

Constraint Closure and Entropic Decimation: Toward a Unified Field Theory of Computation, Cognition, and Cosmology

Flyxion

April 21, 2026

Abstract

Contemporary models of computation and cognition are dominated by predictive paradigms that optimize local consistency without guaranteeing global realizability. This essay develops an alternative framework in which physical fields, cognitive processes, and computational systems are understood as instances of constraint accumulation followed by entropy-respecting decimation. Building on the Integral Decimation Method and tensor-train representations, we reinterpret inference as a process of reconstructing globally admissible configurations through sequential application of local operators and nonlinear compression.

Within this framework, the Relativistic Scalar–Vector Plenum (RSVP) theory is formalized as a field system in which a scalar field encodes constraint density, a vector field encodes directional organization, and an entropy field regulates admissibility rather than participating symmetrically in dynamics. The resulting action functional introduces an explicit entropy floor dependent on local field intensity, producing a soft but enforceable admissibility condition.

A computational scheme is derived by discretizing the action and applying a Trotter-split decomposition into structural and admissibility gates, yielding a tensor-train evolution governed by sequential local updates and singular value truncation. This decimation process is interpreted as an entropy-filtering mechanism that preserves globally consistent structure while eliminating incoherent modes.

We propose a convergence conjecture stating that the resulting truncated sweep map is contractive under appropriate conditions, supported empirically by stabilization of bond dimension across sweeps. Categorical, information-theoretic, dynamical-systems, and geometric analyses of the scheme are developed, connecting it to sheaf-theoretic gluing, Landauer’s principle, projected gradient flow, and manifold curvature. This establishes a unifying principle: stable structure in physics, cognition, and computation arises not from prediction but from constraint closure under admissibility filtering. The essay thereby connects tensor-network numerics, thermodynamic field theory, and models of intelligence within a single geometric and informational framework.

Contents

1	Introduction: From Prediction to Reconstruction	3
2	The RSVP Action with Entropy as Regulator	4
2.1	Field Variables and Their Asymmetric Roles	4
2.2	The Entropy Floor	5
2.3	The Admissibility Potential	5
2.4	The Formal Action	6
2.5	Structural Terms	6
2.6	Entropy as Regulator: The Asymmetric Terms	7
2.7	Variational Equations of Motion	7
2.8	Weight Functional	7
3	The Prototype Decimation Scheme	8
3.1	From Formal Action to Discrete Lattice	8
3.2	Local Action Decomposition	8
3.3	Trotter Splitting of the Local Action	9
3.4	The Admissibility Gate	9
3.5	The Structural Gate	10
3.6	The Decimation Step	10
3.7	Sweep Protocol	11
4	Implementation and Observables	11
4.1	Simulator Architecture	11
4.2	Local Marginals and Field Expectations	11
4.3	Bond Dimension as Closure Diagnostic	12
4.4	Simulation Results	12
5	The Convergence Conjecture	13
5.1	The Truncated Sweep Map	13
5.2	Component Properties	14
5.3	Conjecture Statement	15
5.4	Empirical Support	16
5.5	Relation to Renormalization	17
6	Categorical Formulation of Constraint Closure	17
6.1	The Category of Configurations	17
6.2	The Admissibility Subcategory	17
6.3	Constraint Closure as a Fixed Point	18
6.4	Relation to Sheaf-Theoretic Gluing	18

7	Information-Theoretic Interpretation	18
7.1	Entropy as Computational Resource	18
7.2	Decimation as Information Compression	19
7.3	Connection to Predictive Models	19
8	Dynamical Systems Perspective on Constraint Closure	20
8.1	State Space and Evolution Operator	20
8.2	Invariant Sets and Attractors	20
8.3	Lyapunov Functional	20
8.4	Entropy Production and Irreversibility	21
9	Geometric Interpretation	21
9.1	Configuration Manifold	21
9.2	Projection onto the Admissible Manifold	21
9.3	Geodesic Interpretation of Sweeps	22
9.4	Curvature and Complexity	22
10	Extensions and Open Problems	22
10.1	Higher-Dimensional Lattices	22
10.2	Continuous Limits and Field Reconstruction	22
10.3	Adaptive Truncation and Dynamic Resolution	23
10.4	Toward a Proof of Convergence	23
10.5	Cognitive and Physical Implications	23
11	Conclusion: Constraint Closure as a Unifying Principle	23

1 Introduction: From Prediction to Reconstruction

The dominant paradigm in contemporary artificial intelligence treats intelligence as prediction. Systems are trained to estimate the next token, the next pixel, or the next state in a sequence, and performance is measured by the accuracy of these local estimates. This approach has produced remarkable practical successes, yet it exhibits a fundamental structural limitation. Prediction ensures local plausibility, but it does not guarantee global consistency. A sequence of locally plausible steps may fail to assemble into a coherent whole.

This limitation is not merely technical but conceptual. Predictive systems operate on fragments of structure without enforcing the constraints that make those fragments mutually compatible. The resulting outputs often exhibit internal contradictions, discontinuities, or failures of realizability that cannot be repaired by improving local accuracy alone. The problem is not insufficient prediction, but the absence of a mechanism that enforces global admissibility.

In contrast, physical systems do not operate by prediction. They evolve under constraints. A configuration of a physical field is not accepted because it is locally plausible, but because it satisfies a set of governing relations simultaneously. The global state of the system is therefore a solution to a constraint problem rather than a sequence of guesses. Stability emerges when these constraints are satisfied in a self-consistent manner.

This essay develops a framework in which computation, cognition, and physical dynamics are unified under this constraint-first perspective. The central claim is that intelligence and structure arise from the process of reconstructing globally admissible configurations through the accumulation of local constraints and the elimination of incoherent degrees of freedom. This process is formalized as a combination of local operator application and entropy-respecting decimation.

Recent advances in tensor-network methods, particularly the Integral Decimation Method of Grimm, Staat, and Eaves [1], provide a concrete computational realization of this idea. High-dimensional integrals, which are intractable when treated as monolithic objects, can be decomposed into sequences of local transformations followed by controlled compression. The key insight is that the exponential weight $e^{iS[\psi]}$ over an action S can be reconstructed as a product of gate operators acting sequentially on an initially unentangled state, changing the computational complexity of integration from exponential to polynomial [1, 2].

At the same time, the RSVP framework provides a field-theoretic language in which constraint density, directional organization, and entropy are represented explicitly. By combining these perspectives, we obtain a unified description in which the evolution of a system is governed by local interactions while global consistency is enforced through decimation.

The key shift is therefore from prediction to reconstruction. Rather than asking what the next

state should be, we ask which configurations can exist at all under the imposed constraints. The computational task becomes one of assembling a globally consistent structure from local components while discarding those components that cannot be integrated into a coherent whole. This shift aligns naturally with tensor-network representations, in which a global object is never stored explicitly but assembled from locally structured cores whose internal dimensions — the bond dimensions — measure the retained correlations between neighboring regions.

