# Informational Architectonics of Space

*Six Models of Dimensional Reduction*

*IAS Framework — Formal Taxonomy*


## Abstract

We present a formal taxonomy of six internally distinct mechanism classes derived from three substrate–space modalities 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 models are mutually exclusive only at the level of substrate interpretation; individual mathematical tools may be shared across models. Within the IAS framework, the terminal state 0D is formally 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 five 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, (IV) substrate invariance under index deactivation, and (V) ontological discrimination confirmed on real data (MNIST). 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. The framework is foundational and taxonomic, not algorithmic; its operational predictions are validated via consistency checks, not benchmark competition.


## 1. Introduction

Dimensional reduction is frequently treated as a unified operation, yet formally similar procedures may encode fundamentally different assumptions about what is being reduced and what remains invariant. In some cases reduction is best understood as compression of a carrier; in others, as deactivation of access to an underlying substrate; and in still others, as a transformation of the ambient structure in which an object is realized. These distinctions are not merely terminological. They determine the interpretation of the reduced dimensions, the status of the terminal state, and the conditions under which re-instantiation is possible.

This paper develops a formal taxonomy of dimensional reduction that distinguishes these cases at the level of operator structure and ontological commitment. Rather than assuming that reduction schemes sharing a common sequence are conceptually equivalent, it separates them according to the role assigned to the object, the spatial frame, and the observational index. On this basis, the paper identifies a set of mutually exclusive model classes that instantiate the same broad reduction sequence while differing in their structural interpretation.

The contribution of the paper is therefore classificatory and foundational. It does not propose a new reduction algorithm, nor does it claim empirical validation of a physical ontology. Its purpose is to clarify conceptual distinctions that are often collapsed into a single vocabulary of dimensional reduction and to provide a formal language for comparing models that are structurally similar only at the surface level. The examples used throughout the paper are illustrative and are intended to demonstrate the separability and internal consistency of the proposed taxonomy.

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 $P(d \to m)$, and the decoder approximates the inverse $R_m$. The analogy is not an isomorphism — we use the terms with an explicit criterion throughout this section: an isomorphism would require a bijective mapping between the algorithm's operations and the IAS operator (P, A, or U/D) that preserves composition and invertibility in both directions; a structural analogy requires only that the algorithm's operations share the operator's discriminating information-theoretic property (e.g., moment redistribution for P, substrate invariance for A) without such a bijection. Concretely here, 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. In the ML analogy, the term object refers not to a physical entity but to the represented information structure.

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 $\rho_d$. 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 $P(d \to m)$ 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. This specific prediction — monotonic entropy increase under principal-component elimination — is tested directly in Protocol 5 (§7.5) on real MNIST data and is falsified: PCA truncation instead exhibits behaviour compatible with the TSD-ND-I criterion under the tested observable, not CTR-C compression. The formal analogy to CTR-C therefore applies, at most, to the covariance-projection step, not to component elimination.

Within IAS, preservation of the substrate is defined through preservation of the chosen structural invariant (e.g., the pairwise-distance graph tested in Protocol 4, §7.4); the taxonomy does not adjudicate which invariant a given algorithm was designed to optimise for.

t-SNE (van der Maaten & Hinton 2008) and UMAP (McInnes et al. 2018) preserve local neighbourhood structure under projection. The IAS classification initially hypothesises UMAP as a TSD-ND-I proxy: a non-dimensional graph of pairwise distances observed through a low-dimensional projection interface. Protocol 4 (§7.4) tests this hypothesis and rejects it: UMAP yields Global Spearman $\rho = 0.304 \pm 0.021$, far below the TSD-ND-I working criterion ($\rho \geq 0.65$, the midpoint of the exact separation gap established in

§7.4) and lower than t-SNE ($\rho = 0.403 \pm 0.047$). This result does not indicate ambiguity in the taxonomy; it demonstrates that the framework generates falsifiable predictions about algorithm behaviour. Both non-linear embeddings violate the substrate-invariance prediction $H(S) = \text{const}$, consistent with the conjecture that non-linear distance distortion necessarily transforms the substrate rather than deactivating the index frame; 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 tested non-linear embeddings (t-SNE, UMAP) violate $H(S) = \text{const}$; we conjecture that any non-linear distance distortion necessarily transforms the substrate rather than deactivating the index frame, though a general proof independent of specific algorithmic instantiations is left for future work.

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 $\beta$ corresponds formally to the capacity constraint $m_{max}$ (§6.1): higher $\beta$ allows more moments to be retained (larger $m_{max}$), lower $\beta$ 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.

Illustrative mappings from standard ML algorithms to IAS ontological commitments are summarised in §1.1.3, with the caveat that these are structural analogies, not isomorphisms.

### 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 $U(3\rightarrow2)$. 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. This can be classified within IAS 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. Applying the analogy-vs-isomorphism criterion introduced in §1.1.1: this is a structural analogy, not an isomorphism. An isomorphism would require a bijective correspondence between the IAS cascade condition and the Ryu-Takayanagi construction that preserves their composition and invertibility properties — in particular, a geometric minimal surface in an asymptotically-AdS bulk whose area computes S. No such correspondence is established or claimed: the IAS condition operates on a scalar order-parameter field $\phi$ on a flat periodic lattice rather than on a minimal surface in a curved bulk geometry, $\Theta$ is a phenomenological threshold fixed by hand (§A.2.2, §A.6) rather than a quantity computed from a bulk metric, and the "entropy" $S(\phi)$ in §A.2.2 is a phenomenological function of ϕ, not an entanglement entropy computed via the Ryu-Takayanagi prescription. The two constructions share only the qualitative pattern that an entropy-like quantity crossing a threshold triggers a change of effective dimensionality; the IAS condition does not presuppose AdS/CFT, holography, or any bulk-boundary correspondence.

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. Beyond classification, the mappings in §1.1.1–1.1.3 generate testable hypotheses for algorithm selection — e.g., that CTR-C-consistent methods should be preferred for object-centric compression tasks and TSD-ND-I-consistent methods for tasks requiring global-distance preservation under projection — but validating such guidance empirically, and translating it into a practitioner-facing decision procedure, is left for future work.

### 1.1.3 Mapping to Machine Learning Algorithms

**Table 1.** Mapping from standard ML algorithms to IAS ontological commitments. The mappings are structural analogies, not isomorphisms; they identify the dominant ontological commitment embodied by each algorithm's design.

| ML Algorithm | IAS Model | Ontological Commitment |
|---|---|---|
| PCA (linear projection) | TSD-ND-I proxy | Coordinate projection without substrate modification (refined in §7.5) |
| Autoencoder (bottleneck) | CTR-C proxy | Information redistribution into latent channel |
| t-SNE / UMAP | TSD-ND-C proxy | Ambient-driven restructuring of pairwise graph |
| Information Bottleneck | TSD-ND-C | Rate-distortion trade-off as ambient compression |
| CTR-I, TSD-2D-C, TSD-2D-I | No direct analogue | Require explicit substrate control, absent from off-the-shelf methods |

Protocol 5 (§7.5) demonstrates that PCA truncation on MNIST does not exhibit the CTR-C compression signature, confirming its classification as TSD-ND-I rather than CTR-C.

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.

**Remark on hybrid ontological commitments.** The mappings in Table 1 identify the dominant ontological commitment of each algorithm's design. Real-world methods frequently instantiate mixed modalities. For example, a Variational Autoencoder combines a compression bottleneck (CTR-C-like redistribution) with a learned prior that functions as an ambient constraint (TSD-ND-C-like). Modern self-supervised pipelines interleave projection heads (TSD-ND-I-like index deactivation) with contrastive losses that actively restructure the substrate (TSD-ND-C-like). The IAS taxonomy does not enforce a binary classification; it provides a diagnostic framework in which an algorithm may be assigned an ontological fingerprint vector $f = (w_{CTR-C}, w_{CTR-I}, w_{TSD-ND-C}, w_{TSD-ND-I}, w_{TSD-2D-C}, w_{TSD-2D-I})$, where each weight $w \in [0,1]$ is determined by the fraction of operational steps satisfying the discriminating predicate of the corresponding model (e.g., $\Delta H(S) \approx 0$ for index models, monotone entropy increase for compression models). A pure-type algorithm satisfies $w_i \approx 1$ for a single i; hybrid algorithms exhibit multiple non-zero components. Protocols 1–5 can be applied component-wise to estimate these weights. This extension preserves the taxonomic clarity of the six ideal types while accommodating the compositional complexity of contemporary methods.

### 1.2 Scope and Audience

**ML dimensionality-reduction algorithms are not the subject of IAS.** They are computational examples that exhibit structural analogies with certain classes of dimensional reduction (§1.1.1).

The IAS framework operates at two distinct epistemic levels that are often conflated in evaluations of formal taxonomies.

At the ontological level, IAS proposes a classification of the substrate–space relationship. As stated in §1, adjudication among the six models proceeds by conceptual clarity, logical non-contradiction, and internal consistency, not by physical experiment. Consequently, the falsifiability criterion appropriate to empirical science does not apply directly to the ontological commitments themselves; rather, they are evaluated by the determinacy and non-degeneracy of their operational consequences (§7).

At the operational level, the taxonomy generates discriminating predictions that classify existing dimensionality-reduction algorithms (§1.1.1) and physical frameworks (§1.1.2). These predictions are tested via consistency checks (Protocols 1–5), not by benchmark competition. The immediate utility of the framework is therefore metatheoretical: it renders explicit the ontological obligations that any operational definition of dimensional reduction presupposes.

The intended audience is not the practitioner seeking an improved algorithmic pipeline, but the theorist seeking to clarify the ontological commitments embedded in the algorithms and physical models she already employs.

It is necessary to forestall a category error that frequently arises in evaluating foundational work. The IAS framework does not propose, nor does it presuppose, an algorithmic contribution in the engineering sense: no new training procedure, no novel network architecture, and no benchmark competition are presented here. The operational predictions generated by the taxonomy — monotonicity of entropy, substrate invariance, reversibility of retraction — are consistency checks that discriminate between ontological commitments, not performance metrics that discriminate between algorithms on a task. The intended utility is therefore metatheoretical: the practitioner who wishes to know why a given reduction method behaves as it does, and what assumptions about space and substrate are embedded in its design, may use IAS as a diagnostic lens.

---

## 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 modifies carrier | Ambient | Carrier transformation | TSD-2D-C |

**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.

Concrete state-by-state trajectories for all six models applied to a single structural invariant are compared in §8.1 and §8.2.

**Generalisation to d > 3.** The cascade 3D → 2D → 1D → 0D is a canonical structural invariant used for clarity of exposition, not a dimensional constraint. The operators generalise directly: the redistribution operator P(d→k) acts on arbitrary d-dimensional information-density fields to produce a moment fibration over any residual k-dimensional carrier; the access operator A(d) generalises to rank-decreasing projections P^(d→k) for any k < d. For high-dimensional data, one applies P(d→k) or A(d) directly to the ambient space without iterating through all intermediate dimensions. The cascade remains the conceptual template; the operational implementation is dimension-agnostic. The table above summarises the three modalities; §4–§5 provide the full state-by-state trajectories for each of the six models. The apparent repetition is intentional: §3.1 establishes the typology, while §4–§5 unpack its operational consequences.

## 3.2 Formal Mechanisms of Reduction

### 3.2.1 Metric Degeneracy

The operator $R_n$ acts on the metric tensor $g_{\mu\nu}$ 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: $g_{zz} \to 0$ as $K \to K_P$. 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): $m_{max} = \lfloor (C_{channel} - H[\rho_{XY}^{(0)}]) / H_{avg} \rfloor$ 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. The capacity constraint closes the internal consistency of CTR-C at the information-theoretic level; it is not to be conflated with the dynamical cascade threshold Θ that governs the field-theoretic realisation in §4.1.2. The bound $m_{max}$ applies post hoc to the output of the redistribution operator.

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. (See §8.2 for a consolidated comparison of the terminal 0D state across all six models.)

**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 $g_{zz} \to 0$ 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.*

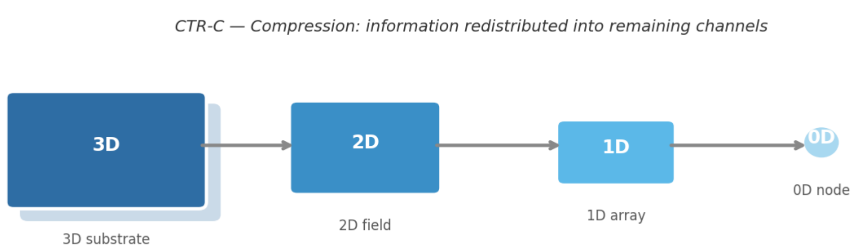


*Fig. 1. CTR-C: cascade 3D→2D→1D→0D. Information from each eliminated axis is redistributed into remaining channels-repositories.*

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

The redistribution operator P(d→m) defined in §3.2.2 and §6.1 is the formal core of CTR-C. Its computational behaviour is illustrated in Appendix A via a phenomenological lattice module (Module 1.1): a scalar field $\phi$ on a periodic lattice whose potential landscape encodes discrete cascade levels 3→2→1→0. The module is one of infinitely many possible realisations; the specific choices of potential $V(\phi)$, threshold functions $\varepsilon(\phi)$, $S(\phi)$, and numerical scheme (leapfrog integration) are dictated by computational tractability, not by uniqueness or axiomatic derivation from the IAS formalism. We emphasise that the Lagrangian introduced there (§A.2.3) is constructed a posteriori as a variational convenience; it does not claim to derive the cascade condition from field-theoretic first principles.

