CalcuArt v2.3 formal calculus

Calculus-only edition prepared 7 November 2026. This document contains the formal kernel from the prepared CalcuArt Theorem v2.3, with the inherited numerical realism anchors supplied in Appendix A. It preserves the mathematical definitions, conditional proofs, modeling contracts and explicit limits. Artwork reproductions, the 50-work catalogue, study narrative, source assets and the terminal illustration request are excluded.

The three theorems are conditional results for the stated mathematical problems. They do not establish universal artistic quality, a complete implementation, full-system stability, exact scene reconstruction or biological validity. Proposed and fictional modules retain their qualifications. This is a reformatted calculus extract of v2.3, not a new scientific validation or a new version.

0. Claims, clocks, domains, and normalized coordinates

Use two independent labels. Ontological status: E = empirical-world proposition; P = plausible extrapolation; X = declared Etheros fiction. Evidence status: documented, observed in the present study, conditional/proved, proposed, or unknown. E does not mean experimentally verified; X does not waive internal consistency. Every evidence record contains claim, source/version, observation method, scope, uncertainty, and permitted use.

The realism coordinate is q = r/100 in [0,1], an authored depiction coordinate, not a measured percentage of reality. The user-specified S_0, scene content, and contextual priorities remain independent of q. Separate:

• τ in seconds: world time; τ* is the committed snapshot time.

• σ in seconds: playback/display time for temporal or interactive media.

• s = t_solver/t_ref: dimensionless solver time; k in the nonnegative integers indexes accepted/proposed solver updates; Δs > 0 is a declared numerical step.

• p in D = $[0,1]^2$: normalized image/support coordinates. Physical support coordinates are a = (W_s p_1,H_s p_2) in meters. Raster indices are integers and are not interchangeable with p.

• x_world in X ⊂ R^3: world position in meters. ξ = log(ℓ/ℓ_ref) for 0 < ℓ_min ≤ ℓ ≤ ℓ_max, with ℓ and ℓ_ref in meters. A logarithm never receives a dimensional quantity.

• i in a finite registry or graph V: counting measure. At a fixed time, dν = dp dx_world dξ times counting measure on V, with only the relevant coordinate factors used by each field. A state norm does not integrate over solver time. Spatial norms use declared physical reference scales or quadrature weights, so differently dimensioned domains are not silently combined.

Normalize every optimization feature as y_j = (f_j − f_j,ref)/s_j, where s_j > 0 has the same units as f_j. Then y, objective weights, and the finite-dimensional theorem below are dimensionless. Record the scales; changing them changes the optimization problem. Missing measurements are unavailable, not zero. Confidence is not precision unless calibrated as such.

1. Typed ArtIR and immutable versus editable state

Replace the ambiguous interpretation of ART as a numerical product by the typed record

ArtIR = (SourceArchive, Context, WorldIR, SceneCommit, DepictionIR,
MaterialProcessIR, PatternIR, TemporalIR, EvidenceIR, EvaluationIR).

The original ⊗ denotes constrained composition of compatible records. It is neither scalar multiplication nor a tensor-product theorem. A typed operator O_h : D_h → T_h has a declared input domain, output type, units, preconditions, source, and limitations. Combine operators only after adapters place their outputs in the same vector space or compatible record fields. A sum of anatomy,

psychology, ornament, and perspective has no meaning merely because each has been given a symbol.

SceneCommit S*. Keep the v2.2 canonical scene: persistent IDs; geometry/topology and UVs; rigs, poses, expressions and action; object-material identity and physical parameters; fitted garments; causal/context records; relationships; camera intrinsics/extrinsics; lights, atmosphere/weather and exposure interval; world time or trajectory; random-coordinate fields; dependency and renderer manifests. Declare canonical serialization and schema version. Commit before the first render at any q. New identity, geometry, pose, world time, camera, causal event, or physical material state requires an explicit new scene/version.

DepictionIR θ. A finite, typed control vector for value/color transforms, edge/stroke response, depiction-media blend, texture presentation, stylized shading and permitted depth-of-field response. The artist's chosen design S_0 is a fixed prior/contract; θ is its realization for the requested q. Physical object material does not change because a painted depiction changes its brush model. Camera pose, sensor size, focal length and geometric visibility ordering remain locked. If focus breathing or another optical change would alter projection, exclude or compensate it explicitly.

Representation modes:

1. Exact asset mode: render the committed authoritative assets; no geometry-generating postprocess; separately certify locked scene data and rendering repeatability.

2. Constrained inference mode: estimated scene/conditioned generation with measured deviations and stated missing information. It cannot inherit exact-lock certification merely from a seed, depth guide, or reference image.

3. Concept mode: a prompt or design proposal with no asset-level preservation guarantee.

The original Ψ and Z remain legacy interfaces. Replace unexplained single-letter Ψ fields with named typed fields, and publish any legacy adapter explicitly rather than guessing their meanings. Preserve $Z_{core} = (\Xi,\kappa,N,\varnothing)$, where Ξ in [0,1] is a declared coherence/control score and κ,N,∅ are declared nonnegative diagnostic coordinates, not measurements of consciousness. Retain e in R^{57}, c in R^{16}, s_{latent} in R^{16}, τ_{latent} in R^{13} and g in R^{15} if backward compatibility is wanted. Their 117 coordinates need a feature dictionary; the dimensions themselves confer no scientific status. Use τ_{latent} and s_{latent} to avoid collisions with world and solver time.

The 16 v2.2 layers remain modular interfaces: foundation, agent, knowledge, cybernetic, GUI, harmonic, quantum, information, temporal, topological, emergent, network, learning, crypto, stochastic and metamath. Each must declare input/output types, state dimension, update law, error bounds and measured role. Names such as quantum, affect, autonomy or emergent do not establish corresponding physical or psychological phenomena.

