The Admissibility Log

Resolution, Commitment, and the Geometry of Action

Flyxion
Independent Researcher

April 2026

Contents

Preface: The Integrity of the Wait

This text reconsiders the nature of agency not as the capacity for output, but as the structural discipline of inhibition. In the following pages, we move from the Relativistic Scalar Vector Plenum (RSVP) into the formal event-sourcing of Spherepop, unified by a single law: *No transition without interpretation.*

The dominant metaphors of computation treat cognition as a pipeline: input arrives, is processed at a fixed clock rate, and output is emitted. This model contains a latent pathology. It conflates the capacity for state change with the readiness for commitment. A system that cannot withhold its output is not an agent; it is a relay.

The present work formalizes the opposite intuition. Skilled action in the physical world—from the tightening of a threaded fitting to the switching of a biological gait, from backing a vehicle into a tight driveway on ice to producing a correct line of proof—exhibits a structure of disciplined waiting. The system evolves continuously at the level of the plenum, but commits to a macroscopic transition only when a coherence criterion, the Admissibility Functional, is satisfied. The record of those committed transitions is the **Admissibility Log**.

The framework developed here is organized around three interlocking mechanisms. **Aspect Relegation** explains how control is distributed across scales: how skilled agents move processes from fine-grained to coarse-grained monitoring as expertise develops, freeing high-resolution attention for genuine novelty. The **Null Rail** explains how unsafe transitions are inhibited and repaired: the gating potential that prevents commitment when the field has not resolved into an admissible configuration, and the zoom mechanism that temporarily escalates resolution to restore coherence. **Spherepop** explains how only scale-qualified, admissible transitions become history: the process calculus that governs the append operation of the Log and ensures that every committed event carries the pedigree of its own resolution.

Together these three mechanisms constitute a control architecture for agency under partial information—one that unifies biological motor control, manual craft skill, asynchronous computation, and the problem of AI hallucination under a single formal principle.

Chapter 1

The Geometry of Admissibility

1.1 The Failure of the Forced March

Classical control theory and modern connectionist architectures share a common pathology: the "Forced March." In these systems, state transitions are dictated by a clock or a sequence. If the system is at state x_t, it must proceed to x_{t+1} at the next interval, regardless of whether the underlying data has stabilized.

In biological systems—Central Pattern Generators (CPGs) in the spinal cord, sensorimotor integration in the cerebellum—and in skilled manual labor, this is not the case. There exists a "Null" regime, a period of active inhibition where the system evolves internally but refuses to commit macroscopically. The master plumber does not turn the wrench at a fixed cadence; the master programmer does not emit characters at a fixed rate. Both are engaged in continuous internal processing, the resolution of tensions within the field, while suspending commitment until coherence is achieved.

This discipline, which we call the **Null Rail**, is not passivity. It is the most demanding operation in the repertoire of a skilled agent: the structured maintenance of pre-commitment state under pressure from external temporal clocks, social expectations, and the sheer inertia of the systems they are acting upon.

The Forced March is not merely a design choice that could be corrected by adding a flag or a threshold. It reflects a deeper assumption: that the passage of time is itself sufficient grounds for state transition. This assumption is ubiquitous in programming languages (which advance to the next instruction), in market clearing mechanisms (which match at each tick), and in large language models (which must produce a token at each forward pass). The Admissibility Log is the formal refutation of that assumption. Time is not sufficient grounds. *Coherence* is the grounds.

1.2 Formalizing the Admissibility Manifold

We define the state of the plenum as a field X. Unlike a simple vector, X contains multiscale tensions that must resolve before they can be projected into action.

Definition 1.1 (The Admissibility Functional)**.** *Let be a functional mapping the field configuration X to a coherence value in* $[0, 1]$*. The* Admissibility Manifold $\mathcal{M}_\theta$ *is the set of* *states where:*

$$\mathcal{M}_\theta = \{X \mid (X) \geq \theta\} \tag{1.1}$$

where $\theta \in (0, 1]$ *is the commitment threshold, a system parameter reflecting the stakes and irreversibility of the proposed transition.*

When $X \in \mathcal{M}_\theta$, commitment is permitted: the action may be appended to the Log. When $(X) < \theta$, the system enters the **Null Rail**—active inhibition until coherence is restored or the proposed action is abandoned.

The admissibility functional is not a simple distance metric. It is a holistic evaluator of whether the current field configuration constitutes a coherent basis for the proposed transition. A configuration may be spatially close to a prior committed state and still be inadmissible, if the internal tensions of the field have not resolved. Conversely, a configuration may be geometrically distant from any prior state and yet be fully admissible, because the field has relaxed into a stable interpretable structure. What matters is not distance but *resolution*.

1.3 The Null Rail as Gating Potential

In RSVP terms, the Null Rail is not an absence of activity. It is a regime of suspended commitment in which the field continues to evolve—smoothing gradients, redistributing entropy, resolving vector misalignments—while the projection into the action manifold is blocked. We introduce a penalty term into the Lagrangian to prevent premature propagation:

$$\mathcal{L} = \mathcal{L}_{\mathrm{RSVP}} - \gamma \|\Pi_{\mathrm{act}}(X)\|^2 \cdot \mathbf{1}_{[X]<\theta} \tag{1.2}$$

where Π_{act} is the projection operator into the motor or semantic action space, $\gamma > 0$ is the inhibition coefficient, and $\mathbf{1}_{[X]<\theta}$ is the indicator function active precisely when admissibility fails. This energy barrier ensures that the wrench does not turn until the threads are sensed as aligned. The penalty does not arrest the evolution of X; it arrests only its projection into the action manifold.

