Constraint, Reachability, and Relaxation

Across Biological, Cognitive, Industrial, and Cosmological Systems

---

A Theoretical Monograph

Flyxion

June 2026

Incorporates analysis of six papers published in *Nature Neuroscience*, *Nature Immunology*, *Nature Methods*, *Nature Sustainability*, *Nature Astronomy*, and *Scientific Reports*, June 2026.

The relevant question is never *what is the system doing*
but always *what can it still do*.

Admissibility geometry, first principle

Sleep pressure is the integral of a deficit.
Off-period induction is the path that clears it.
Tonic suppression is the path that does not.

After Driessen et al. (2026)

The PIP joint is not more vulnerable because of what it is now.
It is more vulnerable because of what it was.

After Davidson et al. (2026)

Contents

Preface

This monograph originated as a comparative reading of six papers published simultaneously in June 2026. The six span neuroscience (Driessen et al., 2026), immunology (Davidson et al., 2026), genomics (Spiro et al., 2026), sustainability science (Tian et al., 2026), cognitive science (Okitsu and Sakai, 2026), and astrophysics (Leung et al., 2026). Their simultaneous appearance is coincidental; the structural observation that links them is not.

Two results forced the central claim into a sharper form than it held at the outset. The halorhodopsin control in Driessen et al. (2026) showed that matched firing-rate reduction does not produce admissibility relaxation unless the trajectory crosses a specific dynamical minimum. The cytokine-stimulation results in Davidson et al. (2026) showed that the maintaining mechanism of the joint's admissibility structure can be inverted—becoming a gateway to pathology rather than a barrier against it. Together, these forced a distinction this monograph treats as central: relaxation, which is reversible and trajectory-dependent, versus collapse, which is irreversible and mechanism-dependent.

On Bibliography Scope

The bibliography clusters into six empirical domains that directly correspond to the six papers: sleep homeostasis and local sleep (Alfonsa et al., 2023; De Vivo et al., 2017; Diering et al., 2017; Funk et al., 2017; Huber et al., 2004; Massimini et al., 2004; Miyamoto et al., 2016; Tononi and Cirelli, 2014, 2020; Vyazovskiy et al., 2009, 2011); fibroblast immunology and arthritis (Buckley et al., 2021; Croft et al., 2019; Donlin et al., 2018; Zhang et al., 2019); sequence-to-function genomics (Avsec et al., 2021; de Boer and Taipale, 2024; Huang et al., 2023; Karollus et al., 2023; Kelley, 2018; Linder et al., 2025; Sasse et al., 2023; Zhou and Troyanskaya, 2015); battery recycling and circular economy (Cheng et al., 2024; Ciez and Whitacre, 2019; Geissdoerfer et al., 2017; Harper et al., 2019; Zhang et al., 2025); reinforcement learning and habit formation (Balleine and O'Doherty, 2010; Dayan and Balleine, 2002; Dezfouli and Balleine, 2014; Gershman and Daw, 2017; Keramati et al., 2011; Momennejad et al., 2017; Sutton and Barto, 2018); and quasar accretion and reverberation mapping (Abramowicz et al., 1988; Burke et al., 2021; Cackett et al., 2020; Fan et al., 2023; Fausnaugh et al., 2016; McHardy et al., 2018; Netzer, 2013; Peterson, 1993; Shakura and Sunyaev, 1973).

The theoretical background cites the primary RSVP and CLIO framework papers

rather than general dynamical systems literature, since the framework is novel rather than derived. The bibliography deliberately excludes entropic-gravity and free-energy references that appeared in earlier drafts, since neither thermodynamic spacetime (Jacobson, 1995) nor the free-energy principle (Friston, 2010) is developed substantively in the text. A citation that is not discussed is not a foundation.

Chapter 1

The Unifying Claim

Each system studied here is best described not by its current configuration but by the structured field of transitions it can still make.

1.1 Six Papers, One Structure

Six papers published in June 2026—spanning neuroscience (Driessen et al., 2026), immunology (Davidson et al., 2026), computational genomics (Spiro et al., 2026), sustainability science (Tian et al., 2026), cognitive science (Okitsu and Sakai, 2026), and astrophysics (Leung et al., 2026)—share a mathematical structure their respective literatures do not name. In each case, the operative quantity governing system function is not the current state but a **reachability structure**: a field of admissible transitions shaped by past trajectories, spatial distributions, and relaxation timescales.

The claim is not analogical. Analogies say: *this looks like that*. The claim here is structural: the same mathematical object—an admissibility field, a set-valued map over state space—appears as the functionally determinative quantity in each domain.

1.2 The Central Thesis

Definition 1.1 (Admissibility Field — Informal)**.** For a system with local state $s(x,t)$, the **admissibility field** $\mathcal{A}(x,t)$ is the set of transitions from $s(x,t)$ that the system can make while remaining consistent with its boundary conditions, history, and coupling to neighbouring regions.

Function—synaptic restoration (Driessen et al., 2026), joint homeostasis (Davidson et al., 2026), regulatory gene expression (Spiro et al., 2026), recycling viability (Tian et al., 2026), behavioral efficiency (Okitsu and Sakai, 2026), quasar luminosity (Leung et al., 2026)—is determined by the geometry of $\mathcal{A}(x,t)$, not by $s(x,t)$ alone.

The central thesis:

Constraint relaxation is a path-dependent operation on the admissibility field. The field must traverse specific regions of state space—threshold surfaces—for relaxation to occur. Furthermore, in systems where the admissibility field is actively maintained by a specific structural mechanism, perturbation of that mechanism can produce field collapse rather than mere relaxation.

1.3 Two Contributions from the Arthritis and RL Papers

The arthritis paper (Davidson et al., 2026) contributes two concepts absent from earlier analysis. First, it demonstrates that admissibility fields can be actively *maintained* by specific cell populations, not merely shaped by passive history: the $PI16^+$ fibroblast is the maintenance mechanism for the joint's admissibility structure. Second, it establishes that this maintenance is stoichiometric: the PIP joint's enrichment in $PI16^+$ cells relative to the DIP joint is quantitative, with measurable consequences for inflammatory susceptibility. This extends and grounds the work of Croft et al. (2019) and Zhang et al. (2019), who established that synovial fibroblast subsets drive distinct inflammatory outcomes, by tracing the stoichiometric bias back to embryonic origin.

The RL paper (Okitsu and Sakai, 2026) contributes the most mathematically explicit model of admissibility compression in any of the six papers. The plan-until-habit framework—a unification of goal-directed and habitual control that extends earlier work on the tension between model-based and model-free systems (Dayan and Balleine, 2002; Dezfouli and Balleine, 2014; Keramati et al., 2011) and on successor representations as intermediate strategies (Gershman and Daw, 2017; Momennejad et al., 2017)—is a formal, computable CLIO projection.

1.4 Theoretical Framework: RSVP and CLIO

The RSVP framework (Relativistic Scalar-Vector Plenum) models systems as coupled fields: a scalar potential $\Phi(x,t)$ encoding configurational tension; a lamphrodyne flow velocity $\mathbf{v}(x,t)$ encoding directed constraint propagation; and a configurational entropy density $S(x,t)$ encoding local degrees of freedom. The CLIO framework (Constraint-Local Inference and Observation) extends RSVP to inference dynamics. Projection $\pi : \mathcal{A} \to \hat{\mathcal{A}}$ compresses the admissibility field; the residual $\mathcal{A} \setminus \hat{\mathcal{A}}$ is the information lost; inference is its partial reconstruction.

1.5 Organization

Chapter 2 develops the mathematics of reachability fields formally. Chapter 3 introduces the Threshold Hypothesis and the distinction between relaxation and collapse. Chapter 4 is the primary neuroscience anchor. Chapter 5 is the primary immunology anchor. Chapter 6 develops residual geometry and historical encoding across all six papers. Chapter 7 bridges CLIO and machine learning through the RL and genomics papers. Chapter 8 unifies relaxation across spatial and temporal scales, with the quasar paper as cosmological anchor. Chapter 10 closes with a research program and twelve falsifiable predictions.

Chapter 2

The Mathematics of Reachability Fields

A field is not a collection of values. It is a geometry.

2.1 Formal Definition

Let $\mathcal{S}$ be a state space and $x \in \Omega$ a spatial index over domain Ω. At each (x, t), the system occupies local state $s(x, t) \in \mathcal{S}$.

Definition 2.1 (Admissibility Field)**.** An **admissibility field** is a set-valued map

$$\mathcal{A} : \Omega \times \mathbb{R}_{\geq 0} \longrightarrow 2^{\mathcal{S}}, \tag{2.1}$$

where $\mathcal{A}(x, t) \subseteq \mathcal{S}$ is the set of states reachable from $s(x, t)$ by local transitions consistent with the system's constraints at (x, t).

