Constraint Geometry Across Scales:
Field-Space Autoencoders, Spin–Valley Modes, and the RSVP / CLIO / TARTAN Research Program

Convergent Structure in Climate Compression,
Condensed Matter Physics, and Constraint-First Field Theory

Flyxion
Independent Researcher

May 2026

Abstract. *Two recent papers, arriving from entirely different empirical domains, independently converge on a shared structural principle: stable physical organization is carried by hidden dynamical manifolds, and observable outputs are secondary projections of that deeper constraint geometry. The Field-Space Autoencoder demonstrates this computationally, in climate emulation: geometry-aware, hierarchically explicit, constraint-preserving compression dramatically outperforms flat representational approaches. A contemporaneous study of spin–valley collective modes in twisted* WSe_2 *moiré superlattices demonstrates it physically: the physically meaningful excitations of the system propagate through charge-decoupled hidden modes that are invisible to standard local observables and require direct space-time imaging to detect. This essay examines both convergences in relation to the* RSVP *(Relativistic Scalar-Vector Plenum),* CLIO *(Constraint-Leveraged Inference and Optimization), and* TARTAN *(Trajectory-Aware Recursive Tiling with Annotated Noise) frameworks — not as direct empirical confirmation, but as evidence of a broader structural pressure forcing multiple independent research lines toward a geometry-first, admissibility-preserving ontology. We develop the formal correspondences, generalize their consequences, and identify the conceptual gap that separates operationally sophisticated systems from a fully field-theoretic account of constraint-geometry propagation.*

Contents

1. Introduction: The Structural Pressure

There is a recurring pattern in the history of scientific computation. A field develops a representational convention that works well within a particular parameter regime, and then encounters systematic failure as the physical systems it models become increasingly demanding of geometric fidelity. The failure is not usually attributable to insufficient data or insufficient parameter count. It is attributable to a mismatch between the topology of the representation and the topology of the underlying process. When that mismatch becomes severe enough, the field eventually reorganizes around a new representational ontology — one that takes the geometry of the process as primary rather than emergent.

1.1. *Representation Failure as Topological Mismatch*

The distinction between numerical error and structural failure is critical. When a climate model accumulates bias near the poles, or when a quantum simulation loses long-range coherence, the conventional response is to increase resolution, add parameters, or improve the loss function. These interventions address the symptom. The underlying cause is that the representational substrate — the grid, the token sequence, the flat latent space — is topologically incompatible with the process it is meant to model. A representational container whose geometry differs from the physical geometry does not merely introduce approximation error; it introduces a systematic bias toward trajectories that are admissible in the representation but inadmissible in the physics. No increase in parameter count resolves a topological mismatch. Only a change of representational ontology can.

This distinction — between insufficient scale and structural incompatibility — is the conceptual pivot of the entire essay. The two empirical papers examined here matter not because they involve large models or impressive benchmarks, but because they both, independently and from very different directions, arrive at architectures and measurement strategies whose design logic is driven by the recognition that the representational topology must be compatible with the physical topology. That recognition, when forced by engineering or experimental necessity, is what we call a structural pressure.

A further structural feature unifies both empirical systems: both involve long-horizon dynamics where local correctness is insufficient for global admissibility. A climate model that is accurate at each time step can still produce a trajectory that is globally inadmissible — that drifts to climatologically impossible states — because local step-by-step reconstruction does not enforce global constraint coherence. A measurement strategy that accurately records local charge density at each point in space can still be entirely blind to the collective mode structure, because the relevant physics is carried by globally organized hidden fields rather than local observables. Whenever admissibility is globally path-dependent, short-horizon optimization fails. This is the deep unifying constraint across climate systems, quantum collective modes, cognition, and semantic memory.

Climate modelling is currently undergoing exactly this transition. The dominant approach for decades has been to discretize atmospheric and oceanic dynamics onto regular grids, apply convolutional or spectral operators to those grids, and treat the resulting array of values as a compressed description of the physical state. This approach has been productive, but it has a structural defect: the sphere is not the plane, and forcing spherical dynamics into Euclidean representational containers introduces systematic distortions that accumulate across temporal integration and spatial downscaling. The distortions are not merely numerical. They reflect a genuine topological incompatibility between the representational substrate and the physical substrate.

The Field-Space Autoencoder (FSA) paper, which we treat throughout as our primary empirical reference, is a response to exactly this failure mode. Its core architectural innovation is to process climate fields natively on their natural geometric domains rather than on flattened approximations, to decompose dynamics into an explicit multi-scale hierarchy rather than leaving scale structure to be inferred by the network, and to treat cross-scale operators as structure-preserving morphisms rather than as unconstrained learned transforms. These are not merely engineering refinements. They reflect a recognizable shift in representational ontology: from storage-first to geometry-first, from flat coding to constraint-preserving projection.

The RSVP (Relativistic Scalar-Vector Plenum), CLIO (Constraint-Leveraged Inference and Optimization), and TARTAN (Trajectory-Aware Recursive Tiling with Annotated Noise) frameworks have been developed from a different direction — theoretical rather than empirical, field-theoretic rather than engineering-motivated — but they have been organized around precisely the same structural commitments. The convergence is therefore worth examining carefully. It suggests that these principles are not artifacts of any particular theoretical programme but rather structural necessities that emerge independently wherever the representational demand for geometric fidelity becomes severe enough.

1.2. *Convergent Rediscovery Rather Than Derivation*

The epistemic framing of this essay requires explicit statement, because the risk of misreading is real. The correspondence between the empirical architectures examined here and the RSVP/CLIO/TARTAN framework is not a claim of priority, derivation, or retrospective validation. Neither the FSA paper nor the moiré spectroscopy group arrived at their results by consulting these frameworks. They arrived there under independent empirical pressure from the physical systems they were studying. The theoretical frameworks were developed independently, under different pressures, from a different direction.

What the convergence demonstrates is that the structural commitments in question are being driven by something real in the domain rather than by the preferences of any particular research group. When engineers building climate emulators and physicists imaging quantum collective modes arrive at overlapping structural principles without consulting each other or any shared theoretical source, that overlap is evidence of a structural necessity. The theoretical frameworks provide a vocabulary for stating that necessity precisely; they do not constitute its source.

This essay proceeds through five stages. In the first, we examine the FSA architecture on its own terms and identify the specific ways in which it departs from flat representational conventions. In the second, we develop the formal correspondence between the FSA and the RSVP/CLIO frameworks, arguing that the architecture is an engineering instance of admissibility-preserving projection. In the third, we generalize the consequences through TARTAN, showing that the FSA approach implements several TARTAN principles in a concrete domain. In the fourth, we examine the spin–valley moiré paper and argue that it provides a physical analogue of the same structural principles: hidden collective modes propagating through constraint manifolds that are invisible to standard local observables. In the fifth, we draw the unified thesis from both papers, identify the conceptual gap that separates operationally powerful systems from a fully field-theoretic account, and develop the sheaf and cohomological formalization of that gap.

2. The Field-Space Autoencoder: Architecture and Failure Modes Addressed

2.1. *The Standard Architecture and Its Defects*

The dominant representational pipeline in data-driven climate modelling follows a by-now familiar pattern: map the physical field to a fixed discretization, apply learned operators to that discretization, and decode the result back to a physical prediction. The appeal of this pipeline is its generality. Convolutional networks, transformers, and their hybrids can all be applied within it without modification, and the engineering infrastructure that has been developed for flat image processing transfers without fundamental change.

The defects are structural. A climate field is defined on the sphere S^2, and the natural symmetries of that domain — rotational invariance, the absence of a canonical pole, the curvature of geodesics — do not survive the transition to a regular latitude-longitude grid. Near the poles, grid cells shrink in area while remaining equal in index space, distorting both the statistics of the representation and the effective receptive fields of convolutional operators. Interpolation to a flat grid introduces artifacts that compound across temporal integration steps. Most significantly, the representational geometry diverges from the physical geometry: the allowed transformations in index space are not the allowed transformations in field space.

The FSA paper identifies this as the central failure and organizes its architecture around avoiding it. Rather than projecting the physical field onto a flat representational substrate, it operates natively on the HEALPix grid [15] — an equal-area spherical discretization that avoids polar distortion by treating all spatial locations uniformly. The architectural consequence is that the network never operates on a representational object whose topology contradicts the topology of the process it is trying to model. At 64× compression, the Field-Space models achieve reconstruction errors (RMSE $\approx 0.28\,^{\circ}$C) lower than the HEALPix convolution baseline achieves at only 16× compression (RMSE $\approx 0.31\,^{\circ}$C), representing a fourfold increase in compression efficiency at equivalent accuracy. The baseline degrades sharply at extreme compression ratios (256×, 1024×) where the Field-Space models remain viable.

2.2. *Why Flat Latent Spaces Fail*

It is worth distinguishing three distinct fidelity criteria that are frequently conflated in representation learning: *reconstruction fidelity*, *dynamical fidelity*, and *admissibility fidelity*. A model with high reconstruction fidelity produces outputs that closely resemble the training data in pixel-level or token-level statistics. A model with high dynamical fidelity produces outputs whose temporal evolution is statistically consistent with the underlying physical dynamics. A model with high admissibility fidelity produces only outputs that are structurally consistent with the constraint geometry of the domain — outputs that lie on or near the manifold of physically realizable states, not merely near the manifold of statistically frequent ones.

Standard flat latent space architectures optimize for reconstruction fidelity and, to a lesser extent, dynamical fidelity. They do not optimize for admissibility fidelity, and they provide no mechanism for enforcing it. The result is that their outputs can be locally accurate — close to the training distribution in the sense of standard loss metrics — while being globally inadmissible: on trajectories that no physical system could actually follow. The failure mode is not large error but wrong error: the model produces states that look plausible in isolation but are collectively inconsistent with the admissibility structure of the domain.

The FSA architecture's geometric commitments move it toward admissibility fidelity without explicitly naming it as such. By operating on the native geometric domain and enforcing explicit cross-scale compatibility, the architecture makes it structurally harder to produce globally inadmissible outputs. The latent space organization that results is evidence that this structural hardening is working: the compressed representations organize themselves according to the dynamical geometry of the physical system, not merely according to reconstruction efficiency.

2.3. *Multi-Scale Residual Decomposition*

The second major architectural innovation is the explicit multi-scale decomposition strategy. Rather than asking a single network to learn all spatial scales simultaneously from raw data, the FSA architecture decomposes the field into a hierarchy of scale-specific representations:

$$X = X_0 \oplus \bigoplus_{k=1}^{n} R_k,$$

where X_0 is a coarse global structure and each R_k is a residual correction at progressively finer spatial resolution.

This decomposition is not merely a computational convenience. It encodes a specific ontological commitment: large-scale structure is primary, and fine-scale structure exists conditionally upon it. The coarse manifold X_0 acts as an admissibility regulator for the residual layers. A fine-scale refinement R_k is constrained to be compatible with the coarser structures at all lower resolutions, not merely with the immediately preceding layer. This creates a cascade of compatibility constraints that strongly limits the space of admissible high-resolution reconstructions.

The architectural specification of the cross-scale operators makes this explicit. The compression operator takes fields at multiple resolutions as input and produces representations at an intermediate resolution:

$$\mathcal{F}_{k\to j} : \{r^{(k)}, r^{(k+1)}, \ldots\} \to \{r^{(j)}\}.$$

The decompression operator inverts this, reconstructing finer resolutions from coarser inputs:

$$\mathcal{F}^{-1}_{j\to k} : \{r^{(j)}, r^{(j-1)}, \ldots\} \to \{r^{(k)}\}.$$

The explicit specification of these operators as functions of multiple scale strata, rather than as point-to-point mappings between adjacent scales, is a strong architectural commitment to the idea that compatibility is a global rather than local constraint.

2.4. *Latent Space Organization*

The most philosophically significant result in the FSA paper is the spontaneous organization of the latent space. The authors analyse the compressed representations of 84 years of daily ERA5 near-surface air temperature data (1940–2024) using t-SNE dimensionality reduction. The two Field-Space Autoencoder variants — convolutional and transformer-based — produce latent spaces that are, without any explicit temporal or climatological supervision, organized into cyclic seasonal trajectories when coloured by month, and a smooth directional drift from earlier to more recent states when coloured by year. Both architectures were trained *exclusively on residual fields*, deliberately excluding the coarse base component that contains most of the seasonal and warming signal. The organized latent structure therefore cannot be attributed to exposure to large-scale climatological information during training; it emerges from the architecture's representational commitments alone.

In sharp contrast, the HEALPix convolution baseline — trained on the full raw fields including the seasonal cycle and warming trend — produces a latent space with no comparably coherent structure in the t-SNE projection. The Field-Space architecture, operating only on fine-scale residuals, recovers the dynamical geometry of the full climate system. The baseline, operating on the full field, loses it. This reversal is the empirical result that most directly confirms the essay's central argument: surface fidelity and dynamical fidelity are orthogonal properties of a representation, and optimizing for one can actively impede the other.

2.5. *Compression as Constraint Transport*

The FSA field-space operators do something subtler than dimensionality reduction. They transport the compatibility structure of the field across scale boundaries. Each operator $\mathcal{F}_{k\to j}$ takes as input the field representations at multiple resolution strata and produces an output that is simultaneously consistent with all of them. This is not compression in the information-theoretic sense of finding a shorter description of the same content. It is compression in the constraint-geometric sense of finding a representation that respects the admissibility conditions imposed by the coarser scales.

The distinction matters because it determines what generalizes. A system that compresses by removing information discards the fine-scale structure and cannot recover it without additional data. A system that compresses by transporting compatibility structure retains the *conditions* under which fine-scale structure is admissible, and can therefore reconstruct compatible fine-scale structure for any input that satisfies those conditions — including inputs at resolutions never seen during training. Zero-shot super-resolution is the empirical signature of constraint transport rather than information reduction.

2.6. *Zero-Shot Super-Resolution*

The FSA paper demonstrates a capability with direct theoretical significance. The transformer-based model, trained at a fixed 64× compression ratio on ERA5 data at HEALPix level 8 (approximately 30 km resolution), can produce outputs at HEALPix levels 7 and 6 — corresponding to 4× and 16× super-resolution — by setting finer residual levels to zero at inference time. The model has never seen the pairing of level-6 input with level-8 output during training. The transformer exhibits only a small decrease in skill at these novel resolution configurations and no spatial artifacts. The convolutional variant, by contrast, degrades rapidly and produces prominent HEALPix grid artifacts at level-6 input, suggesting it overfit to the specific residual resolutions seen during training rather than learning transferable compatibility conditions.

