

Learn the Law, Not the Atlas

Toward Compositional UV Learning for Generative Humans

Position paper, mathematical framework, and reproducible synthetic diagnostics

Working manuscript v0.1 | 8 October 2026

Abstract

We argue that the central learning object in UV-based generation should not be a packed image of a finished surface, but a family of conditional probability laws that remain coherent under changes of chart, partition, and resolution. This shift reframes data efficiency as the reuse of learned interactions rather than the multiplication of complete-asset examples. We call the resulting research target *probability-preserving UV composition*. Its essential components are informative separators, calibrated boundary free energies, and stochastic refinement that preserves specified coarse states. We formalize the distinction between geometric and statistical boundaries, give an exact decomposition of gluing error, show why conditional interior models cannot identify boundary probabilities, and state a constructive coarse-state preservation result. These are standard mathematical identities organized into a UV-learning design, not claims of new probability theory. Re-executed synthetic diagnostics expose the consequences: in a 19,683-state field, exact gluing reproduces the reference law to numerical precision, whereas minimum-energy boundary selection produces total-variation distance 0.2794 despite exact interior conditionals. A Gaussian refinement test loses 74.17% of total variance when conditional variability is omitted. A three-parameter, correctly specified local model transfers from one small grid to another after from-scratch estimation; this is a mechanism check, not evidence of human-level few-shot learning. We propose a jointly trained architecture and a falsifiable evaluation program targeting complete, editable humans. The central hypothesis is that the reusable learned description can be governed by interaction complexity rather than atlas size, while acknowledging the unavoidable cost of producing the output itself.

Keywords: UV representation; generative geometry; probabilistic composition; boundary free energy; multiscale learning; data efficiency; editable digital humans.

1 Introduction: changing what is learned

Consider one digital person represented by two different UV atlases. The charts may be cut differently, reordered, rotated, or packed into different locations. Those choices change the storage representation, not the person. A useful generative model should recognize that distinction. It should also preserve an established identity when the surface is refined, and preserve unedited properties when a local edit is made. These requirements concern the internal organization of a generative model, not merely the appearance of a rendered sample.

Surface parameterizations already provide a route from irregular geometry to structured learning domains. AtlasNet generates surfaces as collections of parametric elements, and multi-chart generative surface modeling uses coordinated parameterizations for shape generation [1, 2]. Object Images packages geometry, appearance, and patch structure into an image-like representation [3]. These works make it untenable to present “learning in UV space” alone as a new research principle.

Our position is more specific: *a UV region should represent how it can participate in a complete asset, not only what it looks like in isolation*. This requires learning the distribution of admissible interface states and the distribution of interiors conditional on those states. A model of isolated patches, even a perfect one under a supplied boundary, does not determine how often the different boundaries should occur. Nor does a geometric cut automatically remove dependence between its two sides.

This distinction matters for the original production objective: generating a complete human whose body, face, skin, hair, clothing, and performed expression remain coherent through editing, motion, relighting, and export. We do not reduce this objective to a collection of unrelated pretrained components. The proposed learned system is to be trained from scratch, with shared representations of interaction and explicit global conditions. Geometry and numerical algorithms remain legitimate structural knowledge; from-scratch learning does not mean discarding known mathematics.

The ambition is not to claim that local observations uniquely specify an unobserved person. Identity, topology, design, and performance contain genuinely global information. Rather, we ask whether reusable local laws can remove the need to observe every high-resolution realization of every global combination. This changes the proposed source of data efficiency: the model learns how a component responds to its context, then reuses that response law in contexts for which its assumptions remain valid.

The manuscript makes a position-level contribution, a formal diagnostic contribution, and a reproducibility contribution. It states a precise learning target; arranges established probability and linear-algebra facts into tests of that target; and reports verified finite-state experiments that expose failure mechanisms. The neural operator, adaptive separator learner, and full-human model remain proposals. This distinction is part of the scientific claim, not a reduction of the intended scope.

2 The position: a surface is not its storage layout

2.1 Coordinates are a computational choice, not an identity

UV coordinates are two-dimensional coordinates assigned to surface regions. A packing layout is a further choice of where those regions are stored. Neither choice should silently become a semantic variable. We therefore distinguish a physical or authored asset from any particular atlas used to represent it. An identity-preserving reparameterization must transform the associated tensor fields, material directions, units, and probability reference measures, rather than simply rotating an array.

This is not a demand to erase meaningful asymmetry. A scar, a garment closure, or an asymmetric hairstyle is part of the asset. Moving its coordinates without changing the asset is different from moving the feature on the body. The model must preserve this distinction at the level of correspondences and distributions.

2.2 Composition, not patch multiplication, is the source of reuse

Cutting one observed asset into ten thousand patches does not produce ten thousand independent people. It produces correlated views of the same underlying observation. Their value is that they expose repeated relations: how fields continue across a cut, how a local material responds to a surrounding state, and which information is needed to reconcile neighboring regions.

Our proposed statistical unit is therefore an *interaction law conditional on context*. A region is useful when its learned law can be reused without relearning its entire interior for each possible whole. The whole remains indispensable for learning global variables and for verifying that local explanations coexist. Locality is an assumption to test, not a synonym for using small patches.

2.3 Validity and editability belong inside the representation

A production-oriented model should distinguish a plausible rendering from a usable asset. Surface connectivity, intended boundaries, material coordinates, attachment relationships, and edit locks must survive changes in resolution and state. This is a design requirement; this paper does not establish that any UV representation automatically guarantees a clean mesh.

