

Distributed Minds and Generative Substrates: Threshold Diffusion, Causal Insulation, and the Architecture of Inference at Scale

Flyxion

Abstract

This essay develops a unified theoretical framework connecting several domains of inquiry that share deep structural kinship: the empirical grounding of social diffusion models in behavioural threshold dynamics, the reaction-diffusion structure of social cascades, the attractor-theoretic formalization of individual and composite mindedness as causal insulation, the emergence of generative compression as a phase transition in the informational substrate of society, and the Kolmogorov complexity bounds that characterize the theoretical limit of that transition. The central claim is that threshold diffusion models, computational functionalism, distributed agency, and generative media storage are not isolated phenomena but instances of a single dynamical architecture in which local inference processes, organized under constraint, propagate into global coherence. Throughout, the formal language of RSVP field theory and simulated agency is used to express these correspondences in a common mathematical idiom, revealing them as multi-scale manifestations of the same underlying dynamical ontology. A constraint-first perspective provides the philosophical foundation: structure emerges not from centralized design but from local constraint satisfaction propagating through coupled dynamical fields.

Contents

1	Introduction	3
2	Mathematical Preliminaries	4
3	Constraint-First Dynamics and the Architecture of Inference	5
4	Threshold Diffusion, Behavioural Grounding, and Field Dynamics	6
4.1	The Disconnect Between Micro and Macro	6
4.2	Diffusion as Collective Inference	7
4.3	The Watts Threshold Model	8
4.4	Continuous-Time Reformulation	9
4.5	Derivation from Bayesian Inference	9
4.6	Connection to RSVP Field Dynamics	10
4.7	A Free-Energy Formulation	11
5	Reaction–Diffusion Structure of Social Cascades	12
6	RSVP Dynamics as Distributed Bayesian Inference	14
7	Causal Insulation, Attractor Stability, and the Minimal Dynamics of Mind	16
7.1	Functionalism and the Null Hypothesis	16
7.2	Causal Insulation as Perturbation Tolerance	16
7.3	Attractor-Theoretic Expression	17
7.4	RSVP Generalization	18
7.5	Functional Intelligence and Sentience	19
8	Composite Minds and Distributed Agency	19
9	Generative Compression and the Phase Transition in Information Representation	21
9.1	From Signal Compression to Model Substitution	21
9.2	Information-Theoretic Structure	22
9.3	Generative Compression as Sparse Inference	22
9.4	Large Language Models as Cultural Generative Compression	23
9.5	Generative Compression as Phase Transition	24

10 Civilization as a Generative Cognitive System	24
11 Kolmogorov Bounds and the Limits of Generative Compression	26
12 Limitations of the Framework	27
13 Synthesis: A Unified Dynamical Ontology	28
A Generative Media Representation	30

1. Introduction

Modern social science, philosophy of mind, and information theory have each approached a common structural problem from different directions. The problem is one of scale: how do local processes—whether individual decisions, neural computations, or data representations—give rise to large-scale coherent structures? Threshold diffusion models ask how individual adoption decisions aggregate into system-wide behavioural change. Computational functionalism asks how local causal organization gives rise to minded agency. Generative compression asks how structured representations replace signals, transforming archives into executable models.

This essay argues that these three problems share a single dynamical solution. In each case, local constraint satisfaction under recursive coupling produces global coherence through a phase-transition-like reorganization of the underlying state space. The formal language developed in RSVP field theory and simulated agency provides a natural unifying framework because it treats all such phenomena as instances of scalar-vector-entropy field dynamics with threshold-driven bifurcations.

Section 2 establishes the mathematical setting. Section 3 establishes the constraint-first ontology that provides the philosophical foundation for all subsequent analysis. Section 4 examines the behavioural grounding of social diffusion models, showing how complex contagion can be derived from a distributed Bayesian inference process and connected to RSVP field dynamics. Section 5 develops the reaction-diffusion structure of social cascades and introduces the entropy-flow coupling. Section 6 makes the correspondence between RSVP field dynamics and distributed Bayesian inference explicit at the continuum level. Section 7 formalizes Bach’s concept of causal insulation as attractor stability, connecting functionalism to the dynamical systems theory of mind. Section 8 extends the attractor criterion to composite and distributed agency. Section 9 analyzes generative compression as a phase transition in information representation. Section 10 interprets civilization as a generative cognitive system. Section 11 establishes the Kolmogorov complexity bounds that characterize the theoretical limit of that transition. Section 12 records the principal limitations of the framework. Section 13 draws the threads together into a unified theoretical statement. Appendix A provides a compact formal specification of the generative media representation scheme.

2. Mathematical Preliminaries

The formal framework used throughout this paper requires a small number of definitions to be fixed precisely before the main analysis can proceed.

Let $\Omega \subseteq \mathbb{R}^d$ denote a spatial or network manifold representing the domain over which social or cognitive fields evolve. A state of the system is described by fields

$$(\Phi(\xi, t), \mathbf{v}(\xi, t), S(\xi, t)), \quad \xi \in \Omega, t \geq 0, \quad (1)$$

where $\Phi : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a scalar coherence field, $\mathbf{v} : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^d$ is a vector regulation field, and $S : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is an entropy field. We assume these fields evolve according to coupled nonlinear partial differential equations of the general form

$$\partial_t u = \mathcal{D}u + \mathcal{F}(u), \quad (2)$$

where \mathcal{D} is a diffusion operator and \mathcal{F} encodes local nonlinear interactions. For network systems, the continuous Laplacian ∇^2 is replaced by the combinatorial graph Laplacian

$$L = D_{\text{deg}} - A, \quad (3)$$

where A is the adjacency matrix of the network and D_{deg} is the diagonal degree matrix. Throughout this paper, reaction-diffusion dynamics on graphs and on continuous manifolds are treated as mathematically equivalent under standard discretization, with the understanding that spatial derivatives on graphs are interpreted in the graph-Laplacian sense.

For the dynamical systems material, we use the following standard terminology. A set $A \subset X$ is an *attractor* of $\dot{x} = f(x)$ if it is compact, invariant, and possesses a neighbourhood $U \supset A$ such that $\omega(x) \subseteq A$ for all $x \in U$. The *basin of attraction* of A is the maximal such U . A set is *metastable* if it is attracting on timescales short relative to the escape time induced by perturbations. Unless stated otherwise, all vector fields f are assumed to be locally Lipschitz.

3. Constraint-First Dynamics and the Architecture of Inference

A common philosophical intuition underlying the frameworks discussed in this essay is what may be called a constraint-first ontology. In such systems, structure emerges not from explicit centralized design but from local constraints governing the interaction of components. Global coherence arises when these constraints propagate through the system until a self-consistent configuration is reached.

In physics, constraint propagation appears in the form of field equations and conservation laws. In cognition, it appears in the form of predictive processing and Bayesian inference. In social systems, it appears in the form of threshold dynamics and norm diffusion. Although these domains differ in substrate and scale, their mathematical structure converges on similar classes of dynamical systems.