The paper also applies the pretrained transformer to MPI-ESM1.2-HR historical simulations natively produced at approximately 100 km resolution — a domain the model was never trained on. After encoding at HEALPix level 6 and decoding to level 8, the model synthesizes high-frequency spatial structures entirely absent from the source simulations, as confirmed by spectral analysis showing maintained power at high multipoles where the climate model falls off. This is sparse reconstruction from constraint geometry operating across both resolution and model boundaries. The model has not memorized fine-scale ERA5 patterns; it has learned the admissibility conditions relating coarse and fine scales, and those conditions transfer.

2.7. *Feature Learning Versus Constraint Learning*

Most representation-learning systems are implicitly optimised to learn *features*: statistical regularities present on the observable surface of the training distribution. A feature is a pattern that recurs across training examples; feature learning is the process of identifying and encoding those patterns in a form that supports downstream prediction. The FSA results suggest a fundamentally different interpretation of what successful generalisation requires.

The architecture does not succeed because it has memorised high-frequency climate features. It succeeds because it has learned the *compatibility relations* that govern which fine-scale structures are admissible given coarse-scale conditions. These compatibility relations are not surface patterns; they are structural constraints that hold across domains, resolutions, and model boundaries because they are properties of the physical system rather than of the training corpus.

The distinction matters because feature-learning systems and constraint-learning systems fail differently. A feature-learning system fails whenever the statistical surface of the input diverges from the training distribution: when the domain shifts, the features no longer match, and generalization collapses. A constraint-learning system fails only when the underlying admissibility structure itself changes — a much more severe condition that corresponds to genuine physical change rather than mere distributional shift. The former is distributionally fragile; the latter is structurally robust.

The convolutional baseline's behaviour at zero-shot resolution makes this contrast sharp. It overfits to the specific HEALPix residual resolutions present during training and produces artifacts when these change, because it has learned features of the training residual distribution. The transformer exhibits no such artifacts because it has learned the compatibility conditions that relate scale strata, and those conditions do not depend on which specific strata were used during training. This asymmetry is the operational signature of constraint learning. It also provides a cleaner conceptual target for the scaling critique developed later: the question is not how many features a system has memorised but whether it has learned the constraint structure of the domain.

3. The RSVP Framework: Fields, Admissibility, and Trajectory Geometry

3.1. *Foundational Ontology*

The Relativistic Scalar-Vector Plenum framework takes as its fundamental ontological commitment the idea that the relevant objects of a physical theory are not states but *fields over possibility spaces*. A state is a momentary cross-section of a field; a trajectory is a sequence of such cross-sections satisfying admissibility constraints; and the accessible future is the set of trajectories that are reachable given the current field configuration and accumulated history.

Formally, the RSVP state is a triple

$$X_t = (\varphi_t, \mathbf{v}_t, S_t),$$

where φ_t is the scalar plenum density governing long-range coherence, $\mathbf{v}_t$ is the vector flow field governing directional transport, and S_t is the local entropy density governing the volume of admissible refinements. The dynamics of each component are coupled:

$$\partial_t \varphi = D_\varphi \nabla^2 \varphi - \nabla \cdot (\varphi \mathbf{v}) + R_\varphi - \lambda_\varphi S, \tag{1}$$

$$\partial_t \mathbf{v} = -(\mathbf{v} \cdot \nabla)\mathbf{v} + \nu \nabla^2 \mathbf{v} - \nabla \varphi + \mathcal{T}(\mathbf{v}), \tag{2}$$

$$\partial_t S = D_S \nabla^2 S + \sigma(\varphi, \mathbf{v}) - \gamma S. \tag{3}$$

The coupling term $\lambda_\varphi S$ in equation (1) is significant: local entropy density suppresses scalar coherence. Regions of high entropy are regions where many trajectories remain admissible, and coherent structure is correspondingly weaker. The scalar field φ should therefore be understood not merely as a density field but as a coherence regulator.

The transport operator $\mathcal{T}(\mathbf{v})$ in equation (2) is equally important. It acts as a geometric connection on the flow field: it absorbs local topological defects — vortices, phase slips, singular points in the vector field — that would otherwise break the global coherence of the flow. In the spin–valley context developed later in this essay, $\mathcal{T}(\mathbf{v})$ plays the role of the gauge-like term that couples the superfluid velocity to the underlying moiré superlattice potential, ensuring that the collective flow remains globally admissible despite local geometric irregularities. This structural symmetry — between $\mathcal{T}(\mathbf{v})$ as defect absorber in RSVP and the phase mode $\mathbf{v}_t \sim \nabla\theta$ as coherence carrier in the spin–valley superfluid — is one of the most precise formal correspondences between the theoretical framework and the condensed matter experiment.

3.2. *Admissibility and the Accessibility Structure*

The central theoretical construct in RSVP is the admissibility function $\kappa(\omega, H_t)$, which assigns a viability score to each candidate transition ω given the accumulated event history H_t. The admissible set at time t is

$$\mathcal{A}_{\text{dm}}(X_t) = \{\omega \in \Omega_t \mid \kappa(\omega, H_t) \geq \kappa_*\},$$

where κ_* is the viability threshold below which transitions are physically or structurally prohibited. The event history $H_t = \{e_0, e_1, \ldots, e_t\}$ is append-only: transitions from Ω_t

select events into H_{t+1}, and this selection is irreversible.

The global model space $\mathfrak{M} = \{s_t \mid t \in \mathbb{R}\}$ supports a family of history-conditioned accessibility slices

$$\mathcal{A}_t = \{s \in \mathfrak{M} \mid s \text{ reachable given } H_t\}.$$

Temporal asymmetry is encoded in the strict monotone shrinkage

$$\mathcal{A}_{t+1}^{\text{consistent}} \subsetneq \mathcal{A}_t^{\text{counterfactual}} :$$

the consistent slice of states compatible with H_{t+1} is always strictly smaller than the counterfactual envelope of states that were accessible before the transition at t.

This structure provides a formal account of irreversibility that does not depend on thermodynamic entropy alone. A transition is irreversible not merely because it increases entropy but because it eliminates branches from the accessibility structure that cannot be recovered by any subsequent sequence of admissible transitions.

3.3. *Projection, Compression, and Admissibility Preservation*

The central epistemological problem in RSVP is what happens when a field-theoretic system is projected onto a lower-dimensional representational manifold. Let $\pi : X_t \to M$ denote such a projection, where X_t is the trajectory space and M is the compressed manifold. The projection is admissibility-preserving if and only if admissible transitions in X_t map to admissible transitions in M:

$$\omega \in \mathcal{A}_{\text{dm}}(X_t) \;\Rightarrow\; \pi(\omega) \in \mathcal{A}_{\text{dm}}(M).$$

When this condition fails, the compressed representation supports trajectories that would be inadmissible in the underlying physical system. The model learns to navigate the compressed manifold rather than the physical one, and generalization breaks down precisely at the boundaries where the two manifolds diverge.

This failure mode is not merely a theoretical concern. It is the failure that the FSA paper is specifically designed to avoid. The insistence on processing fields on their natural geometric domain, rather than on a flattened approximation, is precisely an insistence on admissibility preservation: the representational geometry must be compatible with the physical geometry.

3.4. *Entropy as Refinement Volume*

It is worth pausing on a potential confusion about the role of entropy in RSVP. The entropy density S_t is not thermodynamic disorder in the standard sense, nor is it Shannon information-theoretic entropy over probability distributions. It is *geometric admissibility volume*: the local measure of how many fine-scale refinements are consistent with the current coarse-scale state and accumulated history.

In a region of low entropy, the coarse-scale field φ_t tightly constrains the admissible fine-scale states: the compatibility conditions at that location leave very little room. Reconstruction of fine-scale detail from coarse inputs is nearly deterministic in such regions, because almost all admissible refinements are mutually close. In a region of high entropy, the coarse-scale structure leaves the fine-scale states largely unconstrained: many very different fine-scale configurations are admissible, and reconstruction becomes an act of selection among a large admissible set rather than a nearly unique recovery.

This reinterpretation has a precise implication for the FSA architecture. The regions where the climate model has low reconstruction uncertainty — where its predictions are most reliable and consistent — correspond to low-entropy regions in the RSVP sense: regions where the global constraint structure tightly regulates the admissible fine-scale states. The regions where the model has high uncertainty correspond to high-entropy regions: where the constraint structure leaves many equally admissible fine-scale configurations. The architecture's zero-shot super-resolution capability is strongest precisely in the low-entropy regime, where constraint transport alone is sufficient to determine the admissible refinements.

3.5. *Scale Hierarchy as Constraint Cascade*

The RSVP framework treats scale decomposition as a constraint cascade rather than a feature hierarchy. At the coarsest scale, the scalar field φ encodes global accessibility structure: the long-range coherence that determines which large-scale trajectories are admissible. At progressively finer scales, the vector field $\mathbf{v}$ and entropy density S encode local perturbative refinements of that structure. A region of low entropy (S small) is a region where coarse structure tightly constrains fine structure: few refinements are admissible, and the compression is nearly lossless. A region of high entropy (S large) is a region where coarse structure leaves fine structure relatively unconstrained: many refinements are admissible, and the compression is inherently lossy. This is not a limitation to be overcome; it is a structural feature of the physical system.

3.6. *Constraint Geometry Is Not Reducible to Optimization*

Many contemporary theoretical frameworks describe physical and cognitive systems as optimization processes: free-energy minimisation, variational inference, least-action trajectories, or energy-landscape descent. The RSVP framework is compatible with such descriptions but is not reducible to them, and the distinction is important for reading the empirical results correctly.

Optimization presupposes a space over which optimization occurs. Constraint geometry specifies the topology of that space itself: which trajectories exist, which transitions are admissible, which refinements remain globally consistent, and which regions of state space are structurally inaccessible regardless of energetic cost. The admissibility function $\kappa(\omega, H_t)$ is logically prior to any objective functional; it defines the feasibility region over which objectives are then minimised.

A system may therefore occupy a locally low-energy configuration while remaining globally inadmissible. Conversely, a trajectory may be energetically expensive while remaining the unique admissible option given the accumulated history. Constraint geometry regulates the structure within which optimization takes place; it cannot itself be reduced to a particular optimization objective without circularity.

This distinction becomes operationally visible wherever systems exhibit long-horizon coherence that cannot be explained through local optimization alone. The FSA latent organization, the spin–valley Goldstone propagation, and the trajectory-conditioning principles of CLIO all depend on global compatibility conditions that exceed local energetic minimisation. The convergence of these systems toward geometry-aware design is not driven by finding a better objective function; it is driven by the recognition that the geometry of the feasibility space matters independently of any objective defined over it.

3.7. *Degenerate Observables and Projection Equivalence*

A projection observable $\mathcal{O} : X \to Y$ is *degenerate* whenever dynamically distinct states in X map to the same observable state in Y:

$$x_1 \neq x_2, \qquad \mathcal{O}(x_1) = \mathcal{O}(x_2).$$

In such cases, observational equivalence does not imply dynamical equivalence. Systems that appear identical on the observable surface may possess fundamentally different admissibility structures and evolve differently under perturbation.

Degeneracy is not a correctable measurement error; it is a structural property of the projection. When the observable $\mathcal{O}$ eliminates an entire subspace X_{hid} of dynamically significant degrees of freedom, no improvement in the precision of $\mathcal{O}$-based measurements can recover the lost structure. Recovery requires a fundamentally different observational basis — one that is not degenerate with respect to X_{hid}.

The spin–valley experiment demonstrates this directly. Charge observables collapse dynamically distinct collective configurations into nearly indistinguishable observable states; the hidden spin–valley modes are not merely poorly measured but structurally eliminated by charge projection. The resolution requires space-time imaging of the spin–valley fluctuations directly, not a more precise charge measurement.

The same phenomenon appears in machine learning latent spaces. Two models may achieve nearly identical endpoint accuracy while possessing radically different latent trajectory organization and radically different generalization behaviour. Observable equivalence at the metric level does not imply equivalence of learned constraint structure. Benchmark evaluations that rely exclusively on endpoint metrics are therefore degenerate observables of the underlying representational system: they collapse structurally distinct models into the same observable class.

4. The CLIO Framework: Constraint-Leveraged Inference

4.1. *The Inference Principle*

The CLIO (Constraint-Leveraged Inference and Optimization) framework extends the ontological commitments of RSVP into a computational account of how admissibility-preserving projection systems perform inference. The central claim is that effective inference does not operate by exhaustively searching the space of possible outputs and selecting the most probable. It operates by first identifying the constraint structure of the target domain and then projecting from that constraint structure onto the space of admissible outputs.

Formally, the CLIO inference pipeline is

$$\omega \xrightarrow{\Pi} p \overset{\mathcal{R}}{\rightsquigarrow} \hat{\ell} \xrightarrow{\mathcal{V}} \omega',$$

where $\omega \in \mathcal{S}_{\text{valid}}$ is a constraint-valid structure, Π is the projection functor mapping it to a linguistic or computational representation p, $\mathcal{R}$ is a stochastic rendering operator producing an output $\hat{\ell}$, and $\mathcal{V}$ is a verification map returning either a validated structure ω' or failure $\perp$.

The rendering operator $\mathcal{R}$ is stochastic rather than deterministic because the projection Π is in general many-to-one: multiple constraint structures may project to the same representational point, and multiple admissible outputs may be consistent with a single constraint structure. The role of $\mathcal{V}$ is to check whether the rendered output lands within the admissible region rather than on the boundary of the representational manifold where the projection has become unreliable.

4.2. *Sparse Reconstruction and Latent Geometry*

The most distinctive CLIO prediction is that effective inference systems should exhibit sparse reconstruction: the ability to produce detailed, coherent outputs from compressed constraint representations, without exhaustive memorization of the fine-scale detail. This prediction follows directly from the RSVP account of entropy: in low-entropy regions where constraint structure is tight, detailed reconstruction follows from the constraints alone, without additional information. The constraint structure is sufficient.

This is a strong claim, and one that the FSA paper corroborates in a specific domain. The zero-shot super-resolution capability of the FSA is precisely sparse reconstruction: the model produces fine-scale detail that it has never directly observed by applying learned constraint compatibility relations. The coarse state is sufficient to determine the admissible fine-scale refinements; the model need not memorize those refinements explicitly.

Proposition 4.1. *Let $\pi : X \to M$ be an admissibility-preserving projection, and let $x \in X$ be a state with $S(x) < \epsilon$ for some small $\epsilon > 0$. Then the pre-image $\pi^{-1}(\pi(x))$ has volume* *bounded by a function of ϵ that tends to zero as $\epsilon \to 0$. In the limit of zero entropy, the coarse representation determines the fine-scale state uniquely.*

The proposition formalizes the intuition that low-entropy regions are regions of tight constraint. The FSA zero-shot super-resolution result is an empirical demonstration that the trained architecture has learned to identify these regions and exploit the tight constraints they impose.