In particular, an artificial chart cut and a physical sewing seam are different relations. The former duplicates the description of one surface. The latter connects distinct material pieces and can support

folds, orientation changes, and discontinuities of material. A universal seam-smoothing loss would confuse those meanings. Similarly, a scalp attachment indexes a hair curve; it does not make that curve a two-dimensional skin surface.

2.4 Sparse-data learning is structured inference, not missing information by decree

The goal is to learn more reusable knowledge from each independent observation. It is not to infer arbitrary hidden details with certainty. If several material or identity hypotheses explain the observations, a model should retain that ambiguity or obtain informative new observations. Self-consistency across UV layouts adds constraints on a model’s use of existing evidence; it does not manufacture new empirical evidence about people.

The operational philosophy is consequently: *preserve the meaning, account for the uncertainty, and make reuse conditional on a testable factorization*. The full-human goal and the ambition of a major improvement remain intact, but the improvement must be measured rather than asserted.

3 Relationship to existing research

Atlas and pattern representations. Multi-chart methods establish that several coordinated parameterizations can support shape generation [2]. GarmageNet connects panel geometry images, sewing relationships, and garment initialization [4]; Garment Particles couples two-dimensional pattern coordinates with three-dimensional geometry [5]. These are direct antecedents, not weaker substitutes for the proposal. The distinction sought here is consistency of the *probability law* under alternative decompositions and refinements of the same asset. The reported limitations of Garment Particles concerning particle resolution and repeated-edit detail preservation motivate this requirement but do not prove that our approach solves it.

Intrinsic geometry and coordinate handling. Neural Jacobian Fields predicts differential information and recovers a mapping through a differentiable Poisson solve [6]. Statistical shape modeling with discrete fundamental forms uses metric and curvature information [7]. Gauge-equivariant mesh networks handle orientation-dependent features across local frames [8]. Accordingly, neither differential fields nor coordinate equivariance alone constitutes the proposed contribution. They are candidate tools for making a surface law geometrically meaningful.

Graphical models and multiscale probability. Marginalization, separator-based factorization, and sum-product computation are classical [9]. Wavelet-conditional renormalization learns distributions of fine-scale variables conditioned on coarse fields and estimates energies across scales [10]. Our proposal does not rename those ideas as new mathematics. It asks how to learn and compress them across variable UV partitions, geometry-dependent interfaces, and changing interaction graphs while tracking approximation error.

Function-space generation and domain decomposition. Function-space diffusion gives a rigorous setting for distributions over functions [11]. Learning-based domain decomposition reuses an operator trained on simple domains in more complex PDE domains [12]. Diffusion Last Layer learns stochastic coefficient distributions for neural operators [13]. None of these facts makes a fixed-size representation of every UV boundary law available by default. A compact, cross-partition-consistent construction remains the research target here.

4 Formal setting: laws over equivalent UV descriptions

Let $\mathcal{A}$ denote an atlas, h a finite discretization, τ a typed connectivity and attachment structure, and z a global condition. The condition can include identity, dimensions, design intent, or a specified performance

state. Let Φ contain geometry, appearance, and other material-coordinate fields. We write

$$\mu_{\mathcal{A},h}(d\Phi \mid z, \tau) \quad (1)$$

for a distribution on explicitly specified valid states. The global structure can itself be learned through $p_\theta(z, \tau)$; the local construction does not determine it from nothing.

A compiler $C_{\mathcal{A},h}$ maps a representation into an asset or a common comparison space. Two atlases of the same underlying asset class should satisfy the target relation

$$(C_{\mathcal{A},h})_{\#}\mu_{\mathcal{A},h} \approx (C_{\mathcal{A}',h'})_{\#}\mu_{\mathcal{A}',h'}. \quad (2)$$

The symbol $\#$ denotes push-forward of a probability measure. Finite discretizations generally introduce approximation; equality is not claimed merely because both outputs are called UV maps. For a pure invertible relabeling on a fixed grid, exact equality is a reasonable algebraic requirement.

A candidate surface state could use an embedding and its derivatives, or a first fundamental form a , a second fundamental form b , and material fields m . These options impose different reconstruction and compatibility conditions. Positive metric, integrability, global closure, and absence of self-intersection are separate requirements. Local compatibility does not establish a globally embedded, collision-free surface. We therefore leave the geometry parameterization open while requiring that its reference measure and valid support be explicit.

4.1 The separator is an information object

For a cut into interiors L and R , let B contain the shared state. The familiar factorization

$$P(L, R \mid B, z) = P(L \mid B, z)P(R \mid B, z) \quad (3)$$

is a statement of conditional independence. It need not follow from adjacency in the UV image. A bending interaction can require a boundary collar rather than a single row of values. A contact can connect regions far apart in UV. Global variables may couple regions that never touch.

A separator can therefore include a collar, suitable derivative degrees of freedom, selected nonlocal interaction endpoints, and explicit global information. It must not independently prescribe incompatible displacement and traction data. Which variables are sufficient depends on the field law and its discretization. A useful separator is sufficiently informative without encoding the complete interior.

4.2 Distributional consistency is not pathwise identity

Two sampling procedures can produce the same distribution while producing different samples for a shared random seed. Persistent asset editing requires more: stable correspondences, preserved realized variables, and coupled randomness for untouched details. The target in equation (2) is distributional. Pathwise preservation of specified fields is imposed separately through the refinement construction in section 6.

