

Higher E-Graphs: Cartesian Factorization, Proof Spaces, and Topos Semantics

Anonymous

October 21, 2026 · Version 1.0

Abstract

We extend initial-functor semantics for cartesian term diagrams to homotopy-coherent diagrams and to models in an infinity-topos. We prove product preservation for space-valued and topos-valued left Kan extensions, obtain cartesian comprehensive factorization, and characterize the transformations invisible to truncated observations. A proof-space interpretation separates congruence from coherence. Slice semantics supplies parameter dependence and descent, while explicit examples identify the limits of truncation and finite-product reasoning.

1. Question, contribution, and conventions

Patterson [1] interprets equational reasoning through finite-product diagrams and initial changes of indexing category. The present paper asks what happens when solutions include witnesses of equality, paths between witnesses, and further coherence. Its contribution is a synthesis with explicit proofs and limitations: a cartesian comprehensive factorization in $(\infty, 1)$ -categories, a truncation-sensitive soundness criterion, and a semantics of parameterized constraints in slices of an ∞ -topos. These are applications and extensions of established categorical machinery, not a claim to originate comprehensive factorization, higher rewriting, or descent.

An **∞ -category** means an $(\infty, 1)$ -category. Its mapping objects are spaces; morphisms of dimension at least two are invertible. Write $\mathcal{S}$ for spaces, $|A| = \operatorname{colim}_A *$ for groupoid completion, and $\tau_{\leq n}$ for truncation. All indexing categories are small relative to fixed universes. Limits, colimits, and commutative diagrams are homotopy coherent. “Cartesian” means finite-product-preserving; it does not mean finite-limit-preserving. We distinguish this usage from cartesian fibrations.

For $u : A \rightarrow B$, set $A_{/b} = A \times_B B_{/b}$, the comma ∞ -category of arrows $u(a) \rightarrow b$. Call u **initial** when $|A_{/b}| \simeq *$ for every b . This is the convention appropriate to limits, also called left cofinal [4]. Initiality is stronger than an equivalence after groupoid completion.

The **left fibration** $\int_B Y \rightarrow B$ classified by $Y : B \rightarrow \mathcal{S}$ has fiber $Y(b)$ and covariant transport. This variance is essential: limits are computed by sections of left fibrations, whereas their classification uses covariant functors.

2. The cartesian Kan-extension theorem

Lemma 2.1 (Products in the elements construction). If C has finite products and $X : C \rightarrow \mathcal{S}$ preserves them, then $E = \int_C X$ has finite products. The projection $p : E \rightarrow C$ preserves and reflects product cones.

Proof. For finitely many (c_i, x_i) , choose $c = \prod_i c_i$. The equivalence $X(c) \simeq \prod_i X(c_i)$ gives a contractible space of lifts of the tuple (x_i) , including its comparison paths. These data define the product cone. To check its universal property, for any (d, z) the space of arrows to (c, x) is the homotopy fiber, over x , of

$$\mathrm{Map}_C(d, c) \longrightarrow X(c), \quad f \longmapsto X(f)(z).$$

The product comparisons identify this fiber with the product of the corresponding fibers for (c_i, x_i) . For the empty family, $X(1_C) \simeq *$. The same calculation shows that any cone projecting to a product cone is a product cone. $\square$

Lemma 2.2 (Homotopy cosiftedness). If E has finite products, then E^{op} is sifted in the ∞ -categorical sense.

Proof. It has an initial object and finite coproducts. For a pair of objects, its category of common targets has the coproduct cocone as an initial object, hence contractible groupoid completion. The siftedness criterion [5] applies. Equivalently, each finite diagonal is final because it has a left adjoint. Thus colimits of this shape in spaces commute with finite products. Connectedness alone would not suffice here. $\square$

Theorem 2.3 (Product-preserving Kan extensions). Let C, D be small cartesian ∞ -categories, let $X : C \rightarrow \mathcal{S}$ preserve finite products, and let $F : C \rightarrow D$ be any functor. Then $\mathrm{Lan}_F X$ preserves finite products. No product-preservation assumption on F is required.