This perspective has implications beyond numerical methods. It suggests that cognition itself may operate as a process of constraint closure, in which hypotheses are generated and then filtered by their compatibility with an evolving internal model. It also suggests that physical structure arises not from expansion or accumulation, but from the progressive elimination of inadmissible configurations.

The remainder of the essay develops this framework in detail. Section 2 formalizes the RSVP action with entropy as a regulator of admissibility, deriving the equations of motion and establishing the asymmetric role of the entropy field. Section 3 derives a prototype decimation scheme based on a Trotter-split decomposition of the action, leading to a tensor-train implementation with exact commutativity of the Trotter factors on the local basis. Section 4 describes the implementation and the observables extracted from it, focusing on bond-dimension diagnostics as signatures of constraint closure. Section 5 presents empirical evidence for convergence and formulates the convergence conjecture, relating it to the renormalization group. Sections 6–9 develop categorical, information-theoretic, dynamical-systems, and geometric perspectives on constraint closure. Section 10 surveys extensions and open problems. Section 11 concludes.

2 The RSVP Action with Entropy as Regulator

2.1 Field Variables and Their Asymmetric Roles

Let $\Omega \subset \mathbb{R}^d$ be a compact domain. We define the RSVP state as a triple

$$X = (\Phi, \mathbf{v}, S) \in C^\infty(\Omega) \times C^\infty(\Omega; \mathbb{R}^d) \times C^\infty(\Omega; \mathbb{R}_{\geq 0}). \quad (1)$$

We distinguish the three RSVP fields by function rather than by symmetry. The scalar field $\Phi : \Omega \rightarrow \mathbb{R}$ encodes local constraint density or potential burden. It measures how strongly a region of the domain is required to satisfy structural relations. The vector field $\mathbf{v} : \Omega \rightarrow \mathbb{R}^d$ encodes directional transport or organizing flow. It carries information about how structure propagates, aligns, or redistributes across the domain.

The entropy field $S : \Omega \rightarrow \mathbb{R}_{\geq 0}$ plays a fundamentally different role. It does not participate symmetrically in the dynamics alongside Φ and \mathbf{v} . Instead, it regulates the admissibility of

configurations produced by them. A configuration of (Φ, \mathbf{v}) is not automatically allowed; it must be supported by sufficient entropy to remain stable and realizable. In this sense, S functions as a certification field, determining which configurations can exist rather than contributing directly to their formation.

This asymmetry is essential. If S were treated as an ordinary dynamical field on equal footing with Φ and \mathbf{v} , the distinction between structure and admissibility would collapse. The purpose of the RSVP formulation is precisely to preserve this distinction.

Definition 2.1 (Admissibility). A configuration X is *admissible* at $x \in \Omega$ if

$$S(x) \geq S_*(\Phi(x), \mathbf{v}(x)), \quad (2)$$

where $S_* : \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is the *entropy floor function*. The configuration is globally admissible if (2) holds for almost every $x \in \Omega$.

2.2 The Entropy Floor

We define the entropy floor as a quadratic in local field intensity:

$$S_*(\Phi, \mathbf{v}) = \gamma(\Phi^2 + \|\mathbf{v}\|^2), \quad \gamma > 0. \quad (3)$$

This expression defines the minimal entropy required to encode the local constraint complexity. Regions with large scalar density or strong flow require greater entropy to remain admissible. This establishes a scaling relation between structural intensity and informational overhead.

Proposition 2.2. *The entropy floor (3) is the unique quadratic function $S_*(\Phi, \mathbf{v}) = \gamma_1 \Phi^2 + \gamma_2 \|\mathbf{v}\|^2$ that is non-negative, vanishes at $(\Phi, \mathbf{v}) = (0, 0)$, and is isotropic in \mathbf{v} . Setting $\gamma_1 = \gamma_2 = \gamma$ is the minimal-parameter choice consistent with equal weighting of scalar and vector burden.*

Proof. Non-negativity and vanishing at the origin require both coefficients to be non-negative. Isotropy in \mathbf{v} requires the coefficient of $\|\mathbf{v}\|^2$ to be scalar. Equal weighting collapses two free parameters to one. \square

2.3 The Admissibility Potential

Rather than enforcing (2) as a hard constraint, we introduce a smooth penalty. Define

$$V_{\text{adm}}(S; \Phi, \mathbf{v}) = \log\left(1 + \exp(\alpha(S_*(\Phi, \mathbf{v}) - S))\right), \quad \alpha > 0. \quad (4)$$

This softplus barrier is smooth and differentiable everywhere. When S is well above the entropy floor, the potential is negligible. When S falls below S_* , the potential grows approx-

imately linearly, imposing a strong penalty. The parameter α controls the sharpness of the transition, while ν controls the strength of the penalty in the action.

Proposition 2.3 (Properties of V_{adm}). *The potential (4) satisfies: $V_{\text{adm}} > 0$ for all finite arguments; $V_{\text{adm}} \rightarrow 0$ as $S \rightarrow +\infty$; $V_{\text{adm}} \sim \alpha(S_* - S)$ as $S \rightarrow -\infty$; and $\partial V_{\text{adm}}/\partial S < 0$, so the penalty is strictly decreasing in entropy. Moreover $\partial V_{\text{adm}}/\partial \alpha > 0$: larger α sharpens the transition at $S = S_*$.*

Proof. Direct computation from the softplus definition $\text{sp}(u) = \log(1+e^u)$ with $u = \alpha(S_* - S)$. The derivative with respect to S is $\partial_S V_{\text{adm}} = -\alpha \sigma(\alpha(S_* - S))$, where σ is the logistic function, which is strictly positive. \square

This term gives entropy genuine regulatory power over the action. It does not merely modulate coupling between Φ and \mathbf{v} ; it determines whether a configuration is admissible at all. Together with the entropy floor (3), it ensures that high field intensity demands high entropy, driving the system toward a regime where structural complexity and informational overhead are in balance.

2.4 The Formal Action

We write the full RSVP action as

$$\mathcal{A}[X] = \int_{\Omega} \left(\mathcal{L}_{\Phi} + \mathcal{L}_{\mathbf{v}} + \mathcal{L}_S + e^{-\lambda S} \mathcal{L}_{\Phi\mathbf{v}} + \mu (\nabla S \cdot \mathbf{J}_{\Phi,\mathbf{v}}) + \nu V_{\text{adm}}(S; \Phi, \mathbf{v}) \right) dx, \quad (5)$$

where $X = (\Phi, \mathbf{v}, S)$. The action is composed of three classes of terms. The first consists of standard single-field contributions governing local magnitude and smoothness. The second consists of structural coupling between Φ and \mathbf{v} , modulated by entropy. The third consists of admissibility terms in which S acts as a regulator.

