance should not automatically imply religious, ethnic, political or ceremonial meaning.

### 7 Cloth and pan species fit

The energy contains recognizable membrane, bending, gravity, seam and contact terms and is a credible starting point for cloth mechanics.

Static energy minimization is insufficient for every depicted garment. Moving cloth needs inertia, damping and trajectory history. Friction is dissipative and often path-dependent. Contact and self-contact are nonsmooth. Constitutive coefficients need calibration: woven, knitted, leather and laminated materials behave differently. Thermal/moisture models require boundary conditions and transport laws. Safe pressure and all-poses-in-ROM claims require organism-specific criteria.

Implicit integration is established in cloth simulation; newer work includes differentiable dry-friction contact. Baraff and Witkin [77], DiffCloth [78]

**Repair:** static fitting for resting clothes, dynamic simulation for movement, continuous collision detection, explicit seam topology, material calibration and pose-envelope testing. Pan-species garments are engineering extrapolations, not experimentally validated comfort or safety.

### 8 Organic composition and series diversity

The conditional feasible-layout distribution is a good antidote to repetitive compositions. Debiasing against historically expected feasible frequencies is more defensible than enforcing identical quotas regardless of context.

However, J is undefined when no primary subject exists. Left/center/right and facing bins omit camera elevation, subject type, group structure, scale, interior/exterior, emotional event and socioeconomic context. The reference probability measure on the constrained feasible set must be defined because constrained sets can have zero ambient volume. The debiasing expression is a heuristic without a proved diversity or mixing guarantee. Narratives can repeat even when left/right placement changes.

**Repair:** maintain multidimensional scene history: subject class, age mix, society, task, setting, weather, camera height/distance/direction, crowd structure and narrative tension. Include no-primary-subject and ensemble modes. Evaluate semantic as well as visual repetition.

### *9* Causal world snapshots

Separating world time from optimization time and realism is excellent design. Local causal graphs, persistent IDs, event history, trajectories and resource bookkeeping can make scenes feel inhabited.

Requiring complete histories and futures for every entity is neither computationally necessary nor a guarantee of realism. Argmin E_world remains conceptual until candidate representations, constraints, feasibility and optimization are specified.

**Repair:** use bounded causal context: an immediate prior event, present intention or process, relevant relationships/resources and plausible next event. Maintain off-screen state at needed fidelity. Authored plausible explanations are not a validated simulation of an entire society or planet.

### 10 Mathematical and numerical gaps

Compared with Praxium v1.0, the qualified claims are much better, but gaps remain:

- Feedback contains first and second time derivatives of coherence, which generic mild solutions need not possess.
- The stall indicator is discontinuous, conflicting with blanket Lipschitz-reaction assumptions.
- Redistribution weights are undefined when every need score vanishes; define zero flux or a fallback distribution.
- Bounded individual fluxes do not imply a bounded total when the source population is unbounded.
- Persistent forcing requires equilibrium balance; convergence to a chosen state is not automatic.
- Metzler-Hurwitz component matrices and positive damping do not prove full coupled nonlinear stability.
- State-dependent divergence-form diffusion requires coefficient regularity, boundary conditions and a specified operator domain.
- The proposed reaction/diffusion/Prony/projection/nonconvex-feasibility sequence is not automatically a convergent splitting of one well-defined evolution.
- Lie-Trotter first-order accuracy is conditional; unbounded operators, nonsmooth projections and inexact nonconvex steps need their own analysis. Splitting analysis [79]

**Repair:** prove a separate modest theorem for a finite-dimensional, fixed-scene, convex constrained subproblem. Place nonconvex world generation and aesthetic optimization outside that theorem. Introduce smooth or hysteretic stall detection and filtered derivative states.

