

Boundaries Are Computational Objects

Screening, Pursuit, and Refusal as Three Theories of What a Boundary Does

Flyxion

July 2026

Abstract

Three independent research programs have converged on the claim that the boundary of a living or cognitive system is not a passive geometric fact about where one region of space ends and another begins, but an object that does computational work. The free-energy principle formalizes an organism's boundary as a Markov blanket: a statistical partition inducing conditional independence between internal and external states. Developmental bioelectricity formalizes a tissue's boundary as a bioelectric pattern-memory layer that stores a target morphology and drives correction toward it. A separate formalization of refusal, developed within an admissibility-theoretic account of irreversible computation, treats a commitment's boundary as the deletion of a region of future trajectory-space from what remains admissible at all. This paper argues that these are not three notations for one underlying mechanism. They are three different mechanisms, screening, pursuit, and deletion, answering three different questions about what a boundary does, and at least one pair of them is in genuine tension rather than harmless variation: refusal, formalized as an inadmissibility prior that sets the support of a trajectory distribution to zero, is shown by its own source theory to be incompatible with free-energy minimization unless an active-inference generative model is extended with commitment variables as latent structure. The paper develops each account on its own terms, states the point of friction precisely, and argues that the right response is not to unify the three mechanisms under a single metaphor but to treat boundary-formation as a genuine category of computational object with more than one known implementation.

Contents

1	Boundaries That Compute	2
2	Friston: The Boundary as Statistical Screen	3
3	Levin: The Boundary as Stored Pursuit	3
4	RSVP and Refusal: The Boundary as Admissibility Deletion	4
4.1	Refusal as Deletion, Not Preference	5
5	Screening, Pursuit, and Deletion Are Not the Same Operation	6
6	Where the Accounts Actually Conflict	6
7	A Category, Not a Metaphor	8
	References	9

1 Boundaries That Compute

A boundary, in the everyday sense, is where one thing stops and another begins: a geometric or material fact, prior to and independent of anything the boundary might be said to do. Three research programs, developed for unrelated reasons and largely without reference to one another, have converged on a different picture, in which a boundary is not merely where a system stops but an object that actively determines what the system is permitted to become. The free-energy principle treats the boundary separating an organism from its environment as a Markov blanket, a statistical structure without which active inference cannot be defined at all. Developmental bioelectricity treats the boundary of a regenerating tissue as the visible surface of a pattern-memory layer that specifies, and actively pursues, a target anatomical configuration. An admissibility-theoretic account of refusal, developed independently of either literature, treats the boundary drawn by a genuine commitment as the deletion of an entire region of an agent's future trajectory-space, an operation the source theory proves cannot be represented as an ordinary preference held over a fixed space of options.

The three accounts are frequently gestured at together, under headings like “boundaries as active constraints” or “the computational role of biological boundaries,” as though naming the shared intuition were most of the work. This paper takes the harder position that the shared intuition is not, on inspection, a shared mechanism. Section 2 states Friston's account precisely: a boundary as a statistical screen, defined by which causal edges are absent. Section 3 states Levin's account precisely: a boundary as the visible edge of a stored target state that a collective actively pursues. Section 4 states the admissibility-theoretic account precisely, at three scales that turn out to compose into one picture, and then develops its treatment of refusal specifically. Section 5 argues that screening, pursuit, and deletion are genuinely different operations, not differently-dressed versions of one operation. Section 6 states the point of real friction between the refusal account and Friston's own framework as precisely as the source theory allows. Section 7 argues for treating boundary-formation as a category with multiple known members rather than searching for the single mechanism underlying all three.

One clarification is worth making before any of the three accounts is developed, since the word “boundary” will otherwise appear to change referent from section to section. The comparison that follows is deliberately functional rather than spatial: the three boundaries inhabit different mathematical spaces (a space of conditional dependencies, a space of anatomical configurations, a space of admissible future trajectories), but each is being compared on the same functional grounds, namely that it partitions a system by regulating what can directly influence it, what configuration it corrects toward, or what futures remain available to it. The paper is not claiming that a cell membrane, a Markov blanket, an anatomical contour, and an excluded trajectory set are the same kind of entity. It is claiming that each answers a structurally analogous question, posed against a different mathematical object, and that comparing the answers is what shows the questions are not, after all, the same question.

2 Friston: The Boundary as Statistical Screen

Within the free-energy principle, the boundary of a self-organizing system is formalized as a Markov blanket: a partition of a system’s states into internal states μ , external states η , and blanket states, the blanket states further divided into sensory states s and active states a (Friston, 2019). The defining property is not a distance or a membrane but an absence: internal states do not directly cause sensory states, and external states do not directly cause active states. This specific pattern of missing causal edges is sufficient to induce conditional independence between internal and external states: each is screened off from the other, in the technical sense of probabilistic graphical models, given the blanket states through which any influence must pass (Kirchhoff et al., 2018). The boundary is therefore defined entirely by what does not directly connect across it, not by any positive property of a membrane, wall, or skin.