Proof. Covariant density and the Yoneda lemma give

$$X \simeq \mathrm{colim}_{(c,x) \in (\int_C X)^{\mathrm{op}}} \mathrm{Map}_C(c, -),$$

and therefore

$$\mathrm{Lan}_F X \simeq \mathrm{colim}_{(c,x) \in (\int_C X)^{\mathrm{op}}} \mathrm{Map}_D(Fc, -). \quad (1)$$

Indeed, left Kan extension takes the displayed covariant representable to $\mathrm{Map}_D(Fc, -)$ by its defining adjunction. Each summand preserves products. Lemmas 2.1 and 2.2 make the indexing category sifted, so its colimit preserves finite products, including the terminal object. $\square$

This proof is the space-valued analogue of the classical argument associated with Borceux and Day [2]. Its essential replacement is homotopy siftedness rather than ordinary siftedness. The ambient density and sifted-completion results are standard [3, Sections 5.1 and 5.5.8].

Corollary 2.4 (Topos-valued form). Theorem 2.3 holds with any ∞ -topos $\mathcal{X}$ in place of $\mathcal{S}$.

Proof. Present $\mathcal{X}$ as an accessible left exact localization $L : \mathcal{P}(K) \rightleftarrows \mathcal{X} : i$. The right adjoint i preserves products. Evaluate iX at each object of K and apply Theorem 2.3; the pointwise left Kan extension in $\mathcal{P}(K)$ preserves products. Applying the product-preserving left adjoint L gives the left Kan extension in $\mathcal{X}$. This follows from $LiX \simeq X$ and the Kan-extension adjunction. $\square$

3. Cartesian comprehensive factorization

Let Cart_∞ have small cartesian ∞ -categories as objects and product-preserving functors as morphisms. An orthogonal factorization system here means that every morphism factors into the two classes and each compatible lifting problem has a contractible space of fillers.

Theorem 3.1. The initial/left-fibration factorization system on Cat_∞ restricts to Cart_∞ . Explicitly, for a cartesian $F : C \rightarrow D$, put

$$Y_F(d) = \text{colim}_{(c, Fc \rightarrow d) \in C_{/d}} * = |C_{/d}|, \quad E_F = \int_D Y_F.$$

There is a factorization $C \xrightarrow{\eta_F} E_F \xrightarrow{p_F} D$, with η_F initial, p_F a left fibration, and both functors cartesian. It is unique up to contractible choice within the factorization system.

Proof. The higher comprehensive factorization is standard [6, Section 7.4]; the displayed construction is its covariant Kan-extension description. Namely $Y_F = \text{Lan}_F 1_C$, and the unit supplies η_F over D . The defining universal property is

$$\text{Map}_{/D}(E_F, \int_D Z) \simeq \text{Nat}(Y_F, Z) \simeq \text{Nat}(1_C, ZF) \simeq \text{Map}_{/D}(C, \int_D Z)$$

for every $Z : D \rightarrow \mathcal{S}$. This is the left-fibration reflection underlying comprehensive factorization.

Theorem 2.3 makes Y_F cartesian. Lemma 2.1 makes E_F cartesian and makes p_F preserve and reflect products. Since $p_F \eta_F \simeq F$ preserves products, reflection implies that η_F does too.

For completeness, orthogonality restricts as well. A left fibration $p : E \rightarrow D$ between cartesian categories that preserves products has a product-preserving classifier: its product cones identify the fibers over $\coprod d_i$ with the product of the fibers over d_i , including the empty case. Lemma 2.1 then shows that p reflects product cones. In a lifting square of cartesian functors, the ambient filler h has ph cartesian; hence h is cartesian. Consequently the space of cartesian fillers is the same contractible space as the ambient one. Closure under retracts and the ambient characterization of the two classes finish the restriction. $\square$

Corollary 3.2 (Ordinary shadow). For an ordinary cartesian functor $F : C \rightarrow D$, the covariant set-valued functor $d \mapsto \pi_0 |C_{/d}|$ classifies its ordinary discrete-opfibration factor. Thus the classical construction is recovered by taking components of the classifier, not by identifying discrete opfibrations with all left fibrations.

