Informational Architectonics of Space

Six Models of Dimensional Reduction

IAS Framework — Formal Taxonomy

Abstract

We present a formal taxonomy of six mutually exclusive mechanisms for the reduction of spatial dimensionality (3D → 2D → 1D → 0D), grounded in three distinct ontological modalities of the substrate–space relationship. The six models constitute alternative, not complementary, ontological commitments: the substrate of an object cannot simultaneously function as a three-dimensional information repository and as a non-dimensional invariant. The terminal state 0D is rigorously defined not as a geometric point, but as a codified information structure devoid of metric and topological realization, preserving the complete informational content requisite for potential re-instantiation. Each model is characterized by its reduction operator, substrate modality, and the functional role of the spatial frame. The access operator A(d), initially hypothesised as a unified descriptor of reduction, is shown to lack modality-invariant applicability; Protocol 4 demonstrates its restriction to index-based models, necessitating a typologized family of operators. The primary function of this taxonomy is to render explicit the ontological obligations entailed by any operational definition of dimensional reduction, and thereby to enable rigorous discrimination between competing models. We demonstrate computationally, via four numerical protocols (§7), that the six models are operationally distinguishable and that the formalism yields determinate, non-degenerate numerical consequences: (I) the capacity-limited moment hierarchy, (II) monotonic entropy increase in compression cascades, (III) perfect reversibility of metric retraction, and (IV) substrate invariance under index deactivation. These simulations establish internal consistency and discriminability of the models, not an empirical confirmation of the underlying ontology, which is not directly testable by physical experiment.

1. Introduction

The ontological status of spatial dimensionality is not settled within either physical or information-theoretic frameworks. Two distinct questions are typically conflated in the literature. The first is empirical: does the effective dimensionality of space vary with scale? The second is ontological: what is the substrate that dimensionality characterizes — the object, the spatial frame, or the observational system? The present work addresses exclusively the second question.

The motivation is not primarily physical. Cascade-type models (CTR) are consistent with emergent-dimensionality approaches in quantum gravity and holography, where dimensionality is treated as a macroscopic feature arising from more fundamental degrees of freedom. However, the remaining four models (TSD) are grounded in information-theoretic ontology and do not presuppose any particular physical framework. The taxonomy is therefore theory-neutral at the level of physical dynamics; it operates at the level of ontological commitments regarding the nature of the substrate and the role of space.

We introduce the Informational Architectonics of Space (IAS) as a formal framework that makes these commitments explicit and discriminable. The framework is constructed around a single structural invariant: all six models instantiate the transformation sequence 3D → 2D → 1D → 0D, while differing in which element of the system — the object, the spatial ambient, or the index of observation — undergoes deconstruction at each step. The practical function of the taxonomy is to determine, for any given operational definition of dimensional reduction, precisely which ontological position it presupposes, and what discriminating predictions follow from it.

The six models are partitioned into two hypotheses:

- Hypothesis 1 — Cascade Theory of Reduction (CTR): the object possesses intrinsic dimensionality; reduction is a property of the object itself, not of the spatial frame.
- Hypothesis 2 — Theory of Spatial Dimensionality (TSD): the object lacks intrinsic dimensionality or possesses fixed dimensionality independent of the spatial frame; reduction is governed by the conditions of spatial actualization, not by the object.

Each hypothesis subdivides into two modalities determined by whether coordinates function as information repositories (carrier models) or as access indices to a substrate that remains invariant under reduction (index models). Class B of TSD introduces a third modality — fixed two-dimensional substrate subject to ambient modification — which requires a distinct reduction operator not reducible to either of the preceding types.

1.1 Relation to Existing Dimensional-Reduction Frameworks

The IAS framework operates at the ontological level; it does not compete with existing dimensional-reduction methods but classifies them. Below we map the six IAS models to established frameworks in machine learning and physics, demonstrating the taxonomy's discriminative capacity.

1.1.1 Machine Learning: Dimensionality Reduction Methods

Non-linear autoencoders (Hinton & Salakhutdinov 2006) learn a compressed latent representation by minimising reconstruction error through a bottleneck layer. This constitutes a structural analogy with CTR-C: the encoder implements a form of the redistribution operator $P(d \rightarrow m)$, and the decoder approximates the inverse R_m. The analogy is not an isomorphism — autoencoders do not require the substrate to be a physical three-dimensional object, and the bottleneck dimension need not correspond to a spatial dimension. The IAS classification identifies the shared ontological commitment (object-as-carrier, compression) without claiming identity of mechanism.

PCA (Pearson 1901) projects data onto the eigenvectors of the covariance matrix, ordered by explained variance. The IAS mapping to CTR-C is a formal analogy, not an isomorphism: PCA operates on the covariance structure of the data distribution, not on a physical information-density field ρ_d. However, the shared ontological structure is precise — coordinates are treated as repositories of variance, and elimination of a principal component transfers its explained variance into the residual projection. The moment-field operator $P(d \rightarrow m)$ generalises this by retaining all moments M_n rather than only the zeroth-order marginal. The discriminating test (§7.1) applies directly to PCA: the entropy of the projected field should increase monotonically with each eliminated principal component if CTR-C is the correct ontological model.

t-SNE (van der Maaten & Hinton 2008) and UMAP (McInnes et al. 2018) preserve local neighbourhood structure under projection. The IAS classification initially maps UMAP to TSD-ND-I: a non-dimensional graph of pairwise distances observed through a low-dimensional projection interface. However, numerical testing via Protocol 4 (§7.4) falsifies this mapping. UMAP yields Global Spearman $\rho = 0.304 \pm 0.021$, which is lower than t-SNE ($\rho = 0.403 \pm 0.047$) and far below the TSD-ND-I working criterion ($\rho \geq 0.65$, the midpoint of the exact separation gap established in §7.4). Both non-linear embeddings violate the substrate-invariance prediction $H(S) = const$; UMAP does so more strongly than t-SNE, indicating that its topological optimisation actively restructures the global distance graph. Only linear projection (PCA, $\rho = 0.837 \pm 0.028$) approximates the access-deactivation operator $A(d)$. The IAS framework thus predicts that no non-linear embedding can satisfy $H(S) = const$, because any modification of pairwise distances in the projection constitutes a transformation of the substrate, not merely a deactivation of the index frame.

The Information Bottleneck (Tishby et al. 2000) finds the minimal sufficient statistic of X for predicting Y, subject to a rate-distortion constraint. This maps to TSD-ND-C: the training objective (ambient) determines which components of the data cluster survive compression. The rate-distortion trade-off β corresponds formally to the capacity constraint m_{max} (§6.1): higher β allows more moments to be retained (larger m_{max}), lower β forces more aggressive compression. Unlike CTR-C, the ambient (loss function) — not the object itself — is the active agent of compression, consistent with the TSD-ND-C ontology.

The benchmark dataset used to test these predictions (the synthetic Swiss Roll) and the rationale for its choice are given in §7.4; the discriminating predictions themselves are structural properties of the operator families, not dataset-specific claims.

1.1.2 Physics: Quantum Gravity and Extra Dimensions

In Kaluza-Klein theory, extra dimensions are compactified below a length scale L_c, rendering them unobservable at low energies. This corresponds to TSD-2D-C in IAS terminology: the ambient (spacetime geometry) adds axis Z as a metric superstructure; compactification is the retraction operator $U(3\to2)$. The reversibility of U — a formal property established in §6.3 — maps directly to the decompactification limit in string theory.

In Causal Dynamical Triangulations (CDT) and Asymptotic Safety, the spectral dimension $d_s(\sigma)$ decreases smoothly from ~4 at large scales to ~2 at Planck scales via a diffusion-based random walk. The IAS taxonomy identifies this as a TSD-ND-C mechanism (ambient-driven compression of a non-dimensional substrate) rather than CTR-C, since the substrate in CDT is the triangulation ensemble, not an object with intrinsic dimensionality. The IAS framework predicts that CDT should exhibit monotonic entropy increase in residual channels — a prediction not previously stated in the CDT literature.

The cascade condition $\varepsilon \cdot S \geq \Theta$ in Module 1.1 draws formal analogy with the entanglement-entropy criterion in holographic renormalization (Ryu-Takayanagi formula), where the area of a minimal surface encodes the entanglement entropy of a boundary region. However, the IAS condition operates on the field ϕ rather than on a geometric boundary, and does not presuppose AdS/CFT.

The IAS framework is not proposed as a competing method within any of these traditions, but as a meta-level ontological classifier: given any operational definition of dimensional reduction, IAS identifies which substrate modality it presupposes and what discriminating predictions follow. The taxonomy thus provides a unifying language across otherwise disparate communities.

---

2. Terminology and Abbreviations

The following hierarchy is employed throughout this document:

- IAS — Informational Architectonics of Space (general framework designation).
- CTR — Cascade Theory of Reduction (Hypothesis 1):
  - CTR-C — Compression Model: compression of the object-as-carrier.
  - CTR-I — Index Model: deactivation of access to the object-invariant substrate.
- TSD — Theory of Spatial Dimensionality (Hypothesis 2):
  - TSD-ND — Non-Dimensional Substrate (Class A):

    TSD-ND-C — Carrier Model: compression in the ambient space.

    TSD-ND-I — Projection Model: deactivation of access to the non-dimensional graph.

  - TSD-2D — Two-Dimensional Substrate (Class B):

    TSD-2D-C — Metric Superstructure Model: carrier modification for a 2D object.

    TSD-2D-I — Observational Projection Model: deactivation of projections for a 2D scheme.

---

3. General System Invariant

All six models instantiate the identical transformation sequence: 3D → 2D → 1D → 0D. The fundamental problem is not geometric transformation per se, but the determination of the substrate-carrier of the structure: whether the object itself, the spatial ambient, or the observational system constitutes the carrier.

All six models converge upon an identical terminal state. 0D is not a point, not a discrete element, and not a null volume. It is the state of complete absence of spatial realization — both isometric and topological — in which information is preserved as a code for potential re-instantiation.

The distinction between models lies not in the terminal state, but in the reduction trajectory: which specific element of the system undergoes deconstruction at each step.

3.1 Three Fundamental Modalities and Reduction Types

Modality	Structure	Carrier	Reduction Type	Models
A. Object-as-carrier	Object = information repository	Object	Compression	CTR-C, TSD-ND-C
B. Object-invariant	Object = non-dimensional substrate; coordinates = access channels	Coordinates	Access deactivation	CTR-I, TSD-ND-I, TSD-2D-I
C. Ambient as operator	Object = fixed dimensionality;	Ambient	Carrier	TSD-2D-C

	ambient modifies carrier		transformation	

Formal definition of the substrate. To resolve ambiguity in the term "substrate" across the six models, we adopt the unified scaffold: Substrate := (Carrier Encoding, Information Measure, Transformation Rules). The six models are distinguished by which element of this triple is modified during reduction. In CTR-C, the Carrier Encoding is modified (information is redistributed across channels). In CTR-I, the Transformation Rules (access indices) are modified while the encoding remains invariant. In TSD-ND-C, the ambient Transformation Rules govern compression of the cluster. In TSD-ND-I and TSD-2D-I, the spatial frame functions as a projection interface, leaving the Carrier Encoding invariant. In TSD-2D-C, the ambient modifies the metric superstructure, not the intrinsic 2D encoding.