2. Distillation: retain the source; expose what features omit

A source record A_j contains the original retrieved bitstream b_j, media/container type, dimensions or duration, color/profile metadata where available, acquisition/source URL and time, source/object ID, rights/provenance record, transformations performed, and extraction software/version. Preserve the earliest available bitstream; calling it original means original to this archive, not necessarily the camera negative, artist's file, or physical object. A thumbnail, scan or photograph is an observation of an artwork, not the artwork's entire material history.

A derived record $\varphi(b_j)$ may contain descriptors, annotations, symbolic motifs, palette, spatial hierarchy, rhythm, material/process hypotheses, and uncertainty. Retain their links to the source and exact regions/times. A useful concise descriptor need not be a reversible representation. Rendering an attractive reconstruction is not a proof of reversibility.

Theorem 1. Exact recovery and task sufficiency are different properties

Let S be any set of admissible source bitstreams and E : S → Z an encoder.

(a) There exists a decoder D : E(S) → S with D(E(b)) = b for every b in S if and only if E is injective.

(b) Let T : S → Y be a declared task/query. There exists g : E(S) → Y with T = g ∘ E if and only if $E(b_1)=E(b_2)$ implies $T(b_1)=T(b_2)$. Thus an encoding can be sufficient for one task while losing information needed for another.

(c) The archive encoder E_archive(b) = (b, φ(b), metadata) admits exact source recovery by projection onto b, or by a verified lossless codec's inverse. No correctness claim about φ or the metadata is needed for this byte-recovery result.

Proof. For (a), if $E(b_1)=E(b_2)$, application of a left inverse D gives $b_1=b_2$. Conversely, injectivity makes D(E(b))=b well-defined on E(S). For (b), the forward implication follows by applying g. For the reverse implication, define g(z)=T(b) for any b with E(b)=z; the stated condition makes this independent of the representative. Part (c) follows by the coordinate projection or the stipulated codec inverse. ∎

Consequences. A descriptor with m bits cannot exactly encode every n-bit source when m<n: there are at most 2^m codes for 2^n sources. Variable-length lossless compression can shorten some sources but cannot strictly shorten every source of a given finite length. A fixed-length hash is not a lossless archive; it has collisions on a larger source domain. Cryptographic hashes support practical integrity checks under stated assumptions, not a universal proof of identity or artistic correctness. No finite feature list is automatically sufficient for every future artistic question. Information-theoretic sufficiency here concerns a specified deterministic task, not an unsupported claim that all historical or artistic meaning has been captured.

v2.3 archival contract. Keep source bytes plus derived descriptors plus uncertainty, not descriptors instead of sources. For a scene, keep dependency-closed assets and canonical state. For exact output recovery, keep the final output bytes too. Atomic write → read back → dependency validation → digest comparison → explicit commit is an engineering procedure, not a guarantee of indefinite storage. Re-rendering requires a separately documented deterministic execution environment. Clean-environment restoration must be tested before claiming it worked.

3. Expanded material, process, attention, pattern, and digital variables

The following are proposed fields and models to be calibrated, not observations made in the present corpus study.

3.1 Material and making process

For each object or physical artwork/support, record substrate geometry, thickness h in meters, density ρ in kg/m³, roughness parameters with their chosen model, wavelength λ in meters or explicitly nanometers, spectral reflectance R(λ) in [0,1] for the stipulated passive model, transmission, scattering, and emission where applicable. BRDF f_r has units sr^{-1} and, for a passive surface, obeys nonnegativity and hemispherical energy conservation. Translucency and participating media require BSSRDF/volume models rather than an unjustified BRDF-only substitution.

For a depicted making process, define M_proc = (support, tools, ordered operations, deposition, transport, drying/curing, abrasion, stochastic field, calibration). Example process state η(a,t_p) = (h_p,c_p,d_p,v_p), where h_p ≥ 0 is film thickness [m], c_p ≥ 0 is pigment concentration [kg/m³], d_p in [0,1] is a declared dry/cure fraction, and v_p is flow velocity [m/s]; t_p is process time [s]. Viscosity has units Pa·s and surface tension N/m. Parameters may be inferred only with uncertainty unless measured. A photograph seldom identifies layer sequence, binder, or pigment uniquely.

A generic process evolution ∂η/∂t_p = F_material(η,tool,environment) is a model interface, not a solved PDE or evidence of authentic technique. Conservation, constitutive laws, initial/boundary data,

parameter identification and numerical error must be specified for each executable model. Digital brush radius, opacity and texture can imitate visible effects without simulating real fluid chemistry; label the operational mode.

For cloth, retain v2.2 seams, pattern topology, contact, motion and fit constraints, but make dimensional conventions explicit. With dimensionless membrane strain ε, a membrane stiffness C_s in N/m, bending stiffness B in N·m, curvature II in m^{-1}, mass density ρ in kg/m^3, thickness h in meters and gravitational potential Φ in m^2/s^2:

$$E_{cloth} = \int_\Omega [\tfrac{1}{2}\, \varepsilon^T C_s\, \varepsilon + \tfrac{1}{2}\, B\, ||II - II_0||^2 + \rho\, h\, \Phi(y)]\, dA + E_{seam} + E_{contact} + E_{actuation}$$

has units joules. If C is instead a three-dimensional modulus [Pa], the membrane term must include thickness appropriately. Static minimization addresses equilibrium, not moving cloth. Dynamic cloth requires mass, damping, contact/friction history and an integrator. Use unilateral/signed contact constraints with a declared surface orientation; unsigned distance alone does not distinguish the inside of a closed body.