4.3. *Latent Manifold as Trajectory Space*

The spontaneous organization of the FSA latent space into cyclic seasonal trajectories and historical warming drift is, from a CLIO perspective, a confirmation that the architecture has succeeded in preserving trajectory geometry during compression. The latent manifold is not a cloud of independently learned feature vectors; it is a geometric representation of the dynamical system's phase space, complete with the qualitative structure — the cycles, the drift, the topological organization — that characterizes the underlying physics.

This is what CLIO means by admissibility-preserving compression: the compressed representation preserves not merely the instantaneous state but the geometry of the space of admissible transitions between states. A model that has learned this geometry can generalize to novel states and novel time horizons in a way that a model organized around reconstruction efficiency cannot.

Definition 4.2 (Trajectory-Preserving Compression)**.** A compression $\pi : X \to M$ is *trajectory-preserving* if, for every admissible trajectory $\gamma : [0, T] \to X$ in the original space, the projected trajectory $\pi \circ \gamma : [0, T] \to M$ is admissible in the compressed space, and the mapping $\gamma \mapsto \pi \circ \gamma$ is a homeomorphism from the space of admissible trajectories in X to the space of admissible trajectories in M.

Trajectory-preserving compression is stronger than the basic admissibility-preserving projection condition. The latter requires only that individual admissible transitions survive compression; the former requires that the global structure of the trajectory space be preserved. The FSA latent space organisation suggests the architecture achieves something close to the stronger condition, at least within the training distribution regime.

4.4. *The Inference Failure of Markov Systems*

The CLIO framework provides a precise diagnosis of why standard autoregressive systems fail at long-horizon coherence. A standard next-step predictor operates without explicit constraint structure:

$$P(x_{n+1} \mid x_n), \quad H_t = \varnothing, \quad \mathcal{P}_t = \varnothing, \quad \mathcal{V} = \varnothing.$$

Without history coupling, provenance continuity, or verification, the output distribution at each step samples from the prior manifold of the training distribution rather than from the admissible future of the current trajectory. Short-horizon predictions may be locally accurate, but errors accumulate because the system has no mechanism for maintaining global constraint coherence.

The FSA architecture avoids this failure by operating on temporally structured compressed states rather than isolated frames. The coarse manifold provides the global constraint that prevents the system from drifting to locally plausible but globally inadmissible states. This is a structural advantage over autoregressive approaches, and it reflects a CLIO inference principle rather than merely an engineering choice.

5. Formal Correspondence: FSA as RSVP/CLIO Instance

5.1. *The Mapping*

Having developed both the FSA architecture and the RSVP/CLIO frameworks, we can now state the formal correspondence systematically. The mapping is not exact — the FSA is an engineering system developed without reference to the theoretical frameworks, and it lacks several of their structural features — but it is close enough to be theoretically significant.

The coarse climate manifold X_0 corresponds to the low-frequency component of the RSVP scalar field φ: the global accessibility structure that constrains all admissible fine-scale refinements. The residual layers R_k correspond to the perturbative refinement structure of $\mathbf{v}$ and S: localized corrections that are constrained by the coarser levels. The field-space operators $\mathcal{F}_{k\to j}$ correspond to constraint-preserving projection morphisms in the CLIO pipeline. The compressed latent manifold corresponds to the entropy-weighted accessibility geometry:

$$S_{\text{lat}} = \log|\mathcal{A}(X)|,$$

where $|\mathcal{A}(X)|$ denotes the volume of the accessibility slice.

More precisely:

Field-Space Autoencoder	RSVP / CLIO
Coarse climate manifold X_0	Low-frequency accessibility field φ
Residual scale R_k	Perturbative refinement of $\mathbf{v}$, S
Cross-scale operator $\mathcal{F}_{k\to j}$	Approximation of projection functor Π and verification $\mathcal{V}$
Latent manifold geometry	Entropy-weighted accessibility slice
Zero-shot super-resolution	Sparse reconstruction from constraint
Seasonal cycle in latent space	Admissible trajectory curvature
Historical warming drift	Irreversible shrinkage $\mathcal{A}_{t+1}^{\text{con}} \subsetneq \mathcal{A}_t^{\text{ctf}}$
Sphere-preserving processing	Admissibility-preserving projection

5.2. *Operational Versus Structural Preservation*

The significance of the correspondence in the table above lies not in identity of formalism but in convergence of architectural pressure. The FSA was not designed to instantiate admissibility geometry explicitly; rather, admissibility-preserving structure emerged as a practical requirement for stable multi-scale representation on physically constrained domains. The table should therefore be read as mapping the engineering solutions that pressure forced onto the theoretical concepts that the RSVP/CLIO framework formalizes — not as evidence that the engineers were implementing the framework.

This distinction is the key to reading the correspondence correctly. An architecture that *operationally* preserves admissibility is one whose outputs satisfy physical admissibility conditions because the training data is organized around such conditions. The preservation is a consequence of the training distribution, not a designed feature. An architecture that *structurally* preserves admissibility is one that explicitly encodes the admissibility conditions and enforces them regardless of distributional proximity to the training data. The FSA achieves operational preservation; the RSVP/CLIO/TARTAN framework specifies what structural preservation would require.

The practical consequence of this distinction becomes visible at distribution boundaries. An operationally admissibility-preserving system generalizes within the training distribution and begins to degrade at its edges. A structurally admissibility-preserving system generalizes to any input satisfying the same formal admissibility conditions as the training data, regardless of its distributional proximity. When an operational system encounters an out-of-distribution perturbation severe enough to push it off the training manifold, it does not merely degrade gracefully; it undergoes *topological delamination* — the compressed representational trajectory separates from the physical admissibility manifold entirely, because the system has no closed-loop mechanism for detecting when it has left the admissible region. Structural preservation would require precisely such a mechanism: a cohomological constraint checker capable of computing the obstruction class $[\delta\sigma]$ and flagging non-trivial values before they propagate through the model.

5.3. *The Projection Failure Theorem*

The FSA paper's diagnosis of why flat convolutional architectures fail on spherical domains corresponds to a general theorem about projection failures in RSVP.

Theorem 5.1 (Projection Failure)**.** *Let $\pi : S^2 \to \mathbb{R}^2$ be any continuous surjection from the sphere to the plane. Then π cannot be admissibility-preserving for any physical field dynamics defined by rotational symmetry on S^2.*

Sketch. The failure is both topological and metric. Topologically, S^2 is compact and boundaryless, while $\mathbb{R}^2$ is non-compact; by invariance of domain, no continuous surjection $\pi : S^2 \to \mathbb{R}^2$ can be a homeomorphism, so π must either collapse distinct points or destroy the boundary-free structure of S^2. Metrically, there exists no isometry between S^2 (constant positive curvature) and $\mathbb{R}^2$ (zero curvature), and no conformal equivalence either, since a conformal map would require the Riemann sphere $\hat{\mathbb{C}} \cong S^2$ to be conformally equivalent to $\mathbb{C}$, which it is not (they have different conformal types). Consequently, any π introduces unavoidable metric distortion. The pushforward of the invariant vector fields generating $\mathrm{SO}(3)$ rotations on S^2 cannot map to a closed, smoothly invariant symmetry group under the Euclidean transformation metric without introducing coordinate singularities or metric divergence at at least one point. Since rotationally symmetric dynamics on S^2 are defined by these invariant vector fields, and their pushforward is metrically ill-behaved in $\mathbb{R}^2$, any physically admissible transition along a great circle geodesic maps to a transition that is generally inadmissible under the Euclidean symmetry group. Admissibility cannot be preserved. □

This theorem is not a new result in differential geometry, but its interpretation in the RSVP framework gives it a new significance: flat projection failure is not a numerical artifact to be managed by preprocessing but a structural incompatibility between the representational and physical geometries. The FSA paper's insistence on processing on the native geometric domain is therefore not an engineering preference but a structural necessity.

5.4. *Latent Organization as Fixed-Point Signature*

The spontaneous organization of the FSA latent space is, in RSVP terms, the signature of a fixed-point structure in the compressed accessibility geometry. A compression π that preserves trajectory geometry produces a latent manifold that is dynamically equivalent to the original state space: the attractors, cycles, and drift directions of the original dynamics appear as attractors, cycles, and drift directions in the latent space.

The RSVP fixed-point formalism (developed in the KES extension) treats this as a stability condition:

$$\mathfrak{S}^* = \mathcal{G}(\mathfrak{S}^*), \quad \rho(D\mathcal{G}_{\mathfrak{S}^*}) < 1,$$

where $\mathcal{G}$ is the recursive evolution operator and ρ denotes spectral radius. A compression is trajectory-preserving in the sense of Definition 4.3 if and only if the compressed dynamics have a fixed-point structure that is topologically equivalent to the original. The FSA latent space organization is empirical evidence that the training process has found such a compression.

5.5. *Where the Correspondence Breaks Down*

It is important to be precise about the limits of this correspondence. The FSA architecture achieves trajectory-preserving compression operationally, in the sense that its compressed representations exhibit the qualitative geometric properties expected of trajectory-preserving systems. But it does not *enforce* these properties explicitly. The correspondence is a consequence of the architecture's geometric commitments, not a designed feature.

More significantly, the FSA lacks an explicit account of admissibility. It learns which transitions are physically plausible from data, but it has no formal structure for distinguishing between transitions that are inadmissible due to physical law and transitions that are merely improbable due to the training distribution. In RSVP terms, the architecture conflates $\mathcal{A}_{\mathrm{dm}}(X_t)$ with the high-probability region of $P(x_{t+1} \mid x_t)$. For in-distribution inputs, these sets may be nearly identical. For out-of-distribution inputs, they can diverge substantially, and the architecture provides no mechanism for detecting or correcting this divergence.

This is the primary theoretical gap that the RSVP/CLIO framework addresses but the FSA does not.

6. The TARTAN Framework: Architectural Generalization

6.1. *From Instance to Architecture*

The RSVP/CLIO account of the FSA explains why the architecture works as well as it does and identifies the structural features that account for its success. The TARTAN (Trajectory-Aware Recursive Tiling with Annotated Noise) framework generalizes these features into a principled design vocabulary for admissibility-preserving architectures across domains.

TARTAN takes the following as foundational design principles. Hierarchical structure must be explicit: the architecture must specify the scale decomposition rather than leaving it to be inferred. Residuals must be semantically annotated: each residual correction must carry information about its admissibility conditions, not merely its magnitude. Coarse-to-fine reconstruction must be constraint-regulated: fine-scale detail must be generated as what is admissible given coarse structure, not as what is most probable given the training distribution. And trajectory conditioning must be global:

the architecture must maintain consistency with the full accessible history, not merely with the immediately preceding state.

6.2. *Tiling and Residual Annotation*

The TARTAN framework's central architectural commitment is to recursive tiling with annotated noise. A coarse field X_0 tiles the domain into regions of approximately homogeneous admissibility structure. Each tile is then recursively subdivided, and each subdivision is assigned a residual correction R_k that carries semantic annotation: information about what class of physical process the residual represents and what admissibility constraints govern it.

Formally, the tile decomposition is

$$\mathcal{D}_k : \Omega_t \to \{(U_i^{(k)}, R_i^{(k)})\}_i,$$

where $U_i^{(k)}$ is the i-th tile at scale k and $R_i^{(k)}$ is its annotated residual. The annotation carries both the semantic label of the residual (which physical process it represents) and the admissibility condition it must satisfy given the coarser tiles.

The FSA architecture implements an unannotated version of this structure. Its multi-scale residual decomposition tiles the domain into scale strata and assigns cross-scale operators to the boundaries between strata. The absence of explicit semantic annotation is one of the primary differences between the FSA implementation and the full TARTAN design.

6.3. *Trajectory Conditioning and the Annotation Function*

The trajectory-conditioning component of TARTAN is the feature that most directly addresses the gap identified in the previous section. Rather than conditioning each residual correction only on the current coarse state, TARTAN conditions each residual on the full admissible history:

$$R_i^{(k)}(x,t) = f\big(X_0(t), H_t, \Psi_t, \kappa(\cdot, H_t)\big),$$

where Ψ_t is the constraint field and $\kappa(\cdot, H_t)$ is the admissibility function. This ensures that fine-scale reconstructions are not merely locally consistent with the coarse state but globally consistent with the entire accessible history.

The FSA implements a version of temporal conditioning through its diffusion architecture, which operates over temporally structured compressed states. This is structurally similar to TARTAN trajectory conditioning, but without the explicit admissibility constraint. The FSA conditions on the recent past; TARTAN conditions on the admissible past, which is a strictly stronger condition.

6.4. *Comparison with the Field-Space Operators*

The FSA field-space operators $\mathcal{F}_{k \to j}$ are the architectural component that most closely resembles the TARTAN residual annotation machinery. In both cases, cross-scale operators take multiple scale strata as input, ensuring that cross-scale compatibility is a multi-way rather than pairwise constraint. In both cases, the operators are learned rather than specified, but the architectural context constrains their function.

The primary difference is that the FSA operators are learned from reconstruction objectives, while the TARTAN annotation function is trained to satisfy explicit admissibility criteria. The FSA operators will converge to whatever cross-scale compatibility structure minimizes reconstruction loss on the training distribution. The TARTAN annotation function is explicitly constrained to respect the admissibility conditions derived from the physical or structural domain. In practice, on in-distribution data, these may converge to similar solutions. On out-of-distribution data, they will diverge in ways that reflect the presence or absence of explicit admissibility enforcement.

6.5. *Why Constraint Hierarchies Generalize*

The zero-shot super-resolution capability of the FSA is an instance of a more general principle: constraint hierarchies generalize across resolution and domain in ways that direct input-output mappings do not. The reason is structural. A direct mapping from coarse input to fine output encodes the statistical regularities of the training distribution at the specific resolution pair presented during training. It has no mechanism for extrapolating to resolution pairs it has not seen. A constraint hierarchy, by contrast, encodes the compatibility conditions between scale strata: what fine-scale structures are admissible given what coarse-scale states, and why. These compatibility conditions are properties of the physical system, not of the training distribution. They hold at resolution strata that were never presented during training, because the physics does not change when the measurement resolution does.

This is the TARTAN generalization of the FSA zero-shot result. Any architecture that learns constraint compatibility rather than direct input-output mappings should exhibit zero-shot generalization to novel constraint configurations. The prediction is domain-independent: it applies equally to protein folding dynamics (where the constraints are thermodynamic admissibility conditions on backbone geometry), to turbulent flow (where the constraints are Navier–Stokes admissibility conditions across Reynolds number regimes), and to semantic coherence in language (where the constraints are logical and narrative admissibility conditions that hold across surface paraphrase variants). In each case, the generalizing object is the constraint structure, not the surface form.

6.6. *Generalizing Zero-Shot Capability*

One of the most practically significant features of the FSA — its zero-shot super-resolution capability — follows naturally from the TARTAN design principles. The capability arises because the architecture has learned the constraint compatibility relations between scale strata independently of the specific resolution targets used during training. Given those relations, it can apply them to novel resolution combinations.