### 11 Verification and implementation

PASS mixes exact data assertions with uncertain semantic judgments. Different evidence standards are required:

1. Scene integrity: exact hashes for assets, transforms, IDs, pose and camera.

2. Projection lock: identical geometric ID, depth and normal passes across q.

3. Rendering repeatability: same-environment pixel hash; cross-environment error tolerance.

4. Slider behavior: endpoints, all anchors, near-anchor samples and random q values.

5. Physical checks: penetration, support, force/energy residuals, garment strain and collision.

6. Causal checks: local graph consistency and resource budgets.

7. Historical checks: source-backed metadata and uncertainty labels.

8. Artistic checks: blinded/user-specific assessment, not a claimed theorem.

9. Archive recovery: restore on a clean environment and verify dependencies.

10. Failure handling: identify failed constraints rather than asserting PASS.

Finite sampling alone cannot establish a universal all-real-q claim. An implementation proof or analytic bound is needed, supplemented by testing.

Recommended development order: locked scene bundle and registry; reproducible physical/ink renderers sharing geometry; calibrated slider and invariance tests; garment/physics and causal annotations; multi-scene diversity and provenance registry; then richer biological/ecological simulation and generative assistance.

### Conclusion

Preserve the creative ambition and scene-lock contract. Present CalcuArt as a formal design specification with a small provable computational core, empirical artistic modules and explicitly fictional Etheros axioms. This is more defensible and buildable than treating the entire system as one proved theorem.

# Recommended integrated development sequence

## Phase one Define the common contract

Create a versioned claim ledger and a typed symbol dictionary before adding more equations. For each variable, record its domain, units, measurement procedure, update rate and interpretation. Separate physical time, optimization steps, narrative time and style-slider position. Define interfaces between the controller, creative representation, renderer or synthesizer, evaluators and archive.

Use an explicit capability vocabulary. A requested property is an intention. A prompted property is communicated to a generator. A conditioned property is represented in a model input. A structurally locked property is stored in authoritative data and preserved by the execution path. A verified property has passed a defined test. These are different states and should never be collapsed into a single PASS label.

## Phase two Build a small provable controller

Start with a finite-dimensional deterministic system, smooth bounded gains, explicit derivative filters and an equilibrium that actually balances all forcing. Include memory and adaptation in the state. Establish a constrained well-posedness result and a quantitative disturbance estimate. Use the counterexamples in this report as regression tests rather than relying on successful ordinary runs.

Add continuum PDE, boundary-control and stochastic variants only after their separate domains, regularity and admissibility conditions are specified. The theorem for one variant should not be silently transferred to another.

## Phase three Demonstrate one creative instance per domain

For music, use one language, a limited ensemble, explicit notes and phonemes, a small style set, and retained stems. For visual art, use one authoritative scene with a few identifiable objects, a fixed camera, stored lights and geometry, and physical plus ink renderers sharing the same scene. For poetry, use a constrained revision task with explicit semantic and formal requirements.

The objective is not to make the simplest possible art. It is to make failures attributable: a missing word should map to a lyric or vocal-rendering defect; a moved object should map to a broken scene contract; a damaged metaphor should map to a semantic revision failure.

## Phase four Test the distinctive contribution

Compare a competent conventional baseline with controller-only, asymmetry-only and combined variants under equal compute and selection budgets. Record constraint failures, editing effort, run cost, latency, listener or viewer preference, and disagreement. Include failed outputs and avoid selecting only favorable examples.

Test intentional asymmetry against random perturbation while matching overall complexity. Test scene identity on canonical data and geometric passes, not only embedding similarity. Test lyrics with listening and transcript checks; forced alignment alone cannot establish that

words were sung. Validate cultural claims with tradition-specific evidence and qualified practitioners.

## Phase five Expand with evidence

Broaden instruments, traditions, bodies, clothing, world simulation and artistic styles only when each addition has a documented representation, source provenance and acceptance test. Rich ontologies are useful planning tools, but listing a capability does not implement it.

Publish a versioned evidence package with the specification, implementation, benchmark inputs, settings, outputs, failure cases and analysis where rights permit. Keep the small proved core distinct from empirical modules and Etheros fictional axioms. That separation gives the framework a credible path toward stronger claims without reducing its creative ambition.

## Final assessment

The research supports continuing this project as a modular, testable creative architecture. It does not support the existing broad claims that the entire fused system is mathematically proved, universally convergent, sentient, culturally exhaustive, or capable of guaranteed artistic excellence. The newer Praxymuse integration already moves toward the more defensible approach: preserve source material, state conditional mathematics precisely, record execution evidence, and disclose capability limits.