In the heat equation $\rho\, c_p\, \partial T/\partial \tau = \nabla\cdot(k_{th}\nabla T)+Q$, T is in kelvin, c_p in J/(kg·K), k_{th} in W/(m·K), and volume sources Q in W/m^3. Convective, radiative and evaporative exchanges are boundary fluxes [W/m^2] unless converted by a specified volume model. Organism-specific comfort/safety thresholds remain unknown until justified; an optimization label cannot certify them.

3.2 Mark, pattern and multiscale organization

For each mark, record tool/type, centerline γ(l), physical or image-space width, pressure proxy, pigment/opacity response, orientation, curvature, spatial scale, support layer, stroke order and source. A width is in meters on the support or pixels in a declared raster, never a unitless mix of both. Orientation modulo π is represented by (cos 2θ,sin 2θ), avoiding an artificial jump between equivalent line orientations.

PatternIR = (motif IDs, support charts, adjacency/topology, repetition group, periods, phase fields, orientation fields, scale hierarchy, variation budget, material/process linkage, cultural provenance, semantic uncertainty).

A physical repetition period has units meters; spatial frequency has cycles/meter; phase has radians. For a scalar illustrative channel,

$$P(a) = \Sigma_j\, A_j(a) \cos(2\pi\, k_j\cdot a + \varphi_j),$$

where k_j has cycles/meter, A_j has the chosen output-channel units, and φ_j is dimensionless radians. This is a local texture model, not a universal formula for ornament or meaning. Arbitrary motif topology, tessellation rules, embroidery paths and symbol syntax require discrete/symbolic records. A geometric symmetry group does not identify a culture or authorize a symbol's use.

Separate surface ornament physically present on a locked garment/object from representational strokes used to depict it. The former remains part of S* across q. Macro/meso/micro/nano are explicit scale ranges, not instructions to maximize detail everywhere. “Intricacy without clutter” is a contextual design aim evaluated by task, hierarchy and reader/viewer response, not an identity between two numerical variables.

3.3 Attention and perceptual organization

For a declared viewer/task population, $a(p,\sigma) \geq 0$ with $\int_D a(p,\sigma)\, dp = 1$ is an attention-density hypothesis on normalized image coordinates. It is dimensionless under normalized dp; physical support density would have units m^{-2}. Let a* be a user-authored desired attention distribution, and let uncertainty distinguish intended, model-predicted and measured attention.

One finite-grid proxy is $a_i = \exp(h_i/T_a)/\Sigma_j \exp(h_j/T_a)$, $T_a>0$, where h_i are dimensionless logits. Define a proposed mismatch $E_{att} = \Sigma_i\, (a_i - a^*_i)^2$, or a transport discrepancy only after specifying

grid geometry and units. Logits may depend on contrast, scale, edges, face/body cues, narrative task and context. The formula does not establish that viewers actually look there. Gaze traces, task performance or blinded comparisons would be empirical evidence; a saliency model alone is not.

Extend composition features to subject/ensemble/no-primary modes; direction, height and distance of camera; task/action; scale and spacing; positive and negative shape; edge hierarchy; value and color grouping; rhythms and interruptions; environmental/contextual variation. Define the reference measure for feasible-layout sampling, including discrete choices and coordinates on constraint manifolds. The v2.2 history-based debiasing is retained as a heuristic, with explicit no-primary and ensemble bins. It does not prove independence, uniform diversity or absence of semantic repetition.

3.4 Temporal, procedural and interactive digital media

TemporalIR = (mode, duration, world trajectory or snapshot, time mapping, frame/sample rate,

event log, transition policy, executable/dependency manifest,

external-input capture, replay contract, accessibility/fallback).

For a snapshot, $W^*(\tau^*)$ is fixed and σ controls only declared presentation effects. For animation, commit a trajectory $W^* : [\tau_0,\tau_1] \to \text{WorldIR}$ and a time map $\tau=\psi(\sigma)$; both are independent of q. Then $I(q,\sigma)=\text{Render}(W^*(\psi(\sigma)),\theta(q,\sigma))$. Animation permits world motion across σ while forbidding a hidden change of world motion across q. A scene version can represent a locked trajectory rather than just a single frame, but it must declare that different contract.

For interactive media, an event $e_j = (\sigma_j,\text{type},\text{payload},\text{origin})$ is a typed input; transition $W_{j+1}=T(W_j,e_j)$ is deterministic only under the specified runtime and captured dependencies. Live feeds, clocks, network responses and model APIs break exact replay unless their relevant inputs/outputs are captured. A generative system's code, model/weights, settings, seed, random state, color pipeline, assets and environment are dependencies, not interchangeable labels. Time-based work also needs timing, transitions, installation/display context and interaction rules; one representative still is insufficient for its temporal behavior.

A parameter q controlling medium does not authorize a different action, hidden event or story. An interactive choice that changes content creates a new state/epoch; retain the event log and parent version rather than pretend the snapshot is unchanged.

4. Fixed-scene realism path and its actual guarantee

Retain the ten v2.2 anchors $q_j=j/9$, $j=0,\ldots,9$, and hat weights $w_j(q)=[1-|9q-j|]_+$. They are nonnegative, sum to one, and reproduce q. Retain the original R_j labels and seven-component B_j weights, reproduced in Appendix A. The inherited line-control vector is specified there, but its numerical anchor values were not supplied. Each B_j lies on the medium simplex; $B(q)=\Sigma_j w_j(q)B_j$ remains on it. Interpolation is piecewise affine and normally not differentiable at anchors.

Theorem 2. Continuity of aligned fixed-scene image interpolation

Let H be a normed real vector space of finite-resolution, scene-linear image channels before output quantization, and let $I_0,\ldots,I_9$ in H be fixed anchor outputs. Define $I(q)=\Sigma_j w_j(q)I_j$. Then $I(q_j)=I_j$ and

$$\|I(q)-I(q')\| \le L_I\,|q-q'|,$$

$$L_I = 9 \max_{j=0,\ldots,8} \|I_{j+1}-I_j\|.$$

If all anchors are rendered from the same immutable S^*, use of this interpolation does not change the authoritative scene record. The theorem does not infer that the anchors share geometry from their pixels; that premise must be checked independently.