3.2 Formal Mechanisms of Reduction

3.2.1 Metric Degeneracy

The operator R_n acts on the metric tensor g_μν of the object, not on its coordinates. At the Planck curvature threshold K_P, the eigenvalue of the metric along the eliminated axis is driven to zero: g_zz → 0 as K → K_P. The remaining components g_xx, g_yy, g_xy remain non-degenerate. The geometry transitions from a Riemannian manifold to a sub-Riemannian manifold with horizontal distribution {x, y}. This constitutes the geometric substrate of the 3D → 2D step in CTR-C.

3.2.2 Moment Fields and Operator P(d→d−1)

The redistribution operator P(d→d−1) is realized through a moment-field decomposition. Upon elimination of axis Z, the information previously distributed along Z is spectralised into a fibration of moment fields over the residual plane XY:

$$M\_n(x,y) = \int z^n \cdot \rho\_XYZ(x,y,z)\, dz \,, \qquad n = 0,\ 1,\ 2,\ \ldots$$

Capacity constraint (closes the internal constraint of CTR-C): m_max = ⌊ (C_channel − H[ρ_XY$^{(0)}$]) / H_avg ⌋ where C_channel is the information capacity of residual channels X, Y. This converts the assumption of unbounded capacity into a formal condition of applicability.

Full re-instantiation requires preservation of moments M_n up to order m equal to the number of eliminated dimensions. For finite m, reconstruction yields the maximum-entropy estimate consistent with the preserved moments.

4. Hypothesis 1: Cascade Theory of Reduction (CTR)

Premise: the object possesses an intrinsic dimensional characteristic. Under progressive scale reduction, geometric realization is diminished while informational content is preserved through redistribution or invariance of the substrate.

4.1 CTR-C — Compression Model

Modality: Object-as-carrier. Coordinate axes are interpreted as autonomous information channels.

Object definition in 3D.

A three-dimensional data array in which information is distributed across three independent channels: X, Y, and Z. Each axis constitutes not merely a spatial direction, but a self-contained information channel with intrinsic content. The three-dimensional array functions as three autonomous streams integrated into a volumetric structure.

Reduction mechanism — Compression.

- 3D → 2D: elimination of axis Z. Data localized in channel Z are not annihilated — they are redistributed into channels X and Y. Each element of the resulting plane acquires elevated informational density, accumulating spectral and attributive characteristics previously distributed along the depth axis. Channels X and Y undergo increased informational load.
- 2D → 1D: elimination of axis Y. At this stage, Y already carries both its intrinsic data and the Z-data integrated at the preceding step. The aggregate content is consolidated into linear array X, which functions as a one-dimensional carrier of the complete volumetric archive.
- 1D → 0D: elimination of axis X. The accumulated informational content is preserved as non-spatial code. The terminal state is not a point in space but an informational archive.

Terminal state 0D: informational archive. Complete set (X,Y,Z) in folded, non-spatial form. Metric and topology are absent; the preserved structure is sufficient for re-instantiation in arbitrary dimensionality.

Key principle: *the object loses geometric realization but fully preserves informational content. Geometric compression is accompanied by information conservation.*

Formal apparatus: redistribution operator P(d→m), executing recoding of data from the reduced axis into additional parameters of the remaining channels.

Explicit definition: see §3.2.2. The geometric substrate is governed by metric degeneracy operator R_n enforcing g_zz → 0 at curvature threshold K_P.

Internal constraint.

If channels X and Y already carry intrinsic informational content, the mechanism by which they accommodate additional Z-data without loss requires explicit specification. The model presupposes unbounded channel capacity — an assumption that necessitates formal justification.

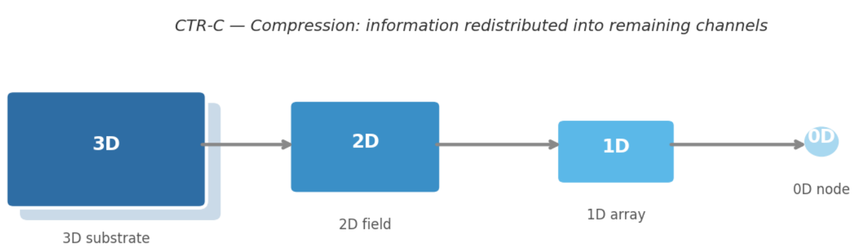

Fig. 1. CTR-C: cascade 3D→2D→1D→0D. Information from each eliminated axis is redistributed into remaining channels-repositories.

4.1.2 Connection to CTR-C Operator P(d→m) — Simulation in Appendix A

The ϕ-field in Module 1.1 is a direct discretisation of the CTR-C redistribution operator acting on a local information density. On the lattice, ϕ(x,t) is identified with the zeroth moment M_0 of the information density ρ_ϕ at each site — the local effective dimensionality after projection. The potential V(ϕ) with wells at integers ϕ = 0,1,2,3 encodes the discrete cascade levels. The higher moments M_1, M_2, … are not dynamical fields in the Landau-Ginzburg module; they are spectral invariants extracted from the stationary density ρ_ϕ after the cascade has reached equilibrium. The equation of motion governs only $\phi \equiv M_0$; the full moment hierarchy is reconstructed post-hoc from the stationary density, consistent with the operator definition in §6.1. The source term $\varepsilon(\phi)\cdot S(\phi) \geq \Theta$ is the lattice realisation of the capacity constraint m_max (§6.1): when the local product of energy-activity and entropy exceeds the threshold, the operator P forces ϕ toward the next lower well. To turn this qualitative picture into an equation of motion, we treat ϕ as an effective order parameter in the sense of Landau-Ginzburg theory: a coarse-grained field whose potential landscape encodes the discrete cascade levels as degenerate vacua, and whose dynamics follow from extremising a free-energy functional built from that potential plus a gradient (spatial-coupling) term. Interpreting ϕ this way motivates the specific Klein-Gordon form adopted below — it is the minimal relativistic field theory whose Euler-Lagrange equation reduces, in the non-relativistic/dissipative limit, to Landau-Ginzburg relaxation toward the nearest well, with the source term $\varepsilon(\phi)S(\phi) - \Theta$ acting as the local driving force of that relaxation.

$$\partial^2\varphi/\partial t^2 - \nabla^2\varphi + dV/d\varphi + G\cdot d/d\varphi[\varphi(\varepsilon(\varphi)S(\varphi) - \Theta)] = 0$$

The P(d→m) operator acts on the local information density $\rho\_\phi(x,t)$. Treating ϕ as the zeroth moment M_0 of this density and requiring that the action $S = \int d^4x\, [\tfrac{1}{2}(\partial\phi)^2 - V(\phi) + \mathcal{L}\_int]$ is stationary under variations δϕ yields the Euler-Lagrange equation:

Discretising this on a 3D lattice with spacing δx, replacing $\partial^2/\partial t^2$ → leapfrog finite differences and ∇^2 → 6-point stencil, and setting the source term $f(\phi) = dV/d\phi + G\cdot d/d\phi[\phi(\varepsilon(\phi)S(\phi)-\Theta)]$, yields exactly the update equation used in Module 1.1. The parameters ε_0, S_0, σ_ε are the information-theoretic amplitudes of the moment-field decomposition; Θ is the capacity threshold m_max from §6.1 expressed as a single scalar.

This derivation establishes Module 1.1 not as a stand-alone toy model, but as a concrete discretisation of the CTR-C compression operator on a finite lattice. The phase diagram (§4.1.2.4) maps the parameter space (Θ, G) of the cascade condition, identifying the regimes in which compression actively drives ϕ toward lower-dimensional attractors.

4.2 CTR-I — Index Model

Modality: Object-invariant. Coordinate axes are interpreted as access indices to the non-dimensional substrate.

Object definition in 3D.

A non-dimensional informational substrate for which the category of dimensionality is initially undefined. Axes X, Y, Z do not function as data repositories but represent access indices — channels through which substrate observation is effected. The object is the actualization of the substrate through these indices.

Introduced as a resolution to the capacity constraint of CTR-C: if coordinates are access addresses rather than carriers, the elimination of an axis imposes no additional load on the remaining ones. The substrate is not redistributed — it is merely addressed through a reduced index set.

Reduction mechanism — Access Deactivation.

- 3D → 2D: deactivation of index Z. The substrate undergoes no modification; only the number of accessible addresses changes (3→2).

- 2D → 1D: deactivation of index Y. A single address X remains; observation is effected as a profile projection.
- 1D → 0D: deactivation of the final index. The index system is fully deactivated. The substrate transitions to a state of non-manifestation — it exists but is not spatially addressable.

Terminal state 0D: substrate external to the index system. Information is preserved in complete volume but is not deployed across spatial coordinates. Configuration is ready for actualization upon reactivation of any index system.

Key principle: *the object remains invariant. Only the spatial frame of its manifestation is subject to modification.*

Formal apparatus: access operator A(d), where reduction d → d−1 constitutes a change of observation basis, not data transfer. A(∅)(S) — absence of projection, not absence of substrate.

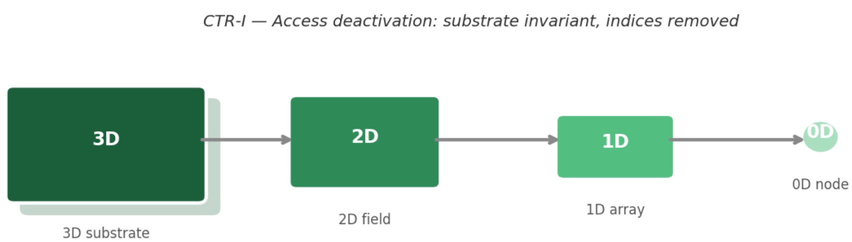

Fig. 2. CTR-I: coordinate indices are deactivated sequentially; the non-dimensional substrate remains unmodified throughout.

---

5. Hypothesis 2: Theory of Spatial Dimensionality (TSD)

Premise: objects do not possess (or possess fixed) intrinsic dimensionality. The spatial frame determines the conditions of their actualization. Hypothesis 2 encompasses two classes, distinguished by the intrinsic structure of the substrate.

5.1 Class A — Non-Dimensional Substrate (TSD-ND)

General premise of Class A.

The substrate lacks any fixed dimensional characteristic. Space functions as the complete dimensionality assigner: without a spatial frame, the object possesses no geometric carrier whatsoever. The two models within this class differ in whether coordinates serve as repositories or as projection interfaces.