Proposition 2.1. $s(x, t) \in \mathcal{A}(x, t)$ for all (x, t). The map $t \mapsto |\mathcal{A}(x, t)|$ need not be monotone: it narrows under constraint accumulation and expands under relaxation.

2.2 Constraint Curvature and Bottlenecks

Definition 2.2 (Constraint Curvature)**.**

$$\kappa(x, t) = -\frac{\partial^2 \log |\mathcal{A}(x, t)|}{\partial t^2}. \tag{2.2}$$

Positive curvature: accelerating narrowing. Negative curvature: accelerating expansion.

Under the RSVP identification $|\mathcal{A}(x, t)| \propto e^{S(x,t)}$, constraint curvature equals $-\partial^2 S / \partial t^2$.

Definition 2.3 (Bottleneck)**.** A **bottleneck** at (x, t) is a local minimum of $|\mathcal{A}(x, t)|$ through which the system's trajectory must pass to reach subsequent configurations.

Bottlenecks appear empirically as: global cortical silence (50–400 ms) in the sleep system (Driessen et al., 2026); the viable-recycling threshold in the battery system

(Tian et al., 2026); and the minimum planning depth $D \geq 2$ in the RL system (Okitsu and Sakai, 2026).

2.3 Threshold Surfaces and Bipartite Decomposition

Definition 2.4 (Threshold Surface)**.** A **threshold surface** $\Sigma \subset \mathcal{S}$ partitions state space such that trajectories crossing Σ and returning produce expansion of the high-constraint admissibility component $\mathcal{A}_H$, while trajectories remaining on one side do not—regardless of proximity to Σ.

Definition 2.5 (Bipartite Admissibility Decomposition)**.**

$$\mathcal{A}(x,t) = \mathcal{A}_H(x,t) \cup \mathcal{A}_F(x,t), \tag{2.3}$$

where $\mathcal{A}_H$ is the **high-constraint component** (narrow, slow-relaxing, threshold-mediated, history-sensitive) and $\mathcal{A}_F$ is the **free component** (wide, fast-relaxing, history-insensitive).

The cortical data in Driessen et al. (2026) establish this decomposition at 67%/33% (HR/non-HR unit pairs). The arthritis data in Davidson et al. (2026) provide its mechanistic basis: $\mathcal{A}_H$ corresponds to the portion of the joint's admissibility field actively maintained by $PI16^+$ fibroblasts.

2.4 Relaxation Operators and Field Collapse

Definition 2.6 (Relaxation Operator)**.** $\mathcal{R}_\tau[\mathcal{A}](x,t) \supseteq \mathcal{A}(x,t+\tau)$, with equality when no relaxation occurs.

Definition 2.7 (Admissibility Collapse)**. Admissibility collapse** at (x,t) is the loss of the mechanism that actively maintains $\mathcal{A}_H$:

$$\mathcal{A}_{\text{collapse}}(x,t) \subsetneq \mathcal{A}(x,t), \tag{2.4}$$

where the reduction is irreversible on the natural system timescale and is driven by the functional inversion of the maintaining structure rather than by constraint accumulation.

Relaxation and collapse are qualitatively distinct. Relaxation expands $\mathcal{A}_H$ through threshold crossing and is reversible on the timescale of the accumulation process. Collapse reduces $\mathcal{A}_H$ through destruction of the maintaining mechanism—the $PI16^+$ fibroblast transition from insulator to amplifier under cytokine stimulation (Davidson et al., 2026)—and is irreversible without external intervention.

2.5 Spatial Coupling

The admissibility field couples across space through the RSVP flow:

$$\frac{\partial \mathcal{A}(x,t)}{\partial t} = f[\mathcal{A}(x,t), s(x,t)] + \int_{\Omega} K(x,x')\,\mathbf{v}(x',t)\,dx'. \tag{2.5}$$

The coupling kernel $K(x,x')$ determines the spatial range of relaxation propagation. In mouse cortex, the contralateral probe shows no response to ipsilateral off-period induction, implying $K \approx 0$ at inter-hemispheric distances (Driessen et al., 2026)—consistent with local sleep being genuinely local (Huber et al., 2004; Vyazovskiy et al., 2011). In quasar accretion, K propagates at the speed of light, yielding the wavelength-dependent lag $\tau \propto \lambda^{4/3}$ (Cackett et al., 2020; Fausnaugh et al., 2016; Leung et al., 2026).

Chapter 3

The Threshold Hypothesis

It is not enough to reduce activity.
The system must cross the silence.

3.1 Motivation: The Halorhodopsin Experiment

The idea that slow-wave generation requires specific circuit conditions rather than mere firing reduction has been established through work on SOM$^+$ interneuron contributions to cortical bistability (Funk et al., 2017) and on chloride homeostasis as a mediator of local sleep pressure (Alfonsa et al., 2023). The halorhodopsin control in Driessen et al. (2026) makes this general principle experimentally decisive.

Halorhodopsin mice received tonic inhibition reducing MUA firing by an amount statistically indistinguishable from the SOM$^+$ and ACR off-period induction protocols (Extended Data Fig. 8 of the source paper: one-way Welch's ANOVA, NS). No downstream effects followed: no SWA reduction, no STTC reduction, no GluA1 or pGluA1 reduction, no memory rescue. Within the same SOM$^+$ mice, off-period induction outperformed tonic inhibition in every animal tested.

The operative variable is the bistable pattern: global population silence traversal (50–400 ms, $\geq$12/16 channels) and return.

3.2 Statement of the Threshold Hypothesis

Hypothesis 3.1 (Threshold Hypothesis)**.** Many admissibility fields contain a threshold surface $\Sigma \subset \mathcal{S}$ such that:

1. Relaxation of $\mathcal{A}_H$ requires the trajectory to cross Σ.
2. Trajectories remaining on one side of Σ produce no expansion of $\mathcal{A}_H$.
3. Crossing Σ is followed by a return trajectory; relaxation occurs during or after the return.

Proposition 3.1 (Threshold Relaxation Condition)**.** Let the crossing indicator $\chi(\gamma) = 1$ iff trajectory γ crosses Σ, else 0. Then:

$$\mathcal{R}_\tau[\mathcal{A}_H](x,t) \supsetneq \mathcal{A}_H(x,t) \iff \chi(\gamma) = 1. \tag{3.1}$$

The relationship between $\mathcal{A}_H$-expansion and proximity to Σ is discontinuous.

3.3 Relaxation vs. Collapse: A Formal Distinction

In the cortical system, the threshold-maintaining mechanism (SOM^+ interneuron networks; intracellular chloride regulation (Alfonsa et al., 2023; Funk et al., 2017)) is intact and can perform threshold crossings on demand (Driessen et al., 2026). In the inflamed joint (Davidson et al., 2026), the *PI16*$^+$ fibroblast population—which maintains the admissibility structure under homeostatic conditions through *WNT*, *BMP*, and *FGF* signalling—inverts its function under TNF and IL-1β stimulation. The maintaining mechanism is not bypassed; it becomes actively destabilizing.

Definition 3.1 (Threshold Maintainer)**.** A **threshold maintainer** is a structural element whose presence ensures that Σ exists and that the bipartite decomposition $\mathcal{A} = \mathcal{A}_H \cup \mathcal{A}_F$ is stable. Its functional inversion constitutes admissibility collapse.

In the cortical system, SOM^+ interneurons are threshold maintainers (Funk et al., 2017). In the joint system, *PI16*$^+$ fibroblasts are threshold maintainers (Davidson et al., 2026), extending the fibroblast-subset literature (Croft et al., 2019; Zhang et al., 2019) to a mechanistic role in admissibility field maintenance.

3.4 Threshold Surfaces Across All Six Domains

Cortical system. Σ = boundary of global population silence. Mechanism: SOM^+ interneuron recruitment (Funk et al., 2017); intracellular chloride regulation (Alfonsa et al., 2023). Demonstrated by halorhodopsin control (Driessen et al., 2026).

Joint system. Σ = cytokine concentration at which *PI16*$^+$ fibroblasts transition from WNT/BMP/FGF maintenance to chemotactic amplification. Mechanism: TNF and IL-1β suppression of homeostatic gene cassette, upregulation of CCL2, CCL3, CCL5, IL6, IL1B (Davidson et al., 2026).

Genomic admissibility. Σ = MAF threshold ($\approx$0.05) below which variants contribute negligibly to model predictions (Spiro et al., 2026). The regulatory grammar encoded in reference sequence models such as Enformer (Avsec et al., 2021) and Borzoi (Linder et al., 2025) is not defined for variants below this frequency.

Accretion disk. Σ = thin-to-slim disk transition in Eddington-ratio space. Below: thin disk, $\lambda f_\lambda \propto \lambda^{-4/3}$ (Shakura and Sunyaev, 1973). Above: slim disk, $\lambda f_\lambda = \text{const}$ (Abramowicz et al., 1988). J0439+1634 at 60% Eddington is below Σ (Leung et al., 2026).