Proof. On $[q_j,q_{j+1}]$, only w_j and w_{j+1} are nonzero, and $I(q)=(1-a)I_j+aI_{j+1}$ with $a=9q-j$. Its slope has norm $9\|I_{j+1}-I_j\|\le L_I$. For points in different intervals, partition $[q,q']$ at the finitely

many anchors, apply the interval estimate and the triangle inequality, and sum the subinterval lengths. Adjacent expressions agree at every anchor. ∎

Scope. This proves a continuous crossfade, not physical mixtures of ink, oil and photographic media, monotone perceived realism, conservation of apparent silhouettes, artistic success, or a universal continuous renderer. If θ(q) is L_θ-Lipschitz and a fixed-scene renderer F is L_F-Lipschitz over that parameter path, composition gives L_I≤L_F L_θ; proving the renderer premise is a separate task. Thresholded strokes, stochastic branching, denoising, quantization and changed sample visibility may violate it. With a quantizer satisfying ||Q(v)−v||≤ε_Q, one obtains only

$$||Q(I(q))-Q(I(q'))|| \le L_I\,|q-q'| + 2\varepsilon_Q.$$

Quantized pixels are generally discontinuous. For a locked animation, a uniform bound over σ yields the same q-continuity estimate uniformly over playback; this adds no theorem about σ-continuity at cuts or interactive jumps.

Geometric versus apparent lock. Check exact geometry/ID/depth/normal passes from S* independently of beauty images. Defocus and stylized marks can broaden or alter apparent boundaries even when geometric visibility is unchanged. Define what is invariant rather than promising identical perceived segmentation at all media settings.

Optical endpoint repair. Keep the D3500-format 23.5×15.6 mm sensor, 6000×4000 nominal output and fixed 200 mm focal length when that endpoint is intended. For distances and focal length in common units, the ideal thin-lens blur diameter is $c=f^2|s_{obj}-s_{focus}|/[N_f\, s_{obj}(s_{focus}-f)]$, where $N_f>0$ is the f-number and $s_{obj},s_{focus}>f$. Horizontal field of view is approximately 6.72°. This is an ideal photographic rendering specification, not proof a physical camera captured the image. Real lens, aperture, spectral response, exposure, demosaicing, aberrations and tone response need calibration. Depth blur generally requires occlusion-aware lens integration rather than one global convolution. At q=0 no photographic defocus; at q=1 no painted or ink marks, as original design endpoint contracts.

5. Mixed discrete/continuous synthesis, constraints and objectives

Let d in a finite or explicitly countable admissible design set encode scene graph, topology, materials' identities, seam topology, motif grammar, historical/fiction labels, event structure and tool/process family. Let x_d in $R^{n(d)}$ be normalized continuous coordinates conditional on d. The complete feasible space is the disjoint union

$$F = \bigcup_{d\in D_{admissible}} \{d\} \times F_d.$$

Here and in the controller, x is a normalized finite control vector, distinct from x_{world}. Different topology may change dimension. Some F_d are nonconvex, disconnected, empty, or known only through an expensive validator. A generic argmin symbol neither produces a feasible design nor proves a minimum exists.

Separate pre-commit generation of world, body, clothing, composition and causal history from post-commit depiction control. Discrete changes after commit branch the scene/version. General nonconvex contact, anatomy, inherited structure, cultural interpretation and history cannot be repaired by an ordinary Euclidean projection onto one convex set. Euclidean projection onto a nonconvex set may be multivalued and discontinuous. For $K=\{-1,+1\}$, projection at zero has two values and any nearest-point selection jumps across zero. This analytic counterexample is not a corpus observation.

Hard constraints include declared scene identity, source requirements, budgets and any testable mechanical/biological/world-law bounds. Unknown hard constraints prevent certification. Soft objectives include composition, expression, craft, stylistic fit, novelty, detail hierarchy, attention and subjective preference. Preserve v2.2 $E_{composition}$, E_{org}, E_{world}, E_{orn}, E_{symbol}, $E_{physics}$,

E_dev/ecology/reproduction, E_fit/era/culture/fabric, E_social, E_artifact and resource diagnostics as named components, each with scale, evaluator, evidence and uncertainty.

Use either a documented weighted sum of normalized terms or a multiobjective/Pareto comparison. Keep the original Ω_art product only as an optional user-calibrated heuristic; it is not a universal measurement and a zero factor annihilates the product. Do not trade a failed hard constraint for a higher aesthetic score. An empirical judge or reward model may be biased or optimized against; hold out independent evaluation and compare with a simpler baseline.

Theorem 3. Contractive projected quadratic depiction control

Fix a scene/version, a discrete design branch, a q, a nonempty closed convex $K\subset R^n$, and a symmetric matrix H with $\mu I \preccurlyeq H \preccurlyeq LI$, $0<\mu\le L<\infty$. Let

$$J_q(x)=\tfrac{1}{2}x^{T}Hx-b(q)^{T}x+c(q),$$

$$T_q(x)=\mathrm{Proj}_K[x-\eta(Hx-b(q))],\ 0<\eta<2/L.$$

Then J_q has a unique constrained minimizer $x^*(q)$, it is the unique fixed point of T_q, and

$$\|T_q(x)-T_q(y)\| \le \rho\,\|x-y\|,$$

$$\rho=\max\{|1-\eta\mu|,\ |1-\eta L|\}<1.$$

Consequently the exact iteration converges geometrically: $\|x_k-x^*\|\le\rho^k\|x_0-x^*\|$. For the disturbed update