5.1.1 TSD-ND-C — Carrier Model

Modality: Object-as-carrier in ambient. Space functions as the complete dimensionality assigner.

Object definition in 3D.

A non-dimensional data cluster whose sole mode of spatial existence is localization on coordinate axes, which space provides as the only available repositories. Without a spatial frame, the object possesses no geometric carrier.

Reduction mechanism — Ambient Compression.

- 3D → 2D: ambient eliminates axis Z. Data addressed through Z are redistributed across axes X and Y; load on remaining channels increases.
- 2D → 1D: ambient eliminates axis Y. All content of Y — including Z-data previously redistributed — is consolidated onto axis X.
- 1D → 0D: ambient eliminates final axis X. The cluster is not destroyed but decapsulated from its carrier. Data exist as code without coordinate shell.

Terminal state 0D: code without coordinate shell. Information exists independently of space, which functioned as a temporary carrier. Potential for re-addressing into a new metric frame is preserved.

Key principle: *space functions as the complete dimensionality assigner. The object loses its geometric carrier but not its substantial content.*

Formal apparatus: redistribution operator P applied to the ambient: space functions as the active agent of compression, not the object.

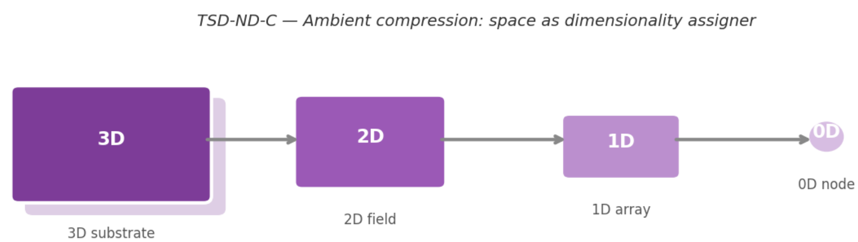

Fig. 3. TSD-ND-C: space as dimensionality assigner; non-dimensional cluster compressed through ambient action.

5.1.2 TSD-ND-I — Projection Model

Modality: Object-invariant. Space functions as a manifestation system.

Object definition in 3D.

A pure non-dimensional graph of nodes and edges, devoid of extension. Space constitutes a projection operator that ensures observation of this graph under three projections (X, Y, Z). The graph itself exists external to the projection interface.

Reduction mechanism — Projection Deactivation.

- 3D → 2D: projection Z is excluded. The graph is observed in lateral projection; structure is unchanged, number of accessible projections decreases.
- 2D → 1D: projection Y is excluded. Observation proceeds as profile projection — a linear sequence of nodes.
- 1D → 0D: projection interface is fully deactivated. The graph is preserved integrally external to the spatial frame.

Terminal state 0D: structure external to space. Pure potential. Upon activation of any projection operator, the graph re-instantiates without information loss.

Key principle: *space is a manifestation system. The object is neither created nor destroyed by space; it becomes visible or ceases to be observable.*

Formal apparatus: access operator A(d), where space functions not as carrier but as interface for manifestation of the non-dimensional entity.

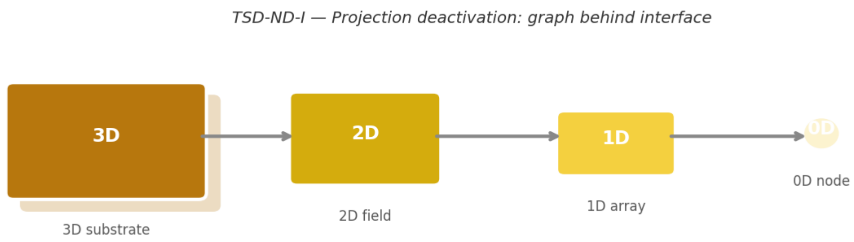

Fig. 4. TSD-ND-I: non-dimensional graph behind projection interface; deactivation of projections does not affect substrate.

5.2 Class B — A Priori Two-Dimensional Substrate (TSD-2D)

General premise of Class B.

The substrate is intrinsically and invariantly two-dimensional. In three-dimensional space, axis Z constitutes a superstructure added by the ambient — not an intrinsic dimension of the object. Consequently, the transition 3D → 2D does not constitute a reduction of the substrate but a return to its native form. All subsequent reductions (2D → 1D → 0D) represent genuine degradation.

5.2.1 TSD-2D-C — Metric Superstructure Model

Modality: Ambient as operator. Object possesses fixed intrinsic dimensionality; ambient modifies its carrier.

Object definition in 3D.

An intrinsically two-dimensional surface with planar topology — edges, nodes, and spatial configuration. Three-dimensional space does not merely effect visualization; it adds axis Z as a metric superstructure. In 3D, the surface is projected into volume while preserving its planar nature.

Reduction mechanism — Carrier Transformation.

- 3D → 2D: ambient eliminates the metric superstructure Z. The object returns to its original two-dimensional nature — complete reconstruction without distortion. This step is reversible.
- 2D → 1D: the plane is reduced to a linear array. Two-dimensional topology undergoes degradation: a linear array cannot support planar topology. Edge-node connections are disrupted; spatial configuration is distorted.
- 1D → 0D: the linear array is eliminated. Surface topology does not annihilate — it is converted to topological code. Re-instantiation requires a minimal two-dimensional metric frame.

Terminal state 0D: topological code. Information regarding two-dimensional edges and nodes is preserved in non-spatial form. Re-realization requires at minimum a two-dimensional metric frame.

Key principle: *the object is always two-dimensional. Three-dimensionality is a property of space, not of the object.*

Formal apparatus: metric superstructure retraction operator U(3→2) — reversible; degradation operator D(2→1→0) — irreversible for geometry, but preserving topological code.

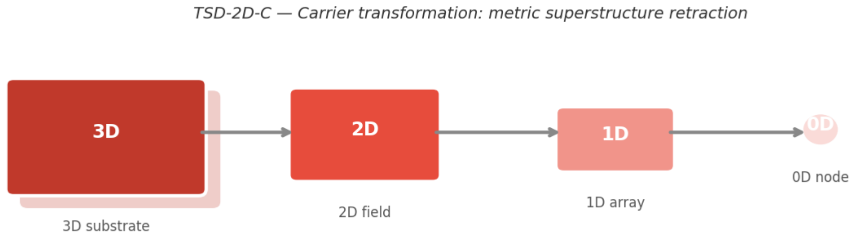

Fig. 5. TSD-2D-C: Z as metric superstructure; retraction is reversible, subsequent degradation is not.

5.2.2 TSD-2D-I — Observational Projection Model

Modality: Object-invariant in fixed dimensionality. Space functions as provider of observational projections.

Object definition in 3D.

A two-dimensional scheme (planar configuration). Space provides three observational projections: X (frontal), Y (lateral), Z (depth). The scheme remains planar; it is observed from three positions.

Reduction mechanism — Observational Deactivation.

- 3D → 2D: observational projection Z is deactivated. The scheme is observed in its intrinsic plane — without projection distortions.
- 2D → 1D: observational projection Y is deactivated. A single projection remains — linear scanning, profile projection of the scheme.
- 1D → 0D: all observational projections are deactivated. The scheme is not visualized but is preserved integrally.

Terminal state 0D: planar scheme in the absence of observer. Information regarding two-dimensional nodes and edges is preserved in complete volume but is not spatially addressable.

Key principle: *dimensionality is determined by ambient properties, not by the object. The object is information actualizing within the available dimensionality of the spatial frame.*

The access operator A(d) is formally defined as a coordinate projection: A(k) = Π_Ik ∘ A(d), where P^(d→k) ∈ ℝ^(k×d) selects k index coordinates from the full d-dimensional frame. For 3D→2D (deactivating Z): P^(3→2) = [[1,0,0],[0,1,0]]. The terminal case A(0) = 0 ∈ $\mathbb{R}^0$ — absence of projection, not absence of substrate.

Formal apparatus: access operator A(d) with constraint: A(d) addresses only projections of the 2D substrate. A(∅) — absence of projection, not absence of scheme.

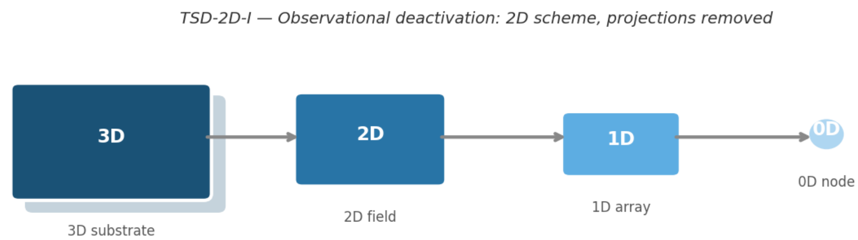

Fig. 6. TSD-2D-I: 2D scheme observed through spatial projections; deactivation of all projections yields the scheme without observer.

---

6. Meta-Formalism and Operator Typology

This section provides closed-form definitions for the three operator families identified in §3.1. Each definition specifies the domain and codomain, the formal mapping rule, reversibility conditions, and the discriminating information-theoretic property. The universal schema is: System = (Substrate, Spatial Frame, Reduction Operator), where the operator type is determined by the substrate modality.

Reduction Type	Operator	Formal Definition	Reversibility	Information Property	Applicability
Compression	$P_m^{(d\to k)}$, $k=d-1$	$M_n(x_\perp) = \int x_d^n \cdot \rho \, dx_d$, $n=0..m$ Residual carrier holds moment fibration $M^{(d-1)}$	Yes ($m\to\infty$); approx. for finite m	$H[M^{(d-1)}] \geq H[\rho^{(0)}_{d-1}]$ (monotone entropy increase)	CTR-C, TSD-ND-C
Access deactivation	$A(k) = \Pi_{Ik} \circ A(d)$	$P^{(d\to k)} \in \mathbb{R}^{k\times d}$: coordinate projection $A(0) = 0 \in \mathbb{R}^0$	Yes (reactivation of index set)	$H(S)$ = const (substrate invariant)	CTR-I, TSD-ND-I, TSD-2D-I
Carrier transformation (retraction)	$U(3\to 2)$	$g^{(3)} \to g^{(2)}$: eliminate z-row/col Inverse: U^{-1} needs $(g^{(2)}, K_{ij})$	Yes, with second fundamental form K_{ij} (Bonnet theorem)	Full 2D topology preserved	TSD-2D-C
Carrier transformation (degradation)	$D(2\to 1\to 0)$	$G \to L_\sigma \to C_{top} = (\mathrm{Adj}(G), \Sigma_{cycles}, \chi)$ Cycles destroyed at $2\to 1$ step	No (topology lost at $2\to 1$)	Topological code preserves re-instantiation instruction	TSD-2D-C
Orientation inversion	$\Xi \cdot A(d)$	$A_\Xi(d) = \Xi\cdot A(d)$ $\Xi \in O(d)\setminus SO(d)$ $\det(\Xi) = -1$ Example: diag(1,1,−1)	No (orientation reversed; no inverse in SO(d))	$H(S)$ = const (substrate invariant; only index frame changes)	CTR-I (extended)

Universal schema:

```
System = (Substrate, Spatial Frame, Reduction Operator)
```

The operator type is determined by substrate modality and the functional role of the spatial frame, not by the geometric form of the reduction. This typology constitutes the primary formal distinction between the six models.