This statistical boundary is what makes active inference definable in the first place. Because internal states are conditionally independent of external states, a system’s internal dynamics can be interpreted as performing approximate Bayesian inference about external causes, mediated entirely through the sensory and active states the blanket comprises. Minimizing variational free energy, an upper bound on surprise, jointly a function of model fit and complexity, becomes the single imperative governing both perception, which updates internal beliefs to better fit sensory evidence, and action, which changes sensory evidence to better fit existing beliefs. The incompatibility argument developed in Section 6 concerns a specific, simplified case of this process: free-energy minimization over a generative model and policy space that are not themselves being revised. This is a simplifying assumption, not a claim about what active inference is in general capable of; real active-inference accounts routinely include learning, model selection, policy pruning, habit formation, and hierarchical inference, all of which alter the model or the policy space over time. What matters for the comparison is not whether the model can change, since it plainly can, but how a change of that kind is represented once it happens: as an inferential event evaluated within the model’s existing terms, using the model’s existing machinery for updating beliefs and revising policies, or as an operation that constitutively redefines the domain over which subsequent inference takes place. Section 4’s admissibility account claims the latter for refusal specifically, and it is that narrower, sharper distinction, not fixed versus changing, that Section 6 turns on.

3 Levin: The Boundary as Stored Pursuit

Developmental bioelectricity locates a different kind of boundary in a different kind of substrate. Endogenous voltage gradients, established across cell membranes by ion channels, pumps, and gap junctions, function as a bioelectric prepatter that instructs cell behavior during embryogenesis, regeneration, and the ongoing maintenance of tissue identity, operating alongside and sometimes independently of transcriptional and genetic control (Levin, 2021). Terms like storage, target, and memory are functional interpretations of this dynamic rather than directly measured entities, and are best read that way here: what is actually observed is that bioelectric dynamics behave as a distributed pattern-regulation system, supporting error correction toward stable, reproducible large-scale anatomical outcomes even after the outcome has been disturbed. When a developing or regenerating structure's normal configuration is disrupted, by injury, by a developmental perturbation, or by an induced mutation, altering the bioelectric state alone can, in documented cases, drive correction back toward the structure's typical configuration, restoring correct large-scale anatomy without any direct repair of the underlying genetic or biochemical lesion that caused the original disturbance (Levin, 2023).

On this account, the boundary of a developing or regenerating structure is not where growth happens to stop; it is the visible outcome of a pattern-regulation process the tissue's bioelectric layer supports, correcting toward and capable of re-encoding when the regulated configuration itself changes. This is a boundary defined by what a system's error-correction dynamics are organized around, not by which states are screened from which. It shares with the Markov blanket a rejection of pure passivity, since the boundary is doing work rather than merely marking a stopping point, but the kind of work is different in a way Section 5 makes precise: a Markov blanket's defining property is a pattern of absent causal edges within a joint distribution, which need not itself encode any historical or goal-directed information; a bioelectric boundary's defining property is a correction process evaluated against a configuration the system does not currently occupy and is actively closing the distance to. The existence of such a correction process does not by itself entail the specific conditional-independence structure that defines a Markov blanket, nor does a Markov blanket by itself entail anything about correction toward a prior configuration; the two properties can come apart in either direction, which is enough to establish that pursuit and screening are not the same operation without requiring any claim about how tightly either system is coupled to its surroundings.

4 RSVP and Refusal: The Boundary as Admissibility Deletion

A third account, developed within an admissibility-theoretic treatment of irreversible computation, locates the boundary neither in an absence of causal edges nor in a stored target but in what has been permanently removed from a system's space of possible futures.¹ The account operates at three scales that share a single underlying schema rather than merely resembling one another. Let $\mathcal{A}_t$ denote the set of admissible continuations available to a system after its history up to time t . The schema shared across all three source theories is that admissible continuations only shrink:

$$\mathcal{A}_{t+1} \subseteq \mathcal{A}_t.$$

¹This section departs from the practice, followed elsewhere in this series, of not citing the author's own prior work, for the reason given in earlier papers in the same series: the argument depends on naming specific prior results precisely enough to compare them against an external literature, which citation without attribution would make impossible to check.

A structural plateau is a region where $\mathcal{A}_t$ has stopped shrinking under further ordinary dynamics; a Pop or Refuse event is a single, discrete step at which $\mathcal{A}_{t+1}$ becomes a proper subset of $\mathcal{A}_t$; exclusion distance measures how obstructed the remaining path between two states is once $\mathcal{A}_t$ has already contracted around them; and refusal, developed below, is the specific event that performs such a contraction over trajectories beginning with a given action. What follows works through each scale in the source theory’s own terms, with the shared schema stated here so that the connection between them need not be taken on the paper’s word.