$$x_{k+1}=\mathrm{Proj}_K[x_k-\eta(Hx_k-b(q))+\eta u_k],$$

one has

$$\|x_k-x^*\| \le \rho^k\|x_0-x^*\| + \eta\sum_{j=0}^{k-1}\rho^{k-1-j}\|u_j\|.$$

In particular, $\|u_j\|\le u_{max}$ gives $\limsup_k \|x_k-x^*\|\le\eta u_{max}/(1-\rho)$; $u_j\to 0$ gives convergence to x^*. If the same H and K apply for all q and $\|b(q)-b(q')\|\le L_b|q-q'|$, then

$$\|x^*(q)-x^*(q')\| \le (L_b/\mu)|q-q'|.$$

Proof. Strong convexity and coercivity of J_q on nonempty closed K give existence and uniqueness. The variational inequality $\langle Hx^*-b(q),y-x^*\rangle\ge 0$ for all $y\in K$ is equivalent to $x^*=\mathrm{Proj}_K[x^*-\eta(Hx^*-b(q))]$. Projection onto a closed convex set is nonexpansive; the spectral norm of $I-\eta H$ is at most ρ, so T_q is a contraction. Compare the disturbed update with $T_q(x^*)$ and use projection nonexpansiveness to obtain $e_{k+1}\le\rho e_k+\eta\|u_k\|$. Unrolling gives the stated bound. The bounded-input limit follows by a geometric series. For vanishing input, split that convolution into a finite prefix tending to zero and an arbitrarily small tail. For parameter sensitivity, add the two optimality inequalities with each minimizer as the other's test point: $\mu\|x^*(q)-x^*(q')\|^2\le\langle b(q)-b(q'),x^*(q)-x^*(q')\rangle$. Apply Cauchy-Schwarz and divide when the difference is nonzero. ∎

Useful extension, same assumptions. When q changes at iteration boundaries, define $e_k=\|x_k-x^*(q_k)\|$ and use $b(q_k)$ in the update. Then

$$e_{k+1} \le \rho e_k+\eta\|u_k\|+(L_b/\mu)|q_{k+1}-q_k|.$$

A changing target creates a tracking error; the theorem does not permit claiming convergence to a single fixed optimum while q keeps changing. Feedback $u=u(x)$ with Lipschitz constant L_u yields a contraction of the perturbed map only if $\rho+\eta L_u<1$; its fixed point need not equal the unforced minimizer. Additional dynamic memory/controller states require a certificate for the full augmented system.

Scope and explicit exclusions. This theorem covers a declared quadratic reference subproblem with convex bounds, not the complete v2.2 objective, the biological/world constructor, neural synthesis, perception, a nonconvex mesh optimizer, or the original quasilinear PDE. A locally quadratic approximation is not globally the same function. If H is semidefinite, indefinite, q-dependent, K varies,

or projection is inexact, a new analysis is needed. The added regularization necessary to make H positive definite changes the design objective and must be reported.

6. Bounded Praxium feedback and reversible refinement

Use the v2.2 affect, telemetry, autonomy, memory, PSR and layer ideas as explicit control interfaces. They are not actual feelings, sentience, biological drives or unrestricted autonomy. Telemetry must be measured and normalized; affect coordinates are expressive/control variables with declared mappings.

6.1 Finite bounded state and regular gates

Replace the merely nonnegative unbounded core by declared operational limits:

$$K_z=[0,1]\times[0,\kappa_{max}]\times[0,N_{max}]\times[0,D_{max}]$$
$$\times B_2(R_e)\times B_2(R_c)\times B_2(R_s)\times B_2(R_\tau)\times B_2(R_g),$$

with finite nonnegative caps and finite radii. This is a nonempty compact convex set. The caps are design limits, not evidence that the unclamped legacy dynamics stay bounded. Add every derivative filter, accumulator and memory state to the implemented state record.

For $\Delta s>0$ and α_v,α_a in $(0,1]$, define bounded derivative estimates

$$v_{k+1}=(1-\alpha_v)v_k+\alpha_v\ \mathrm{clip}((\Xi_k-\Xi_{k-1})/\Delta s,-v_{max},v_{max}),$$

$$a_{k+1}=(1-\alpha_a)a_k+\alpha_a\ \mathrm{clip}((v_{k+1}-v_k)/\Delta s,-a_{max},a_{max}).$$

With bounded initial states, induction preserves the stated bounds. These are filtered first and second differences per normalized solver time, not guaranteed derivatives of a mild PDE solution. Replace the legacy Ξ and $\ddot{\Xi}$ gates by v and a. Smooth stall weight, for example

$$St_k = \mathrm{sigmoid}((\varepsilon_E-Ex_k)/T_E)\ \mathrm{sigmoid}((\Xi_{stall}-\Xi_k)/T_\Xi),$$

has positive temperatures and declared thresholds; it is not the same logical predicate as the original indicator. Hysteresis is an alternative, but then a switching-system analysis is required.

Use bounded sigmoid gain scheduling $m_v=m_v^-+(m_v^+-m_v^-)\mathrm{sigmoid}(g_v)$, finite $0<m_v^-\le m_v^+$, and explicitly capped H and Γ. Evaluate $A(s)=\exp(\min(\log A_{max},\Lambda s))$ with $A_{max}\ge 1$ and $\Lambda\ge 0$ to avoid intermediate exponential overflow. Its saturation makes it bounded, not necessarily differentiable at the cap. Positive resource/physiology constraints may instead set a gain to zero when the corresponding capability is absent; do not simultaneously demand a strictly positive lower bound there.