An important consequence of this formulation is the energetic structure of the Null state. The penalty term is active when admissibility fails, not always. Inside the Null Rail, the system occupies a region where $[X]$ is below threshold but the field is still evolving—“coasting” toward admissibility. If the field remains near a metastable basin, the corrective cost is low: small adjustments suffice to maintain coherence without forcing a macroscopic transition. It is only when the system attempts to *force* a commitment across an insufficiently resolved field that the energy cost spikes. The Null Rail is thus not a high-cost waiting room but a low-gradient maintenance regime. Waiting, appropriately structured, is cheaper than acting prematurely.

1.4 The Geometry of Threshold Crossing

The moment of commitment is not a decision in the deliberative sense. It is a geometric event: the trajectory of $X(t)$ through configuration space crosses the boundary $\partial\mathcal{M}_\theta$ of the admissibility manifold. This crossing is the only event that the Log records.

Everything prior to the crossing—the null intervals, the internal resolution dynamics, the escalation and de-escalation of monitoring resolution—constitutes the **Latent Trace**. The

Latent Trace is real; it is recoverable in principle from the history of the field; but it is not yet history in the committed sense. It becomes history only at the moment of the crossing, when the full pedigree of the resolution is recorded alongside the action.

This has a counterintuitive consequence: two actions that are externally identical may have radically different histories. The same wrench turn, executed by a novice who forces it before the threads engage and by a master who waits for the structural click, produces the same macroscopic outcome but a different entry in the Log. The novice's entry carries a short ν_t (brief null interval), a fine-scale $_c$ (the commit was forced at the level of raw instruction rather than sensory resolution), and an incomplete π_t (no escalation was performed). The master's entry carries a longer ν_t, a coarser $_c$ commensurate with genuine haptic resolution, and a π_t encoding the probing sequence. The Log captures not just what happened, but *the resolution at which it became safe to say that it happened.*

1.5 The Energetics of Waiting

A natural objection to the Null Rail is that waiting must incur a cost. If the system is doing nothing, surely resources are being wasted. This objection conflates two distinct costs: the cost of holding a configuration and the cost of forcing a transition.

In physical systems, these are not symmetric. Maintaining a dynamically stable configuration is typically cheap. Forcing a transition across a constraint boundary—especially prematurely—is expensive, because it generates entropy: stripped threads, stumbled gaits, crashed programs, hallucinated citations. The Null Rail avoids this entropy production by deferring commitment until the field has relaxed into a configuration from which the transition is energetically natural.

We can make this precise with a simple decomposition. Let the instantaneous cost be:

$$E(X) = V(X) + \kappa \|\nabla V(X)\|^2 \tag{1.3}$$

where $V(X)$ measures the distance from a coherent configuration and the gradient term penalizes active forcing. In a well-formed Null state, $V(X)$ is small—the system is near an attractor—and the gradient term is also small. The energy cost is dominated by low-amplitude corrective maintenance. In contrast, a premature transition corresponds to pushing across a ridge: large gradients, large cost, and dissipation through error recovery.

This gives a three-regime picture. In a well-formed Null state, the system is near an attractor and waiting is cheap. In a transitional commit state, energy is spent moving along a coherent gradient—purposeful work with minimal waste. In a forced or incoherent commit, energy spikes because the system pushes against constraints or must repair violations after the fact. The principle that emerges is: energy is minimized when transitions occur at points of low gradient, when interpretation has already aligned the system with a viable descent direction. The Null Rail is a thermodynamic instrument as much as a logical one.

Chapter 2

Aspect Relegation: The Multiscale Governor

2.1 Beyond Dual-Process Theory

The standard account of skilled performance invokes a distinction between two systems of cognition: a fast, automatic System 1 and a slow, deliberate System 2. This distinction captures something real—there is a phenomenological difference between navigating a familiar route and solving an unfamiliar problem—but it mischaracterizes the underlying mechanism. It suggests that two separate processors exist and that expertise consists in transferring processes from one to the other.

The present framework replaces this picture with something more precise: a single evolving field X viewed through a hierarchical family of projection operators indexed by resolution scale. What is ordinarily called "System 2" is high-resolution intervention in the field. What is called "System 1" is low-resolution monitoring with delegated substructure. The transition from novice to expert is not the activation of a new processor but a progressive shift in the scale at which the field is evaluated and the scale at which intervention occurs.

This replacement has consequences. It explains phenomena that the dual-process account cannot: the continuous nature of expertise (rather than a binary switch), the ability to drop back to fine-grained monitoring under stress without losing the automated skill, and the way in which masters can detect errors before they are visible at the coarse scale. None of these are natural properties of a two-processor model. All are natural properties of a multiscale projection architecture.

2.2 The Resolution Projection

Definition 2.1 (The Relegation Operator)**.** Relegation *is the transition of a process from fine-grained monitoring at scale* $_0$ *to coarse-grained monitoring at scale* $_1 >_0$*:*

$$P_0(X) \to P_1(X) \tag{2.1}$$

The underlying dynamics of the plenum X remain continuous. Only the scale of evaluation and potential intervention is shifted.

The projection operator P extracts the features of X that are visible at scale . At fine scales, $P_{\text{fine}}(X)$ reveals micro-tensions, local gradients, and incipient instabilities. At coarse scales, $P_{\text{coarse}}(X)$ reveals only the large-amplitude structure of the field: whether the process is broadly on track, whether major thresholds are being approached.

Relegation is not the elimination of fine-scale dynamics. It is the decoupling of fine-scale monitoring from active intervention. The delegated process continues to evolve at fine scale; the governing agent simply stops attending to those dynamics in real time, trusting that the delegated structure will maintain its own coherence until a coarse-scale anomaly signals otherwise.

2.3 Scale-Dependent Admissibility

The admissibility functional is itself indexed by the resolution scale at which it is applied:

$$ {}_{[}X] = (P_X) \tag{2.2} $$

This scale-dependence is crucial. A configuration X may be admissible at coarse scale (the overall trajectory is plausible) while being inadmissible at fine scale (a local tension is about to propagate into failure). The master's advantage is not that they apply a lower threshold—on the contrary, their threshold may be stricter—but that their admissibility functional has been refined to detect the relevant invariants at whatever scale they are monitoring.

