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,