2.5 Structural Terms

The one-field contributions are quadratic in the fields and their gradients:

$$\mathcal{L}_{\Phi} = a_{\Phi} \Phi^2 + b_{\Phi} |\nabla \Phi|^2, \quad (6)$$

$$\mathcal{L}_{\mathbf{v}} = a_v \|\mathbf{v}\|^2 + b_v \|\nabla \mathbf{v}\|^2, \quad (7)$$

$$\mathcal{L}_S = a_S S^2 + b_S |\nabla S|^2. \quad (8)$$

These terms ensure that each field remains bounded and exhibits controlled variation. The coupling between scalar and vector sectors is

$$\mathcal{L}_{\Phi\mathbf{v}} = c_{\Phi v} \Phi (\nabla \cdot \mathbf{v}) + c'_{\Phi v} |\nabla \Phi|^2 \|\mathbf{v}\|^2. \quad (9)$$

The divergence term captures how scalar accumulation responds to compression and expansion of flow; the gradient interaction term couples spatial variation in Φ to the magnitude of \mathbf{v} .

2.6 Entropy as Regulator: The Asymmetric Terms

The defining feature of the RSVP action is the asymmetric role of entropy. The factor $e^{-\lambda S} \mathcal{L}_{\Phi\mathbf{v}}$ modulates the scalar–vector coupling. In regions of high entropy, coupling is suppressed; in regions of low entropy, it is amplified. This is not a symmetric interaction. It is a gating mechanism in which S determines how strongly the structural fields are allowed to interact.

The interaction current $\mathbf{J}_{\Phi,\mathbf{v}}$ is the variational current of the coupling term:

$$\mathbf{J}_{\Phi,\mathbf{v}} = \frac{\partial \mathcal{L}_{\Phi\mathbf{v}}}{\partial(\nabla\Phi)} + \frac{\partial \mathcal{L}_{\Phi\mathbf{v}}}{\partial(\nabla\mathbf{v})}. \quad (10)$$

Evaluating from (9) with $c'_{\Phi v} = 0$, the current reduces to $\mathbf{J}_{\Phi,\mathbf{v}} = \Phi \mathbf{v}$, the simplest scalar-weighted flow vector. The term $\mu(\nabla S \cdot \mathbf{J}_{\Phi,\mathbf{v}})$ penalizes configurations in which entropy gradients are misaligned with the flow of scalar–vector coupling energy: when entropy increases in the direction of structural transport, the configuration is supported; when entropy gradients oppose the direction of coupling, the configuration is penalized.

2.7 Variational Equations of Motion

Taking functional derivatives of (5) yields the Euler–Lagrange equations for each field. The entropy equation is the most revealing:

$$-b_S \Delta S + 2a_S S - \lambda e^{-\lambda S} \mathcal{L}_{\Phi\mathbf{v}} + \mu \nabla \cdot \mathbf{J}_{\Phi,\mathbf{v}} - \nu \alpha \sigma(\alpha(S_* - S)) = 0, \quad (11)$$

where $\sigma(u) = (1 + e^{-u})^{-1}$ is the logistic function. The term $-\lambda e^{-\lambda S} \mathcal{L}_{\Phi\mathbf{v}}$ drives S upward where coupling is strong, while $-\nu \alpha \sigma(\alpha(S_* - S))$ drives S upward when it falls below the entropy floor. Both forces push S in the same direction when the field is intense: the system self-regulates toward admissibility. This is not imposed externally but emerges from the structure of the action.

2.8 Weight Functional

The configuration weight is

$$W[X] = e^{-\beta \mathcal{A}[X]}. \quad (12)$$

Observables are computed as weighted averages over the configuration space. However, direct evaluation of W is exponentially intractable, as it requires integration over all possible

field configurations. The remainder of the essay develops a method for constructing and compressing this weight through local operations.

Remark 2.4 (Double-exponential and linearization). When the gate is derived via $G_k \sim e^{-\beta \Delta \mathcal{A}_k}$, the entropy-gated coupling term $e^{-\lambda S} \mathcal{L}_{\Phi_{\mathbf{v}}}$ contributes $\exp(-\beta e^{-\lambda S} \Delta \mathcal{L}_{\Phi_{\mathbf{v}}})$ to the gate: a double exponential in S . This is analytically correct but numerically hazardous, as the dynamic range across local basis states can be extreme and degrades SVD truncation quality. The prototype scheme of Section 3 separates the admissibility and structural sectors via a Trotter split, avoiding this issue while preserving the conceptual asymmetry.

3 The Prototype Decimation Scheme

3.1 From Formal Action to Discrete Lattice

To obtain a computationally tractable representation of the RSVP weight functional, we discretize Ω as a one-dimensional chain of N sites. Each site n carries a local state

$$x_n = (\Phi_n, v_n, S_n),$$

drawn from a finite local basis of size $d = d_\Phi \times d_v \times d_S$. The global weight $W(x_1, \dots, x_N) = e^{-\beta \mathcal{A}[X]}$ is not assumed to factorize a priori. Instead, it is constructed through a sequence of local operator applications derived from the discretized action. This construction induces an approximate matrix product state (MPS) / tensor-train representation [2, 4]:

$$W(x_1, \dots, x_N) \approx \sum_{\{\alpha\}} C_1^{\alpha_0 \alpha_1}(x_1) C_2^{\alpha_1 \alpha_2}(x_2) \dots C_N^{\alpha_{N-1} \alpha_N}(x_N), \quad (13)$$

where each core C_n has shape (χ_{n-1}, d, χ_n) and $\chi_0 = \chi_N = 1$ (open boundary conditions). The bond dimension χ_n measures retained correlations between neighboring sites.

The chain begins in a trivially factorized state, analogous to the unentangled initial state of the Integral Decimation Method [1]. All nontrivial dependence between sites is accumulated through the sequential application of local operators. This representation emerges from the same principle as the reference method: a global object is built incrementally through local transformations, with compression applied at each step to maintain tractability.

3.2 Local Action Decomposition

The discrete action decomposes as

$$\mathcal{A}_{\text{discrete}}[X] = \sum_n \mathcal{A}_n^{(1)}(x_n) + \sum_n \mathcal{A}_{n,n+1}^{(2)}(x_n, x_{n+1}), \quad (14)$$

where the one-site term is

$$\mathcal{A}_n^{(1)}(x_n) = a_\Phi \Phi_n^2 + a_v v_n^2 + a_S S_n^2, \quad (15)$$

and the two-site term is split into structural and admissibility sectors:

$$\mathcal{A}_{n,n+1}^{(2)}(x_n, x_{n+1}) = \mathcal{A}_{\Phi\mathbf{v}}(x_n, x_{n+1}) + \mathcal{A}_{\text{adm}}(x_n, x_{n+1}). \quad (16)$$

3.3 Trotter Splitting of the Local Action

The local contribution of the action on a pair of neighboring sites $(n, n+1)$ is denoted by $\Delta\mathcal{A}_{n,n+1}$. Rather than exponentiating the full local action directly, we decompose it according to (16) and apply a first-order Trotter splitting.