Battery recycling. Σ = minimum interprovincial coordination level for viable formal recycling. Below: 39.3% utilization. Above (neighbouring-province coordination): 65.68% (Tian et al., 2026).

RL agent. Σ = minimum forward planning depth $D \geq 2$. Below: PIT fails entirely. Above: PIT emerges robustly and persists for $D = 2, 3, 4$ (Okitsu and Sakai, 2026). This is the most structurally explicit threshold in any of the six papers: proven mathematically by the authors.

Chapter 4

The Cortical Anchor: Bistability and Local Phase Operations

> Sleep is not a global state that the brain enters.
> It is a local phase operation that the brain performs.

4.1 Background: The Local-Sleep Lineage

The local-sleep interpretation of cortical dynamics did not emerge suddenly. Evidence for spatially heterogeneous sleep pressure accumulated through learning-dependent slow-wave modulation (Huber et al., 2004), the characterization of cortical firing homeostasis (Vyazovskiy et al., 2009), the identification of local sleep intrusions during wakefulness (Vyazovskiy et al., 2011), the synaptic homeostasis hypothesis (Tononi and Cirelli, 2014, 2020), ultrastructural synaptic scaling across the wake/sleep cycle (De Vivo et al., 2017), molecular scaling mechanisms involving Homer1a (Diering et al., 2017), chloride-mediated local sleep regulation (Alfonsa et al., 2023), and the travelling-wave structure of cortical slow oscillations (Massimini et al., 2004). The contribution of Driessen et al. (2026) is to demonstrate that this local sleep pressure can be selectively discharged during wakefulness by inducing the characteristic bistable on/off pattern, without requiring global sleep onset.

4.2 Experimental Design

Driessen et al. (2026) implanted adult mice ($>$ 8 weeks) with bilateral homotopic 16-channel silicon probes (50 μm spacing). Two transgenic strategies induced local on/off oscillations: SOM$^+$ mice ($n = 10$, ChR2 in somatostatin interneurons, positive-going slow waves) and ACR mice ($n = 8$, soma-targeted stGtACR1 in CaMKIIα neurons, negative-going slow waves). A third line—halorhodopsin mice—provided the matched-firing-rate tonic inhibition control. The staged induction protocol (1 Hz/180 ms for 10 min; 2 Hz/140 ms for 10 min; 3 Hz/100 ms for 5 min; then 3 Hz/80

ms until sleep onset) was designed to mimic the natural decline in off-period duration across sleep-pressure states (Vyazovskiy et al., 2009).

4.3 Results and RSVP Interpretation

Measure	SOM+	ACR	Halo	Interpretation
SWA reduction (1-h recovery)	$p<0.0005$	$p<0.005$	NS	$\Delta P(x) < 0$
STTC, HR unit pairs	$p<0.05$	$p<0.05$	NS	$\mathcal{A}_H$ expansion
GluA1 and pGluA1 levels	$p<0.005$	$p<0.0005$	NS	Synaptic renormalization
Memory rescue (FTR task)	Yes	—	—	$\mathcal{A}_H$ stabilization

Table 4.1: Downstream effects under matched firing-rate reduction. NS: $p > 0.05$. Halo: halorhodopsin tonic inhibition. HR: homeostatically regulated unit pairs. All significance values from Supplementary Table 1 of Driessen et al. (2026).

The halorhodopsin result establishes Hypothesis 3.1 in the cortical domain: the trajectory pattern, not the magnitude of firing reduction, determines whether $\mathcal{A}_H$ expands. This is consistent with the physiology of bistability: somatostatin interneurons act as master regulators of cortical excitability (Funk et al., 2017), and their activation produces synchronized transitions to and from the off-period state that mere amplitude reduction cannot replicate.

The bipartite structure (67% HR / 33% non-HR unit pairs, consistent across mouse lines and probe identity) grounds the formal decomposition $\mathcal{A} = \mathcal{A}_H \cup \mathcal{A}_F$ with biological specificity. The fraction of HR pairs tracks whether a unit pair is homeostatically regulated—whether its synchrony reflects the waking-activity history. Only these pairs participate in $\mathcal{A}_H$ dynamics.

4.4 The RSVP Sleep Pressure Functional

Define sleep pressure at location x over duration T:

$$P(x,T) = \int_0^T \max\big[0,\ S^*(x) - S(x,t)\big]\, dt, \tag{4.1}$$

where $S^*(x)$ is the homeostatic target entropy density. Under waking activity, $S(x,t)$ decreases (synaptic potentiation narrows $\mathcal{A}_H$ (Tononi and Cirelli, 2014)) and $P(x,T)$

accumulates. The SWA in the first hour of recovery sleep is the empirical observable for P (Driessen et al., 2026; Vyazovskiy et al., 2009).

An off-period oscillation traverses Σ and produces $\Delta S > 0$, reducing P by approximately $\Delta S \cdot (\tau_2 - \tau_1)$. The halorhodopsin constraint:

$$\chi(\gamma_{\text{tonic}}) = 0 \implies \Delta S_{\text{tonic}} \approx 0, \tag{4.2}$$

constraining the RSVP potential landscape: $\Phi(x, t)$ must have a barrier structure such that the off-period minimum must be visited for relaxation to occur.

4.5 Memory Consolidation as CLIO Stabilization

Sleep's role in memory consolidation has been established through a series of studies (Huber et al., 2004; Miyamoto et al., 2016), with the current paper extending this to the key result that off-period induction—without global sleep—is sufficient for consolidation (Driessen et al., 2026). In CLIO terms: acquisition projects the texture discrimination onto a compressed representation in M2 and S1 (Miyamoto et al., 2016). Without consolidation, $\mathcal{A}_H$ around the representation remains wide, permitting interfering traces. Off-period induction crosses Σ and narrows $\mathcal{A}_H$ locally, stabilizing the projection against interference.

Chapter 5

The Joint System: Stoichiometry, Topology, and Collapse

> The PIP joint is not more vulnerable because of what it is now.
> It is more vulnerable because of what it was.

5.1 Background: Fibroblast Heterogeneity and Inflammatory Disease

The recognition that synovial fibroblasts are not a homogeneous support population but comprise distinct functional subsets capable of driving either inflammation or tissue damage was established by Croft et al. (2019), who identified spatially distinct fibroblast lineages within the mouse joint, and by Zhang et al. (2019), who applied single-cell transcriptomics and mass cytometry to human rheumatoid arthritis tissue to define inflammatory fibroblast states. Subsequent work has emphasized that the tissue microenvironment—including mechanical forces, vascular proximity, and cytokine gradients—shapes fibroblast function at the local level (Buckley et al., 2021). The contribution of Davidson et al. (2026) is to trace the basis of these local differences back to their embryonic origins, demonstrating that site-specific inflammatory susceptibility is pre-patterned *in utero*.

5.2 Architectural and Cellular Asymmetry

High-resolution synchrotron X-ray microtomography of human fetal joints at 14–15 post-conception weeks reveals fundamental differences between PIP and DIP anatomy established before birth (Davidson et al., 2026):

- The DIP synovium is a thin layer, prominent dorsally. The PIP possesses a substantially larger "baggy" synovial volume forming a palmar apron.
- The PIP tendon architecture is dramatically more complex: the flexor digitorum

superficialis (FDS) splits and inserts at the PIP joint, forming a structural tunnel through which the flexor digitorum profundus (FDP) traverses. Only the FDP inserts at the DIP.

- Across early (8–9 pcw) and late (12–14 pcw) fetal development, joint spaces are predominantly (96.5%) *COL1A1*$^+$ early fibroblasts and *COL2A1*$^+$ chondrocytes.
- Critically, the PIP joint's larger palmar volume and interstitial spaces are enriched in *PI16*$^+$ fibroblasts relative to the DIP joint—a stoichiometric asymmetry established embryonically.

Single-cell RNA sequencing identifies 16 distinct embryonic stromal clusters partitioning into three macro-regions: chondrocytes (*COL9A2*$^+$), soft tissue fibroblasts (STFs, *COL1A1*$^+$ *ZFHX4*$^+$), and cartilage zone stromal cells (CZSCs, *COL1A1*$^+$ *CLU*$^+$). Cross-pair correlation modelling confirms that *PI16*$^+$ cells (cluster F2) occupy two precise spatial niches: perivascular/adventitial positions associated with *vWF*$^+$ blood vessels, and interstitial tissue interfaces along complex tendon-ligament borders.