5 The mathematics of probability-preserving gluing

5.1 An exact error ledger

All identities below concern finite-dimensional distributions with suitable absolute continuity and finite KL divergence. Global conditions can be fixed and then averaged over. The identities are standard consequences of factorization and marginalization; their role is to organize a learnable UV system.

Proposition 1 (Gluing error decomposition). *For a true law $P(L, R, B)$ and a normalized approximation*

$$\widehat{P}(L, R, B) = \widehat{P}_B(B) \widehat{P}_L(L | B) \widehat{P}_R(R | B), \quad (4)$$

the following identity holds:

$$\begin{aligned} D_{\text{KL}}(P \| \widehat{P}) &= I_P(L; R | B) + D_{\text{KL}}(P_B \| \widehat{P}_B) \\ &\quad + \mathbb{E}_{P_B} D_{\text{KL}}(P_{L|B} \| \widehat{P}_{L|B}) \\ &\quad + \mathbb{E}_{P_B} D_{\text{KL}}(P_{R|B} \| \widehat{P}_{R|B}). \end{aligned} \quad (5)$$

Proof. Insert $P_B P_{L|B} P_{R|B}$ into the density ratio $P/\widehat{P}$, take logarithms, and integrate under P . The first term is conditional mutual information; each remaining term integrates to the stated marginal or conditional KL divergence. $\square$

The first term is a structural error floor for this conditional-independence family. Increasing the capacity of the other three factors cannot remove it. The second term diagnoses how often boundary states are selected, while the last two measure the accuracy of filling each interior. Thus a large network is not a general remedy for a badly chosen separator.

Applied recursively, this accounting requires expectations under the true distribution of ancestor states. A low average error per node does not imply a small global error when the number of nodes grows. Nor does the identity make high-dimensional conditional mutual information easy to estimate. Practical dependence certificates remain part of the proposed method.

5.2 Why boundary matching and local consistency are insufficient

A shared geometric boundary does not specify its probability. Even compatible lower-order marginals need not define a joint law on a cyclic graph. For a simple counterexample, let three binary variables have pairwise laws that demand disagreement with probability one. Every singleton marginal can be uniform, so the pairwise laws agree on their overlaps. Yet three binary variables cannot all disagree pairwise. A global joint law with those pairwise specifications does not exist.

Our gluing operation therefore refers to factors from a coherent energy decomposition or to a calibrated joint hierarchy, not an arbitrary collection of normalized patch marginals. This distinction also prevents double-counting shared boundary terms. Exact inference already makes these distinctions in graphical models [9]; learned UV composition must preserve them rather than ignore them behind seam blending.

5.3 Boundary free energy preserves the mass of possible interiors

For a region Ω with interior x and boundary q , define

$$W_\Omega(q | z) = \int \exp[-E_\Omega(x, q | z)] d\nu_{\Omega, q}(x), \quad (6)$$

$$F_\Omega(q | z) = -\log W_\Omega(q | z). \quad (7)$$

The integral is over admissible states with an explicit reference measure. Hard equality constraints require coordinates or a measure on their feasible set; they cannot be imposed by integrating over a measure-zero subset of ordinary ambient volume without further care. We make no continuum path-integral claim.

E can be a statistical energy, not the thermodynamic energy of a human at equilibrium. The distinction between $F_\Omega(q)$ and $\min_x E_\Omega(x, q)$ is nevertheless essential. For

$$E(x, q) = \frac{1}{2}k(q)[x - m(q)]^2 + V(q), \quad k(q) > 0, \quad (8)$$

Gaussian integration gives

$$F(q) = V(q) + \frac{1}{2} \log k(q) - \frac{1}{2} \log(2\pi). \quad (9)$$

At equal V , stiffnesses 1 and 100 give relative boundary weights 10 to 1. Their minimum energies are identical. Selecting boundaries by minimum energy discards the probability mass carried by interior variability.

Proposition 2 (Conditional interiors do not identify boundary weights). *Knowledge of $p(x | q)$ for every q does not determine the boundary law. Adding a finite function $g(q)$ to $E(x, q)$ leaves all interior conditionals unchanged and replaces $F(q)$ by $F(q) + g(q)$, whenever the relevant normalizers exist.*

Proof. For fixed q , the factor $e^{-g(q)}$ appears in both the numerator and the normalizer of the conditional density and cancels. It does not cancel between different boundary states. $\square$

Consequently, a boundary-conditioned flow trained only on normalized interior likelihoods is insufficient. Boundary or joint observations, a shared calibrated energy model, or trustworthy free-energy targets must supply the missing information. Cross-layout consistency alone cannot identify the true data distribution.

5.4 The gluing operator and its conditions

Suppose all cross-region dependence is represented by a shared variable s , and the factor energy has been split without duplicating terms. Then

$$W_{A \cup B}(q_{\text{ext}}) = \int W_A(q_A, s) W_B(s, q_B) d\nu(s). \quad (10)$$

For a physical seam or interaction with distinct interface variables, an explicit compatibility factor may instead be needed. The simple equality form in equation (10) corresponds to the specified shared-state factorization. Nonnegative integrable factors admit order-independent exact elimination; a neural approximation is not automatically order-independent.

The research problem is to approximate these messages while controlling their growth in complexity. Eliminating a nonlinear interior can induce long-range, high-order interactions on its boundary. A fixed-width embedding is not a universal sufficient statistic. Exact discrete inference can become expensive as separator size increases, as its graphical-model counterpart already makes clear [9].