The TARTAN framework predicts this capability as a general consequence of design that preserves admissibility. Any architecture that learns constraint compatibility rather than direct input-output mappings should exhibit zero-shot generalisation to novel constraint configurations. This is a testable prediction extending beyond the climate domain. An architecture trained on constraint-compatible multi-scale representations of protein folding dynamics, or of turbulent flow, or of language at multiple granularities, should exhibit analogous zero-shot capability in its respective domain.

7. Sheaf-Theoretic Formalization and the Mathematical Horizon

7.1. *From Engineering Constraints to Cohomology*

The FSA field-space operators enforce cross-scale compatibility operationally: the architecture is designed so that outputs at finer resolutions are consistent with inputs at coarser resolutions, by construction. This engineering constraint has a precise mathematical expression. Whenever a system must enforce local conditions that are simultaneously consistent with all other local conditions across a covered domain, the natural language for that requirement is sheaf theory.

A sheaf is exactly the data structure that encodes local-to-global compatibility: it assigns information to each open set in a cover, and requires that the assignments on overlapping sets agree on their intersection. The FSA cross-scale operators are engineering implementations of sheaf sections: they assign field representations to each scale stratum, and the multi-strata input requirement ensures that these assignments are mutually consistent. The shift to sheaf language is therefore not an imported mathematical metaphor but an explicit formalization of the same compatibility structure already implemented operationally by the field-space operators. The categorical machinery makes visible the structure that the engineering is already enforcing.

This also explains why the zero-shot super-resolution capability has a natural cohomological interpretation. The capability corresponds to the ability to extend a cocycle from observed scale strata to unobserved ones. That extension is trivial — uniquely determined — precisely when the local compatibility conditions at the observed scales are tight enough to constrain the admissible extensions. The sheaf language turns a benchmark result into a structural theorem about the cohomological complexity of the domain.

7.2. *Local-Global Consistency as Sheaf Condition*

The compatibility constraints that govern the TARTAN tiling structure — and that the FSA implements operationally — can be given a precise mathematical formulation in terms of sheaf theory. A sheaf over the scale hierarchy encodes the requirement that local admissibility conditions be globally consistent: what is admissible at a fine scale must be compatible with what is admissible at all coarser scales, and the compatibility conditions must satisfy a gluing axiom.

Let $\mathfrak{U} = \{U_i\}$ be an open cover of the domain $\mathfrak{M}$ by admissibility regions. The local accessibility sheaf $\mathscr{A}$ assigns to each open set U_i the admissible set of transitions within that region:

$$\mathscr{A}(U_i) = \mathcal{A}_{\mathrm{dm}}(X_t) \cap U_i.$$

The gluing condition requires that locally admissible transitions be globally consistent:

$$\mathscr{A}(U_i)\big|_{U_i \cap U_j} = \mathscr{A}(U_j)\big|_{U_i \cap U_j}.$$

When this condition fails, the system has a *global section obstruction*: there is no globally admissible trajectory consistent with the local admissibility conditions at each scale.

7.3. *Obstruction Classes and Irreducible Inconsistency*

The obstruction to finding a global section is measured by the first Čech cohomology class:

$$[\delta\sigma] \in \check{H}^1(\mathfrak{U}, \mathscr{A}).$$

Vanishing cohomology $[\delta\sigma] = 0$ corresponds to a globally consistent history: the local admissibility conditions are compatible and a global section exists. Non-trivial cohomology $[\delta\sigma] \neq 0$ encodes an irreducible inconsistency: the local conditions are mutually incompatible in a way that cannot be resolved by refining the cover.

In the climate domain, this obstruction structure has a physical interpretation. A non-trivial cohomology class would correspond to a climate state in which the locally admissible refinements at different scales are globally inconsistent — a situation that the physical system cannot actually reach. The FSA architecture avoids such states by construction, because its field-space operators enforce cross-scale compatibility. But it does so without explicitly computing or checking the cohomology class.

The RSVP/TARTAN framework provides a principled path toward explicit obstruction detection: an architecture equipped with a cohomological constraint checker could identify globally inadmissible states before generating them, rather than relying on the training distribution to have excluded them.

7.4. *Holonomy and Path Dependence*

A deeper sheaf-theoretic feature of the RSVP framework is the holonomy of the constraint connection. If the constraint field Ψ_t is treated as a connection on the admissibility bundle over $\mathfrak{M}$, then the holonomy of a closed loop γ in the model space is

$$\mathrm{Hol}_\gamma(\Psi) = \mathcal{P}\exp\Big(\oint_\gamma A_\Psi \, d\ell\Big) \in \mathrm{Aut}(\Omega_t),$$

where A_Ψ is the connection form associated with Ψ. Non-trivial holonomy implies that the admissibility structure is path-dependent: the set of admissible transitions depends not only on the current state but on the history of states through which the system has passed.

This is the formal expression of the intuition, central to RSVP, that history matters intrinsically rather than merely instrumentally. The FSA architecture implicitly exhibits path dependence through its temporal conditioning — the compressed state at time t depends on the sequence of compressed states that preceded it, not merely on the last state. But it lacks an explicit account of the holonomy structure that generates this dependence.

7.5. *Refinement Cohomology*

The zero-shot super-resolution capability of the FSA suggests a cohomological interpretation of the reconstruction process. The admissibility conditions at different scales define a cochain complex:

$$\cdots \to C^{k-1}(\mathfrak{U}, \mathscr{A}) \xrightarrow{\delta} C^k(\mathfrak{U}, \mathscr{A}) \xrightarrow{\delta} C^{k+1}(\mathfrak{U}, \mathscr{A}) \to \cdots$$

A super-resolution query at a novel resolution k^* corresponds to extending a cocycle from the scales $\{0, 1, \ldots, k^* - 1\}$ to a cocycle that includes k^*. The query is answerable if and only if the extension is cohomologically trivial: if the admissibility conditions at scales 0 through $k^* - 1$ uniquely determine an admissible refinement at scale k^*.

The FSA zero-shot super-resolution result is empirical evidence that the training process has learned a representation in which this extension is approximately trivial within the training distribution. The TARTAN framework predicts that explicitly enforcing the cohomological constraint during training would extend the zero-shot capability beyond the training distribution.

7.6. *Constraint Frustration and Obstruction Structure*

Non-trivial cohomology corresponds physically to *frustrated* constraint systems: collections of locally admissible conditions that cannot be simultaneously satisfied globally. Spin glasses, geometrically frustrated magnets, and incompatible boundary conditions in fluid dynamics all exhibit this structure. In each case, every local subsystem has a well-defined preferred state, but the local preferences are mutually incompatible across the full domain.

The significance of the sheaf-theoretic formalism is that it treats such failures not as numerical instability or optimization failure but as *topological obstruction*. A non-vanishing obstruction class $[\delta\sigma] \neq 0$ does not indicate that the optimizer has not converged or that more training data is required. It indicates that no globally admissible section exists for the local conditions supplied: the incompatibility is structural and cannot be resolved by any continuous local adjustment.

This provides a principled distinction between two failure modes that existing machine learning systems frequently conflate. An optimization failure is recoverable: with more capacity, more data, or a better objective, the system can approach the globally consistent solution. A topological obstruction is irrecoverable within the given constraint structure: the system must change its admissibility conditions (i.e., the physical model) rather than its optimization. Architectures lacking an explicit obstruction formalism cannot distinguish between these cases. They interpret topological failures as optimization failures, add capacity, and discover that the capacity has no effect — because the problem is not insufficient optimization but structural inadmissibility of the local conditions.

8. Spin–Valley Modes in Moiré Superlattices: Hidden Constraint Propagation

8.1. *The Physical Setting*

Moiré superlattices formed by stacking two atomically thin layers of transition metal dichalcogenides at a small twist angle constitute one of the richest experimental platforms in contemporary condensed matter physics. The periodic interference pattern between the two lattices creates an effective superlattice potential with a unit cell orders of magnitude larger than the atomic scale, and the resulting flat bands concentrate kinetic energy suppression in precisely the regime where interaction effects become dominant. Twisted WSe_2 bilayers, in particular, have recently been shown to host superconductivity [6, 7], correlated insulator phases, and topological states, making them a setting where multiple competing orders — excitonic, magnetic, superconducting — are simultaneously accessible.

The paper we treat as our second empirical reference [2] studies a distinct but related phenomenon: the propagation of collective spin–valley modes in a moiré-hosted spin–valley superfluid. The theoretical background was established by Bi and Fu [4], who predicted an excitonic density wave and spin–valley superfluid in bilayer transition metal dichalcogenides arising from interlayer Coulomb attraction between electrons in opposite spin–valley sectors. The key feature of this state is that it supports a charge-neutral order parameter: the condensed excitons carry no net charge, only correlated spin and valley quantum numbers.

8.2. *Hidden Does Not Mean Unphysical*

The language of hidden fields and hidden manifolds risks a misreading that must be addressed directly. Throughout this essay, "hidden" does not mean metaphysically inaccessible, unobservable in principle, or existing in a realm beyond the physical. The spin–valley modes are hidden only relative to a particular measurement basis: ordinary charge-sensitive probes. Once the measurement strategy changes — from static charge response to direct space-time imaging of spin–valley fluctuations — the previously hidden structure becomes fully and precisely empirically accessible. The Xiong et al. paper does exactly this: it makes the hidden field visible by choosing a non-degenerate observational basis.

The RSVP framework makes the same ontological claim. It does not posit an inaccessible substrate beyond observation. It argues that different observational projections preserve different subsets of the underlying admissibility structure. A projection may be operationally convenient — charge measurement is simpler than space-time spin–valley imaging — while still eliminating the dynamically relevant organizational degrees of freedom. The hidden field is hidden relative to a representational basis that was optimised for different observables, not hidden in any absolute or metaphysical sense.

This matters for interpreting the broader convergence thesis. The claim that "hidden dynamical manifolds are more physically fundamental than observable state variables" is not a mystical claim about hidden realities. It is a structural claim about which representational choices preserve the information necessary for long-horizon prediction and global constraint coherence. The manifest structure of any projection is real; the claim is that it is not *sufficient* for the physical purposes under discussion.

8.3. *Observable Surfaces and Hidden Dynamics*

The experimental and theoretical situation in twisted WSe_2 presents with unusual clarity a distinction that the RSVP framework treats as structurally fundamental: the distinction between the observable projection surface and the hidden constraint dynamics beneath it. Standard condensed matter measurement tools — transport, optics, scanning probes — are sensitive to charge. They record the charge density, the charge current, the charge compressibility. These are real physical quantities, and measuring them carefully yields genuine information about the system.

The difficulty is that the most important collective organization of the spin–valley superfluid state is charge-neutral. The order parameter carries no net charge; the Goldstone modes that propagate through the system carry no charge current; the Higgs fluctuations that stabilize the condensate produce no charge response. The observable surface is real, but it is not the physical system. The physical system includes both the observable projection and the hidden constraint geometry that the projection does not access.

This situation is formally identical to the RSVP decomposition $X_t = X_t^{\text{obs}} \oplus X_t^{\text{hid}}$ with

$\mathcal{O}(X_t^{\text{hid}}) = 0$. The observable projection $\mathcal{O}$ kills precisely the degrees of freedom that carry the system's collective organization. Recovering those degrees of freedom requires a measurement strategy that is sensitive to the hidden field — one that does not project onto the observable surface first and attempt reconstruction from that projection, but that accesses the constraint geometry directly.

8.4. *Collective Modes as Constraint Propagation*

The Goldstone modes of the spin–valley superfluid are not merely soft fluctuations associated with a broken symmetry. They are the physical realization of constraint propagation: the mechanism by which the local admissibility conditions imposed by the moiré superlattice potential propagate coherently across the sample and maintain long-range collective order.

In the language of RSVP, the moiré superlattice potential plays the role of the constraint field Ψ_t: it specifies, at each spatial location, which configurations of the spin–valley degrees of freedom are energetically admissible. The spin–valley superfluid state is the global section of the admissibility sheaf $\mathscr{A}$ over the moiré domain: the unique (up to gauge) configuration that is locally admissible everywhere simultaneously. The Goldstone modes are the gapless perturbations of this global section that cost no energy to the leading order, because they correspond to global rotations of the order parameter that respect all the local admissibility conditions uniformly. The Higgs mode is the perturbation that costs energy because it locally violates the amplitude constraint: it attempts to change $|\psi|$ from its condensate value, which the local admissibility condition resists.

The propagation velocity of the Goldstone mode is therefore not merely a material parameter but a measure of the stiffness of the global constraint structure: how strongly the local admissibility conditions resist relative phase fluctuations between spatially separated regions of the sample. A high propagation velocity indicates a stiff constraint structure, one where the local admissibility conditions are tightly coupled across space. A low velocity indicates a soft constraint structure, where local admissibility conditions are relatively independent.

8.5. *The Measurement Problem and Its Resolution*

The central experimental challenge is that the spin–valley modes are charge-decoupled. Ordinary transport measurements, which detect charge currents, are insensitive to the propagating excitations because the condensed excitons carry no net charge. Optical measurements that probe charge density are similarly blind. Standard steady-state measurements show only the total response, in which charge-coupled excitations dominate and the charge-neutral spin–valley modes are eclipsed entirely.

The experimental resolution is a pump–probe space-time imaging technique. A shaped pump laser injects a localized packet of excitations at time $t = 0$; a delayed wide-field probe then tracks their spatial evolution as a function of time delay Δt, recording the full space-time trajectory rather than a static integrated response. Near the van Hove singularity — a region of the moiré phase diagram with particularly high density of electronic states — the technique resolves two propagating collective modes. The modes carry *opposite* spin–valley currents and propagate at drastically different velocities: one faster than 3 km s^{-1} with partially ballistic behaviour consistent with a gapless excitation, the other slower and diffusive, consistent with a gapped excitation. The fast, ballistic mode is interpreted as the Goldstone mode of the intervalley coherent (IVC) ground state; the slow, diffusive mode is interpreted as the amplitude (Higgs) mode with an energy gap. The authors describe their mission as developing the capacity to directly record the "complete space-time evolution" of many-body quantum systems — language that is almost exactly the RSVP prescription for trajectory recovery rather than state-snapshot measurement.

8.6. RSVP *Interpretation of Phase and Amplitude Modes*

The Goldstone–Higgs mode decomposition in the spin–valley superfluid maps naturally onto the RSVP field decomposition $X_t = (\varphi_t, \mathbf{v}_t, S_t)$. The phase (Goldstone) modes correspond to the vector flow $\mathbf{v}_t$: they govern coherent directional propagation through the constraint manifold, are massless (gapless), and carry long-range organizational information about the system's collective state. The amplitude (Higgs-like) modes correspond to fluctuations of the scalar density φ_t: they are energetically costly, localized in effect, and reflect the local cost of deviating from the condensate density.