6.2 Total flux, zero need, memory and leakage

Use a finite set of sources U and destinations V, or declare volume weights for a finite discretization. Let $b_u\ge 0$ be bounded source readiness and $n_v\ge 0$ destination need. Let $D_{uv}\ge 0$ be dimensionless normalized distance and $\varepsilon_D>0$. Set

$$a_{uv}=n_v/(D_{uv}+\varepsilon_D),$$

$$w_{uv}=a_{uv}/(\Sigma_v a_{uv}) \text{ when } \Sigma_v a_{uv}>0;\ \text{otherwise } w_{uv}=0,$$

$$F_{uv}=F_{budget}\,[b_u/(\varepsilon_b+\Sigma_u b_u)]\,w_{uv},$$

where $F_{budget}\ge 0$ and $\varepsilon_b>0$ are fixed dimensionless per-update control bounds. Thus all entries are nonnegative and $\Sigma_{uv} F_{uv}\le F_{budget}$, including the zero-need case. This supplies the missing total bound; a bound on each source alone does not. These are control allocations, not physical energy transfers. If interpreted as a conserved physical/resource flow, use matched inflow/outflow bookkeeping and account for projection/saturation losses explicitly. Destination capacities can be additional transport constraints.

A leaky nonnegative stall accumulator can use

$$g_{k+1}=(1-\alpha_g)g_k+\alpha_g G_k,\ 0<\alpha_g\le 1,$$

with $G_k\in[0,G_{max}]$ and g_0 in the same interval. Memory channels with a frozen input x_k-x_{ref} over the step use the exact exponential update

$$m_{j,k+1}=\exp(-\beta_j\Delta s)m_{j,k}$$
$$+(a_j/\beta_j)(1-\exp(-\beta_j\Delta s))(x_k-x_{ref}),$$

where $\beta_j>0$, $a_j\ge 0$. This is exact only for the declared finite exponential kernel and held-constant input; general kernels require approximation error and compatible initial prehistory. Bounded input yields bounded memory. Positive memory feedback can still destabilize a coupled system: the illustrative continuous system $\dot{x}=-x+m$, $\dot{m}=-m+2x$ has eigenvalues $-1\pm\sqrt{2}$, including a positive one, despite stable isolated decay and an integrable exponential memory kernel. This is an analytic counterexample, not a simulated artistic result.

An implementable bounded controller is

$$z_{k+1}=\mathrm{Proj}_{K_z}[z_k+\Delta s(-\Lambda_z(z_k-z_{home})+G_z(x_k,z_k,m_k,telemetry_k))],$$

$$u_k=\mathrm{Proj}_{B_2(u_{max})}[G_u(z_k,m_k,layer_readouts_k,PSR_k)].$$

Finite bounds, typed domains and finite gains are mandatory. Projection establishes feasibility/boundedness of z and u; it does not prove contraction or improve art. PSR is a declared bounded input here. A genuine spatial boundary control needs its own weak formulation, trace space, boundary conditions and discretization scaling; it is not added again as an ordinary volume source.

If each augmented subsystem has a proved incremental bound $e_{next}\le G\,e$ with a nonnegative full gain matrix G and spectral radius below one, a suitable weighted norm yields a contraction certificate. This statement requires the subsystem estimates and includes all memory, layer, filter and controller loops. Checking only an isolated raw coupling matrix, a pair of gains, or Metzler-Hurwitz diagonal blocks does not supply those estimates.

6.3 Acceptance, rollback and version boundaries

Maintain a verified checkpoint $C_k=(S^*,x_k,z_k,m_k,filters,PRNG\ state,dependencies,outputs,validation\ record)$. The scene portion is read-only. An update is proposed in an isolated working copy and is not a committed improvement merely because the controller requested it.

1. Freeze the current objective/evaluator, feature scales, q and admissible branch for the refinement epoch. Compute a candidate and a change summary.

2. Verify scene/dependency integrity and all required hard constraints. A required unknown is not pass. Evaluate independent artistic/contextual criteria with their uncertainty.

3. When a scalar technical descent criterion is appropriate, accept only if $J(candidate)\le J(current)-c_{acc}\|candidate-current\|^2$, $c_{acc}>0$, or use another explicitly declared acceptance policy. Do not treat noisy estimated improvement as certified descent.

4. On failure, restore the last verified computational checkpoint, reduce the permitted step/gain or change the proposal, and record why. Cap retries and budget; if no acceptable update exists, return the last verified result with a stall/limitation report. A larger request that changes scene meaning creates a new explicit version.

5. Save the new accepted checkpoint and its parent link. Rollback covers this recorded computational state; it cannot promise reversal of external transactions, published communications or lost source data.

For x in K and the exact undisturbed quadratic step with $0<\eta<2/L$, the projection optimality inequality gives

$$J(x_{next})-J(x)\le -(1/\eta-L/2)\|x_{next}-x\|^2,$$

so an appropriate descent threshold is justified. With arbitrary feedback, inexact projection or a different nonconvex objective, acceptance must be checked rather than inferred from that formula. Rejection preserves the last verified state but does not prove a future improving step will be found. Optimizing one metric does not guarantee improved beauty, originality, cultural appropriateness or user preference.

7. PDE, biological, cultural and causal limits preserved

The original reaction-advection-diffusion, graph, nonlocal, memory and boundary architecture is retained as a research extension, not silently asserted as the same system proved above. To claim well-posedness or discretization order, specify: function/state spaces; operator domains; physical/numerical boundary conditions; ellipticity on diffusive channels; state-dependent coefficient regularity; admissible boundary control; graph/kernel bounds; normal-cone constraints; local existence/continuation conditions; full a priori estimates; and the exact solver. Generic Lie-Trotter order and CFL bounds do not validate a pipeline containing arbitrary inexact nonconvex repair steps.