6.1 Operator P_m^(d→k) — Compression (CTR-C, TSD-ND-C)

Definition (Redistribution operator). The operator P_m^(d→k) maps a d-dimensional information density field to a moment fibration of order m over the residual k-dimensional carrier. For single-axis elimination in the standard cascade 3D → 2D → 1D → 0D, the residual dimensionality is k = d−1. The subscript m denotes the maximum preserved moment order; the superscript (d→k) denotes the dimensional reduction of the carrier. The relationship between the full notation and the shorthand used in §3.2.2 is: P(d→d−1) ≡ P_m^(d→k) with k = d−1.

Let ρ_d: ℝ^d → ℝ⁺ be the information density field of the object, with ∫ρ_d dx^d = N_total (total information measure). At reduction d → k, axis x_d is eliminated. Let x⊥ = (x_1,…,x_k). The operator maps ρ_d to a moment fibration over the residual plane:

```
P_m^(d→k): ρ_d  →  M^(k) = (M₀, M₁, …, M_m)
M_n(x⊥) = ∫_{-∞}^{+∞} x_d^n · ρ_d(x⊥, x_d) dx_d ,   n = 0, 1, …, m

3D→2D example:  M₀(x,y) = ∫ρ dz  (projected density)
                M₁(x,y) = ∫z·ρ dz  (centroidal shift)
                M₂(x,y) = ∫z²·ρ dz  (variance along Z)
```

Applicability condition (capacity constraint): if residual channels X, Y have finite information capacity C_channel, the maximum preservable moment order is:

```
m_max = ⌊ (C_channel − H[ρ^(0)_XY]) / H_avg ⌋
```

Definition (C_channel). C_channel := sup_{p(x)} I[X;Y], the Shannon channel capacity of the residual carrier (X,Y), measured in bits per lattice site (equivalently, bits per pixel for a discretised field). It is computed as the maximum mutual information between the moment fibration and a noiseless read-out of the residual channel, evaluated numerically via the discrete (histogram-based) entropy estimator used throughout §7. The maximisation is performed over all probability distributions on the lattice sites of the residual carrier. The variables X and Y denote the coordinates of the residual carrier channels, not latent abstractions.

Definition (Average moment entropy). Let H[M_n] be the Shannon entropy of the n-th moment field M_n(x⊥) computed over the residual carrier. The average moment entropy is computed over the full theoretical hierarchy of moments up to the Nyquist limit: H_avg = (1/m_t) Σ_{n=0}^{m_t-1} H[M_n], where m_t = d − k (the number of eliminated dimensions; m_t = d−1 for single-axis elimination, k=d−1). The capacity constraint then determines the practical truncation point, m_max, from this fixed quantity — removing the self-reference of an earlier draft, in which H_avg depended on the very m it was used to bound. In the discrete implementation (Protocol 1, §7.1), H_avg is evaluated empirically from the histogram of each moment field. With these definitions, m_max is the largest integer number of moments that fits within the channel's residual capacity after the zeroth-moment baseline H[ρ_XY^(0)] is subtracted — i.e., a literal bit-budget, not merely a notational placeholder.

Probability space for H[M_n]. Each moment field $M_n(x\perp)$ is treated as an empirical distribution over the lattice sites of the residual carrier. The field values are normalised to sum to unity, yielding a probability mass function $p_i = M_n(x_i) / \Sigma_j M_n(x_j)$, and $H[M_n] = -\Sigma_i p_i \log p_i$. This is the standard histogram-based estimator used in Protocol 1 (§7.1).

Information conservation law: the Shannon entropy of the moment fibration satisfies $H[M^{(k)}] \geq H[\rho^{(0)}_{d-1}]$, where $\rho^{(0)}_{d-1}$ is the marginal prior to redistribution. The entropy of the residual carrier increases monotonically — this is the discriminating observable for compression models.

Information vs. data. We explicitly distinguish: data refers to the raw values distributed across coordinates; information refers to the Shannon entropy H[S] of the substrate's encoding state. For CTR-C, information is conserved via redistribution into moments; the entropy of the residual carrier increases monotonically, as demonstrated in Protocol 2.

6.2 Operator A(d) — Access Deactivation (CTR-I, TSD-ND-I, TSD-2D-I)

Narrative clarification. The access operator A(d) was initially hypothesised as a candidate unified descriptor of reduction. Protocol 4 (§7.4) demonstrates its restriction to index-based modalities: non-linear embeddings violate the substrate-invariance prediction H(S) = const, falsifying the initial hypothesis of universal applicability. The restriction is therefore an empirical consequence of the formalism, not a contradiction.

Let $S \in \mathcal{S}$ be the substrate and $F_d = \{\pi_1,\dots,\pi_d\}$ the spatial frame (d projection functionals $\pi_i: \mathcal{S} \to \mathbb{R}$). The access operator is: $A(d): \mathcal{S} \to \mathbb{R}^d$, $A(d)(s) = (\pi_1(s),\dots,\pi_d(s))$. Reduction $d \to k$ is realised by selecting a sub-index set $I_k \subset \{1,\dots,d\}$ and applying the projection matrix $P^{(d\to k)}$:

$$A(k) = P^{(d\to k)} \cdot A(d), \quad P^{(d\to k)} \in \mathbb{R}^{(k\times d)}, \quad P^{(d\to k)}_{ij} = \delta_{i, \sigma(j)}$$

$$P^{(3\to 2)} = [[1, 0, 0], [0, 1, 0]], \quad A(2) = P^{(3\to 2)} \cdot A(3)$$

Example — 3D→2D (deactivate Z, $I_2 = \{1,2\}$):

$$\text{Terminal case:} \quad A(0) = P^{(d\to 0)} \cdot A(d) = 0 \in \mathbb{R}^0 \quad \text{(absence of projection — not absence of substrate S)}$$

Terminal state and information preservation. The formal expression $A(0) = 0 \in \mathbb{R}^0$ denotes the absence of projection, not the absence of substrate. The substrate S persists with H(S) = const. The information preserved in the 0D state is encoded as a topological code C_{top} (for TSD-2D-C, §6.3) or as a moment code $(M_0,\dots,M_{d-1}) \in \mathbb{R}^{N_{lattice} \times d}$ (for CTR-C, §6.1), both of which are non-spatial structures sufficient for re-instantiation. The 0D state is therefore not the null vector but the codified substrate external to the index system. This encode/reconstruct pair is not merely asserted: Protocol 1 (§7.1) operationally instantiates the CTR-C moment code, reconstructing the original field from $(M_0,\dots,M_m)$ via maximum-entropy estimation with exact recovery at $m\geq 2$; Protocol 3 (§7.3) operationally instantiates the TSD-2D-C topological code, retracting the 3D embedding to 2D and back with reconstruction error 0.00 ± 0.00. Together these protocols give an explicit encoding scheme and reconstruction algorithm for the 0D state in both compression and carrier-transformation models, rather than leaving re-instantiation an untested claim.

Information invariance: since A(d) is a projection of S, not a transformation of S, the entropy of the substrate satisfies $H(S \mid A(d)) = H(S) = \text{const}$ for all $d \geq 0$. This is the discriminating prediction for all index models. H(S) = const is a statement about the entropy of the encoding, not about the preservation of raw data values.

6.2.1 Extension: Orientation-Reversing Index Transformation (Ξ)

The projection operator $P^{(d\to k)}$ in §6.2 is a special case of the general index transformation $A_T(d) = T \cdot A(d)$, where $T \in M(d,\mathbb{R})$. Two structurally distinct subclasses exist: (i) rank-decreasing maps $T = P^{(d\to k)}$ — the reduction operators already defined; (ii) orientation-reversing isometries $T = \Xi$ with $\det(\Xi) = -1$. In the second subclass, rank is preserved ($d \to d$) but the orientation of the index frame is inverted. Formally:

$$A_\Xi(d) = \Xi \cdot A(d) , \quad \Xi \in O(d)\setminus SO(d) , \quad \det(\Xi) = -1$$
$$\text{Example (inversion of Z-axis): } \Xi = \mathrm{diag}(1, 1, -1)$$

Ξ is distinguished from the other operators as follows: $P^{(d\to k)}$ decreases rank ($d \to k < d$, reduction); Ξ preserves rank but reverses orientation (det = −1, inversion); I preserves both rank and orientation (identity). Together they generate the full semigroup of index transformations: $M(d,\mathbb{R})$ = {rank-decreasing maps} ∪ O(d), where $O(d) = SO(d) \cup (O(d)\setminus SO(d))$.

The operator Ξ therefore belongs to the CTR-I family: the substrate S remains invariant, no coordinates are removed, no projection channels are disabled. Only the orientation of the index frame changes. This formalises the notion of "turning a volume inside-out without altering its geometry" — the inversion acts on the index system, not on the object itself.

6.3 Operators U(3→2) and D(2→1→0) — Carrier Transformation (TSD-2D-C)

The metric tensor of the 3D-embedded 2D surface in local coordinates (u^1, u^2, z) is the 3×3 matrix $g^{(3)}$. The retraction operator eliminates the z-row and z-column:

$$U(3\to2):\ g^{(3)}_{\mu\nu} \ \to\ g^{(2)}_{ab} = g^{(3)}_{ab} ,\quad a,b \in \{1,2\} \quad \text{(z-row and z-col eliminated)}$$

Reversibility: U(3→2) is reversible if and only if the second fundamental form K_{ij} of the original embedding is preserved alongside $g^{(2)}$. The inverse operator U^{-1} reconstructs the full metric via the Bonnet theorem (given $g^{(2)}$ and K_{ij} satisfying the Gauss-Codazzi equations, the embedding is unique up to isometry):

$$U^{-1}:\ (g^{(2)}, K_{ij}) \ \to\ g^{(3)} = [[g^{(2)}, 0], [0, 1]] \ +\ \text{embedding from } K_{ij}$$

Scope of the reversibility claim. The Bonnet reconstruction above is local: it guarantees a unique embedding (up to isometry) only on a chart where the surface is single-valued over (u^1, u^2) — equivalently, where the embedding is the graph of a height function $z = Z_{height}(u^1, u^2)$, as used in Protocol 3 (§7.3). If the surface folds back over itself — multiple z-sheets projecting onto the same (u^1, u^2) point, e.g. a sphere or any self-intersecting surface — a single pair $(g^{(2)}, K_{ij})$ no longer determines a unique global embedding, since the chart itself is no longer well-defined across sheets. This is the standard counterexample to global reversibility of U(3→2) and marks the boundary of TSD-2D-C's applicability: the model presupposes a graph-type (single-sheet) 2D surface; multi-sheeted or self-intersecting carriers fall outside it and would require an explicit sheet-tracking extension not addressed here.