More formally, let $\psi = \sqrt{\rho}\, e^{i\theta}$ denote the spin–valley order parameter in polar decomposition, where ρ is the condensate density and θ is the phase. The effective action in the vicinity of the condensate has the form

$$\mathcal{L} \approx \tfrac{1}{2}(\partial_\mu \theta)^2 + \tfrac{1}{2} m_H^2 (\delta\rho)^2 + \cdots ,$$

where the first term governs gapless phase propagation and the second governs gapped amplitude fluctuations with Higgs mass $m_H \propto \sqrt{\rho_0}$. In RSVP notation:

$$\text{Phase modes:} \quad \mathbf{v}_t \sim \nabla\theta, \quad \partial_t \theta = -\nabla\varphi + \mathcal{T}(\mathbf{v}_t), \tag{4}$$

$$\text{Amplitude modes:} \quad \delta\varphi_t \sim \delta\rho, \quad \partial_t(\delta\varphi) = -m_H^2 \delta\varphi + D_\varphi \nabla^2(\delta\varphi). \tag{5}$$

The phase modes preserve coherence and carry the system's organizational information across space; the amplitude modes stabilize the local condensate density against perturbative fluctuations. This is precisely the RSVP distinction between low-entropy coherence-preserving flows (which govern accessibility geometry) and local energetic stabilization perturbations (which regulate the scalar field against collapse).

8.7. *Hidden Modes and Observable Projections*

The most theoretically significant feature of the spin–valley experiment is the explicit demonstration that the dynamically relevant structure of the system is not accessible through its standard observable projection. This is the condensed matter version of the RSVP critique of purely observational epistemology, and it is worth stating precisely.

Let $\mathcal{O}$ denote the observable projection: the map from the physical state X_t to the set of measurable quantities (charge density, conductance, optical response). The charge-decoupled spin–valley modes lie in the kernel of $\mathcal{O}$: they contribute nothing to the observable signal. The physically significant dynamics — the propagating Goldstone modes that organize the collective state — are, in this sense, invisible. They belong to the hidden part of the field,

$$X_t = X_t^{\text{obs}} \oplus X_t^{\text{hid}}, \quad \mathcal{O}(X_t^{\text{hid}}) = 0.$$

This decomposition is not merely a technical accident of the measurement apparatus. It reflects a structural feature of the physics: the symmetry that protects the spin–valley modes from charge hybridization is the same symmetry that makes them dynamically significant. The hidden modes are hidden precisely because they are charge-neutral; and they are the carriers of collective order precisely for the same reason.

In RSVP terms, the observable projection $\mathcal{O}$ plays the role of a lossy compression that eliminates admissibility-relevant structure from the representation. The standard condensed matter toolkit — transport, optics — operates exclusively on X_t^{obs} and is therefore blind to the constraint dynamics of X_t^{hid}. Recovering those dynamics requires a measurement strategy that is sensitive to the hidden field: direct space-time imaging of the spin–valley fluctuations rather than inference from charge observables.

This is the physical analogue of the CLIO prescription for constraint-leveraged inference: do not project the constraint structure onto a lower-dimensional observable manifold and attempt reconstruction from that projection. Instead, develop measurement or inference strategies that are sensitive to the admissibility structure directly.

8.8. *Space-Time Imaging as Trajectory Recovery*

The experimental technique developed in the moiré paper is a form of trajectory recovery. Rather than recording a static snapshot of the collective state, the authors record its complete space-time evolution: how the spin–valley fluctuations propagate across the sample as a function of both position and time. The result is a direct observation of the trajectories of the collective modes rather than merely their instantaneous configuration.

This maps onto the CLIO argument about Chain-of-Memory versus static symbolic traces with unusual precision. The authors are not asking "what is the charge density at time t?" — the static observable projection. They are asking "how does the spin–valley order parameter evolve from t_0 to t?" — the trajectory in the hidden field space. The physically meaningful answer is the trajectory, not the snapshot.

The RSVP formalization of this distinction is the provenance structure $\mathcal{P}_t = \mathcal{P}(\omega_t)$. Provenance records the history of admissible transitions that produced the current state: it is the trajectory, not the state, that carries the organizational information. A system whose collective modes propagate coherently over macroscopic distances is a system whose provenance trace extends over those distances; the coherence is a property of the trajectory rather than of any instantaneous configuration.

8.9. *Emergent Macroscopic Organization from Local Constraints*

A further feature of the spin–valley experiment resonates with the RSVP/TARTAN framework: the emergence of macroscopic collective organization from purely local interaction constraints. The moiré potential provides a periodic local constraint — the superlattice — that frustrates the electrons into a correlated ground state. The spin–valley superfluid emerges not from any single interaction but from the global consistency of local admissibility conditions across the entire sample.

This is the physical realization of the sheaf-theoretic structure developed in Section 7. The local admissibility condition at each moiré unit cell is the constraint that interlayer excitons occupy the energetically favoured spin–valley sectors. The global section is the ordered condensate state. The Goldstone modes are the gapless fluctuations of the global section; the Higgs mode is the cost of a section that violates the local amplitude constraint.

The experimental observation that these modes propagate coherently across the sample — rather than decaying into local incoherence — is evidence that the sheaf gluing condition is satisfied: the local admissibility conditions are globally consistent, and a global section exists. The absence of a propagating Goldstone mode would correspond to a non-trivial Čech cohomology class: the local constraints would be mutually incompatible in a way that prevents the formation of long-range collective order.

9. Novel Theoretical Conclusions

The formal correspondences developed in the preceding sections are not merely analogical. When pressed, they yield conclusions that are more specific than the general thesis of convergent rediscovery, and that have implications extending beyond either empirical paper. This section states five such conclusions explicitly.

9.1. *Compression and Measurement as Formal Duals*

The FSA paper studies a compression problem: reduce a high-dimensional physical field to a compact latent representation while preserving admissibility structure. The moiré experiment studies a measurement problem: reconstruct hidden collective dynamics from an observable projection surface. These are dual operations in a precise sense.

Let $\pi : X \to M$ denote the compression map and $\mathcal{I} : M_{\text{obs}} \to X_{\text{hid}}$ denote the reconstruction map from observables to hidden dynamics. The two papers demonstrate, from opposite directions, that both operations must respect the same geometric constraint: the admissibility structure of the physical system. The FSA shows that compression fails when it violates the geometry of the underlying process. The moiré experiment shows that measurement fails when it projects onto a surface that eliminates the physically significant degrees of freedom.

The formal duality is:

$$X \xrightarrow{\pi} M \qquad \text{(compression: preserve admissibility downward)}$$

$$M_{\text{obs}} \xrightarrow{\mathcal{I}} X_{\text{hid}} \qquad \text{(measurement: recover admissibility upward)}$$

Both operations succeed when they respect constraint geometry and fail when they do not. This suggests a unified principle: *effective measurement and effective compression obey the same admissibility-preserving geometric constraints*. The design of good scientific instruments, the design of good latent representations, and the design of good inference systems may all reduce to the same underlying problem.

Proposition 9.1 (Compression–Measurement Duality)**.** *Let X be a field space with admissibility structure $\mathcal{A}_{\text{dm}}(X)$. A compression $\pi : X \to M$ is admissibility-preserving iff the reconstruction map $\mathcal{I} = \pi^{-1}|_{\mathcal{A}_{\text{dm}}}$ exists and maps M back into $\mathcal{A}_{\text{dm}}(X)$. The existence of $\mathcal{I}$ is equivalent to π being injective on $\mathcal{A}_{\text{dm}}(X)$: the admissible set is not collapsed by the compression.*

The proposition states that a compression is good — in the admissibility-preserving sense — if and only if it supports a good reconstruction of the admissible region. The two operations are not merely analogous; they are inverses, and their quality is determined by the same geometric property of the map π.

9.2. *Hidden Dynamical Manifolds as Primary*

Both papers provide empirical evidence for a claim that is stronger than the convergence thesis: that hidden dynamical manifolds may be more physically fundamental than the observable state variables that standard theory and measurement treat as primary.

The moiré experiment demonstrates this in the clearest possible form. The physically significant organization of the spin–valley superfluid — the collective order, the propagating modes, the broken symmetry that defines the phase — exists entirely in the hidden field X_{hid}. The observable surface X_{obs} carries none of it. An account of the system that begins with observable state variables and attempts to derive the collective organization from them is not merely incomplete; it is organized around the wrong primary objects.

The FSA result makes the same point computationally. The Field-Space architecture, trained only on residuals and never exposed to the coarse seasonal and warming signal, spontaneously recovers the full dynamical geometry of the climate system in its latent space. The convolution baseline, trained on the full field including the large-scale climatological structure, loses it. The organized latent geometry emerges from the constraint structure of the residuals, not from the surface statistics of the observable field.

Together these results support the RSVP ontological claim:

$$X_{\text{obs}} = \Pi(X_{\text{constraint}}),$$

where the observable field is a secondary projection of the hidden constraint manifold rather than a primary object. This is structurally close to the programme of generalized effective field theory, where macroscopic observables are understood as integrations-out of microscopic degrees of freedom. The difference is that in the RSVP account, the hidden degrees of freedom are not microscopic in the conventional sense but *constraint-geometric*: they are the admissibility conditions that govern which trajectories are physically realizable, rather than high-energy modes that happen to be integrated out.

9.3. *Deep Learning as Implicit Field Theory*

The FSA architecture, examined from the perspective of what it actually computes, is structurally much closer to discretized field evolution than to classical symbolic AI or to the lookup-table intuition that motivates most large-scale ML. It preserves topology by operating on the HEALPix sphere. It evolves latent states continuously through temporally structured diffusion. It transports compatibility constraints across scale boundaries rather than mapping surface statistics. It encodes trajectory organization rather than point-to-point correspondences.

These are precisely the structural commitments of a numerical field theory. The architecture's success — its superior compression efficiency, its organized latent space, its zero-shot generalization — depends directly on respecting geometry, preserving continuity, and maintaining cross-scale constraint consistency: the properties that define well-posed field theories. The architecture succeeds to the extent that it behaves like a dynamical field, and fails to the extent that it behaves like a lookup system.

A more precise formulation: under optimization pressure from physically constrained domains, deep learning architectures increasingly become discretized field theories without explicitly recognizing or naming this transition. The FSA is a data point in what appears to be a systematic empirical trend. GraphCast [10], Pangu-Weather [11], FourCastNet [12] — the leading neural weather prediction systems — all converge on architectures that preserve spherical geometry, propagate information continuously across spatial scales, and maintain temporal coherence through structured latent dynamics. The transition from storage-first to field-first representation is being driven by physics, not by theoretical preference.

9.4. *Irreversibility as Geometric Contraction*

Both papers contain evidence that the physically important object is not instantaneous state but constrained historical trajectory, and that irreversibility is not primarily a thermodynamic property but a geometric one.

The FSA latent space visualization is particularly revealing. Encoded daily climate states from 1940 to 2024 drift smoothly and *directionally* through the latent manifold when coloured by year, with earlier states concentrated in one region and more recent states displaced along an outer arc. This drift is not a random walk; it has a consistent direction in latent space that mirrors the directional character of global warming. The trajectory through the latent manifold is irreversible not because entropy increases in the thermodynamic sense, but because the accessible region of the latent manifold at late times is a proper subset of the region accessible at early times: certain latent states that were reachable in 1940 have become unreachable by 2024, and not merely because they are unlikely but because the constraint structure of the system has shifted.

This is the RSVP account of temporal asymmetry:

$$\mathcal{A}_{t+1}^{\text{consistent}} \subsetneq \mathcal{A}_{t}^{\text{counterfactual}},$$

stated in empirical rather than formal terms. The latent manifold contracts under historical constraint accumulation. The contraction is geometric — a reduction in the volume of the accessible region — rather than thermodynamic. Irreversibility is not the growth of disorder but the shrinkage of possibility.

The moiré experiment provides a physical analogue: the Goldstone modes propagate directionally, carrying spin–valley current in a specific direction across the sample. The propagation is not time-reversible in the experimental sense: exciting a localized packet and watching it spread produces a specific asymmetric space-time evolution. The constraint geometry of the IVC state selects a direction of propagation; the trajectory through the collective mode space is path-dependent in exactly the sense the RSVP framework predicts.

9.5. *A Generalization Theorem*

The FSA zero-shot capability across both resolution domains (HEALPix levels 6–8) and model domains (ERA5 to MPI-ESM at 100 km resolution) supports a general theoretical claim that is stronger than the observation of a single benchmark result.

Theorem 9.2 (Generalization from Constraint Geometry)**.** *Let $\pi : X \to M$ be an admissibility-preserving compression and let C denote the constraint compatibility structure learned during training on distribution $\mathcal{D}$. Then the system generalizes zero-shot to any input $x \notin \mathcal{D}$ satisfying the same admissibility conditions as $\mathcal{D}$, in the sense that the reconstruction $\pi^{-1}(\pi(x))$ lands in $\mathcal{A}_{\text{dm}}(X)$.*

Sketch. If π is admissibility-preserving, then π maps $\mathcal{A}_{\text{dm}}(X)$ injectively into M. The learned compatibility structure $\mathcal{C}$ is a representation of this map restricted to the training distribution $\mathcal{D} \cap \mathcal{A}_{\text{dm}}(X)$. For any $x \in \mathcal{A}_{\text{dm}}(X)$ (regardless of whether $x \in \mathcal{D}$), the admissibility conditions satisfied by x are the same as those encoded by $\mathcal{C}$, since admissibility is a property of the physical system rather than of the training distribution. Therefore $\pi^{-1}(\pi(x)) \in \mathcal{A}_{\text{dm}}(X)$, and the reconstruction is admissible. Distributional proximity is irrelevant; admissibility membership is sufficient. □

The theorem makes precise what the FSA zero-shot results demonstrate empirically: it is the admissibility conditions, not the training distribution, that determine the scope of generalization. A system that learns constraint compatibility generalizes wherever those constraints apply. A system that learns distributional statistics generalizes only where those statistics hold. The practical consequence is that geometry-aware, constraint-preserving architectures are qualitatively more robust than distribution-matching architectures at any finite scale.

This reframes the scaling debate directly. The relevant question is not how many parameters or how much data but whether the system has learned the constraint structure of the domain. A small system that has learned admissibility geometry may generalize more robustly than a large system that has memorized distributional surface statistics. This is a strong prediction that the current convergence of architectural pressure toward field-theoretic design makes increasingly testable.

9.6. *Trajectory Coherence Precedes Semantic Interpretation*

A subtler but important observation emerges from both papers: coherent dynamical organization appears to be detectable before and independently of semantic interpretation. In the FSA, the latent space organizes into seasonally structured cyclic trajectories and historically directional drift without any explicit semantic labelling. The geometric structure emerges from the architecture's constraint commitments alone, prior to any downstream task that would require semantic interpretation of the compressed state.