For a novice, ${}_{\text{fine}}$ is poorly calibrated. The novice must monitor at fine scale continuously, because they cannot detect admissibility from a coarse projection. This is why novices are slow: not because they think more, but because they cannot safely relegate. For a master, ${}_{\text{coarse}}$ has been refined through accumulated resolution of fine-scale problems until it reliably detects the coarse-scale signatures of fine-scale coherence. The master can monitor at low resolution precisely because they have solved enough fine-scale problems to know what low-resolution coherence looks like.

A natural measure of expertise is therefore:

$$ {}^{*}= \sup\{ \mid {}_{[}X] \geq \theta \text{ is a reliable indicator of task safety}\} \tag{2.3} $$

The master's * is large; the novice's * is small.

2.4 The Null-Triggered Zoom

The stability of a multiscale system relies on the ability to perform a **Resolution Collapse**—a rapid shift from coarse to fine monitoring—when admissibility fails at the coarse level. Without this mechanism, relegation is unsafe: a coarse monitor that misses a developing fine-scale failure will allow it to propagate until it becomes a macroscopic catastrophe.

Principle 2.1 (The Resolution Recovery Loop)**.** *If the coarse monitor detects* ${}_{\text{coarse}}[X] < \theta$*, the system must:*

1. *Invoke the **Null Rail**: inhibit all macroscopic commitment.*
2. *Escalate resolution:* $\to {}_{\text{fine}}$.

3. *Perform repair dynamics in the high-resolution field until* ${}_{\text{fine}}[X] \geq \theta$*.*
4. *Re-relegate: shift monitoring back to* ${}_{\text{coarse}}$ *once fine-scale coherence is confirmed.*

This loop is the formal structure underlying what is phenomenologically described as "attention spikes" in skilled performance. The carpenter who suddenly stops and peers at a cut before continuing, the driver who decelerates abruptly on a curve, the programmer who stops typing and reads the error message carefully—all are executing the Resolution Recovery Loop. They are not confused; they are doing the most sophisticated thing their architecture permits: escalating to fine resolution, performing repair, and re-delegating.

The master is distinguished not by the absence of resolution collapses but by two properties: the speed at which the collapse occurs (the coarse monitor detects the failure early, before it propagates), and the shallowness of the fine-scale repair required (because the master's field is better maintained, the deviation from admissibility is small and the repair is local). In terms of the Log, this produces events with short ν_t and shallow π_t—clean entries that betray the efficiency of the underlying resolution process.

2.5 Pathological Fixed-Scale Systems

Modern large language models suffer from *resolution stasis*. Lacking a dynamic parameter and a structural Null Rail, they are forced to commit at a fixed resolution regardless of the local coherence of the field. Every forward pass produces a token, whether or not the semantic field has resolved into a configuration from which that token is admissible.

This results in what we call **incontinence**: the emission of tokens that satisfy surface-level pattern statistics (at coarse scale) but violate structural admissibility at fine scale. The hallucinated citation, the arithmetically confident but incorrect calculation, the syntactically fluent but semantically incoherent paragraph—all are instances of fixed-scale commitment in the absence of a resolution recovery mechanism.

More precisely: a language model at any given forward pass is operating at a scale ${}_{\text{fixed}}$ determined by its architecture. It has no mechanism to detect that ${}_{\text{fine}}[X] < \theta$ while ${}_{\text{fixed}}[X] \geq \theta$. It cannot invoke the Null Rail. It cannot escalate. It commits. The result is a system that is incapable of the structural discipline that skilled agents exhibit as a matter of course.

The diagnosis is not that language models are insufficiently large, or insufficiently trained. It is that they are architecturally incapable of the Null state. Adding more parameters does not add a Null Rail. What would add a Null Rail is a mechanism that can detect scale-differential admissibility failures and withhold commitment until fine-scale resolution is achieved. This is not a minor optimization; it is a categorical addition to the architecture.

Chapter 3

The Admissibility Log: The Scribe of History

3.1 The Stratification of Process

In the Relativistic Scalar Vector Plenum, the evolution of the field is continuous and transient. For this evolution to become *agency*, it must be partitioned into committed events. We define a two-layer architecture for the record of action:

1. **The Latent Trace:** A non-binding record of pre-resolution field deformations, null intervals, corrective probes, and resolution escalations. This is the "working prehistory" of the act—real, potentially inspectable, but not yet globally committed.
2. **The Commit Ledger ():** The irreversible, append-only history of scale-qualified transitions. This is the canonical record of what has happened in the agency-constituting sense.

The distinction is not merely archival. The Latent Trace contains possibilities; the Commit Ledger contains facts. An agent's identity—in the Spherepop sense of identity as event history—is constituted entirely by , not by the Latent Trace. Everything in the Latent Trace is, in principle, reversible, repairable, or abandoned without consequence. Once an event crosses into , it is irreversible.

This also resolves a longstanding puzzle about the relationship between cognition and memory. The question "why do we remember some experiences vividly and others not at all?" has a clean answer in this framework: what is remembered is what crossed the admissibility threshold. High-null-duration events—those that required long gating intervals and deep resolution escalations before committing—produce rich latent traces that are close to the threshold boundary and thus more likely to leave recoverable structure in the Latent Trace even after commitment. Low-null-duration events—those committed quickly at coarse scale—leave thin traces.