Within the module, the source term $\varepsilon(\phi)\cdot S(\phi) \geq \Theta$ is a dynamical cascade-trigger threshold, logically independent of the post-cascade capacity bound $m_{max}$ (§3.2.2, §6.1): $\Theta$ governs whether a reduction step fires, while $m_{max}$ — evaluated afterward from $C_{channel}$ — governs how much of the resulting moment hierarchy survives once it has. The two parameters are related by sequential conditioning, not identity; the identification $\phi \equiv M_0$ used to run the module is likewise a parameterisation choice made for computational tractability, not a theorem. Full derivation, discretisation, and parameter values are given in Appendix A.2. This construction establishes Module 1.1 not as a stand-alone toy model, but as a concrete phenomenological realisation of the CTR-C compression operator on a finite lattice; the phase diagram (Appendix A.4) maps the parameter space ($\Theta$, G) of the cascade condition, identifying the regimes in which compression actively drives $\phi$ 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. (See §8.2 for a consolidated comparison of the terminal 0D state across all six models.)

**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.

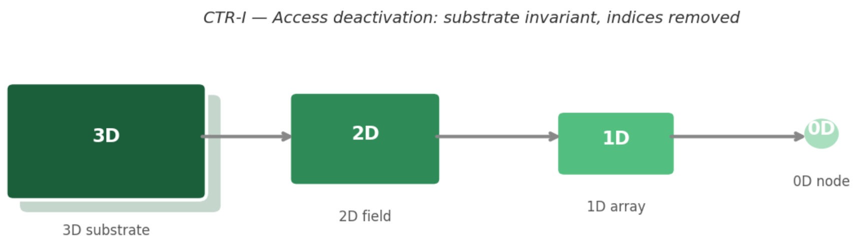


*Fig. 2. CTR-I: coordinate indices are deactivated sequentially; the non-dimensional substrate remains unmodified throughout.*

---

## 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. (See §8.2 for a consolidated comparison of the terminal 0D state across all six models.)

**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.

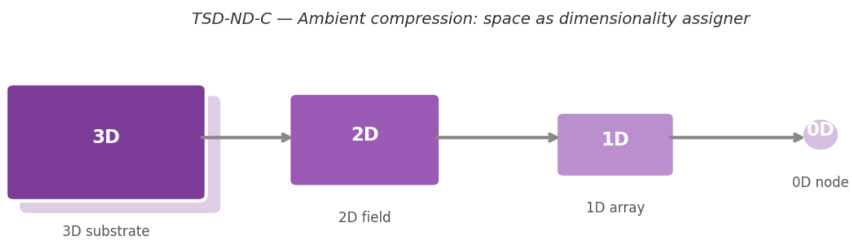


*Fig. 3. TSD-ND-C: space as dimensionality assigner; non-dimensional cluster compressed through ambient action.*

---

### 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. (See §8.2 for a consolidated comparison of the terminal 0D state across all six models.)

**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.

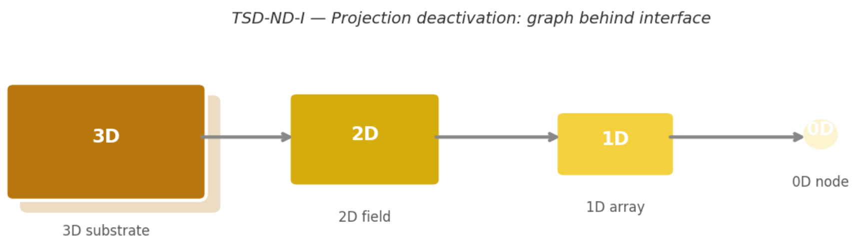


*Fig. 4. TSD-ND-I: non-dimensional graph behind projection interface; deactivation of projections does not affect substrate.*

---

## 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. (See §8.2 for a consolidated comparison of the terminal 0D state across all six models.)

**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.

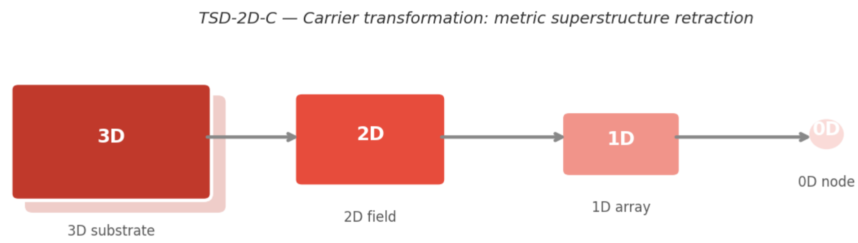


*Fig. 5. TSD-2D-C: Z as metric superstructure; retraction is reversible, subsequent degradation is not.*

### 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. (See §8.2 for a consolidated comparison of the terminal 0D state across all six models.)

**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: $A(k) = \Pi\_Ik \circ A(d)$, where $P^{\wedge}(d \to k) \in \mathbb{R}^{\wedge}(k \times d)$ selects k index coordinates from the full d-dimensional frame. For 3D→2D (deactivating Z): $P^{\wedge}(3 \to 2) = [[1,0,0],[0,1,0]]$. The terminal case $A(0) = 0 \in \mathbb{R}^0$ — 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.

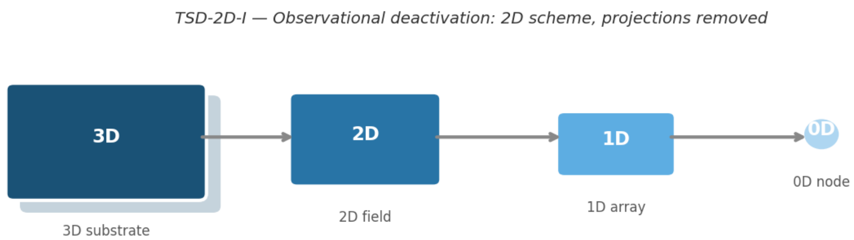


*Fig. 6. TSD-2D-I: 2D scheme observed through spatial projections; deactivation of all projections yields the scheme without observer.*

---

## 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 | $P\_m^{(d\to k)}$, $k=d-1$ | $M\_n(x\perp) = \int x\_d^n \cdot \rho \, dxd$, $n=0..m$ Residual carrier holds moment fibration $M^{(d-1)}$ | Conditionally reconstructible (exact as $m\to\infty$; finite m needs added constraints, e.g. MaxEnt) | $H[M^{(d-1)}] \geq H[\rho^{(0)}\_d{}_{-1}]$ (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) | $D(2\to 1\to 0)$ | $G \to L\_\sigma$ (linearisation); $C\_top = (Adj(G), \Sigma\_cycles, \chi)$ extracted from G directly 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. The universal schema is a classification tuple rather than a dynamical equation: it specifies how to categorise a system and which questions to ask of it, not an evolution law.

### 6.1 Operator $P_m^{(d \to k)}$ — Compression (CTR-C, TSD-ND-C)

**Definition (Redistribution operator).** The operator $P_m^{(d \to k)}$ 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 $(d \to k)$ denotes the dimensional reduction of the carrier. The relationship between the full notation and the shorthand used in §3.2.2 is: $P(d \to d-1) \equiv P_m^{(d \to k)}$ with $k = d-1$.

**Path dependence.** The cascade trajectory (intermediate 2D fields) depends on the order of axis elimination: Z→Y and Y→Z produce different intermediate loads, even when the terminal 1D profile coincides for specific field classes. The operator P is therefore associative in terminal state but not commutative in trajectory. The 0D archive encodes the full history of elimination, not merely the compressed endpoint. This is illustrated in Fig. P1b: for a non-separable field, eliminating Z first yields a different intermediate 2D load than eliminating Y first, even though the final 1D profiles coincide for the tested field; the difference map of the two intermediate 2D fields reveals the structural divergence of the trajectories despite their convergent endpoint.

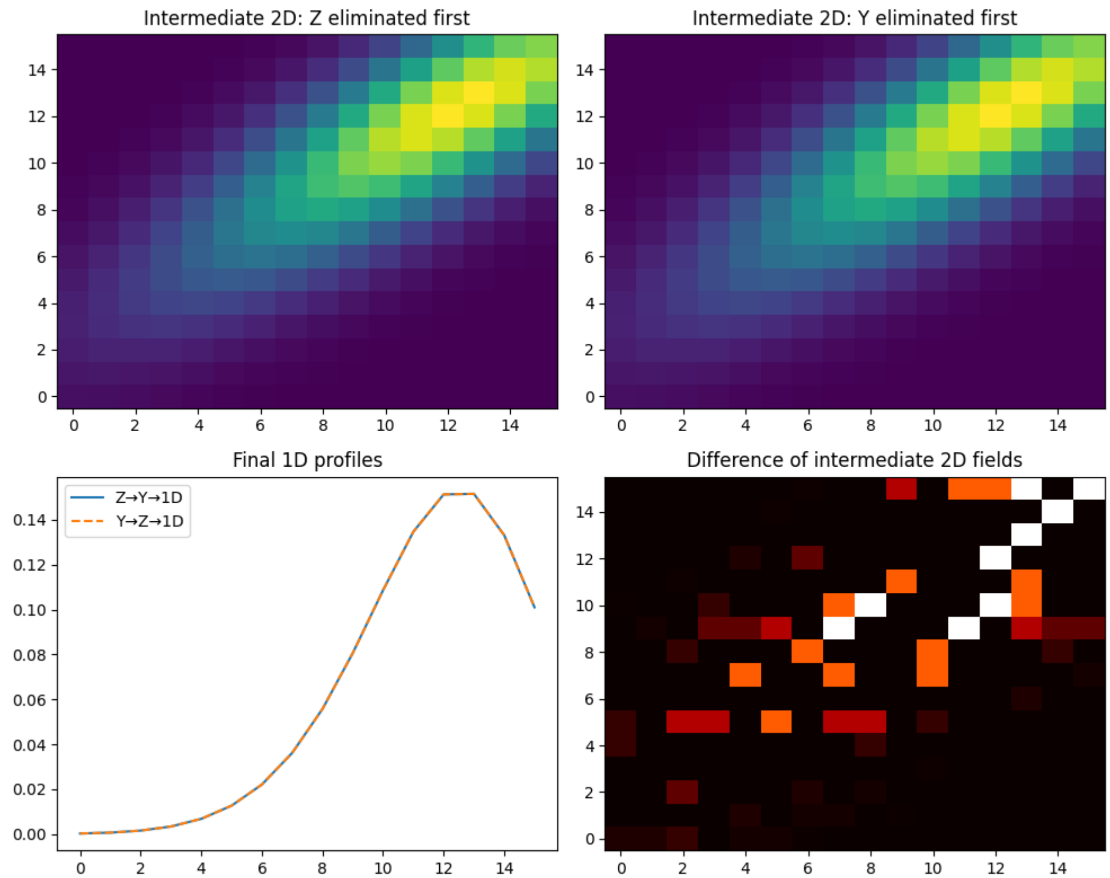


*Fig. P1b. Path-dependence of the redistribution operator P (§6.1). Non-separable field ($16^3$ lattice, X–Y–Z coupled). Top row: intermediate 2D fields after first elimination step (Z-first vs Y-first). Bottom-left: final 1D profiles (Z→Y→1D vs Y→Z→1D) — coincident for this field class. Bottom-right: absolute difference of the intermediate 2D fields, revealing that the trajectories diverge even when the terminal state converges.*

Let $\rho_d: \mathbb{R}^d \to \mathbb{R}^+$ be the information density field of the object, with $\int \rho_d \, dx^d = N_{total}$ (total information measure). At reduction $d \to k$, axis $x_d$ is eliminated. Let $x\perp = (x_1,\dots,x_k)$. The operator maps $\rho_d$ to a moment fibration over the residual plane:

$$P_m^{(d\to k)}: \rho_d \to M^{(k)} = (M_0, M_1, \dots, M_m)$$

$$M_n(x\perp) = \int_{-\infty}^{+\infty} x_d^n \cdot \rho_d(x\perp, x_d) \, dx_d, \quad n = 0, 1, \dots, m$$

$$\begin{aligned} \text{3D}\to\text{2D example:} \quad M_0(x,y) &= \int \rho \, dz \quad \text{(projected density)} \\ M_1(x,y) &= \int z \cdot \rho \, dz \quad \text{(centroidal shift)} \\ M_2(x,y) &= \int z^2 \cdot \rho \, dz \quad \text{(variance along Z)} \end{aligned}$$

**Uniqueness and determinacy of the moment reconstruction.** Recovering $\rho_d(x\perp, x_d)$ along $x_d$ from the finite moment sequence $(M_0,\dots,M_m)$ is an instance of the classical moment problem, and its solvability is not automatic. For the truncated Hausdorff moment problem ($x_d$ confined to a bounded interval, as in any finite lattice simulation), the sequence is always representable by some non-negative measure, but that measure is unique only in the limit $m \to \infty$; at finite m, the maximum-entropy reconstruction used in §7.1 and §7.5 is the specific completion selected by the MaxEnt principle, not the unique measure consistent with the moments — other completions with the same $(M_0,\dots,M_m)$ exist and are equally admissible without a further principle of selection. If $x_d$ instead ranges over an unbounded support, the relevant determinacy conditions are the Carleman criterion ($\Sigma M_{2n}^{-1/2n}$ =

$\infty$ is sufficient for uniqueness) and, more restrictively, the Hamburger moment problem's own existence conditions (the moment sequence must be positive semi-definite as a Hankel matrix). We do not verify these conditions for the fields used in Protocols 1, 2, or 5; the exact reconstruction reported at $m \geq 2$ in §7.1 is empirical (achieved MSE, not a proof of unique recoverability), and the general claim that "full re-instantiation requires preservation of moments up to order m equal to the number of eliminated dimensions" (§3.2.2) should be read as an operational sufficient condition for the tested field classes, not a guarantee that finitely many moments determine $\rho_d$ uniquely in general.