Proof. The left adjoint $\pi_0 : \mathcal{S} \rightarrow \text{Set}$ preserves colimits and finite products. Therefore $\pi_0 Y_F$ is the ordinary left Kan extension of the terminal set-valued functor. Its category of elements is precisely the ordinary comprehensive factor [7]. $\square$

4. Solutions and levels of observational equality

Fix a cartesian ∞ -category T and a model $M : T \rightarrow \mathcal{X}$ preserving products. A **higher term diagram** is a cartesian $D : J \rightarrow T$. Its solution object is

$$\text{Sol}_M(D) = \lim_{j \in J} MD(j).$$

A transformation consists of a cartesian $u : J \rightarrow K$ and a specified equivalence $D \simeq D'u$. Its comparison on solutions goes in the reverse direction. Points of a solution object include compatible values and all required coherence; merely solving equations in the homotopy category loses this information.

Theorem 4.1 (Soundness). If u is initial, its comparison on solutions is an equivalence. A zigzag of such transformations preserves the solution object up to equivalence.

Proof. Apply the initial-functor theorem to MD' . Alternatively, apply every generalized-point functor $\text{Map}_{\mathcal{X}}(U, -)$, use the space-valued theorem, and invoke Yoneda. Compose the resulting equivalences along the zigzag. $\square$

Definition 4.2. For $n \geq 0$, call $u : A \rightarrow B$ **n -initial** if $\tau_{\leq n}|A/b| \simeq *$ for every b . This definition measures the truncation level of observations; it does not assert a separate factorization system for every n .

Theorem 4.3 (Exact truncation criterion). The following are equivalent:

1. u is n -initial.
2. For every $H : B \rightarrow \mathcal{S}_{\leq n}$, restriction induces an equivalence $\lim_B H \simeq \lim_A Hu$.

They imply the same statement for diagrams of n -truncated objects in any ∞ -topos, and in any of its slices.

Proof. In $\text{Fun}(B, \mathcal{S}_{\leq n})$, the left Kan extension of the terminal diagram is

$$Y_n(b) = \tau_{\leq n} \text{colim}_{A/b} * = \tau_{\leq n}|A/b|.$$

Restriction on limits is the map

$$\text{Map}(1_B, H) \longrightarrow \text{Map}(Y_n, H)$$

induced by $Y_n \rightarrow 1_B$, using the Kan-extension adjunction. It is an equivalence for every H exactly when $Y_n \rightarrow 1_B$ is an equivalence, by Yoneda. Generalized points of an n -truncated object are n -truncated spaces and preserve limits, proving the topos and slice assertions. $\square$

Thus 0-initiality recovers the nonempty-connected comma criterion of ordinary category theory, while full initiality requires contractibility. This distinction is already visible for $BG \rightarrow *$, where G is a nontrivial discrete group. It is 0-initial, but the comparison $BG \rightarrow \text{Map}(BG, BG)$ for the constant diagram with value BG is not an equivalence: its image consists of constant maps, and the identity is not null-homotopic. This example concerns general indexing categories; BG is not being presented as a cartesian category.

5. Proof spaces and the lens calculation

For a fixed context Γ and output sort A , terms belong to a mapping space $\text{Map}_T(\Gamma, A)$. A chosen rewrite becomes a path in that space. Rewriting inside an operation or substitution is induced by composition and finite products. A path between parallel derivations is a two-dimensional coherence in the mapping space, hence one dimension higher than a rewrite. Directed reductions can instead be retained in an $(\infty, 2)$ -categorical presentation; passing to equality semantics inverts them.

