Informational Architectonics of Space

Six Models of Dimensional Reduction

IAS Framework — Formal Taxonomy

AL_78 | Independent Research Group | September 2026

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 a representation sufficient for re-instantiation under the stated reconstruction assumptions. 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) formal reversibility of metric retraction under preservation of the required embedding data, (IV) substrate invariance under index deactivation, shown to hold on isotropic synthetic data and to require substrate isotropy (§7.4.3–§7.4.4), and (V) persistence of the predicted operational distinction in a real-data benchmark (MNIST). Five of these protocols (§7.1–§7.5) test internal consistency among the six ideal types via discriminating operational predictions — they establish that the taxonomy's distinctions are non-degenerate, not that any algorithm outperforms another. A sixth protocol (§7.6) serves a categorically different purpose: it demonstrates, on a single hybrid case (a β-VAE), that the fingerprint vector proposed for mixed ontological commitments is directly computable — not a discriminating test, a benchmark, or evidence of practical utility for algorithm selection. 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. A further test (§7.8) provides an operational negative-control check against circular

classification in the tested setting: candidates with no principled relation to any of the six operators return a null result rather than a forced classification.

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. Mutual exclusivity holds at the level of ontological interpretation of an ideal type, not at the level of the observable signatures a composite procedure may exhibit; a procedure combining several operational steps may satisfy the discriminating predicate of more than one model without contradiction, since each predicate is tested against its own step rather than against a single shared observable (see the hybrid fingerprint discussion in §1.1.3).

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→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.

Definition (Structural Analogy). Let $\mathcal{A} = (A_1, \ldots, A_n)$ be the operations of an algorithm, and let $A_r \in \mathcal{A}$ denote the *reduction step* — the operation identified by the algorithm's own specification (not chosen post hoc by the analyst) as mapping the d-dimensional representation to the k-dimensional one, $k < d$. Let $\mathcal{O}$ be an IAS operator with discriminating property $\Phi_\mathcal{O}$. Then $\mathcal{A}$ is a *structural analogy* to $\mathcal{O}$ if and only if A_r satisfies $\Phi_\mathcal{O}$ under the substrate representation S chosen for the algorithm. $\mathcal{A}$ is an *isomorphism* to $\mathcal{O}$ if and only if, additionally, there exists a bijection ϕ between $\{A_1,\ldots,A_n\}$ and a

generating set of $\mathcal{O}$ that preserves composition and invertibility in both directions. Isomorphism therefore implies analogy but not conversely; anchoring the test to the algorithm's own reduction step (rather than to an arbitrary subset of operations chosen after the fact) is what makes the criterion checkable rather than satisfiable by construction.

Worked example: PCA truncation. Let $\mathcal{A}$_PCA = (A_1, A_2, A_3), where A_1 = covariance projection onto the leading eigenvectors, A_2 = truncation (discarding components k+1,…,d), and A_3 = inverse transform with zero-filling. The reduction step is A_2: A_1 is a change of basis at fixed dimension, and A_2 is the operation that actually removes dimensions. *Testing A_2 against Φ_P* (moment redistribution; monotonic entropy increase under elimination): Protocol 5 (§7.5) shows Var_rec decreasing monotonically and H_rec peaking then declining under truncation, rather than rising — Φ_P fails. $\mathcal{A}$_PCA is therefore not a structural analogy to P (CTR-C), confirming §1.1.1's and §7.5's conclusion as a checkable procedure rather than a verbal one. *Testing A_2 against Φ_A* (substrate invariance, H(S) = const): the absence of compensating redistribution is itself the index-deactivation signature identified in §7.5's "Remark: what would satisfy CTR-C?" — consistent with the existing TSD-ND-I classification of PCA truncation. This worked example is illustrative, not exhaustive: the definition above is stated generally enough to apply to other algorithms in Table 1, but the paper does not carry out that check for every entry.

Worked example: Autoencoder bottleneck. Let $\mathcal{A}$_AE = (A_1, A_2), where A_1 = encoder (mapping the input to a lower-dimensional latent code through the bottleneck) and A_2 = decoder (mapping the latent code back to the input space). The reduction step is A_1: the bottleneck is where dimensionality is actually reduced, and A_2 acts on the already-reduced representation. *Testing A_1 against Φ_P* (moment redistribution; information-theoretic compression with retention, not deletion): the encoder minimises reconstruction error by re-encoding the variance eliminated at the bottleneck into the surviving latent coordinates rather than discarding it without compensation — Φ_P holds on A_1 in this sense. No bijection exists between A_1 and P: the decoder A_2 is a learned, approximate inverse, not an algebraic one, so composition and invertibility are not preserved. *$\mathcal{A}$_AE is therefore a structural analogy to CTR-C, not an isomorphism* — confirming the classification already given above via the same checkable procedure used for PCA truncation, rather than by verbal assertion alone.

PCA (Pearson 1901) projects data onto the eigenvectors of the covariance matrix, ordered by explained variance. Two distinct operations are both called "PCA" below and should not be conflated: PCA as a projection (computing the full set of principal components, used in the CTR-C analogy immediately following) and PCA truncation (discarding some components, the operation actually tested in Protocol 5, §7.5, against TSD-ND-I). The IAS mapping to CTR-C is a formal analogy, not an isomorphism: PCA operates on the covariance structure of the data distribution, not on a physical information-density field ρ_d. However, the shared ontological structure is precise — PCA treats principal coordinates as carriers of variance; truncation removes the selected variance-bearing coordinates rather than redistributing their information into a residual moment hierarchy (Protocol 5, §7.5, confirms this: Var_rec decreases monotonically under truncation, rather than the increase a redistribution signature would predict). The moment-field operator P(d→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 fails at the operational level (a specific testable prediction is falsified; the ontological commitments themselves are not directly

falsifiable — see §1.2): 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. In short: IAS does not assign PCA as a whole to a single model — the covariance-projection step retains, at most, a formal analogy to CTR-C, while the tested operation, component truncation, is classified as TSD-ND-I-like (Protocol 5), not CTR-C-like.

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 ($\rho \geq 0.65$ — not a universal statistical threshold, but the dataset-specific 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 discriminating operational predictions about algorithm behaviour. Both non-linear embeddings violate the substrate-invariance prediction H(S) = const; UMAP does so more strongly than t-SNE, indicating that its topological optimisation actively restructures the global distance graph. Only linear projection (PCA, $\rho = 0.837 \pm 0.028$) approximates the access-deactivation operator A(d). The tested non-linear embeddings (t-SNE, UMAP) violate H(S) = const, consistent with the conjecture stated below.

Conjecture (substrate transformation under non-linear distortion). Under the IAS substrate representation used in Protocol 4, a non-linear transformation that changes the global distance structure changes the substrate rather than merely deactivating the index frame. The tested embeddings (t-SNE, UMAP) provide operational evidence consistent with this conjecture; a general proof independent of algorithmic instantiation remains open.

The Information Bottleneck (Tishby et al. 1999) finds the minimal sufficient statistic of X for predicting Y, subject to a rate-distortion constraint. This provides a structural analogy to TSD-ND-C, not a formal equivalence: the training objective (ambient) determines which components of the data cluster survive compression. The rate-distortion trade-off β plays a role analogous to the capacity constraint m_{max} (§6.1) — higher β allows more moments to be retained (larger m_{max}), lower β forces more aggressive compression — but this is a qualitative correspondence of function, not a derived mathematical identity between β and m_{max}. Unlike CTR-C, the ambient (loss function) — not the object itself — is the active agent of compression, consistent with the TSD-ND-C ontology.

Positioning relative to established ML theory. IAS is orthogonal to, not competing with, information geometry, manifold learning, and rate-distortion theory. Manifold learning asks how to recover low-dimensional structure from data; information geometry studies the geometry of statistical model families; rate-distortion theory quantifies information loss under a coding constraint. IAS asks a different, prior question: what substrate–space relation does a given reduction procedure presuppose, independently of whether it succeeds at any of the above tasks. A method's success by manifold-learning or rate-distortion criteria is compatible with any of the six IAS ontological commitments; IAS classifies the commitment, not the performance. Where existing frameworks classify by geometry, representation,

coding efficiency, or statistical structure, IAS classifies by the ontological role assigned to substrate and spatial frame during reduction; this is offered as a difference in the axis of classification, not as a claim that no prior ontological classification of dimensional reduction exists.

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 (Kaluza 1921; Klein 1926), extra dimensions are compactified below a length scale L_c, rendering them unobservable at low energies. This is discussed only as a structural analogy to metric retraction, not as a claim that IAS's retraction operator is instantiated in physical spacetime: physical spacetime is not, in general, a Euclidean manifold equipped with an information-density field ρ_d in the sense §3.1 defines. Under this analogy, the ambient (spacetime geometry) adds axis Z as a metric superstructure; compactification is analogous to the retraction operator U(3→2). The reversibility of U — a formal property established in §6.3 for the IAS operator — is noted here only as a structural parallel to the decompactification limit in string theory, not as a derivation of one from the other.

In Causal Dynamical Triangulations (CDT; Ambjørn, Jurkiewicz & Loll 2004) and Asymptotic Safety (Reuter & Saueressig 2012), the spectral dimension d_s(σ) decreases smoothly from ~4 at large scales to ~2 at Planck scales via a diffusion-based random walk (Ambjørn, Jurkiewicz & Loll 2005). 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. As a conjectural operational prediction — not a consequence derived from CDT dynamics itself, but a candidate test suggested by the IAS classification — the framework proposes that CDT should exhibit monotonic entropy increase in residual channels: a prediction not previously stated in the CDT literature and not yet derived from the CDT action or path-integral measure.

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; Ryu & Takayanagi 2006), 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 ϕ on a flat periodic lattice rather than on a minimal surface in a curved bulk geometry, Θ 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. None of the physics mappings in this section constitute a physical test of the IAS ontologies; they identify structural correspondences between the IAS formalism and existing physical frameworks. Physical adjudication between the ontological commitments themselves — e.g., which experiment would favour CTR-C over TSD-ND-C — would require an independently specified physical state space and observable, which this paper does not provide.

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. PCA projection and PCA truncation are listed as separate rows because they are distinct operations with distinct classifications (§1.1.1, §7.5); the taxonomy classifies the operation performed, not the algorithm name as a whole.

ML Algorithm	IAS Model	Ontological Commitment
PCA projection (covariance step)	formal analogy to CTR-C only	Analogy only; not itself tested as TSD-ND-I or CTR-C (see §1.1.1)
PCA truncation (as tested)	TSD-ND-I proxy	Coordinate projection without substrate modification (tested in Protocol 5, §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, showing behaviour consistent with its classification as TSD-ND-I rather than CTR-C under the tested observable.

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 (Kingma & Welling 2014) 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 $f = (w_{CTR\text{-}C}, w_{CTR\text{-}I}, w_{TSD\text{-}ND\text{-}C}, w_{TSD\text{-}ND\text{-}I}, w_{TSD\text{-}2D\text{-}C}, w_{TSD\text{-}2D\text{-}I})$, 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; §7.6 (Protocol 6) provides a worked example computing f for a β-VAE (Higgins et al. 2017). Hybrid classification cannot yield contradictory predictions, because each w_j is computed independently against its own discriminating predicate (§6): a step either satisfies signature j's criterion or it does not, and a non-zero w_j for one signature does not require a zero w_k for another, since the predicates are not mutually exclusive tests on the same observable but separate tests on separate operational steps. This extension preserves the taxonomic clarity of the six ideal types while accommodating the compositional complexity of contemporary methods. Computing f directly from a single data matrix, without prior knowledge of the generating algorithm — inferring the fingerprint from the transformation's observable effect alone, rather than from Protocols 1–5 applied to a known algorithmic step — is outside the present scope: fingerprinting as defined here operates over known operational steps, not over input-output behaviour inferred without access to those steps, so a black-box variant would require a different construction, not an extension of the one given.

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. A classification's operational signature is thereby challenged — independently of the ontological commitment it instantiates — whenever it fails to distinguish the predicted behaviour from an appropriate null or baseline case, or when two models classified as ontologically distinct yield observationally identical predictions under the same intervention; §7.4's PCA-vs-non-linear separation and §7.5's PCA reclassification are instances of this operational scrutiny already applied. 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.

Epistemic architecture, stated once. Every distinction drawn throughout this paper — analogy vs. isomorphism (§1.1.1), ontological vs. operational (above), protocol vs. benchmark (§7), classification vs. algorithmic superiority — reduces to the same two-level architecture. Level I (ontological): the six substrate–space commitments are adjudicated by conceptual clarity and internal consistency, not physical experiment. Level II (operational): each commitment entails a discriminating observable (§6), tested by consistency protocols (§7) that establish non-degeneracy — computational evidence bears on Level II, never directly on Level I. Where a later section restates this distinction in local terms (e.g., "analogy not isomorphism", "proxy not substitute"), it is applying this same architecture to that section's specific operator or observable, not introducing a new epistemic principle.

---

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 on a common formal limit: absence of spatial realization, while differing in the ontological status and encoding of that limit. 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 phenomenological curvature threshold K_P (see below), 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. The name "Planck curvature threshold" is descriptive rather than derived: K_P is not computed

from any quantum-gravity theory here, but is introduced, in parallel with the phenomenological status already noted for Θ (§A.2.2, §A.6), as a free parameter marking where this formal mechanism activates.