Definition 3.1 (Trotter split). Given a two-site action of the form (16), the *first-order Trotter split* of the corresponding gate is

$$G_{n,n+1} = e^{-\beta\mathcal{A}_{n,n+1}^{(2)}} \approx G_{\text{adm}} G_{\Phi\mathbf{v}} = e^{-\beta\mathcal{A}_{\text{adm}}} e^{-\beta\mathcal{A}_{\Phi\mathbf{v}}}. \quad (17)$$

Proposition 3.2 (Trotter error and exact commutativity). *The error in (17) is bounded by $\frac{\beta^2}{2} \|[\mathcal{A}_{\Phi\mathbf{v}}, \mathcal{A}_{\text{adm}}]\|_F + O(\beta^3)$. In the prototype, \mathcal{A}_{adm} depends only on single-site S values and $\mathcal{A}_{\Phi\mathbf{v}}$ depends only on (Φ, v) pairs; on the local discrete basis these operators are simultaneously diagonal, so the commutator vanishes exactly and the Trotter split is exact.*

Proof. The Baker–Campbell–Hausdorff formula gives $e^{A+B} = e^A e^B e^{-[A,B]/2} \dots$, so the leading error is $\frac{1}{2}[A, B]$ in the exponent. Since both gates are diagonal on the local basis, they commute on every basis element and the commutator is zero. \square

The Trotter splitting therefore introduces zero first-order error in the prototype regime, making it an exact operator factorization rather than an approximation. This also preserves the conceptual separation between structure and admissibility at the level of the numerical scheme.

3.4 The Admissibility Gate

The admissibility sector acts independently on each site within the pair:

$$\mathcal{A}_{\text{adm}}(x, y) = \nu V_{\text{adm}}(x) + \nu V_{\text{adm}}(y),$$

where V_{adm} is evaluated using the full local triplet (Φ, v, S) . The corresponding gate is diagonal:

$$G_{\text{adm}}(x', y'; x, y) = \delta_{x'x} \delta_{y'y} \exp(-\beta \mathcal{A}_{\text{adm}}(x, y)).$$

Because this gate depends only on local variables and does not couple neighboring sites directly, it is computationally inexpensive. Its role is purely regulatory: it suppresses configurations that violate the entropy floor and leaves admissible configurations largely unchanged.

3.5 The Structural Gate

The structural sector captures the interaction between scalar and vector fields across neighboring sites:

$$\mathcal{A}_{\Phi\mathbf{v}}(x, y) = b_{\Phi}(\Phi_x - \Phi_y)^2 + b_v(v_x - v_y)^2 + c_{\Phi v} \Phi_x v_y.$$

This term enforces coherence between neighboring values while allowing directed interaction through the asymmetric cross-term $\Phi_x v_y$. The corresponding gate is also diagonal:

$$G_{\Phi\mathbf{v}}(x', y'; x, y) = \delta_{x'x} \delta_{y'y} \exp(-\beta \mathcal{A}_{\Phi\mathbf{v}}(x, y)).$$

Although diagonal, this gate induces correlations between neighboring sites by reweighting configurations according to their structural compatibility.

3.6 The Decimation Step

After applying a two-site gate to cores C_n and C_{n+1} , the update proceeds in four steps.

Step 1: Contract.

$$\Theta^{\ell, i, j, r} = \sum_m C_n^{\ell i m} C_{n+1}^{m j r}. \quad (18)$$

Step 2: Apply gate. Since the combined gate (19) is diagonal,

$$\tilde{\Theta}^{\ell, i, j, r} = G_{n, n+1}(i, j; i, j) \Theta^{\ell, i, j, r}. \quad (19)$$

Step 3: Reshape and SVD.

$$\Theta_{\text{mat}} = \tilde{\Theta}.\text{reshape}(Ld, dR), \quad \Theta_{\text{mat}} = U\Sigma V^\dagger. \quad (20)$$

Step 4: Truncate and split. Retain $k = \min(\max(1, |\{i : \sigma_i > \varepsilon\}|), \chi_{\max})$ singular values:

$$C_n^{\text{new}} = (U_{:,1:k} \text{diag}(\sqrt{\sigma_{1:k}})).\text{reshape}(L, d, k), \quad C_{n+1}^{\text{new}} = (\text{diag}(\sqrt{\sigma_{1:k}}) V_{1:k,:}^\dagger).\text{reshape}(k, d, R). \quad (21)$$

The update rule is:

$$W_{k+1} = \mathfrak{D}_\varepsilon(G_k W_k), \quad (22)$$

where \mathfrak{D}_ε denotes the truncation operator and G_k the gate at bond k . This is the core decimation equation of the scheme. It eliminates components of the representation that

contribute negligibly to the global weight, enforcing an entropy-respecting compression of the system.

3.7 Sweep Protocol

The full update proceeds through alternating sweeps. A left-to-right pass applies one-site and two-site gates sequentially across all sites and bonds, followed by a right-to-left pass in reverse order. After each pass, the tensor train is brought into left-canonical form by QR decomposition, ensuring that contractions used in observable computation remain well-conditioned. The full sweep map is:

$$\mathcal{T} = \mathcal{D}_\varepsilon \circ G_{\text{sweep}}, \quad (23)$$

where G_{sweep} denotes the full left-to-right-to-left gate sequence. This mirrors the time-evolving block decimation (TEBD) algorithm [6], adapted here to operate on a reweighted configuration space rather than a unitary quantum state.

4 Implementation and Observables

4.1 Simulator Architecture

The simulator implements the decimation scheme as a composition of modular components that mirror the mathematical structure of the method. The `lattice` module defines the discrete domain, including the enumeration of local states and the adjacency structure of the chain. The `core` module stores the tensor-train representation, providing operations for reshaping, contracting, and truncating cores. The `operators` module constructs one-site and two-site gates from the Trotter-split action, maintaining the separation between structural and admissibility contributions. The `sweeps` module executes alternating passes, applying gates and invoking truncation after each local update. The `observables` module computes marginal distributions and derived quantities from the current tensor-train state.

This modular decomposition reflects the separation of concerns present in the theoretical construction. The action determines the operators, the operators drive the evolution, and the tensor-train representation mediates between local updates and global structure.

4.2 Local Marginals and Field Expectations

Left environments $L^{(n)} \in \mathbb{R}^{\chi_n \times \chi_n}$ are computed recursively:

$$L^{(0)} = [1], \quad L_{bb'}^{(n)} = \sum_{a,a',i} L_{aa'}^{(n-1)} C_n^{aib} C_n^{a'ib'}. \quad (24)$$

Right environments $R^{(n)}$ are computed analogously from the right boundary. The local marginal probability at site n is then:

$$p_n(i) = \frac{1}{Z} \sum_{a,a',b,b'} L_{aa'}^{(n-1)} C_n^{aib} C_n^{a'ib'} R_{bb'}^{(n)}, \quad (25)$$

where $Z = \sum_i p_n(i)$ before normalization. These environments are most naturally computed in a canonical gauge, where orthogonality conditions prevent uncontrolled growth or decay of intermediate quantities.

Field expectations at site n are computed as

$$\langle \Phi_n \rangle = \sum_i p_n(i) \Phi^{(i)}, \quad \langle v_n \rangle = \sum_i p_n(i) v^{(i)}, \quad \langle S_n \rangle = \sum_i p_n(i) S^{(i)}. \quad (26)$$

4.3 Bond Dimension as Closure Diagnostic

The bond dimension χ_n at bond $(n, n+1)$ is the number of singular values retained after truncation at that bond. It measures the amount of correlation between neighboring sites that cannot be compressed below threshold ε .