One concrete, deliberately limited construction starts with a fixed finite graph G of represented terms and witnessed rewrites. Its free proof space is the realization $|G|$. Specified comparisons between derivations attach two-cells along the corresponding loops; higher specified comparisons attach further cells. This construction describes the represented fragment, not a complete algorithm for all substitution instances. Compatibility with operations must be supplied by maps of these spaces or by a coherent categorical presentation. Polygraphs provide established machinery for that additional structure [8].

Proposition 5.1 (Congruence does not determine coherence). The set $\pi_0|G|$ is the equivalence closure of the edges of G . It does not determine the homotopy type of a proof space, even for finite presentations.

Proof. Two vertices lie in the same component exactly when joined by a finite unoriented edge path. Consider vertices s, a, b, t with edges $s \rightarrow a \rightarrow t$ and $s \rightarrow b \rightarrow t$. Their realization is a circle. Attaching a disk along the diamond yields a contractible space. Both have one component, but their fundamental groups are $\mathbb{Z}$ and 0 . Attaching cells of dimension at least two never changes components. Hence equality of classes cannot certify equality of proof spaces. $\square$

Even a termination-and-confluence argument producing common normal forms supplies only part of the desired coherence. Results such as Squier’s theorem require precise presentation hypotheses and construct specified homotopy bases; they do not say that arbitrary equality saturation automatically yields a contractible space of derivations [8].

The lens example admits a particularly explicit witness. In the internal language of $\mathcal{X}$, suppose maps $g : S \rightarrow V$, $p : S \times V \rightarrow S$ come with homotopies

$$\alpha_{s,v} : g(p(s, v)) = v, \quad \beta_s : p(s, g(s)) = s.$$

Concatenation $r \cdot q$ means first r , then q .

Proposition 5.2 (Witnessed PutTwice). There is a specified homotopy

$$\delta_{s,v} : p(p(s, v), v) = p(s, v), \quad \delta_{s,v} = (\text{ap}_{p(p(s,v), -)} \alpha_{s,v})^{-1} \cdot \beta_{p(s,v)}. \quad (2)$$

Proof. Put $x = p(s, v)$. Applying $p(x, -)$ to $\alpha_{s,v}$ gives a path from $p(x, g(x))$ to $p(x, v)$. Reverse this path and concatenate with $\beta_x : p(x, g(x)) = x$. The endpoints are exactly those in (2). Path functoriality makes the construction natural in (s, v) . $\square$

No uniqueness of δ follows. For example, let $V = 1$, let $S = BG$ for a group with nontrivial center, and take $p = \text{id}_S$. The first homotopy is unique, while homotopies $\text{id}_{BG} \simeq \text{id}_{BG}$ can differ by central elements of G ; the formula retains this choice in β . Ordinary set semantics suppresses it. Nor does (2) prove the full PutPut law for two distinct views. A three-state countermodel is included in the computational supplement.

6. Topos semantics, substitution, and descent

An ∞ -topos is an accessible left exact localization of a presheaf ∞ -category. Its objects support dependent families, identity objects, and effective descent. Its 0-truncated objects form an ordinary Grothendieck topos. For an object U , the slice $\mathcal{X}_{/U}$ interprets a parameter context; its generalized points need not arise from global points of $\mathcal{X}$.

The category $\text{Mod}_T(\mathcal{X}) = \text{Fun}^\times(T, \mathcal{X})$ is a category of models *in* this topos. It is not generally itself a topos. For the theory of a pointed object and $\mathcal{X} = \mathcal{S}$, it is pointed spaces, a nontrivial category whose initial and terminal objects coincide. A nontrivial ∞ -topos cannot have this property because its initial object is strict. No topology is silently assigned to a term graph, and no arbitrary localization of diagram categories is asserted to be left exact.

Proposition 6.1 (Parameterized soundness). For $M : T \rightarrow \mathcal{X}_{/U}$, initial transformations preserve $\text{Sol}_M(D)$ in the slice. If $f : V \rightarrow U$, then

$$f^* \text{Sol}_M(D) \simeq \text{Sol}_{f^*M}(D). \quad (3)$$

Theorem 2.3 also applies to Kan extension with values in $\mathcal{X}_{/U}$.