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:

$$M_n(x,y) = \int z^n \cdot \rho_{XYZ}(x,y,z)\, dz , \quad n = 0, 1, 2, \ldots$$

Capacity constraint (closes the internal constraint of CTR-C): $m_{max} = \lfloor (C_{channel} - H[\rho_{XY}^{(o)}]) / H_{avg} \rfloor$ where $C_{channel}$ is the Shannon capacity of the residual channel, measured in bits per lattice site (formally, $C_{channel} := \sup_{p(x)} I[X;Y]$; full definition in §6.1); in the protocols of §7 it is treated as an externally specified sweep parameter, not estimated from data. 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.

For the tested field class, preservation of moments M_n up to order m equal to the number of eliminated dimensions provided an operationally sufficient reconstruction condition; this is not established as a general requirement for arbitrary ρ_d (see the moment-problem caveat in §6.1). For finite m, reconstruction yields the maximum-entropy estimate (Jaynes 1957) 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 → 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.

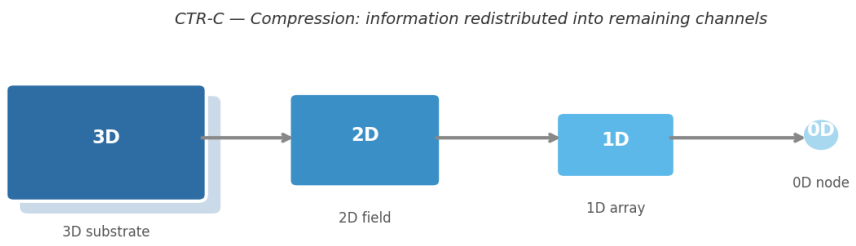


Fig. 1. CTR-C: cascade 3D→2D→1D→0D. Information from each eliminated axis is redistributed into remaining channels-repositories.

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 ϕ 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(ϕ), threshold functions ε(ϕ), S(ϕ), and numerical scheme (leapfrog integration; Hairer, Lubich & Wanner 2003) 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 ε(ϕ)·S(ϕ) ≥ Θ is a dynamical cascade-trigger threshold, logically independent of the post-cascade capacity bound m_max (§3.2.2, §6.1): Θ 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 analogy $\phi \approx M_0$ used to run the module is likewise a parameterisation choice made for computational tractability, not a theorem — ϕ is a discrete field of effective dimensionality and M_0 a continuous density projection, and treating one as standing in for the other is a computational convenience, not a claim that dimensionality (a topological invariant) and density (a measure) are the same kind of object. 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 (Θ, G) of the cascade condition, identifying the regimes in which compression actively drives ϕ toward lower-dimensional attractors.

4.2 CTR-I — Index Model

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

Object definition in 3D.

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

CTR-I avoids the carrier-capacity problem of CTR-C by adopting a different ontological premise, not by improving on it: 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.

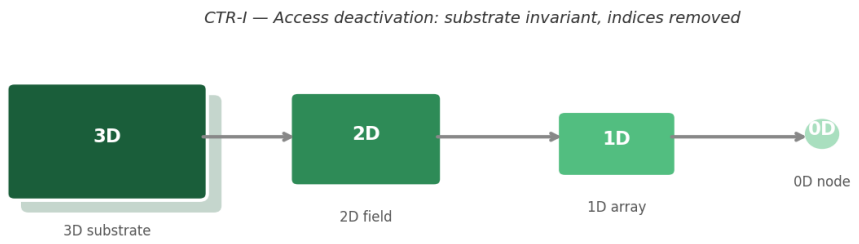


Fig. 2. CTR-I: coordinate indices are deactivated sequentially; the non-dimensional substrate remains unmodified throughout.

---

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.

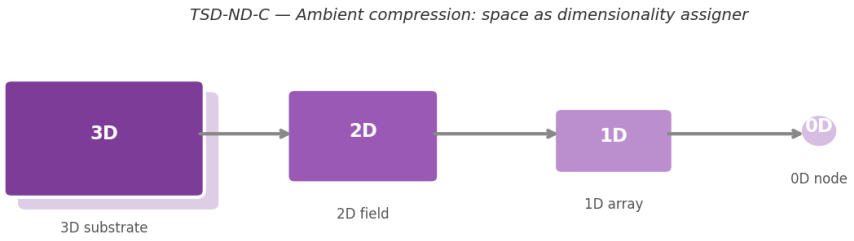


Fig. 3. TSD-ND-C: space as dimensionality assigner; non-dimensional cluster compressed through ambient action.

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.

Information Measure for this model. Following the substrate scaffold of §3.1, the Information Measure of this model's substrate is canonically taken to be S := Adj(G), the graph's adjacency structure (which nodes connect to which), with H(S) := H(Adj(G)). This is not one illustrative choice among several equally valid options; it is the encoding fixed for this model instantiation, against which the operator-family criterion of §6.2 (S′ = S for A(d), S′ ≠ S for P, U) is evaluated throughout §6.2's worked example and Protocol 4 (§7.4).

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.

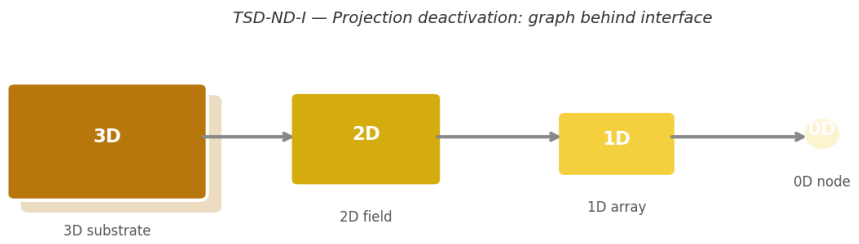


Fig. 4. TSD-ND-I: non-dimensional graph behind projection interface; deactivation of projections does not affect substrate.

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.

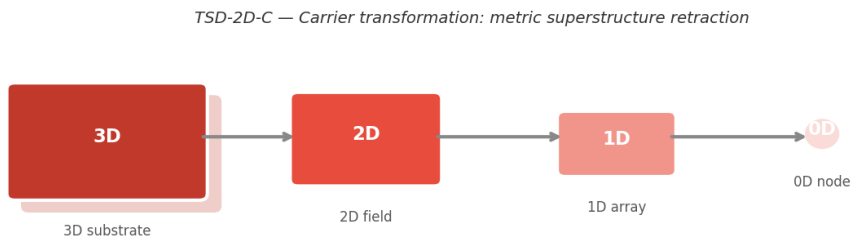


Fig. 5. TSD-2D-C: Z as metric superstructure; retraction is reversible, subsequent degradation is not.

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

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

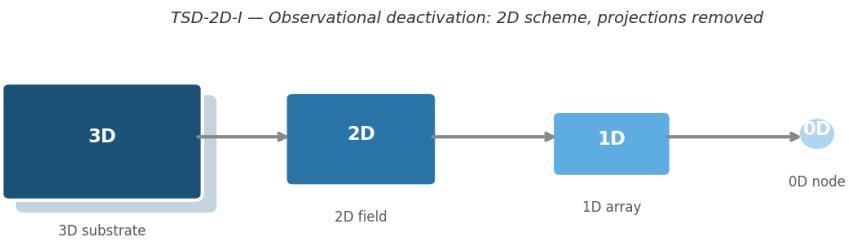


Fig. 6. TSD-2D-I: 2D scheme observed through spatial projections; deactivation of all projections yields the scheme without observer.

---

6. Meta-Formalism and Operator Typology

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

Reduction Type	Operator	Formal Definition	Reversibility	Information Property	Applicability
Compression	$P_m^{(d\to k)}$, $k=d-1$	$M_n(x\perp) = \int x_d^n \cdot \rho\, dxd$, $n=0..m$ Residual carrier holds moment fibration $M^{(d-1)}$	Conditionally reconstructible (exact as $m\to\infty$; finite m needs added constraints, e.g. MaxEnt)	$H[M^{(d-1)}] \geq H[\rho^{(0)}_d_1]$ (empirical entropy signature)	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 $A(0) = 0 \in \mathbb{R}^o$	Yes (reactivation of index set)	$H(S) = \text{const}$ (substrate invariant)	CTR-I, TSD-ND-I, TSD-2D-I
Carrier transformation (retraction)	$U(3\to2)$	$g^{(3)} \to g^{(2)}$: eliminate z-row/col Inverse: U^{-1} needs $(g^{(2)}, K_ij)$	Yes, with second fundamental form K_ij (Bonnet theorem, given Gauss–Codazzi compatibility; do Carmo 1976)	Full 2D topology preserved	TSD-2D-C
Carrier transformation (degradation)	$D(2\to1\to0)$	$G \to L_\sigma$ (linearisation); $C_top = (Adj(G), \Sigma_cycles, \chi)$ extracted from G directly Cycles destroyed at $2\to1$ step	No (topology lost at $2\to1$)	Topological code preserves re-instantiation instruction	TSD-2D-C
Orientation inversion	$\Xi \cdot A(d)$	$A_\Xi(d) = \Xi\cdot A(d)$ $\Xi \in O(d)\backslash SO(d)$ $\det(\Xi) = -1$ Example: $\text{diag}(1,1,-1)$	No (orientation reversed; no inverse within SO(d), though $\Xi^{-1} = \Xi$)	$H(S) = \text{const}$ (substrate invariant; only index	CTR-I (extended)

	frame changes)

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 an ontological 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.

Towards axiomatic closure. The definitions above already fix domain, codomain, and reversibility for each operator (table above; §6.1–§6.3); what remains open is a fully closed algebraic calculus. Summarising what is established and what is not: P_m^(d→k): S_d → M_m has domain the space of substrate density fields on d coordinates and codomain the order-m moment fibration on the residual k coordinates; it is a partial map (admissible whenever the field admits the stated moments) and its partial inverse (maximum-entropy reconstruction) is unique only under the sufficient condition given in §3.2.2, not in general. A(d): S → $\mathbb{R}$^d is a total, everywhere-defined coordinate projection with partial inverse A(k)$^{-1}$ existing only up to the re-instantiation ambiguity of the deactivated indices (§6.2). U(3→2) is defined on single-sheet embeddings satisfying Gauss–Codazzi compatibility (§6.3) and is invertible exactly on that domain. Composition across operator families (e.g. P then A) is not treated as a closed operation here — each protocol applies a single operator to a fixed substrate rather than chaining operators — and extending the calculus to closed composition is left for future work; the classification claims of §7 do not depend on such closure.

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

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

Lemma (entropy monotonicity — sufficient condition). Let ρ_d factorise as ρ_d(x⊥, x_d) = ρ^(0)_{d-1}(x⊥) · q(x_d | x⊥), where q is the conditional density along the eliminated axis. If the conditional variance Var_{x_d}[q(·|x⊥)] is non-zero on a set of positive measure in x⊥, then the redistribution operator P_m^(d→k) with m ≥ 1 satisfies H[M^(k)] > H[ρ^(0)_{d-1}]. If q is uniform on its support and independent of x⊥, then H[M^(k)] = H[ρ^(0)_{d-1}] (saturation). Proof sketch: a non-uniform q introduces additional degrees of freedom into the moment fibration M_1, M_2, …, which — by the maximum-entropy property used in reconstruction (§7.1) — strictly increase the entropy of the residual carrier beyond the zeroth-moment baseline. The equality case follows because a uniform conditional contributes no information beyond the marginal. Remark: this condition is sufficient, not necessary. The non-separable field tested in §7.2.1 violates monotonicity because cross-axis coupling redistributes

information in a manner not captured by this marginal–conditional decomposition; a general necessary condition remains open.

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.

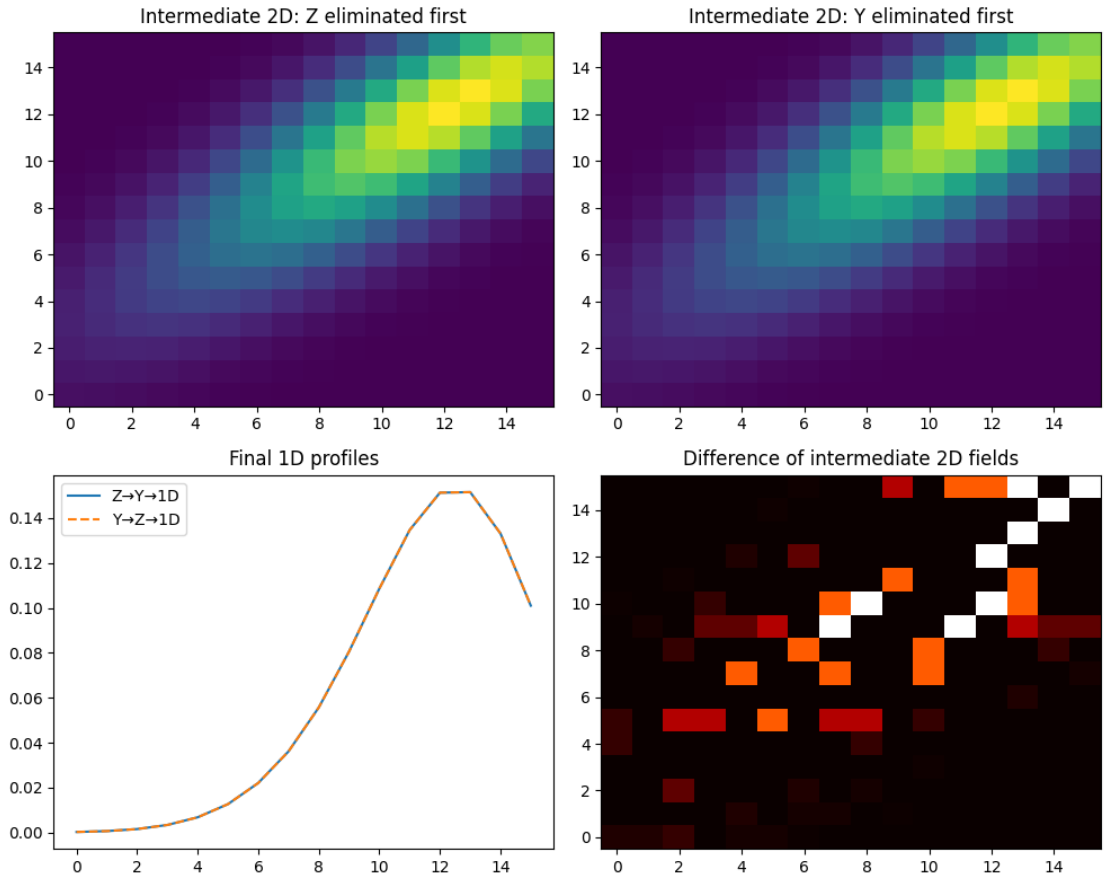


Fig. P1b. Path-dependence of the redistribution operator P (§6.1). Non-separable field (16^3 lattice, X–Y–Z coupled). Top row: intermediate 2D fields after first elimination step (Z-first vs Y-first). Bottom-left: final 1D profiles (Z→Y→1D vs Y→Z→1D) — coincident for this field class. Bottom-right: absolute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