Keep the world, life, evolution, ecology, pan-species couture and society modules. Empirical heredity/development is a model family requiring mechanisms and data. A finite energy budget or feasible mesh is necessary for some physical claims but insufficient to prove organism viability. Chimeras, mosaics, hybrids and cyborgs remain distinct. Sex phenotype, sex system, identity, gender and social role remain separate records; none is inferred from costume.

Universal Etheros lineage-level reproductive compatibility remains X, a world axiom involving Λ_0, adapters, developmental/gestational support and inherited consequences. It is not real biomedical capability. Viability and fertility thresholds have no empirical meaning until assays, populations, units and uncertainty are supplied. Consent, autonomy and welfare are constraints on the fictional depiction, not biochemical derivations. Living architecture's $\beta_{primary} \geq 0.90$ remains an authored fictional design threshold with a declared denominator and load case; it is not a natural constant or a safety proof. Clarify whether β measures gravity-load share, structural mass, or another quantity before using it.

Preserve regionally overlapping couture timelines, historical construction/tool/material causality, garment access/maintenance and body-specific mechanics. Artist/culture names remain research anchors, not exhaustive definitions or permission to assign symbols by resemblance. Each actionable cultural claim needs source/object, date range, region/community, social use and confidence; documented meaning and stylistic inference are separate.

Causal world snapshots use bounded relevant context: prior event, present activity/process, relationships/resources, material condition and plausible next event. Inert objects have purpose/function, not invented consciousness. Conditional future trajectories are predictions or authored possibilities, not known futures. Off-screen state persists at a declared fidelity; these records do not claim a complete simulated planet.

8. Verification and evidence ledger

Replace the unqualified product PASS=Q_A…Q_I by a result record for each check: status in {pass, fail, unknown, not applicable}, method, tolerance, source, timestamp, scope and rationale. Record what is exact, tested numerically, judged empirically, asserted by fictional axiom or not checked.

• A / scene: exact identity of canonical asset and state records; distinguish hashes from dependency completeness and canonical-byte comparison.

• B / world biology: E/P/X-specific logical and physical criteria; no biological certification from fiction or symbolic placeholders.

• C / cloth: unit-consistent materials, penetrations/contact, strain, motion/pose envelope and stated comfort evidence.

• D / provenance: historical and symbolic claims with linked sources, not merely visual similarity.

• E / coordinate: exact declared q and medium-mixture constraints, not measured percent realism.

• F / optics: calibrated or idealized camera model, projection and occlusion, stated approximation error.

• G / causality: bounded causal/event consistency and resource/flow residuals at declared fidelity.

• H / composition: context-feasible ensemble variety, semantic repetition and attention evaluation; aesthetic ratings remain population/task-dependent.

• I / archive and slider: source/dependency recovery; output-bit preservation; deterministic or tolerance-based replay; anchor/near-anchor tests and applicable analytic continuity bound.

• J / v2.3 material-process-pattern: physical versus depiction material separation, process provenance, field units, topology and multiscale linking.

• K / v2.3 temporal: frame/event timing, captured external inputs, trajectory lock, deterministic replay scope and interaction branches.

• L / v2.3 controller: bounded state/total flux, objective/evaluator version, observed acceptance/rejection, rollback test and full-state stability evidence where claimed.

Issue a certification only for the explicit relevant scope when every required hard check passes. Otherwise report the available artifact with the narrower mode and failed/unknown checks. Finite q sampling is useful regression testing but not a proof over all real q. A theorem about the pipeline does not verify that a particular implementation satisfies its assumptions.

Evaluation plan, not results. Once a corpus is available, compare the baseline v2.2-derived descriptors with v2.3 typed material/process, attention, pattern and temporal fields on held-out works. Track missingness and uncertainty. Use matched-resource ablations, source-referenced reconstruction tasks, expert/user review and independent metrics. The same metric used to select a candidate should not be the only evidence that it improved. Never label a proposed addition “discovered in the corpus” until its cited observations exist.

9. Integration map to every original module

Original v2.2 section | v2.3 treatment and location

0. Notation/contract | Retain q, immutable S*, S_0, source domains and typed synthesis. Repair logarithm units, measure/time separation, variable collisions and named field schemas; §§0-2.

1. Mastercraft/symbol/history | Retain all mastercraft, cross-cultural, design/media and architectural registers. Type operators; separate construction evidence from style attribution; add process/attention/pattern fields; §§1,3,5,7.

2. World/life/evolution [E→X] | Retain world, organisms, anatomy, evolution, ecology, kinship and Etheros compatibility. Keep empirical biology distinct from plausible design and fictional axiom; §§0,5,7-8.

3. Couture chronology/causal operators | Retain H_0…H_9 overlapping regional history, causal garment grammar and every stitch/process family. Add provenance/confidence and typed material/process records; §§3,7.

4. Pan-species fit/garment physics | Retain patterns, seams, body access, fit, thermomechanics and dignity. Repair stiffness/source units; separate static/dynamic and tested/speculative safety; §§3.1,5,7-8.

5. Social fashion/pan-species society | Retain individual agency, class/context, nonuniform society, living architecture and social objectives. Define fictional load-share threshold and separate identity from inferred biology; §§5,7.

5A. Organic composition/series debias | Retain feasible layouts and contextual debias heuristic. Add reference measure, no-primary/ensemble bins, attention and semantic series dimensions; §3.3.

5B. Causal world in progress | Retain prior/present/future consistency, relevant off-screen state and bounded policy. Define fidelity and prevent invented prescience or total-world simulation claims; §§1,3.4,7.

5C. Canonical scene/exact slider lock | Retain pre-render commit, exact asset lock and persistent UID registry. Add modes, dependency-closed archive, source retention, output retention and temporal/event contracts; §§1-2,3.4,4.

6. Realism/optics/line | Retain ten anchors and all original medium weights. Prove only the declared interpolation result; separate geometric and apparent lock; repair optical claims and units; §§3.2,4.