Proposition 3 (A local calibration bound). *Let $p(q) \propto e^{-F(q)}$ and $\hat{p}(q) \propto e^{-\hat{F}(q)}$ have the same support and reference measure, with finite normalizers. If, after an irrelevant additive constant is removed, $|\hat{F}(q) - F(q)| \leq \epsilon$ on their support, then*

$$D_{\text{KL}}(p \| \hat{p}) \leq 2\epsilon. \quad (11)$$

Proof. Writing $\Delta = \hat{F} - F$, the ratio of normalizers lies in $[e^{-\epsilon}, e^{\epsilon}]$. Hence $\log(p/\hat{p}) = \Delta + \log(\hat{Z}/Z) \leq 2\epsilon$ pointwise. Integrating proves the bound. $\square$

This is a limited finite-domain bound, not a sample-complexity theorem. If approximate log-messages are added, their worst-case errors can accumulate. Resolution-independent global accuracy does not follow from a uniform local error held fixed while the number of factors increases.

6 Projective refinement and persistent editing

Let R restrict a fine state to a coarse state. A projectively consistent family aims to satisfy

$$R_{\#} \mu_f = \mu_c. \quad (12)$$

This concerns a distribution, not only agreement of mean surfaces. Deterministic interpolation can preserve coarse values while eliminating most of the conditional variability.

Proposition 4 (Exact coarse-state preservation). *Assume a fixed linear restriction $R : \mathbb{R}^n \rightarrow \mathbb{R}^m$ has full row rank. Choose H and N with $RH = I$ and $RN = 0$. Then*

$$x_f = Hc + NT_\theta(\xi; c, B, z) \quad (13)$$

satisfies $Rx_f = c$ for every parameter value, input noise ξ , and conditioning state.

Proof. Multiplying equation (13) by R gives $RHc + RNT_\theta = c$. □

When the columns of N span $\ker R$, the construction can express all states on the affine fiber. It does not guarantee that the learned distribution on that fiber is correct. Sampling $c \sim \mu_c$ and then a normalized conditional supported on that fiber does, however, establish equation (12) for the chosen hierarchy. Other partition trees require additional consistency.

Edit locks can be included as further independent linear constraints, provided they are compatible. A skin-material edit should not change locked geometry. A body-shape edit with garment refitting intentionally releases the affected fit and contact variables. Nonlinear constraints and topology changes require different local coordinates or explicit structural updates; the proposition is not a universal guarantee for arbitrary remeshing.

Persistent details also need persistent realization. Reusing the same integer seed across unrelated parameterizations is insufficient. We propose storing stable material-coordinate correspondences and the realized residual variables for unedited regions, then conditionally regenerating only the released degrees of freedom. A new high-frequency realization may be plausible without being the same identity-specific detail.

Finally, partitioning, refinement, and coordinate conversion should commute at the level of their compiled laws. A model trained on a single refinement tree cannot claim this property for alternative cuts merely because it samples successfully. Cross-tree and cross-atlas comparisons are explicit obligations of the proposed learning program.

7 A proposed model: jointly learning sufficiency, mass, and detail

The architecture described here has not been trained on human assets. We propose one shared, from-scratch geometric backbone with typed field inputs and three coupled functions. No existing body, face, garment, or language model is assumed as a frozen source of pretrained knowledge.

7.1 An adaptive separator mechanism

For a candidate partition, the model estimates whether its shared state sufficiently explains cross-region dependence, and how much the shared representation costs. A conceptual objective is

$$\min_{\pi, B} I_P(L; R \mid B, z, \pi) + \lambda \text{Cost}(B, \pi), \quad (14)$$

where π denotes the partition. Actual optimization requires estimators or bounds and should validate them on held-out data. Their numerical values cannot be treated as certified mutual information without such validation.

The candidate state includes intrinsic geometry, reference material directions, and typed contact or attachment links. It can expand a collar or expose a nonlocal interaction when an inadequate factorization is detected. A global latent must not become a hidden encoding of the complete asset; its information and computational cost are part of the accounting.

Adaptive partitions need special care. Selecting a partition using hidden interior information and then ignoring that selection event can bias the resulting law. A generative system must fix the partition from permitted conditions or include the partition-selection distribution in its joint model. Partition choice is not statistically free preprocessing.

7.2 A calibrated boundary-message model

The second function approximates unnormalized interface messages or their free energies. To avoid treating freely chosen local functions as a coherent global energy, training must connect them through joint observations or a shared normalized hierarchy. On small calibration domains, exact enumeration or validated integration can supply targets. Larger domains require an explicitly evaluated approximation strategy.

For one partition, a normalized conditional $k_\theta(x | q)$ and a free energy $F_\theta(q)$ define

$$E_\theta(x, q) = F_\theta(q) - \log k_\theta(x | q). \quad (15)$$

Integrating e^{-E_θ} over the interior recovers e^{-F_θ} . This construction resolves normalization for that partition, not agreement with all other partitions. Enforcing cross-partition coherence without intractable contraction is a central unresolved algorithmic question.

7.3 A constraint-preserving detail model

The third function learns a conditional transport on the allowed residual space in equation (13). Geometry, albedo, roughness, and deformation detail need not share identical output distributions. They share the requirement that prescribed states remain fixed and uncertainty is retained where information is absent.