In the moiré experiment, the coherent space-time propagation of the collective modes is observed and characterized — velocity, dispersion relation, directionality, ballistic versus diffusive behaviour — before the exact nature of the IVC ground state is identified. The Goldstone mode is recognizable as coherent propagation long before the question of whether the underlying state is an incommensurate Kekulé spiral order or some other IVC variant is resolved. Coherence geometry is a more primitive experimental observable than phase classification.

This convergence suggests a broader epistemic principle: coherence detection may be prior to semantic classification in the order of scientific knowledge. Dynamical systems organize themselves coherently before human observers interpret what the coherence means. The implication for machine learning is that training coherence — trajectory consistency across temporal and spatial scales — may be a more fundamental objective than semantic correctness, and that systems evaluated primarily on semantic endpoint accuracy may systematically miss the coherence structure that makes correct endpoints possible at long horizons.

9.7. *Trajectory Coherence Under Compression*

The following theorem makes precise the sense in which admissibility-preserving compression protects trajectory coherence, and characterises what happens when it fails.

Theorem 9.3 (Trajectory Coherence Preservation). *Let $\pi : X \to M$ be a compression map that preserves admissibility and cross-scale compatibility. If $\gamma : [t_1, t_2] \to X$ is a globally admissible trajectory in X, then the projected trajectory $\pi \circ \gamma$ preserves coherence under temporal composition:*

$$(\pi \circ \gamma)(t_2) \circ (\pi \circ \gamma)(t_1) = \pi\big(\gamma(t_2) \circ \gamma(t_1)\big),$$

up to bounded reconstruction error determined by the local entropy density S_t. Conversely, if π fails to preserve admissibility structure, temporal composition becomes non-associative under projection, producing cumulative trajectory drift and eventual representational delamination*: the projected trajectory separates from the physical admissibility manifold and cannot be returned to it by any sequence of admissible local corrections.*

Sketch. Admissibility-preservation implies that π is injective on $\mathcal{A}_{\mathrm{dm}}(X)$, so the composition law in X transfers to M: $\pi(\gamma(t_2) \circ \gamma(t_1)) = (\pi \circ \gamma)(t_2) \circ (\pi \circ \gamma)(t_1)$ with error bounded by the reconstruction error of π, which is controlled by S_t (Proposition 4.1). If π is not admissibility-preserving, π collapses some admissible transitions: there exist $\omega_1 \neq \omega_2 \in \mathcal{A}_{\mathrm{dm}}(X)$ with $\pi(\omega_1) = \pi(\omega_2)$. Temporal composition applied to such collapsed transitions produces a sequence of projected states that satisfies the projected composition law but not the original one; the error accumulates over the trajectory and is not bounded by any local quantity. Topological delamination follows once the accumulated error exceeds the admissibility threshold κ_*. □

The phrase *representational delamination* deserves emphasis. It names the specific failure mode that distinguishes structurally admissibility-violating systems from merely noisy ones. A noisy system produces small random errors that may be correctable locally. A delaminated system has left the physical admissibility manifold entirely; its trajectory in representation space no longer corresponds to any trajectory in the physical system's state space, and no amount of local correction can restore the correspondence. Climate trajectory drift that exits the physical climate manifold, latent space collapse that loses the seasonal trajectory structure, semantic incoherence that accumulates over long text generation, and hallucination that produces locally plausible but globally inadmissible outputs are all instances of representational delamination. The FSA architecture resists delamination by construction; the convolutional baseline exhibits the precursor to it in its zero-shot resolution experiments.

9.8. *Constraint Curvature and Historical Drift*

The latent warming drift observed in the FSA representations is not merely a visualisation artifact. It suggests that the admissibility manifold possesses an intrinsic curvature induced by historical constraint accumulation. The climate trajectory does not translate uniformly through latent space; it follows a curved path whose geometry is history-dependent in a way that distinguishes accessible future states from counterfactually accessible past states.

In the RSVP framework this curvature arises from the progressive contraction of the accessibility structure:

$$\mathcal{A}_{t+1}^{\text{consistent}} \subsetneq \mathcal{A}_{t}^{\text{counterfactual}}.$$

The resulting geometry is non-Euclidean in a precise sense. Parallel transport of trajectories through the admissibility manifold produces path-dependent displacement analogous to holonomy in curved connections. Two trajectories that begin arbitrarily close may diverge permanently if their histories induce different admissibility contractions — the topological manifestation of irreversibility discussed in §9.

Constraint curvature therefore provides a unified geometric interpretation of four phenomena that otherwise appear as separate empirical observations: the directional warming drift in the FSA latent space (curvature in the temporal direction of the accessibility manifold); the path-dependent propagation of the spin–valley Goldstone modes (curvature in the spatial direction of the constraint connection); the non-vanishing holonomy of the sheaf formalism in §7 (topological curvature of the admissibility bundle); and the long-horizon coherence failures of systems that ignore constraint geometry (geodesic deviation in a curved manifold when a flat approximation is assumed). The latent drift is not a visualisation curiosity. It is a geometric record of the curvature that the physical system has accumulated.

9.9. *The Structural Pressure Hypothesis*

The formal correspondences established in the preceding sections are not, by themselves, evidence that the RSVP/CLIO/TARTAN framework is correct. What makes them significant is their character and their independence. The FSA paper was designed to solve a specific engineering problem: scalable climate emulation. The spin–valley moiré paper was designed to solve a specific measurement problem: directly imaging hidden collective dynamics in a correlated quantum material. The RSVP/CLIO/TARTAN framework was developed to address foundational questions in the philosophy of physics and computation. All three arrived at overlapping structural commitments through entirely different channels.

The FSA demonstrates computationally that effective compression requires geometry-aware, constraint-preserving, hierarchically explicit architecture. The moiré experiment demonstrates physically that effective understanding of collective dynamics requires sensitivity to hidden constraint manifolds that are invisible to standard observable projections. Together they constitute two independent instances of the same structural principle: the dynamically significant organization of a physical system is carried by its constraint geometry, and observable outputs are secondary projections of that geometry rather than primary objects of study.

We call this the *structural pressure hypothesis*: geometry-aware, constraint-preserving, trajectory-sensitive representations emerge independently wherever the physical demands of the target domain make observable-surface approaches fail. The FSA paper is one instance; the moiré experiment is another; and similar pressures are visible in developments in protein structure prediction [10, 11], neural operator theory, and long-horizon planning in reinforcement learning. The convergence is not coincidental. It reflects a structural feature of physical reality: the dynamically significant organization of a system is carried by its constraint geometry, not by its instantaneous observable configuration.

9.10. *What the Convergence Does Not Imply*

It is equally important to be explicit about what the convergence does not imply. It does not imply that the FSA architecture is a complete engineering implementation of the RSVP programme, or that the spin–valley experiment constitutes a laboratory demonstration of RSVP field dynamics. Both empirical papers operate within their own conceptual frameworks and make no reference to the theoretical structures developed here. The correspondences are structural, not definitional.

More importantly, the convergence does not imply that the engineering and experimental approaches are approaching a theoretical limit that the formal framework has already reached. Both the theoretical and empirical traditions are works in progress. The theoretical framework identifies structural features that the empirical approaches lack — explicit admissibility formalisms, provenance structures, cohomological constraint checking, holonomy-aware path conditioning. The empirical approaches provide constraints and exemplars that the theoretical framework has not yet been tested against in those specific physical regimes. The relationship is one of productive tension rather than one-sided validation.

9.11. *Long-Horizon Coherence as the Common Constraint*

The structural pressure hypothesis identifies a specific common constraint that unifies the climate system, the spin–valley moiré experiment, and the broader domains of cognition, reasoning, and semantic memory: *local correctness is insufficient for global admissibility*. This is the deep unifying principle beneath the essay's entire argument, and it deserves explicit statement as a subsection rather than a remark distributed across the text.

In the climate system, a model that is locally accurate — that predicts the correct atmospheric state at each individual time step — can nonetheless produce a globally inadmissible trajectory, one that drifts to climatologically impossible states over long integration horizons. The local correctness of each step provides no guarantee that the global trajectory remains on the physical manifold. Enforcing global admissibility requires a representational structure that is sensitive to the full constraint geometry of the system's possible trajectories, not merely to the local statistics of adjacent states.

In the spin–valley system, a measurement strategy that is locally accurate — that correctly records the charge density at each spatial point — can nonetheless be entirely blind to the globally organized collective modes that carry the system's physical order. The local measurement correctness provides no access to the global constraint structure that the Goldstone modes propagate. Accessing that structure requires a measurement approach that is sensitive to the system's trajectory in hidden field space, not merely to its projection onto the observable surface.

In cognition and language, a generative system that is locally accurate — that produces grammatically and semantically plausible continuations at each step — can nonetheless produce globally incoherent outputs: texts that are locally reasonable but collectively contradictory, plans that are individually achievable but jointly infeasible, arguments that are locally valid but globally fallacious. Enforcing global coherence requires precisely the trajectory-conditioning and provenance-tracking mechanisms that the RSVP/CLIO framework formalizes.

The structural unity of these cases is not coincidental. In each, the system has a long-horizon constraint structure that is not reducible to any finite set of local constraints. Short-horizon optimization fails not because the local objectives are wrong but because global admissibility is a *path-dependent* condition: whether a current state is admissible depends on the entire history of transitions that produced it, not merely on its current configuration. This is the RSVP account of accessibility shrinkage made precise as an engineering and experimental constraint. Systems that do not maintain provenance-sensitive access to the full history of admissible transitions cannot enforce global admissibility, regardless of their local accuracy.

9.12. *The Remaining Conceptual Gap*

The conceptual gap that separates the current engineering approaches from the full theoretical framework can be stated precisely. Existing architectures, including the FSA, achieve operational admissibility preservation: their outputs tend to satisfy physical admissibility conditions because the training data is organized around such conditions. They do not achieve structural admissibility preservation: they lack explicit formal mechanisms for enforcing admissibility in the event of distribution shift, novel boundary conditions, or compositionally novel inputs.

The distinction matters because it determines generalization. An architecture that achieves only operational admissibility preservation will generalize within the training distribution and fail at its boundary. An architecture that achieves structural admissibility preservation will generalize to any input that satisfies the same admissibility conditions as the training data, regardless of distributional proximity. The FSA's zero-shot super-resolution capability is a partial demonstration of structural generalization, which is why it is theoretically significant. But it remains partial: the architecture has learned constraint compatibility within a specific domain and a specific distribution, not the general admissibility structure of the underlying physics.

The full theoretical framework provides a research direction for bridging this gap. Architectures equipped with explicit admissibility formalisms, sheaf-theoretic constraint checkers, and holonomy-aware trajectory conditioning would, in principle, achieve structural rather than merely operational admissibility preservation. Building and evaluating such architectures is a concrete research programme that the FSA paper helps to motivate.

9.13. *Why Endpoint Metrics Systematically Mislead*

Most contemporary evaluation protocols assess endpoint states rather than admissible trajectories. Reconstruction RMSE, next-token prediction loss, classification accuracy, and benchmark pass rates all compare final outputs against target outputs while remaining largely insensitive to the constraint geometry that produced them. These metrics are degenerate observables in the sense of §3: they collapse structurally distinct systems into the same observable equivalence class.

The FSA paper demonstrates this directly and precisely. The HEALPix convolution baseline achieves competitive RMSE at low compression ratios — it is not a poor model by the standard metric. Yet its latent space is disorganised, its zero-shot generalisation fails, and its parameter scaling exhibits instability (validation loss diverging as parameters increase beyond 36M). The Field-Space architecture achieves lower RMSE at higher compression while simultaneously producing organised latent trajectories and stable parameter scaling. Both dimensions of improvement are missed by endpoint RMSE alone.

The asymmetry is structural: endpoint metrics reward local correctness without penalising global inadmissibility. A model that accurately reconstructs each individual state while violating the dynamical geometry of the trajectory between states scores well. A model that preserves trajectory geometry while differing superficially from the training distribution may score poorly. Benchmark culture therefore systematically selects for representational systems that optimise projection-surface fidelity at the expense of hidden constraint coherence.

The appropriate evaluation metric for an admissibility-preserving system is not endpoint RMSE but the quality of the learned constraint structure, measured by properties such as latent space organisation, zero-shot generalisation to novel admissibility conditions, and long-horizon trajectory coherence. These metrics are harder to compute but correspond more directly to what the physical system actually requires.

9.14. *Broader Implications for Representation Theory*

The convergence also has implications beyond the specific domains of climate modelling and field theory. If the structural pressure hypothesis is correct, then the ongoing failure of large flat models to achieve long-horizon coherence in language, planning, and reasoning is not primarily a scaling problem but a structural one. The representational ontology of flat token sequences, like the representational ontology of flat climate grids, is incompatible with the geometry of the processes it is trying to model. Increasing the scale of a fundamentally mismatched representation does not resolve the mismatch; it may, in some cases, amplify it by increasing the number of degrees of freedom available for overfitting to surface regularities [24, 25].

A direct consequence concerns benchmark culture. Most evaluations assess endpoint accuracy — reconstruction RMSE, classification accuracy, next-token prediction loss — without examining admissibility-preserving trajectory organization. The FSA paper contains an implicit critique of this practice: the HEALPix convolution baseline achieves competitive RMSE at low compression ratios, yet its latent space is disorganized and its zero-shot generalization fails. A model can appear accurate while structurally violating the geometry of the underlying process. Surface metrics reward local correctness without penalizing global inadmissibility. The appropriate evaluation metric for an admissibility-preserving system is not endpoint RMSE but the quality of the learned constraint structure — measured by properties such as latent space organization, zero-shot generalization to novel admissibility conditions, and long-horizon trajectory coherence.

The geometry-first, constraint-preserving, hierarchically explicit approach suggests a different design and evaluation strategy: rather than increasing the depth and width of flat representations, increase the explicitness and precision of the constraint structure, and evaluate on the properties that depend on that structure. This is not a proposal to abandon statistical learning. It is a proposal to provide statistical learning with a representational substrate that is geometrically compatible with the target domain, so that what is learned reflects the admissibility structure of the domain rather than the distributional structure of the training corpus.

9.15. *Representation Collapse as a General Failure Mode*

The recurring failure identified across climate modelling, quantum measurement, and large-scale machine learning is not merely loss of accuracy but *collapse of representational fidelity to the underlying constraint geometry*. A representation collapses when it preserves local surface statistics while destroying the admissibility structure that organises the physical process globally. The collapse is not always visible in endpoint metrics because the metrics themselves are evaluated on the projection surface rather than on the hidden trajectory geometry of the process.