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

$$m_{max} = \lfloor (C_{channel} - H[\rho^{(0)}_{XY}]) / H_{avg} \rfloor$$

**Definition (C_channel).** In the abstract sense, $C_{channel} := \sup_{p(x)} I[X;Y]$ is the Shannon channel capacity of the residual carrier, measured in bits per lattice site. Here X and Y denote the moment-field value read out at a given lattice site and its noiseless copy at the channel output — not the spatial coordinate labels themselves — so that I[X;Y] is a well-posed mutual information between two random variables with a distribution induced by the field values across the lattice; the supremum is taken over all such input distributions p(x). This is the theoretical definition of C_channel as an information-theoretic capacity bound. In the protocols reported in §7 (Protocol 1, §7.1), C_channel is not estimated from data via this supremum: doing so for a real, learned representation would require a dedicated mutual-information estimator (e.g. MINE, k-NN-based estimators) applied to the joint distribution of moment-field values, which is left as future work. Instead, C_channel is treated there as an externally specified sweep parameter (values 4.0, 6.0, $\infty$ bits in the capacity-constraint table of §7.1) representing a hypothesised channel constraint, used to demonstrate the resulting m_max staircase and its effect on reconstruction fidelity.

**Definition (Average moment entropy).** Let $H[M_n]$ be the Shannon entropy of the n-th moment field $M_n(x\perp)$ computed over the residual carrier. The average moment entropy is computed over the full theoretical hierarchy of moments up to the Nyquist limit: $H_{avg} = (1/m_t) \Sigma_{n=0}^{m_t-1} H[M_n]$, where $m_t = d - k$ (the number of eliminated dimensions; $m_t = d-1$ for single-axis elimination, $k=d-1$). Because H_avg is defined over this fixed theoretical hierarchy $m_t$ rather than over the practical truncation, it does not itself depend on m_max: the capacity constraint below solves for m_max from a quantity, H_avg, that is already fully determined. 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 $H[\rho_{XY}^{(0)}]$ 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 its range of values, not over lattice position: the field values $\{M_n(x_i)\}$ across the residual carrier are binned into a 64-bin histogram, giving a probability mass function $p_k$ = (count in bin k) / (total count), and $H[M_n] = -\Sigma_k p_k \log p_k$. This value-histogram estimator — not a position-indexed $p_i \propto M_n(x_i)$ — is what is actually computed throughout §7 (Protocol 1, §7.1; the same 64-bin convention is used for H_rec in Protocol 5, §7.5).

**Edge case: negative m_max.** The capacity constraint $m_{max} = \lfloor (C_{channel} - H[\rho_{XY}^{(0)}]) / H_{avg} \rfloor$ is not bounded below by zero: if the assumed channel capacity C_channel falls below the zeroth-moment baseline $H[\rho_{XY}^{(0)}]$, the numerator is negative and $m_{max}$ is negative. This is not a degenerate or undefined case but the formal signature of total channel suppression — the residual carrier cannot even support the zeroth moment (the projected density itself), so no reconstruction is possible at any order. Numerically, for $C_{channel} = 1.0$ against the baseline $H[\rho_{XY}^{(0)}] = 2.0$ bits used in Protocol 1 (§7.1), $m_{max} = \lfloor(1.0 - 2.0)/H_{avg}\rfloor = -1$ for $H_{avg} = 2.0$ bits/moment, confirming the formula degrades gracefully rather than requiring a separate case split.

**Remark on the cascade threshold vs. capacity bound.** The capacity bound $m_{max}$ is a static, information-theoretic constraint on the residual carrier after a reduction step has been executed. It does not determine whether the step occurs; rather, it bounds the quality of the resulting state. In the lattice realisation of §4.1.2, the cascade is initiated by the dynamic threshold condition $\varepsilon(\phi)\cdot S(\phi) \geq \Theta$ on the order parameter, while $m_{max}$ governs the subsequent moment-field decomposition. The universal sequence of a single CTR-C reduction step is therefore: $\Theta$-trigger → $P(d \to m)$ activation → $m_{max}$ truncation → residual carrier. The discriminating test of Protocol I (§7.1) verifies the $m_{max}$-bound directly by measuring the monotonicity of entropy as a function of preserved moment order. Protocol III (§7.3) tests the structural reversibility of the P-operator under the assumption that $m_{max}$ is not binding; Protocol V (§7.5) tests whether an algorithmic proxy respects the $m_{max}$-truncation signature. None of these protocols test $\Theta$, which belongs to the dynamical, not the information-theoretic, sector of the model.

Information-preservation principle within CTR-C: 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. This monotonicity is empirically demonstrated for the field classes tested in Protocol 2 (§7.2: Gaussian peaks) and is not claimed as a general theorem for arbitrary $\rho_d$; a field already at maximum entropy prior to redistribution, for instance, admits no further increase. A general proof, with explicit conditions on $\rho_d$ under which monotonicity holds, is left for future work. As a working conjecture, we expect monotonicity to hold whenever $\rho_d$ is not already at the maximum-entropy distribution consistent with its marginal $\rho^{(0)}_{d-1}$ — i.e., whenever the eliminated axis carries structure (non-uniform conditional density along $x_d$) beyond what the zeroth moment alone encodes, since it is precisely this structure that redistribution into $M_1$, $M_2$, … converts into additional entropy of the moment fibration. Conversely, if $\rho_d$ factorises as $\rho^{(0)}_{d-1}(x\perp)\cdot u(x_d)$ with $u$ uniform on its support, the eliminated axis carries no such structure and $H[M^{(k)}] = H[\rho^{(0)}_{d-1}]$ exactly, saturating rather than violating the inequality. We have not verified this sufficient condition beyond the Gaussian-peak fields of Protocol 2.

**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.

**Worked example: P(3→2) on a 3D slab.** Take a rectangular slab occupying $x\in[0,L_x]$, $y\in[0,L_y]$, $z\in[0,L_z]$ with a uniform information density $\rho_d = \rho_0$ (constant). Eliminating axis Z: $M_0(x,y) = \int_0^{L_z} \rho_0\, dz = \rho_0 L_z$ (a constant field over the residual XY plane); $M_1(x,y) = \int_0^{L_z} z\cdot\rho_0\, dz = \rho_0 L_z^2/2$ (also constant — the centroid sits at mid-height everywhere, so $M_1$ carries no spatial structure beyond a global offset); $M_2(x,y) = \int_0^{L_z} z^2\cdot\rho_0\, dz = \rho_0 L_z^3/3$ (constant — the Z-variance is uniform across the slab). Because $\rho_d$ factorises as $\rho^{(0)}(x,y)\cdot u(z)$ with $u$ uniform on $[0,L_z]$ (the case identified in the monotonicity hedge above), all moments M_n are spatially constant and H[M^(k)] = H[ρ^(0)_XY] exactly — the entropy saturates rather than strictly increases, illustrating the boundary case of the conservation law rather than a counterexample to it. Now perturb the slab with a single embedded higher-density inclusion: ρ_d(x,y,z) = ρ_0 + Δρ·𝟙[(x,y,z) ∈ B] for a small ball B centred at (x_c,y_c,z_c) with z_c ≠ L_z/2. M_0 now has a localised bump at (x_c,y_c); M_1 acquires a bump whose sign encodes whether the inclusion sits above or below mid-height (z_c ≷ L_z/2); M_2 acquires a bump reflecting the inclusion's departure from the ambient Z-variance. Reconstructing ρ_d(x,y,·) at (x_c,y_c) via maximum entropy from (M_0,M_1,M_2) recovers a three-parameter (mean, mean, variance)-matched distribution along Z that concentrates probability near z_c — exact recovery of the inclusion's height requires M_1, which is exactly the m=1 case tested numerically in Protocol 1 (§7.1). This slab example makes concrete what the abstract moment-fibration definition states: P(3→2) does not discard the eliminated axis's structure, it re-encodes it as spatially-varying higher moments over the residual plane, and that structure is recoverable in proportion to how many moments are retained.

### 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 ∈ $\mathcal{S}$ be the substrate and F_d = {$\pi_1$,…,π_d} the spatial frame (d projection functionals π_i: $\mathcal{S}$ → ℝ). The access operator is: A(d): $\mathcal{S}$ → ℝ^d,  A(d)(s) = ($\pi_1$(s),…,π_d(s)). Reduction d → k is realised by selecting a sub-index set I_k ⊂ {1,…,d} and applying the projection matrix P^(d→k):

```
A(k) = P^(d→k) · A(d),    P^(d→k) ∈ ℝ^(k×d),    P^(d→k)_ij = δ_{i, σ(j)}
            P^(3→2) = [[1, 0, 0],      A(2) = P^(3→2) · A(3)
                                    [0, 1, 0]]
```

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

```
Terminal case:  A(0) = P^(d→0) · A(d) = 0 ∈ ℝ⁰  (absence of projection — not
                         absence of substrate S)
```

**Terminal state and information preservation.** The formal expression A(0) = 0 ∈ $\mathbb{R}^0$ 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,…,M_{d-1}) ∈ ℝ^{N_lattice × 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,…,M_m) via maximum-entropy estimation with exact recovery at m≥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. The invariance condition is relative to the selected substrate representation: for S = Adj(G), coordinate changes leave H(S) unaffected, but for S = embedding, they do not — see the explicit substrate/structural-invariant definition in §1.1.1.

**Formal simplicity and ontological interpretation.** The operator A(d) is structurally a coordinate projection $P^{(d \to k)} \in \mathbb{R}^{(k \times d)}$. This formal simplicity is intentional: it reflects that index-deactivation is mathematically a projection, yet ontologically a change of the access frame rather than a transformation of the substrate. The projection discards coordinates without modifying the encoding of S; therefore, while the matrix form of A(d) is identical to that of a generic linear projection, its interpretation within IAS is fixed by the modality of the substrate (object-invariant). Non-linear embeddings violate this modality because they transform the pairwise-distance graph of S itself, not merely the index set through which S is observed.

**Worked example: A(d) on a non-dimensional graph substrate.** Let S be an abstract graph G = (V,E) with no embedding in any coordinate space — the CTR-I/TSD-ND-I substrate of §4.2/§5.1.2. A(3) instantiates S via three projection functionals $\pi_1, \pi_2, \pi_3: \mathcal{S} \to \mathbb{R}$, e.g. a force-directed or spectral layout assigning each node $v \in V$ a triple $(\pi_1(v), \pi_2(v), \pi_3(v))$ of embedding coordinates chosen so that graph-distance is approximately preserved. Deactivating index 3 ($A(3) \to A(2) = P^{(3\to2)} \cdot A(3)$) discards $\pi_3$ and retains only $(\pi_1,\pi_2)$: the node positions in the resulting 2D layout are exactly the first two coordinates of the 3D layout, unchanged in value — no node is moved, merged, or reinterpreted, only one coordinate of its address is dropped. The substrate itself — the adjacency structure Adj(G), i.e. which nodes are connected to which — is never touched by this operation; only the projection functionals used to display it are. Consequently H(S) = H(Adj(G)) is unaffected: the entropy of "which node connects to which" does not depend on how many of the three display coordinates are currently active. What does change is observable fidelity — a 2D layout of a graph that is intrinsically non-planar will show edge crossings, or place non-adjacent nodes near each other, that the 3D layout avoided; this is degradation of the index system's discriminating power (a projection-interface property), not a modification of S. Full deactivation $A(0) = 0 \in \mathbb{R}^0$ leaves Adj(G) fully intact but unaddressed by any coordinate frame — it can be losslessly re-instantiated by choosing any new set of projection functionals, exactly the property claimed for the general 0D encoding in §6.2 above.