Definition 4.1 (Effective bond dimension and closure burden).

$$\chi_{\text{eff}} = \frac{1}{N-1} \sum_{n=1}^{N-1} \chi_n, \quad \chi_{\text{max}} = \max_{1 \leq n \leq N-1} \chi_n. \quad (27)$$

The *closure burden* at bond n is defined as χ_n/χ_{max} . A uniform burden indicates homogeneous long-range correlation; localized peaks indicate sites where the field resists compression.

A small χ_{eff} indicates that the global configuration can be expressed with low inter-site correlation, corresponding to a state of low torsion and high admissibility. A large χ_{max} localized at a particular bond indicates a region of persistent structural interaction that resists decimation, often associated with sharp field gradients or competing constraints.

From the perspective of the framework, stabilization of bond dimensions under repeated sweeps indicates that the system has reached a regime in which constraint closure has been achieved at the chosen resolution. Information flow across bonds becomes bounded, and the representation ceases to grow in complexity. This boundedness is the empirical signature of admissibility.

4.4 Simulation Results

Figures 1 and 2 show χ_{eff} per sweep for varying β and ν respectively, on a chain of $N = 20$ sites with $d = 27$ local states, $\chi_{\text{max}} = 24$, and truncation threshold $\varepsilon = 10^{-5}$. Bond dimension grows initially as correlations are established, then stabilizes within approximately ten sweeps

across all tested parameter values. Figures 3 and 4 show the spatial profile of bond dimensions at the final sweep and the sweep-by-sweep evolution.




Figure 1: Effective bond dimension χ_{eff} per sweep for $\nu = 0.5$ and varying β . Stabilization within ten sweeps supports the contraction conjecture.

5 The Convergence Conjecture

5.1 The Truncated Sweep Map

The full update procedure can be expressed as a map on the space of tensor-train states. Let $\mathcal{M}_{\chi,N}$ denote the space of N -site tensor trains with bond dimension at most χ in left-canonical form, equipped with the Frobenius norm $\|W\|_F^2 = \sum_{x_1, \dots, x_N} W(x_1, \dots, x_N)^2$.

Let G denote the operator applying all gates in a complete sweep and Π_ε the nonlinear projection truncating the tensor train at threshold ε . The truncated sweep map is

$$\mathcal{T} = \Pi_\varepsilon \circ G. \tag{28}$$




Figure 2: Effective bond dimension χ_{eff} per sweep for $\beta = 1.0$ and varying ν . Stronger admissibility coupling does not prevent stabilization.

5.2 Component Properties

Lemma 5.1 (Nonexpansiveness of Π_ε). *For any two tensor trains $W, W' \in \mathcal{M}_{\chi, N}$ in left-canonical form,*

$$\|\Pi_\varepsilon(W) - \Pi_\varepsilon(W')\|_F \leq \|W - W'\|_F. \quad (29)$$

Proof sketch. SVD truncation is the best rank- k approximation in the Frobenius norm (Eckart-Young theorem). The map $W \mapsto \Pi_\varepsilon(W)$ is therefore a projection onto a closed set in the space of $(Ld) \times (dR)$ matrices. Projections onto closed convex sets are nonexpansive in Hilbert space norms. Applying this bond-by-bond preserves nonexpansiveness over the full chain. \square

Lemma 5.2 (Contractiveness of G near identity). *For sufficiently small β , the gate map G satisfies $\|G(W) - G(W')\|_F \leq (1 - c\beta)\|W - W'\|_F$ for some $c > 0$ depending on the action coefficients.*

Proof sketch. Each gate $G_k(i, j) = e^{-\beta \mathcal{A}(x^{(i)}, x^{(j)})} \leq 1$ since $\mathcal{A} \geq 0$. The Frobenius norm



Figure 3: Closure burden profile at the final sweep ($\beta = 1.0$, $\nu = 0.5$). Localized peaks mark regions of persistent structural correlation.

therefore cannot increase under gate application. For strictly positive \mathcal{A} , at least some gate values are strictly less than 1, giving a strict contraction on the support. \square

5.3 Conjecture Statement

Conjecture 5.3 (Contraction and fixed point). *Under the conditions of Lemmas 5.1 and 5.2, the truncated sweep map $\mathcal{T} = \Pi_\varepsilon \circ G$ is contractive on $\mathcal{M}_{\chi_{\max}, N}$. It has a unique fixed point W^* satisfying*

$$W^* = \mathcal{T}(W^*) = \Pi_\varepsilon(G(W^*)), \tag{30}$$

which is the admissible compressed representation of the RSVP weight functional at parameters $(\beta, \varepsilon, \chi_{\max})$.

Remark 5.4. The main obstacle to a full proof is that Π_ε is nonlinear, so standard Banach fixed-point arguments do not apply directly. A complete proof would require either showing Π_ε is strictly contractive, or working in a metric space where the composition can be shown contractive directly. The empirical evidence supports the conjecture across all tested




Figure 4: Bond dimension evolution per sweep ($\beta = 1.0, \nu = 0.5$). Rows are sweeps; columns are bonds. Stabilization is visible within the first ten sweeps.

parameter regimes.

The conjecture rests on two components. First, Π_ε must be nonexpansive in the chosen norm, which is expected when truncation error is dominated by discarded singular values. Second, G must be contractive, which occurs when β is small enough that the gates are close to identity and do not amplify deviations.

5.4 Empirical Support

Across parameter regimes with $\beta \in [0.5, 2.0]$, $\varepsilon = 10^{-5}$, $\chi_{\max} \in \{24\}$, and $\nu \in [0.2, 0.8]$, stabilization of χ_{eff} is observed within approximately ten sweeps (Figures 1 and 2). The stabilized value χ_{eff}^* decreases monotonically as ε increases or β decreases, consistent with the expectation that contraction strengthens as gate effects weaken and truncation becomes more aggressive. These trends align with the predictions of Conjecture 5.3.

5.5 Relation to Renormalization

The truncated sweep map \mathcal{T} bears a structural resemblance to renormalization group transformations in statistical mechanics [7, 5]. In both cases, local interactions are combined with a reduction of degrees of freedom to produce an effective description of the system.

However, there is a crucial difference. Traditional renormalization proceeds by integrating out short-range degrees of freedom to obtain a coarse-grained theory at larger scales [7]. The decimation scheme presented here preserves the full configuration space while compressing its representation. No degrees of freedom are eliminated in principle; instead, their contribution is encoded in a reduced set of correlations controlled by the bond dimension.

The fixed point W^* of \mathcal{T} is therefore not an RG fixed point — a scale-invariant theory — but a *compression fixed point*: the most compact tensor-train representation of the RSVP weight that survives repeated entropy-respecting decimation at threshold ε . This is a new kind of fixed point whose physical interpretation is precisely the notion of admissible closure developed throughout this monograph.