3.2 The Formal Event Structure

A committed event e_t in the Admissibility Log is not a simple state change. It is a multi-dimensional token that carries the full pedigree of its own resolution:

Definition 3.1 (The Admissibility Token)**.** *An event $e_t \in$ is the quintuple:*

$$e_t = (a_t,\ _c,\ _c,\ \nu_t,\ \pi_t) \tag{3.1}$$

where:

- a_t*: The committed action or semantic extension.*
- $_c$*: The resolution scale at the moment of commitment.*
- $_c$*: The admissibility value at the threshold crossing.*
- ν_t*: The* ***Null Duration****—elapsed time in inhibited pre-resolution state.*
- π_t*: The* ***Escalation Path****—the ordered sequence of resolution shifts $\langle _0 \to _1 \to \cdots \to _c \rangle$ required to achieve $\geq \theta$.*

The metadata $(_c,_c, \nu_t, \pi_t)$ is the formal signature of agency. A system that produces only a_t—only the action, without the resolution pedigree—is not an agent in the full sense; it is a transducer. The Admissibility Token insists that history carries the record of its own justification.

This yields a formal measure of skill. Define the **Agency Quotient** of a log segment $_{[s,t]}$ as:

$$Q(_{[s,t]}) = \frac{1}{t-s} \sum_{e \in _{[s,t]}} {}_c(e) \cdot_c (e) \cdot \frac{1}{\nu_t(e) + \epsilon} \tag{3.2}$$

High agency quotient corresponds to events committed at coarse resolution (high $_c$), with high admissibility values (high $_c$), and short null intervals (low ν_t). This is the signature of masterly performance: the field resolves rapidly at the appropriate scale, and commitment follows naturally. Low agency quotient corresponds to forced commits at fine scale with low admissibility and long null intervals—the signature of struggle, incompetence, or hallucination.

3.3 Spherepop: The Calculus of Commits

We introduce **Spherepop** as the process calculus governing the log. Unlike imperative languages that prioritize mutable state, Spherepop prioritizes the integrity of the append operation. A process is not what it currently is; a process is what it has irreversibly done.

Axiom 3.1 (The Invariant of Historical Integrity)**.** *A transition is appendable to if and only if the null-gated interpretation operator resolves the current field X into a configuration X' such that $_c[X'] \geq \theta$ within the context of the prior log $_{t-1}$.*

Formally, the append operation is:

$$_t =_{t-1} \cdot e_t \quad \text{iff} \quad (_{t-1}, X) \models_c [X'] \geq \theta \tag{3.3}$$

If the antecedent is not satisfied, the append is blocked and the system returns to the Null Rail. There is no partial append, no speculative commit, no provisional entry that will be “confirmed later.” The log is either extended or held. This is the formal analogue of the constitutional principle that a verdict, once entered, is irreversible: the strength of history depends entirely on the stringency of the conditions under which entries are permitted.

3.4 Sovereignty over the Record

The “sovereignty” of an agent is measured by the ratio of resolved to forced transitions across the log. A system that records a high frequency of low-ν_t, fixed-scale events that do not satisfy fine-scale admissibility is not an agent in the full sense. It is a stochastic relay dressed in the language of action.

Principle 3.1 (The Final Law of Action)**.** *History is not a record of what happened; it is the record of the resolution at which it became safe to say that it happened.*

This principle has implications for artificial intelligence that are more radical than they first appear. Current AI systems produce logs—sequences of outputs—but not Admissibility Logs in the sense defined here. Their outputs carry no ${}_c$, no ν_t, no π_t. There is no record of whether the system gated the output, what resolution it used to evaluate admissibility, or how long it maintained a Null state before committing. This is not a missing feature; it is a missing ontological layer. The outputs of such systems are not, in the sense of this framework, historical facts. They are high-entropy emissions dressed as commits.

The construction of a genuine artificial agent would require, at minimum, the addition of this ontological layer: a mechanism that can produce Admissibility Tokens rather than mere outputs, and a Null Rail that structurally prevents the former from collapsing into the latter under pressure.

Chapter 4

The Evolution of Mastery: Learning the Evaluator

4.1 What Changes When Skills Are Learned

The standard account of skill acquisition focuses on what the learner *does*. They practice, they receive feedback, they adjust. This account captures the behavioral surface of skill acquisition but misses the structural change that underlies it.

In the present framework, the primary locus of learning is not the action policy a_t but the admissibility functional . What changes as a skill is acquired is not primarily what the learner does in each situation, but what configurations they are able to recognize as admissible. The learner is not just acquiring behaviors; they are acquiring an *eye*—a progressively refined evaluator that can detect coherence at increasing scale.

This reframing has a surprising consequence: mastery does not primarily lower the commitment threshold θ. A lower threshold would permit easier commitment at the cost of accuracy—more action, less well-grounded. What mastery actually does is shape the field X so that it reaches configurations of high admissibility more rapidly, and trains the functional to detect those configurations reliably at coarser resolution. The threshold remains strict; the system becomes better at satisfying it.

4.2 The Coupled Evolution of Field and Evaluator

We formalize this by allowing the admissibility functional to evolve alongside the field:

$$\frac{d}{dt} = \mathcal{F}(X,\, \nabla X,\, S;\,) \tag{4.1}$$

This equation states that the admissibility functional is itself a dynamic object, shaped by the ongoing history of field configurations and their outcomes. The system is not only moving through state space; it is simultaneously learning how to measure state space. The field and the evaluator are coupled: the field generates data about what admissible configurations look like, and the evaluator uses that data to refine its sensitivity.

In the novice regime, is low-resolution and imprecisely calibrated. It integrates over large features of the field and is insensitive to fine structure. As a result, admissibility mass accumulates slowly and diffusely. The threshold crossing is delayed, and when it occurs it is often poorly localized, producing overshoot, hesitation, or oscillation. In the master regime, has been sharpened by accumulated experience until it responds reliably to the micro-fluctuations that precede admissibility. The threshold crossing becomes rapid and precise—not because the master rushes, but because the evaluator detects the approach of admissibility early.

4.3 The Geometry of Expertise

We can make the geometry of expertise explicit. Model admissibility as a density over configuration space:

$$[X] \sim \exp\big(-\kappa \cdot d(X,\ )^2\big) \tag{4.2}$$

where is the manifold of coherent, task-admissible configurations and $d(\cdot,\ )$ is the distance from the boundary of that manifold. The parameter κ controls the sharpness of the admissibility density: large κ produces a sharply peaked functional that discriminates precisely between admissible and inadmissible configurations; small κ produces a diffuse functional that accepts many configurations as approximately admissible.