7. Praxium affect/autonomy | Retain bounded gain/affect interfaces and telemetry purpose. Add finite core caps, filtered derivatives, bounded total allocation and explicit no-sentience interpretation; §6.

8. Fusion/boundary/memory | Retain core/latent/layer, graph, nonlocal, Prony and PSR modules. Define finite controller first; audit entire coupling before stability claims; preserve PDE research with required hypotheses; §§1,6-7.

9. Synthesis/invariants/optimality | Retain every named hard/soft objective module. Separate discrete/continuous branches and nonconvex feasibility from the proved convex reference subproblem; §5.

10. Stability/numerics/verification | Replace broad theorem language with three conditional proofs, bounded-input limits and explicit implementation assumptions. Expand QA into scoped evidence records; §§2,4-8.

11. Execution | Preserve generate → compose → develop/fit → commit → depict → verify → return. Add archive ingestion, declared mode, isolated proposals, bounded feedback, rollback and evidence-bearing delivery; §10 below.

10. Revised execution contract

```
CALCUART_v2.3(request, r, context, source_records, scene_id?, version_id?):
  q ← clip(r/100,0,1); disclose any clipping
  declare exact-asset / constrained-inference / concept mode
  ingest and preserve source bytes; validate metadata and dependency links
  if scene_id and version_id exist:
    load authoritative commit; validate dependencies and requested scope
  else:
    choose typed discrete world/process/design branch
    construct bounded causal context, body/garments and organic composition
    solve/check applicable nonconvex feasibility; record unknowns
    commit immutable snapshot or declared trajectory before any q render
  freeze style contract, feature scales, evaluators and q-refinement epoch
  derive ten-anchor controls; initialize bounded Praxium/controller state
  repeat within declared resource budget:
    propose a typed depiction update in an isolated copy
    verify hard constraints and the declared acceptance test
    accept and checkpoint, or restore and reduce/revise the proposal
  render in the declared mode and execution environment
  preserve output bytes, dependency manifest and scoped evidence record
  return artifact + version + what is exact/tested/proposed/unknown
```

Same scene/version implies the same declared world state or trajectory across q. A new time, angle, action, object, physical material, or event branch creates a new explicit version. The complete system remains an ambitious art-making architecture with a small proved computational core, empirical modules and explicit fiction, not a proof that art, biology, history or perception can be reduced without remainder to one scalar calculus.

Appendix A inherited realism controls

These authored control values are transcribed from section 6 of the supplied CalcuArt v2.2 source and retained by v2.3. They are design choices, not measured percentages of realism or a calibrated physical mixture. They support section 4 of this calculus.

For $j = 0,\ldots,9$, $q_j = j/9$ and $w_j(q) = [1 - |9q - j|]_+$. The medium vector is ordered as $B = (\text{ink}, \text{watercolor}, \text{gouache}, \text{acrylic}, \text{oil}, \text{physical}, \text{DSLR})$, with $B(q) = \Sigma_j w_j(q) B_j$. Each row is nonnegative and sums to one.

j	Depiction anchor R_j	Seven-component weight B_j
0	Monochrome engraving	(1, 0, 0, 0, 0, 0, 0)
1	Ink and watercolor	(.75, .25, 0, 0, 0, 0, 0)
2	Ornamental watercolor	(.45, .40, .15, 0, 0, 0, 0)
3	Ink and gouache	(.35, .35, .20, .10, 0, 0, 0)
4	Classical paint	(.15, .15, .20, .20, .30, 0, 0)
5	Luminous naturalism	(.08, .07, .10, .15, .60, 0, 0)
6	Cinematic oil	(.04, .06, 0, 0, .70, .20, 0)
7	Hyperreal paint	(.02, .03, 0, 0, .45, .50, 0)
8	Cinematic photo	(0, 0, 0, 0, .05, .25, .70)
9	Nikon D3500 at 200 mm	(0, 0, 0, 0, 0, 0, 1)

Line control and optical parameters

The inherited line-control vector is $\mathcal{I} = (\text{orientation}, \text{width}, \text{curvature}, \text{density}, \text{rhythm}, \text{phase}, \text{pressure}, \text{variance})$, with $\mathcal{I}(q) = \Sigma_j w_j(q)\mathcal{I}_j$. Ink is fully enabled at $q = 0$ and absent at $q = 1$; geometric and material edges may remain. The source supplies no numerical $\mathcal{I}_j$ values or component scales. Implementations must declare them before the interpolation can be evaluated. The associated source coordinate $F = q$ has $dF/dr = 0.01$; this is a coordinate identity, not a perceptual law.

The declared endpoint uses a 23.5 × 15.6 mm sensor, nominal 6000 × 4000 output and fixed focal length f = 200 mm. The inherited default f-number is $N_f = 5.6$. The ideal thin-lens image distance is $s_{image} = f\, s_{obj}/(s_{obj} - f)$, with all lengths in a common unit and $s_{obj} > f$. The depth-blur diameter and optical limitations are those in section 4. A camera label alone does not establish calibrated capture.

Source and scope of this extract

The body preserves formal-kernel sections 0-10, clarifying that numerical line-control anchors were not supplied. Appendix A transcribes the inherited R_j labels and B_j weights. Typographic dashes are normalized; equations and proof assumptions are otherwise retained.

The original v2.2 historical appendix is omitted, so this is not a verbatim archive of every legacy equation. The v2.3 kernel narrows conflicting legacy claims. The full nonlinear/PDE, biological, cultural and rendering architecture remains unverified.

Source edition: CalcuArt Theorem v2.3, prepared 7 November 2026. Parent source: CalcuArt v2.2-Ω^9, supplied message bbb214e7-b909-4dc2-9f65-4b58cc359efa.