5.3 The PI16$^+$ Cell as Admissibility Maintainer

Because tenosynovitis is a known early propagator of rheumatoid arthritis (Buckley et al., 2021), the enrichment of *PI16*$^+$ cells around the PIP joint's complex tendon matrix provides a structural explanation for site-specific inflammatory vulnerability. Under homeostatic conditions, *PI16*$^+$ fibroblasts express a gene cassette maintaining structural integrity and vascular stabilization: *IGF1*, *FGF2*, *FGF9*, *WNT* pathway components, *BMP* pathway components, and cell adhesion molecules (Davidson et al., 2026). This cassette is the molecular implementation of the threshold-maintaining function described in Definition 3.3.

When stimulated with TNF and IL-1β:

- **Shared proinflammatory activation**: massive upregulation of CCL2, CCL3, CCL5, IL6, IL1B—responses shared with other fibroblast populations, consistent with Zhang et al. (2019).
- **Unique chemotactic activation**: upregulation of active cellular migration and leukocyte chemotaxis cascades not observed in other fibroblasts.
- **Homeostatic collapse**: dramatic downregulation of the entire WNT, BMP, and FGF signalling network and cell adhesion molecules.

The *PI16*$^+$ cell inverts from structural insulator to active gateway—admissibility collapse in the sense of Definition 2.7.

5.4 Convergent Lining Geometry

A conceptually important finding concerns the adult synovial lining layer (LL). Davidson et al. (2026) demonstrate that the embryonic LL emerges convergently from two separate developmental lineages:

1. **Interzone lineage**: GDF5$^+$ CZSCs adjacent to the articular surface differentiate into *CRTAC1*$^+$ *PRG4*$^+$ *HBEGF*$^+$.
2. **Non-interzone lineage**: STFs from *PI16*$^+$ (F2) and *ALDH1A1*$^+$ (F1) progenitors migrate and polarize into *TPPP3*$^+$ *PRG4*$^+$ *HBEGF*$^+$.

Both converge on a shared *SOX5*$^+$ *HBEGF*$^+$ lining phenotype driven by localized hypoxia and EGFR signalling. Matching embryonic signatures against adult AMP2 datasets confirms deep transcriptomic homology: embryonic lining cells cluster with adult arthritic lining cells, while embryonic *PI16*$^+$ and *ALDH1A1*$^+$ subsets map onto adult CD34$^+$ and CXCL12$^+$ SFRP1$^+$ sublining fibroblasts (Davidson et al., 2026).

In admissibility terms, this is trajectory-independent residual geometry: two different developmental paths produce the same admissibility structure. The convergent lining phenotype is an attractor of the developmental dynamics under hypoxia and EGFR signalling, regardless of precursor lineage.

5.5 Formal Development: Stoichiometric Admissibility

Let $\rho(x)$ denote the local density of *PI16*$^+$ fibroblasts at joint site x, and $c(x,t)$ the local cytokine concentration. Under homeostasis:

$$\frac{d|\mathcal{A}_H(x,t)|}{dt} \approx 0 \quad \text{(actively maintained by } PI16^+ \text{ cassette).} \tag{5.1}$$

Under inflammatory perturbation, as $c(x,t) \to c_{\text{threshold}}$, the homeostatic function inverts:

$$\frac{d|\mathcal{A}_H(x,t)|}{dt} < 0, \quad \text{with rate} \propto \rho(x). \tag{5.2}$$

The PIP joint, with higher ρ, collapses faster and more completely than the DIP joint under equal inflammatory challenge. This generates a quantitative prediction: collapse rate should be a monotone function of baseline *PI16*$^+$ density across joint sites, holding cytokine stimulus constant.

Chapter 6

Historical Encoding and Residual Geometry

> The past does not merely precede the present.
> It is written into the geometry of what can happen next.

6.1 Three Kinds of Memory

Three distinct mechanisms by which prior trajectories constrain current admissibility appear across the six papers:

1. **Trajectory accumulation** (reversible, fast): recent activity narrows $\mathcal{A}_H$ through synaptic potentiation (Tononi and Cirelli, 2014); threshold crossing restores it (Driessen et al., 2026). Encoding timescale: hours. Reversal: hours.

2. **Developmental encoding** (irreversible, slow): embryonic morphogenetic trajectory encodes a residual geometry biasing adult admissibility (Davidson et al., 2026). Encoding timescale: months. Reversal: never.

3. **Stoichiometric encoding** (irreversible without intervention): the density of threshold-maintaining cells varies by developmental history, setting the collapse vulnerability of adult tissue (Davidson et al., 2026). Encoding timescale: months. Reversal: cell therapy only.

6.2 The Arthritis Paper: Residual Geometry at Two Levels

Davidson et al. (2026) demonstrate historical encoding operating at both architectural and cellular levels simultaneously. At the architectural level, the palmar synovial apron, the FDS-FDP tendon tunnel, and the interstitial volumes of the PIP joint are established during embryogenesis and persist into adulthood as structural biases.

At the cellular level, the higher baseline density of $PI16^+$ fibroblasts at the PIP joint reflects the developmental allocation of perivascular and interstitial niches during morphogenesis.

Definition 6.1 (Residual Geometry)**.** The **residual geometry** $\mathcal{G}_{\text{res}}(x)$ is the constraint structure encoded by prior trajectories no longer actively occurring:

$$\mathcal{G}_{\text{res}}(x) = \mathcal{A}(x, t_{\text{adult}}) - \mathcal{A}_{\text{free}}(x, t_{\text{adult}}), \tag{6.1}$$

where $\mathcal{A}_{\text{free}}$ is the admissibility that would exist without historical encoding.

The matching of embryonic fibroblast signatures against adult AMP2 datasets (Davidson et al., 2026; Zhang et al., 2019) is an empirical instance of residual geometry recovery: the developmental trajectory is reconstructed from its adult residual.

6.3 Non-Markovian Dynamics and the Yarncrawler Structure

Systems with residual geometry are not Markovian at the observable level. The current state $s(x, t_{\text{adult}})$ does not determine $\mathcal{A}(x, t_{\text{adult}})$; the developmental trajectory γ_{dev} is also required. This is the epistemological structure of the Yarncrawler framework: world-state reconstruction as constraint closure. Bayesian inference on such systems requires a prior over developmental trajectories, not merely over current states.

This non-Markovian structure appears in the cortical system on shorter timescales: the recent waking trajectory encodes into $\mathcal{A}_H$ through synaptic potentiation, and the SWA in recovery sleep depends on the integrated history of recent waking activity (Tononi and Cirelli, 2014; Vyazovskiy et al., 2009) rather than on the instantaneous state at sleep onset. The cortical and joint systems differ only in the timescale of encoding and the reversibility of the result.

6.4 Battery Recycling: Dynamic Residual Geometry

The spatial distribution of EOL batteries in 2030 is encoded by the diffusion trajectory of EV adoption from 2016 to 2030 (Tian et al., 2026). The hotspot migration pattern (northeast $\rightarrow$ southwest 2020–2026; southeast $\rightarrow$ northwest 2026–2030) is the residual geometry of the adoption trajectory, projected onto the admissibility field of the recycling system. Future recycling infrastructure must be sited to anticipate this relocation rather than serve the current geography of battery retirement (Cheng et al.,

2024; Tian et al., 2026).

System	Encoding timescale	Reversal timescale
Cortical (synaptic potentiation)	Hours	Hours (sleep)
SAGE-net (training)	Days–weeks	N/A (model)
Battery adoption diffusion	Years	Decades
Joint (architectural)	Months (embryonic)	Never
Joint ($PI16^{+}$ density)	Months (embryonic)	Never without cell therapy
Quasar disk (corona-mediated)	Days (rest-frame)	Days–weeks

Table 6.1: Timescale hierarchy of historical encoding.

Chapter 7

Constraint Compression and Predictive Efficiency

> To plan efficiently is to compress the admissibility field without losing the transitions that matter.

7.1 The CLIO Projection Cycle

A CLIO projection $\pi : \mathcal{A} \to \hat{\mathcal{A}}$ compresses the admissibility field. The compression is efficient but lossy: the residual $\mathcal{A} \setminus \hat{\mathcal{A}}$ is the admissibility information sacrificed. Inference is its partial reconstruction from new observations.

7.2 Habit Formation as Admissibility Compression

The distinction between habitual and goal-directed control has a long history in computational neuroscience (Balleine and O'Doherty, 2010; Dayan and Balleine, 2002; Dezfouli and Balleine, 2014; Keramati et al., 2011). More recent work has emphasized successor representations and compressed future-state prediction as intermediate strategies between model-free and model-based learning (Gershman and Daw, 2017; Momennejad et al., 2017). The hybrid-architecture tradition treats the two systems as running in parallel with an explicit gating or weighting parameter. The contribution of Okitsu and Sakai (2026) is to unify habit formation and Pavlovian-instrumental transfer within a single planning architecture.