Let ρ_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 → k, axis x_d is eliminated. Let $x\perp = (x_1,\ldots,x_k)$. The operator maps ρ_d to a moment fibration over the residual plane:

$$P_m^{(d\to k)}: \rho_d \;\to\; M^{(k)} = (M_0, M_1, \ldots, M_m)$$
$$M_n(x\perp) = \int_{-\infty}^{+\infty} x_d^n \cdot \rho_d(x\perp, x_d)\, dx_d\,, \quad n = 0, 1, \ldots, m$$

Discrete form. On a finite lattice, as used throughout §7, the integral above reduces to a sum over the eliminated axis's discrete index j:

$$M_n(x\perp) = \Sigma_j \; x_d(j)^n \cdot \rho_d(x\perp, j) \;, \quad j = 1, \ldots, N_d \;, \quad n = 0, 1, \ldots, m$$

where N_d is the number of lattice sites along the eliminated axis and x_d(j) is the coordinate value at site j. This is the form actually evaluated by Protocols 1, 2, and 5 (§7.1, §7.2, §7.5); the continuous integral above is the formal definition it discretises.

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

Uniqueness and determinacy of the moment reconstruction. Recovering ρ_d(x⊥, x_d) along x_d from the finite moment sequence (M_0,…,M_m) is an instance of the classical moment problem, and its solvability is not automatic. For the truncated Hausdorff moment problem (Shohat & Tamarkin 1943; 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 → ∞; 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,…,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 (Σ M_{2n}^{-1/2n} = ∞ is sufficient for uniqueness) and, more restrictively, the Hamburger moment problem's own existence conditions (Akhiezer 1965; 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≥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 ρ_d uniquely in general.

Applicability condition (capacity constraint). The capacity constraint below is an operational closure condition — a phenomenological bit-budget rule — not a bound derived from an information-theoretic optimisation problem. If residual channels X, Y have finite information capacity C_channel, the maximum preservable moment order is (m_max denotes the highest additional moment order retained beyond M_0, so m_max=2 means the triple M_0,M_1,M_2 is preserved, not merely two moments):

$$m_max = \lfloor \; (C_channel - H[\rho^{(0)}_XY]) \; / \; H_avg \; \rfloor$$

Definition (C_channel). In the abstract sense, C_channel := sup_{p(x)} I[X;Y] is the Shannon channel capacity of the residual carrier (Cover & Thomas 2006), 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, Belghazi et al. 2018; 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, ∞ 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 (Shannon 1948) of the n-th moment field M_n(x⊥) computed over the residual carrier. The average moment entropy is computed over the prescribed theoretical moment hierarchy: H_avg = (1/m_t) Σ_{n=0}^{m_t-1} H[M_n], where m_t = d − k is the model-selected hierarchy depth (the number of eliminated dimensions; m_t = 1 for the canonical single-axis elimination, k=d−1, in which case H_avg reduces to H[M_0]). 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[ρ_XY^(0)] is subtracted — i.e., a literal bit-budget, not merely a notational placeholder.

Note on the m_t = 1 reduction. For the canonical single-axis elimination, H_avg = H[M_0] and, since M_0 is by definition the zeroth-moment field (M_0 ≡ ρ^(0)), this coincides exactly with the baseline H[ρ_XY^(0)] already subtracted in the capacity constraint — as confirmed numerically in the Protocol 1 worked example below, where both quantities equal 2.000 bits. This is not a degeneracy: the two occurrences of the same quantity enter the m_max formula differently (once as a subtrahend, once as a divisor), so the formula simplifies to m_max = ⌊C_channel/H[M_0] − 1⌋ in this case rather than becoming undefined.

Multi-axis cascades. When the cascade is executed as sequential single-axis eliminations, C_channel, H[ρ^(0)], and H_avg are re-evaluated at each step on the residual carrier produced by the previous step, since the moment fields at step n+1 are defined on the already-reduced substrate from step n. For a direct multi-axis reduction (e.g. 3D → 1D performed in one operator application rather than as two single-axis steps), m_t equals the total number of eliminated dimensions and H_avg is computed over the full hierarchy of moments generated by all eliminated axes simultaneously, not per-step. The two procedures are not asserted to be equivalent; which one applies depends on how the operator P(d→k) is instantiated, and is left as an explicit modelling choice rather than resolved here.

Probability space for H[M_n]. Each moment field M_n(x⊥) 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] = −Σ_k p_k log p_k. This value-histogram estimator — not a position-indexed p_i ∝ 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 = ⌊(C_channel − H[ρ_XY^(0)]) / H_avg⌋ is not bounded below by zero: if the assumed channel capacity C_channel falls below the zeroth-moment baseline H[ρ_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[ρ_XY^(0)] = 2.0 bits used in Protocol 1 (§7.1),

m_max = ⌊(1.0 − 2.0)/H_avg⌋ = −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 ε(ϕ)·S(ϕ) ≥ Θ 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: Θ-trigger → P(d→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 Θ, which belongs to the dynamical, not the information-theoretic, sector of the model.

Empirical entropy signature of CTR-C: the Shannon entropy of the moment fibration satisfies H[M^(k)] ≥ H[ρ^(0)_{d-1}], where ρ^(0)_{d-1} is the marginal prior to redistribution. The entropy of the residual carrier increases monotonically for the field classes tested below — 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 ρ_d; a field already at maximum entropy prior to redistribution, for instance, admits no further increase. A general proof, with explicit conditions on ρ_d under which monotonicity holds, is left for future work. As a working conjecture, we expect monotonicity to hold whenever ρ_d is not already at the maximum-entropy distribution consistent with its marginal ρ^(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 ρ_d factorises as ρ^(0)_{d-1}(x⊥)·u(x_d) with u uniform on its support, the eliminated axis carries no such structure and H[M^(k)] = H[ρ^(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∈[0,L_x], y∈[0,L_y], z∈[0,L_z] with a uniform information density ρ_d = ρ_0 (constant). Eliminating axis Z: M_0(x,y) = ∫_0^{L_z} ρ_0 dz = ρ_0 L_z (a constant field over the residual XY plane); M_1(x,y) = ∫_0^{L_z} z·ρ_0 dz = ρ_0 L_z²/2 (also constant — the centroid sits at mid-height everywhere, so M_1 carries no spatial structure beyond a global offset); M_2(x,y) = ∫_0^{L_z} z²·ρ_0 dz = ρ_0 L_z³/3 (constant — the second Z-moment is uniform across the slab). Because ρ_d factorises as ρ^(0)(x,y)·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 second Z-moment. Reconstructing ρ_d(x,y,·) at (x_c,y_c) via maximum entropy from (M_0,M_1,M_2) recovers a maximum-entropy distribution constrained by M_0, M_1, and M_2 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 = {π_1,…,π_d} the spatial frame (d projection functionals π_i: $\mathcal{S}$ → $\mathbb{R}$). The access operator is: A(d): $\mathcal{S}$ → $\mathbb{R}$^d, A(d)(s) = (π_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):

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

$$P^{(3\to 2)} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \qquad A(2) = P^{(3\to 2)} \cdot A(3)$$

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

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

Terminal state and information preservation. The formal expression A(0) = 0 ∈ $\mathbb{R}^0$ 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}) ∈ $\mathbb{R}$^{N_lattice × d} (for CTR-C, §6.1), both of which are non-spatial structures sufficient for re-instantiation under the stated reconstruction assumptions — an operational sufficiency condition for the modelled class, not a claim of general finite-moment uniqueness (§3.2.2). 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 (empirical, for the tested field class — see §3.2.2 for the general non-uniqueness caveat); 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 interface applied to S, not a transformation of S, the substrate itself is unchanged: S′ = S, and hence H(S′) = H(S) = const for all d ≥ 0 (not a conditional entropy

H(S | A(d)), which for a deterministic function of S would generally satisfy H(S|A(d)) ≤ H(S), not equality). 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→k) ∈ ℝ^(k×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.

Proposition 1 (operator-family criterion). Let S denote the substrate's Information Measure (§3.1) immediately before a reduction step, and S′ the corresponding Information Measure immediately after. The reduction belongs to the index-type operator family A(d) if and only if S′ = S: the substrate's information measure is invariant under the reduction, and only its accessible frame is narrowed. The reduction belongs to the carrier-transformation operator family (P, U) if and only if S′ ≠ S: the information measure itself changes, because the content of the eliminated dimension is redistributed into, transformed within, or discarded from the substrate. Substrate-measure invariance is therefore the necessary and sufficient criterion for assigning any single operational step to one of these two operator families, independently of which of the six specific models that step instantiates; the further assignment among CTR-C, CTR-I, TSD-ND-C, etc. is then fixed by which modality (object-as-carrier, object-invariant, ambient-as-operator) the transformed or preserved S belongs to.

Remark. The criterion concerns the substrate's Information Measure, not its observable representation or coordinate frame. A change of coordinates, projection frame, or accessibility does not by itself imply S′ ≠ S. Conversely, preservation of the observable representation does not imply S′ = S if the underlying information measure has changed. Holding the choice of Information Measure fixed for a given substrate (as fixed by its Carrier Encoding in §3.1) is a precondition for applying this criterion; changing which functional is used to quantify S is a change of measurement convention, not a reclassification of the operator family.

This is hypothesis revision, not post-hoc classification. A(d) was initially hypothesised as a universal descriptor of reduction (§1). Protocol 4 (§7.4) tested this hypothesis directly and found that non-linear embeddings violate the predicted invariance H(S) = const, falsifying universal applicability. Proposition 1 above is the resulting, narrower criterion: A(d) applies exactly when S′ = S, a restriction derived from the falsification, not fitted to it after the fact. The same pattern governs PCA's reclassification in Protocol 5 (§7.5): a specific mapping was proposed, tested, found not to hold under truncation, and the classification was revised accordingly. In both cases the sequence is hypothesis → protocol → falsification → restricted criterion, the standard falsificationist pattern, not a framework that assigns whichever label happens to fit observed behaviour after the fact.

Scope under change of basis. The operator A(d) as defined above selects axis-aligned coordinates via P^(d→k) (§6.2). This leaves open the status of a linear projection performed in a different, non-axis-

aligned basis — e.g., projection onto the leading eigenvectors of a covariance matrix rather than deletion of a named coordinate (the case already exercised by PCA in Protocol 4, §7.4). Such a projection remains index deactivation in the IAS sense provided the basis is fixed independently of the reduction step and the substrate is represented basis-invariantly (e.g. as a pairwise-distance graph): the discriminating criterion H(S) = const does not depend on which orthonormal basis the retained directions are expressed in. What distinguishes A(d) from the redistribution operator P(d→m) is therefore not the choice of basis, but whether eliminated directions are discarded (A(d)) or folded into the remaining directions as higher moments (P(d→m)).

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 π_1, π_2, π_3: $\mathcal{S} \to \mathbb{R}$, e.g. a force-directed or spectral layout assigning each node v ∈ V a triple (π_1(v), π_2(v), π_3(v)) of embedding coordinates chosen so that graph-distance is approximately preserved. Deactivating index 3 (A(3) → A(2) = P^(3→2)·A(3)) discards π_3 and retains only (π_1,π_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 ∈ $\mathbb{R}^0$ 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→k) in §6.2 is a special case of the general index transformation A_T(d) = T · A(d), where T ∈ M(d,$\mathbb{R}$). Two structurally distinct subclasses exist: (i) rank-decreasing maps T = P^(d→k) — the reduction operators already defined; (ii) orientation-reversing isometries T = Ξ with det(Ξ) = −1. In the second subclass, rank is preserved (d → d) but the orientation of the index frame is inverted. Formally:

$$A_\Xi(d) = \Xi \cdot A(d) \;, \quad \Xi \in O(d)\backslash SO(d) \;, \quad \det(\Xi) = -1$$

$$\text{Example (inversion of Z-axis):} \quad \Xi = \mathrm{diag}(1, 1, -1)$$

Ξ is distinguished from the other operators as follows: P^(d→k) decreases rank (d → k < d, reduction); Ξ preserves rank but reverses orientation (det = −1, inversion); I preserves both rank and orientation (identity). Together they cover the index transformations used elsewhere in this taxonomy: P^(d→k) removes rank, while Ξ and I generate O(d) acting on the frame without removing coordinates. This is not the full matrix monoid M(d,ℝ): general invertible, non-orthogonal maps (e.g., anisotropic scaling, shears) are full-rank yet outside O(d), and are not classified by this operator set, since none of this paper's protocols require them.