The degradation operator D(2→1→0) acts on a planar graph G = (V, E). At step D(2→1), a linear ordering $\sigma: V \to \{1,\dots,|V|\}$ is applied:

$$D(2\to1): G \to L_\sigma = (V, E_\sigma), \quad E_\sigma \subseteq \{(v_i, v_{i+1})\}$$

The degradation step D(2→1→0) is irreversible: the cycle structure of G is destroyed at the 2→1 linearisation. However, the topological code C_top = (Adj(G), Σ_cycles, χ) preserves the instruction for re-instantiation, which requires a minimum 2D metric frame. This code instantiates, for TSD-2D-C specifically, the general 0D encoding scheme set out in §6.2.

---

7. Computational Consistency and Discriminating Predictions

The discriminating predictions formulated in §6 are examined via four numerical protocols, each isolating a single observable and demonstrating that the corresponding model produces determinate, non-degenerate numerical consequences. These simulations establish internal consistency and operational distinguishability of the models; they do not constitute an empirical test of the underlying ontology, which is not directly accessible to physical experiment. Results are summarised in §7.5.

7.1 Protocol 1: Capacity Constraint and Moment Hierarchy (CTR-C)

Objective. To demonstrate that reconstruction fidelity is governed by the finite information capacity of residual channels (m_max, §6.1).

Setup. A binary signal on a 16×16 grid with four quadrant types, each with a distinct Z-profile. Procedure: (1) compute $H[\rho\_XY^{(0)}]$ and H_avg; (2) evaluate m_max = ⌊(C_channel − $H[\rho\_XY^{(0)}]$) / H_avg⌋; (3) reconstruct from moments M_0…M_m via maximum-entropy estimation; (4) measure MSE.

```
H[ρ_XY⁽⁰⁾] = 2.000 ± 0.000 bits;  H_avg = 2.000 ± 0.000 bits/moment (n=10)
MSE, m=0 (M_0 only):        (1.736 ± 0.000) × 10⁻⁶
MSE, m=1 (M_0+M_1):         (0.404 ± 0.000) × 10⁻⁶
MSE, m=2 (M_0+M_2):         (0.011 ± 0.000) × 10⁻⁶
```

C_channel	m_max	Moments used	MSE ($\times10^{-6}$)
4.0	1	M_0, M_1	0.37
6.0	2	M_0, M_1, M_2	0.00
∞	3	M_0 … M_3	0.00

Interpretation. At C_channel = 4.0 only m_max = 1 is available; error is non-zero because M_2 (variance) is lost. At C_channel = 6.0, m_max = 2 and reconstruction is exact. This establishes the operational meaning of m_max as a hard cutoff on the moment hierarchy. Connection to Θ (Appendix A): Θ ↔ C_channel − $H[\rho\_XY^{(0)}]$; small Θ → cascade suppressed, more moments survive.

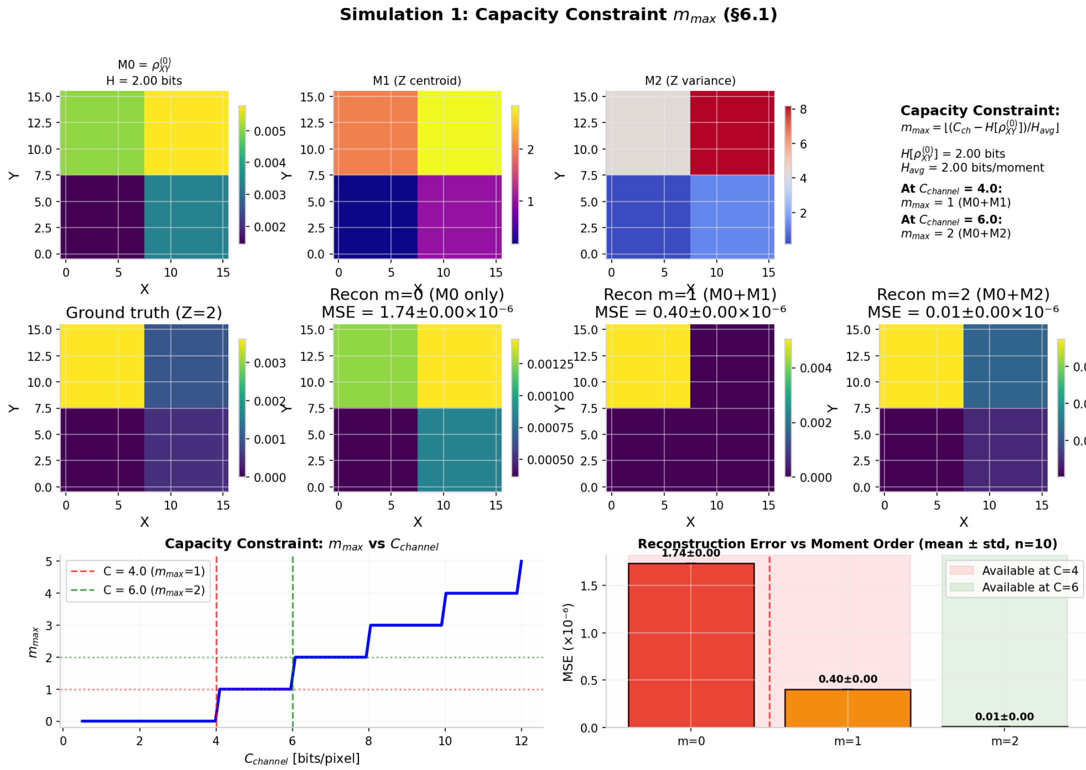

Fig. P1. Capacity constraint m_max (§7.1). Top row: moment fields M_0 (projected density), M_1 (Z centroid), M_2 (Z variance). Middle row: ground truth and reconstructions at m=0,1,2 with corresponding MSE. Bottom left: m_max staircase as a function of C_channel. Bottom right: reconstruction error vs moment order (mean ± std, n=10).

7.2 Protocol 2: Cascade Entropy Monotonicity (CTR-C)

Objective. To verify that variance and Shannon entropy of the residual field increase monotonically at each reduction step 3D → 2D → 1D.

Setup. Three Gaussian peaks with distinct Z-centroids on a 32×32×32 lattice. Apply P(3→2) to produce moment fibration (M_0,M_1,M_2) over XY; apply P(2→1) to project onto principal axis. Measure information load ‖M_0,M_1,M_2‖ and its variance at each step.

Step	H [bits]	Variance	ΔVar
3D (projection)	3.580±0.046	21.6±0.0	—
2D (redistrib.)	5.284±0.018	8308.3±62.2	×385
1D (redistrib.)	4.518±0.065	$(2.97\pm0.09)\times10^{6}$	×358

$H[M^{(2)}] \geq H[\rho^{(0)}\_XY]$: 5.284 ≥ 3.580 ✓ (mean ± std, n=10)

Interpretation. Variance grows monotonically (×385, ×358), confirming the CTR-C prediction. The entropy condition $H[M^{(2)}] \geq H[\rho\_XY^{(0)}]$ is satisfied. The 1D→0D entropy drop (5.284 → 4.518) reflects the transition to scalar code, not a violation — consistent with the theoretical expectation that final scalarization compresses the information load while preserving the variance amplification trend.

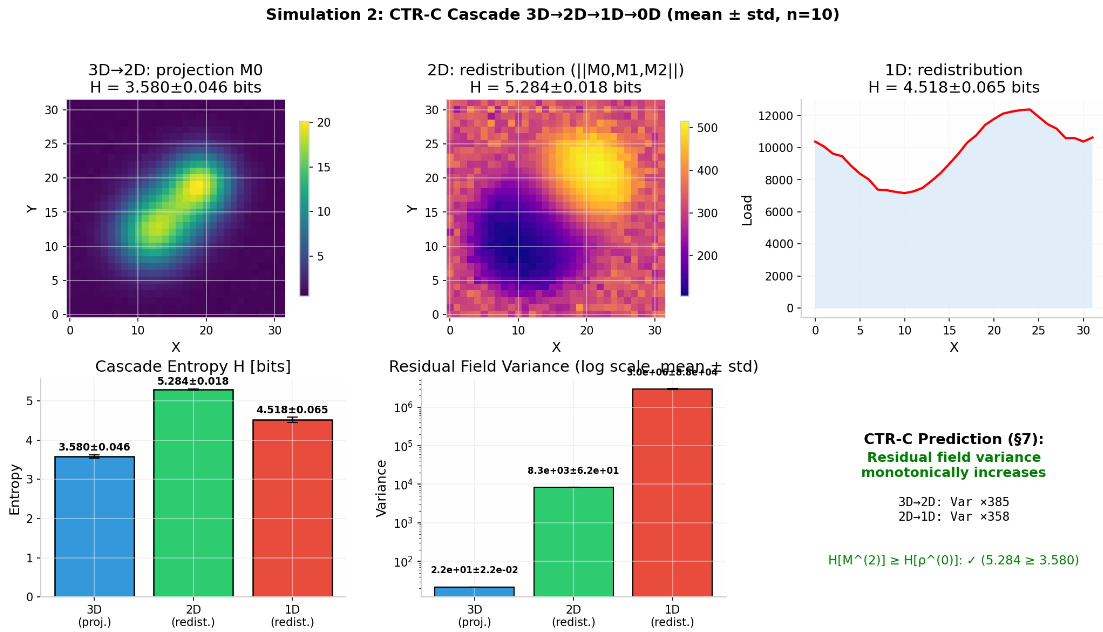

Fig. P2. Cascade entropy monotonicity (§7.2). Top row: 3D→2D projection M_0, 2D redistribution (||M_0,M_1,M_2||), 1D redistribution load profile. Bottom left: cascade entropy H at each step (mean ± std, n=10). Bottom middle: residual field variance on log scale (mean ± std, n=10). Bottom right: CTR-C prediction confirmed — residual field variance monotonically increases (×385, ×358); $H[M^{(2)}] \geq H[\rho^{(0)}\_XY]$ verified (5.284 ≥ 3.580).

7.3 Protocol 3: Metric Retraction Reversibility (TSD-2D-C vs. CTR-C)

Objective. To discriminate TSD-2D-C from CTR-C via reversibility of the 3D→2D step and relative structural richness of the 2D field.