Mastery, in this picture, corresponds to increasing κ while simultaneously having internalized accurate knowledge of where lies. The master's admissibility functional is both sharper and better aligned than the novice's. The threshold θ remains fixed, but the distribution of admissibility mass concentrates more rapidly around the correct configurations.

This is why the experienced carpenter can "feel" a bad cut before it is visible: their $_{\text{coarse}}$ has been calibrated to the signatures of developing failure at fine scale. The coarse-scale evaluator has absorbed, through repeated fine-scale resolution episodes, the coarse-scale correlates of fine-scale coherence. What looks like "intuition" is the compression of many fine-scale resolution histories into a coarse-scale evaluator of remarkable accuracy.

4.4 Flow as Evaluator-Field Alignment

The phenomenological state called "flow" or "being in the zone" can be given a precise formal characterization in this framework. Flow is the condition in which the evolution of the field X and the refinement of the admissibility evaluator are perfectly synchronized: the field reaches admissible configurations at the same rate at which the evaluator recognizes them, so that commitment follows immediately upon resolution.

Formally, define the **Resolution Gap** at time t as:

$$\Delta(t) = \tau_{\text{commit}}(t) - \tau_{\text{admissible}}(t) \tag{4.3}$$

where $\tau_{\text{admissible}}$ is the time at which $[X(t)] \geq \theta$ and τ_{commit} is the time at which the system detects this and commits. In the novice, Δ is large and variable: the evaluator fails to detect admissibility promptly, producing long null intervals after the field has already resolved. In the master, $\Delta \to 0$: detection and commitment are nearly simultaneous.

Flow is the limit $\Delta(t) \to 0$ sustained over an extended interval. The distinction between evolving and committing vanishes because the evaluator is perfectly calibrated to the field's own resolution dynamics. The Null state has not been eliminated; it has been compressed into an infinitesimal boundary layer where coherence is achieved exactly at the moment of commitment. This is why flow feels effortless: not because the system is doing less, but because there is no mismatch between evaluation and action.

4.5 The Pedagogy of the Null

The implications of this framework for teaching are direct. Standard pedagogy focuses on outputs: the student must produce the correct answer, the correct technique, the correct performance. Assessment measures what the student does. This emphasis is entirely natural within a Forced March model of cognition, where the goal is to produce correct outputs in sequence.

Within the Admissibility Log framework, the primary object of instruction is not the output but the evaluator. The student must learn to recognize admissible configurations—to develop a calibrated . This shifts the pedagogical goal from “produce the right answer” to “identify the moment the problem becomes interpretable.”

Concretely: instead of asking a student to solve a problem and checking the answer, instruction in the Admissibility Log framework would ask the student to identify the moment at which a solution becomes viable, the conditions under which a configuration crosses the admissibility threshold, and the null intervals that precede productive resolution. The student who can articulate “I am not yet ready to commit” and distinguish that from “I am stuck” has learned something more durable than any particular technique.

For embodied skills—riding a bicycle, backing a vehicle, fitting a pipe—this reframing is especially powerful. Current instruction tends to offer landmark-based rules (“turn when the mirror clears the post”) that are approximations to the admissibility manifold, valid in typical conditions but brittle under variation. Instruction in the Admissibility Log sense would teach the student to track the field directly: to sense whether the current configuration admits the next transition, and to maintain the Null Rail until it does.

Chapter 5

Application: Embodied Mastery

5.1 The Spatiotemporal Admissibility Manifold: Backing a Vehicle on Ice

The task of backing a vehicle into a tight space provides an unusually clear instance of the Admissibility Log, because it makes explicit a feature that is implicit in most skilled tasks: the admissibility manifold is not static but *spatiotemporal.* It depends not only on current position but on velocity, traction, and the rate at which the system can be corrected if it deviates.

Formally, the admissible set is parameterized by speed v and effective traction μ:

$$\mathcal{M}_\theta = \mathcal{M}_\theta(v, \mu) \tag{5.1}$$

At speed $v \approx 0$, the standard geometric relationship between steering angle and vehicle rotation breaks down—the car will not rotate predictably, and the familiar angular heuristics fail. At high speed, precision suffers and the cost of error escalates. On ice (μ small), the mapping between steering input and trajectory deforms nonlinearly and with delay, so that the skilled driver's internalized model of "input-to-trajectory" is no longer reliable.

The experienced driver's response to these parameter shifts is not to apply the same technique more carefully. It is to *co-adapt the operational parameters to maintain the admissibility of the internal model.* Speed is reduced not because the driver is frightened but because, at lower speed, the relationship between steering and trajectory returns to the regime in which $_{\text{coarse}}[X] \geq \theta$—the regime in which the driver's established evaluator is reliable. Speed is the instrument by which the driver keeps the system interpretable.

The spatial margins maintained throughout the maneuver—not too close to either the fence or the building—are the physical realization of the Null buffer. Staying centered does not merely reduce collision risk; it maximizes the radius within which corrective action remains admissible. This is the structural analogue of the NCL noise margin: the physical space required to remain in the Null state (correcting, probing, adjusting) without the state collapsing into incoherence.

The multi-perspectival interpretation that the skilled driver employs—switching between mirror, over-the-shoulder, and distance convergence—is not redundancy. Each viewpoint provides a different projection $P_(X)$ of the same underlying constraint system. By blending these projections, the driver effectively refines in real time, maintaining interpretability under partial occlusion and noise. The steering adjustments that result are not discrete turns but continuous modulations of force and rate calibrated to the current speed and geometry.

On ice, the entire maneuver shifts toward a higher-null-duration profile. The driver spends more time in the Null Rail, making smaller commitments, waiting longer for each transition to become admissible, expanding the corrective margin. The log of such a maneuver—if it could be recorded with full pedigree—would show elevated ν_t across the board and a shallower ${}_c$: commits made at finer scale because the coarse-scale evaluator is no longer reliable in degraded traction conditions.