Representation collapse has a precise formal signature. A compressed representation

M has undergone constraint collapse relative to the original space X if

$$\dim\big(\mathcal{A}(M)\big) \ll \dim\big(\pi(\mathcal{A}(X))\big),$$

where $\mathcal{A}(M)$ is the accessibility structure of the compressed representation and $\pi(\mathcal{A}(X))$ is the image of the original accessibility structure under compression. The compression has destroyed the accessible future structure of the physical system: states that should be reachable in the compressed representation are not, and states that should not be reachable are.

This failure mode explains why systems can appear operationally successful under short-horizon evaluation while remaining structurally incapable of long-horizon coherence. Over short horizons, the collapsed accessibility structure is approximately consistent with the original; the errors are small enough to fall within acceptable metric tolerances. Over long horizons, the divergence between $\mathcal{A}(M)$ and $\pi(\mathcal{A}(X))$ accumulates, and the system eventually generates trajectories that are inadmissible in the original space. This is representational delamination in slow motion.

The FSA results provide a controlled empirical demonstration. The convolutional baseline at high compression ratios undergoes partial constraint collapse: it loses the organised accessibility structure visible in the Field-Space latent trajectories. The short-horizon RMSE metric does not detect this collapse until the extreme compression regime (1024×) where accuracy degrades sharply. The latent space organisation metric detects it at every compression ratio.

10. Physics Education Research as a Third Instantiation

A third domain provides unexpected confirmation of the structural pressure hypothesis from a direction entirely orthogonal to both climate emulation and condensed matter physics. A recent comparative study of active learning methods in introductory physics and astronomy courses [26] examined four established instructional approaches across 31 implementations at 28 institutions: Peer Instruction [27], ISLE [28], Tutorials [29], and SCALE-UP [30]. The study reported that all four methods improved short-term conceptual understanding as measured by concept inventories, with SCALE-UP exhibiting significantly larger gains. Its observational analysis suggested that time spent in genuine collaborative work, rather than in lecture-plus-clicker cycles, predicted larger conceptual gains.

This is a well-executed study within its design assumptions. It is precisely those assumptions that make the paper theoretically interesting in the present context. The study implicitly treats physics learning effectiveness as a scalar quantity measurable through short-term standardized instruments. The argument of this essay predicts, on structural grounds, that this compression should produce exactly the failure mode described in preceding sections: representational collapse of the underlying admissibility structure, invisibility of the collapse to the metric used, and systematic misleading of long-horizon inference from short-horizon evaluation.

10.1. *Concept Inventories as Degenerate Observables*

A concept inventory is a multiple-choice instrument designed to assess conceptual understanding of a particular domain, typically administered before and after an instructional intervention. Force Concept Inventories, conceptual surveys of electricity and magnetism, and similar instruments have been widely validated for reliability and discriminant validity within their intended short-term measurement context. The question at issue is not their reliability as instruments but their adequacy as observables of the underlying process they are meant to track.

In the formal vocabulary developed in §3, a concept inventory score is a degenerate observable of intellectual development. The projection

$$\mathcal{O}_{\text{inv}} : X_{\text{dev}} \to \mathbb{R}$$

maps the full developmental state X_{dev} of a learner — which includes conceptual structure, tacit knowledge [33], problem-solving heuristics, analogical repertoire, motivational trajectory, and the accumulated pattern of domain-crossing experience — onto a single scalar score. The map is degenerate in a structurally precise sense: it collapses dynamically distinct developmental configurations into the same observable value.

Two students who achieve identical pre-to-post gains on a force concept inventory may have undergone radically different developmental trajectories. One may have consolidated a robust, transferable conceptual schema with rich connections to adjacent domains. Another may have memorised the vocabulary sufficient to pass the instrument without acquiring the underlying structural understanding. Both map to the same observable, but their accessible futures — the set of admissible developmental trajectories available to them at the next stage — are entirely different. The degenerate observable $\mathcal{O}_{\text{inv}}$ cannot distinguish between them.

This degeneracy is not a correctable measurement error that better instrument design could eliminate. It is structural: the dimensionality of X_{dev} far exceeds the dimensionality of $\mathbb{R}$, and no scalar projection can avoid the collapse. The study by Sundstrom et al. is careful and rigorous within its design, but it is asking about the image of $\mathcal{O}_{\text{inv}}$ rather than about X_{dev} itself. The structural pressure hypothesis predicts that this will produce operationally useful but long-horizon-misleading conclusions: the gains are real, the measurement is coherent, and the inference to long-term scientific capability is structurally unwarranted.

10.2. *The Reflexivity Complication*

The educational domain compounds the degeneracy problem with a second structural feature absent from climate and condensed matter physics: reflexivity. Physical fields do not modify their behaviour in response to being measured. Climate systems do not read papers about climate modelling and adjust their dynamics accordingly. Spin–valley modes do not reorganise themselves in response to observational protocols.

Human learners do all of these things. Students respond to grading structures, institutional incentives, perceived performance standards, and the cultural meanings attached to different disciplines. A concept inventory is not merely an instrument applied to a stable system; it is an intervention that modifies the system it measures. Students who know they are being evaluated on conceptual understanding may allocate effort differently than those who are not. Instructors who know their effectiveness is being measured by pre-to-post gains on a specific instrument may align their teaching toward that instrument. The act of measurement modifies the accessibility structure of the measured system in ways that have no analogue in physical measurement.

This reflexivity problem is formally analogous to the difference between a constraint that is merely statistical and one that is structurally enforced. A statistical constraint can be violated by a sufficiently motivated agent; a structural constraint cannot be violated by any continuous adjustment. Human learning systems operate under statistical constraints that bend under institutional pressure. The result is Goodhart's Law [32]: when a measure becomes a target, it ceases to be a good measure. The concept inventory score, once used to evaluate instructional effectiveness, partially decouples from the construct it was designed to track, because the measurement itself becomes part of the constraint structure that shapes instructional decisions.

The RSVP framework handles reflexive systems by tracking the provenance of the accessibility structure: $\mathcal{P}_t = \mathcal{P}(\omega_t)$ records not merely the current state but the history of how the admissibility conditions themselves have evolved. A non-reflexive physical system has a stable admissibility structure that the measurement does not perturb. A reflexive social system has an admissibility structure that is partially constituted by the measurement protocol. Without explicit provenance tracking, an analysis of a reflexive system cannot distinguish between genuine improvement in the underlying process and improvement in the system's capacity to optimise for the measurement.

10.3. *Intellectual Development as Admissibility Geometry*

The most theoretically productive framing of the education paper's results is as evidence about accessible futures rather than about scalar quantities. What SCALE-UP produces, if the observational analysis is correct, is an instructional environment that expands the accessible developmental trajectory space more effectively than the alternatives. Students who spend more time in genuine collaborative work on complex problems develop a richer set of admissible subsequent trajectories: they have more ways to proceed, more analogical connections available, more problem-solving strategies in their repertoire. The concept inventory gain is a degenerate projection of this expanded accessibility; it is real evidence of the underlying change but not a measure of it.

This reframing dissolves an apparent tension in the literature. Studies occasionally find that active learning methods improve standardized test scores without producing equivalent improvements in other measures of scientific capability [31]. The RSVP interpretation explains this: the accessibility structure expanded by good instruction is multidimensional, and a scalar projection will detect expansion in some dimensions while being blind to others. Different instruments project different dimensions; the apparent inconsistency across instruments reflects genuine multidimensionality in the underlying developmental space rather than measurement error.

The implications for educational research design are directly analogous to the implications for computational and experimental design drawn from the FSA and moiré papers. Just as effective climate compression requires operators sensitive to the admissibility geometry of the atmospheric system rather than to pixel-level statistics, and just as effective measurement of spin–valley collective modes requires space-time imaging of hidden trajectories rather than static charge response, effective assessment of intellectual development requires instruments sensitive to the accessibility structure of the learner's developmental space rather than to the endpoint coordinates of a standardized projection.

Concretely: rather than measuring pre-to-post gains on a fixed instrument, assess the structure of the learner's analogical transfer capacity across novel domains, the coherence of their problem decomposition strategies over long time horizons, and the stability of their conceptual representations under perturbation by unusual examples. These are not merely harder measurements; they are measurements of a different object — the constraint geometry of the developmental space rather than a degenerate scalar projection of it. The analogy with the FSA's latent space organisation metric versus its RMSE metric is direct: the latter is easier to compute and partially informative; the former is more informative about the structural property that determines long-horizon performance.

11. Constraint Geometry and the Failure of Flat Epistemology

11.1. *Why Surface Statistics Become Misleading*

The recurring pattern across every empirical domain examined in this essay is that systems initially appear tractable through surface statistics alone, and then progressively fail as the temporal horizon, geometric complexity, or coupling structure of the domain increases. Flat representations work well in the weak-coupling regime because local statistics partially shadow the deeper admissibility geometry. They fail in the strong-coupling regime because the local statistics become progressively decoupled from the constraint structure that actually governs long-horizon evolution.

This transition can be formalised geometrically. Let X denote the full constraint space of the physical system and let $\pi : X \to M$ be a projection onto an observable manifold M. In weakly constrained systems, the fibers $\pi^{-1}(m)$ for $m \in M$ remain small relative to the scale of the dynamics. Observable states therefore approximately determine the underlying physical state, and surface statistics are informative. In strongly constrained systems, the fibers become large and structurally heterogeneous: many dynamically distinct trajectories project to the same observable state. The observable surface ceases to determine the physical trajectory.

The result is a characteristic epistemic illusion. The system appears predictable over short horizons because the observable projection still contains enough information to interpolate locally. Over longer horizons, prediction degrades because the hidden admissibility geometry dominates the dynamics while remaining invisible to the representation. This is the failure mode visible in long-context language generation, in climate trajectory drift, in hidden quantum collective modes, and in educational evaluation. Local correctness survives while global coherence collapses.

11.2. *Constraint Geometry as the Carrier of Persistence*

The deeper implication is that persistence itself is a geometric property. A system persists through time not because its local state is stable but because its admissibility structure constrains the space of possible perturbations tightly enough that coherent trajectories remain available.

This suggests a reinterpretation of stability theory. Traditional dynamical systems theory treats stability in terms of local perturbative return: a trajectory is stable if nearby trajectories converge back toward it under evolution. The present framework suggests a stronger criterion. A trajectory is structurally stable if perturbations preserve admissibility membership: $x + \delta x \in \mathcal{A}_{\mathrm{dm}}(X)$ for all perturbations δx within some admissible neighbourhood.

The distinction matters because a trajectory may be locally stable while globally inadmissible. Such trajectories are operationally persistent over short horizons but eventually undergo representational delamination from the physical manifold. This is exactly what occurs in systems optimised exclusively for local loss functions without explicit admissibility preservation. The hidden field is therefore not merely explanatory overhead. It is the carrier of persistence itself.

11.3. *Constraint Transport and the Nature of Generalisation*

The standard interpretation of generalisation in machine learning is that models interpolate statistical regularities beyond the training set. The evidence examined in this essay suggests a more precise interpretation: systems generalise robustly when they learn transport laws for admissibility constraints rather than surface correlations.

This can be stated formally. Let $\mathcal{C}(x)$ denote the constraint structure associated with a state x. A system generalises structurally if $\mathcal{C}(x_{\text{train}}) = \mathcal{C}(x_{\text{test}})$ implies successful reconstruction or prediction even when $x_{\text{test}} \notin \mathcal{D}_{\text{train}}$. This is exactly what the FSA transformer demonstrates in its transfer from ERA5 fields to MPI-ESM outputs: the surface statistics differ, the constraint structure does not, and the architecture transfers.

The implication reframes the scaling problem. Increasing parameter count without improving constraint transport primarily increases interpolation capacity inside the observable manifold. Improving admissibility preservation increases the volume of structurally reachable trajectories outside the training distribution. The former scales memory. The latter scales understanding.

11.4. *Scientific Revolutions as Representational Topology Change*

The argument developed throughout this essay provides a structural analogy — offered not as a causal explanation but as a formal parallel — for why scientific paradigms shift. A representational framework remains productive as long as its observable manifold remains approximately compatible with the admissibility structure of the domain. As empirical precision increases, the mismatch accumulates. Anomalies appear because trajectories that are admissible in the representation are inadmissible in the physical system. Eventually the representation becomes topologically unstable relative to the domain, and adding parameters, corrections, or local refinements no longer restores coherence. The representational substrate itself must change.

This structural description applies, with appropriate domain-specific differences, to the shift from Euclidean to relativistic spacetime, from local hidden-variable pictures to Hilbert-space quantum mechanics, from symbolic AI to geometry-aware latent dynamics, and from flat climate grids to field-space representations. In each case the observable surface eventually ceases to support the admissible trajectories required by the physics.

The structural parallel is this: scientific revolutions share with these engineering transitions the feature that they cannot be resolved by incremental local improvement. They require a change of representational topology — a change in the class of objects the framework treats as primary, not merely a change in the parameters of a fixed framework.

11.5. *Toward a Constraint-First Epistemology*

Traditional empiricism treats observables as primary and hidden structure as derivative inference. The repeated convergence examined in this essay suggests the reverse ordering may be more fundamental: constraint geometry is primary, while observables are projections that preserve only part of the underlying organizational structure.

Knowledge, under this view, is not fundamentally the accumulation of surface measurements. It is the progressive reconstruction of the hidden admissibility manifold from the traces left by observable projections. The aim of science changes subtly but decisively: the goal is not merely to predict observations but to recover the constraint geometry that makes those observations cohere over long horizons.

This explains why long-horizon coherence repeatedly emerges as the deepest criterion of successful representation. A representation that captures the constraint geometry continues to produce admissible trajectories under perturbation, extrapolation, and scale transfer. A representation that captures only surface regularities fails precisely when the system leaves the narrow manifold occupied by the training distribution. The difference between the two is the difference between prediction and understanding.

12. Future Directions

The convergence documented throughout this essay suggests several concrete research directions whose importance extends beyond the specific domains examined here.

Explicit admissibility-aware architectures. Existing geometry-aware systems such as the FSA preserve admissibility operationally but not formally. Architectures equipped with explicit constraint-checking operators, sheaf-theoretic compatibility verification, and trajectory-level coherence regularization would allow admissibility to be enforced structurally rather than statistically. Such systems would constitute a qualitatively different category of representation learner: not distribution-matching systems but admissibility-preserving systems, whose generalization failures correspond to genuine constraint violations rather than distributional novelty.

Observable degeneracy detection. Many scientific domains likely contain hidden dynamical manifolds analogous to the spin–valley modes in twisted WSe_2: physically organizing structures that standard observables systematically project away. A general theory of projection degeneracy would provide criteria for identifying when an observable basis is insufficient to recover the underlying constraint geometry. The formal tools are already present in §3 and the sheaf formalism of §7; what remains is their systematic application to domain selection.