6 Categorical Formulation of Constraint Closure

6.1 The Category of Configurations

We formalize the reconstruction process categorically. Let \mathcal{C} be a category whose objects are RSVP configurations represented as tensor-train states and whose morphisms are admissible transformations induced by local operators.

An object $W \in \text{Ob}(\mathcal{C})$ is a tensor-train representation of a weight functional. A morphism $f : W \rightarrow W'$ corresponds to the application of a finite sequence of local gates followed by truncation. Composition of morphisms corresponds to successive sweeps of the decimation process.

This category is not arbitrary. It is endowed with additional structure that reflects the tensor-product decomposition of the system. In particular, \mathcal{C} carries a symmetric monoidal structure $(\mathcal{C}, \otimes, I)$ in which the tensor product corresponds to concatenation of independent subsystems and I is the trivial single-site system.

6.2 The Admissibility Subcategory

Define a full subcategory $\mathcal{C}_{\leq \varepsilon} \subset \mathcal{C}$ consisting of tensor-train states whose bond dimensions are bounded by χ_{\max} and whose truncation error is below ε at every bond. These objects represent admissible compressed configurations. The truncation operator Π_ε acts as a reflector:

$$\Pi_\varepsilon : \mathcal{C} \rightarrow \mathcal{C}_{\leq \varepsilon}, \tag{31}$$

mapping each $W \in \mathcal{C}$ to the closest admissible object in $\mathcal{C}_{\leq \varepsilon}$ with respect to the Frobenius norm.

Proposition 6.1. Π_ε is idempotent: $\Pi_\varepsilon \circ \Pi_\varepsilon = \Pi_\varepsilon$.

Proof. Once a tensor-train state satisfies the truncation constraints, further application of Π_ε does not alter it, since no singular values below threshold remain. \square

6.3 Constraint Closure as a Fixed Point

Let $G : \mathcal{C} \rightarrow \mathcal{C}$ denote the endofunctor induced by a full sweep of gate applications. The truncated sweep map is the composite endofunctor $\mathcal{T} = \Pi_\varepsilon \circ G$, and constraint closure corresponds to a fixed point:

$$W^* \simeq \mathcal{T}(W^*). \quad (32)$$

In categorical terms, W^* is an algebra over the endofunctor \mathcal{T} : a configuration invariant under the process of local generation followed by admissibility filtering. This invariance expresses global consistency. The configuration contains no components that would be removed by further decimation, and therefore no components that fail to belong to a consistent global section of the field.

6.4 Relation to Sheaf-Theoretic Gluing

The categorical formulation admits an interpretation in terms of sheaf theory. Each local tensor core defines a section over a site, and the bond indices encode compatibility conditions between neighboring sections. Constraint closure corresponds to the existence of a global section obtained by gluing local data. Failures of closure correspond to obstructions where local compatibility conditions cannot be satisfied globally.

The truncation step Π_ε can therefore be interpreted as removing cohomological defects: components of the representation that prevent the existence of a consistent global section. Decimation is a process of enforcing sheaf consistency by eliminating non-glueable contributions. This connects to the Baez–Stay Rosetta Stone [23], in which logical, physical, and computational structures share a common categorical skeleton; the RSVP decimation scheme realizes that skeleton concretely in a field-theoretic setting.

7 Information-Theoretic Interpretation

7.1 Entropy as Computational Resource

The entropy field S acquires a direct interpretation in information theory. It measures the information required to encode a local configuration of (Φ, \mathbf{v}) , and the entropy floor S_* represents the minimal information necessary to represent the local constraint structure.

The admissibility condition $S \geq S_*$ expresses a fundamental principle: no structure can exist without sufficient information to encode it. This aligns with Landauer’s principle [10], which relates information processing to physical entropy, and with the broader statistical-mechanical framework in which entropy constrains accessible states [13]. Bennett’s analysis of reversible computation [11] and Lloyd’s bound on computational capacity [12] further ground this correspondence: the entropy floor $S_*(\Phi, \mathbf{v}) = \gamma(\Phi^2 + \|\mathbf{v}\|^2)$ can be interpreted as a local Lloyd bound, setting the minimum informational overhead required for a region of field intensity to sustain a dynamical structure.

7.2 Decimation as Information Compression

The truncation operator Π_ε performs lossy compression of the tensor-train representation. Singular values below threshold ε correspond to correlations that contribute negligibly to the global weight and can be discarded without significantly altering the representation. This is analogous to rate-distortion theory [14], in which a signal is compressed subject to a constraint on reconstruction error. Here, the distortion measure is the Frobenius norm, and the rate is determined by the bond dimension.

Proposition 7.1. *The bond dimension χ_n at bond $(n, n + 1)$ provides an upper bound on the mutual information between the two halves of the system separated by that bond.*

Proof. In tensor-network representations, the entanglement entropy across a bond is bounded by $\log \chi_n$ [4]. Since mutual information is bounded by entanglement entropy, the result follows. \square

Decimation therefore enforces an information constraint. Only correlations representable within the available information budget are retained, reinforcing the interpretation of constraint closure as an information-limited reconstruction process.

7.3 Connection to Predictive Models

Predictive models optimize local likelihoods without enforcing global constraints. From the present perspective, this corresponds to constructing local marginals without ensuring they can be assembled into a consistent joint distribution. The failure mode is well-known: locally coherent outputs that are globally incoherent, a pathology not correctable by improving local accuracy alone [18, 19].

The tensor-train framework, by contrast, constructs a global object directly, with local marginals derived from it. Decimation ensures that this global object remains within an admissible complexity class. Robust intelligence therefore requires mechanisms for maintaining global consistency under information constraints, rather than optimizing local predictions in isolation.

8 Dynamical Systems Perspective on Constraint Closure

8.1 State Space and Evolution Operator

We reinterpret the truncated sweep map $\mathcal{T} = \Pi_\varepsilon \circ G$ as a discrete-time dynamical system acting on the state space $\mathcal{M}_{\chi_{\max}, N}$. Each iteration

$$W_{k+1} = \mathcal{T}(W_k) \tag{33}$$

defines a trajectory in this space. The evolution is composed of a smooth transformation G followed by a nonlinear projection Π_ε , placing the system in the class of *projected dynamical systems*, where evolution is constrained to remain within an admissible subset.

8.2 Invariant Sets and Attractors

An invariant set $\mathcal{A} \subset \mathcal{M}_{\chi_{\max}, N}$ satisfies $\mathcal{T}(\mathcal{A}) \subseteq \mathcal{A}$. The fixed point W^* of Conjecture 5.3 is a trivial invariant set. Empirically, no non-trivial periodic or quasi-periodic behavior is observed in the prototype regime; trajectories converge monotonically toward a stable configuration, suggesting a global attractor consisting of a single fixed point.

Proposition 8.1. *If \mathcal{T} is contractive, it admits a unique global attractor consisting of the single fixed point W^* [21].*

Proof. Standard contraction mapping theory implies that all trajectories converge to the unique fixed point, and no other invariant sets can exist. \square

The absence of oscillatory behavior is consistent with the interpretation of decimation as an entropy-increasing process at the level of representation. Each iteration removes incoherent components, reducing the effective dimensionality of the state, and this monotonic reduction prevents the recurrence of previous states.