7.2.1 Mathematical Structure of the Okitsu-Sakai Model

Let the environment state at time t be a feature vector $s_t \in \mathbb{R}^n$ with $s_0 \equiv 1$ (intercept). Immediate reward:

$$r_t = u_{a_t}^\top s_t, \tag{7.1}$$

where u_{a_t} is an action-dependent parameter vector implementing the classic Rescorla-Wagner reward representation (Rescorla and Wagner, 1972). One-step forward transition:

$$\hat{s}_{t+1} = W_{a_t} s_t. \tag{7.2}$$

Chaining over a planned sequence $\{a_{t+1}, \ldots, a_{t+\tau-1}\}$:

$$\hat{s}_{t+\tau} = \Big(\prod_{k=1}^{\tau-1} W_{a_{t+k}}\Big) W_{a_t} s_t. \tag{7.3}$$

For lookahead depth D, the unified action-value is:

$$Q(s_t, a_t) = \max_{a_{t+1:t+D-1}} \left[\sum_{\tau=0}^{D-1} \gamma^\tau r^{\mathrm{MB}}_{t+\tau} + \gamma^D V(\hat{s}_{t+D})\right], \tag{7.4}$$

where $V(s) = v^\top s$ is updated by temporal-difference learning (Sutton and Barto, 2018):

$$v \leftarrow v + \alpha_v \, \delta_t \, s_t, \qquad \delta_t = r_t + \gamma V(s_{t+1}) - V(s_t). \tag{7.5}$$

7.2.2 CLIO Interpretation

In CLIO terms: the model-based rollout is explicit computation of $\mathcal{A}(s,t)$ for $t \in [0, D]$ (within the planning horizon). The model-free tail $v^\top s$ is the compression $\hat{\mathcal{A}}(s, t > D)$—a low-dimensional approximation to the long-range admissibility field, corresponding to the successor representation (Momennejad et al., 2017) at convergence. The habit is $\hat{\mathcal{A}}$; the Pavlovian response is the residual $\mathcal{A} \setminus \hat{\mathcal{A}}$: sensitivity to stimuli predicting states outside the habitual compression.

Proposition 7.1. The unification claim of Okitsu and Sakai (2026)—habits and Pavlovian responses from a single computational substrate—is equivalent in CLIO terms to the claim that $\hat{\mathcal{A}}$ and $\mathcal{A} \setminus \hat{\mathcal{A}}$ are projection and residual of a single underlying field, not outputs of separate systems.

7.2.3 The Critical Depth Threshold

Okitsu and Sakai (2026) prove that PIT fails entirely for $D < 2$. A minimum planning depth of $D = 2$ is required for the matrix multiplications to bridge the gap between the initial state-action pair and the terminal outcome nodes. This is the RL instance of Hypothesis 3.1: a discrete condition ($D \geq 2$) separates zero-PIT from robust-PIT, with no smooth interpolation.

7.3 SAGE-net: Limits of Admissibility Compression in Genomics

The limitations observed by Spiro et al. (2026) extend a broader pattern in sequence-to-function modelling. Earlier architectures such as DeepSEA (Zhou and Troyanskaya, 2015), Basenji (Kelley, 2018), Enformer (Avsec et al., 2021), and more recent RNA-seq coverage models (Linder et al., 2025) all demonstrated impressive regulatory prediction while struggling to capture the full structure of personal transcriptomic variation (Huang et al., 2023; Karollus et al., 2023; Sasse et al., 2023). The challenge of learning generalizable regulatory grammar—the constraint structure that transfers across genomic loci—has been framed explicitly as a roadmap for the field (de Boer and Taipale, 2024).

The SAGE-net contrastive architecture decomposes gene expression into mean component ($\hat{\mathcal{A}}$, cross-locus regulatory grammar) and personal component ($\mathcal{A} \setminus \hat{\mathcal{A}}$, individual deviation). Personal genome training improves $\hat{\mathcal{A}}$ for seen genes but does not improve it for unseen genes (Spiro et al., 2026): the compression learns locus-specific admissibility boundaries but not the global constraint geometry. DNA methylation partially generalizes at 100,000 training regions ($p = 4.96 \times 10^{-4}$), suggesting a globally smoother—lower curvature—epigenomic admissibility field (Spiro et al., 2026).

The seqlet distance analysis (p-SAGE-net mean 5,403 bp vs. r-SAGE-net 3,117 bp from TSS; $p = 2 \times 10^{-80}$, two-sample K-S test) shows that personal genome training extends the effective range of the spatial coupling kernel $K(x, x')$, capturing distal regulatory interactions (Karollus et al., 2023) that the reference-genome model misses. The failure to generalize across loci reflects locus-specificity of K: the distal coupling structure does not transfer.

Chapter 8

Relaxation Across Scales

The same propagation structure appears
at the scale of milliseconds and of millennia.

8.1 Three Propagation Regimes

Cortical (ms to hour). Off-period induction produces local SWA reduction that does not propagate to the homotopic contralateral hemisphere: $K(x, x') \approx 0$ at inter-hemispheric distances in mouse cortex (Driessen et al., 2026). This is consistent with local sleep being a genuine spatial phenomenon (Huber et al., 2004; Vyazovskiy et al., 2011) rather than a global state modulation. Recovery to contralateral equivalence occurs within 3–5 hours.

Quasar (light-crossing, days to years). The lamp-post model predicts corona perturbations propagate through the accretion disk at the speed of light, with lag $\tau(\lambda) \propto \lambda^{4/3}$ (Cackett et al., 2020; Fausnaugh et al., 2016; Peterson, 1993). Reverberation mapping has established this relation in nearby AGNs (Cackett et al., 2020; Fausnaugh et al., 2016; McHardy et al., 2018); the characteristic variability timescale in astrophysical accretion disks has been characterized statistically in large samples (Burke et al., 2021). Leung et al. (2026) extend this to $z = 6.51$, constraining the W2–W1 lag to < 160 days (observed frame), consistent with the thin-disk prediction of $\sim$20 days.

Battery network (year scale). The hotspot migration propagates EV adoption patterns through the recycling system's admissibility field on year timescales (Tian et al., 2026), consistent with the circular-economy dynamics literature (Ciez and Whitacre, 2019; Geissdoerfer et al., 2017).

8.2 The Thin-Disk Variable SED as Field Diagnostic

The broader population of high-redshift quasars at $z \approx 6$ accretes at substantially higher Eddington ratios than local AGNs, and understanding accretion disk structure at these redshifts is a key open problem in SMBH growth (Fan et al., 2023; Netzer,

2013). The variable SED of J0439+1634 follows $\lambda f_\lambda \propto \lambda^{-1.58\pm0.25}$, consistent with the thin-disk prediction of $-4/3$ (Shakura and Sunyaev, 1973) and inconsistent with slim-disk expectations (Abramowicz et al., 1988). At 60% Eddington, the system is on the thin-disk side of the threshold Σ in accretion parameter space—the admissibility field retains its narrow, well-constrained geometry.

The X-ray variability amplitude ($8.2 \pm 3.7\times$) vastly exceeds the IR amplitude ($\sim 1.15\times$ in W1), consistent with a more compact X-ray-emitting corona (Kelly et al., 2009). In admissibility terms: amplitude $\propto |\mathcal{A}|^{-1}$. The corona's smaller admissibility volume makes it more sensitive to perturbations in Φ—the same bipartite compactness effect that makes HR unit pairs ($\mathcal{A}_H$) more responsive than non-HR pairs ($\mathcal{A}_F$) in the cortical system (Driessen et al., 2026).

8.3 The RSVP Cosmological Connection

The thin-disk reverberation relation $\tau \propto \lambda^{4/3}$ is a propagation-speed constraint. The RSVP reinterpretation of cosmological redshift as entropic field response predicts a slight steepening of this relation at high redshift: $\tau \propto \lambda^{\alpha}$ with $\alpha > 4/3$ at $z \gg 1$. The current result (J0439+1634 consistent with $\alpha = 4/3$ within uncertainties (Leung et al., 2026)) does not discriminate between standard and RSVP interpretations. Population-level analysis with Rubin/Roman ($N > 100$ variable quasars at $z > 4$ (Burke et al., 2021; Fan et al., 2023)) will provide the statistical power to test this.

Chapter 9

Admissibility Display: An Applied Principle

The spinner exposes $s(x, t)$.
What the user wants is $\mathcal{A}(x, t)$.

9.1 A Modern Instance of the General Problem

The monograph has argued that in each of its empirical domains—cortical circuits, developing joints, genomic regulatory landscapes, industrial recycling networks, reinforcement-learning agents, and quasar accretion disks—the operative quantity governing system function is not the current state but the *reachability structure*: the field of admissible transitions from that state. The opening chapter states this as a principle of theoretical analysis. This chapter applies it to a designed artifact.

The artifact in question is the conversational AI interface. The analysis is brief, but the application is precise enough to be useful.