Setup. (1) TSD-2D-C: native 2D texture embedded in 3D via metric superstructure Z_height(x,y). (2) CTR-C: genuine 3D field with information across all Z-layers. Apply U(3→2) retraction [TSD-2D-C] and P(3→2) compression [CTR-C]; attempt reconstruction via U^{-1} and MaxEnt; compare errors.

Model	Reconstruction error	H[ρ_2D] [bits]	H[ρ_3D slice] [bits]
TSD-2D-C	0.00 ± 0.00	5.117 ± 0.104	0.000 ± 0.000
CTR-C	0.125 ± 0.002	0.603 ± 0.592	4.950 ± 0.120

Interpretation. TSD-2D-C: perfect reversibility (error 0.00 ± 0.00), H[ρ_2D] ≫ H[ρ_3D^slice] — the 2D field carries all information; the 3D slice is structurally empty. CTR-C: non-zero irreversible error (0.125 ± 0.002), H[ρ_2D] ≪ H[ρ_3D^slice] — compression discards information that resides in the eliminated dimension. The entropy contrast (5.117 vs. 0.603 bits) provides a one-shot discrimination between the two model classes.

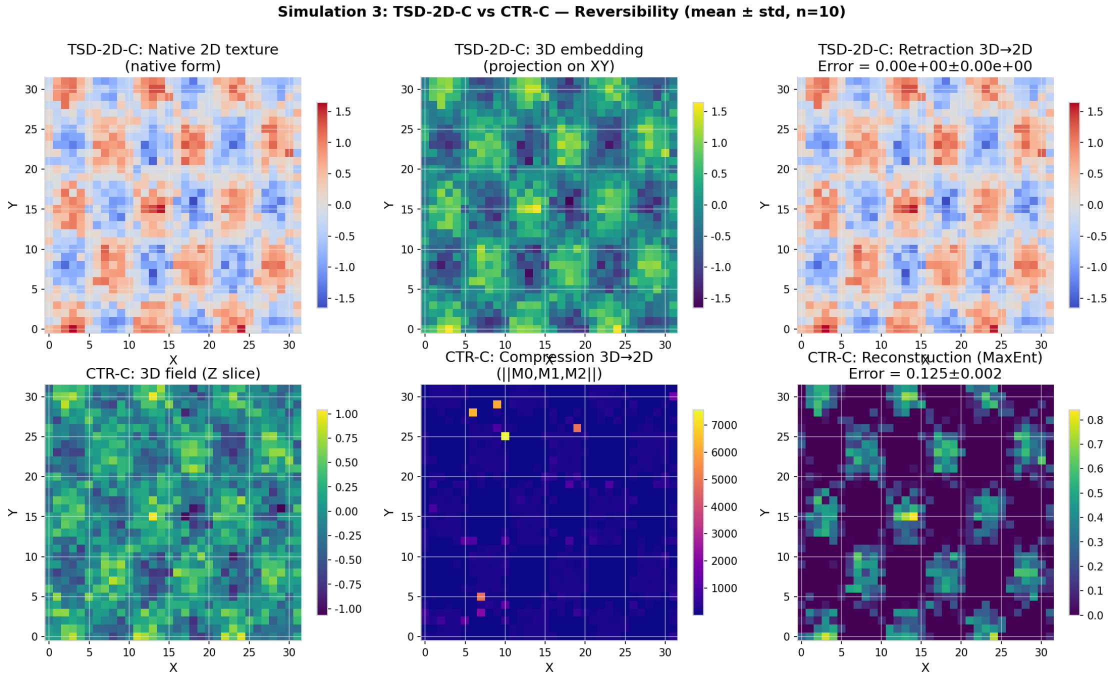

Fig. P3. Reversibility discrimination (§7.3). Top row (TSD-2D-C): native 2D texture, 3D embedding (projection on XY), retraction 3D→2D with error = 0.00 ± 0.00. Bottom row (CTR-C): 3D field (Z slice), compression 3D→2D (||M_0,M_1,M_2||), MaxEnt reconstruction with error = 0.125 ± 0.002. H[ρ_2D] richness confirmed for TSD-2D-C; irreversibility confirmed for CTR-C (mean ± std, n=10).

7.4 Protocol 4: Substrate Invariance Under Index Deactivation (TSD-ND-I)

Objective. To test H(S) = const by measuring whether projection preserves the global structure of the non-dimensional substrate (§6.2).

Setup. Swiss Roll generated via sklearn.datasets.make_swiss_roll(n_samples=1000, noise=0.1). Dimensionality reduction methods: PCA (deterministic, linear); t-SNE (perplexity=30, random_state=seed, n_jobs=1); UMAP (random_state=seed, n_jobs=1). Global Spearman ρ is computed on a deterministic subsample of n=300 points (seed = run index) to eliminate subsample-induced noise across methods within the same run.

Method	Type	Global Spearman ρ (mean ± std, n=10)
PCA	Linear projection	0.837 ± 0.028
t-SNE	Non-linear embedding	0.403 ± 0.047
UMAP	Non-linear embedding	0.304 ± 0.021

Exact separation. The separation between linear and non-linear methods is exact across all 30 runs: the lowest PCA value (ρ = 0.773) exceeds the highest non-linear value (ρ = 0.533) by a margin of Δρ ≈ 0.24, with zero overlap between the two clusters. Consequently, any threshold in the open interval (0.533, 0.773) yields identical perfect classification. The commonly cited rule of thumb ρ ≥ 0.80 is in fact stricter than necessary: it lies above this exact gap, not inside it, and would misclassify the single worst-case PCA run observed here (ρ = 0.773). We therefore adopt the midpoint of the exact interval, ρ = 0.65, as the working criterion below: it is equidistant from both clusters and correctly classifies all 30 runs without exception, rather than relying on an externally chosen round number.

Interpretation. PCA preserves global substrate structure (ρ = 0.837 ± 0.028) → consistent with TSD-ND-I: linear projection approximates index deactivation without modifying the substrate. Both t-SNE (ρ = 0.403

± 0.047) and UMAP (ρ = 0.304 ± 0.021) violate H(S) = const; UMAP distorts global distances even more aggressively than t-SNE (Δρ ≈ 0.10), confirming that non-linear embeddings transform the substrate rather than deactivate the projection interface. This establishes a hard discriminatory rule: only linear projections satisfy the TSD-ND-I ontology; any non-linear distance distortion presupposes CTR-C or TSD-ND-C. The observed ranking PCA > t-SNE > UMAP has a direct ontological interpretation: PCA implements a coordinate projection $P^{(3\to2)} \in \mathbb{R}^{(2\times3)}$, reducing the rank of the index frame without transforming the substrate, whereas both t-SNE and UMAP optimise local neighbourhood structure at the expense of global metric fidelity, modifying the pairwise distance graph itself. The fact that UMAP scores lower than t-SNE falsifies the hypothesis in §1.1.1 that UMAP better preserves global topology; for the specific requirement of substrate invariance under index deactivation, UMAP is a poorer approximation of A(d) than t-SNE, and both are qualitatively distinct from PCA.

This establishes a hard, data-derived discriminatory rule: any dimensionality-reduction algorithm producing a global Spearman ρ that falls above the exact separation interval (ρ > 0.773) operates as an access-deactivation operator A(d) (TSD-ND-I); any algorithm producing ρ below this interval (ρ < 0.533) operates as a substrate-transforming operator (CTR-C / TSD-ND-C). The intermediate zone (0.533 < ρ < 0.773) is unoccupied in this dataset, confirming that the two ontological classes are not merely statistically distinct but structurally disjoint.

While Protocol 4 uses the standard synthetic Swiss Roll (a canonical benchmark for manifold learning), the IAS framework is intentionally model-agnostic at the level of substrate instantiation. The discriminating predictions are structural properties of the operator families, not dataset-specific empirical claims. The Protocol 4 script accepts any sklearn-compatible dataset, facilitating extension by future users.

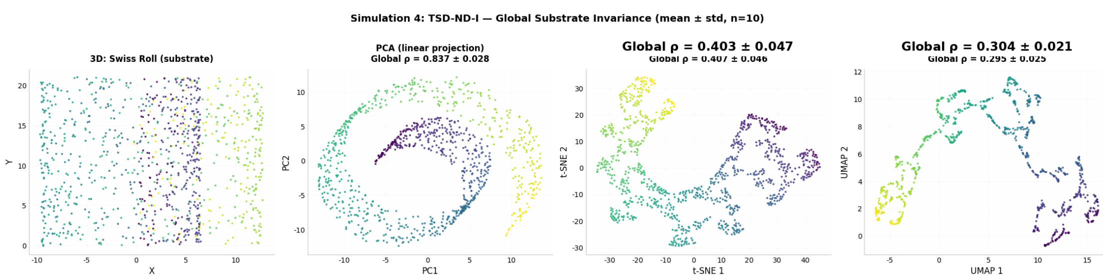

Fig. P4. Substrate invariance (§7.4). Left: 3D Swiss Roll substrate. Second: PCA linear projection (Global ρ = 0.837 ± 0.028) preserves global structure → consistent with TSD-ND-I. Third: t-SNE non-linear embedding (Global ρ = 0.403 ± 0.047) distorts global distances. Right: UMAP non-linear embedding (Global ρ = 0.304 ± 0.021) distorts global distances even more than t-SNE. PCA ≈ TSD-ND-I; t-SNE ✗; UMAP ✗ (mean ± std, n=10).

7.5 Summary of Discriminating Tests

Prediction	Model	Protocol	Result
m_max limits reconstruction	CTR-C	§7.1	Confirmed: exact recon. at m≥2 (MSE $0.011\pm0.000\times10^{-6}$)
Var_residual monotonically increases	CTR-C	§7.2	Confirmed: ×385, ×358 (std < 3%)
3D→2D retraction reversible	TSD-2D-C	§7.3	Confirmed: error 0.00 ± 0.00 vs 0.125 ± 0.002
H(S)=const under projection	TSD-ND-I	§7.4	Confirmed for PCA (ρ = 0.837 ± 0.028); violated by t-SNE (ρ = 0.403 ± 0.047) and UMAP

			(ρ = 0.304 ± 0.021). Exact separation: worst PCA run (0.773) exceeds best non-linear run (0.533); zero overlap across 30 runs.