Proof. The slice is an ∞ -topos. Apply Theorem 4.1 and Corollary 2.4 there. Pullback f^* has both a left adjoint Σ_f and a right adjoint Π_f , so it preserves limits and colimits; preservation of the defining limit proves (3). $\square$

For a more concrete constraint language, suppose $A \rightarrow U$ is a family of assignments and $r, s : A \rightrightarrows B$ are maps over U . The family of assignments with an equality witness is the homotopy pullback

$$\text{Eq}(r, s) = A \times_{B \times_U B} B,$$

using (r, s) and the diagonal of B . Its fiber is the identity type $r(a) = s(a)$. Further constraints between proofs are formed by iterating this construction. Finite systems use finite limits and hence are preserved by inverse-image functors of geometric morphisms. Arbitrary infinite limits need not be so preserved; this statement is distinct from the all-limit preservation of pullback between slices in (3).

Proposition 6.2 (Descent for witnesses). If $q : U' \rightarrow U$ is an effective epimorphism with Čech nerve $U'_\bullet$, and $Z \rightarrow U$ is any solution family, then

$$\text{Map}_{\mathcal{X}_{/U}}(U, Z) \simeq \lim_{[k] \in \Delta} \text{Map}_{\mathcal{X}_{/U}}(U'_k, Z). \quad (4)$$

Equivalences between solution families can be checked after pullback along q .

Proof. Effectiveness expresses U as the colimit of its Čech nerve in the slice. Mapping into Z changes this colimit to the limit (4). Descent identifies the slice over U with compatible families over that nerve, and therefore makes pullback along q conservative. $\square$

Equation (4) explains why locally chosen rewrite proofs need compatibility on overlaps, compatibility between those compatibilities, and higher descent data. A local proof merely existing everywhere does not automatically choose a global proof. The construction here internalizes values, constraints, and parameters. Comprehensive factorization for internal ∞ -categories is already available in [6, Corollary 7.23], with related module calculus in [9]. We do not additionally claim an internal product-restriction theorem or a syntax-completeness theorem for dependent e-graphs.

7. Exactness boundaries and a finite-limit criterion

Proposition 7.1 (Truncation is not left exact). In spaces, π_0 preserves products but need not preserve pullbacks. For the basepoint maps into BG ,

$$* \times_{BG} * \simeq G, \quad * \times_{\pi_0 BG} * \simeq *.$$

For nontrivial G the component sets differ.

Proof. The first pullback is the based loop space of BG , equivalent to the discrete group G . The connected space BG has one component. $\square$

Consequently one cannot recover a space of constrained solutions by computing an ordinary quotient first and then adding unspecified higher structure. Nor does Theorem 4.3 assert that $\tau_{\leq n} \lim H$ equals $\lim \tau_{\leq n} H$ for unrestricted diagrams.

Proposition 7.2 (A sufficient exactness condition). Suppose C, D have finite limits, $X : C \rightarrow \mathcal{S}$ preserves products, and $(\int_C X)^{\text{op}}$ is filtered. Then $\text{Lan}_F X$ preserves finite limits for every $F : C \rightarrow D$.

Proof. In (1), representables preserve finite limits, and filtered colimits of spaces commute with finite limits [3, Section 5.3.3]. $\square$

Corollary 7.3 (Finite-limit refinement). If C, D have finite limits and $F : C \rightarrow D$ preserves them, its comprehensive factorization consists of finite-limit-preserving functors between finitely complete ∞ -categories.