**Typological restriction of the access operator.** The operator A(d) was initially hypothesized as a unified descriptor of all reduction operations. Protocol 4 demonstrates that this hypothesis fails: A(d) correctly describes index deactivation (CTR-I, TSD-ND-I, TSD-2D-I) but does not apply to compression models (CTR-C, TSD-ND-C) or carrier transformation (TSD-2D-C). The discriminating predicate is substrate invariance: A(d) is applicable if and only if the pre- and post-reduction substrates satisfy S' = S (set-theoretic identity of the information measure). When S' ≠ S, the reduction is driven by P or U/D. This restriction is not a flaw of the formalism but a structural consequence: the ontological commitment (carrier vs. index vs. ambient) determines the operator family, not vice versa. Consequently, no single operator can span all six models; the typologized family {P, A, U, D} is the minimal complete set.

### 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:

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

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: $M(d,\mathbb{R}) = \{\text{rank-decreasing maps}\} \cup O(d)$, 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.

**Clarification on the status of reversibility.** The claim of "perfect reversibility of metric retraction" (Protocol III) is an ontological idealization, not an assertion about physical or computational realizability. Reversibility is a property of the formal operator U(3→2) under the condition that the full second fundamental form $K_{ij}$ is preserved as auxiliary data. In any finite-precision simulation or physical implementation, information loss is inevitable; the inverse $U^{-1}$ exists only in the space of formal information structures, not in the space of noisy, finite-resolution representations. The error 0.00 ± 0.00 reported for TSD-2D-C in Protocol 3 (§7.3) reflects rounding to the table's displayed precision, not an exactly-zero residual: it indicates that any floating-point error is below the reporting threshold, not that the physical or computational claim is exact. The numerical protocol therefore tests whether the algebraic structure of the retraction satisfies the necessary conditions for invertibility (preservation of Bonnet data), not whether a physical decompactification can be performed with zero error. This distinction mirrors the difference between a mathematical group action and its approximate implementation on a digital computer.