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→2): g^(3)_μν  →  g^(2)_ab  =  g^(3)_ab ,  a,b ∈ {1,2}   (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)  →  g^(3) = [[g^(2), 0], [0, 1]]  +  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:

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

The degradation step D(2→1→0) is irreversible: the cycle structure of G is destroyed at the 2→1 linearisation. However, the topological code C_top = (Adj(G), Σ_cycles, χ) preserves the instruction for re-instantiation, which requires a minimum 2D metric frame. This code instantiates, for TSD-2D-C specifically, the general 0D encoding scheme set out in §6.2. The linear ordering σ used to construct the intermediate state L_σ is not itself specified by a formal criterion (e.g. minimum edge length, lexicographic order on vertex labels) — any choice suffices for the argument here, since C_top is defined directly from G (via Adj(G), Σ_cycles, and χ) rather than from L_σ, and so does not depend on σ by construction; different choices of σ therefore produce different intermediate linearisations L_σ but the identical topological code.

---

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. The protocols are constructive demonstrations of non-degeneracy for the specified observables and the tested model instantiations; they are not intended as general uniqueness theorems establishing that collisions between operator families are impossible for every possible implementation. Results are summarised in §7.7.

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

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

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

$H[\rho_XY^{(0)}]$ = 2.000 ± 0.000 bits; H_avg = 2.000 ± 0.000 bits/moment (n=10)
MSE, m=0 (M_0 only): $(1.736 \pm 0.000) \times 10^{-6}$
MSE, m=1 (M_0+M_1): $(0.404 \pm 0.000) \times 10^{-6}$
MSE, m=2 (M_0+M_2): $(0.011 \pm 0.000) \times 10^{-6}$

C_channel	m_max	Moments used	MSE ($\times 10^{-6}$)
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 (the second raw moment) 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. Θ and m_max are independent parameters (Appendix A): Θ determines whether the dynamical cascade activates; m_max determines how many moments can be retained after activation. No numerical mapping between them is assumed. 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: the residual carrier cannot represent even the zeroth-moment baseline, so reconstruction is unavailable at any retained moment order. This confirms that m_max functions as a hard bit-budget cutoff, not merely a notational device.

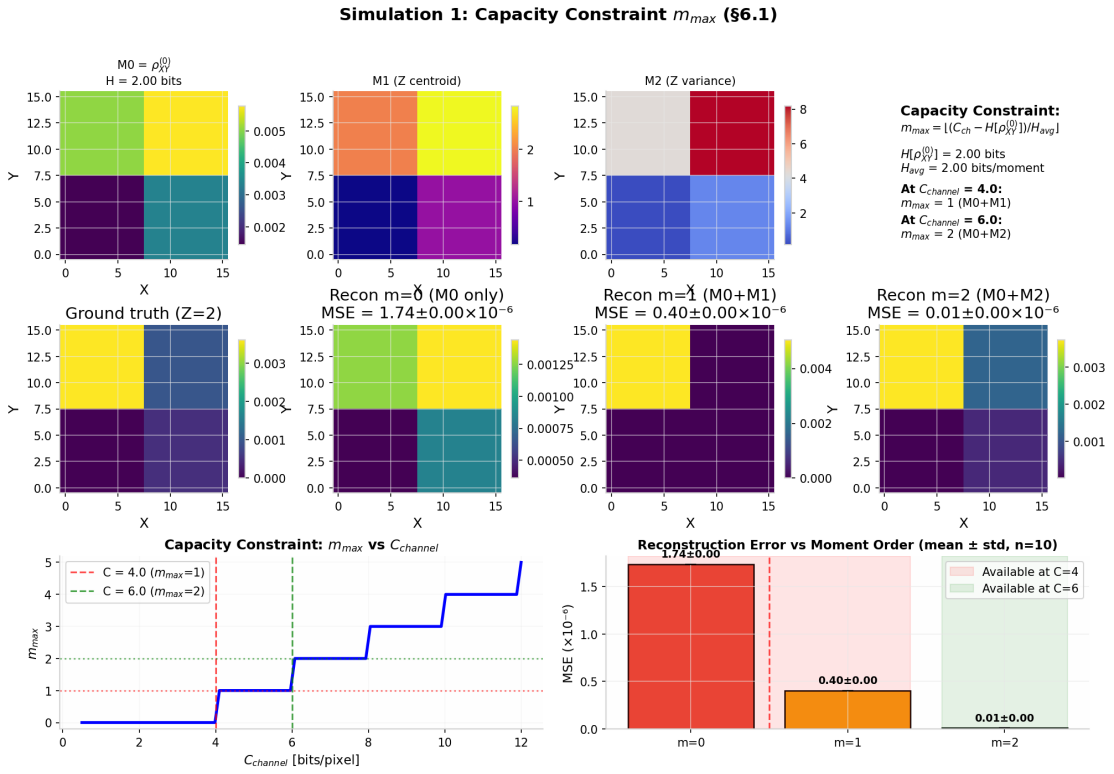


Fig. P1. Capacity constraint m_max (§7.1). Top row: moment fields M_0 (projected density), M_1 (Z centroid), M_2 (second Z-moment). Middle row: ground truth and reconstructions at m=0,1,2 with corresponding MSE. Bottom left: m_max staircase as a function of C_channel. Bottom right: reconstruction error 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: $(8.146 \pm 1.056) \times 10^{-7}$ at m=0, $(6.847 \pm 0.898) \times 10^{-7}$ 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 variance amplification of the residual field across the cascade, and to test the conditional entropy inequality at the 3D→2D redistribution step.

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	$(2.97 \pm 0.09) \times 10^{6}$	×358

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

Interpretation. Residual variance grows monotonically (×385, ×358) across the tested cascade, confirming the CTR-C prediction. The entropy inequality $H[M^{(2)}] \geq H[\rho_XY^{(0)}]$ is satisfied for the 3D→2D redistribution step. The subsequent 2D→1D entropy drop (5.284 → 4.518) reflects the transition to scalar code, not a violation of the tested inequality — consistent with the theoretical expectation that final scalarization compresses the information load while preserving the variance amplification trend. Monotonic entropy increase is therefore not claimed for the complete 3D→2D→1D trajectory, only the tested 3D→2D inequality and the variance amplification across all steps.

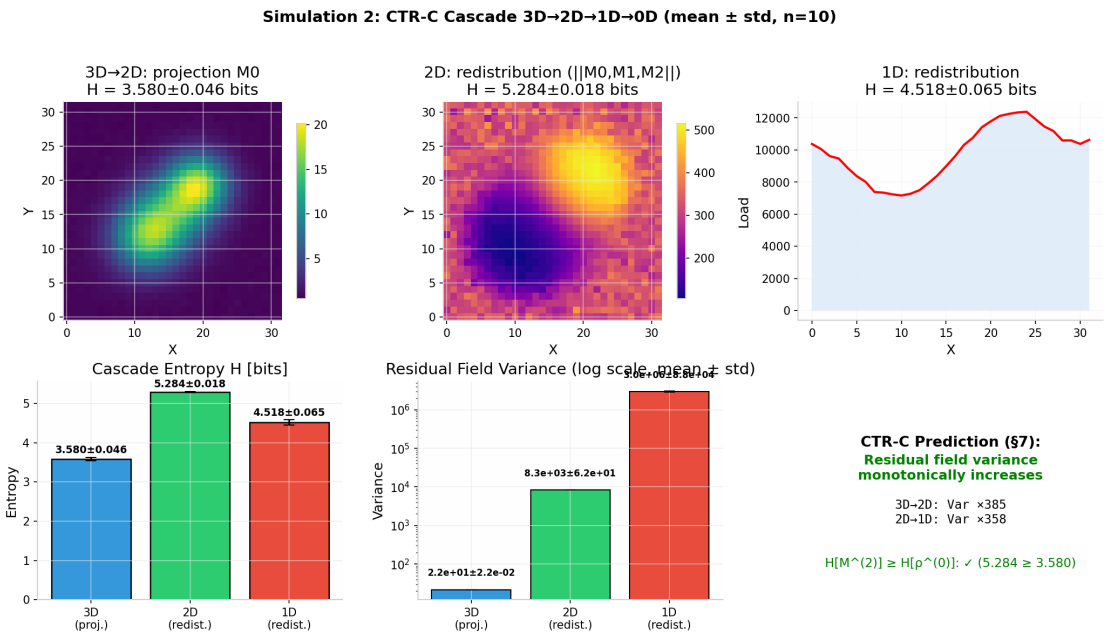


Fig. P2. Cascade entropy monotonicity (§7.2). Top row: 3D→2D projection M_0, 2D redistribution (||M_0,M_1,M_2||), 1D redistribution load profile. Bottom left: cascade entropy H at each step (mean ± std, n=10). Bottom middle: residual field variance on log scale (mean ± std, n=10). Bottom right: CTR-C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

A note on parameter robustness. Apparent agreement between an information-theoretic observable and a structural invariant can be regime-dependent: a parameter- or scaling-level coincidence does not by itself establish a structural correspondence. Accordingly, an entropy-based discriminator such as the one tested above should not be regarded as ontologically diagnostic solely on the basis of agreement observed for a single parameter regime or field configuration.

Parameter-robustness requirement. An entropy-based discriminator must not be regarded as ontologically diagnostic solely on the basis of agreement observed for a single parameter regime or field configuration. Its discriminating behaviour must remain stable under controlled variation of perturbation amplitude, spatial regularity, correlation structure, and cascade parameters. Agreement at a single parameter value may represent an exponent or scaling coincidence rather than a structural consequence of the underlying operator.

Boundary-of-validity summary. Table 2 collects the status of each entropy-monotonicity-related claim established above, following the parameter-robustness requirement just stated.

IAS claim	Status
CTR-C predicts entropy increase	Conditional empirical result
Holds for Gaussian-peak fields (Protocol 2)	Demonstrated
Holds for arbitrary ρ_d	False (§7.2.1: fails for non-separable fields)
Independent of correlation structure of ρ_d	False (§7.2.1: non-separable field shows ΔH = −0.20 bits)
Agreement at a single parameter regime is sufficient grounds for a diagnostic claim	Not established (see parameter-robustness requirement, above)
Entropy monotonicity is operator-invariant	False (§7.2.1: depends on correlation structure of ρ_d)

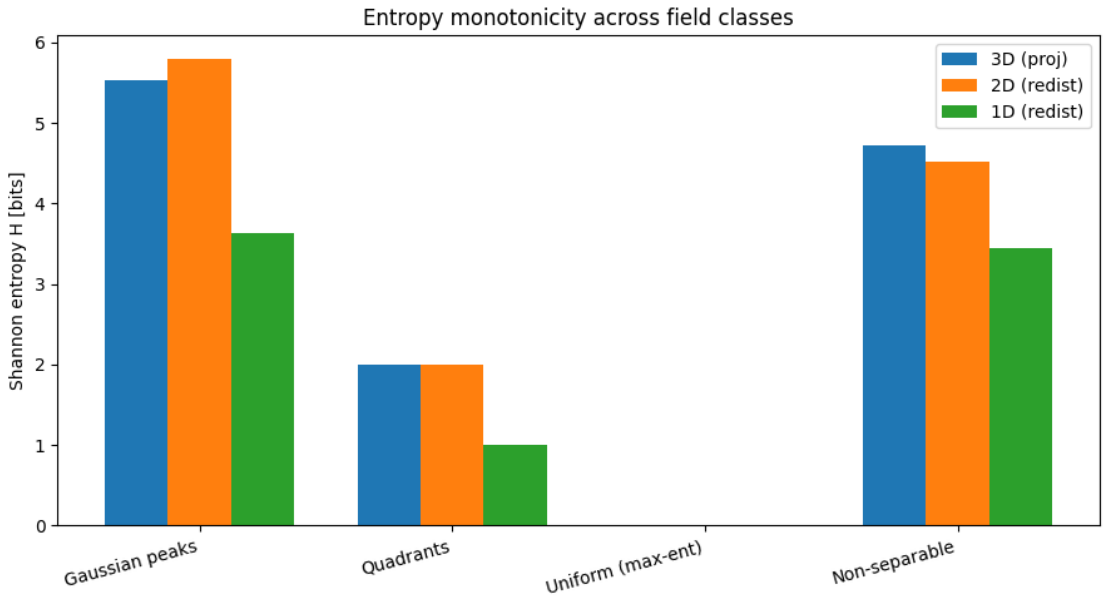


Fig. P2b. Boundary of the entropy monotonicity claim (§7.2.1). Shannon entropy H of the residual carrier at each cascade step for four field classes (16^3 lattice, m = 2). Gaussian peaks: monotonic increase confirmed. Quadrants and uniform (max-ent): saturation/ceiling effect (ΔH = 0). Non-separable (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 to displayed precision — see §6.3 for the distinction between this rounding and an exactly-zero residual), 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.