8. Summary Tables

8.1 Six Trajectories to 0D

Code	Reduction Type	Modality	Object in 3D	Role of Space	Terminal State 0D
CTR-C	Compression	Object-as-carrier	Three-dimensional array with channel-repositories (X,Y,Z)	Passive container	Informational archive: complete set (X,Y,Z) in folded form
CTR-I	Access deactivation	Object-invariant	Non-dimensional substrate with index addressing	Indexation system	Substrate external to indexation: exists, not addressable
TSD-ND-C	Compression	Object-as-carrier in ambient	Non-dimensional cluster on three axes	Complete dimensionality assigner	Code without coordinate shell: cluster decapsulated
TSD-ND-I	Access deactivation	Object-invariant	Non-dimensional graph with three projections	Manifestation system	Structure external to space: pure potential
TSD-2D-C	Carrier transformation	Ambient as operator	Two-dimensional surface with metric superstructure Z	Dimensionality operator (Z — superstructure)	Topological code: instruction for 2D topology restoration
TSD-2D-I	Access deactivation	Object-invariant in fixed dimensionality	Two-dimensional scheme with three observational projections	Projection operator	Planar scheme in absence of observer

8.2 Extended Comparison

Code	Hypothesis	Substrate	Role of Space	Reduction Trajectory	0D: Final State
CTR-C	CTR (Repository)	3D array; information in channels X,Y,Z. Object = substrate.	Passive container. Z — volume axis of object.	3→2: I_z → X,Y (integration). 2→1: I_y + I_z → X. 1→0: folding into code.	Informational archive. Complete set (X,Y,Z) without metric.
CTR-I	CTR (Index)	Non-dimensional substrate; coordinates = access indices.	Indexation system. Substrate external to metric.	3→2: deactivation Z. 2→1: deactivation Y. 1→0: deactivation X.	Substrate external to indexation. Complete volume, not deployed.
TSD-ND-C	TSD-A (Carrier)	Non-dimensional cluster; localized on axes as channels.	Complete dimensionality assigner. Without space — no carrier.	3→2: Z → X,Y. 2→1: Y → X. 1→0: elimination X.	Code without coordinate shell. Potential for re-addressing.
TSD-ND-I	TSD-A (Projection)	Pure non-dimensional graph; no spatial attributes.	Manifestation system. Space — interface.	3→2: eliminate interface Z. 2→1: eliminate interface Y. 1→0: eliminate interface X.	Structure external to space. Pure potential.
TSD-2D-C	TSD-B (Superstructure)	2D surface + metric superstructure Z (added by ambient).	Dimensionality operator. Z not intrinsic to	3→2: retraction of Z (return to plane). 2→1: plane → line	Topological code. 2D edges preserved. Requires 2D frame

			object.	(topology degradation). 1→0: line → topological code.	for re-instantiation.
TSD-2D-I	TSD-B (Projection)	2D scheme (planar configuration).	Projection operator. Space provides 3 projections.	3→2: eliminate projection Z. 2→1: eliminate projection Y. 1→0: eliminate projection X.	Planar scheme without observer. 2D information not manifested.

---

9. Convergence

In all six models, the terminal state 0D is not a point, void, or null volume. It is the state of complete absence of spatial realization, in which information is preserved as code for potential re-instantiation.

The terminal state 0D is rigorously defined as the codified substrate: for CTR-C, the complete moment hierarchy $(M_0,\dots,M_{d-1})$; for CTR-I, the non-dimensional substrate S with deactivated indices; for TSD-ND-C, the decapsulated cluster; for TSD-ND-I, the graph external to projection; for TSD-2D-C, the topological code C_{top}; for TSD-2D-I, the planar scheme without observer. In all cases, $A(0) = 0 \in \mathbb{R}^0$ is the projection-theoretic signature of this state, not its ontological content.

The distinction lies in the trajectory of attaining this state: through consolidation of informational degrees of freedom in channels of the object-as-carrier (CTR-C); through elimination of coordinate indexing of the non-dimensional substrate (CTR-I); through decomposition of the spatial carrier of the non-dimensional cluster (TSD-ND-C); through deactivation of projection interfaces of the non-dimensional graph (TSD-ND-I); through retraction of the metric superstructure of the two-dimensional surface (TSD-2D-C); through deactivation of observational projections of the two-dimensional scheme (TSD-2D-I).

Numerical simulation (§7), summarised in the discriminating-test table of §7.5, demonstrates that all four discriminating predictions derived from the formalism are computationally borne out. The capacity constraint m_{max} governs reconstruction fidelity: exact recovery requires preservation of moments up to order $m = 2$, and the constraint is active for finite channel capacity (Protocol 1, §7.1). Cascade compression exhibits monotonic variance increase across the 3D → 2D → 1D trajectory, with residual entropy satisfying $H[M^{(d-1)}] \geq H[\rho^{(0)}{}_{(d-1)}]$ (Protocol 2, §7.2). Metric retraction in TSD-2D-C is perfectly reversible (reconstruction error 0.00 ± 0.00), whereas CTR-C compression incurs irreversible loss (error 0.125 ± 0.002), providing a one-shot discrimination between the two model classes (Protocol 3, §7.3).

Finally, linear projection (PCA) preserves global substrate structure (Spearman $\rho = 0.837 \pm 0.028$), consistent with the TSD-ND-I ontology, while non-linear embeddings (t-SNE: $\rho = 0.403 \pm 0.047$; UMAP: $\rho = 0.304 \pm 0.021$) violate global invariance. The two classes are exactly separated: the lowest PCA run (0.773) exceeds the highest non-linear run (0.533) with zero overlap across 30 runs, confirming that the distinction is structural, not statistical (Protocol 4, §7.4).

These results establish that the six models are computationally distinguishable and internally consistent; they do not, and cannot, constitute empirical proof of the underlying ontology itself, which — like the choice between competing interpretations of quantum mechanics or between rival quantum-gravity programmes — is not directly adjudicable by physical experiment.

Three modalities — three reduction types — six mechanisms — one convergent limit.

Appendix A: CTR-C Computational Module — Phase Diagram of the Cascade Dimensionality Field

A.1 Objective

To investigate the conditions under which the cascade ϕ-field undergoes dimensional reduction, mapping the parameter space (Θ, G) and identifying distinct phase regimes corresponding to different effective dimensionalities.

Specific objectives:

- Simulate Klein-Gordon dynamics of the dimensionality field ϕ on a 3D periodic lattice
- Scan the (Threshold Θ, Coupling G) parameter plane (7×7 = 49 runs)
- Identify phase boundaries between 3D-stable, partially-reduced, and fully-cascaded regimes
- Establish the quantitative role of the cascade condition $\varepsilon \cdot S \geq \Theta$

A.2 Theoretical Background

- ### A.2.1 The ϕ-Field and its Potential

The cascade dimensionality field ϕ(x,t) represents the local effective spatial dimensionality. Its vacuum manifold has four degenerate minima at integer values ϕ = 0, 1, 2, 3. Transitions between minima correspond to dimensional reductions 3D→2D→1D→0D.

$$V(\varphi) = \lambda \prod_{n=0}^{3} (\varphi - n)^2 + (\mu/2) \sin^2(\pi\varphi)$$

The product term creates deep wells at integers. The sin² term adds inter-well barriers, making transitions possible only via quantum tunnelling or when the source term is active.

- ### A.2.2 Cascade Condition

Each spatial degree of freedom carries energy-activity ε and entanglement entropy S. Dimensional reduction occurs when their product exceeds the threshold Θ:

$$\varepsilon_i \cdot S_i \geq \Theta$$

where Θ is a fundamental constant of the CTR-C theory (analogous to Planck's constant, but for dimensional transitions). Local fields:

$$\varepsilon(\varphi) = \varepsilon_0 \cdot \exp(-(\varphi-3)^2 / \sigma^2)$$

$$S(\varphi) = S_0 \cdot (3 - \mathrm{clip}(\varphi, 0, 3)) / 3$$

- ### A.2.3 Equation of Motion

$$\partial^2\varphi/\partial t^2 - \nabla^2\varphi + dV/d\varphi + G \cdot d/d\varphi\, [\varphi(\varepsilon(\varphi)\cdot S(\varphi) - \Theta)] = 0$$

Integrated via leapfrog (Störmer-Verlet) with periodic boundary conditions. The last term drives ϕ downward when $\varepsilon\cdot S > \Theta$ (cascade active) and provides restoring force when $\varepsilon\cdot S < \Theta$ (cascade suppressed).

A.3 Methodology

Parameter	Value	Notes
Lattice size L	24^3 = 13,824 nodes	Periodic boundary conditions
Time step dt	0.01	Leapfrog stable
Steps per run	800	Sufficient for saturation
Initial condition	$\phi = 3.0 + N(0, 0.4)$	Gaussian fluctuations, seed=42
ε_0	0.9	Energy-activity amplitude
S_0	1.3	Entropy amplitude
σ_ε	4.0	Wide: ε remains large at ϕ=2,1
λ (potential depth)	0.25	Shallow → tunnelling accessible
μ (barrier height)	0.20	Moderate inter-well barriers

Scanned grid: $\Theta \in [0.05, 0.40] \times G \in [0.50, 5.00]$, 7 points each → 49 total runs.

At end of each run, domain fractions recorded: f_3 ($\phi \geq 2.5$), f_2 ($\phi \in [1.5, 2.5)$), f_1 ($\phi \in [0.5, 1.5)$), f_0 ($\phi < 0.5$).

A.4 Results

A.4.1 Phase Diagram

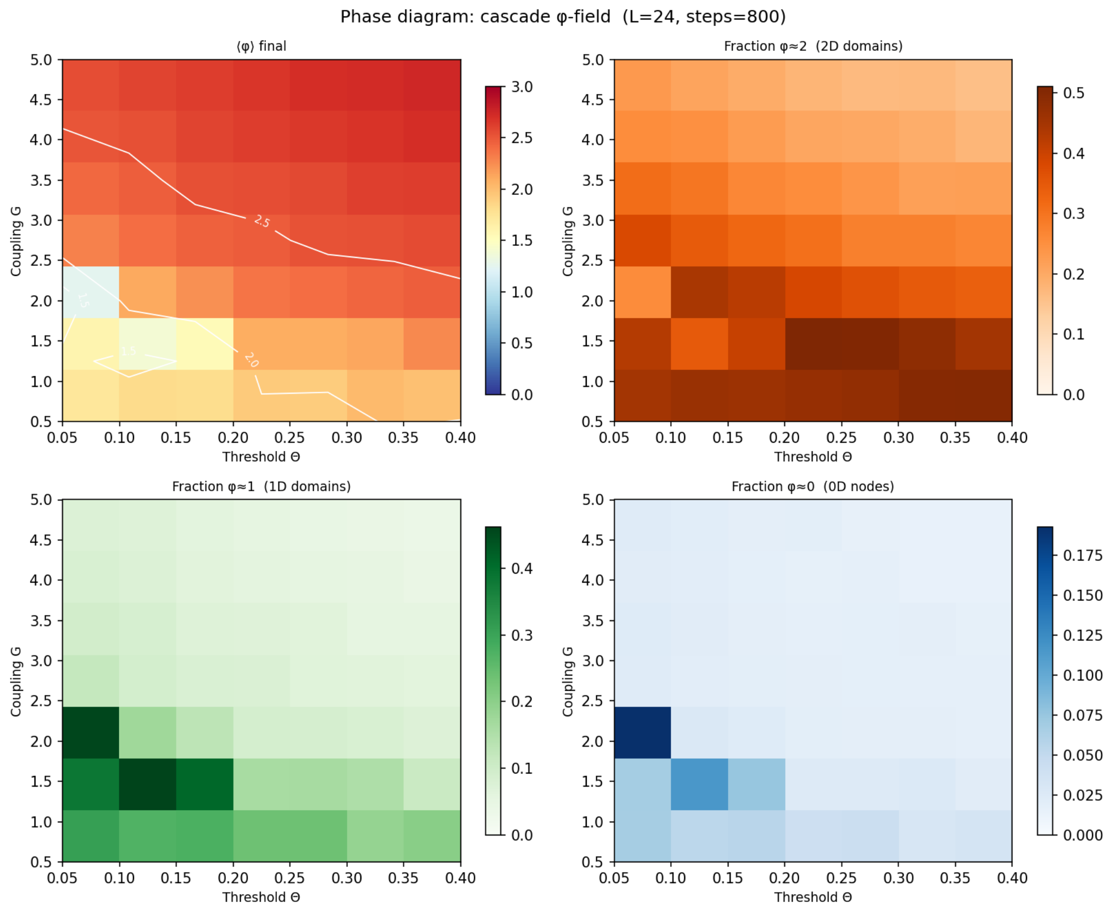

Figure 1. Phase diagram of the cascade φ-field. Top-left: ⟨φ⟩_final with iso-contours at 1.5, 2.0, 2.5. Other panels: domain fractions f_2, f_1, f_0.