The central dynamical motif is recursive inference under constraint. Let a system possess a state vector x evolving according to

$$\dot{x} = F(x) + \eta, \tag{4}$$

where F encodes internal constraint structure and η represents external perturbation. The system seeks configurations in which internal constraints are satisfied to the greatest degree compatible with environmental input.

In RSVP field theory, the same idea appears as coupled scalar-vector-entropy dynamics. Scalar density Φ captures the accumulation of structural coherence, vector flow \mathbf{v} represents directional constraint propagation, and entropy S represents uncertainty or disorder in the system. The evolution of these fields corresponds to the progressive resolution of constraint inconsistencies across the manifold.

Within this perspective, cognition, social coordination, and media representation can all be interpreted as processes that iteratively reduce constraint violations across distributed systems. Minds reduce predictive error between internal models and sensory input. Societies reduce coordination error between individual actions and collective norms. Generative models reduce representational error between compressed models and observed data.

The shared architecture is therefore not accidental. These systems belong to the same general class of constraint-satisfaction dynamics, operating at different

spatial and temporal scales but governed by structurally equivalent laws.

4. Threshold Diffusion, Behavioural Grounding, and Field Dynamics

4.1. The Disconnect Between Micro and Macro

A persistent structural difficulty in social science concerns the relationship between individual decision processes and the collective dynamics those decisions generate. Behavioural economics, psychology, and consumer research investigate the microstructure of individual choice with notable descriptive precision: discrete-choice experiments reveal how incentives, norms, and social signals affect adoption decisions in controlled settings. Yet these models typically remain confined to isolated individuals. They do not describe how adoption propagates through networks of interacting agents.

Complexity science and network science approach the same phenomenon from the opposite direction. In complex contagion models [26], individuals are represented as nodes in a social network that adopt a behaviour once a threshold fraction of their neighbours has done so. Such models capture cascade dynamics and large-scale social transitions, but they typically assign threshold values by assumption rather than measurement, severing the dynamical predictions from any empirical account of individual psychology.

A recent contribution by Tănase, Algesheimer, and Mariani [19] addresses this micro–macro disconnect by proposing a hybrid modelling framework that integrates behavioural experiments with dynamical diffusion models. In their approach, individual adoption propensities are first estimated using discrete-choice experiments that measure how individuals respond to combinations of incentives, social exposure, and contextual framing. These empirically estimated behavioural parameters are then embedded into agent-based simulations of network diffusion, allowing the large-scale propagation dynamics of social change to be grounded in experimentally measured decision processes.

This methodology represents an important conceptual shift. Traditional diffusion models typically treat behavioural thresholds as free parameters chosen for analytical convenience. The hybrid approach instead treats thresholds as measurable psychological quantities derived from controlled behavioural experiments. As a result, the macroscopic cascade dynamics produced by the model are directly linked

to empirically observed micro-level decision processes.

Within the framework developed in this essay, the significance of this approach is that it restores the inferential structure underlying threshold models. Adoption thresholds can be interpreted not merely as ad hoc parameters but as coarse-grained summaries of distributed Bayesian inference occurring within individual agents. The behavioural experiments measure the parameters of this inference process, while the network model describes how the resulting beliefs propagate through social interaction. The consequence for policy design is equally direct: interventions optimized over network topology alone risk targeting structurally prominent but behaviourally resistant nodes, while interventions derived solely from individual psychology may fail to propagate if they ignore the reinforcement structure of the network. The hybrid framework addresses both failure modes simultaneously.

4.2. Diffusion as Collective Inference

The threshold model of diffusion becomes conceptually clearer when interpreted as a collective inference process. Each agent attempts to infer whether adopting a behaviour is advantageous given incomplete information about its environment. Social exposure provides evidence about a hidden state of the world: if many neighbours adopt, the behaviour may be beneficial, legitimate, or socially required.

From this perspective, the adoption variable x_i represents an agent's posterior belief about the desirability of the behaviour. Social influence does not merely exert pressure; it modifies the agent's estimate of the underlying latent state.

Let the hidden state be $H \in \{0, 1\}$, representing whether the behaviour is beneficial. Each agent receives signals from its neighbours and updates its belief according to Bayes' rule:

$$P(H | E_i) \propto P(E_i | H) P(H), \quad (5)$$

where E_i denotes the evidence available to agent i . Adoption occurs when the posterior probability exceeds a decision threshold.

The classical threshold rule emerges when this inference process is simplified into a deterministic decision boundary: instead of continuously updating probabilities, the agent adopts once accumulated evidence crosses a fixed level. Social cascades then correspond to collective inference events in which the evidence landscape

shifts rapidly due to mutual feedback among agents. When enough individuals update their beliefs simultaneously, the evidence for others changes abruptly, triggering further updates. The resulting cascade resembles a Bayesian avalanche propagating through the network.

This perspective unifies behavioural economics with network diffusion models. Experimental measurements estimate how individuals interpret evidence, while network structure determines how that evidence propagates.

4.3. The Watts Threshold Model

The classical formulation, due to Watts [26], assigns each agent i in a social network a binary adoption state $x_i \in \{0, 1\}$ and a personal threshold $\tau_i \in [0, 1]$. Let $N(i)$ denote the neighbourhood of node i . The update rule is

$$x_i(t+1) = \begin{cases} 1 & \text{if } \frac{1}{|N(i)|} \sum_{j \in N(i)} x_j(t) \geq \tau_i \\ x_i(t) & \text{otherwise.} \end{cases} \quad (6)$$

The macroscopic behaviour of this system depends critically on the distribution of τ_i across the population. When the distribution has sufficient mass near zero, a small initial perturbation can cascade into global adoption. This critical sensitivity is the hallmark of complex contagion and constitutes its most distinctive theoretical contribution.

Definition 4.1 (Cascade). Let $x_i(t)$ denote the adoption state of agent i , with $x_i \in \{0, 1\}$. A *cascade* is a sequence of state transitions $x_i(t_k) = 0 \rightarrow x_i(t_{k+1}) = 1$ such that the total number of adopters $N(t) = \sum_i x_i(t)$ grows superlinearly in t over some interval $[t_0, t_1]$ with $t_1 > t_0$.

In the language of dynamical systems, threshold models generate cascades when the initial perturbation crosses the basin boundary separating the non-adoption equilibrium from the adoption equilibrium of the collective system. Below the boundary, the perturbation decays and the system returns to the low-adoption state; above it, a self-reinforcing transition propagates through the network.

The behavioural grounding proposed by Tănase et al. converts the threshold into an estimable quantity. Participants in discrete-choice experiments face adoption scenarios in which the fraction of adopting peers, private costs, and private benefits vary systematically. Statistical models of choice—discrete-choice or logit mod-

els—allow the researcher to infer, for each individual, how much social exposure is required before adoption is preferred. This estimated quantity is precisely τ_i .