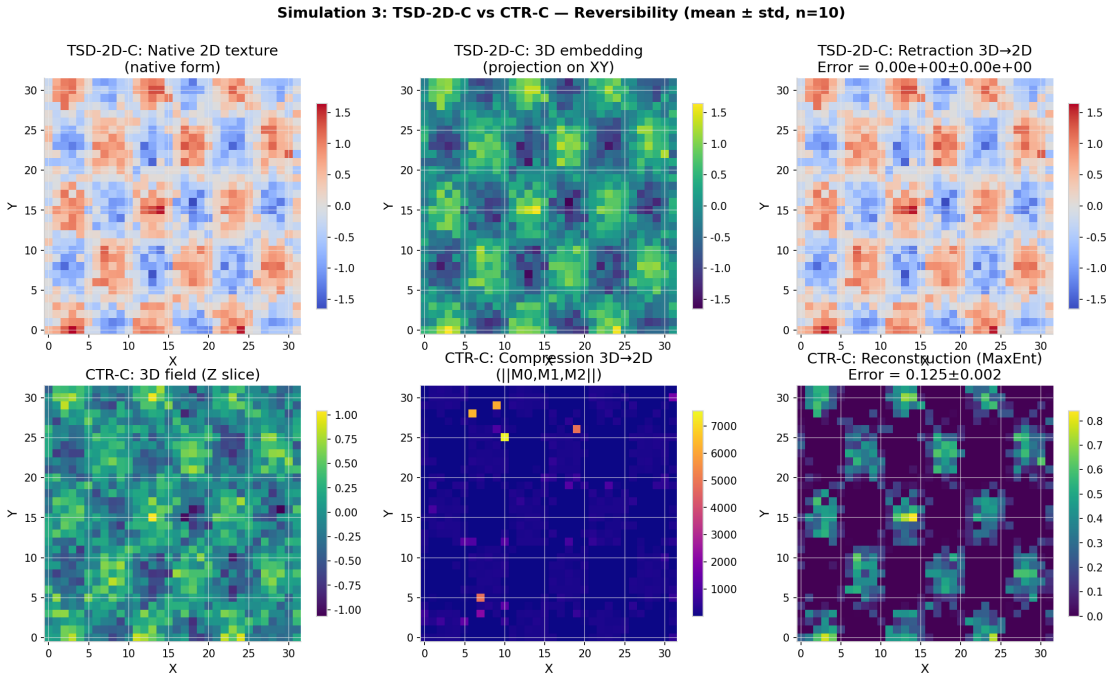


Fig. P3. Reversibility discrimination (§7.3). Top row (TSD-2D-C): native 2D texture, 3D embedding (projection on XY), retraction 3D→2D with error = 0.00 ± 0.00. Bottom row (CTR-C): 3D field (Z slice), compression 3D→2D (||M_0,M_1,M_2||), MaxEnt reconstruction with error = 0.125 ± 0.002. H[ρ_2D] richness 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). The fundamental criterion is S′ = S (Proposition 1, §6.2); the Global

Spearman ρ statistic used below is an operational proxy for this criterion, not a substitute for it, and the working threshold on ρ is an operational criterion specific to this dataset and protocol, not a universal ontological constant.

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. On the synthetic Swiss Roll benchmark used in this protocol, 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, the separation is confirmed by exact non-parametric tests (two-sided Mann-Whitney U, exact and Monte Carlo permutation tests, and percentile-bootstrap confidence intervals) on all pairwise and pooled comparisons; every test rejects equality at $p < 10^{-4}$ or lower. Full statistics are given in Appendix C.1.

Interpretation. PCA preserves global substrate structure (ρ = 0.837 ± 0.028) → consistent with TSD-ND-I: linear projection approximates index deactivation without modifying the substrate. Both t-SNE (ρ = 0.403 ± 0.047) and UMAP (ρ = 0.304 ± 0.021) violate H(S) = const; UMAP distorts global distances even more aggressively than t-SNE (Δρ ≈ 0.10), confirming that non-linear embeddings transform the substrate rather than deactivate the projection interface. This establishes a dataset- and protocol-specific empirical separation rule at the target dimension tested here (2D): only linear projections satisfy the TSD-ND-I ontology under this observable; any non-linear distance distortion presupposes CTR-C or TSD-ND-C. The rule is operational, not ontological: it classifies observable behaviour under the tested proxy, not metaphysical substrate status. The fundamental criterion remains S′ = S (Proposition 1); ρ is an operational proxy for it, not a substitute for it — and, as the MNIST dimension sweep below shows, the exactness of this proxy's separation is itself dimension-dependent. 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 $P^{(3\to2)} \in \mathbb{R}^{(2\times3)}$, reducing the rank of the index frame without transforming the substrate, whereas both t-SNE and UMAP optimise local neighbourhood structure at the expense of global metric fidelity, modifying the pairwise distance graph itself. The fact that UMAP scores lower than t-SNE falsifies the hypothesis in §1.1.1 that UMAP better

preserves global topology; for the specific requirement of substrate invariance under index deactivation, UMAP is a poorer approximation of A(d) than t-SNE, and both are qualitatively distinct from PCA.

This establishes a hard, data-derived discriminatory rule: any dimensionality-reduction algorithm producing a global Spearman ρ that falls above the exact separation interval ($\rho > 0.773$) 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 empirical separation 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.

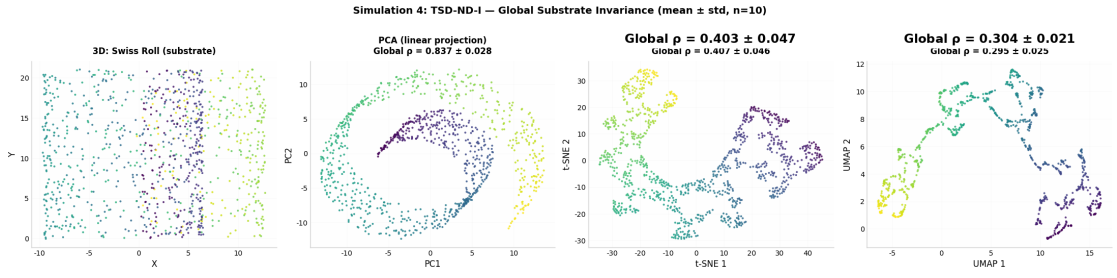


Fig. P4. Substrate invariance (§7.4). Left: 3D Swiss Roll substrate. Second: PCA linear projection (Global ρ = 0.837 ± 0.028) preserves global structure → consistent with TSD-ND-I. Third: t-SNE non-linear embedding (Global ρ = 0.403 ± 0.047) distorts global distances. Right: UMAP non-linear embedding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

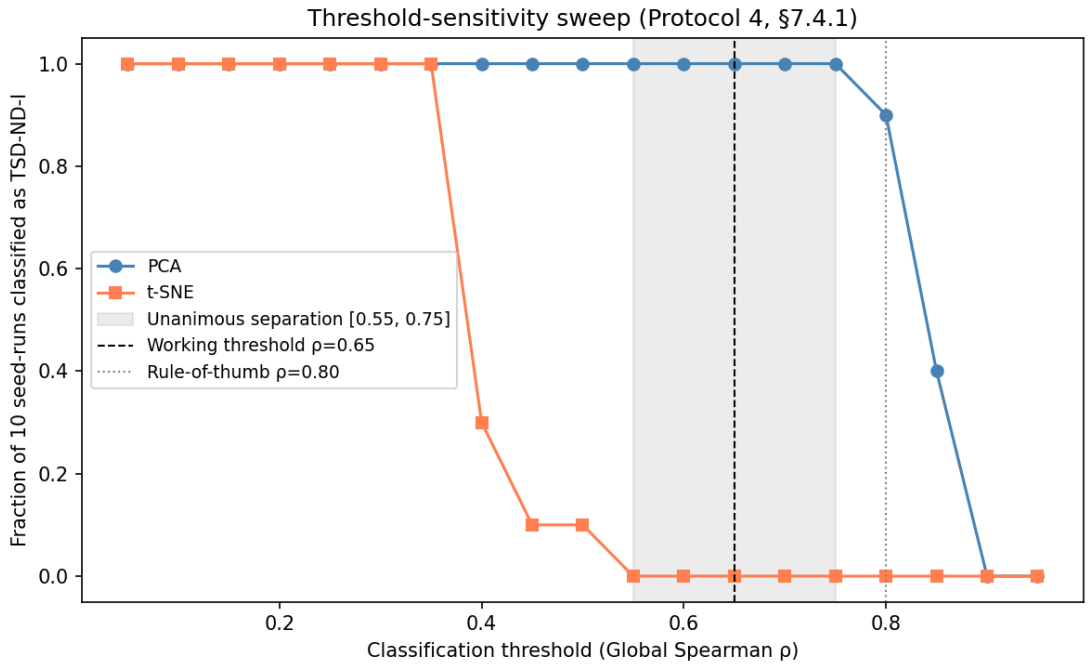


Fig. P4b. Threshold-sensitivity sweep (§7.4.1). Fraction of the 10 PCA and 10 t-SNE seed-runs (Swiss Roll, n=1000, subsample n=300 for the ρ calculation) classified as exceeding a given Global-Spearman-ρ cutoff, swept from 0.05 to 0.95. PCA (blue) stays at 100% through ρ=0.75, declining only above ρ=0.80; 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 ρ = 0.65 used throughout this paper remains below this narrower interval and continues to classify all runs without exception.

Statistical significance. The PCA-vs-Kernel-PCA and PCA-vs-pooled-non-linear separations are confirmed by exact Mann-Whitney U, exact and Monte Carlo permutation tests, and percentile-bootstrap confidence intervals, all rejecting equality at $p < 10^{-4}$ or lower. Full statistics are given in Appendix C.2.

Gamma-sensitivity of Kernel PCA. Varying the RBF bandwidth γ reveals a continuous spectrum between the two ontological classes (Fig. P4c): at γ=0.001, Kernel PCA converges to PCA (ρ = 0.828 ± 0.025), recovering near-complete substrate invariance; at γ=0.003, ρ = 0.797 ± 0.023, still above the exact-separation interval; at γ=0.01 (the value used above), ρ = 0.657 ± 0.026, inside the narrow ambiguous band; at γ=0.03, ρ = 0.297 ± 0.061, converging toward the t-SNE/UMAP cluster; at γ=0.1, ρ = −0.162 ± 0.027, indicating aggressive metric distortion with no residual substrate correlation. This result suggests that the operational distinction between index preservation and substrate transformation can form a continuous spectrum in concrete implementations, even though IAS treats the corresponding ontological commitments (TSD-ND-I and CTR-C) as distinct ideal types: only the γ→0 limit (strict linearity) recovers full substrate invariance. In the tested implementation, sufficiently small γ produced behaviour close to the linear-PCA baseline; the exact γ→0 limit is not identified here with PCA without additional kernel-normalisation assumptions.

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

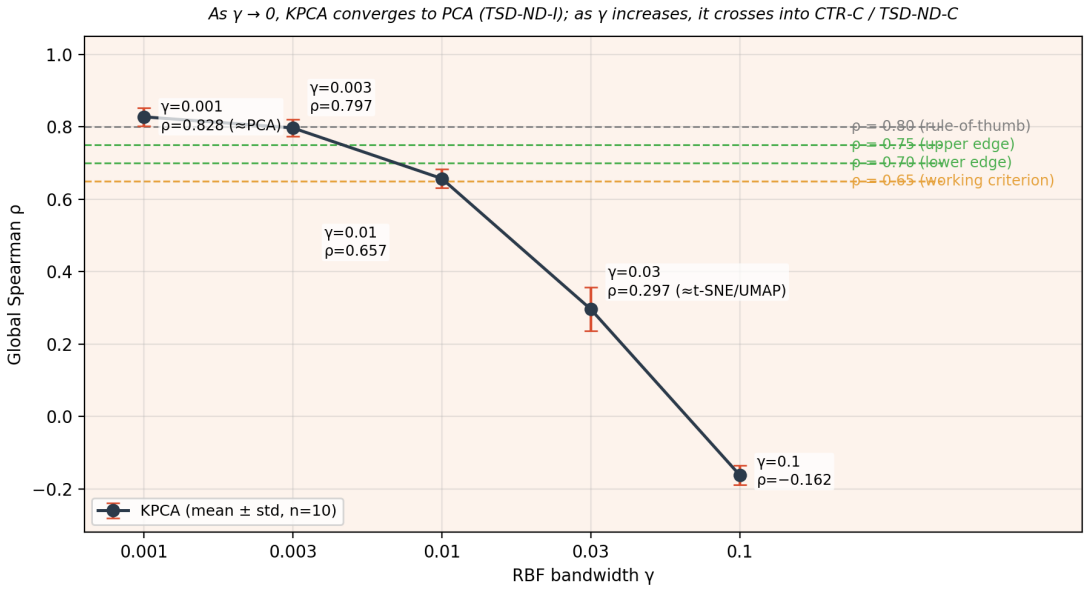


Fig. P4c. Kernel PCA gamma-sensitivity (§7.4.2). Global Spearman ρ as a function of RBF bandwidth γ (mean ± std, n=10 seeds per γ). Dashed lines mark the §7.4 rule-of-thumb (ρ=0.80), the §7.4.2 stable-interval edges (ρ=0.70–0.75), and the working criterion (ρ=0.65) used throughout this paper.

7.4.3 Anisotropic and Correlated Substrate Robustness

Objective. Protocol 4 (§7.4) and its Kernel PCA extension (§7.4.2) establish the PCA/nonlinear separation on an isotropic synthetic substrate. This raises a fair question: is the separation a property of the operator families being compared, or an artefact of the isotropic toy data? This test asks whether the separation survives when the substrate has anisotropic or correlated variance structure.