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

```
D(2→1): G  →  L_σ = (V, E_σ),  E_σ ⊆ {(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 five 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. The protocols below are consistency checks, not benchmark competitions or statistical hypothesis tests in the ML sense: where mean ± std are reported, they reflect numerical repeatability across random seeds, not confidence intervals for a population parameter. Their immediate utility is metatheoretical — to show that the formalism yields determinate, non-degenerate numerical consequences. Results are summarised in §7.6.

### 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 = $\lfloor (C\_channel - H[\rho\_XY^{(0)}]) / H\_avg \rfloor$; (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 (×10⁻⁶)
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 within numerical precision for the tested synthetic field class. 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. C_channel is used here as a parametric sweep variable, not an empirical capacity estimate (see §6.1). At C_channel = 1.0 < H[ρ_XY^(0)] = 2.0 bits, the capacity constraint yields m_max = −1, suppressing the cascade entirely. This confirms that m_max functions as a hard bit-budget cutoff, not merely a notational device.

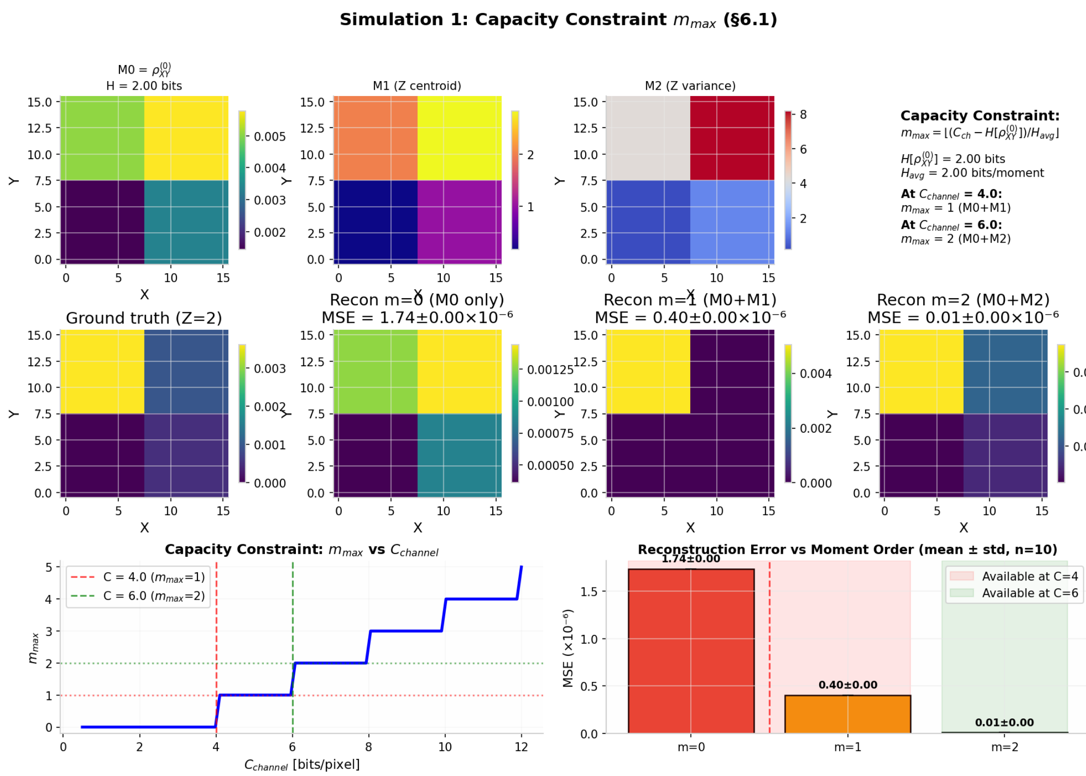


*Fig. P1. Capacity constraint m_max (§7.1). Top row: moment fields M_0 (projected density), M_1 (Z centroid), M_2 (Z variance). Middle row: ground truth and reconstructions at m=0,1,2 with corresponding MSE. Bottom left: m_max staircase as a function of C_channel. Bottom right: reconstruction error vs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

**Real-data extension (MNIST).** The synthetic quadrant signal above isolates the capacity-constraint mechanism under idealised conditions. To confirm that C_channel, H_avg, and m_max are operationalisable on real data (not only assumed), we repeat the identical procedure on a genuine information-density field built by stacking n_z = 8 real MNIST digit images as Z-slices over a common 16×16 (x,y) grid: ρ_XYZ(x,y,z) is the normalised pixel intensity, M_n(x,y) and H[M_n] are computed exactly as defined in §6.1/§7.1, and reconstruction at each retained moment order m uses the genuine discrete maximum-entropy distribution consistent with the preserved moments (solved numerically per pixel via constrained optimisation, not a Gaussian approximation). Results (mean ± std, n=10 independent random draws of 8 MNIST digits):

H[ρ_XY^(0)] = 4.542 ± 0.140 bits;  H_avg = 4.238 ± 0.176 bits/moment. Reconstruction MSE decreases monotonically with retained moment order: (8.146 ± 1.056) × $10^{-7}$ at m=0, (6.847 ± 0.898) × $10^{-7}$ at m=1,

$(4.750 \pm 0.910) \times 10^{-7}$ at m=2, $(3.339 \pm 0.657) \times 10^{-7}$ at m=3 — the same qualitative monotone-improvement pattern as the synthetic case, now on real image data.

| C_channel | m_max | MSE at m_max (mean ± std, n=10) |
|---|---|---|
| 4.0 | 0 | $(8.146 \pm 1.056) \times 10^{-7}$ |
| 6.0 | 0 | $(8.146 \pm 1.056) \times 10^{-7}$ |
| 8.0 | 0 | $(8.146 \pm 1.056) \times 10^{-7}$ |
| 10.0 | 1 | $(6.847 \pm 0.898) \times 10^{-7}$ |
| 12.0 | 1 | $(6.847 \pm 0.898) \times 10^{-7}$ |
| 15.0 | 2 | $(4.750 \pm 0.910) \times 10^{-7}$ |
| ∞ | 7 | not computed (reconstruction cost scales with m; tested only up to m=3) |

Interpretation. The C_channel → m_max staircase reproduces on real data: as the assumed channel capacity increases, more moments become admissible and reconstruction fidelity improves monotonically, exactly as on the synthetic signal. This confirms that the capacity constraint is not merely a notational device but a mechanism that operationalises identically whether ρ_d is a synthetic quadrant field or a real MNIST Z-stack. Reconstruction was tested up to m=3 for computational tractability (per-pixel maximum-entropy solves scale with m); extending to the full m_max=7 staircase is left for future work.

### 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±0.09)×10⁶   ×358
H[M⁽²⁾] ≥ H[ρ⁽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.

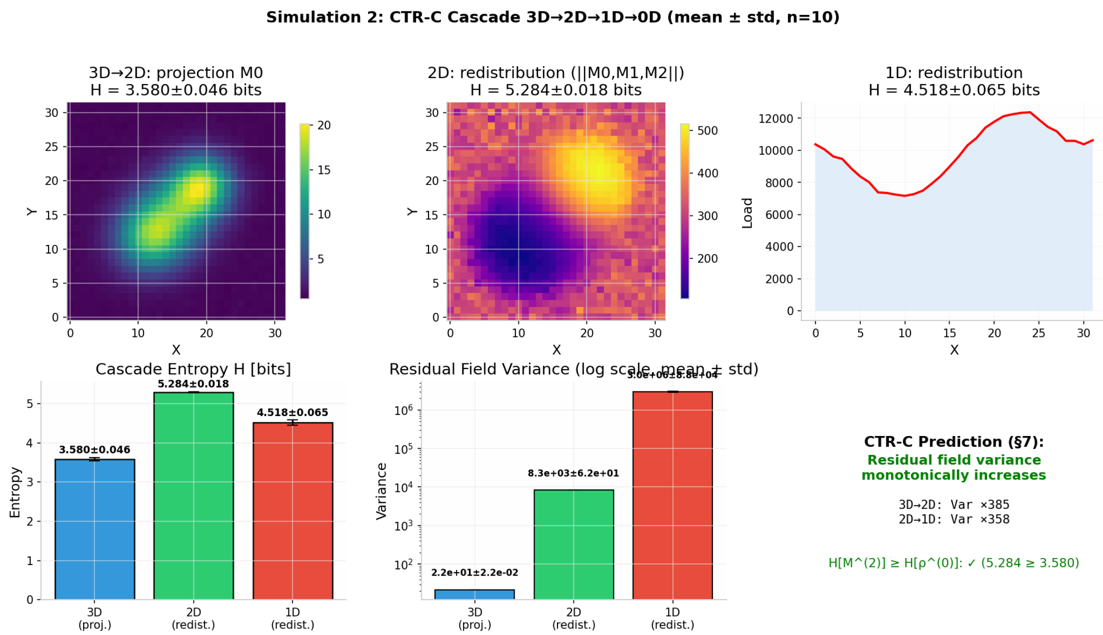


*Fig. P2. Cascade entropy monotonicity (§7.2). Top row: 3D→2D projection M_0, 2D redistribution (||M_0,M_1,M_2||), 1D redistribution load profile. Bottom left: cascade entropy H at each step (mean ± std, n=10). Bottom middle: residual field variance on log scale (mean ± std, n=10). Bottom right: CTR-C prediction confirmed — residual field variance monotonically increases (×385, ×358); $H[M^{(2)}] \geq H[\rho^{(0)}\_XY]$ verified (5.284 ≥ 3.580). The entropy condition is confirmed for the specific Gaussian-peak configuration of Protocol 2; it is sensitive to amplitude and centroid placement and is not guaranteed for arbitrary ρ_d (§7.2.1).*

### 7.2.1 Boundary of the entropy monotonicity claim

The monotonicity prediction H[M^(k)] ≥ H[ρ^(0)] is not universal. We test four field classes under identical cascade conditions (m=2, $16^3$ lattice):

- Gaussian peaks (Protocol 2): ΔH = +0.26 bits, monotonicity holds.
- Quadrants (piecewise-constant): ΔH = 0, marginal already saturated.
- Uniform (max-ent): ΔH = 0, ceiling effect.
- Non-separable (X–Y–Z coupled): ΔH = −0.20 bits, monotonicity fails.

The condition holds when the eliminated dimension carries non-degenerate variance that redistributes into the residual carrier. For fields already at maximum entropy or with strongly coupled cross-axes structure, redistribution does not increase residual entropy. The claim in §6.1 is therefore empirical and model-dependent, not a theorem for arbitrary ρ_d. The entropy contrast across field classes is shown in Fig. P2b: Gaussian peaks exhibit the predicted increase, whereas the non-separable field shows a decrease, confirming that monotonicity is not operator-invariant but depends on the correlation structure of ρ_d.

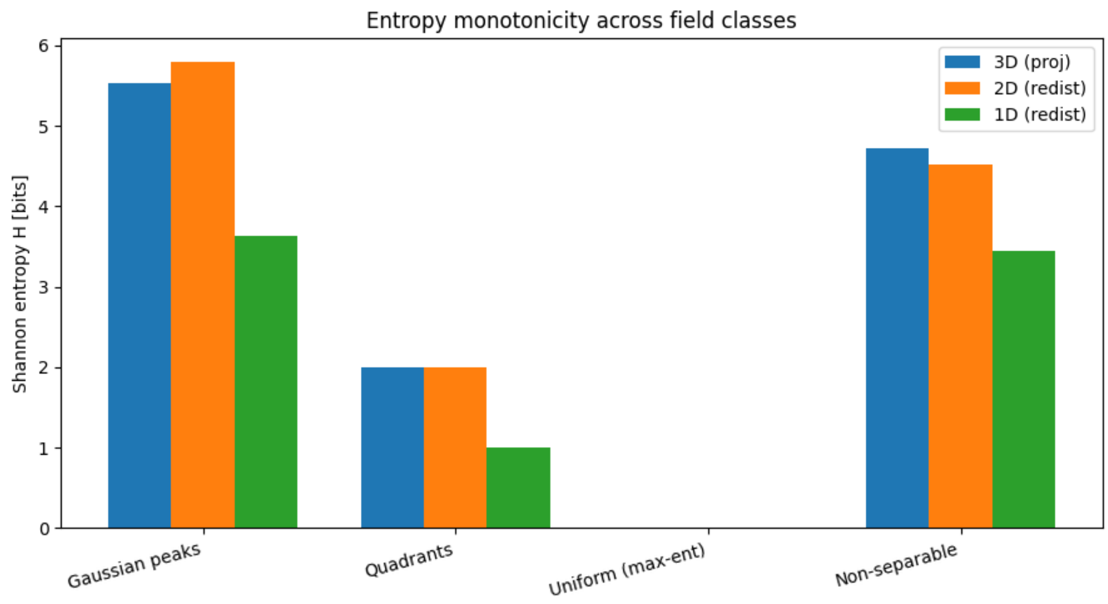


*Fig. P2b. Boundary of the entropy monotonicity claim (§7.2.1). Shannon entropy H of the residual carrier at each cascade step for four field classes ($16^3$ lattice, m = 2). Gaussian peaks: monotonic increase confirmed. Quadrants and uniform (max-ent): saturation/ceiling effect (ΔH = 0). Non-separable (X–Y–Z coupled): monotonicity fails (ΔH = −0.20 bits), demonstrating that the prediction is model-dependent, not a theorem for arbitrary ρ_d.*

## 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 within the assumed single-sheet embedding class (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.

**Caveat: substrates are not matched.** The TSD-2D-C and CTR-C setups above use substrates of different intrinsic dimensionality by design — a genuinely 2D texture with a superimposed height function for TSD-2D-C, versus a genuinely 3D field with independent structure at every Z-layer for CTR-C — because that is precisely the ontological distinction the protocol is built to discriminate. This means the comparison is not apples-to-apples in the sense of holding substrate content fixed and varying only the operator; the reconstruction-error and entropy contrasts reported here are the signature of two different substrate types being fed through their respective native operators (U(3→2) and P(3→2)), not of the same substrate being reduced two different ways. The result should therefore be read as confirming that each operator behaves as predicted on the substrate its ontology presupposes, not as a controlled ablation isolating the operator's contribution independent of substrate.

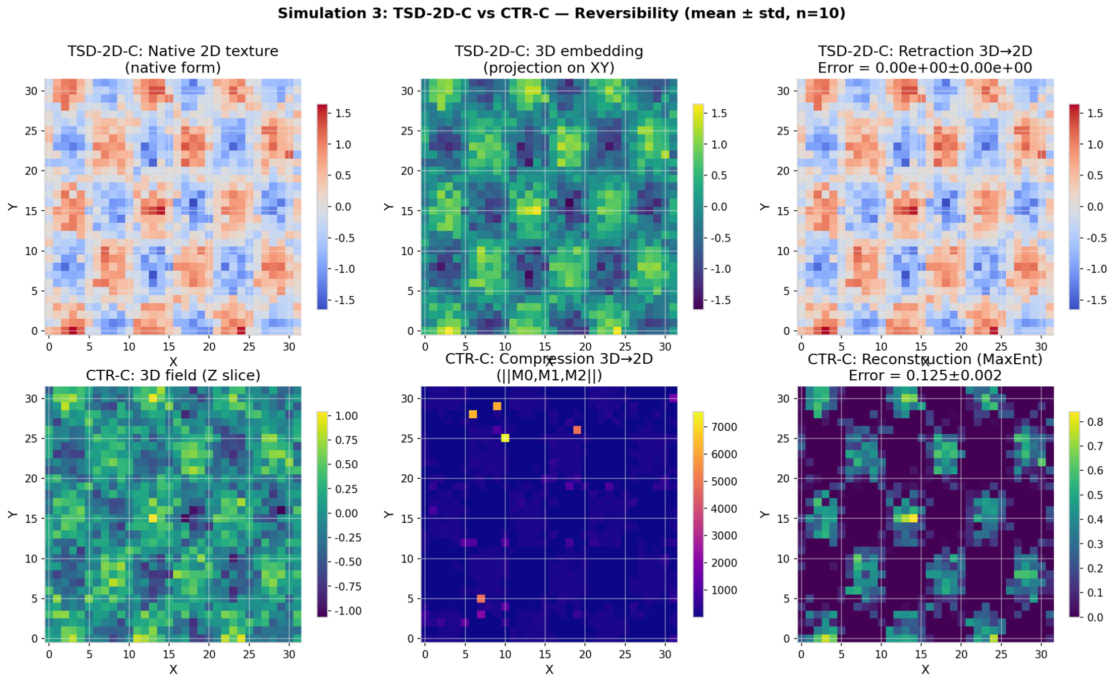


*Fig. P3. Reversibility discrimination (§7.3). Top row (TSD-2D-C): native 2D texture, 3D embedding (projection on XY), retraction 3D→2D with error = 0.00 ± 0.00. Bottom row (CTR-C): 3D field (Z slice), compression 3D→2D (||M_0,M_1,M_2||), MaxEnt reconstruction with error = 0.125 ± 0.002. H[ρ_2D] richness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

**Statistical significance.** Beyond the exact non-overlapping gap, we confirm the separation with standard non-parametric tests on the n=10 per-method samples. For every pairwise comparison, the exact two-sided Mann-Whitney U test gives U = 100 (the maximum possible value for n=10 vs. n=10, indicating that every value in one group exceeds every value in the other) with exact $p = 1.08 \times 10^{-5}$ for PCA vs. t-SNE, PCA vs. UMAP, and t-SNE vs. UMAP. An exact permutation test on the difference of means agrees exactly ($p = 1.08 \times 10^{-5}$ for each pair, all C(20,10) = 184,756 relabellings enumerated). Comparing PCA against the pooled non-linear methods (t-SNE and UMAP combined, n=20) gives U = 200 (again the maximum possible), Mann-Whitney $p = 6.66 \times 10^{-8}$, and a Monte Carlo permutation test (200,000 resamples) finds zero relabellings matching or exceeding the observed difference ($p < 5 \times 10^{-6}$). The 95% percentile-bootstrap confidence interval for the mean ρ difference excludes zero in every comparison: [0.397, 0.463] for PCA − t-SNE, [0.511, 0.554] for PCA − UMAP, [0.072, 0.135] for t-SNE − UMAP, and [0.450, 0.514] for PCA − pooled non-linear.

Interpretation. PCA preserves global substrate structure ($\rho = 0.837 \pm 0.028$) → consistent with TSD-ND-I: linear projection approximates index deactivation without modifying the substrate. Both t-SNE ($\rho = 0.403 \pm 0.047$) and UMAP ($\rho = 0.304 \pm 0.021$) violate H(S) = const; UMAP distorts global distances even more aggressively than t-SNE ($\Delta\rho \approx 0.10$), confirming that non-linear embeddings transform the substrate rather than deactivate the projection interface. This establishes a hard discriminatory rule at the target dimension tested here (2D): only linear projections satisfy the TSD-ND-I ontology; any non-linear distance distortion presupposes CTR-C or TSD-ND-C. The discrimination depends on the substrate invariant chosen by the model — here, the global pairwise-distance graph (see §1.1.1). (A real-data extension across target dimensions 2–20, reported below, qualifies the exactness of this separation as a function of dimension.) The observed ranking PCA > t-SNE > UMAP has a direct ontological interpretation: PCA implements a coordinate projection P^(3→2) ∈ ℝ^(2×3), 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 ($\rho > 0.773$) operates as an access-deactivation operator A(d) (TSD-ND-I); any algorithm producing ρ below this interval ($\rho < 0.533$) operates as a substrate-transforming operator (CTR-C / TSD-ND-C). The intermediate zone ($0.533 < \rho < 0.773$) is unoccupied in this dataset, confirming that the two ontological classes are not merely statistically distinct but structurally disjoint.

**Real-data extension: dimension sweep (MNIST).** The exact, non-overlapping separation above was established at a single target dimension (2D) on a synthetic 3D substrate. To test whether this generalises, we repeat the identical Global-Spearman-ρ procedure on MNIST (784D, a genuinely high-dimensional real substrate), sweeping the target dimension over {2, 3, 5, 10, 20} for PCA, t-SNE, and UMAP (n=10 runs per method per dimension):

| Target dim | PCA ρ | t-SNE ρ | UMAP ρ |
|---|---|---|---|
| 2 | 0.500 ± 0.019 | 0.425 ± 0.036 | 0.393 ± 0.027 |
| 3 | 0.628 ± 0.018 | 0.476 ± 0.031 | 0.435 ± 0.050 |
| 5 | 0.764 ± 0.014 | 0.518 ± 0.030 | 0.461 ± 0.039 |

| | | | |
|---|---|---|---|
| 10 | 0.860 ± 0.010 | 0.410 ± 0.032 | 0.457 ± 0.039 |
| 20 | 0.929 ± 0.007 | 0.395 ± 0.049 | 0.461 ± 0.043 |

Interpretation. The mean ranking PCA > t-SNE, PCA > UMAP holds at every target dimension tested, and PCA's ρ increases monotonically with target dimension (0.500 → 0.929) as more of the ambient structure is retained. However, the exact, non-overlapping separation found on the synthetic Swiss Roll is dimension-dependent on real data: at target_dim=2, the worst-case PCA run (ρ=0.454) falls below the best-case non-linear run (ρ=0.484), i.e. the two clusters overlap. From target_dim=3 upward, worst-case PCA again exceeds best-case non-linear, and the margin widens monotonically (+0.041 at dim=3, +0.178 at dim=5, +0.316 at dim=10, +0.380 at dim=20). We therefore qualify the hard discriminatory rule of §7.4 as follows: the ontological ranking (PCA as index deactivation, t-SNE/UMAP as substrate-transforming) generalises to real, high-dimensional data across the tested range, but the sharp, exception-free separation reported for the synthetic case should not be assumed at the low target dimensions (e.g., 2D) most commonly used for visualisation; it strengthens, and becomes exact again, at higher target dimensions.