The training objective combines observed-data likelihood, calibrated boundary predictions, conditional detail likelihood or a validated surrogate, cross-cut consistency, cross-resolution consistency, and representation cost. A consistent but wrong distribution can satisfy all representation-only losses. Observational fit is therefore indispensable. The exact KL ledger provides a diagnostic organization of these losses, not an immediately available high-dimensional training oracle.

7.4 One human, heterogeneous fields

The full-human application retains a global anatomy and performance state while placing skin and clothing fields on their appropriate material domains. Hair curves are indexed by scalp coordinates and curve parameters; their nonlocal interactions are explicit. Body and face topology, sewing graphs, and attachment structures are distinct typed variables. We do not claim that every component is literally a two-dimensional sheet.

Skin tone belongs to a relightable material description rather than the color of one photograph. Illumination and camera response are separate nuisance variables, with uncertainty when observations cannot distinguish them. Expression and emotion controls denote an intended performance, not a claim to read private mental states from appearance. These are representation commitments for the complete system, not capabilities demonstrated by the scalar experiments.

8 The sparse-data hypothesis and its boundaries

Our central hypothesis is *controlled compositional closure*: over a specified class of human-asset laws, there exists a family of separator messages and conditional refinements that remains approximable after gluing, recutting, and coarsening, with a measurable complexity budget. If that family can be learned from fewer independent assets than a whole-atlas model needs at the same quality, the intended gain is real.

There are three possible sources of improvement. Coordinate handling can remove nuisance variation. Parameter sharing can reuse a conditional interaction law across many admissible contexts. Alternative cuts and resolutions can expose contradictions in the model’s explanation of the same observation. None of these mechanisms creates independent evidence about an unobserved identity or an unmeasured material.

The hypothesis requires shared mechanisms, adequately observed boundary contexts, identifiable global conditions, and controlled approximation error. It can fail when identity-specific detail has little reusable structure, when contacts create dense dependencies, or when learned boundary laws are asked to extrapolate far outside their observed support. Increasing the number of patches can also increase the number of dependent observations without increasing effective information.

A major improvement must be reported against independent assets, identities, design families, and capture sessions, not crop counts. A hundred-fold data-efficiency goal means reaching a predefined complete-asset quality target with one hundredth as many independent real training units as a strong comparator. It is an experimental target, not a conclusion of this paper. Simulation, language supervision, compute, and manual curation remain separately counted resources.

The intended complexity claim also excludes a basic impossibility: outputting N explicit samples requires at least output-proportional work. We seek to reduce learned description complexity, independent supervision, and unnecessary global reasoning. We do not claim to emit a high-resolution asset at zero cost as its size grows.

9 Verified synthetic diagnostics

The supplied diagnostic implementation was inspected and re-executed for this manuscript with seed 20261008 and eight learning replicates per sample size. Tables below use the verification run. Numerical results differ from the earlier saved run only at floating-point scale; the largest absolute difference among corresponding numeric entries is below 6×10^{-14} . The experiments are scalar-field tests, not geometric-human benchmarks.

9.1 Gaussian fields: sufficient separators and stochastic detail

On a 17×17 four-neighbor grid, let L_g be the graph Laplacian and

$$Q = 0.3I + 0.9L_g + 0.15L_g^2 \quad (16)$$

be the precision of a zero-mean Gaussian field. The positive mass removes the constant null mode. We compare 289 fine, 81 intermediate, and 25 coarse degrees of freedom. Marginalization is computed with Schur complements.

Diagnostic	Verification result
Direct versus nested marginal precision, relative error	2.94×10^{-16}
Marginal precision after vertex reindexing, relative error	3.36×10^{-16}
Full stochastic reconstruction of covariance, relative error	6.83×10^{-16}
Conditional-mean-only covariance, relative error	0.651434
Total variance omitted by conditional-mean-only reconstruction	74.1700%
Maximum change to locked coarse values in the sampled refinement	0

Table 1: Gaussian consistency diagnostics. Relative errors use the Frobenius norm. The missing-variance percentage belongs to this scalar model, not to human surface detail. Reindexing is not a test of arbitrary smooth UV reparameterization.

The L_g^2 interaction spans farther than a nearest-neighbor term. A one-column separator leaves conditional mutual information 0.0142245 nats. Two columns remove it to numerical precision. Adding a spring-like nonlocal factor between two distant sites raises it to 0.0761028 nats; placing one of those endpoints in the separator removes it again. Values around 10^{-13} below zero are log-determinant roundoff, not negative mutual information. The extra spring illustrates omitted dependence; it is not a frictional-contact simulator.

The refinement test also distinguishes deterministic correctness from distributional correctness. The conditional mean preserves the coarse state but omits covariance on the unresolved degrees of freedom. Restoring the conditional residual recovers the full covariance in table 1. A mean-only reconstruction lives on a lower-dimensional set, so we do not replace this covariance diagnostic by an unjustified finite joint-density KL calculation.

9.2 A non-Gaussian field: entropy changes the assembled law

For $x_i \in \{-1, 0, 1\}$ on a 3×3 grid, define

$$E(x) = \frac{0.65}{2} \sum_{(i,j)} (x_i - x_j)^2 + 0.85 \sum_i (x_i^2 - 1)^2 - 0.12 \sum_i x_i. \quad (17)$$

All $3^9 = 19,683$ states are enumerated. A middle column separates the two side columns. Its onsite and within-column interaction terms are divided evenly between the two factors, avoiding duplicate energy.