Setup. Points drawn from a multivariate Gaussian N(0,Σ) in $\mathbb{R}^3$ under four covariance structures: (C1) isotropic, σ=(1,1,1); (C2a) moderate anisotropy, σ=(1,3,0.3), 10× variance ratio; (C2b) extreme anisotropy, σ=(1,10,0.1), 100× variance ratio; (C3) same variances as C2b but rotated by a random SO(3) rotation so the axes of stretch do not align with the coordinate axes. n=500 points per run, n=100 independent seeds per condition. PCA and t-SNE reduce each substrate to 2D; Global Spearman ρ is computed between original and embedded pairwise distances, exactly as in §7.4.

Condition	rho_PCA	rho_tSNE	tSNE>=0.65	Gap
C1 Isotropic (1,1,1)	0.82 ± 0.02	0.64 ± 0.03	39%	+0.17
C2a Moderate (10x)	1.00 ± 0.00	0.94 ± 0.01	100%	+0.06
C2b Extreme (100x)	1.00 ± 0.00	0.96 ± 0.01	100%	+0.04
C3 Correlated/tilted	1.00 ± 0.00	0.96 ± 0.01	100%	+0.04

(mean ± std, n=100 seeds per condition)

Interpretation. The separation demonstrated in §7.4 holds cleanly only in the isotropic case: t-SNE clears the ρ≥0.65 criterion in 39% of seeds under C1 (exact 95% CI: 30–49%) — a genuinely unstable boundary. Under anisotropy — even the moderate 10× condition — t-SNE clears the same threshold in 100% of seeds (exact 95% CI: at least 96%), and the PCA/t-SNE gap collapses from +0.17 to +0.04–+0.06. C2b and C3 are numerically identical, indicating the collapse depends on the variance ratio alone, not on whether

the stretched axes align with the coordinate frame. A follow-up diagnostic (§7.4.4) isolates the mechanism.

7.4.4 Diagnostic: Mechanism of the Anisotropic Collapse

Objective. §7.4.3 leaves open why the PCA/t-SNE separation collapses under anisotropy. This diagnostic tests one specific hypothesis: that the collapse occurs because the substrate becomes effectively one-dimensional once one axis dominates the variance, at which point any method — not just PCA or t-SNE — preserves pairwise distances well.

Setup. A trivial baseline is constructed by projecting onto the single highest-variance coordinate axis alone — no PCA, no orthogonalization, no method-specific optimisation. Global Spearman ρ is computed for this baseline, for PCA, and for t-SNE, across a finer sweep of variance ratios ($\sigma_y/\sigma_z \in \{1, 2, 3, 5, 10, 20, 50, 100\}$, σ_x fixed at 1). n=500 points, n=100 seeds per ratio.

Ratio	rho_dominant-axis	rho_PCA	rho_tSNE
1x	0.53	0.82	0.64
2x	0.88	0.99	0.90
3x	0.95	1.00	0.94
5x	0.99	1.00	0.95
10x	1.00	1.00	0.96
20x	1.00	1.00	0.95
50x	1.00	1.00	0.95
100x	1.00	1.00	0.94

(mean, n=100 seeds per ratio)

Interpretation. The trivial single-axis baseline tracks the rise in ρ_tSNE closely: both climb sharply between ratio 1× and 5×, and both saturate by ratio 10×–20×. By ratio 10×, projecting onto one coordinate axis alone already reaches ρ=1.00 — the same value PCA reaches with a full two-dimensional embedding. This supports the hypothesis stated in the Objective: t-SNE's high ρ under anisotropy does not reflect genuine nonlinear structure-preservation; it reflects that preserving pairwise distances becomes a nearly trivial task once the substrate's variance concentrates on a single axis. The ρ≥0.65 criterion (§7.4) is therefore diagnostic in the isotropic-to-mildly-anisotropic regime; it should not be read as evidence of PCA/nonlinear separability under substrates with strong axis-dominance, where the criterion is satisfied by construction rather than by genuine structure-preservation. This bears on the empirical demonstration in §7.4 only — not on the underlying invariant S'=S (Proposition 1), for which ρ is already stated to be an operational proxy, not a substitute (§7.4).

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 ± 0.01) × 10^{-2}
40	2.019 ± 0.005	(2.65 ± 0.01) × 10^{-2}
30	2.046 ± 0.005	(2.39 ± 0.01) × 10^{-2}
20	2.066 ± 0.005	(2.04 ± 0.01) × 10^{-2}
10	2.028 ± 0.003	(1.51 ± 0.01) × 10^{-2}
5	1.941 ± 0.005	(1.03 ± 0.01) × 10^{-2}
2	1.664 ± 0.011	(5.31 ± 0.12) × 10^{-3}
1	1.375 ± 0.008	(3.33 ± 0.13) × 10^{-3}

Operational null condition. Let H_r and V_r denote reconstruction entropy and variance at retained dimension r. The CTR-C signature requires $H_{r_{i+1}} \geq H_{r_i}$ and $V_{r_{i+1}} \geq V_{r_i}$ for every successive step of the elimination cascade. The operational null baseline is the absence of this joint monotone increase; the CTR-C signature fails to hold when either observable exhibits a non-monotone trajectory under the predefined cascade. This is an operational null condition, not a population-level statistical null hypothesis: Protocol 5 has no sampling distribution to test against, only a determinate prediction about the trajectory of a computed quantity across a fixed cascade.

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 provides operational evidence consistent with 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 $\rho = 0.837 \pm 0.028$, TSD-ND-I). Protocol 5 extends that result from synthetic manifolds to a standard ML benchmark, demonstrating that the same operational discrimination extends from synthetic manifolds to real data. The absence of the CTR-C signature in PCA truncation further illustrates the taxonomy's discriminating power at the operational level: not every dimensionality-reduction algorithm exhibits a compression signature; the IAS framework provides the formal apparatus to determine which operator family's operational signature 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.

Scope of this result. The CTR-C-falsification claim above rests on a single dataset (MNIST); the cross-dataset confirmation available for PCA's TSD-ND-I classification (Protocol 4: Swiss Roll; Protocol 5: MNIST) applies to the projection step, not specifically to this truncation test. Nor is the entropy/variance observable uniquely diagnostic: a different observable could in principle be compatible with more than one model class for the same operation. The claim should be read as this operation, under this observable, on this dataset, failing the CTR-C signature — not as a general proof that PCA truncation cannot instantiate CTR-C under any measurement.

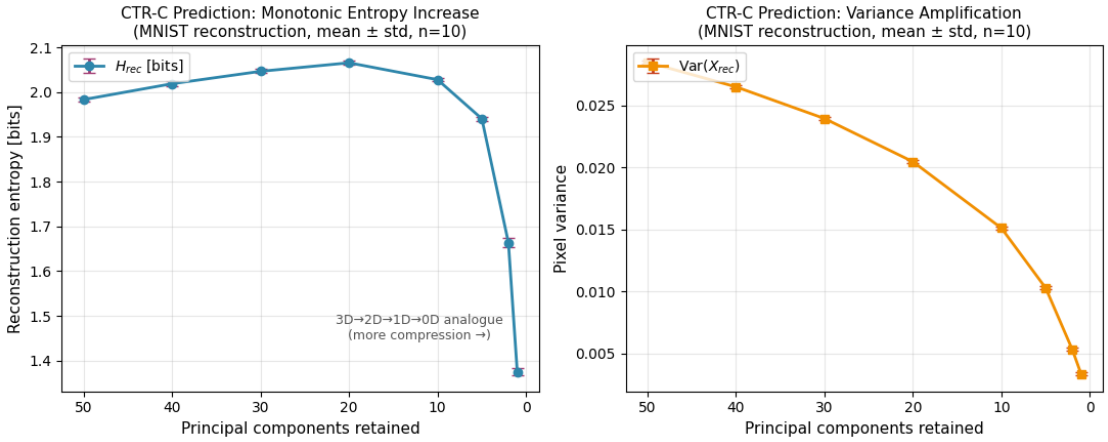


Fig. P5. Ontological discrimination via PCA truncation on MNIST. Left: reconstruction entropy H_rec [bits] versus retained principal components (mean ± std, n=10); entropy peaks at ~20 PCs and declines, violating the CTR-C monotonic-increase prediction. Right: pixel variance Var_rec decreases monotonically 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 Protocol 6: Operational Demonstration — VAE Ontological Fingerprint

Objective. To demonstrate that the ontological fingerprint vector f proposed in §1.1.3 is computable on an existing algorithm, not merely a conceptual extension of the taxonomy, by instantiating it on a β-VAE and testing the hybrid-commitment expectation raised there for that architecture.

Setup. A β-VAE (MLP encoder 784→256→2×latent_dim, MLP decoder latent_dim→256→784 with sigmoid output; Adam optimiser, learning rate 1e-3, batch size 128, 20 epochs; MNIST subset n=5,000, 10 seeds per configuration) is probed with three measurements, each reusing an estimator already established elsewhere in this paper: (1) H[μ] — the 50-bin density-histogram entropy of latent activations, same estimator family as Protocol 2 (§7.2) — measured across latent_dim ∈ {50,20,10,5,4,3,2} at β=1, testing for the CTR-C redistribution signature (entropy rising as the bottleneck narrows); (2) the count of "active" latent units (per-dimension KL(q(z_i)‖p(z_i)) exceeding 0.01 nats, a standard VAE diagnostic) measured across β ∈ {0.1,0.5,1,4,10,25} at latent_dim=10, testing for a TSD-ND-C capacity-limit signature analogous to the m_max staircase of §6.1/§7.1; (3) Global Spearman ρ between input and latent-μ pairwise distances, same n=300 subsample estimator as Protocol 4 (§7.4), at

the baseline configuration (β=1, latent_dim=10), testing for the TSD-ND-I substrate-invariance signature against the working criterion ρ > 0.65 established in §7.4.1.

Statistical treatment. Rather than counting every consecutive step by the direction of its point estimate, each step's bootstrap 95% confidence interval (2,000 resamples) is checked for overlap with its neighbour; only steps whose intervals do not overlap are counted as confidently distinguishable from noise. The reported fingerprint weight w_j is the fraction of confidently-distinguishable steps consistent with signature j, excluding plateau steps (CI-overlap) from both numerator and denominator rather than assigning them a direction by default. Fingerprint weights are conditional on the specified operational probes (H[μ], active-unit count, Global Spearman ρ) and should therefore be read as probe-dependent diagnostic weights, not as intrinsic, probe-independent properties of the algorithm.

Results (mean ± std, n=10, reproduced across three independent runs including CPU and GPU backends). H[μ] rises confidently across three consecutive steps (latent_dim 50→20→10→5: 4.13→4.58→4.86→5.05 bits, all CI95 non-overlapping), plateaus at latent_dim 5→4→3 (5.05→5.07→5.09 bits, CI-overlapping — not distinguishable from noise), then drops confidently at the final step (latent_dim 3→2: 5.09→4.98 bits, CI95 non-overlapping). Active units remain saturated at 10.00±0.00 across β∈{0.1,0.5,1,4,10}, an uninterrupted CI-overlapping plateau, then collapse confidently to 1.00±0.7 at β=25 — the only confidently-distinguishable step in this measurement. Global Spearman ρ = 0.754 ± 0.017 (range 0.722–0.791 across runs), with its full 95% CI comfortably above the working criterion of 0.65.

The components below are fractions of confidently-distinguishable operational steps, not Bayesian weights or posterior probabilities of ontological membership: they quantify behavioural signatures, not metaphysical truth-values.

```
f_VAE = (w_CTR-C=0.75, w_TSD-ND-C=1.00, w_TSD-ND-I=1.00, w_CTR-I=—, w_TSD-2D-C=—,
                                                    w_TSD-2D-I=—)
```

Interpretation. Each w_j is the fraction of confidently-tested operational criteria satisfied for that signature, not a probability of ontological membership; "—" marks a component this protocol did not test (not applicable to this substrate type), distinct from an observed weight of 0. The result confirms the hybrid-commitment expectation raised in §1.1.3 for Variational Autoencoders with a concrete, reproducible number rather than an unverified claim: this VAE configuration exhibits CTR-C-consistent redistribution over most of its tested range with a genuine breakdown only at the smallest bottleneck, a sharp TSD-ND-C-consistent capacity collapse concentrated at high β rather than a gradual staircase, and TSD-ND-I-consistent preservation of global relational structure throughout. This experiment is not presented as an ML benchmark or as evidence that IAS constitutes an ML methodology; it demonstrates that the fingerprint proposed in §1.1.3 can be operationally instantiated, with the same statistical care as Protocols 1–5, on an existing computational reduction system.

7.6a Experimental Reproducibility Summary

Table 1a collects, in one place, the dataset, sample size, seed count, observable, and statistical test used by each protocol above; full parameter values remain in each protocol's own Setup paragraph.