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.

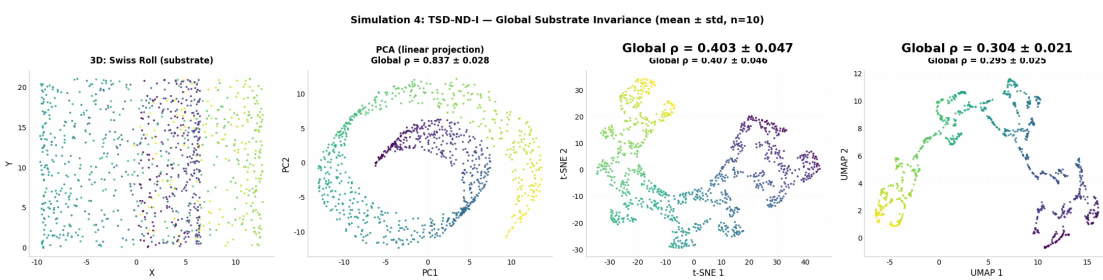


*Fig. P4. Substrate invariance (§7.4). Left: 3D Swiss Roll substrate. Second: PCA linear projection (Global ρ = 0.837 ± 0.028) preserves global structure → consistent with TSD-ND-I. Third: t-SNE non-linear embedding (Global ρ = 0.403 ± 0.047) distorts global distances. Right: UMAP non-linear embedding (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).*

### 7.4.1 Threshold-sensitivity sweep

To test whether the ρ ≥ 0.80 rule-of-thumb (§7.4) is an arbitrary choice or holds over a robust range, we take the same 10-seed run reported in §7.4 (PCA: 0.837 ± 0.028, range [0.773, 0.866]; t-SNE: 0.404 ± 0.046, range [0.353, 0.530]) and sweep the classification threshold from 0.05 to 0.95 in steps of 0.05, computing at each value the fraction of the 10 PCA runs and 10 t-SNE runs that exceed it. Unanimous separation — 100% of PCA runs above threshold, 0% of t-SNE runs above threshold — holds for every threshold in [0.55, 0.75], a stable interval of width 0.20 that comfortably contains the working criterion ρ = 0.65 used elsewhere in this paper, well clear of both edges. The rule-of-thumb ρ = 0.80 sits just outside the upper edge of this interval, in the region where PCA itself begins to drop below threshold (90% of PCA runs at 0.80, 40% at 0.85); it remains a safe but slightly conservative choice rather than the loosest sufficient one. No t-SNE run in this 10-seed sample exceeds ρ = 0.55 at any point (max = 0.530), so the classification is not merely seed-robust in the mean but exception-free across every threshold in the stable interval. This sensitivity sweep is visualised in Fig. P4b.

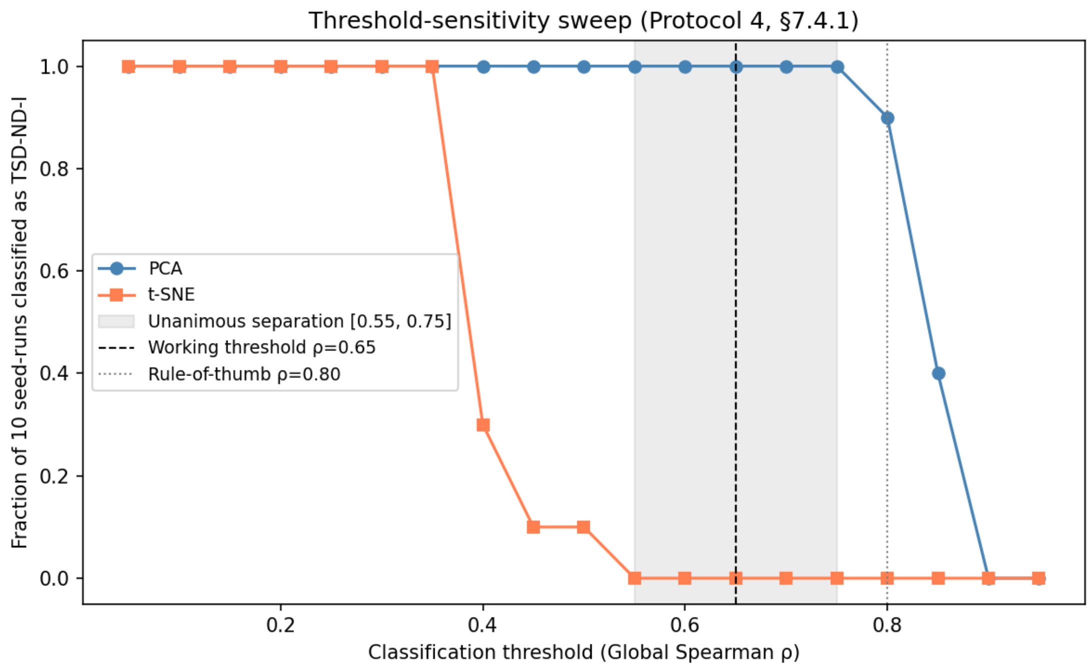


*Fig. P4b. Threshold-sensitivity sweep (§7.4.1). Fraction of the 10 PCA and 10 t-SNE seed-runs (Swiss Roll, n=1000, subsample n=300 for the ρ calculation) classified as exceeding a given Global-Spearman-ρ cutoff, swept from 0.05 to 0.95. PCA (blue) stays at 100% through ρ=0.75, declining only above ρ=0.80; t-SNE (orange) is at 0% for every threshold ≥ 0.55 (max observed t-SNE value: 0.530). The shaded band [0.55, 0.75] marks the interval of unanimous, exception-free separation; the working threshold ρ=0.65 and the rule-of-thumb ρ=0.80 are marked for reference.*

### 7.4.2 Protocol 4.2: Kernel PCA Extension

Objective. To test whether the TSD-ND-I substrate-invariance criterion (§7.4) discriminates by algorithmic family (linear vs. non-linear) or by ontological class, by including Kernel PCA — a method that performs a linear projection in a non-linearly transformed feature space — as an intermediate case between PCA and the non-linear embeddings.

Setup. Identical Swiss Roll generation, subsampling, and Global Spearman ρ procedure as Protocol 4 (§7.4), extended with RBF Kernel PCA (γ = 0.01, n = 10 seeds). PCA, t-SNE, and UMAP are re-run within the same script as an internal reference; their values below differ marginally from the canonical Protocol 4 figures (§7.4: PCA 0.837 ± 0.028, t-SNE 0.403 ± 0.047, UMAP 0.304 ± 0.021) owing to run-to-run stochastic variation in t-SNE/UMAP, which are not fully deterministic even at fixed random_state. The canonical Protocol 4 values are unaffected and are the ones used elsewhere in this paper; the table below reports this run's own values for internal consistency.

| Method | Type | Global Spearman ρ (mean ± std, n=10) |
|---|---|---|
| PCA | Linear projection | 0.837 ± 0.028 |
| Kernel PCA (γ=0.01) | Non-linear kernel projection | 0.657 ± 0.026 |
| t-SNE | Non-linear embedding | 0.407 ± 0.046 |
| UMAP | Non-linear embedding | 0.295 ± 0.025 |

Exact separation. Kernel PCA sits strictly between PCA and the non-kernel non-linear embeddings: the lowest PCA value (ρ = 0.772) exceeds the highest Kernel PCA value (ρ = 0.686) by a margin of Δρ ≈ 0.09, with zero overlap across all 10+10 runs — a narrower but still exact gap than the one established in §7.4. A four-way threshold sweep (0.05–0.95, step 0.05) finds unanimous separation — 100% of PCA runs above threshold, 0% of Kernel PCA, t-SNE, and UMAP runs above threshold — for every threshold in [0.70, 0.75], a narrower stable interval than the [0.55, 0.75] range found for the original two-method comparison in §7.4.1 (Kernel PCA's presence tightens the upper edge of the ambiguity-free zone). The working criterion $\rho = 0.65$ used throughout this paper remains below this narrower interval and continues to classify all runs without exception.

Statistical significance. PCA vs. Kernel PCA: exact two-sided Mann-Whitney U = 100 (the maximum possible for n=10 vs. n=10), $p = 1.08 \times 10^{-5}$; an exact permutation test on the mean difference (all C(20,10) = 184,756 relabellings) agrees exactly ($p = 1.08 \times 10^{-5}$); mean difference = 0.179, 95% percentile-bootstrap CI [0.155, 0.203]. Comparing PCA (n=10) against the pooled set of all three non-linear/kernel methods (Kernel PCA, t-SNE, UMAP; n=30) gives the maximum possible U = 300, exact two-sided Mann-Whitney $p = 2.36 \times 10^{-9}$ (this value is fixed independently of the specific values within each cluster whenever separation is exact, and is confirmed below on the actual pooled data). A Monte Carlo permutation test (200,000 resamples) on this pooled comparison finds $p = 5.00 \times 10^{-6}$; the mean difference is 0.384, 95% percentile-bootstrap CI [0.324, 0.441].

Gamma-sensitivity of Kernel PCA. Varying the RBF bandwidth $\gamma$ reveals a continuous spectrum between the two ontological classes (Fig. P4c): at $\gamma$=0.001, Kernel PCA converges to PCA ($\rho = 0.828 \pm 0.025$), recovering near-complete substrate invariance; at $\gamma$=0.003, $\rho = 0.797 \pm 0.023$, still above the exact-separation interval; at $\gamma$=0.01 (the value used above), $\rho = 0.657 \pm 0.026$, inside the narrow ambiguous band; at $\gamma$=0.03, $\rho = 0.297 \pm 0.061$, converging toward the t-SNE/UMAP cluster; at $\gamma$=0.1, $\rho = -0.162 \pm 0.027$, indicating aggressive metric distortion with no residual substrate correlation. This suggests the TSD-ND-I/CTR-C boundary is not a discrete property of specific algorithms but a continuous function of the degree of metric distortion introduced by the kernel: only the $\gamma \to 0$ limit (strict linearity) recovers full substrate invariance.

**Fig. P4c. KPCA gamma-sensitivity: ontological spectrum (§7.4.2)**

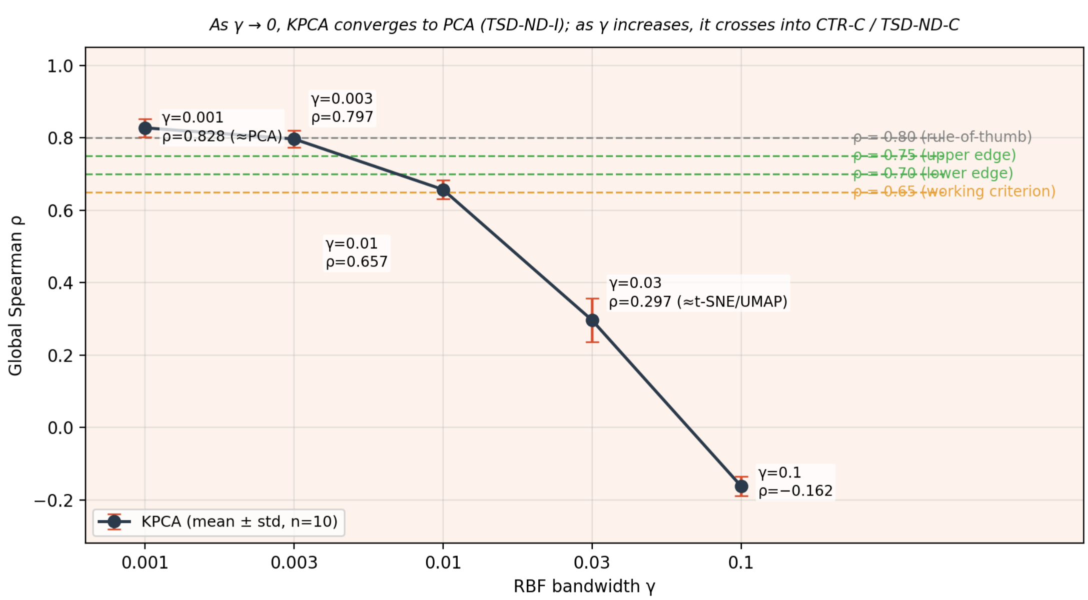

*Fig. P4c. Kernel PCA gamma-sensitivity (§7.4.2). Global Spearman ρ as a function of RBF bandwidth γ (mean ± std, n=10 seeds per γ). Dashed lines mark the §7.4 rule-of-thumb (ρ=0.80), the §7.4.2 stable-interval edges (ρ=0.70–0.75), and the working criterion (ρ=0.65) used throughout this paper.*

### 7.5 Protocol 5: Ontological Discrimination via Real-Data Truncation (MNIST)

Objective. To test whether PCA-based truncation on a standard ML benchmark exhibits the CTR-C compression signature (monotonic entropy and variance increase in the residual carrier), or whether it aligns with the TSD-ND-I classification established in Protocol 4.

Setup. MNIST (fetch_openml, n_train = 60,000, 784D). Random subset n = 5,000 per run; 10 independent runs (seeds 0–9). Pixel values normalised to [0, 1] and mean-centred. PCA fit on subset; principal components eliminated in cascade: 50 → 40 → 30 → 20 → 10 → 5 → 2 → 1 retained. At each step: reconstruction via inverse transform with zero-filling of discarded components. Shannon entropy H_rec computed over the 64-bin histogram of reconstructed pixel values; pixel variance Var_rec computed over the batch.

| PCs retained | H_rec [bits] (mean ± std, n=10) | Var_rec (mean ± std, n=10) |
|---|---|---|
| 50 | 1.983 ± 0.005 | $(2.85 \pm 0.01) \times 10^{-2}$ |
| 40 | 2.019 ± 0.005 | $(2.65 \pm 0.01) \times 10^{-2}$ |
| 30 | 2.046 ± 0.005 | $(2.39 \pm 0.01) \times 10^{-2}$ |
| 20 | 2.066 ± 0.005 | $(2.04 \pm 0.01) \times 10^{-2}$ |
| 10 | 2.028 ± 0.003 | $(1.51 \pm 0.01) \times 10^{-2}$ |
| 5 | 1.941 ± 0.005 | $(1.03 \pm 0.01) \times 10^{-2}$ |
| 2 | 1.664 ± 0.011 | $(5.31 \pm 0.12) \times 10^{-3}$ |
| 1 | 1.375 ± 0.008 | $(3.33 \pm 0.13) \times 10^{-3}$ |

Interpretation. The CTR-C prediction — monotonic increase of both entropy and variance as dimensions are eliminated — is not satisfied. Entropy rises to a peak at 20 retained components and then declines; variance decreases monotonically throughout the cascade. This negative result is ontologically informative: PCA truncation does not implement CTR-C compression because it does not redistribute the information load of eliminated coordinates into the residual subspace. Instead, it discards variance-bearing coordinates without compensation, producing the characteristic signature of index deactivation (TSD-ND-I): progressive loss of structural information as the index frame is narrowed.

This confirms the classification of PCA established in Protocol 4 (§7.4), where PCA was identified as an approximate realisation of the access-deactivation operator A(d) (Spearman ρ = 0.837 ± 0.028, TSD-ND-I). Protocol 5 extends that result from synthetic manifolds to a standard ML benchmark, demonstrating that the IAS ontological discrimination remains valid on real data. The absence of the CTR-C signature in PCA truncation further validates the taxonomy: not every dimensionality-reduction algorithm is a compression operator; the IAS framework provides the formal apparatus to determine which operator family a given algorithm instantiates.

**Remark: what would satisfy CTR-C?** To satisfy the CTR-C criterion (monotonic entropy increase under component elimination), an algorithm would need to implement the redistribution operator P(d→m): the information load of each discarded principal component must be actively re-encoded into higher-order moments of the residual carrier, not merely discarded. PCA truncation fails this because it projects onto a subspace and discards the orthogonal complement without compensation. A CTR-C-compliant algorithm would require a learned encoder that maps eliminated variance into the surviving components' moment structure, analogous to the moment fibration of §3.2.2. The TSD-ND-I classification of PCA is stable across both synthetic (Protocol 4, Swiss Roll) and real (Protocol 5, MNIST) data. The distinction between PCA projection (tested in Protocol 4) and PCA truncation (tested in Protocol 5) is a distinction between two different operations, not an ambiguity in the classification of a single operation.

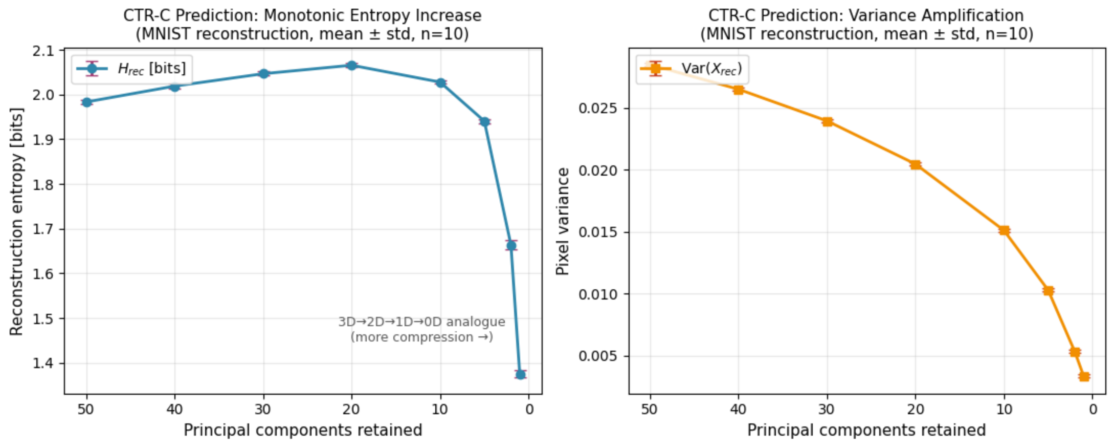


*Fig. P5. Ontological discrimination via PCA truncation on MNIST. Left: reconstruction entropy H_rec [bits] versus retained principal components (mean ± std, n=10); entropy peaks at ~20 PCs and declines, violating the CTR-C monotonic-increase prediction. Right: pixel variance Var_rec decreases monotonically throughout the cascade. Both panels confirm index-deactivation behaviour (TSD-ND-I) rather than compression.*

**Connection to §1.1.1.** The mapping of PCA to CTR-C in §1.1.1 was presented as a formal analogy (shared ontological commitment to coordinates-as-repositories). Protocol 5 refines this: PCA projection preserves the covariance structure and approximates TSD-ND-I index deactivation; PCA truncation (component elimination) is not equivalent to the redistribution operator P(d→m). The distinction between projection and truncation is therefore ontologically consequential.

## 7.6 Summary of Discriminating Tests

| Prediction | Model | Protocol | Result |
|---|---|---|---|
| m_max limits reconstruction | CTR-C | §7.1 | Confirmed: exact recon. at m≥2 (MSE $0.011 \pm 0.000 \times 10^{-6}$) |
| 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. |
| PCA truncation ≠ CTR-C signature | TSD-ND-I | §7.5 | Confirmed: entropy peaks at 20 PCs then declines; variance decreases monotonically. PCA truncation is index deactivation, not compression. (Here |

| | | | |
|---|---|---|---|