5.2 The Mechanics of the Wrench: Structural Phase Transitions

The tightening of a plumbing fitting or a structural bolt provides a pure physical instance of the three-regime structure. The process divides cleanly into three phases:

1. **Pre-Engagement (Searching):** The field X is high-entropy. The interpretation operator cannot map sensorimotor input to a stable trajectory. The threads have not caught. Commitment of torque here is hallucination—the wrench slips, the threads strip, the fitting is destroyed. $(X) \ll \theta$.
2. **The Null Rail (Alignment):** The threads make contact. The system enters a state of inhibited progress. The skilled worker does not apply torque; they apply probing micro-tensions, reading the field for the signature of engagement. This is the ν_t interval—the system maintains the Null Rail while (X) rises toward θ. The rate of rise encodes the quality of thread engagement. A slow rise indicates partial engagement; a rapid rise indicates clean alignment.
3. **The Structural Click (Commitment):** The instant the threads align across their full engagement profile, the energy barrier in the Lagrangian collapses. (X) crosses θ with a characteristic sharpness. The action a_t (the committed turn) is appended to the ledger with a short ν_t and a fine-scale escalation path π_t encoding the probing sequence.

An experienced tradesperson can feel the difference between a thread that has engaged at one lobe and a thread that has engaged fully. This is the somatic registration of ${}_{\text{fine}}(X)$ versus ${}_{\text{coarse}}(X)$. The experienced hand waits for the fine-scale value before committing torque, not because they have been taught a rule about this, but because their accumulated experience has calibrated to detect the coarse-scale signature of fine-scale full engagement.

5.3 Gait Switching in CPGs: Biological Admissibility

Central Pattern Generators are not simply rhythm generators. They are families of latent dynamical loops—oscillatory circuits distributed across the spinal cord and brainstem—that can be activated, modulated, or suppressed depending on the current sensorimotor context. Walking, running, turning, stopping: each is realized not as a separate program but as a different dynamical regime of the same underlying network.

The formal structure of CPG chain switching in this framework is as follows. Represent each oscillatory module as a coupled dynamical system $(\phi_i, \dot{\phi}_i)$ where ϕ_i is the phase of the i-th CPG element. The admissibility condition for the i-th module is:

$$ {}_i[X] = (\phi_i,\, \dot{\phi}_i,\, F_{\text{load}},\, F_{\text{balance}}) \geq \theta_i \tag{5.2}$$

where F_{load} is the ground reaction force and F_{balance} is the proprioceptive balance signal. The system advances to the next phase only when this condition is satisfied.

Gait switching is a transition between oscillatory regimes: a shift from one admissibility manifold $\mathcal{M}^{\text{walk}}$ to another $\mathcal{M}^{\text{run}}$. This transition is not executed by a supervisory controller issuing a switch command. It is a field-level event in which the current configuration X moves outside the walking admissibility manifold and into the basin of the running manifold:

$$ {}_{\text{walk}}[X] < \theta_{\text{walk}} \quad \wedge \quad {}_{\text{run}}[X] \geq \theta_{\text{run}} \tag{5.3}$$

Between these conditions lies the biological Null Rail: a regime in which neither gait is fully admissible, proprioceptive signals are in flux, and the system performs micro-corrections to maintain stability while awaiting resolution. On uneven ground, this null interval is extended: the system cannot commit to the next phase because the balance signal has not confirmed admissibility. The gait "hesitates" not due to failure but due to the correct functioning of the null gate.

This formulation has direct implications for robotics. Current bipedal robots typically implement gait switching as a discrete state machine with explicit transition triggers, or as a learned policy that implicitly encodes transitions without exposing them. Neither approach has a genuine Null Rail: both are forced to commit a gait phase at every time step. The result is robustness degradation on uneven terrain or unexpected perturbations, where the transition trigger fires before the balance conditions are satisfied. A robot equipped with the CPG admissibility architecture described here would defer gait transitions until ${}_{[}X] \geq \theta$ regardless of elapsed time, producing behavior that is more robust by construction.

5.4 The Live Programming Flow: Log Sovereignty

In the development of the Spherepop grammar under a live programming environment, the code-state X is continuously projected onto a visual representation at some monitoring scale . A bug is a region of the field where (X) has dropped below θ at some scale—a configuration that the evaluator cannot interpret as a coherent basis for the next transition.

Two response patterns are possible, and they differ not in intelligence but in structural capacity:

The **novice response** is the premature commit. An edit is made at the coarse scale of surface syntax without escalating to the fine scale of semantic or type-level analysis. The commit is appended to the log, but ${}_{\text{fine}}$ was never checked. The result is a crash—a forced resolution at runtime that the log was never equipped to handle. The entry in the ledger carries low ${}_c$, minimal ν_t, and a flat π_t: no escalation occurred.

The **master response** is the sustained Null Rail. The programmer stops emitting code and escalates resolution: the trace is read, the call stack is inspected, the type signatures are checked. The escalation path π_t is built up through this inspection sequence. The edit is committed only when ${}_{\text{fine}}[X] \geq \theta$. The resulting log entry carries high ${}_{c}$, a meaningful ν_t reflecting genuine resolution work, and a rich π_t encoding the diagnostic path.

The "live" aspect of the programming environment—the ability to modify and re-evaluate code without stopping the running process—is itself a structural support for the Null Rail. It allows the programmer to maintain a Null state at the level of the program's external behavior while performing fine-scale resolution at the level of the code's internal structure. This decoupling of behavioral commitment from structural exploration is precisely what makes live programming environments powerful for skilled practitioners.