Boundary selection, with exact interior conditionals	TV	KL (nats)
Exact free-energy gluing	$< 10^{-14}$	$< 10^{-14}$
Minimum-energy boundary weights	0.279386	0.319537
Uniform boundary weights	0.405853	0.611330

Table 2: Exact enumeration separates boundary calibration from interior modeling. TV is half the summed absolute probability difference, not a percentage of geometric accuracy.

Adding $\frac{1.3}{2}(x_{\text{left}} - x_{\text{right}})^2$ across the geometric separator breaks the presumed conditional independence. For the deliberately misspecified glued model, the structural, boundary, left-interior, and right-interior terms in equation (5) are respectively 0.0874431, 0.2359464, 0.0260462, and 0.0260462 nats. Their sum equals the total KL, 0.3754818, to below 10^{-14} . This verifies the error accounting on a known law; it is not an estimator of those terms from scarce human data.

9.3 From-scratch estimation under a supplied local family

We replace the three coefficients in equation (17) by unknown parameters. Independent complete fields from the 3×3 law supply training observations. Parameters are fitted by exact likelihood with a small quadratic regularizer, then reused on a held-out 2×5 domain with the same underlying law. Eight independently seeded repetitions are used for each training-set size.

Independent training fields	Median test KL/site	Minimum–maximum
8	0.0165594	0.0106819–0.1235282
32	0.0060978	0.0001063–0.0205688
128	0.0015739	0.0004257–0.0032776

Table 3: Domain transfer for a correctly specified three-parameter field model. Values are nats per site. The ranges are observed extrema over eight repetitions, not confidence intervals.

All 24 parameter optimizations report success. No pretrained parameters are used, but the correct functional family is supplied by hand. The held-out domain has a different shape, not a new physical law or semantic category. The diagnostic establishes a narrow reuse mechanism. It does not establish a neural discovery of that mechanism or a favorable comparison with contemporary generative models. A deliberately misspecified independent-site baseline is retained in the raw record for inspection and is not used to claim state-of-the-art performance.

10 A falsifiable full-human research program

The next empirical claim must be tested on complete assets rather than inferred from the preceding diagnostics. The initial integrated representation should include body and face surfaces, skin appearance, hair attachment and curve state, garment pattern and material state, and explicit performance conditions. Mechanism tests are small because exact answers are available, not because the final objective is being narrowed.

Repartitioning and coordinate tests. Withhold chart counts, cut placements, and valid coordinate transformations. Compare generated laws after compiling them into a common geometry-and-material space. Separately test pathwise preservation for edits with known correspondences. A model should not receive credit for atlas invariance merely because it ignores local detail.

Compositional generalization. Withhold whole combinations while retaining selected constituent interactions. Also withhold genuinely new local mechanisms to distinguish compositional reuse from unsupported extrapolation. Splits must be grouped by identity, garment design family, hair source, and capture session. Global context leakage is explicitly audited.

Refinement and editing. Measure coarse-state error, conditional detail statistics, and unedited-region drift across repeated edits. Check alternate orders of partitioning, refinement, and compilation. Independent edits may be required to commute on the locked specification; coupled edits and history-dependent physical motion must not be assigned that requirement indiscriminately.

Asset-level acceptance. The evaluation includes intended topology and boundaries, deformation quality, contact failures, relighting consistency, and cross-tool reproducibility. Skin-tone and material-condition strata should be reported separately so aggregate scores do not hide systematic failures. Visual fidelity is necessary but not the only acceptance condition.

Fair resource accounting. Compare with multi-chart generators, fixed-tree multiscale conditionals, function-space models, and strong whole-asset generators. Give comparators access to the same real data, simulation, geometric constraints, and optimization budget. Report memory, sample latency, inference iterations, rejection or repair costs, and human cleanup time independently of sample quality.

A large data-efficiency result requires learning curves at fixed acceptance criteria, not a single small-data example. A large productivity result requires complete end-to-end human labor measurements, not just the first sampling call. Both goals can be ambitious without being conflated.

11 Limitations, risks, and criteria for rejecting the hypothesis

The largest unresolved issue is closure: elimination can make boundary messages as complex as the original interior. Separator size alone is not a sufficient measure of representational difficulty. If compact messages are unavailable at production fidelity, the proposed advantage may disappear. No universal compression or sample-complexity theorem is offered.

The geometry problem is also unresolved. The scalar experiments do not include nonlinear surface integrability, garment construction, skin optics, hair dynamics, or changing topology. Their exact finite-state computations are exponential in the number of sites and cannot be naively scaled. Machine-precision identities are software checks, not perceptual validation.

Data pose additional limits. Sparse observations may not identify material parameters or global semantics. Synthetic observations generated by an incorrect model can reinforce the error. Coordinate and cross-cut consistency permit self-consistent but empirically wrong solutions. A skin-tone edit should not silently alter unrelated identity or behavioral properties; real-person reconstruction requires appropriate consent and provenance.

The key hypothesis should be rejected or restricted if efficient separators effectively encode whole interiors; if each resolution needs a new whole-distribution model; if alternate decompositions cause uncontrolled

drift; or if any apparent data advantage vanishes against equally informed, equally resourced baselines. These outcomes would reveal where interaction complexity prevents the desired reuse.

12 Conclusion

The object of learning proposed here is neither a packed UV image nor a bag of patches. It is a coherent family of conditional surface laws together with the information and probability mass required to compose them. Informative separators determine what must be shared. Boundary free energies determine how likely those shared states are. Conditional stochastic refinement restores what remains uncertain without changing what has already been fixed.