Proof. For $X = 1_C$, its elements category is C . The opposite C^{op} has finite colimits, so every finite diagram has a cocone and C^{op} is filtered. Proposition 7.2 makes Y_F left exact. For a finite diagram in $\int_D Y_F$, its underlying diagram in D has a limit d ; the compatible elements specify a point in $\lim Y_F(d_i) \simeq Y_F(d)$. The mapping-space fiber calculation of Lemma 2.1 proves that this is a limit upstairs and that the projection reflects limit cones. Consequently η_F preserves finite limits because $p_F \eta_F \simeq F$ does. $\square$

The same argument applies to left exact set-valued functors on ordinary finitely complete categories, using ordinary filtered colimits. This sharpens the finite-limit direction raised in [1]: the additional filteredness argument supplies the exactness missing from finite-product hypotheses alone. It does not extend left exactness to arbitrary product-preserving functors or supply a presentation calculus for dependent syntax. Flatness is the relevant general language when the domain lacks the requisite limits.

8. Verification, scope, and further work

The attached standard-library Python program independently checks the finite diamond’s cellular homology before and after adding a two-cell, the normalized bar-complex $H_1(B(\mathbb{Z}/2); \mathbb{F}_2)$, and all instances of the lens laws in the stated three-state countermodel. These are reproducible checks of examples, not machine verification of the general theorems. The proofs above are mathematical arguments using explicitly cited foundational results.

A practical higher e-graph should retain term sharing, proof generators, boundaries of coherence cells, and a declared observational level. Extraction should return both a selected term and its witness. Merging at level zero certifies congruence only; preserving a full solution space requires initiality or another independently proved semantic equivalence. The current paper does not supply a complete equality-saturation engine, a decidability theorem, a finite presentation of all higher coherence, or a complexity guarantee. A natural next implementation is a proof-producing finite fragment with independently checkable two-dimensional certificates.

Authorship and automation. Submitted under the anonymous byline at the responsible submitting researcher’s direction. An AI assistant assisted with literature retrieval, formulation, proofs, source preparation, and finite-example checks. No independent referee review or proof-assistant certification is claimed. Publication and review classification are editorial decisions.

References

1. E. Patterson. *How to prove equations using diagrams, Part 3*. Topos Institute, October 14, 2026. Article.
2. F. Borceux and B. J. Day. *On product-preserving Kan extensions*. Bulletin of the Australian Mathematical Society 17 (1977), 247–255. DOI.
3. J. Lurie. *Higher Topos Theory*. Princeton University Press, 2009; author’s updated text. Especially Sections 4.1, 5.3.3, 5.5.8, and Chapter 6. Full text.
4. J. Lurie. *Kerodon*, Section 7.2.1: Cofinal Morphisms of Simplicial Sets. Tag 02N0.
5. J. Lurie. *Kerodon*, Section 9.7.2: Sifted ∞ -Categories, especially the distinction from ordinary siftedness. Tag 02QD.
6. J. Ruit. *Formal category theory in ∞ -equipments II: Lax functors, monoidality and fibrations*. 2024, Section 7.4. arXiv:2408.15190.
7. R. Street and R. F. C. Walters. *The comprehensive factorization of a functor*. Bulletin of the American Mathematical Society 79 (1973), 936–941. DOI.
8. Y. Guiraud and P. Malbos. *Polygraphs of finite derivation type*. Mathematical Structures in Computer Science 28 (2018), 155–201. arXiv:1402.2587v2.
9. E. Riehl and D. Verity. *Kan extensions and the calculus of modules for ∞ -categories*. Algebraic & Geometric Topology 17 (2017), 189–271. arXiv:1507.01460.
10. A. Tiurin, C. Barrett, D. R. Ghica, and N. Hu. *Equivalence Hypergraphs: DPO Rewriting for Monoidal E-Graphs*. 2025 revision. arXiv:2406.15882v2.
11. A. Tiurin, D. R. Ghica, and N. Hu. *Categorical E-Graphs for Lambda Calculi*. June 2026 revision. arXiv:2505.00807v2.
12. The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. 2013, Chapters 2 and 6. Book and open full text.

References [10–11] give complementary categorical treatments of sharing, rewriting, and binding; they are not claimed to establish the proof-space semantics here. Reference [12] supplies the identity-type language used in (2). Links and bibliographic versions were checked during preparation.