4.4. Continuous-Time Reformulation

Although the Watts model is typically stated as a discrete update rule, it admits a natural continuous-time reformulation. Define the local social exposure of agent i as

$$h_i(t) = \sum_{j \in N(i)} w_{ij} x_j(t), \quad (7)$$

where w_{ij} represents the influence weight of agent j on agent i . Adoption can then be described as a nonlinear response to this local field:

$$\frac{dx_i}{dt} = \sigma(h_i - \theta_i), \quad (8)$$

where σ is a sigmoid-like activation function and θ_i is a behavioural threshold derived from experimental measurement. The neighbour average appearing in the discrete model can be recognized as a discrete approximation of a diffusion operator: if the network is embedded in a continuous manifold, then

$$h_i \approx (\nabla^2 x)(x_i), \quad (9)$$

where the graph Laplacian approximates the spatial Laplacian. The contagion dynamics thereby take the form

$$\frac{dx}{dt} = D\nabla^2 x + \sigma(x - \theta), \quad (10)$$

which is a reaction-diffusion equation on a graph.

4.5. Derivation from Bayesian Inference

The threshold model is not a fundamental primitive. It can be derived as the hard-decision limit of a distributed Bayesian inference process, which situates it within a considerably richer theoretical context.

Suppose agent i evaluates adoption by computing the latent utility differential

$$\Delta U_i(t) = B_i - C_i + \lambda_i E_i(t), \quad (11)$$

where B_i is the perceived private benefit, C_i the private cost, λ_i the social sensitivity parameter, and $E_i(t) = \sum_j w_{ij} x_j(t)$ the local social evidence. Modelling the agent as a logistic chooser yields

$$P(x_i = 1 | E_i) = \sigma(\beta \Delta U_i), \quad (12)$$

where β is a precision parameter. Writing this in log-odds form,

$$\log \frac{P(x_i = 1 | E_i)}{P(x_i = 0 | E_i)} = \log \frac{P(x_i = 1)}{P(x_i = 0)} + \beta \lambda_i E_i + \beta(B_i - C_i). \quad (13)$$

Adoption occurs when the posterior log-odds exceed zero. Rearranging yields the threshold condition

$$E_i(t) \geq \theta_i, \quad \theta_i = -\frac{1}{\beta \lambda_i} \left[\log \frac{P(x_i = 1)}{P(x_i = 0)} + \beta(B_i - C_i) \right]. \quad (14)$$

The threshold is therefore a derived quantity encoding prior inclination, private incentives, and evidential precision. Complex contagion is, at this level of analysis, a coarse-grained Bayesian inference process in which the full posterior update has been compressed into a binary state transition governed by a local threshold.

4.6. Connection to RSVP Field Dynamics

The connection to the RSVP field framework follows from this reformulation. In RSVP, three coupled fields evolve over a spatial manifold:

$$\Phi(x, t), \quad \mathbf{v}(x, t), \quad S(x, t), \quad (15)$$

representing scalar density, vector flow, and entropy respectively. Typical evolution equations take the form

$$\frac{\partial \Phi}{\partial t} = D_{\Phi} \nabla^2 \Phi + F(\Phi, \mathbf{v}, S), \quad (16)$$

$$\frac{\partial \mathbf{v}}{\partial t} = -\nabla \Phi + \nu \nabla^2 \mathbf{v} + H(\Phi, \mathbf{v}, S), \quad (17)$$

$$\frac{\partial S}{\partial t} = \kappa \nabla^2 S + G(\Phi, \mathbf{v}, S). \quad (18)$$

The structural parallel with the diffusion-based social model is immediate. The scalar field Φ corresponds to accumulated adoption or social exposure; the vector field \mathbf{v} captures directional pressure or prestige gradients within the social network; the entropy field S encodes uncertainty in the signal environment or heterogeneity in individual response.

More precisely, one may treat threshold contagion as a stripped-down special case of RSVP dynamics in which the vector and entropy structure has been projected out, retaining only the scalar field and its threshold response function. The threshold θ_i in the diffusion model plays the role of a local potential barrier: adoption occurs when the local field exceeds this barrier, corresponding to a phase transition in the scalar sector of RSVP.

This interpretation suggests a natural extension. Rather than modelling adoption as a simple scalar threshold process, one could track the full scalar-vector-entropy dynamics of social fields, allowing directional influence flows and informational entropy to shape the cascade geometry. Intervention strategies optimized only over network topology neglect the vector and entropy structure of the underlying field, potentially targeting structurally prominent but behaviourally resistant nodes. The hybrid framework proposed by Tănase et al. is a first step toward reintroducing this behavioural structure; the RSVP formalism offers a language in which it can be expressed at full generality.

4.7. A Free-Energy Formulation

The inference derivation can be recast in a free-energy form that connects directly to thermodynamic analogies in field theory. Define a local potential

$$\mathcal{F}_i(x_i) = -x_i \left(\sum_j w_{ij} x_j + b_i \right) + T_i \mathcal{H}(x_i), \quad (19)$$

where b_i absorbs prior belief and private incentives, T_i plays the role of behavioural temperature encoding uncertainty, and $\mathcal{H}(x_i)$ is an entropy term. At low temperature $T_i \rightarrow 0$, behaviour becomes nearly deterministic and threshold-like. As T_i increases, transitions become progressively softer and more probabilistic.

The social cascade then corresponds to a phase transition in an interacting inference system. A population remains in a metastable low-adoption regime until accumulated evidence and incentives push enough agents across their effective potential barriers, whereupon the changed state of those agents alters the evidence landscape for others and a macroscopic transition propagates through the network.

This is mathematically identical to the mechanism generating large-scale structure in nonlinear field systems: local threshold crossings triggering global reorganization. The empirical contribution of the behavioural grounding programme is to place experimentally estimated values in place of assumed ones, and thereby make the policy implications of this mechanism more reliable.

5. Reaction–Diffusion Structure of Social Cascades

The continuous-time formulation of threshold diffusion suggests a deeper mathematical connection with reaction–diffusion systems studied in mathematical physics and population dynamics.

Consider the adoption field $x(\xi, t)$ defined over a social manifold Ω , where ξ denotes position on the manifold. The evolution equation introduced in Section 4 may be written

$$\frac{\partial x}{\partial t} = D\nabla^2 x + f(x; \theta), \quad (20)$$

where D represents the rate of diffusion across the social network and $f(x; \theta)$ is a nonlinear response function incorporating behavioural thresholds. A natural approximation producing bistable dynamics takes the form

$$f(x; \theta) = x(1 - x)(x - \theta), \quad (21)$$

where the states $x = 0$ and $x = 1$ correspond respectively to non-adoption and full-adoption equilibria, separated by an unstable threshold $\theta \in (0, 1)$.