9.2 The Spinner as State Display

When a conversational AI system is generating a response, most current interfaces display a loading indicator of some kind—a spinner, a pulsing cursor, an animated ellipsis. The communicative content of this indicator is minimal: it conveys that the system is in a processing state and that no output is currently available. In the language developed here:

$$\text{The spinner exposes } s(x, t) = \{\textit{thinking}\} \text{ while concealing } \mathcal{A}(x, t).$$

This is not a neutral design choice. It is a specific selection from a large space of possible interface behaviors, and it systematically privileges state information over reachability information.

What the user actually wants to know during a waiting period is not what state the system is in—that is obvious—but what the system's *admissibility field* looks like at this moment. Concretely:

- What interpretations of the query are currently under consideration?
- Which interpretations have already been rejected, and why?
- What assumptions remain unresolved, and which of several plausible resolutions is the system currently pursuing?
- What questions, if answered by the user, would most rapidly narrow the admissibility field toward a useful response?

None of this information is conveyed by the spinner. The interface presents a collapsed admissibility field—a single node: *wait*.

9.3 The CLIO Reading of the Final Response

When the response eventually arrives, it represents a CLIO projection: the output is $\hat{\mathcal{A}}$, the compressed result of a reasoning process whose full admissibility field $\mathcal{A}$ is never exposed.

In CLIO terms, the residual $\mathcal{A} \setminus \hat{\mathcal{A}}$ contains:

- Interpretations that were considered and rejected.
- Lines of reasoning that were partially followed and then abandoned.
- Uncertainty that was resolved by implicit assumption rather than user input.
- Alternative responses that were close to the chosen output but were discarded.

Current interfaces *maximize projection opacity*: the user observes only $\hat{\mathcal{A}}$ and has no access to the geometry of $\mathcal{A}$ from which it emerged. This is not an inevitable feature of language model inference; it is a design decision that treats the conversation as a query-and-response transaction rather than as a shared reasoning process.

The contrast with human collaborative reasoning is sharp. A human researcher working on an unfamiliar problem does not disappear for several minutes and then deliver a finished analysis. They think aloud: "I see three possible interpretations; the first doesn't work because…; let me ask you whether you mean X or Y before I go further." This keeps the shared cognitive workspace alive. The admissibility field of the collaboration remains visible to both parties throughout.

9.4 Productive Waiting vs. Dead Waiting

The distinction between productive and dead waiting maps directly onto the admissibility framework.

Definition 9.1 (Dead Waiting)**. Dead waiting** is a waiting period during which the user has no access to the system's admissibility field: no information about what interpretations are under consideration, what uncertainties remain open, or what transitions are currently reachable.

Definition 9.2 (Productive Waiting)**. Productive waiting** is a waiting period during which the system continues to expose its admissibility field: presenting partial hypotheses, asking clarifying questions, displaying rejected interpretations, or exploring adjacent conceptual territory while heavier computation proceeds.

The difference is not merely psychological, though the psychological effect is real—humans tolerate latency much better when the waiting period contains meaningful content (Miller, 1968; Shneiderman, 1984). The difference is also epistemic: productive waiting allows the user to *contribute* to the admissibility field during computation. A well-placed clarifying question can eliminate half the remaining interpretations; a user correction can redirect the computation before it reaches a wrong conclusion. Dead waiting forecloses this collaboration until it is too late to matter.

A concrete example of productive waiting in a long-running AI analysis:

> *(1 second after query)* "I currently have three interpretations of your proposal. Before I work through the details, let me ask: does the admissibility field in your framework vary continuously with time, or does it undergo discrete transitions?"
>
> *(User answers.)*
>
> *(15 seconds later)* "That eliminates the discrete-transition reading. One more question while I continue: do you treat entropy as a property of field configurations, or as a field itself?"
>
> *(User answers.)*
>
> *(40 seconds later)* "Good. Here is the full analysis, with those two assumptions fixed."

The total latency is the same. The epistemic outcome is better, because the output is now conditioned on user-supplied disambiguation rather than on the system's implicit assumptions.

9.5 The Admissibility Display Principle

The foregoing motivates a general design principle that extends well beyond AI interfaces.

Definition 9.3 (Admissibility Display Principle)**.** A system interacting with a user or agent on timescales exceeding immediate response should expose as much of its current admissibility field as possible rather than merely its current state. Specifically: it should communicate what transitions remain available, what uncertainties remain unresolved, and what trajectories through state space are currently under active consideration.

The Admissibility Display Principle is not new in its practical instances—it is already observed, in varying degrees, by a wide range of interactive systems:

- Navigation software displays alternative routes, not merely the current route.
- Compilers show build progress and intermediate errors, not merely a final success or failure.
- File transfer interfaces show completion estimates and transfer rates, not merely a progress bar.
- Operating systems expose process status, not merely the aggregate load.
- Version control systems expose branch and merge state, not merely the current working tree.

In each case, the system exposes reachability information—what is still open, what paths remain, what the current trajectory implies about future states—rather than merely reporting current state. The informational value of these displays lies precisely in their admissibility content.

The spinner violates this principle because it collapses the visible admissibility field to a single node. It is, in the terms of this monograph, the interface equivalent of halorhodopsin tonic inhibition: the activity level is reduced, but nothing crosses the threshold surface, and no useful information is transmitted to the user about the system's internal dynamics.

Proposition 9.1. The informational value of an interactive interface is proportional not to its representation of present state but to its representation of currently reachable future states.

This proposition restates the central thesis of the monograph—that reachability is more informative than state—in the domain of interface design. It connects to a long

tradition of latency research showing that perceived responsiveness depends less on raw speed than on the quality of feedback during waiting (Miller, 1968; Nielsen, 1994; Shneiderman, 1984).

9.6 Historical Encoding in the Interface

The CLIO projection-opacity critique connects to the historical-encoding argument of Chapter 6. A response that emerges from a black-box waiting period has the structure of residual geometry: the user sees the output $\hat{\mathcal{A}}$ but not the trajectory γ that produced it. The reasoning history is encoded into the response but not exposed.

This has practical consequences. A user who cannot observe the reasoning trajectory cannot identify where the system made an assumption that conflicts with their intent. They can only observe the final output and either accept it or reject it. The interface is non-Markovian from the user's perspective: the current response depends on the system's internal reasoning history, but that history is invisible. This is precisely the Yarncrawler problem—reconstruction from residual alone—applied to conversational AI.

A transparent interface would expose the reasoning trajectory as it develops, allowing the user to interrupt at the point where a wrong turn is taken rather than after the wrong turn has propagated through the entire analysis.

9.7 Admissibility and the "Fly Around" Idea

The suggestion that a system generating a long report could, during the generation period, "fly around to different planets"—present side discussions, related concepts, alternative interpretations, adjacent hypotheses—is a proposal for a specific form of admissibility display. The planets are the neighboring nodes of the admissibility field: regions of the conceptual space that are reachable from the current query but not on the direct path to the primary response.

This is consistent with the semantics of productive waiting. The user is not receiving the primary output; they are touring the admissibility field that surrounds it. When the primary output arrives, it is contextualized by the tour: the user already knows something about why certain interpretations were not chosen, what assumptions the analysis depends on, and how the primary response relates to its conceptual neighbors.

In the RSVP picture, this corresponds to exposing the lamphrodyne flow $\mathbf{v}(x, t)$: the direction and velocity of constraint propagation through the conceptual field,

not merely the final configuration. The user observes the dynamics, not merely the endpoint.

9.8 Scope and Limits

The Admissibility Display Principle is not a claim that all waiting should be eliminated or that interfaces should always be maximally verbose. Admissibility display has costs: it can produce cognitive overload, distract the user from ongoing tasks, or expose uncertainty that is resolved before the user has processed it. The principle is about the direction of design, not about a uniform maximum.

Nor is this chapter a claim that current AI interfaces are uniquely deficient. The spinner is widespread because it is simple, and many interaction contexts genuinely call for minimal interruption. The critique is structural: the spinner is a design choice that systematically prioritizes state display over admissibility display, and this choice has costs that the framework developed in this monograph makes precise.

The deeper point is that the monograph's central claim—that reachability is more informative than state—is not a claim about physics or biology alone. It is a claim about the structure of informative representations in general. The interface design case is an illustration of that generality: the same principle that explains why the cortical admissibility field matters more than the instantaneous firing rate, why the joint's developmental trajectory matters more than its adult cellular composition, and why the quasar's reverberation timescale matters more than its instantaneous luminosity, also explains why an interface that displays admissibility is more useful than one that displays only state.

The spinner is, in this sense, a small instance of a large problem.

Chapter 10

A Research Program

A theory that makes no predictions is not a theory but a vocabulary.