Developmental and cognitive trajectory analysis. If educational and cognitive systems possess rich accessibility structures, then scalar endpoint evaluations systematically underestimate the geometry of learning. Longitudinal trajectory-sensitive metrics, analogical transfer topology, and provenance-aware developmental analysis may become necessary for understanding cognition in the same way that space-time imaging became necessary for understanding the spin–valley superfluid. The Sundstrom et al. study measures concept inventory gains; the analogous advancement would measure the structure of the learner's accessible trajectory space directly.

Latent geometry as an attractor under optimization. The spontaneous emergence of organised latent manifolds in FSA systems trained only for reconstruction suggests that certain geometric structures are attractors under optimization pressure whenever the underlying domain contains persistent admissibility constraints. Understanding this relationship would help explain why increasingly many high-performing systems converge toward field-like internal organizations even when not explicitly designed to do so, and would provide principled guidance for when such convergence should be expected to occur.

Unification across projection types. The present essay has treated physical compression, quantum measurement, and educational assessment as three instances of the same structural problem. If constraint geometry is the primary organizational substrate across these domains, then field dynamics, inference, memory, and measurement may be different projections of a common underlying structure. That possibility remains speculative. The structural convergence motivating it does not.

13. Conclusion: The Hidden Field is the Physical Field

Three empirical papers, from three entirely different domains, each independently discover that the dynamically significant organization of their respective systems is not carried by the obvious observable surface. The Field-Space Autoencoder [1] finds that climate compression fails when it optimises for pixel-level reconstruction fidelity, and succeeds when it preserves the admissibility geometry of the atmospheric field. The spin–valley moiré experiment [2] finds that the physically meaningful collective modes of the quantum material are charge-decoupled and invisible to standard observational bases, becoming accessible only through direct space-time imaging of hidden field trajectories. The physics education study [26] measures short-horizon gains on scalar concept inventories, instruments that are degenerate observables of the multidimensional accessibility structure of intellectual development, producing results that are real but structurally incomplete as accounts of long-horizon scientific capability.

The three cases form a progression from the most physical to the most reflexive, and the structural argument scales across all three. In the climate system, the hidden constraint geometry is the admissibility structure of the atmospheric field, which the FSA learns to compress without destroying. In the quantum material, the hidden constraint geometry is the spin–valley order parameter and its collective modes, which the pump–probe imaging makes directly accessible. In the educational system, the hidden constraint geometry is the multidimensional developmental accessibility structure of the learner, which no scalar instrument can capture and which reflexivity partially constitutes through the measurement process itself.

This progression reveals a gradient in the severity of the representational problem. Physical systems have stable admissibility structures that are in principle fully accessible to appropriate measurement strategies. The problem is choosing the right measurement basis. Human developmental systems have admissibility structures that are historically contingent, reflexively modified by measurement, and entangled with the entire biographical and institutional context of the learner's development. The problem is not merely choosing the right measurement basis but acknowledging that no finite set of standardized measurements can capture the full constraint geometry of a historically embedded developmental trajectory.

The convergence has three implications, at increasing levels of generality.

The first is methodological. Standard observational strategies across all three domains are systematically biased toward the observable projection surface. Climate models that optimize for pixel-accurate reconstruction, condensed matter measurements that record only charge response, and educational studies that assess short-term standardized gains are all measuring their respective projections rather than the underlying processes. The structural pressure hypothesis predicts that this bias will produce systematic failures at precisely the scales and regimes where the constraint geometry diverges most strongly from the observable projection. These failures are not evidence of insufficient data or insufficient parameters; they are evidence of a representational mismatch.

The second is theoretical. The RSVP framework provides a vocabulary in which all three findings can be stated with uniform precision. Constraint-preserving projection, admissibility-preserving compression, degenerate observables, accessible future structure, representational delamination, and topological obstruction apply equally to atmospheric dynamics, collective quantum modes, and intellectual development. The formal vocabulary transfers because the structural problem is the same: how to represent a system with rich constraint geometry using limited observational resources without destroying the geometric structure that makes long-horizon behaviour predictable.

The third is philosophical. The repeated independent discovery of constraint geometry as the carrier of physical and developmental organization — across physical, quantum, and human domains — suggests that this is not a feature of particular theoretical frameworks but a structural feature of reality as systems that persist and organize themselves. Physical systems organize along admissible trajectories through constraint manifolds. The macroscopic collective behaviour we observe is the projection of that hidden organization onto the observational surfaces we have developed to detect it. The surfaces are real, but they are not the primary objects. The primary objects are the constraint manifolds and the trajectories through them.

The geometry is not optional. It is what the physics is made of.

Acknowledgement. This work was developed as an independent research contribution without institutional affiliation. All formal frameworks — RSVP, CLIO, TARTAN — have been developed under the name Flyxion. The present essay situates them in relation to concurrent developments in scientific machine learning, condensed matter physics, and physics education research, none

of which was consulted in the development of the frameworks.

A. Category-Theoretic Formalization

The formal structures deployed throughout this essay can be given a unified category-theoretic interpretation. This appendix states the principal objects and morphisms explicitly, for readers who prefer the categorical perspective as a way of tracking the logical dependencies between the framework's components.

The Trajectory Category

Let **Traj** denote the category whose objects are admissible state spaces X equipped with an accessibility structure $\mathcal{A}_t$ and an admissibility function $\kappa : \Omega \times H \to \mathbb{R}$, and whose morphisms are admissibility-preserving maps:

$$f \in \mathrm{Hom}_{\mathbf{Traj}}(X, Y) \iff \omega \in \mathcal{A}_{\mathrm{dm}}(X) \Rightarrow f(\omega) \in \mathcal{A}_{\mathrm{dm}}(Y).$$

Composition is the usual composition of maps, which preserves admissibility by transitivity. The identity morphism on each object is the identity map, which trivially preserves admissibility.

The Representation Functor

Let **Rep** denote the category of representational manifolds: compact smooth manifolds M equipped with a metric and a latent trajectory structure. The projection functor

$$\Pi : \mathbf{Traj} \to \mathbf{Rep}$$

maps each admissible state space X to a compressed representational manifold $M = \Pi(X)$ and each admissibility-preserving morphism to a structure-preserving map between manifolds. The FSA compression operators $\mathcal{F}_{k\to j}$ are the component maps of Π at each scale stratum; the conditions they must satisfy to be admissibility-preserving are precisely the cross-scale compatibility conditions enforced by the field-space architecture.

A compression $\pi \in \mathrm{Hom}_{\mathbf{Rep}}(M, M')$ is *trajectory-preserving* if the following diagram commutes:

$$\begin{CD} X @>{\Pi}>> M \\ @V{f}VV @VV{\pi}V \\ Y @>>{\Pi}> M' \end{CD}$$

for every admissibility-preserving morphism $f : X \to Y$ in **Traj**.

The Verification Natural Transformation

The CLIO verification map $\mathcal{V} : \widehat{\mathcal{L}} \to \mathcal{S}_{\text{valid}} \cup \{\bot\}$ is a natural transformation between functors:

$$\mathcal{V} : \mathcal{R} \circ \Pi \Rightarrow \mathrm{Id}_{\mathcal{S}_{\text{valid}}} \cup \{\bot\},$$

where $\mathcal{R} : \mathbf{Rep} \to \widehat{\mathbf{Rep}}$ is the stochastic rendering functor. Naturality requires that verification commutes with admissibility-preserving morphisms: if $f : X \to Y$ in **Traj**, then verifying the rendering of $\Pi(X)$ is consistent with verifying the rendering of $\Pi(Y)$ after applying f.

The Accessibility Sheaf

The accessibility sheaf $\mathscr{A}$ is a sheaf of sets over the model space $\mathfrak{M}$, assigning to each open $U \subseteq \mathfrak{M}$ the set of admissible transitions within U. Restriction maps are the canonical inclusions. The gluing axiom requires that locally admissible transitions be globally consistent: for any compatible family of admissible transitions on an open cover $\{U_i\}$, there exists a unique globally admissible transition restricting to each local one.

The Čech cohomology $\check{H}^\bullet(\mathfrak{U}, \mathscr{A})$ measures the failure of global sections to exist. In degree 1, a non-trivial class $[\delta\sigma] \in \check{H}^1(\mathfrak{U}, \mathscr{A})$ corresponds to a globally incoherent transition system: the local admissibility conditions are mutually incompatible in a topologically irresolvable way. The full pipeline of §7 fits into the long exact sequence of cohomology associated to a short exact sequence of sheaves encoding the inclusion of admissible transitions into all transitions.

Power and Holonomy

The holonomy of the constraint connection Ψ_t around a closed loop γ in $\mathfrak{M}$ defines an element of $\mathrm{Aut}(\Omega_t)$:

$$\mathrm{Hol}_\gamma(\Psi) = \mathcal{P}\exp\Big(\oint_\gamma A_\Psi \, d\ell\Big) \in \mathrm{Aut}(\Omega_t).$$

Non-trivial holonomy implies that the admissibility structure is path-dependent: the set of admissible transitions at a point in $\mathfrak{M}$ depends on the history of the path used to reach that point. This is the categorical expression of the RSVP irreversibility thesis: provenance matters, and it cannot be recovered from instantaneous state.

The group of holonomies $\{\mathrm{Hol}_\gamma(\Psi) \mid \gamma \text{ closed loop in } \mathfrak{M}\}$ is the holonomy group of the constraint connection. Flat constraint regions (where $\mathfrak{R}_{ij} = 0$) have trivial holonomy group; the historical ordering of transitions is irrelevant. Non-flat regions have non-trivial holonomy; the historical ordering is dynamically significant. The FSA latent warming drift is evidence of non-trivial holonomy in the climate system's constraint connection over the accessible state space.

References

[1] Meuer, J., Witte, M., Plésiat, É., Ludwig, T. & Kadow, C. Field-space autoencoder for scalable climate emulators. *npj Artificial Intelligence* **2**, 50 (2026).

[2] Xiong, R. *et al.* Observation of propagating collective spin–valley modes in twisted WSe_2. *Nature Physics* (2026). https://doi.org/10.1038/s41567-026-03280-w

[3] Taking snapshots of spin–valley modes in a moiré superlattice. *Nature Physics* (2026). https://doi.org/10.1038/s41567-026-03295-3

[4] Bi, Z. & Fu, L. Excitonic density wave and spin–valley superfluid in bilayer transition metal dichalcogenide. *Nature Communications* **12**, 642 (2021).

[5] Mak, K. F. & Shan, J. Semiconductor moiré materials. *Nature Nanotechnology* **17**, 686–695 (2022).

[6] Xia, Y. *et al.* Superconductivity in twisted bilayer WSe_2. *Nature* **637**, 833–838 (2025).

[7] Guo, Y. *et al.* Superconductivity in 5.0° twisted bilayer WSe_2. *Nature* **637**, 839–845 (2025).

[8] Bernevig, B. A. *et al.* Fractional quantization in insulators from Hall to Chern. *Nature Physics* **21**, 1702–1713 (2025).

[9] Witte, M., Meuer, J., Plésiat, É. & Kadow, C. Field-space attention for structure-preserving earth system transformers. arXiv:2501.xxxxx (2025).

[10] Lam, R. *et al.* GraphCast: learning skillful medium-range global weather forecasting. *Science* **382**, 1416–1421 (2023).

[11] Bi, K. *et al.* Accurate medium-range global weather forecasting with 3D neural networks. *Nature* **619**, 533–538 (2023).

[12] Kurth, T. *et al.* FourCastNet: accelerating global high-resolution weather forecasting using adaptive Fourier neural operators. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 1–14 (2023).

[13] Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-resolution image synthesis with latent diffusion models. *Proceedings of CVPR*, 10684–10695 (2022).

[14] Connor, M., Canal, G. & Rozell, C. Variational autoencoder with learned latent structure. *Proceedings of AISTATS*, 2359–2367 (2021).

[15] Górski, K. M. *et al.* HEALPix: a framework for high-resolution discretization and fast analysis of data distributed on the sphere. *Astrophysical Journal* **622**, 759–771 (2005).

[16] Hersbach, H. *et al.* The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* **146**, 1999–2049 (2020).

[17] Song, J., Meng, C. & Ermon, S. Denoising diffusion implicit models. *International Conference on Learning Representations* (2021).

[18] Lessig, C. *et al.* AtmoRep: a stochastic model of atmosphere dynamics using large scale representation learning. arXiv:2308.13280 (2023).

[19] Friston, K. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience* **11**, 127–138 (2010).

[20] Tononi, G., Boly, M., Massimini, M. & Koch, C. Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience* **17**, 450–461 (2016).

[21] Shaw, T. A. & Stevens, B. The other climate crisis. *Nature* **639**, 877–887 (2025).

[22] Arora, S., Hu, W. & Kothari, P. K. An analysis of the t-SNE algorithm for data visualization. *Conference on Learning Theory*, 1455–1462 (2018).

[23] Angiulli, F., Fassetti, F. & Ferragina, L. LatentOut: an unsupervised deep anomaly detection approach exploiting latent space distribution. *Machine Learning* **112**, 4323–4349 (2022).

[24] Yan, C.-W. *et al.* Fourier amplitude and correlation loss: beyond using L2 loss for skillful precipitation nowcasting. *Advances in Neural Information Processing Systems* **37**, 100007–100041 (2024).

[25] Hunt, K. M. R. Stop using root-mean-square error as a precipitation target! arXiv:2509.08369 (2025).

[26] Sundstrom, M., Gambrell, J., Green, C., Traxler, A. L. & Brewe, E. Relative benefits of different active learning methods to conceptual physics learning. arXiv:2505.04577 [physics.ed-ph] (2025).

[27] Mazur, E. *Peer Instruction: A User's Manual*. Prentice Hall, Upper Saddle River, NJ (1997).

[28] Etkina, E. & Van Heuvelen, A. Investigative science learning environment — a science process approach to learning physics. In *Research-Based Reform of University Physics*, vol. 1, pp. 1–48 (2007).

[29] McDermott, L. C. & Shaffer, P. S. *Tutorials in Introductory Physics*. Prentice Hall, Upper Saddle River, NJ (2002).

[30] Beichner, R. J. *et al.* The student-centered activities for large enrollment undergraduate programs (SCALE-UP) project. In *Research-Based Reform of University Physics*, vol. 1, pp. 2–39 (2007).

[31] Andrews, T. M., Leonard, M. J., Colgrove, C. A. & Kalinowski, S. T. Active learning not associated with student learning in a random sample of college biology courses. *CBE—Life Sciences Education* **10**, 394–405 (2011).

[32] Campbell, D. T. Assessing the impact of planned social change. *Evaluation and Program Planning* **2**, 67–90 (1979).

[33] Polanyi, M. *The Tacit Dimension*. Doubleday, Garden City, NY (1966).