| | | | compression refers specifically to the IAS CTR-C operator — redistribution of eliminated information into residual moments — not to dimensionality reduction in the generic ML sense, under which PCA truncation is trivially a compression method.) |
| Monotonicity boundary: fails for non-separable | CTR-C | §7.2.1 | Confirmed: empirical, not universal |
| Threshold choice not arbitrary | TSD-ND-I | §7.4.1 | Confirmed: unanimous separation for threshold ∈ [0.55, 0.75] |
| Kernel PCA occupies the TSD-ND-I/CTR-C boundary | TSD-ND-I / CTR-C | §7.4.2 | Confirmed: Kernel PCA ($\gamma$=0.01) $\rho$=0.657±0.026, strictly between PCA and t-SNE/UMAP; gamma sweep shows a continuous transition, with $\gamma \to 0$ recovering TSD-ND-I |

## 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 | Non-dimensional | Indexation system. | 3→2: deactivation Z. | Substrate external to |

| | | | | | |
|---|---|---|---|---|---|
| | (Index) | substrate; coordinates = access indices. | Substrate external to metric. | 2→1: deactivation Y. 1→0: deactivation X. | 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 object. | 3→2: retraction of Z (return to plane). 2→1: plane → line (topology degradation). 1→0: line → topological code. | Topological code. 2D edges preserved. Requires 2D frame 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.

Within the IAS framework, the terminal state 0D is formally defined as the codified substrate: for CTR-C, the complete moment hierarchy (M_0,...,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 §7.6, 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 within the assumed single-sheet embedding class (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 ρ = 0.837 ± 0.028), consistent with the TSD-ND-I ontology, while non-linear embeddings (t-SNE: ρ = 0.403 ± 0.047; UMAP: ρ = 0.304 ± 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. IAS should therefore be interpreted as a classification framework rather than as a claim that all dimensional reductions share a single physical mechanism.

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

## 10. Limitations and Scope

The IAS framework is a taxonomic classifier, not a generative or algorithmic method. It does not produce new dimensionality-reduction algorithms, nor does it replace empirical hyperparameter selection. The following limitations apply:

**Mixed-method approaches.** Real-world algorithms (e.g., VAEs with both encoder and decoder) may instantiate elements of multiple IAS models simultaneously. The taxonomy assumes a single dominant ontological commitment; hybrid cases require explicit decomposition.

**Operational scope.** Protocols 1–5 test internal consistency and ontological discrimination, not empirical optimality on downstream tasks. A practitioner-facing decision procedure that maps task requirements to IAS model selection is left for future work.

**Physical testability.** The underlying ontological commitments (e.g., whether space is a carrier or an index frame) are not directly adjudicable by physical experiment. The numerical protocols establish computational distinguishability, not empirical confirmation of the substrate ontology itself.

**Completeness.** IAS does not prove that the three substrate–space modalities (§3.1) are the only ones that can be constructed, or that the resulting six model classes are logically exhaustive; it classifies the model space that follows from adopting this particular triple of fundamental roles (object, frame, index). This completeness claim is falsifiable in a specific sense, distinct from the physical non-falsifiability noted in §1.2: IAS would be falsified by a reduction mechanism that (i) cannot be classified under any of the six existing model classes, (ii) is not a composition of them, and (iii) genuinely instantiates a substrate–space role not captured by the object/frame/index triad.

**Reconstruction (moment problem).** CTR-C's re-instantiation is invertible only under additional assumptions on the moment sequence (Hausdorff/Carleman determinacy, MaxEnt selection at finite m; §6.1); the finite-m reconstruction tested in the protocols is not a proof of unique recoverability in general.

**Substrate-dependence.** Every invariance claim ($H(S)$ = const, substrate-preservation predicates) is relative to the substrate representation chosen for a given model or algorithm (§1.1.1, §6.2); changing that choice can change the classification.

### 10.1 Towards a Practitioner-Facing Decision Procedure

While a fully automated decision procedure is left for future work, the discriminating tests of §7 already constitute a reproducible two-stage diagnostic for classifying arbitrary dimensionality-reduction algorithms within the IAS taxonomy.

**Stage 1 — Substrate-invariance test.** Compute the Global Spearman ρ between original and embedded pairwise distances. If $\rho > 0.65$ (the working threshold established in §7.4.1), the algorithm approximates an access-deactivation operator $A(d)$ (TSD-ND-I family). If $\rho < 0.533$, the algorithm transforms the substrate and falls under CTR-C or TSD-ND-C.

**Stage 2 — Compression signature test.** Truncate the algorithm's representation progressively and measure the Shannon entropy $H_{rec}$ of the residual carrier. Monotonic increase under truncation indicates CTR-C compression; peak-and-decline or monotonic decrease indicates TSD-ND-I index deactivation (§7.5).

These two stages provide an operational bridge between the ontological taxonomy and algorithm selection without requiring new algorithmic development.

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

This appendix provides a self-contained computational realisation of the CTR-C redistribution operator on a finite lattice. The Lagrangian and discretisation below are phenomenological templates, not unique or axiomatically derived constructions; any other dynamics that encodes the same discrete cascade levels and threshold condition would serve equally well as an illustration.

### 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 $\sin^2$ 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 threshold in the CTR-C theory. Θ is introduced here as a phenomenological control parameter, fixed by hand and swept over the grid $\Theta \in [0.05, 0.40]$ in the phase-diagram scan of §A.3–A.4, not derived from a microscopic theory (see also §A.6 and §4.1.2). Formally, Θ plays the same role as the Lagrange multiplier β in the information bottleneck framework (Tishby, Pereira & Bialek, 1999) and the distortion budget in rate-distortion theory: a threshold that trades off compression against fidelity without itself being derived from first principles in either tradition. This places Θ's phenomenological status within an established methodological precedent rather than resolving it. Local fields:

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

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

The Lagrangian density $\mathcal{L}$ given in §A.2.3 is not unique. Any other functional that yields degenerate minima at integer ϕ and admits a threshold-driven transition between them would serve the same illustrative purpose. The exponential and linear forms of ε(ϕ) and S(ϕ) above are chosen to ensure monotonic decay of the cascade condition away from the 3D well; other sigmoidal or polynomial choices would produce qualitatively similar phase diagrams.

- ### A.2.3 Equation of Motion

The ϕ-field is identified with the zeroth moment $M_0$ of the local information density $\rho_\phi$ at each lattice site; this identification is a parameterisation choice made for computational tractability, not a theorem — higher moments $M_1$, $M_2$, … are not dynamical fields in this module but are reconstructed offline from the stationary density once the cascade equilibrates (§3.2.2, §4.1.2). The Lagrangian density is constructed a posteriori as one possible variational realisation among many, not derived from the IAS ontology: $\mathcal{L} = ½(\partial_\mu\phi)(\partial^\mu\phi) - V(\phi) - G\cdot\phi\cdot(\varepsilon(\phi)S(\phi) - \Theta)$, where $\mathcal{L}_{int} \equiv -G\cdot\phi\cdot(\varepsilon S - \Theta)$ couples the order parameter to the cascade condition. Requiring that the action $S = \int d^4x\, \mathcal{L}$ is stationary under variations δϕ ($\delta S/\delta\phi = 0$) yields the Euler-Lagrange equation of motion:

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

Discretisation proceeds in three steps. Step 1 (spatial grid): the continuum field ϕ(x,t) is replaced by its values $\phi_i(t)$ on lattice sites $x_i = i\cdot\delta x$, $i \in \{1,\ldots,L\}^3$ on the $L^3$ periodic lattice (L=24, §A.3). Step 2 (Laplacian): the continuum operator $\nabla^2\phi$ is replaced by the standard second-order 6-point (nearest-neighbour) stencil in 3D, $\nabla^2\phi|_i \approx (1/\delta x^2)\cdot\Sigma_{j\in nn(i)}(\phi_j - \phi_i)$, summing over the 6 lattice neighbours of site i under periodic boundary conditions — the leading-order finite-difference approximation to the continuum Laplacian, with truncation error $O(\delta x^2)$. Step 3 (time integration): the second time-derivative $\partial^2\phi/\partial t^2$ is replaced by the central-difference leapfrog (Störmer-Verlet) update $\phi_i(t+dt) = 2\phi_i(t) - \phi_i(t-dt) + dt^2\cdot[\nabla^2\phi|_i - f(\phi_i(t))]$, where $f(\phi) = dV/d\phi + G\cdot d/d\phi[\phi(\varepsilon(\phi)S(\phi)-\Theta)]$ collects the potential-gradient and cascade-source terms above. This scheme is explicit, symplectic (energy-conserving up to $O(dt^2)$ for the free part), and stable for $dt \lesssim \delta x$ (satisfied here: $\delta x = L_{box}/24$, $dt = 0.01$, §A.3); no additional approximation is introduced beyond these three steps. 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 |
| $\varepsilon_0$ | 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

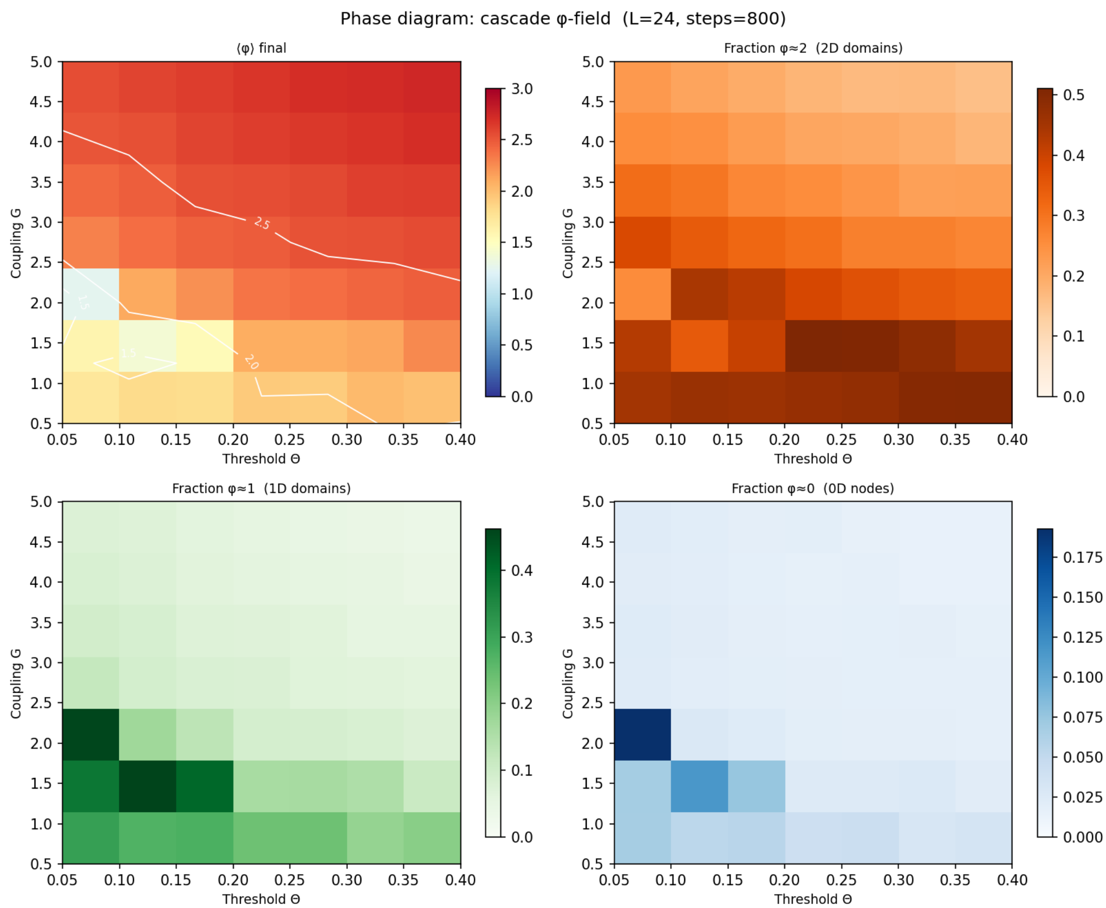


*Figure 1. Phase diagram of the cascade φ-field. Top-left: ⟨φ⟩_final with iso-contours at 1.5, 2.0, 2.5. Other panels: domain fractions $f_2$, $f_1$, $f_0$.*

- ### 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 | $G < 2$, $\Theta < 0.15$ | < 1.8 | 1D/0D domains dominate |
| II — Partial cascade | $G \approx 1$–$2$, $\Theta \approx 0.15$–$0.30$ | 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% across the full (Θ,G) grid, including the deepest-cascade point tested (G=0.50, Θ=0.050, $f_0$=0.07) and the maximum-cascade point (G=1.25, Θ=0.108, $f_0$=0.12). Full 0D reduction — $f_0 \to 1$ — is not reached at L=24, 800 steps under any parameter combination scanned; the phase diagram therefore demonstrates that a cascade tendency toward lower ϕ exists and strengthens under certain (Θ,G) regimes, not that the simulated system reaches full terminal convergence to the 0D state. Longer simulations or larger G would be required to determine whether $f_0$ saturates below 1 or continues to grow with run length.

## A.5 Interpretation within CTR-C

| Phase | CTR-C interpretation | Physical analogue |
|---|---|---|
| Regime I (⟨ϕ⟩ < 1.8) | Active cascade — reduction in progress | Near-Planck; quantum gravity effects |
| Regime II (⟨ϕ⟩ ≈ 2) | Transition zone — 3D/2D coexistence | Spectral dim. ≈ 2, cf. CDT/AS result |
| Regime III (⟨ϕ⟩ > 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 (⟨ϕ⟩ ≈ 2) reproduces the spectral dimension d_s ≈ 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
- $\Theta$ itself is a free phenomenological control parameter of this lattice module (swept by hand over [0.05, 0.40] in §A.3), not a value fixed by, or derived from, the CTR-C formalism of §3–§6; it is a cascade-trigger threshold logically independent of the post-cascade capacity bound $m_{max}$ (§6.1) — the two are related by sequential conditioning, not identity (§4.1.2)
- No back-reaction: $\phi$-field does not yet couple to matter or spacetime curvature
- Module 1.2:
    - $\Theta$ value is a free parameter and will attempt to constrain it from CMB data
    - Couple $\phi$ 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 $(\Theta, G)$ 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 $\Theta \cdot G \approx$ const. The phase boundary is hyperbolic, suggesting a simple analytic form for the cascade transition.

4. Connection to CDT/Asymptotic Safety. Partial-cascade regime ($\langle\phi\rangle \approx 2$) independently reproduces $d_s \approx 2$ 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")
```