At the continuum scale, a field-theoretic formulation of admissible historical structure treats the entropy field S over a system’s configuration space as approaching a maximal admissible value in certain regions; those regions act as effective barriers to further change, producing what the source theory calls structural plateaus and frozen domains, replacing the role classical singularities play in less constraint-first formulations (Flyxion, 2026a). The boundary here is where a system has run out of admissible room to keep changing, a global saturation surface rather than a local event.

At the discrete, event scale, an irreversible rewriting calculus over option-spaces formalizes the elimination of a region U of an option-space by a generator Pop_U , realized geometrically as restriction to a face of a probability simplex; the entropy cost of the elimination is not distributed evenly but concentrated in a bump function supported near the boundary of the eliminated region, in what the source theory calls boundary-sharpening slack (Flyxion, 2026b). At this scale, the boundary is not a surface a system approaches asymptotically but the specific locus where a single irreversible commitment gets priced.

At the level of pairs of states, a third formulation defines the exclusion distance $D_E(x, y)$ between two states as the minimum admissibility barrier separating them under the system’s admissible trajectories; a boundary exists between x and y exactly where $D_E(x, y) > 0$, and identity is characterized as what persists because of, not despite, that obstruction (Flyxion, 2026c). This is the same underlying structure examined at the scale of a single pair of states rather than a whole field or a single event.

4.1 Refusal as Deletion, Not Preference

The sharpest and most fully worked instance of this account concerns refusal specifically: the case in which an agent, capable of pursuing a high-value trajectory, permanently excludes it from consideration rather than merely declining it on a given occasion (Flyxion, 2026d). The source theory proves a non-representability result: refusal cannot be represented as ordinary utility maximization over a fixed triple (Act, U, π) (actions, a utility function, and a policy) without augmenting the state space with event-history as a primitive. The proof turns on separating the utility function U used to represent an agent’s choices from V , the agent’s own recognized instrumental evaluation of a trajectory: an agent who merely discovers that winning has become worthless, and one who refuses to win, are behaviorally indistinguishable under U alone but are not the same event, and only a representation that tracks event-history distinguishes them.

A first approximation states refusal of an action-type α as an inadmissibility prior: a constraint C_α such that $p(\tau \mid C_\alpha) \propto p(\tau) \cdot \mathbf{1}\{\tau \text{ does not begin with } \alpha\}$, setting the probability of an entire class of trajectories to zero rather than merely reweighting it. Stated this way, the formalism is not yet distinguishable from ordinary Bayesian conditioning on a hard constraint, since a conditional distribution can assign zero posterior probability to an event without the underlying sample space itself having changed; a Bayesian who conditions on C_α can always, in principle, revisit the conditioning and recover access to the excluded trajectories, because the sample space was never altered, only reweighted. What the source

theory intends is stronger, and is better made explicit at the level of the sample space itself rather than left to the prose surrounding the equation:

$$\Omega_{t+1} = \Omega_t \setminus A_\alpha, \quad p_{t+1} = \text{Renorm}(p_t \upharpoonright_{\Omega_{t+1}}).$$

Here the domain changes first: A_α , the region of trajectory-space beginning with α , is removed from the sample space Ω_t itself, producing a strictly smaller Ω_{t+1} , and only afterward is the probability measure renormalized over what remains. Ordinary conditioning leaves Ω_t available to the model even for an event assigned posterior probability zero; refusal, on this reading, changes the domain against which all future policies and probabilities are defined, and does not leave the excluded region available to be reconditioned back into consideration by any update internal to the model. The source theory describes this as constitutive rather than evidential, a topological deletion rather than a belief update, and the domain-level notation above is what makes that description a claim about the formalism rather than a claim about how the formalism should be read. This is an instance of the same schema from the opening of this section, applied to the specific case of a single commitment: $\mathcal{A}_{t+1} \subsetneq \mathcal{A}_t$, with the trajectories beginning in α removed permanently at the boundary the refusal draws.