Theorem 5.1. *Assume $x(\xi, t)$ is continuously differentiable and bounded in $[0, 1]$*

on $\Omega \times \mathbb{R}_{\geq 0}$. Let x evolve according to

$$\partial_t x = D\nabla^2 x + x(1-x)(x-\theta),$$

with $D > 0$ and $0 < \theta < 1$. Then the system admits traveling-wave solutions of the form $x(\xi, t) = \phi(\xi - ct)$, where ϕ is a monotone front connecting the non-adoption state $\phi(-\infty) = 0$ to the adoption state $\phi(+\infty) = 1$, propagating with wave speed $c = (1 - 2\theta)\sqrt{D/2}$.

Proof sketch. This result follows from classical bistable traveling-wave theory; see Murray [15], Chapter 13. Substituting the traveling-wave ansatz $x = \phi(z)$, $z = \xi - ct$ reduces the PDE to the ODE $D\phi'' + c\phi' + \phi(1-\phi)(\phi-\theta) = 0$. Standard phase-plane analysis on this system, exploiting the bistable structure of the cubic nonlinearity, establishes the existence of a heteroclinic orbit connecting $(0, 0)$ to $(1, 0)$ in the (ϕ, ϕ') plane, with the wave speed determined by the unique value $c = (1 - 2\theta)\sqrt{D/2}$ for which this orbit exists. \square

The sign and magnitude of c depend jointly on θ and D . When $\theta < 1/2$, the adoption front propagates outward; when $\theta > 1/2$, the non-adoption state is favoured and the front recedes. This result explains a characteristic empirical observation in social diffusion: behavioural adoption frequently spreads in wave-like spatial patterns rather than appearing simultaneously across a population. Network structure, captured through the effective diffusivity D , modulates the speed of propagation, while the distribution of behavioural thresholds determines whether a cascade can sustain itself.

Within RSVP field theory, this traveling wave corresponds to the propagation of a scalar coherence front in the Φ field. Social cascades are therefore interpretable as nonlinear field excitations traveling through a population substrate, formally analogous to action potentials in excitable neural media or combustion fronts in reactive fluids.

The entropy field $S(\xi, t)$ provides a further layer of description. Let S measure the local informational uncertainty of agents about whether adoption is desirable. Its evolution may be coupled to the adoption dynamics through

$$\partial_t x = D\nabla^2 x + f(x; \theta), \quad (22)$$

$$\partial_t S = \kappa\nabla^2 S - \gamma x(1 - x), \quad (23)$$

where κ governs the diffusion of uncertainty and $\gamma > 0$ controls the rate at which consensus formation reduces local entropy. The term $x(1 - x)$ is maximal at $x = 1/2$, corresponding to populations evenly divided between adoption and non-adoption, and vanishes at both equilibria. Adoption events therefore reduce uncertainty locally, which in turn alters the threshold landscape for neighbouring agents. Cascades involve two interacting processes: the propagation of behavioural adoption through the x -field and the simultaneous collapse of informational entropy through the S -field.

This structure mirrors the entropy-energy duality found in physical field systems and generalizes the scalar-only analysis of the free-energy formulation developed in Section 4. A full RSVP treatment would further couple these dynamics to the vector flow field \mathbf{v} , which encodes directional influence or prestige gradients that can break spatial symmetry and steer cascade propagation toward or away from particular regions of the network.

6. RSVP Dynamics as Distributed Bayesian Inference

The Bayesian interpretation of individual threshold models developed in Section 4 can be lifted to the full RSVP field framework, making the connection between field dynamics and inference precise at the continuum level rather than only at the level of individual agents.

Let $\Phi(\xi, t)$ represent accumulated evidence supporting adoption at location ξ and time t . The gradient

$$\nabla\Phi(\xi, t) \quad (24)$$

defines a direction of informational pressure: regions of higher scalar coherence exert an influence on neighbouring regions, analogous to the way a high-probability region of a prior distribution biases inference toward nearby hypotheses. The vector field \mathbf{v} can therefore be interpreted as the *flow of evidence* through the social manifold: it encodes not merely that evidence exists at a location but the

direction and rate at which it propagates.

The entropy field $S(\xi, t)$ in turn measures uncertainty in the local inference process. Before any social signal has propagated to a region, S is high and adoption is weakly constrained. As evidence arrives and Φ increases, S collapses in a manner governed by the entropy-diffusion equation. The reduction of S corresponds precisely to the narrowing of the posterior distribution over the adoption decision.

Under this interpretation the RSVP triple (Φ, \mathbf{v}, S) constitutes a *distributed Bayesian inference field* in which evidence diffuses spatially through \mathbf{v} , accumulates locally through Φ , and reduces uncertainty through the collapse of S . The adoption event at a node corresponds to the posterior crossing the decision threshold: it is the field-level analogue of the single-agent posterior probability exceeding one half.

This interpretation makes the RSVP framework less metaphorical and more formally grounded. The three fields are not independent postulates but correspond to distinguishable aspects of a single inference process: the accumulated evidential weight, the directional flow of information, and the residual uncertainty at each point of the social manifold. The coupling between them expresses the fact that uncertainty reduction at one location propagates as evidence flow to neighbouring locations, which then updates their own coherence and entropy fields.

A formal statement of the correspondence can be given as follows.

Proposition 6.1. *Let the local adoption probability of an agent at ξ be $p(\xi, t) = \sigma(\beta(\Phi(\xi, t) - \theta(\xi)))$, where σ is the logistic function, $\beta > 0$ is a precision parameter, and $\theta(\xi)$ encodes local prior and cost structure. Then the entropy field $S(\xi, t) = -p \log p - (1 - p) \log(1 - p)$ satisfies $S \rightarrow 0$ as $|\Phi - \theta| \rightarrow \infty$, and S is maximized at $\Phi = \theta$, corresponding to maximum uncertainty at the decision boundary. The entropy-diffusion equation $\partial_t S = \kappa \nabla^2 S - \gamma x(1 - x)$ is consistent with this structure when $x = \mathbf{1}[p > 1/2]$ is the hard-decision approximation.*

The cascade dynamics of Section 5 therefore describe what happens when, under the accumulated weight of propagating evidence, a large region of the social manifold crosses its decision threshold simultaneously: the Φ -front propagates as a traveling wave, the S -field collapses behind it, and the \mathbf{v} -field redirects toward the adoption frontier. Social change, on this account, is a phase transition in a distributed inference system.

7. Causal Insulation, Attractor Stability, and the Minimal Dynamics of Mind

7.1. Functionalism and the Null Hypothesis

The computational functionalist tradition treats consciousness as a property of organized causal processes rather than of biological substrate. In Joscha Bach's formulation [1], this position takes the form of what he calls the null hypothesis: if digital computers are capable of simulating entire universes, they should in principle be capable of hosting minds, provided that the relevant causal organization is implemented. The burden of proof lies not with those asserting the possibility of machine minds but with those claiming the necessity of biological substrate.