10.1 Twelve Falsifiable Predictions

10.1.1 Class I: Threshold and Relaxation in Neuroscience

P1. Graded tonic inhibition approaching but not reaching global silence will produce no SWA reduction, regardless of depth. *Test*: halorhodopsin protocol at variable depths; measure SWA as a function of silence fraction.

P2. Single brief crossings of Σ embedded in tonic inhibition will produce discrete SWA reduction proportional to crossing count. *Test*: isolated off periods embedded in the tonic protocol.

P3. The ΔP-versus-crossing-count relationship will be approximately linear, not saturating. *Test*: vary crossing number while holding total firing reduction constant.

10.1.2 Class II: Collapse Dynamics in the Joint

P4. Admissibility collapse rate under TNF/IL-1β will be a monotone function of baseline $PI16^+$ cell density across joint sites, holding inflammatory stimulus constant. *Test*: ex vivo stimulation of fetal joint tissue from PIP, DIP, and other joints; measure chemotaxis induction.

P5. Interventions preserving $PI16^+$ homeostatic function (WNT/BMP/FGF pathway stabilization) under inflammatory challenge will delay or prevent transition from relaxation to collapse. *Test*: $PI16^+$ protection assay under cytokine challenge with WNT/BMP agonists.

P6. The convergent lining phenotype ($SOX5^+$ $HBEGF^+$) can be restored from either interzone or non-interzone precursors under matched hypoxia and EGFR signalling,

restoring homeostatic admissibility regardless of precursor identity. *Test*: directed differentiation of both precursor types under matched conditions.

10.1.3 Class III: Bipartite Admissibility in Genomics

P7. High-heritability genes will show sharper admissibility boundaries in sequence space (more abrupt ISM functional-effect profiles) than low-heritability genes (Spiro et al., 2026). *Test*: ISM saturation mutagenesis across the two gene sets.

P8. The bipartite decomposition will be reproduced in DNA methylation: a minority of CpG sites will show sharp admissibility boundaries, consistent with partial generalization at high training set size (Spiro et al., 2026).

10.1.4 Class IV: The Critical Depth Threshold in RL

P9. The $D \geq 2$ threshold for PIT will be a genuine discontinuity: $D = 1$ and $D = 2$ will differ categorically in PIT magnitude (Okitsu and Sakai, 2026). *Test*: probabilistic depth sampling to interpolate between integer values.

P10. The plan-until-habit framework will reproduce Pavlovian-instrumental competition (PIT reduction with extensive Pavlovian training; increase with extensive instrumental training (Okitsu and Sakai, 2026)) across qualitatively different task architectures.

10.1.5 Class V: High-Redshift Reverberation Scaling

P11. Population-level measurement of the τ-λ slope α at $z > 4$ will show statistically significant deviation from $\alpha = 4/3$ (Shakura and Sunyaev, 1973), with direction discriminating RSVP from standard predictions. *Test*: Rubin/Roman era, $N > 100$ variable quasars at $z > 4$ (Fan et al., 2023).

P12. The variable SED spectral slope will be shallower (closer to slim-disk (Abramowicz et al., 1988)) at $z > 4$ than at $z < 2$ for matched Eddington ratios (Leung et al., 2026). *Test*: same Rubin/Roman sample.

10.2 The Open Mathematical Problem

The deepest open problem: *under what conditions does an admissibility field contain a threshold surface, and when does perturbation of the threshold maintainer produce collapse rather than relaxation?*

The empirical evidence suggests threshold surfaces arise when: (1) constraint accumulation is approximately irreversible on the accumulation timescale; (2) the system has a natural minimum-admissibility state; and (3) the relaxation mechanism requires visiting this minimum and returning (Driessen et al., 2026; Funk et al., 2017; Tononi and Cirelli, 2014).

Collapse arises when the threshold maintainer is subject to a perturbation that inverts its function (Croft et al., 2019; Davidson et al., 2026). A formal characterization of these conditions in terms of the RSVP field equations—conditions on the potential landscape of $\Phi(x, t)$ that generate threshold surfaces and threshold maintainers—is the primary mathematical problem for the next stage of this research program.

Chapter 11

Conclusions

This monograph began with six papers and a structural observation. It ends with a theoretical framework, formal definitions, a central hypothesis, a distinction between relaxation and collapse, a research program, and twelve falsifiable predictions.

The central results:

1. The admissibility field $\mathcal{A}(x, t)$ is the appropriate formal object for constraint-governed systems across all four empirical domains examined here (Davidson et al., 2026; Driessen et al., 2026; Leung et al., 2026; Okitsu and Sakai, 2026; Spiro et al., 2026; Tian et al., 2026).

2. Admissibility fields contain a bipartite structure: high-constraint $\mathcal{A}_H$ (threshold-mediated, slow-relaxing, history-sensitive) and free $\mathcal{A}_F$ (fast-relaxing, history-insensitive). The cortical data establish this at 67%/33% (Driessen et al., 2026); the arthritis data provide its mechanistic basis—*PI16*$^+$ fibroblasts as the cellular implementation of $\mathcal{A}_H$ maintenance (Davidson et al., 2026).

3. The Threshold Hypothesis (Hypothesis 3.1) establishes that relaxation of $\mathcal{A}_H$ requires threshold-surface traversal (Driessen et al., 2026). The $D \geq 2$ condition in the RL paper provides the most mathematically explicit instance (Okitsu and Sakai, 2026). Threshold surfaces appear in all six domains.

4. Admissibility collapse (Definition 2.7) is distinct from relaxation. Collapse occurs when the threshold-maintaining mechanism is functionally inverted. The *PI16*$^+$ fibroblast system is the clearest instance: from structural insulator to immune gateway under cytokine stimulation (Davidson et al., 2026), extending the fibroblast heterogeneity literature (Croft et al., 2019; Zhang et al., 2019) to a mechanistic admissibility framework.

5. Historical encoding operates across a hierarchy of timescales from hours to developmental epochs, with reversal timescales ranging from hours to never (Davidson et al., 2026; Driessen et al., 2026; Tononi and Cirelli, 2014). Convergent admissibility geometry—the same $\mathcal{G}_{\text{res}}(x)$ produced by different developmental trajectories (Davidson et al., 2026)—demonstrates that admissibility structure can be trajectory-independent.

6. The thin-disk reverberation relation $\tau \propto \lambda^{4/3}$ (Shakura and Sunyaev, 1973), confirmed consistent with RSVP field-propagation at $z = 6.51$ (Leung et al., 2026), provides a future probe of RSVP cosmological predictions at population scale (Burke et al., 2021; Fan et al., 2023).

The framework is falsifiable on twelve counts across five experimental domains and four orders of magnitude in spatial scale. That is the appropriate scope for a claim that admissibility geometry is not a metaphor but a mathematical structure with empirical content from cortical circuits to the cosmic dawn.

Bibliography

M. A. Abramowicz, B. Czerny, J. P. Lasota, and E. Szuszkiewicz. Slim accretion disks. *Astrophysical Journal*, 332:646, 1988.

Hannah Alfonsa, Edward M. Merricks, Darpan L. Bhatt, et al. Intracellular chloride regulation mediates local sleep pressure in the cortex. *Nature Neuroscience*, 26:64–78, 2023.

Żiga Avsec, Vikram Agarwal, Daniel Visentin, et al. Effective gene expression prediction from sequence by integrating long-range interactions. *Nature Methods*, 18: 1196–1203, 2021.

Bernard W. Balleine and John P. O'Doherty. Human and rodent homologs in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology*, 35:48–69, 2010.

Christopher D. Buckley, Caroline Ospelt, Steffen Gay, and Kim S. Midwood. Location, location, location: how the tissue microenvironment affects inflammation in RA. *Nature Reviews Rheumatology*, 17:195–212, 2021.

Colin J. Burke et al. A characteristic optical variability time scale in astrophysical accretion disks. *Science*, 373:789–792, 2021.

Edward M. Cackett et al. Supermassive black holes with high accretion rates in active galactic nuclei. XI. accretion disk reverberation mapping of Mrk 142. *Astrophysical Journal*, 896:1, 2020.

Anna L. Cheng, Erica R. H. Fuchs, and Jeremy J. Michalek. US industrial policy may reduce electric vehicle battery supply chain vulnerabilities and influence technology choice. *Nature Energy*, 9:1561–1570, 2024.

Rebecca E. Ciez and Jay F. Whitacre. Examining different recycling processes for lithium-ion batteries. *Nature Sustainability*, 2:148–156, 2019.

Adam P. Croft, João Campos, Kerstin Jansen, et al. Distinct fibroblast subsets drive inflammation and damage in arthritis. *Nature*, 570:246–251, 2019.

S. Davidson, D. Simone, K. Jansen, et al. The embryonic origins of site-specific arthritis. *Nature Immunology*, 2026.

Peter Dayan and Bernard W. Balleine. Reward, motivation, and reinforcement learning. *Neuron*, 36:285–298, 2002.