Protocol	Dataset / n	Seeds	Observable(s)	Statistical test
P1 (§7.1)	16×16 quadrant signal; MNIST real-data extension (n_z=8	n=10	Reconstruction MSE	None (mean ± std only)

	slices)			
P2 (§7.2)	Synthetic Gaussian peaks, 32^3 lattice	n=10	Shannon entropy H, variance	None (mean ± std only)
P3 (§7.3)	Synthetic 2D texture / 3D field	n=10; seed not recorded	Reconstruction error, H[ρ_2D], H[ρ_3D slice]	None (mean ± std only)
P4 / P4.2 (§7.4)	Swiss Roll (n=1,000); MNIST dimension sweep (d=2,3,5,10,20)	n=10 (seeds 0–9)	Global Spearman ρ	Exact Mann-Whitney U, exact + Monte Carlo permutation, bootstrap CI (Appendix C.1–C.2); threshold and γ sensitivity sweeps (§7.4.1, §7.4.2)
P5 (§7.5)	MNIST (n=5,000/run)	n=10 (seeds 0–9)	H_rec, Var_rec	None (mean ± std only)
P6 (§7.6)	MNIST (n=5,000, β-VAE)	n=10 per config	H[μ], active units, Global Spearman ρ	Bootstrap 95% CI + step-by-step CI-overlap significance test

7.7 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 (below reporting precision) 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 classified as TSD-ND-I-like under the tested entropy/variance observables, not CTR-C-like. (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 exhibits an intermediate operational	TSD-ND-I / CTR-C	§7.4.2	Confirmed: Kernel PCA (γ=0.01) ρ=0.657±0.026,

signature between TSD-ND-I and CTR-C			strictly between PCA and t-SNE/UMAP; gamma sweep shows a continuous transition, with γ→0 recovering TSD-ND-I

7.8 Out-of-Taxonomy / Negative Classification Control

Objective. The fingerprint procedure (§1.1.3, §7.6) assigns each tested algorithm a weight in [0,1] for each of the six ontological commitments. This raises a structural question distinct from any single algorithm's classification: does the procedure ever return a negative result — zero weight across the board — or is every input necessarily assigned to at least one of the six classes by construction? A taxonomy that cannot say “none of the above” risks circularity: the classification would then be guaranteed by the test's own design rather than discovered by it.

Setup. Reference substrate: three Gaussian peaks on a 32^3 lattice (§7.2 generator), n=1000 independent seeds. Two positive controls with known ground truth (PCA, a genuine TSD-ND-I operator; moment-order compression P(3→2), a genuine CTR-C operator) and three negative-control candidates with no principled relation to any IAS operator: i.i.d. Gaussian noise substitution, random permutation of the flattened field, and a random (non-PCA) Gaussian linear projection. Only w_TSD-ND-I (Global Spearman ρ≥0.65, swept across a compression-strength schedule) is tested here; two earlier attempts at a generic w_CTR-C / w_TSD-ND-C proxy did not discriminate cleanly on these same controls (PCA scored as high as or higher than the true CTR-C operator on the CTR-C proxy) and are not reported, rather than reporting numbers that would not mean what they claim to.

```
Candidate                              w_TSD-ND-I (mean ± std, n=1000)
P1 PCA (positive control)              1.00 ± 0.00
P2 Moment compression (pos. control)   0.40 ± 0.46
N1 Random noise (negative control)     0.00 ± 0.00
N2 Random permutation (neg. control)   0.00 ± 0.00
N3 Random projection (neg. control)    1.00 ± 0.00
```

Interpretation. Both structurally unrelated negative controls — noise and permutation — return w_TSD-ND-I = 0.00 with zero variance across all 1000 seeds: the criterion is never satisfied by construction for either, confirming a genuine exclusion region rather than a classification forced by the test's design. PCA returns 1.00±0.00, confirming the estimator remains sensitive to a genuine positive case at this sample size. Random projection also returns 1.00±0.00 — not a test failure but a substantive finding consistent with the Johnson–Lindenstrauss lemma: sufficiently high-dimensional random linear projections approximately preserve pairwise distances regardless of any relation to the substrate's principal directions. Moment compression, included as the CTR-C positive control, returns a high-variance 0.40±0.46 on this unrelated axis — unsurprising, since it was never expected to satisfy an index-preservation criterion. This confirms an exclusion region for the one criterion tested here, not for the full six-component fingerprint: no non-tautological generic proxy was found for the remaining five components under this experimental design (§7.4.3–§7.4.4 illustrate a related difficulty for a different axis), and they are not claimed to share this property.

8. Summary Tables

8.1 Six Trajectories to 0D

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

8.2 Extended Comparison

Code	Hypothesis	Substrate	Role of Space	Reduction Trajectory	0D: Final State
CTR-C	CTR (Repository)	3D array; information in channels X,Y,Z. Object = substrate.	Passive container. Z — volume axis of object.	3→2: I_z → X,Y (integration). 2→1: I_y + I_z → X. 1→0: folding into code.	Informational archive. Complete set (X,Y,Z) without metric.
CTR-I	CTR (Index)	Non-dimensional substrate; coordinates = access indices.	Indexation system. Substrate external to metric.	3→2: deactivation Z. 2→1: deactivation Y. 1→0: deactivation X.	Substrate external to indexation. Complete volume, not deployed.
TSD-ND-C	TSD-A (Carrier)	Non-dimensional cluster; localized on axes as channels.	Complete dimensionality assigner. Without space — no carrier.	3→2: Z → X,Y. 2→1: Y → X. 1→0: elimination X.	Code without coordinate shell. Potential for re-addressing.
TSD-ND-I	TSD-A (Projection)	Pure non-dimensional graph; no spatial attributes.	Manifestation system. Space — interface.	3→2: eliminate interface Z. 2→1: eliminate interface Y. 1→0: eliminate interface X.	Structure external to space. Pure potential.
TSD-2D-C	TSD-B (Superstructure)	2D surface + metric superstructure Z (added by ambient).	Dimensionality operator. Z not intrinsic to 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. The convergence is therefore formal rather than ontologically uniform: the six models agree on the disappearance of spatial realization while disagreeing on what, if anything, remains spatially unmanifest.

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 0D terminal state is therefore defined ontologically as a class of non-spatial codifications sufficient for re-instantiation under the model-specific reconstruction assumptions stated above; IAS does not prescribe a single unified encoding scheme or state space common to all six models, since the content of the code is fixed by which element of the substrate–frame–index triad the model in question preserves.

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 §7.7, demonstrates that all four discriminating predictions derived from the formalism are supported within their declared field classes, substrate representations, and embedding assumptions. The capacity constraint m_max governs reconstruction fidelity: exact recovery requires preservation of moments up to order m = 2 for the tested field class (§3.2.2 discusses the general non-uniqueness of this reconstruction), 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,

indicating a strong separation rather than a marginal statistical effect, for this dataset and this parameter regime (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.

This classificatory core (§1–§9) is the paper's central contribution and stands on its own. Appendix A works one model's mechanism (CTR-C) into an explicit lattice simulation; Appendix B goes further, asking a question the core taxonomy deliberately leaves open — which of the six models can be physically instantiated in a space of a given dimensionality — and proposes a separately-flagged, speculative answer for d = 3. Appendix B.1–B.6 states that extension's epistemic status explicitly.

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). The taxonomy is not physically falsifiable in the empirical sense disclaimed in §1.2, but its claimed structural completeness is formally challengeable by a counterexample: this completeness claim would be undermined 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.

Anisotropic and correlated substrates. The discriminating predictions of §6–§7 are derived for substrates in which the eliminated dimension carries information that is locally separable or conditionally independent of the residual carrier (see the sufficient condition in §6.1's entropy-monotonicity lemma). Strongly correlated, anisotropic, or entangled high-dimensional substrates — e.g. moment fields with $n \gg 3$, or dimensions coupled through non-trivial mutual-information chains — may violate the monotonicity or invariance signatures tested here, as already illustrated by the non-separable field counterexample in §7.2.1. The taxonomy does not claim to have exhausted the phenomenology of such substrates; extending the operator families to correlated and anisotropic cases is left for future work.

10.1 Towards a Practitioner-Facing Decision Procedure

This section is a secondary, incidental application, not a redefinition of the paper's audience: §1.2 states plainly that the primary audience is the theorist clarifying ontological commitments, not the practitioner seeking a pipeline improvement. What follows shows that the same classificatory apparatus happens to also support a practical diagnostic, without that possibility being the work's purpose.

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.

Author Contribution Statement

The conceptual framework, theoretical foundation, and core innovations of this work — including the Informational Architectonics of Space taxonomy, the six-model ontological classification (CTR-C, CTR-I, TSD-ND-C, TSD-ND-I, TSD-2D-C, TSD-2D-I), the operator typology (P, A, U/D), and the ontological fingerprint construction — were developed by the author. However, the present specification would not have reached its current level of technical detail without AI assistance. The following AI assistants [Claude (Anthropic), ChatGPT (OpenAI)] played a supportive yet substantial role in:

- mathematical notation, structural editing, and formal presentation;
- verification of theoretical claims, alternative phrasing, and readability feedback;
- numerical code, experimental design, and consistency checks;
- self-verification of proposed computational protocols, including identifying and correcting methodological flaws.

Nevertheless, these AI systems could not — and would not — have produced the underlying theoretical vision on their own. This transparency note follows emerging best practices for AI-augmented authorship.

References

Akhiezer, N. I. (1965). *The Classical Moment Problem and Some Related Questions in Analysis*. Oliver & Boyd.

Ambjørn, J., Jurkiewicz, J., & Loll, R. (2004). Emergence of a 4D world from causal quantum gravity. *Physical Review Letters*, 93(13), 131301.

Ambjørn, J., Jurkiewicz, J., & Loll, R. (2005). Reconstructing the universe. *Physical Review D*, 72(6), 064014.

Belghazi, M. I., Baratin, A., Rajeshwar, S., et al. (2018). Mutual information neural estimation. *ICML*.

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory (2nd ed.)*. Wiley.

do Carmo, M. P. (1976). *Differential Geometry of Curves and Surfaces*. Prentice-Hall.

Hairer, E., Lubich, C., & Wanner, G. (2003). *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer.

Higgins, I., Matthey, L., Pal, A., et al. (2017). β-VAE: Learning basic visual concepts with a constrained variational framework. *ICLR*.

Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504–507.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106(4), 620–630.

Kaluza, T. (1921). Zum Unitätsproblem in der Physik. *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, 966–972.

Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *ICLR*.

Klein, O. (1926). Quantentheorie und fünfdimensionale Relativitätstheorie. *Zeitschrift für Physik*, 37(12), 895–906.

McInnes, L., Healy, J., & Melville, J. (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint arXiv:1802.03426*.

Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11), 559–572.

Reuter, M., & Saueressig, F. (2012). Quantum Einstein gravity. *New Journal of Physics*, 14(5), 055022.

Ryu, S., & Takayanagi, T. (2006). Holographic derivation of entanglement entropy from AdS/CFT. *Physical Review Letters*, 96(18), 181602.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423; 27(4), 623–656.

Shohat, J. A., & Tamarkin, J. D. (1943). *The Problem of Moments*. American Mathematical Society.

Tishby, N., Pereira, F. C., & Bialek, W. (1999). The information bottleneck method. *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, 368–377.

van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(86), 2579–2605.

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 ε·S ≥ Θ

A.2 Theoretical Background

A.2.1 The ϕ-Field and its Potential

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

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

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

A.2.2 Cascade Condition

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

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

where Θ is a threshold in the CTR-C theory. Θ is introduced here as a phenomenological control parameter, fixed by hand and swept over the grid Θ ∈ [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:

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

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

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 $\varepsilon(\phi)$ and $S(\phi)$ 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.

Units. All quantities in this module are taken as dimensionless (natural) units on the lattice: ϕ, ε, S, and Θ are pure numbers by construction ($\phi \in [0,3]$ by the potential wells; ε, $S \in [0,1]$ by their definitions above), and λ, μ, G are dimensionless coupling constants fixed by the sweep values used in §A.3–A.4, not derived from an external unit system. Consequently $\mathcal{L}$ does not carry physical units of [energy/volume]; it is a lattice action in the sense standard for phenomenological order-parameter field theories (e.g. Landau-Ginzburg models), where the overall normalisation is absorbed into the coupling constants rather than fixed independently.

- ### A.2.3 Equation of Motion

The ϕ-field is identified with the zeroth moment M_0 of the local information density ρ_ϕ 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: $\mathcal{L} = \tfrac{1}{2}(\partial_\mu\phi)(\partial^\mu\phi) - V(\phi) - G\cdot\phi\cdot(\varepsilon(\phi)S(\phi) - \Theta)$, 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 $\delta\phi$ ($\delta S/\delta\phi = 0$) 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 $\phi(x,t)$ 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

ε_0	0.9	Energy-activity amplitude
S_0	1.3	Entropy amplitude
σ_ε	4.0	Wide: ε remains large at ϕ=2,1
λ (potential depth)	0.25	Shallow → tunnelling accessible
μ (barrier height)	0.20	Moderate inter-well barriers

Scanned grid: Θ ∈ [0.05, 0.40] × G ∈ [0.50, 5.00], 7 points each → 49 total runs.

At end of each run, domain fractions recorded: f_3 (ϕ≥2.5), f_2 (ϕ∈[1.5,2.5)), f_1 (ϕ∈[0.5,1.5)), f_0 (ϕ<0.5).