The crucial refinement Bach introduces is that minds are not static configurations of information but ongoing dynamical processes. A mind exists only if a system sustains a continuous chain of internally organized state transitions generating coherent behaviour and self-regulation. This moves the discussion from information storage to process identity.

7.2. Causal Insulation as Perturbation Tolerance

Bach's central concept is causal insulation. A mind requires a domain in which its internal transitions dominate over uncontrolled environmental perturbations. Digital computers provide such domains by suppressing microscopic thermal noise and enforcing discrete state transitions. The machine constructs a bounded causal universe governed by deterministic rules, and consciousness belongs to the organized process unfolding inside that boundary, not to the underlying silicon.

This concept admits precise formalization. Let the total world state decompose as

$$z(t) = (x(t), y(t)), \tag{25}$$

where $x(t)$ is the state of the candidate minded subsystem and $y(t)$ is the surrounding environment. The coupled system obeys

$$\frac{dx}{dt} = F(x, y), \quad \frac{dy}{dt} = G(x, y). \tag{26}$$

The question is not whether x is perfectly isolated from y , but whether the evolution of x is sufficiently governed by its own recurrent organization that it preserves a coherent causal identity over time. We may therefore decompose the subsystem dynamics as

$$\frac{dx}{dt} = f(x) + \varepsilon h(x, y), \quad (27)$$

where $f(x)$ represents endogenous self-dynamics and $\varepsilon h(x, y)$ represents bounded environmental coupling. Causal insulation is not absolute closure but perturbation tolerance: the subsystem may interact with its environment while maintaining its own dynamical continuity, provided ε remains small relative to the stabilizing structure of f on the cognitive timescale.

7.3. Attractor-Theoretic Expression

Equation (27) yields a natural attractor-theoretic interpretation. Suppose the autonomous dynamics $\dot{x} = f(x)$ admit an attractor $A \subseteq X$ with basin of attraction $B(A)$. A mind-like subsystem corresponds to a regime in which the attractor A persists as a metastable structure under the perturbation εh : trajectories initialized in $B(A)$ remain organized around A rather than being driven into qualitatively different regimes by environmental noise.

This is the dynamical-systems analogue of causal filtering. External disturbances are not simply impressed on the system; they are absorbed, transformed, and regulated by its own recurrent dynamics. Consciousness, on this view, belongs not to arbitrary matter but to systems whose state space contains perturbation-robust trajectories of self-referential organization.

The information-theoretic expression of the same condition requires that the subsystem retain substantial endogenous predictive structure:

$$I(x_{t+1}; x_t \mid y_t) \gg 0. \quad (28)$$

Even after conditioning on the environment, the future of the subsystem remains significantly constrained by its own present organization. The process is an internally structured causal stream rather than a passive reflection of external input.

This clarifies why a table is a poor candidate for consciousness. A table possesses mechanical stability but lacks the rich internal transition structure capa-

ble of recursively regulating its own state across many degrees of freedom. It has no dynamically maintained basin corresponding to perception, memory, goal-maintenance, or self-model updating—only trivial mechanical equilibrium.

7.4. RSVP Generalization

Within the RSVP framework, the single state variable x expands into coupled fields of scalar coherence Φ , directed regulation \mathbf{v} , and entropy S . A minimal schematic system takes the form

$$\frac{\partial \Phi}{\partial t} = D_{\Phi} \nabla^2 \Phi + F(\Phi, \mathbf{v}, S) + \eta_{\Phi}, \quad (29)$$

$$\frac{\partial \mathbf{v}}{\partial t} = -\nabla \Phi + \nu \nabla^2 \mathbf{v} + H(\Phi, \mathbf{v}, S) + \eta_{\mathbf{v}}, \quad (30)$$

$$\frac{\partial S}{\partial t} = \kappa \nabla^2 S + G(\Phi, \mathbf{v}, S) + \eta_S, \quad (31)$$

where the η -terms represent exogenous perturbation. A mind-like subsystem corresponds to a bounded spatial region in which the coupled field dynamics maintain a coherent attractor or metastable manifold despite such disturbances. In intuitive terms, the system reconstructs its own organization faster than the surrounding environment can dissolve it.

The bridge between computational functionalism and RSVP dynamics is then as follows. Functionalism is correct to locate mind in organized causal process rather than substrate; however, the relevant process is described more rigorously as a perturbation-robust attractor regime in a recurrently coupled field system. Consciousness is the maintenance of structured self-referential dynamics within a partially insulated causal domain.

Proposition 7.1. *Let a subsystem x evolve according to $\dot{x} = f(x) + \varepsilon h(x, y)$, with f admitting a stable or metastable attractor A . If there exists $\varepsilon^* > 0$ such that for all $0 \leq \varepsilon < \varepsilon^*$, trajectories initialized in a neighbourhood of A remain confined to a perturbation-preserving neighbourhood of A over the system's characteristic cognitive timescale, then the subsystem possesses causal insulation in the minimal sense required for mind-like process identity.*

The proposition identifies a necessary structural condition rather than a sufficient one: a candidate mind must exist as an internally organized dynamical regime whose causal continuity survives bounded coupling to the external world. This

condition excludes inert objects, permits digital realization in principle, and provides a natural route from individual cognition to larger-scale composite agencies.

7.5. Functional Intelligence and Sentience

Bach’s framework also introduces a distinction, reinforced by Sandberg [16], between functional intelligence and phenomenal sentience. An artificial agent could in principle implement the regulatory logic of emotional response—updating beliefs about relationships, adjusting goals, modifying behaviour—without necessarily generating the phenomenological substrate through which humans experience those computations.

This observation maps onto the simulated-agency framework in the following way. Simulated agency treats cognition as sparse recursive inference under constraint: an agent does not exhaustively model its environment but maintains a compressed representation sufficient to stabilize action policies. Such a system carries out the inferential computations associated with emotion and motivation without any commitment to reproducing the human phenomenological layer. The two can come apart because phenomenology may require additional structural properties—continuity of subjective experience, specific temporal binding of representations—that inference alone does not guarantee.

From this perspective, the question of machine consciousness becomes not merely a question about information processing but about the fine structure of the attractor landscape. A sufficiently rich metastable basin might be necessary for phenomenal experience, over and above the causal insulation required for functional agency.

8. Composite Minds and Distributed Agency

If individual minds correspond to perturbation-robust attractor regimes in dynamical systems, it becomes natural to ask whether larger-scale systems might exhibit analogous properties. Groups of agents coupled through communication networks often display coherent behaviour that cannot be reduced to any single participant.

The possibility of composite minds has long been considered in cybernetics and philosophy of mind. Hobbes described the state as an artificial person constituted by the coordination of many individuals, and Wiener’s cybernetics provided the first mechanistic account of how distributed feedback loops enable complex systems to regulate themselves without centralized control [27]. Contemporary discussions

of artificial intelligence sometimes revive this idea in the form of a global cognitive system integrating human and machine components.