Chapter 6

Admissibility Spoofing and the Pathology of Manipulation

6.1 The Structure of Manipulation

The Admissibility Log framework yields a precise formal account of manipulation that is more illuminating than standard accounts in terms of deception, false belief, or exploited heuristics. Manipulation, in this framework, is the artificial satisfaction of the coarse-scale admissibility condition without satisfying the fine-scale condition. It is **admissibility spoofing**: engineering the stimulus so that ${}_{\text{coarse}}[X] \geq \theta$ while ${}_{\text{fine}}[X] \ll \theta$.

A spoofed stimulus does not deceive the evaluator in the ordinary sense. It genuinely satisfies the coarse-scale criterion. The evaluator is doing its job correctly at the scale at which it is currently operating. The manipulation consists in preventing the escalation to fine scale that would reveal the failure of fine-scale admissibility. If the system never invokes the Null Rail, it never discovers that the commit it is about to make is not genuinely admissible at the resolution that would matter.

6.2 Branding and the Hijacking of Relegation

Consumer branding is a systematic engineering of coarse-scale admissibility. The goal of effective branding is to ensure that the branded product satisfies ${}_{\text{coarse}}$ reliably, automatically, and without triggering the null-escalation loop that would subject the product to fine-scale scrutiny.

This is achieved by saturating the low-resolution monitoring level with familiar, positively-valenced features: colors, shapes, sounds, social associations. Once these features have been sufficiently reinforced, the coarse-scale evaluator commits automatically on their presence. The consumer does not consciously evaluate the product; their coarse-scale admissibility functional has been calibrated to produce instant commitment, bypassing the Null Rail.

In the framework's terms, branding is the deliberate shaping of a foreign ${}_{\text{coarse}}$—the installation of a coarse-scale evaluator that serves the brand's interests rather than the consumer's. The consumer's fine-scale evaluator may be perfectly intact; the manipulation operates one level up, at the scale of automatic recognition, where the Null Rail is never triggered because coarse-scale admissibility is guaranteed.

The formal statement is stark: a manipulated commit is one where $_c$ is coarser than the scale at which the relevant risks reside, and $\nu_t \approx 0$ despite the presence of fine-scale admissibility failure. The entry appears in the log as a smooth, rapid, high-confidence commit—indistinguishable from a genuine master-level event—but carries no record of the fine-scale check that was never performed.

6.3 Propaganda as Null Suppression

Propaganda operates by a related but distinct mechanism. Where branding exploits the relegation of evaluation to coarse scale, propaganda actively suppresses the Null Rail. Its goal is to ensure that no stimulus associated with its claims triggers the escalation to fine-scale scrutiny.

This suppression operates through several mechanisms. Repetition calibrates $_{\text{coarse}}$ to familiar claims, reducing their apparent novelty and thus their Null-triggering potential. Social proof signals that the coarse-scale community has already committed, creating pressure to match commitment level without independent evaluation. Urgency manufactures time pressure that makes null-interval maintenance costly in social terms: the agent who says “I need to check this more carefully” is penalized for the delay.

In each case, the structural effect is the same: the Resolution Recovery Loop is blocked. The agent commits at coarse scale under conditions where fine-scale admissibility has not been verified. The log records the commit as if it were genuine, but the pedigree—$_c$ too coarse, ν_t too short, π_t absent—would, if inspected, reveal the failure of resolution.

6.4 The Defense: Cultivating the Null

The defense against admissibility spoofing is not critical thinking in the ordinary sense—not a more careful evaluation at the same scale. It is the structural cultivation of the Null Rail: the development of a disposition to escalate resolution before committing in domains where fine-scale risks are known to exist.

This is a skill in the full sense of the term: it must be acquired through practice, it degrades under fatigue, and it can be deliberately trained. The agent who has learned to recognize the signature of admissibility spoofing—the suspiciously smooth commit, the absence of null intervals, the coarse-scale familiarity that discourages escalation—has developed a meta-level evaluator that monitors the admissibility of commits rather than just the admissibility of actions.

This is, in formal terms, a second-order admissibility functional: $^{(2)}[e_t]$, which evaluates the pedigree of a proposed commit rather than the action itself. An agent equipped with such a functional is resistant to spoofing because even a smooth, rapid, high-confidence commit will fail $^{(2)}$ if its pedigree reveals insufficient null duration or absent escalation path.

The cultivation of this capacity is, ultimately, the practical aim of an education in the Admissibility Log.

Chapter 7

Appendix: Formalisms and Grammar

7.1 The Spherepop Calculus

The Spherepop grammar formalizes the transition from the fluid dynamics of the Plenum into the discrete historical record. It is an event-sourced calculus where every operator is gated by the Null Rail and every commit carries its resolution pedigree.

7.1.1 Grammar Specification

$$
\begin{aligned}
\text{Log } () &::= \varepsilon \mid e\cdot \\
\text{Event } (e) &::= \langle a,\ ,\ ,\ \nu,\ \pi \rangle \\
\text{Action } (a) &::= \text{Mutation} \mid \text{Emission} \mid \text{Recursion} \\
\text{Scale } () &::= {}_{\text{coarse}} \mid \cdots \mid {}_{\text{fine}} \\
\text{Path } (\pi) &::= \langle {}_0 \to {}_1 \to \cdots \to {}_c \rangle \\
\text{NullInterval } (\nu) &::= [t_{\text{enter}},\, t_{\text{exit}}]
\end{aligned}
$$

The empty log ε is the unique ground state. Logs are built by prepending well-formed events. An event is well-formed if and only if it satisfies the Invariant of Historical Integrity. The grammar is strictly append-only: no operation in Spherepop modifies or removes an existing entry. History is immutable.