# Source references

The numbered references below are linked where they support the discussion. Counterexamples are direct calculations from the supplied specifications. A source supports the specific subject discussed; it does not validate the complete fused framework. Online licensing statements are version specific and should be rechecked before deployment.


1. github.com. https://github.com/facebookincubator/Meta-open-sound

2. producerspace.com. https://producerspace.com/license

3. github.com. https://github.com/Mcamento8/open-game-sfx-index

4. github.com. https://github.com/SoundSafari/CC0-1.0-Music

5. github.com. https://github.com/m-malandro/CC0-midis

6. creativecommons.org. https://creativecommons.org/publicdomain/zero/1.0/deed.en

7. wiki.creativecommons.org. https://wiki.creativecommons.org/wiki/CC0_FAQ

8. github.com. https://github.com/sergree/matchering

9. github.com. https://github.com/aubio/aubio/blob/master/COPYING

10. github.com. https://github.com/jiixyj/libebur128/blob/master/COPYING

11. essentia.upf.edu. https://essentia.upf.edu/licensing_information.html

12. Consciousness in Artificial Intelligence. https://arxiv.org/abs/2308.08708

13. Barrett et al., 2019. https://pmc.ncbi.nlm.nih.gov/articles/PMC6640856/

14. Official reproducibility documentation. https://github.com/pytorch/pytorch/blob/main/docs/source/notes/randomness.md

15. W3C PROV overview. https://www.w3.org/TR/prov-overview/

16. NASA-STD-7009B. https://standards.nasa.gov/standard/NASA/NASA-STD-7009

17. NIST AI RMF measurement guidance. https://airc.nist.gov/airmf-resources/playbook/measure/

18. Runco and Jaeger, 2012. https://doi.org/10.1080/10400419.2012.650092

19. ITU-R BS.1534. https://www.itu.int/rec/r-rec-bs.1534/en

20. FAIR Guiding Principles. https://www.nature.com/articles/sdata201618

21. Mironchenko, Karafyllis and Krstic, boundary-disturbance ISS. https://arxiv.org/abs/1706.07224

22. Schwenninger, parabolic boundary control. https://ris.utwente.nl/ws/files/250830395/Schwenninger2020_Chapter_Input_to_stateStabilityForPara.pdf

23. Dashkovskiy, Ruffer and Wirth, general-network ISS small-gain theorem. https://arxiv.org/abs/math/0506434

24. LeCrone and Simonett, quasilinear parabolic well-posedness. https://arxiv.org/abs/1808.01032

25. Pruss, Simonett and Wilke, critical spaces. https://arxiv.org/abs/1708.08550

26. Hurtado and Kirosingh, exponential-chain state lifting. https://pmc.ncbi.nlm.nih.gov/articles/PMC6800873/

27. Schur-test research. https://www.sciencedirect.com/science/article/pii/S0022123698933195

28. Brogliato et al., equivalence of projected systems and differential inclusions. https://doi.org/10.1016/j.sysconle.2005.04.015

29. Brogliato and Tanwani review. https://epubs.siam.org/doi/10.1137/18M1234795

30. Faou, Ostermann and Schratz, inhomogeneous parabolic splitting analysis. https://academic.oup.com/imajna/article-abstract/35/1/161/772236

31. reaction-diffusion invariant-region preserving discretization. https://academic.oup.com/imajna/article/39/1/235/4568335

32. Identifying indicators of consciousness in AI systems. https://pubmed.ncbi.nlm.nih.gov/41219038/

33. Hauswirth et al., Optimization algorithms as robust feedback controllers. https://doi.org/10.1016/j.arcontrol.2024.100941