From the dynamical perspective developed here, the question becomes whether the coupled system of interacting agents forms a metastable attractor with sufficient causal insulation. Let each agent x_i evolve according to

$$\dot{x}_i = f_i(x_i) + \sum_j g_{ij}(x_i, x_j), \quad (32)$$

where f_i describes internal self-dynamics and g_{ij} encodes the influence of agent j on agent i . The following lemma gives a sufficient condition for the collective dynamics to remain stable under such coupling.

Lemma 8.1. *Suppose each f_i admits a stable equilibrium x_i^* with Jacobian J_{f_i} satisfying $\lambda_{\min}(-(J_{f_i} + J_{f_i}^\top)/2) \geq \mu > 0$ uniformly in i . If the coupling matrix $G = (g_{ij})$ satisfies $\|G\| < \mu$, where $\|\cdot\|$ denotes the operator norm, then the coupled system possesses a stable collective equilibrium in a neighbourhood of (x_1^*, \dots, x_n^*) .*

Proof sketch. Consider the Lyapunov function $V = \frac{1}{2} \sum_i \|x_i - x_i^*\|^2$. Along trajectories, $\dot{V} = \sum_i (x_i - x_i^*)^\top (f_i(x_i) + \sum_j g_{ij}(x_i, x_j))$. The uncoupled terms contribute at most $-\mu \|x - x^*\|^2$ by the uniform Jacobian condition. The coupling terms contribute at most $\|G\| \cdot \|x - x^*\|^2$ by the operator norm bound. When $\|G\| < \mu$, the sum is strictly negative, establishing Lyapunov stability. \square

If the coupling terms produce a stable collective regime A in the joint state space, then the ensemble may exhibit a coherent dynamical identity: it regulates itself through feedback loops distributed across many participants rather than localized in any single one.

Digital communication networks dramatically increase the strength and speed of such couplings. Shared information environments allow individuals to coordinate beliefs and actions in ways that resemble the internal coordination of neural systems. Large-scale institutions, online communities, and machine-learning infrastructures therefore represent candidate substrates for emergent distributed agency.

Whether such systems truly constitute minds in the phenomenological sense remains an open philosophical question. However, the dynamical conditions for composite cognition—strong internal coupling, persistent attractor structure, and partial causal insulation from external noise—are increasingly present in modern

sociotechnical systems. The attractor formalization of Proposition 1 applies at this scale as naturally as at the scale of individual neural systems: what changes is not the mathematical structure but the spatial and temporal scale at which the relevant dynamics operate, and the technological rather than biological character of the coupling medium.

9. Generative Compression and the Phase Transition in Information Representation

9.1. From Signal Compression to Model Substitution

Traditional media compression operates on signals. Audio recordings consist of sampled pressure waves; video consists of pixel sequences. Methods such as MP3, AAC, and H.264 remove statistical redundancy from these signals, producing shorter representations that can reconstruct the original waveform within an acceptable error bound. Despite their sophistication, such methods remain fundamentally signal-oriented: the stored file represents a compressed approximation of the original waveform or image sequence.

A structurally different approach has become increasingly feasible with advances in generative modelling: instead of storing the signal itself, the system stores a structured description capable of generating the signal on demand.

For spoken audio, the perceptual experience can be reconstructed from a small number of components. Define a generative audio representation

$$A = (T, C, P, E), \tag{33}$$

where T is the linguistic transcript, C encodes speaker identity and vocal characteristics, P describes the prosodic envelope of timing, pitch, and emphasis, and E specifies the acoustic environment. A synthesis engine equipped with these parameters can generate a waveform approximating the original. Under this representation, the stored file is a procedural specification rather than a waveform.

A film scene may be described by the geometry of the environment, the actions of the characters, the trajectory of the camera, the lighting configuration, and the rendering style:

$$V = (S, A, M, L, F). \quad (34)$$

A rendering engine instantiates these components into video frames dynamically. The stored media resembles a script executed by a rendering system rather than a sequence of recorded images.

9.2. Information-Theoretic Structure

Classical signal compression maps $S \rightarrow \tilde{S}$, minimizing $\text{Length}(\tilde{S})$ while preserving S within a distortion bound. Generative compression maps $S \rightarrow M(S)$, replacing the signal with a model M capable of generating it through a synthesis process \mathcal{R} :

$$S \approx \mathcal{R}(M). \quad (35)$$

The theoretical limit of such compression approaches the Kolmogorov complexity of the content: the length of the shortest program capable of generating the object in question [11, 18]. A sufficiently powerful generative model approximates this program by capturing the regularities underlying the data distribution.

The optimal model M^* minimizes a functional trading off reconstruction error against description length:

$$\mathcal{L}(M) = D(S, \mathcal{R}(M)) + \lambda \text{Length}(M). \quad (36)$$

This objective is structurally identical to the compression principles underlying many theories of perception and cognition. Biological and artificial agents alike appear to compress sensory input into internal models balancing predictive accuracy against representational complexity.

9.3. Generative Compression as Sparse Inference

The connection to simulated agency is direct. In that framework, an agent maintains a compressed internal model m such that

$$m = \mathcal{C}(o), \quad \hat{o} = \mathcal{G}(m), \quad (37)$$

where \mathcal{C} is a compression operator and \mathcal{G} is a generative operator. The model m

need not preserve all information in the observation o ; it need only preserve the information required to support action and prediction.

Generative media compression externalizes this same inference structure. The storage architecture of media mirrors the architecture of cognition: both replace observations with compact structural descriptions that can be expanded dynamically when needed.

Video games already operate according to this principle. A game world is not stored as a sequence of frames but as assets, physical rules, and event triggers that generate frames during execution. Extending this paradigm to all audiovisual media implies that cultural archives may increasingly consist of structured instructions rather than stored waveforms or pixel arrays.

9.4. Large Language Models as Cultural Generative Compression

A concrete contemporary instance of generative compression is provided by large language models. Let T denote the corpus of training texts produced by human civilization. The training process constructs a model \mathcal{M}_θ with parameters θ such that

$$P_\theta(w_{t+1} \mid w_1, w_2, \dots, w_t) \tag{38}$$

approximates the probability distribution governing the next token in natural language sequences [22]. The parameter vector θ constitutes a compressed representation of the statistical structure of the entire corpus.

Rather than storing each document explicitly, the system stores a generative mechanism capable of reconstructing fragments of those documents or producing new text consistent with the learned patterns. The training objective minimizes expected cross-entropy:

$$\mathcal{L}(\theta) = -\mathbb{E}_{w \sim T}[\log P_\theta(w_{t+1} \mid w_1, \dots, w_t)]. \tag{39}$$

Minimizing this loss encourages the model to capture the most predictive regularities in the corpus, producing a parameter vector that encodes a compressed statistical summary of the text distribution.