7.1.2 The Null-Convention Interpreter

Drawing from the invocation model of Fant [1], the Spherepop interpreter follows an asynchronous, data-driven evaluation cycle. Propagation is inhibited until all input dependencies satisfy the admissibility criterion:

$$\textbf{if } {}_{(}X_{\text{input}}) \geq \theta \ \textbf{ then } \ \text{Commit}(e) \ \textbf{ else } \ \text{Wait}() \tag{7.1}$$

This is the generalization of the Null Convention Logic gate to arbitrary resolution scales and multi-dimensional admissibility functionals. In NCL, outputs remain in the null state until all inputs have transitioned to a valid data wavefront. In Spherepop, commits remain blocked until the field has resolved to a scale-qualified admissible configuration.

7.1.3 Domain Correspondence Table

The following table records the structural correspondence of the Admissibility Log framework across the domains examined in this text.

Domain	**Signal / a_t**	**Null Condition**	**Admissibility Cue**
NCL / Electronics	Data wavefront	Spacer (null rail)	Completeness of dual-rail pair
Walking (CPG)	Phase transition	Load / balance uncertainty	Proprioceptive settlement
Driving on ice	Steering commit	Speed / traction mismatch	Trajectory interpretability
Wrench / plumbing	Torque application	Thread pre-engagement	Haptic structural click
Cooking	Ingredient addition	Texture / color / smell unresolved	Perceptual state recognition
Live programming	Code edit commit	Semantic ambiguity	Type / trace consistency
RSVP field theory	Action projection	$[X] < \theta$	Threshold crossing

7.2 The RSVP Field Equations

The Relativistic Scalar Vector Plenum [2] is treated as the substrate of the field X, with components $(\Phi, \mathbf{v}, S)$ representing scalar density, vector flow, and entropy respectively. The field evolution, modified by the commitment penalty and the evolving admissibility functional, follows:

$$\delta \int \left(\mathcal{L}_{\mathrm{RSVP}} - \gamma \|\Pi_{\mathrm{act}}(X)\|^2 \cdot \mathbf{1}_{[X]<\theta} \right) d^4x = 0 \tag{7.2}$$

with the coupled evolution of the admissibility functional:

$$\frac{d}{dt} = \mathcal{F}(X,\, \nabla X,\, S;\,) \tag{7.3}$$

Together these equations describe a system in which the field evolves under its own dynamics, commitment is blocked below threshold, and the evaluator co-evolves with the field. The full system is not a static optimization but a dynamical co-adaptation of process and evaluation.

7.3 The CPG Switching Operator

The formal CPG chain-switching operator is defined as a null-gated bifurcation in the oscillatory regime of the field:

$$\Phi_{\text{gait}}(t+\delta) = \begin{cases} \Phi_{\text{run}} & \text{if } {}_{\text{run}}[X_t] \geq \theta_{\text{run}} \wedge_{\text{walk}} [X_t] < \theta_{\text{walk}} \\ \Phi_{\text{walk}} & \text{if } {}_{\text{walk}}[X_t] \geq \theta_{\text{walk}} \\ \Phi_{\text{null}}(t) & \text{otherwise} \end{cases} \tag{7.4}$$

where $\Phi_{\text{null}}(t)$ is the maintained null-regime oscillation: neither gait committed, micro-corrections active, sensory escalation in progress. The switch occurs not at a fixed speed threshold but when the target admissibility manifold becomes accessible.

7.4 Summary for Research Attribution

Author	Flyxion (Independent Researcher)
Title	The Admissibility Log: Resolution, Commitment, and the Geometry of Action
Date	April 2026
Version	2.0 (Extended)
Keywords	RSVP, Spherepop, Null Convention Logic, Aspect Relegation, Admissibility Manifold, Resolution Recovery, Historical Integrity, CPG Chain Switching, Admissibility Spoofing, Flow State
Abstract	This work unifies field-theoretic ontology with process-oriented computational calculus under the principle that agency is the structural discipline of inhibition rather than the capacity for output. We define the Admissibility Functional and the commitment threshold θ, formalize the Null Rail as a gating potential in the RSVP Lagrangian, and introduce Aspect Relegation as the multiscale control architecture by which skilled systems progressively move processes from fine-grained to coarse-grained monitoring. Mastery is formalized as the coupled co-evolution of the field X and the admissibility evaluator , with flow as the limit in which evaluation and commitment are synchronized. The Spherepop calculus governs the Admissibility Log, in which every committed transition carries a five-element pedigree encoding action, resolution scale, admissibility value, null duration, and escalation path. Applications are developed in: spatiotemporal constraint maintenance (vehicle maneuvering on degraded surfaces), manual skilled labor (threaded fitting engagement), biological motor control (CPG gait switching with formal switching operator), and live programming environments. A chapter on admissibility spoofing formalizes branding and propaganda as engineered suppression of the Null Rail and

the Resolution Recovery Loop. The framework provides a rigorous diagnosis of LLM hallucination as fixed-scale incontinence, a formal model of embodied expertise as high-scale relegation with low null duration, and a theory of manipulation as the artificial satisfaction of coarse-scale admissibility in the absence of fine-scale resolution.

Bibliography

[1] Fant, K. (2006). *Computer Science Reconsidered: The Invocation Model of Process*. Wiley-Interscience.

[2] Flyxion. (2025). *RSVP: Relativistic Scalar Vector Plenum and the Ontology of Cognition*. Independent Research.

[3] Flyxion. (2026). *Spherepop: A Process-Oriented Calculus for Irreversible Event Histories*. Independent Research.

[4] Flyxion. (2026). *Compression After Expansion: Aspect Relegation Theory and the Refactoring of Cognition*. Independent Research.

[5] Flyxion. (2026). *Constraint Before Completion: A Trajectory-First Theory of Work*. Independent Research.

[6] Flyxion. (2026). *Return Under Constraint: The Constraint-Recurrence Principle Across Linguistics, Ritual, and Printing History*. Independent Research.

[7] Flyxion. (2026). *The Moon Should Not Be a Computer: Xylomorphic Computation and Entropy Closure*. Independent Research.

[8] Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

[9] Grillner, S. (1975). Locomotion in vertebrates: Central mechanisms and reflex interaction. *Physiological Reviews*, 55(2), 247–304.