8.3 Lyapunov Functional

Define

$$\mathcal{L}(W) = \|W - \mathcal{T}(W)\|_F^2. \tag{34}$$

This measures deviation from invariance under the update map and vanishes at the fixed point.

Proposition 8.2. *Under the contraction assumption, $\mathcal{L}(W_k)$ is nonincreasing along trajectories.*

Proof. Contractiveness implies successive iterates approach W^* , reducing $\|W - \mathcal{T}(W)\|_F$ at each step. \square

This serves as a Lyapunov function for the system, certifying stability of the fixed point. In practice, $\mathcal{L}(W_k)$ can be estimated from the change in χ_{eff} between successive sweeps, providing a convergence diagnostic computable without explicit knowledge of W^* .

8.4 Entropy Production and Irreversibility

The decimation step introduces irreversibility. Information discarded through truncation cannot be recovered by subsequent iterations, distinguishing this system from reversible Hamiltonian dynamics. This connects directly to Landauer’s principle [10]: the discarding of information has a minimum thermodynamic cost of $k_B T \ln 2$ per bit, and the total truncation error accumulated across sweeps represents a lower bound on the thermodynamic cost of the computation.

The entropy field S plays a dual role: it regulates admissibility in the action, and the decimation process itself produces entropy at the level of representation as information about discarded modes is irretrievably lost. Both mechanisms act to eliminate configurations that cannot be sustained within the available informational budget, creating a tight correspondence between physical admissibility and computational irreversibility.

9 Geometric Interpretation

9.1 Configuration Manifold

The space of RSVP configurations can be viewed as a manifold \mathcal{X} parameterized by (Φ, \mathbf{v}, S) . The action functional $\mathcal{A}[X]$ defines a geometry on this manifold, with lower-action regions corresponding to more probable configurations under $W[X] = e^{-\beta \mathcal{A}[X]}$.

The admissibility condition defines a submanifold

$$\mathcal{X}_{\text{adm}} = \{X \in \mathcal{X} \mid S(x) \geq S_*(\Phi(x), \mathbf{v}(x)) \text{ for a.e. } x \in \Omega\}. \quad (35)$$

The soft potential V_{adm} smooths the boundary of this submanifold into a transition region, ensuring that the action remains differentiable and the Euler–Lagrange equations (11) are well-posed everywhere.

9.2 Projection onto the Admissible Manifold

The truncation operator Π_ϵ acts geometrically as a projection onto a low-dimensional manifold of admissible configurations within the space of tensor-train representations. Each truncation step selects the closest point in Frobenius norm within this manifold, removing components that lie outside it. Repeated application drives the system toward the intersection of the image of G and the admissible manifold.

9.3 Geodesic Interpretation of Sweeps

Gate application G can be interpreted as moving along a trajectory in configuration space determined by local gradients of the action. In the continuous limit, this corresponds to gradient flow:

$$\frac{dX}{dt} = -\nabla\mathcal{A}[X]. \quad (36)$$

The discrete sweep process approximates this flow, while the truncation step projects the trajectory back onto the admissible manifold. The overall dynamics is therefore a *projected gradient flow*:

$$X_{k+1} = \Pi_{\text{adm}}(X_k - \eta \nabla\mathcal{A}[X_k]), \quad (37)$$

connecting the tensor-network algorithm to variational optimization methods with admissibility enforced at each step rather than only at convergence.

9.4 Curvature and Complexity

Regions of high bond dimension correspond to regions of high curvature in the configuration manifold, where local constraints interact complexly and require many degrees of freedom to represent accurately. Decimation flattens these regions by removing negligible directions of curvature.

The stabilized configuration lies in a region of reduced curvature, where the manifold is effectively lower-dimensional. Constraint closure is therefore the process of finding a low-curvature submanifold approximating the full configuration space while preserving admissible structure. The bond-dimension profile of Figure 3 is a direct measurement of this curvature distribution along the lattice.

10 Extensions and Open Problems

10.1 Higher-Dimensional Lattices

The prototype scheme is formulated on a one-dimensional chain. Extension to higher-dimensional lattices requires replacing tensor trains with projected entangled pair states (PEPS) or tree tensor networks [3]. The main challenge is computational tractability, as exact contraction costs grow exponentially with dimension. Nevertheless, the conceptual framework remains unchanged: local operators are applied to neighboring regions, decimation controls correlation growth, and the admissibility potential V_{adm} requires no modification.

10.2 Continuous Limits and Field Reconstruction

A key open problem is the recovery of continuous field behavior from the discrete tensor-train representation. In the limit of large N and refined local bases, the discrete model

should approximate solutions to the Euler–Lagrange equations derived from the RSVP action. Establishing this correspondence rigorously would connect the computational scheme to the underlying field theory, analogous to the relationship between lattice field theory and continuum quantum field theory.

10.3 Adaptive Truncation and Dynamic Resolution

The truncation threshold ε and maximum bond dimension χ_{\max} are currently fixed parameters. An adaptive scheme varying these quantities across the lattice would improve efficiency by allocating more resources to regions of high closure burden. This aligns with the interpretation of bond dimension as a curvature measure: dynamically adjusting computational effort according to local constraint complexity.

10.4 Toward a Proof of Convergence

The convergence conjecture remains the central open theoretical problem. Possible approaches include: constructing a weighted norm in which Π_ε becomes strictly contractive; bounding the Lipschitz constant of G in terms of action parameters $(\beta, a_\Phi, a_\nu, a_S, b_\Phi, b_\nu, \nu)$; or reformulating in terms of monotone operator theory, which accommodates nonlinear projections more naturally than classical Banach fixed-point arguments. Each pathway would provide a rigorous foundation for the empirical observations presented here.

10.5 Cognitive and Physical Implications

The framework suggests a reinterpretation of both cognition and physical dynamics. The Friston free-energy principle [17] frames perception as minimization of a variational bound on surprise; constraint closure provides a complementary picture in which consistency, rather than surprise reduction, is the organizing criterion. Thought processes may be understood as iterative construction of admissible configurations, with inconsistencies eliminated through a process analogous to decimation.

In physics, the Jacobson derivation of Einstein’s equations from thermodynamic principles [15] and Verlinde’s entropic gravity program [16] both locate gravitational structure in entropic constraints rather than metric dynamics. The RSVP framework is consistent with this perspective and extends it: the entropy floor S_* connects local field intensity to informational overhead in a manner that could, in principle, be extended to a cosmological setting in which spacetime structure emerges from constraint closure rather than from a background metric.

11 Conclusion: Constraint Closure as a Unifying Principle

The development presented in this essay establishes a continuous chain from formal field theory to concrete computational implementation and finally to an empirically grounded conjecture about convergence. At each stage, the same structural principle appears: configurations are generated through local interactions, but only those that satisfy global admissibility constraints are retained. The mechanism by which this selection occurs is entropy-respecting decimation.