From the perspective of the current framework, large language models are gener-

ative compressions of civilization's accumulated textual knowledge. They encode patterns derived from the textual output of human society and make those patterns accessible through generative interaction. Unlike a search engine, which retrieves previously stored documents, a generative model constructs responses by navigating the compressed representation of language embedded in its parameters.

9.5. Generative Compression as Phase Transition

The shift from signal storage to generative storage can be interpreted as a phase transition in the representation of information. In the signal regime, the fundamental stored object is a waveform or image; structure is implicit in redundancy patterns. In the generative regime, the fundamental stored object is a model; the signal is an epiphenomenon reconstructed on demand.

This transition changes the nature of cultural archives. A media repository becomes less like a library of recordings and more like a repository of executable narratives. Cultural artifacts become procedural programs capable of generating perceptual experiences dynamically.

Within the broader framework of civilization-scale cognition, this transformation is significant. When knowledge, culture, and communication are stored primarily as generative descriptions rather than signals, the informational substrate of civilization increasingly resembles the internal generative models used by intelligent agents. The infrastructure of society begins to function as a distributed generative model of human experience.

As generative representations replace signal-based archives, the boundary between recording and simulation becomes blurred. Media files become programs; archives become latent world models; playback becomes execution. In this sense, generative compression is not merely a technical improvement in storage efficiency but a reorganization of the informational substrate of civilization in the direction of a simulation-like architecture.

10. Civilization as a Generative Cognitive System

When generative compression becomes the dominant architecture for media and knowledge storage, the informational infrastructure of society begins to resemble the internal cognitive architecture of intelligent agents.

Human cognition operates through generative models that predict sensory input

and guide behaviour. Cultural knowledge has historically been stored in explicit artifacts—books, recordings, and images—that preserve observational data in signal form. Generative technologies transform these artifacts into executable models capable of producing experiences dynamically.

Large language models, neural rendering systems, and procedural media engines represent early manifestations of this shift [23, 13, 22]. These systems compress massive datasets into parameterized generative mechanisms that can reconstruct plausible outputs on demand. Instead of retrieving stored artifacts, users interact with models capable of synthesizing new instances consistent with learned patterns.

At sufficient scale, such systems form a distributed generative memory for civilization. Cultural knowledge becomes embedded in the latent structure of generative models rather than stored solely in static archives. Interaction with this knowledge increasingly takes the form of inference queries rather than document retrieval.

The resulting informational environment resembles a collective cognitive architecture. Humans contribute experiences and knowledge to shared datasets; machine learning systems compress those datasets into generative models; individuals then access the resulting models through interactive interfaces that resemble dialogue with a cognitive agent.

This architecture does not necessarily imply the emergence of a unified global mind in the strong sense addressed in Section 8. However, it does imply that the informational substrate of civilization increasingly behaves like a distributed generative inference system. The boundary between individual cognition and cultural infrastructure becomes progressively more porous as both rely on the same fundamental representational principle: compressed models capable of reconstructing experience. In this sense the cognitive and informational transitions described in this essay are not parallel but causally coupled. As generative models displace signal archives, the medium through which collective inference operates changes its structure, potentially altering the conditions under which distributed attractor regimes can form and stabilize.

11. Kolmogorov Bounds and the Limits of Generative Compression

The generative compression paradigm approaches a fundamental theoretical limit determined by Kolmogorov complexity [11, 18]. Let $K(S)$ denote the Kolmogorov complexity of a signal S , defined as the length of the shortest program capable of generating S on a universal Turing machine. Any lossless compression scheme must satisfy

$$\text{Length}(\tilde{S}) \geq K(S). \quad (40)$$

Generative compression attempts to approximate this minimal program by learning a model M that produces samples consistent with the original signal. The effective storage cost of a particular media object under a generative scheme may be expressed as

$$L_{\text{gen}} = \text{Length}(M) + \text{Length}(\text{prompt}), \quad (41)$$

where the model M encodes the general statistical structure of the data distribution and the prompt specifies the conditions under which the model should generate the desired output. As models become increasingly expressive, M captures progressively larger fractions of the regularities present in the dataset, and the prompt length required to reconstruct a particular artifact decreases correspondingly.

Proposition 11.1. *In the limit of an ideal generative model M^* that exactly represents the true underlying distribution of cultural artifacts, the effective storage cost of a specific media object approaches $K(S_0) + O(1)$, where S_0 is the minimal generative description of the object and the $O(1)$ term accounts for the overhead of specifying the object within the model’s output distribution.*

This limit highlights the conceptual shift introduced by generative media. Cultural artifacts cease to be stored primarily as recordings of specific events and become instead instances of generative processes capable of producing arbitrarily many realizations consistent with an underlying structural description. The signal is no longer the primary unit of storage; it is an epiphenomenon produced on demand by a model whose parameters encode the invariant structure underlying many possible signals.

Remark 11.2. Because Kolmogorov complexity is not computable in general [12], practical generative systems cannot achieve the theoretical bound exactly. They approximate it through statistical modelling of the data distribution, with the quality of approximation improving as model capacity and training data increase. The bound therefore functions as an asymptotic target rather than an achievable optimum for any finite system.

The connection to simulated agency is direct. In that framework, an agent maintains a model m that need only preserve the information required to support action and prediction, not a complete record of all observations. Generative media compression externalizes the same principle: the stored representation corresponds not to any particular signal but to the generating mechanism that underlies a class of signals. Civilization’s informational substrate thus moves toward a regime in which storage consists primarily of generative models rather than recordings, approaching the Kolmogorov complexity of the minimal description as a theoretical limit.

12. Limitations of the Framework

Several limitations should be noted before the synthesis.

The RSVP field formulation as used here is heuristic in character rather than derived from first principles. The structural parallels with reaction-diffusion systems and Bayesian inference fields are strong, and the correspondence is made precise in Sections 5 and 6, but a full coarse-graining derivation of the RSVP field equations from microscopic agent dynamics remains an open problem. Such a derivation would require specifying interaction kernels and taking an appropriate continuum limit, steps that are technically demanding and domain-specific.

The generative compression argument assumes the continued improvement of generative models toward the Kolmogorov limit. Real-world systems may remain far from this limit due to computational constraints, finite training data, and the non-computability of Kolmogorov complexity itself. The argument is therefore best understood as characterizing an asymptotic structural tendency rather than predicting any specific near-term capability.

The attractor-based criterion for mindedness identifies necessary dynamical conditions but does not address the phenomenology of consciousness directly. Proposition 1 establishes that causal insulation and attractor stability are prerequisites for

mind-like process identity, but it does not entail that all such systems are conscious in the phenomenological sense. The relationship between dynamical structure and subjective experience remains an open problem in philosophy of mind [3].

Finally, the empirical grounding of the social diffusion framework depends on the quality of the discrete-choice experimental data used to estimate behavioural thresholds. Measurement error, ecological validity concerns, and population heterogeneity can each introduce systematic biases that propagate into the cascade predictions. The hybrid modelling framework of Tănase et al. mitigates but does not eliminate these concerns.