The mathematics explains why these requirements are distinct, and the synthetic experiments verify concrete consequences of violating them. The central engineering and scientific advance remains to be achieved: a compact learned operator that preserves these relations across the geometry, material, and interaction complexity of a complete human.

Our governing question is therefore not simply how to generate a larger or cleaner atlas. It is whether a model can learn a reusable law of surfaces whose meaning survives cutting, gluing, refinement, and editing. A demonstrated positive answer would change what must be learned from each example, rather than merely increase the number of examples consumed.

Evidence, authorship, and reproducibility statement

This is a working position paper with a proposed method, not a completed neural-system benchmark report. The analytic identities are established mathematics and are not claimed as original theorems. Numerical findings refer to the accompanying scalar-field code, re-executed for this draft. Authorship and affiliations are intentionally left unspecified for the research team to establish. Manuscript drafting and computational verification were AI-assisted; no external human study, independent replication team, or institutional endorsement is implied. Sources were checked through 8 October 2026. No priority claim or submission status is asserted.

A Additional derivations and technical qualifications

A.1 Gaussian elimination, normalization, and conditional detail

Partition a positive-definite Gaussian precision into interior and retained variables:

$$Q = \begin{bmatrix} Q_{xx} & Q_{xq} \\ Q_{qx} & Q_{qq} \end{bmatrix}, \quad E(x, q) = \frac{1}{2} [x^\top, q^\top] Q [x^\top, q^\top]^\top. \quad (18)$$

Completing the square yields

$$S = Q_{qq} - Q_{qx} Q_{xx}^{-1} Q_{xq}, \quad (19)$$

$$F(q) = \frac{1}{2} q^\top S q + \frac{1}{2} \log \det Q_{xx} - \frac{n_x}{2} \log(2\pi), \quad (20)$$

$$x \mid q \sim \mathcal{N}(-Q_{xx}^{-1} Q_{xq} q, Q_{xx}^{-1}). \quad (21)$$

The normalizing term is constant in q for fixed Q , but it must not be discarded when it varies with a condition or with a compared structure. The Schur complement is the retained precision. Nested elimination and direct elimination agree when the same variables and measure are retained.

Let $A = -Q_{xx}^{-1} Q_{xq}$ and $\Sigma_q = S^{-1}$. The law of total covariance gives

$$\text{Cov}(x) = A \Sigma_q A^\top + Q_{xx}^{-1}. \quad (22)$$

Mean-only refinement retains the first term and removes the second. In the reported whole-field calculation, the omitted fraction is the trace of the conditional residual divided by the trace of the full covariance.

A.2 The KL decomposition as a best-factorization statement

For any fixed separator with the stated support assumptions, all three model-fitting terms in equation (5) can be zero when their families contain the true marginal and conditionals. Therefore

$$\inf_{\hat{P}_B, \hat{P}_L, \hat{P}_R} \text{D}_{\text{KL}} \left(P \parallel \hat{P}_B \hat{P}_L \hat{P}_R \right) = I_P(L; R \mid B). \quad (23)$$

This identifies the structural floor for that family. It does not say the floor is small, inferable from a few samples, or removable without increasing representation cost. If hard support constraints are lost, the KL may be infinite; a small visual seam error does not preclude that failure.

For a hierarchical factorization, apply the same identity conditionally at each node and use the KL chain rule. Ancestor-state expectations and the number of factors remain explicit. An unweighted mean of local training losses is not a proof of a global accuracy bound.

A.3 Reference measures and coordinate changes

For an invertible finite-dimensional coordinate change $y = T(x)$, densities with respect to Lebesgue coordinates transform as

$$p_y(y) = p_x(T^{-1}y) \left| \det DT^{-1}(y) \right|. \quad (24)$$

Equivalently, one can transform the reference measure itself. A UV reparameterization also transforms tensor components and discrete integration weights. Ignoring these transformations confounds representation changes with changes in probability law. A vertex permutation has unit absolute determinant and is a much weaker test than a general chart transformation.

For nonlinear feasible sets, use an appropriate coordinate chart or disintegrated measure on each constraint fiber. The affine null-space construction is deliberately narrower and avoids claiming a global parameterization of arbitrary valid surfaces.

A.4 Approximate closure: the intended theorem target

For a specified class of regions and conditions, suppose a message family $\mathcal{F}_K$ of complexity K satisfies

$$\inf_{f \in \mathcal{F}_K, a \in \mathbb{R}} \sup_{q \in Q} |F_\Omega(q) - f(q) - a| \leq \epsilon. \quad (25)$$

A useful result would bound the required K , its increase under composition, and the total distributional error over the intended hierarchy. The set Q , the treatment of tails, the geometry class, and the cost of evaluating or sampling a message all matter. This display is a research specification, not an approximation theorem proved in this manuscript.

A theorem about log-message approximation would still not imply a finite-sample bound without assumptions on the data-generating family, identifiability, estimation, and optimization. It would also not remove the cost of writing a high-resolution output.