34. Poliquin, Rockafellar and Thibault. https://sites.math.washington.edu/~rtr/papers/rtr170-LocalDiffDistance.pdf

35. Robbins and Siegmund. https://www.sciencedirect.com/science/chapter/edited-volume/pii/B9780126045505500158

36. Parikh and Boyd. https://web.stanford.edu/~boyd/papers/prox_algs.html

37. Benaïm. https://www.numdam.org/item/SPS_1999__33__1_0/

38. Deep-speare. https://aclanthology.org/P18-1181/

39. Hafez. https://aclanthology.org/P17-4008/

40. Chen et al. 2024. https://aclanthology.org/2024.emnlp-main.1097/

41. Evaluation of Generated Poetry. https://aclanthology.org/2025.eval4nlp-1.9.pdf

42. Reimann and Scheffler 2024. https://aclanthology.org/2024.law-1.9/

43. Cheung et al. 2019. https://doi.org/10.1016/j.cub.2019.09.067

44. Stiff-string modeling. https://www.dsprelated.com/freebooks/pasp/Stiff_String.html

45. digital waveguide models. https://www.dsprelated.com/freebooks/pasp/Digital_Waveguide_Models.html

46. Microtone example. https://www.w3.org/2021/06/musicxml40/musicxml-reference/examples/alter-element-microtones/

47. interoperability warning. https://www.w3.org/2021/06/musicxml40/tutorial/midi-compatible-part/

48. Cowen et al., original study. https://pmc.ncbi.nlm.nih.gov/articles/PMC6995018/

49. UNESCO Iraqi maqām. https://ich.unesco.org/en/RL/iraqi-maqam-00076?RL=00076

50. UNESCO radif. https://ich.unesco.org/en/RL/radif-of-iranian-music-00279?RL=00279

51. UNESCO gamelan. https://ich.unesco.org/en/RL/gamelan-01607

52. Mehr et al., original study. https://pmc.ncbi.nlm.nih.gov/articles/PMC7001657/

53. MusicGen original paper. https://arxiv.org/abs/2306.05284

54. SinTechSVS original research. https://smcnus.comp.nus.edu.sg/archive/pdf/2024/2024_SinTechSVS_IEEETrans.pdf

55. ITU-R BS.1770. https://www.itu.int/rec/R-REC-BS.1770/en

56. DDSP original paper. https://arxiv.org/abs/2001.04643

57. Official model card. https://huggingface.co/facebook/musicgen-large/blob/main/README.md

58. Official repository. https://github.com/ace-step/ACE-Step-1.5/blob/main/README.md

59. Exact model license. https://github.com/multimodal-art-projection/YuE/blob/main/MODEL_LICENSE

60. CC0 explanation. https://creativecommons.org/publicdomain/zero/1.0/deed.en?atype=html

61. Copyright Office AI reports. https://www.copyright.gov/ai/

62. Primary evaluation research. https://arxiv.org/abs/2311.01616

63. ITU-R BS.1534. https://www.itu.int/rec/R-REC-BS.1534/en

64. OpenUSD concepts. https://openusd.org/release/glossary.html

65. ControlNet. https://arxiv.org/abs/2302.05543

66. ControlNet++. https://arxiv.org/abs/2404.07987

67. Mitsuba samplers. https://mitsuba.readthedocs.io/en/stable/src/generated/plugins_samplers.html

68. Nikon specifications. https://onlinemanual.nikonimglib.com/d3500/en/16_technical_notes_07.html

69. PBRT realistic cameras. https://pbr-book.org/3ed-2018/Camera_Models/Realistic_Cameras

70. PBRT volume scattering. https://www.pbr-book.org/4ed/Volume_Scattering

71. Met collection entry. https://www.metmuseum.org/art/collection/search/283253

72. NIMA. https://research.google/pubs/nima-neural-image-assessment/

73. NHGRI. https://www.genome.gov/genetics-glossary/Mosaicism

74. Hybrid incompatibility review. https://pubmed.ncbi.nlm.nih.gov/21910629/

75. V&A kimono. https://www.vam.ac.uk/articles/kimono

76. Met geometric patterns. https://www.metmuseum.org/es/essays/geometric-patterns-in-islamic-art

77. Baraff and Witkin. https://doi.org/10.1145/280814.280821

78. DiffCloth. https://arxiv.org/abs/2106.05306

79. Splitting analysis. https://www.sciencedirect.com/science/article/pii/S0898122124003985