Within the RSVP framework, this principle is encoded directly in the action. The scalar and vector fields generate structure through local interactions, while the entropy field regulates whether that structure can be sustained. The admissibility potential enforces a relationship between local intensity and informational overhead, ensuring that configurations lacking sufficient entropy are suppressed. This introduces a hierarchy in which structure is subordinate to admissibility, rather than the two being treated as coequal components of the dynamics. As the entropy equation of motion (11) reveals, this hierarchy is not imposed externally but emerges from the variational structure of the action itself.

When the action is discretized and translated into a computational scheme, this same hierarchy persists. The Trotter-split decomposition separates structural and admissibility effects into distinct operators, with Proposition 3.2 establishing that this split is exact on the local basis. The sequential application of these operators constructs a global representation through local updates. The decimation step then enforces admissibility by removing components that cannot be integrated into a coherent structure. The tensor-train representation serves as the medium through which this process is realized, allowing global structure to emerge from local operations while remaining computationally tractable.

The convergence conjecture formalizes the observed behavior of this system. Stabilization of bond dimension under repeated sweeps indicates that the representation has reached a regime in which further application of \mathcal{T} does not increase complexity. Lemmas 5.1 and 5.2 establish the two components required for contraction, and Conjecture 5.3 states the expected fixed-point result. The categorical formulation of Section 6 shows that the fixed point is an algebra over the endofunctor \mathcal{T} and that decimation enforces sheaf consistency by removing non-glueable contributions. The information-theoretic analysis of Section 7 connects the entropy floor to Landauer's principle and frames bond dimension as a mutual-information bound. The dynamical-systems perspective of Section 8 supplies a Lyapunov functional and identifies the irreversibility of truncation with thermodynamic entropy production. The geometric interpretation of Section 9 recasts sweeps as projected gradient flow and bond-dimension peaks as regions of high manifold curvature. Although a formal proof of Conjecture 5.3 remains open, the convergence of these four perspectives provides strong evidence that constraint closure is not merely a conceptual idea but a measurable and geometrically grounded property of the

dynamics.

The implications of this framework extend beyond the specific numerical method. In computation, it suggests that the limitations of predictive models arise from their failure to enforce global consistency, and that more robust systems will require mechanisms for constraint closure. In cognition, it suggests that understanding may be better modeled as the process of assembling admissible configurations rather than predicting isolated outcomes. In physics, it provides a perspective in which structure arises from the selective retention of configurations that satisfy both local interactions and global admissibility conditions.

The unifying claim is therefore that stable structure — whether physical, computational, or cognitive — is the result of a two-stage process. First, local operators generate candidate configurations through the accumulation of constraints. Second, a decimation mechanism filters these configurations according to an admissibility criterion, removing those that cannot be integrated into a coherent whole. The combination of these processes yields a system that is both generative and selective, capable of producing complex structure while maintaining global consistency.

Future work will refine each component of this framework. On the formal side, the variational structure of the RSVP action can be extended to incorporate additional fields and symmetries, and the role of the entropy field can be explored in higher-dimensional settings. On the computational side, the tensor-train implementation can be generalized to higher-dimensional lattices and improved through adaptive truncation strategies. On the theoretical side, the convergence conjecture can be investigated more rigorously, with the goal of establishing conditions under which \mathcal{T} is provably contractive. The broader objective is a unified theory in which computation, cognition, and physical dynamics are understood as manifestations of the same underlying process of constraint closure.

References

- [1] R. T. Grimm, A. J. Staat, and J. D. Eaves, *The Integral Decimation Method for Quantum Dynamics and Statistical Mechanics*, *Quantum*, 10, 2064 (2026). DOI: 10.22331/q-2026-04-13-2064.
- [2] I. V. Oseledets, *Tensor-Train Decomposition*, *SIAM Journal on Scientific Computing*, 33(5), 2295–2317 (2011).
- [3] R. Orús, *A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States*, *Annals of Physics*, 349, 117–158 (2014).
- [4] U. Schollwöck, *The Density-Matrix Renormalization Group in the Age of Matrix Product States*, *Annals of Physics*, 326(1), 96–192 (2011).
- [5] S. R. White, *Density Matrix Formulation for Quantum Renormalization Groups*, *Physical Review Letters*, 69(19), 2863–2866 (1992).
- [6] G. Vidal, *Efficient Simulation of One-Dimensional Quantum Many-Body Systems*, *Physical Review Letters*, 93(4), 040502 (2004).
- [7] K. G. Wilson, *The Renormalization Group: Critical Phenomena and the Kondo Problem*, *Reviews of Modern Physics*, 47(4), 773–840 (1975).
- [8] A. Cichocki et al., *Tensor Networks for Dimensionality Reduction and Large-Scale Optimization*, *Foundations and Trends in Machine Learning*, 9(4-5), 249–429 (2016).
- [9] W. Hackbusch, *Tensor Spaces and Numerical Tensor Calculus*, Springer (2012).
- [10] R. Landauer, *Irreversibility and Heat Generation in the Computing Process*, *IBM Journal of Research and Development*, 5(3), 183–191 (1961).
- [11] C. H. Bennett, *The Thermodynamics of Computation—A Review*, *International Journal of Theoretical Physics*, 21(12), 905–940 (1982).
- [12] S. Lloyd, *Ultimate Physical Limits to Computation*, *Nature*, 406, 1047–1054 (2000).
- [13] E. T. Jaynes, *Information Theory and Statistical Mechanics*, *Physical Review*, 106(4), 620–630 (1957).
- [14] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley (2006).
- [15] T. Jacobson, *Thermodynamics of Spacetime: The Einstein Equation of State*, *Physical Review Letters*, 75(7), 1260–1263 (1995).
- [16] E. Verlinde, *On the Origin of Gravity and the Laws of Newton*, *Journal of High Energy Physics*, 2011(4), 29 (2011).

- [17] K. Friston, *The Free-Energy Principle: A Unified Brain Theory?*, Nature Reviews Neuroscience, 11(2), 127–138 (2010).
- [18] A. Vaswani et al., *Attention Is All You Need*, NeurIPS (2017).
- [19] Y. LeCun, Y. Bengio, and G. Hinton, *Deep Learning*, Nature, 521, 436–444 (2015).
- [20] I. Goodfellow et al., *Generative Adversarial Nets*, NeurIPS (2014).
- [21] S. Mac Lane, *Categories for the Working Mathematician*, Springer (1998).
- [22] J. Lurie, *Higher Topos Theory*, Princeton University Press (2009).
- [23] J. C. Baez and M. Stay, *Physics, Topology, Logic and Computation: A Rosetta Stone*, In New Structures for Physics, Springer (2010).
- [24] A. M. Turing, *On Computable Numbers, with an Application to the Entscheidungsproblem*, Proceedings of the London Mathematical Society, 42(2), 230–265 (1936).
- [25] A. Church, *An Unsolvable Problem of Elementary Number Theory*, American Journal of Mathematics, 58(2), 345–363 (1936).
- [26] C. E. Shannon, *A Mathematical Theory of Communication*, Bell System Technical Journal, 27, 379–423 (1948).