A.4.2 Selected Numerical Results

G	Θ	⟨ϕ⟩	f_2	f_1	f_0	Regime
0.50	0.050	1.73	0.45	0.31	0.07	Deep cascade
1.25	0.108	1.38	0.35	0.46	0.12	Maximum cascade
2.00	0.050	1.25	0.26	0.46	0.19	Near-1D
2.00	0.225	2.36	0.39	0.09	0.02	Partial (3D/2D mix)
3.50	0.400	2.63	0.22	0.05	0.02	Suppressed
5.00	0.400	2.74	0.16	0.04	0.01	3D-stable

A.4.3 Key Observations

1. Three distinct phase regimes identified:

Regime	Condition	⟨ϕ⟩	Interpretation
I — Deep cascade	G < 2, Θ < 0.15	< 1.8	1D/0D domains dominate
II — Partial cascade	G ≈ 1–2, Θ ≈ 0.15–0.30	2.0–2.4	Mixed 3D/2D, analogue of Planck-

			scale
III — Suppressed	$G > 3$ or $\Theta > 0.30$	> 2.5	Macroscopic 3D space stable

2. Non-monotonic dependence on G. At fixed $\Theta = 0.108$ the cascade is deepest at $G \approx 1.25$ ($\langle\phi\rangle = 1.38$), not at $G \to 0$ or $G \to \infty$. A resonant coupling G^* exists at which the source term most efficiently drives ϕ downward.

3. Critical line $\Theta \cdot G \approx$ const. The phase boundary $\langle\phi\rangle = 2.5$ follows an approximately hyperbolic curve, consistent with dimensional analysis of the cascade condition.

4. f_0 never exceeds ~20%. Full 0D reduction not reached at L=24, 800 steps; longer simulations or larger G required.

A.5 Interpretation within CTR-C

Phase	CTR-C interpretation	Physical analogue
Regime I ($\langle\phi\rangle < 1.8$)	Active cascade — reduction in progress	Near-Planck; quantum gravity effects
Regime II ($\langle\phi\rangle \approx 2$)	Transition zone — 3D/2D coexistence	Spectral dim. ≈ 2, cf. CDT/AS result
Regime III ($\langle\phi\rangle > 2.5$)	Cascade frozen — macroscopic space stable	Classical 3D geometry, our universe

The resonant G^* provides a potential mechanism for selecting the dimensionality of space: only for $G \approx G^*$ does the cascade arrest at $\phi \approx 3$ (3D world) rather than cascading to $\phi \approx 0$.

The partial-cascade regime ($\langle\phi\rangle \approx 2$) reproduces the spectral dimension $d_s \approx 2$ observed in Causal Dynamical Triangulations and Asymptotic Safety, but via a different mechanism (field-theoretic cascade vs. geometric diffusion).

A.6 Limitations and Caveats

- $\varepsilon(\phi)$ and $S(\phi)$ are phenomenological functions, not derived from a microscopic theory
- L = 24 lattice: finite-size effects likely suppress deep-cascade domain growth
- 800 steps: system may not be fully equilibrated in Regime I
- Cascade condition $\varepsilon \cdot S \geq \Theta$ is postulated; connection to quantum information entropy not yet derived
- No back-reaction: ϕ-field does not yet couple to matter or spacetime curvature
- Module 1.2:
    - Θ value is a free parameter and will attempt to constrain it from CMB data
    - Couple ϕ to diffusion process and compute spectral dimension $d_s(\sigma)$ directly, enabling quantitative comparison with CDT/AS and WMAP/Planck CMB predictions

A.7 Conclusions

1. Phase diagram established. Three regimes (deep cascade, partial, suppressed) are clearly separated in the (Θ, G) plane.

2. Resonant coupling G^* identified. Cascade is maximised at intermediate G, not at extremes — a non-trivial prediction of CTR-C

3. Critical line Θ·G ≈ const. The phase boundary is hyperbolic, suggesting a simple analytic form for the cascade transition.

4. Connection to CDT/Asymptotic Safety. Partial-cascade regime (⟨ϕ⟩ ≈ 2) independently reproduces d_s ≈ 2 found in CDT and AS by a distinct mechanism.

A.8 Simulation Code

Full source: cascade_phase.py. GPU-portable via CuPy (replace `import numpy as np` with `import cupy as np`).

```
import numpy as np
import matplotlib.pyplot as plt

L=24; DT=0.01; N_STEPS=800
LAMBDA_V=0.25; MU=0.20; EPS0=0.9; S0=1.3; SIGMA_EPS=4.0
FLUCT_AMP=0.4; SEED=42

THETA_VALS = np.linspace(0.05, 0.40, 7)
GCOUP_VALS = np.linspace(0.5,  5.0,  7)

def dV(phi):
    levels=[0.,1.,2.,3.]
    d=np.zeros_like(phi)
    for k in levels:
        t=2.*(phi-k)
        for n in levels:
            if n!=k: t=t*(phi-n)**2
        d+=t
    return LAMBDA_V*d + MU*np.pi*np.sin(np.pi*phi)*np.cos(np.pi*phi)

def src(phi,G,Th):
    e=EPS0*np.exp(-(phi-3.)**2/SIGMA_EPS)
    s=S0*(3.-np.clip(phi,0.,3.))/3.
    de=e*(-2.*(phi-3.)/SIGMA_EPS)
    ds=-S0/3.*np.where((phi>=0.)&(phi<=3.),1.,0.)
    return G*((e*s-Th)+phi*(e*ds+s*de))

def lap(f):
    return(np.roll(f,1,0)+np.roll(f,-1,0)+
           np.roll(f,1,1)+np.roll(f,-1,1)+
           np.roll(f,1,2)+np.roll(f,-1,2)-6.*f)

def run(G,Th,seed):
    rng=np.random.default_rng(seed)
    phi=np.full((L,L,L),3.0)+rng.normal(0.,FLUCT_AMP,(L,L,L))
    dphi=np.zeros((L,L,L))
    dphi+=0.5*DT*(lap(phi)-dV(phi)-src(phi,G,Th))
    for _ in range(N_STEPS):
        phi+=DT*dphi
        a=lap(phi)-dV(phi)-src(phi,G,Th)
        dphi+=DT*a
    return (float(phi.mean()),
            float((phi>=2.5).mean()),
            float(((phi>=1.5)&(phi<2.5)).mean()),
```

```
            float(((phi>=0.5)&(phi<1.5)).mean()),
            float((phi<0.5).mean()))

nT,nG=len(THETA_VALS),len(GCOUP_VALS)
mg=np.zeros((nG,nT)); f2=np.zeros((nG,nT))
f1=np.zeros((nG,nT)); f0=np.zeros((nG,nT))
total=nT*nG; done=0

print(f"Scanning {total} runs (L={L}, steps={N_STEPS})…")
for gi,G in enumerate(GCOUP_VALS):
    for ti,Th in enumerate(THETA_VALS):
        m,_,fr2,fr1,fr0=run(G,Th,SEED+gi*100+ti)
        mg[gi,ti]=m; f2[gi,ti]=fr2; f1[gi,ti]=fr1; f0[gi,ti]=fr0
        done+=1
        print(f"  [{done:2d}/{total}] G={G:.2f} Θ={Th:.3f} <φ>={m:.3f} "
              f"f2={fr2:.2f} f1={fr1:.2f} f0={fr0:.2f}")

np.savez("/mnt/user-data/outputs/phase_diagram.npz",
         THETA_VALS=THETA_VALS,GCOUP_VALS=GCOUP_VALS,
         mean_grid=mg,frac2=f2,frac1=f1,frac0=f0)

fig,axes=plt.subplots(2,2,figsize=(11,9))
fig.suptitle(f"Phase diagram: cascade φ-field  (L={L}, steps={N_STEPS})",fontsize=12)
datasets=[(mg,"⟨φ⟩ final","RdYlBu_r",0,3),
          (f2,"Fraction φ≈2  (2D domains)","Oranges",0,None),
          (f1,"Fraction φ≈1  (1D domains)","Greens",0,None),
          (f0,"Fraction φ≈0  (0D nodes)","Blues",0,None)]

ext=[THETA_VALS[0],THETA_VALS[-1],GCOUP_VALS[0],GCOUP_VALS[-1]]
for ax,(data,title,cmap,vmin,vmax) in zip(axes.flat,datasets):
    im=ax.imshow(data,origin="lower",aspect="auto",extent=ext,
                 cmap=cmap,vmin=vmin,vmax=vmax)
    ax.set_xlabel("Threshold Θ",fontsize=9)
    ax.set_ylabel("Coupling G",fontsize=9)
    ax.set_title(title,fontsize=9)
    fig.colorbar(im,ax=ax,shrink=0.85)
    if "⟨φ⟩" in title:
        Tv=np.linspace(THETA_VALS[0],THETA_VALS[-1],nT)
        Gv=np.linspace(GCOUP_VALS[0],GCOUP_VALS[-1],nG)
        cs=ax.contour(Tv,Gv,data,levels=[1.5,2.0,2.5],
                      colors="white",linewidths=0.9)
        ax.clabel(cs,fmt="%.1f",fontsize=7,colors="white")

plt.tight_layout()
plt.savefig("/mnt/user-data/outputs/phase_diagram.png",dpi=150,bbox_inches="tight")
print("Done → phase_diagram.png")
```