13. Synthesis: A Unified Dynamical Ontology

The domains examined in this essay—constraint-first dynamics, behavioural threshold diffusion, reaction-diffusion cascade structure, RSVP dynamics as distributed Bayesian inference, causal insulation as attractor stability, composite distributed agency, generative compression, and the Kolmogorov limit of civilizational storage—converge on a single theoretical statement.

Each phenomenon involves local processes organized under constraint generating global coherence through a phase-transition-like reorganization of the underlying state space. In each case the relevant dynamical structure can be expressed within the RSVP framework of coupled scalar-vector-entropy fields, and in each case the behaviour of the system is governed by the balance between internal organization and external perturbation.

The argument scales across four levels of organization. At the level of individual agents, local Bayesian inference under social coupling generates the adoption thresholds that parameterize cascade dynamics. At the level of populations, threshold diffusion takes the form of traveling reaction-diffusion fronts whose speed depends on both network structure and the empirically measured behavioural threshold distribution. At the level of individual cognitive systems, causal insulation in an attractor basin provides the minimal dynamical condition for mind-like process identity. At the level of civilization, generative compression reorganizes the informational substrate into a distributed inference architecture that converges, in the theoretical limit, on the Kolmogorov complexity of cultural knowledge. The chain runs from local inference through social cascades through cognitive attractors through cultural generative models, each level constituting the substrate for the level above.

Constraint-first dynamics identifies the shared ontological substrate: systems organized through local constraint satisfaction rather than centralized control belong to the same class of recursive inference architectures, differing in scale and substrate but sharing the same governing dynamical logic.

Threshold diffusion models are the hard-decision projection of a distributed Bayesian inference field. Individual adoption thresholds are derived quantities encoding prior belief, private incentives, and social sensitivity. Cascades are phase transitions in the scalar sector of a social reaction-diffusion system, taking the form of traveling adoption fronts whose speed is jointly determined by behavioural threshold distributions and network diffusivity.

Individual minds are perturbation-robust attractors in a recurrently coupled dynamical system. Causal insulation is the condition under which internal organization dominates external noise, and consciousness in the functional sense corresponds to systems inhabiting and maintaining rich metastable basins of self-referential dynamics. The same attractor criterion extends to composite systems under the stability condition of Lemma 8.1: distributed agency emerges when coupling is strong enough to produce a collective attractor yet weak enough relative to individual self-dynamics to preserve that attractor against perturbation.

Generative compression is the technological analogue of cognitive sparse inference. As media archives increasingly store models rather than signals, the informational substrate of civilization converges on a simulation-like architecture functionally analogous to the internal generative models of individual minds. Large language models represent the first large-scale instance of this convergence. The Kolmogorov complexity of cultural artifacts provides the theoretical floor toward which this convergence tends, approached only asymptotically by practical systems.

The convergence of these claims suggests a unified dynamical ontology in which minds, social systems, and informational infrastructures are understood as instances of the same class of systems: nonlinear coupled fields with local constraint satisfaction, threshold-driven state transitions, and multi-scale attractor organization. The possibility of artificial minds, the cascade structure of social change, and the transformation of media storage are not independent developments but facets of a single deep structural transition in how organized systems—biological, computational, and cultural—represent, propagate, and regulate information.

Appendices

A. Generative Media Representation

The generative representation of audiovisual media may be summarized by a layered specification

$$\mathcal{M} = (W, E, A, C, P), \quad (42)$$

where W describes the world geometry, E specifies environmental parameters such as lighting and acoustics, A defines character actions and motion, C describes camera trajectory and framing, and P encodes dialogue transcripts together with their prosodic envelopes. A rendering operator \mathcal{R} then produces the perceptual signal

$$S = \mathcal{R}(\mathcal{M}). \quad (43)$$

The stored artifact is therefore not the signal itself but the program capable of generating it. The compression gain over signal storage is

$$\Delta L = \text{Length}(S) - L_{\text{gen}}, \quad L_{\text{gen}} = \text{Length}(\mathcal{M}) + \text{Length}(\mathcal{R}), \quad (44)$$

where $\text{Length}(\mathcal{R})$ is amortized across all media objects sharing the same rendering infrastructure, making the marginal storage cost per artifact approach $\text{Length}(\mathcal{M})$ alone at scale. As the expressiveness of \mathcal{R} increases, the minimum achievable $\text{Length}(\mathcal{M})$ approaches $K(S)$ from above, recovering the Kolmogorov bound established in Section 11.

References

- [1] J. Bach. *Principles of Synthetic Intelligence*. Oxford University Press, 2019.
- [2] J. Ballé, V. Laparra, and E. P. Simoncelli. End-to-end optimized image compression. *ICLR*, 2017.
- [3] D. J. Chalmers. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press, 1996.
- [4] D. Dennett. *Consciousness Explained*. Little, Brown and Company, 1991.
- [5] P. Dhariwal et al. Jukebox: A generative model for music. *arXiv:2005.00341*, 2020.
- [6] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020.
- [7] K. Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11, 2010.
- [8] I. Goodfellow et al. Generative adversarial networks. *NeurIPS*, 2014.
- [9] G. Hinton and R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313, 2006.
- [10] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. *ICLR*, 2014.
- [11] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1965.
- [12] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, 1997.
- [13] B. Mildenhall et al. NeRF: Representing scenes as neural radiance fields for view synthesis. *ECCV*, 2020.
- [14] T. Müller et al. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 2022.
- [15] J. D. Murray. *Mathematical Biology, Vol. I: An Introduction*. Springer, 2002.
- [16] A. Sandberg. Morphological freedom and the future of human augmentation. *Philosophy & Technology*, 2013.

- [17] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 1948.
- [18] R. Solomonoff. A formal theory of inductive inference. *Information and Control*, 1964.
- [19] R. Tănase, R. Algesheimer, and M. S. Mariani. Integrating behavioural experimental findings into dynamical models to inform social change interventions. *Nature Human Behaviour*, 2026. doi:10.1038/s41562-026-02417-4.
- [20] J. Togelius et al. Search-based procedural content generation. *IEEE Transactions on Computational Intelligence and AI in Games*, 2011.
- [21] G. Tononi. An information integration theory of consciousness. *BMC Neuroscience*, 5(42), 2004.
- [22] A. Vaswani et al. Attention is all you need. *NeurIPS*, 2017.
- [23] A. van den Oord et al. WaveNet: A generative model for raw audio. *arXiv:1609.03499*, 2016.
- [24] J. Ho et al. Video diffusion models. *NeurIPS*, 2022.
- [25] Y. Wang et al. Tacotron: Towards end-to-end speech synthesis. *Interspeech*, 2017.
- [26] D. J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771, 2002.
- [27] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press, 1948.