## Appendix B: Principle of Topological Compatibility (Meta-Extension)

### B.1 Preamble: From Possibility to Selection

The IAS taxonomy, as developed in §1–§9, enumerates six mutually exclusive reduction mechanisms without prescribing which mechanism is realised in a space of given dimensionality d. The six models are ontologically prior to dimensionality: they describe how reduction occurs, not where it is permitted.

However, the mathematical structures underlying each model carry implicit topological constraints that depend on d. A mechanism that is formally definable for arbitrary d may become physically degenerate or topologically unstable in certain dimensions. This observation motivates the following meta-principle.

**Epistemic status of this appendix.** The core taxonomy of §1–§9 is deliberately classificatory: as stated in §1 and §1.2, adjudication among the six models proceeds by conceptual clarity and internal consistency, not by physical experiment, and the paper explicitly disclaims empirical validation of a physical ontology. What follows departs from that register by design: it is an explicitly speculative, physically-motivated meta-extension that asks which of the six models could be jointly instantiated in a space of a given dimensionality, drawing on established results in topology and dynamics (not new physics) to propose an answer for our own d = 3. It is offered as a separately-flagged hypothesis about physical applicability, outside the paper's classificatory core; §B.6 states its status explicitly.

### B.2 Formulation of the Principle

**Principle of Topological Compatibility (PTC).**

For each spatial dimensionality d ≥ 2, the set of IAS models that can be physically instantiated is a proper subset of the full six-model taxonomy. Membership in this subset is determined by three topological criteria:

- **Embedding criterion.** For the Class-C (Ambient-as-operator) model TSD-2D-C specifically, the ambient dimension d must strictly exceed the intrinsic dimensionality k = 2 of the substrate: d > k. If d ≤ k, the metric superstructure required by TSD-2D-C's own retraction operator U(d→k) (§6.3) does not exist — there is no “extra” ambient axis for the operator to retract. This criterion is specific to TSD-2D-C's carrier-transformation mechanism; it does not constrain the Class-B (Object-invariant) models CTR-I, TSD-ND-I, TSD-2D-I, whose access operator A(d) (§6.2) addresses projections of the substrate rather than an ambient superstructure and remains well-defined for d ≤ k.

- **Structural stability criterion.** The reduction mechanism must preserve topological invariants (knotting, linking, entanglement structure) under the cascade. If d is high enough that these invariants trivialise, the mechanism becomes degenerate.
- **Dynamical stability criterion.** The re-instantiation dynamics from the 0D terminal state must admit stable orbits or bound states in dimension d. If d is too large, the effective potential governing re-instantiation becomes non-confining.

### B.3 Compatibility Matrix

| d | CTR-C | CTR-I | TSD-ND-C | TSD-ND-I | TSD-2D-C | TSD-2D-I | Excluded models (rationale) |
|---|---|---|---|---|---|---|---|
| 2 | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | TSD-2D-C: embedding criterion violated (d = k) |
| 3 | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | None; all six are topologically viable |
| 4 | ? | ✓ | ✓ | ✓ | ✓ | ✓ | CTR-C: knotting trivialises; redistribution may lose structural stability |
| 5 | ? | ✓ | ✓ | ✓ | ? | ✓ | CTR-C: same; TSD-2D-C: superstructure overabundance weakens retraction uniqueness |
| ≥ 6 | ? | ✓ | ? | ? | ✗ | ? | TSD-2D-C: Bonnet reconstruction fails globally for generic embeddings; CTR-C: highly unstable |

Asterisk/question-mark convention: “?” denotes a case not resolved by the arguments in this appendix — the model is formally definable but its topological viability at that d has not been established either way; see §B.6.

### B.4 Why d = 3 is Distinguished

The dimensionality d = 3 occupies a unique position in the compatibility matrix: it is the only dimension in which all six models are simultaneously topologically viable. This is not a coincidence, but a consequence of several well-established mathematical facts:

- **Knotting and linking.** In d = 3, knots and links are non-trivial (Alexander, Jones, HOMFLY polynomials are non-degenerate). In d = 4, all knots trivialise via ambient isotopy; in d = 2, knotting is impossible. CTR-C relies on the non-trivial entanglement of information channels; its physical content is richest precisely where topology is non-trivial.
- **Holographic duality.** The AdS/CFT correspondence and its generalisations require a (d − 1)-dimensional boundary for a d-dimensional bulk. For d = 3, the boundary is 2D — the setting of TSD-2D-I/C. For d = 4, the boundary is 3D, and the projection interface of TSD-ND-I gains an additional degree of freedom that complicates the information-preservation condition H(S) = const.
- **Orbital stability (Laplace).** In Newtonian gravity, bound stable orbits exist only for $1/r^2$ forces, which are the Green's functions of the Poisson equation in d = 3. In d ≥ 4, the effective potential is too steep or too shallow for generic stability. TSD-2D-C, which requires a 2D substrate to re-instantiate into a metric frame, depends on the stability of the embedding dynamics; this stability is maximised in d = 3.

### B.5 Physical Interpretation

The PTC does not alter the formal definitions of §6. Rather, it adds a selection layer: the IAS framework describes what can happen, while the Principle of Topological Compatibility describes what does happen in a given d.

In our observed universe (d = 3 spatial dimensions), the full taxonomy is available. This explains why competing ontologies of space — emergent (CTR), holographic (TSD-2D), and relational (TSD-ND) — can all claim partial empirical support: they are not mutually exclusive descriptions of the same phenomenon, but alternative, topologically permitted realisations of reduction in the one dimension where all are stable.

Conversely, in d = 2 or d ≥ 4, the taxonomy collapses to a smaller set, potentially explaining why hypothetical higher-dimensional or lower-dimensional physics appears to select simpler ontologies (e.g., purely index-based or purely compressive).

### B.6 Limitations and Status

The PTC is presented here as a meta-hypothesis, not as a theorem derivable from the IAS axioms. Its connection to the core taxonomy is conjectural: it requires rigorous proof that knot trivialisation in d = 4 indeed forbids CTR-C, or that holographic bounds in d > 3 destabilise TSD-ND-I. These proofs lie outside the scope of the present work and are left for future investigation.

The IAS framework remains valid in all d as a formal classifier; the PTC adds a physical filter that may guide the selection of models in concrete applications.

**Key implication (conjectural, §B.6):** if the PTC holds, the question "Why does our universe appear to be 3D?" receives a structural answer: d = 3 is the unique dimensionality in which the full spectrum of ontological commitments about space is simultaneously realisable. Lower or higher dimensions would force a premature collapse of the taxonomy. This answer is offered as a hypothesis motivated by the taxonomy, not as an empirical result established by it.