B Reproduction protocol

The accompanying package contains the inspected `experiments.py`, the original results, the new verification results, and the execution log. The re-execution used Python 3.13.5, NumPy 2.3.5, and SciPy 1.17.0 on Linux, with BLAS-related thread limits set to one. Numerical differences in the final digits can depend on the linear-algebra runtime.

```
python -m pip install -r requirements.txt
export OPENBLAS_NUM_THREADS=1
export MKL_NUM_THREADS=1
export OMP_NUM_THREADS=1
python experiments.py --seed 20261008 --replicates 8 \
  --out results_local.json
```

The three-parameter fit minimizes the empirical mean energy plus the exact log normalizer and $0.0005 \|\theta\|_2^2$. It uses L-BFGS-B, initialization $(0.4, 0.4, 0)$, bounds $[0.01, 3]$ on the first two coefficients and $[-1.5, 1.5]$ on the third, and at most 200 iterations. Ground truth is $(0.65, 0.85, 0.12)$. The test grid has $3^{10} = 59,049$ states, permitting exact probability comparison. Replicate r at sample size n uses the seed $20261008 + 10000n + r$.

The implementation asserts the Schur, covariance, reindexing, lock, conditional-independence, exact-gluing, and KL-decomposition checks. The independent-site fit in the raw file is intentionally unable to represent pair interactions. Its internal code comment should not be read as describing a state-of-the-art or correctly matched generative baseline.

Claim class	Status in this manuscript
Probability identities	Derived under explicit assumptions; established mathematics.
Finite scalar diagnostics	Executed again; raw outputs and code included.
Local parametric estimation	Executed from scratch in a supplied three-parameter family.
Neural UV composition	Proposed; no trained human generator or neural closure result.
Full-human production quality	Evaluation target; not demonstrated by the scalar tests.
Hundred-fold improvement	Falsifiable ambition; no measured result claimed.

Table 4: Claim ledger: what the evidence supports and what remains to be established.

References

- [1] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. AtlasNet: A papier-mâché approach to learning 3D surface generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. URL <https://arxiv.org/abs/1802.05384>.
- [2] Heli Ben-Hamu, Haggai Maron, Itay Kezurer, Gal Avineri, and Yaron Lipman. Multi-chart generative surface modeling. *ACM Transactions on Graphics*, 37(6):215:1–215:15, 2018. doi: 10.1145/3272127.3275052. URL <https://arxiv.org/abs/1806.02143>.
- [3] Xingguang Yan, Han-Hung Lee, Ziyu Wan, and Angel X. Chang. An object is worth 64x64 pixels: Generating 3D object via image diffusion. *arXiv preprint arXiv:2408.03178*, 2024. URL <https://arxiv.org/abs/2408.03178>.
- [4] Siran Li, Ruiyang Liu, Chen Liu, Zhendong Wang, Gaofeng He, Yong-Lu Li, Xiaogang Jin, and Huamin Wang. GarmageNet: A multimodal generative framework for sewing pattern design and generic garment modeling. *arXiv preprint arXiv:2504.01483*, 2025. URL <https://arxiv.org/abs/2504.01483>.
- [5] Kiyohiro Nakayama, I-Chao Shen, Ruofan Liu, Yiming Wang, Gordon Wetzstein, and Takeo Igarashi. Garment particles: A 2D–3D symmetric garment representation for generation and editing. *arXiv preprint arXiv:2605.26391*, 2026. URL <https://arxiv.org/abs/2605.26391v2>.
- [6] Noam Aigerman, Kunal Gupta, Vladimir G. Kim, Siddhartha Chaudhuri, Jun Saito, and Thibault Groueix. Neural Jacobian fields: Learning intrinsic mappings of arbitrary meshes. *arXiv preprint arXiv:2205.02904*, 2022. URL <https://arxiv.org/abs/2205.02904>.
- [7] Felix Ambellan, Stefan Zachow, and Christoph von Tycowicz. Rigid motion invariant statistical shape modeling based on discrete fundamental forms. *Medical Image Analysis*, 73:102178, 2021. doi: 10.1016/j.media.2021.102178. URL <https://arxiv.org/abs/2111.06850>.
- [8] Pim de Haan, Maurice Weiler, Taco Cohen, and Max Welling. Gauge equivariant mesh CNNs: Anisotropic convolutions on geometric graphs. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Jnspzp-oIZE>.
- [9] Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008. doi: 10.1561/22000000001. URL <https://people.eecs.berkeley.edu/~jordan/papers/wainwright-jordan-fnt.pdf>.
- [10] Tanguy Marchand, Misaki Ozawa, Giulio Biroli, and Stéphane Mallat. Multiscale data-driven energy estimation and generation. *Physical Review X*, 13:041038, 2023. doi: 10.1103/PhysRevX.13.041038. URL <https://arxiv.org/abs/2207.04941>. Preprint title: Wavelet Conditional Renormalization Group, arXiv:2207.04941.
- [11] Jae Hyun Lim, Nikola B. Kovachki, Ricardo Baptista, Christopher Beckham, Kamyar Azizzadenesheli, Jean Kossaifi, Vikram Voleti, Jiaming Song, Karsten Kreis, Jan Kautz, Christopher Pal, Arash Vahdat, and Anima Anandkumar. Score-based diffusion models in function space. *Journal of Machine Learning Research*, 26(158):1–62, 2025. URL <https://www.jmlr.org/papers/v26/23-1472.html>.
- [12] Rui Wu, Nikola Kovachki, and Burigede Liu. A learning-based domain decomposition method. *arXiv preprint arXiv:2507.17328*, 2025. URL <https://arxiv.org/abs/2507.17328>.
- [13] Sungwon Park, Anthony Zhou, Hongjoong Kim, and Amir Barati Farimani. Generative neural operators through diffusion last layer. In *International Conference on Machine Learning*, 2026. URL <https://arxiv.org/abs/2602.04139>. arXiv:2602.04139.