Carl G. de Boer and Jussi Taipale. Hold out the genome: a roadmap to solving the *cis*-regulatory code. *Nature*, 625:41–50, 2024.

Luisa De Vivo, Michele Bellesi, William Marshall, Eric A. Bushong, Mark H. Ellisman, Giulio Tononi, and Chiara Cirelli. Ultrastructural evidence for synaptic scaling across the wake/sleep cycle. *Science*, 355:507–510, 2017.

Amir Dezfouli and Bernard W. Balleine. Actions, action sequences and habits: evidence that goal-directed and habitual action control are hierarchically organized. *PLOS Computational Biology*, 10:e1003364, 2014.

Graham H. Diering, Raju S. Nirujogi, Richard H. Roth, Paul F. Worley, and Ananya Bhattacharyya. Homer1a drives homeostatic scaling-down of excitatory synapses during sleep. *Science*, 355:511–515, 2017.

Laura T. Donlin et al. Methods for high-dimensional analysis of cells dissociated from cryopreserved synovial tissue. *Arthritis Research & Therapy*, 20:139, 2018.

Kort Driessen, Fabio Squarcio, Giulio Tononi, and Chiara Cirelli. Induction of cortical on/off periods in awake mice fulfills sleep functions. *Nature Neuroscience*, 2026. doi: 10.1038/s41593-026-02318-9.

Xiaohui Fan, Eduardo Bañados, and Robert A. Simcoe. Quasars and the intergalactic medium at cosmic dawn. *Annual Review of Astronomy and Astrophysics*, 61:373–426, 2023.

M. M. Fausnaugh et al. Space Telescope and Optical Reverberation Mapping Project. III. optical continuum emission and broadband time delays in NGC 5548. *Astrophysical Journal*, 821:56, 2016.

Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11:127–138, 2010.

Caitlin M. Funk, Kaspar Peelman, Michele Bellesi, William Marshall, Chiara Cirelli, and Giulio Tononi. Role of somatostatin-positive cortical interneurons in the generation of sleep slow waves. *Journal of Neuroscience*, 37:9132–9148, 2017.

Martin Geissdoerfer, Paulo Savaget, Nancy M. P. Bocken, and Erik Jan Hultink. The circular economy: a new sustainability paradigm? *Journal of Cleaner Production*, 143: 757–768, 2017.

Samuel J. Gershman and Nathaniel D. Daw. Reinforcement learning and episodic memory in humans and animals. *Annual Review of Psychology*, 68:101–128, 2017.

Gavin Harper, Roberto Sommerville, Emma Kendrick, et al. Recycling lithium-ion batteries from electric vehicles. *Nature*, 575:75–86, 2019.

Chengfei Huang et al. Personal transcriptome variation is poorly explained by current genomic deep learning models. *Nature Genetics*, 55:2056–2059, 2023.

Reto Huber, M. Felice Ghilardi, Marcello Massimini, and Giulio Tononi. Local sleep and learning. *Nature*, 430:78–81, 2004.

Ted Jacobson. Thermodynamics of spacetime: the Einstein equation of state. *Physical Review Letters*, 75:1260–1263, 1995.

Alexander Karollus, Thomas Mauermeier, and Julien Gagneur. Current sequence-based models capture gene expression determinants in promoters but mostly ignore distal enhancers. *Genome Biology*, 24:56, 2023.

David R. Kelley. Sequential regulatory activity prediction across chromosomes with convolutional neural networks. *Genome Research*, 28:739–750, 2018.

Brandon C. Kelly, Jill Bechtold, and Aneta Siemiginowska. Are the variations in quasar optical flux driven by thermal fluctuations? *Astrophysical Journal*, 698: 895–910, 2009.

Mehdi Keramati, Amir Dezfouli, and Payam Piray. Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLOS Computational Biology*, 7: e1002055, 2011.

Gene C. K. Leung, Anna-Christina Eilers, Christos Panagiotou, Julien Wolf, Kishalay De, Luke Weisenbach, Minghao Yue, Xiaohui Fan, Yuzo Ishikawa, Erin Kara, Mirko Krumpe, Andrea Merloni, Robert A. Simcoe, Feige Wang, and Jinyi Yang. Discovery of quasar variability and early accretion disk signatures at cosmic dawn. *Nature Astronomy*, 2026. doi: 10.1038/s41550-026-02897-4.

Johannes Linder, Divyanshi Srivastava, Han Yuan, Vikram Agarwal, and David R. Kelley. Predicting RNA-seq coverage from DNA sequence as a unifying model of gene regulation. *Nature Genetics*, 2025. doi: 10.1038/s41588-024-02053-6.

Marcello Massimini, Reto Huber, Fabio Ferrarelli, Sean Hill, and Giulio Tononi. The sleep slow oscillation as a traveling wave. *Journal of Neuroscience*, 24:6862–6870, 2004.

Ian M. McHardy, Stephen D. Connolly, Keith Horne, et al. Mrk 509: Swift and ground-based observations. *Monthly Notices of the Royal Astronomical Society*, 480:2881–2896, 2018.

Robert B. Miller. Response time in man-computer conversational transactions. *Proceedings of the AFIPS Fall Joint Computer Conference*, 33:267–277, 1968.

Daisuke Miyamoto, Daijiro Hirai, Chi Chung Andy Fung, et al. Top-down cortical

input during NREM sleep consolidates perceptual memory. *Science*, 352:1315–1318, 2016.

Ida Momennejad, Evan M. Russek, Jin H. Cheong, Matthew M. Botvinick, Nathaniel D. Daw, and Samuel J. Gershman. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1:680–692, 2017.

Hagai Netzer. *The Physics and Evolution of Active Galactic Nuclei*. Cambridge University Press, 2013.

Jakob Nielsen. *Usability Engineering*. Academic Press, 1994.

K. Okitsu and Y. Sakai. A single reinforcement learning model to unify habit formation and Pavlovian-instrumental interaction. *Scientific Reports*, 2026.

Bradley M. Peterson. Reverberation mapping of active galactic nuclei. *Publications of the Astronomical Society of the Pacific*, 105:247–268, 1993.

Robert A. Rescorla and Allan R. Wagner. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black and W. F. Prokasy, editors, *Classical Conditioning II: Current Research and Theory*, pages 64–99. Appleton-Century-Crofts, 1972.

Alexander Sasse, Maria Chikina, and Sara Mostafavi. Benchmarking of deep neural networks for predicting personal gene expression from DNA sequence highlights shortcomings. *Nature Genetics*, 55:2060–2064, 2023.

N. I. Shakura and R. A. Sunyaev. Black holes in binary systems: observational appearance. *Astronomy and Astrophysics*, 24:337–355, 1973.

Ben Shneiderman. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Addison-Wesley, 1984.

Anna E. Spiro, Xinming Tu, Yilun Sheng, Alexander Sasse, Rezwan Hosseini, Maria Chikina, and Sara Mostafavi. A scalable approach to investigating sequence-to-function predictions from personal genomes. *Nature Methods*, 2026. doi: 10.1038/s41592-026-03124-8.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.

Xi Tian, Fei Peng, Jon McKechnie, Amir F. N. Abdul-Manan, Yaobin Liu, and Fanran Meng. Sustainable battery recycling through spatial and technological alignment. *Nature Sustainability*, 2026. doi: 10.1038/s41893-026-01851-6.

Giulio Tononi and Chiara Cirelli. Sleep and the price of plasticity: from synaptic and cellular homeostasis to memory consolidation and integration. *Neuron*, 81:12–34, 2014.

Giulio Tononi and Chiara Cirelli. Sleep and synaptic down-selection. *European Journal of Neuroscience*, 51:413–421, 2020.

Vladyslav V. Vyazovskiy, Chiara Cirelli, Myriam Pfister-Genskow, Ugo Faraguna, and Giulio Tononi. Cortical firing and sleep homeostasis. *Neuron*, 63:865–878, 2009.

Vladyslav V. Vyazovskiy, Umberto Olcese, Erin C. Hanlon, Yuval Nir, Chiara Cirelli, and Giulio Tononi. Local sleep in awake rats. *Nature*, 472:443–447, 2011.

B. Zhang et al. Lithium-ion battery recycling relieves the threat to material scarcity amid China's electric vehicle ambitions. *Nature Communications*, 16:6661, 2025.

Fan Zhang, Kevin Wei, Kamil Slowikowski, et al. Defining inflammatory cell states in rheumatoid arthritis joint synovial tissues by integrating single-cell transcriptomics and mass cytometry. *Nature Immunology*, 20:928–942, 2019.

Jian Zhou and Olga G. Troyanskaya. Predicting effects of noncoding variants with deep learning-based sequence model. *Nature Methods*, 12:931–934, 2015.