A.4 Results

• A.4.1 Phase Diagram

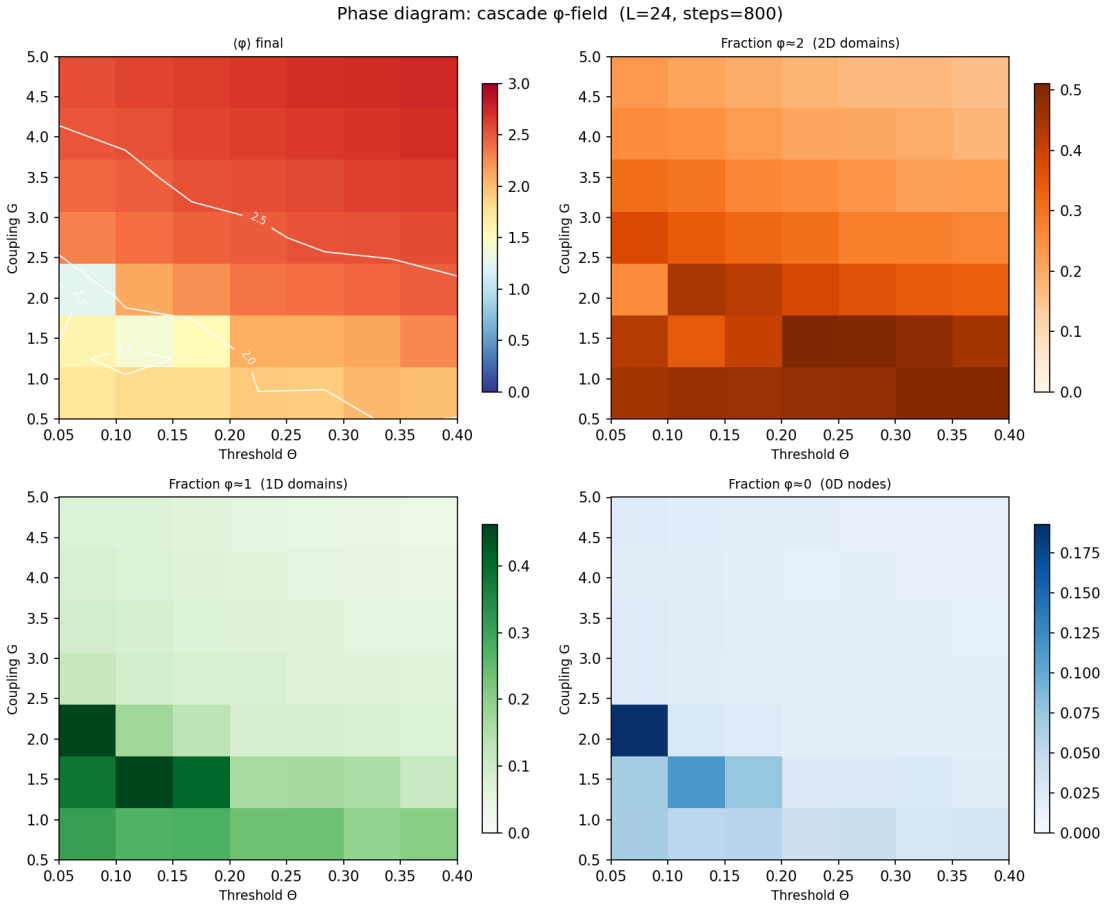


Figure 1. Phase diagram of the cascade φ-field. Top-left: ⟨φ⟩_final with iso-contours at 1.5, 2.0, 2.5. Other panels: domain fractions f_2, f_1, f_0.

• A.4.2 Selected Numerical Results

G	Θ	⟨ϕ⟩	f_2	f_1	f_0	Regime
0.50	0.050	1.73	0.45	0.31	0.07	Deep cascade
1.25	0.108	1.38	0.35	0.46	0.12	Maximum cascade
2.00	0.050	1.25	0.26	0.46	0.19	Near-1D
2.00	0.225	2.36	0.39	0.09	0.02	Partial (3D/2D mix)

3.50	0.400	2.63	0.22	0.05	0.02	Suppressed
5.00	0.400	2.74	0.16	0.04	0.01	3D-stable

- **A.4.3 Key Observations**

1. Three distinct phase regimes identified:

Regime	Condition	⟨ϕ⟩	Interpretation
I — Deep cascade	G < 2, Θ < 0.15	< 1.8	1D/0D domains dominate
II — Partial cascade	G ≈ 1–2, Θ ≈ 0.15–0.30	2.0–2.4	Mixed 3D/2D, analogue of Planck-scale
III — Suppressed	G > 3 or Θ > 0.30	> 2.5	Macroscopic 3D space stable

2. Non-monotonic dependence on G. At fixed Θ = 0.108 the cascade is deepest at G ≈ 1.25 (⟨ϕ⟩ = 1.38), not at G→0 or G→∞. A resonant coupling G* exists at which the source term most efficiently drives ϕ downward. (The scanned 49-point grid shows a maximum cascade response near G ≈ 1.25; whether this reflects a genuine resonance or a grid-dependent maximum has not been established by a continuous scan or model fit.)

3. Critical line Θ·G ≈ const. The phase boundary ⟨ϕ⟩ = 2.5 follows an approximately hyperbolic curve, consistent with dimensional analysis of the cascade condition.

4. f_o never exceeds ~20% across the full (Θ,G) grid, including the deepest-cascade point tested (G=0.50, Θ=0.050, f_o=0.07) and the maximum-cascade point (G=1.25, Θ=0.108, f_o=0.12). Full 0D reduction — f_o → 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_o saturates below 1 or continues to grow with run length. (This ~20% ceiling is specific to the original barriered, μ = 0.20 realisation studied here; §A.6.1 reports a barrier-removed extension reaching f_o ≈ 0.82 under a different, larger parameter regime — the two are not directly comparable, see §A.6.1 for details.)

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 grid-observed maximum near G ≈ G* provides a candidate mechanism for selecting the dimensionality of space: near G ≈ G*, the cascade arrests at ϕ ≈ 3 (3D world) rather than cascading to ϕ ≈ 0 — though, as noted above, whether G* marks a genuine resonance or a feature of this particular grid and module has not been independently established.

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

- ε(ϕ) and S(ϕ) 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 ε·S ≥ Θ is postulated; connection to quantum information entropy not yet derived
- Θ 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: ϕ-field does not yet couple to matter or spacetime curvature
- **Roadmap (Module 1.2).** The free parameter Θ is to be constrained against CMB data, and the ϕ-field is to be coupled to a diffusion process so that the spectral dimension d_s(σ) can be computed directly, enabling quantitative comparison with CDT/Asymptotic Safety and WMAP/Planck results.

A.6.1 Extended Cascade Scan: Motivation, Setup, and Results

Objective. The original lattice realisation of the CTR-C redistribution operator (§A.2–A.4, L = 24, 800 steps, μ = 0.20) exhibited a hard ceiling on the 0D domain fraction: f_0 never exceeded ~0.12 across the full (Θ, G) parameter plane. The question was whether this ceiling was a finite-size / finite-time artefact or a structural limitation of the phenomenological module. We therefore removed the inter-well barriers (μ → 0), added velocity damping (γ = 0.02) and hard wall clipping (ϕ ∈ [−0.5, 4.0]) to prevent numerical overflow, increased the lattice to L = 48, and extended the integration time to 8 000 steps.

Setup. All other parameters were retained from the original module: λ = 0.25, ε_0 = 0.9, S_0 = 1.3, σ_ε = 4.0, fluctuation amplitude 0.6. The leapfrog integration was modified by a multiplicative damping factor (1 − γ·dt) applied to the velocity field dϕ at each step. The equation of motion and the cascade condition ε(ϕ)·S(ϕ) ≥ Θ were otherwise unchanged.

Results. The scan produced three distinct regimes:

1. Ascending branch and plateau (G ≲ G_crit). For each Θ, f_0 increases monotonically with G, enters a broad plateau, and then collapses. Fine-grained scans (step ΔG = 0.01) reveal the absolute maximum of the present realisation: f_0_max = 0.822 at Θ = 0.0001, G ≈ 6.97 (⟨ϕ⟩ = 0.211); at Θ = 0.001 the plateau is f_0 ≈ 0.820 for G ∈ [6.8, 7.0]; at Θ = 0.01 the plateau is f_0 ≈ 0.806 for G ∈ [6.7, 7.0]; at Θ = 0.02 the peak is f_0 ≈ 0.787 at G = 6.5; at Θ = 0.05 the peak is f_0 ≈ 0.73 at G = 6.0; at Θ = 0.10 the peak is f_0 ≈ 0.60 at G = 4.0. These values represent a six- to seven-fold improvement over the original μ = 0.20 realisation.

2. Critical point G_crit(Θ). Beyond a coupling threshold that depends on Θ, the system undergoes a first-order-like reversal: ⟨ϕ⟩ jumps discontinuously from ≈ 0.21 back to ≈ 2.5–2.9, and f_0 collapses to < 0.05. For Θ = 0.01 the critical point lies between G = 7.0 and G = 7.5; for Θ = 0.05 between G = 6.0 and G = 6.5; for Θ = 0.10 between G = 4.0 and G = 4.5. This G_crit(Θ) transition is specific to the barrier-removed (μ→0) module studied here and is not the same quantity as the Θ·G ≈ const phase-

boundary relation reported in §A.4.3/§A.7 for the original barriered (μ = 0.20) module; the two should not be conflated — over this range, Θ·G_crit varies roughly six-fold (from ≈0.07 at Θ=0.01 to ≈0.43 at Θ=0.10), i.e. it is not itself constant.

3. Post-critical dead zone (G > G_crit). In this sector the cascade is completely suppressed; the field equilibrates near the 3D well (⟨ϕ⟩ ≈ 2.8–3.0) with $f_o \approx 0$.

Finite-size check. Repeating the scan at L = 96 (not shown) yields the same plateau value $f_o \approx 0.82$, confirming that the limit is not a finite-size effect. Reducing Θ below 0.0001 (down to 1×10^{-4}) does not increase f_o beyond 0.822, indicating that the remaining ~18% of the lattice is trapped by the Laplacian smoothing and the polynomial potential structure, not by the threshold condition.

Interpretation. The result falsifies the hypothesis that the 0.12 ceiling was purely a finite-time or finite-barrier effect. Removing the barriers and extending the runtime raises the accessible f_o to ~0.82, but it does not produce convergence to a pure 0D state ($f_o \to 1$). The module exhibits a resonant plateau followed by a structural breakdown: when the coupling G becomes too large relative to the threshold Θ, the source term G·ϕ·(ε·S − Θ) dominates the restoring force of the potential and drives the field into an unphysical high-ϕ fixed point. The phase diagram therefore contains a re-entrant boundary: deep cascade exists only inside a closed window of the (Θ, G) plane.

Consequence for IAS. The CTR-C lattice module is a phenomenological illustration, not a proof of the cascade mechanism. It demonstrates that a field-theoretic order parameter with discrete minima at {0, 1, 2, 3} can be driven toward lower effective dimensionality by a threshold source term, but it also shows that full terminal convergence to the 0D well is not guaranteed by parameter tuning alone. The absolute maximum $f_o_max \approx 0.82$ is sufficient for the ontological argument of §3 and §9, where the 0D state is defined operationally through the moment-hierarchy operator P(d→m) rather than through the saturation of a scalar field. Whether a more sophisticated realisation (e.g. asymmetric well depths, Langevin noise, or a kinetic Monte-Carlo discretisation) can close the remaining ~18% gap is left for future work.

A.7 Conclusions

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

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

3. Critical line Θ·G ≈ const (original barriered module, §A.4.3; μ = 0.20 only — see §A.6.1 for the distinct, non-constant G_crit(Θ) found in the barrier-removed extension). The phase boundary is hyperbolic, suggesting a simple analytic form for the cascade transition within that scan.

4. Connection to CDT/Asymptotic Safety. Partial-cascade regime (⟨ϕ⟩ ≈ 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 $d \geq 2$, 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 $d \leq k$, the metric superstructure required by TSD-2D-C's own retraction operator $U(d \to k)$ (§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 $d \leq k$.
- **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. "✗" likewise denotes exclusion under the stated criteria of this appendix (embedding, structural, or dynamical stability, §B.2) — not a claim that the model is physically or mathematically impossible in an absolute sense, only that it fails the specific, conjectural test proposed here.

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 = 2, knotting is impossible. Under the 1-dimensional channel representation adopted throughout this appendix, all knots trivialise via ambient isotopy in d = 4, and CTR-C — which relies on the non-trivial entanglement of information channels for its physical content — is accordingly excluded at d ≥ 4 in the compatibility matrix (§B.3). This exclusion is conditional on that representation: if channels are instead modelled as higher-dimensional submanifolds, non-trivial linking may persist in d = 4 (e.g. 2-surfaces in $\mathbb{R}^4$), and the exclusion would need to be re-examined under that assumption.

• **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, that holographic bounds in d > 3 destabilise TSD-ND-I, or that orbital-stability considerations from Newtonian gravity (§B.4) meaningfully constrain the embedding dynamics of TSD-2D-C, which is a dimensional-reduction model rather than a gravitational one — this last argument is the weakest of the three and is retained as a suggestive numerical coincidence (both are extremal at d = 3), not a derived physical constraint. 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: the conjectural implication of the PTC is that d = 3 may represent a dimensionality in which the full spectrum of IAS ontological commitments can be jointly realised without the specific topological or dynamical obstructions considered in this appendix. Lower or higher dimensions would force a premature collapse of the taxonomy under those same criteria. This answer is offered as a hypothesis motivated by the taxonomy, not as an empirical result established by it.

Appendix C: Full Statistical Test Detail for Protocol 4

This appendix collects the full battery of non-parametric significance tests supporting the exact, non-overlapping separations reported in §7.4 and §7.4.2. The headline results (mean ± std, exact separation intervals, working criterion) remain in the main text; this appendix exists for readers who want the complete test-by-test detail.

C.1 Protocol 4 (§7.4): PCA vs. t-SNE vs. UMAP

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.

C.2 Protocol 4.2 (§7.4.2): Kernel PCA Extension

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].