5 Screening, Pursuit, and Deletion Are Not the Same Operation

Read side by side, the three accounts answer structurally different questions, and the difference survives an attempt to translate each into the others' terms. Friston's account answers: given a system's joint distribution over internal, external, sensory, and active states, which parts are statistically shielded from which, and through what mediating channel. The blanket relation is a property of that joint distribution; it can be embedded in a dynamical, hierarchical, temporally deep model without ceasing to be, at its core, a conditional-independence structure, and nothing about that structure by itself distinguishes a trajectory that is currently improbable from one that has been constitutively removed from consideration. Levin's account answers: what configuration is a system's regulatory dynamics organized around restoring, and how far is its current state from that configuration. It is evaluated against a configuration the system does not currently occupy, and it presupposes an error-correction process that the boundary's behavior can be read off from. The admissibility account answers: which regions of future trajectory-space remain reachable at all, given the system's accumulated history of irreversible commitments. It is evaluated over a history rather than a single state or a single distance-to-target, and its central claim, developed in Section 6, is narrower than a claim about time depth in general: not that the other two accounts fail to represent history or change, but that neither, on its own terms, distinguishes an option that is merely unlikely from one that has been deleted.

These are not three routes to the same computation. A Markov blanket can exist, fully formed, around a system with no error-correction process organized around any particular target and no history of irreversible commitment behind it: conditional independence is a property of a joint distribution, not of a trajectory or a regulatory target. The existence of a bioelectric error-correction process does not logically entail the specific conditional-independence structure that defines a Markov blanket, since Markov blankets identify mediated dependencies rather than requiring any particular degree of insulation, and a system can exchange matter, energy, and signal extensively while still retaining a blanket partition; the two properties are simply not the same claim, and one holding is not evidence that the other does. An admissibility deletion, once made, constrains a system's future regardless of whether that future is being corrected toward any specific configuration or is

statistically screened from anything in particular; the deletion is a fact about the trajectory-space, not about either the system’s joint distribution or its regulatory target. Each account can be true of a system while either of the others is false, under-defined, or simply not the right kind of claim to evaluate against that system at all. This is the sense in which they are three mechanisms rather than three names.

6 Where the Accounts Actually Conflict

The relationship between the admissibility account and Friston’s is sharper than mutual inapplicability, however, and the sharpness is worth stating precisely rather than leaving as a general gesture toward tension. The source theory for refusal states, as a specific and citable claim rather than a general observation, that refusal is incompatible with free-energy minimization unless the generative model includes commitment variables as latent structure. The argument is direct: a genuine refusal permanently excludes a trajectory that may have had high prior probability and low expected surprise under the agent’s existing generative model. Relative to that model, refusing the trajectory increases, rather than decreases, expected free energy, because it removes access to a low-surprise region of trajectory-space without the removal being licensed by anything in the model that would make the removal itself surprise-reducing. Refusal becomes compatible with free-energy minimization only once the generative model is extended with latent variables (the source theory’s example concerns social role stability and mutual recognition) relative to which the previously low-surprise trajectory becomes high-surprise, because pursuing it would destabilize those latent invariants.

This is not merely a claim that the two frameworks emphasize different things, and it is worth being precise about what kind of claim it is, since the more obvious reading is not quite right. The incompatibility is not simply that free-energy minimization reweights where refusal deletes, since an active-inference model can perfectly well assign a policy zero prior probability, or exclude it by construction, without any difficulty; a modeller can always encode a settled refusal after the fact by writing the excluded policy out of the model from the start. The claim is narrower and harder to dismiss on those grounds: standard active inference does not, on its own terms, explain the *event* by which zero support became obligatory, as an irreversible historical act performed at a specific point along a trajectory, rather than merely represent the zero-probability result once it is already settled. A genuine refusal permanently excludes a trajectory that had high prior probability and low expected surprise under the agent’s existing generative model; relative to that model, the exclusion itself increases expected free energy, because it removes access to a low-surprise region of trajectory-space without the removal being licensed by anything in the model that would make performing the removal itself surprise-reducing. Refusal becomes representable within free-energy minimization only once the generative model is extended with latent variables, the source theory’s example concerns social role stability and mutual recognition, relative to which the previously low-surprise trajectory becomes high-surprise, because pursuing it would destabilize those latent invariants. What active inference alone cannot supply is not the zero, but the reason the zero was constituted as an event rather than discovered as a standing fact about the model.

It is worth stating the dependence structure of this claim explicitly, since it is also the paper’s most contestable step and deserves to be contested at the right point rather than the wrong one. The argument runs: (1) Section 4 adopts a particular mathematical definition of refusal, in which $\Omega_{t+1} = \Omega_t \setminus A_\alpha$ constitutes a new, permanently smaller domain rather than merely reweighting probabilities over the existing one; (2) standard active inference,

absent further additions, optimizes expected free energy over a fixed generative model and a fixed domain, and has no internal operation corresponding to step (1)’s domain-level contraction, as opposed to reweighting within a domain that remains available throughout; (3) therefore, under definition (1), the specific event of constituting Ω_{t+1} cannot arise from process (2) without extending the generative model with commitment variables as latent structure. The conclusion follows from the definition adopted in step (1), not from any independent claim about what active inference can or cannot in general represent once a domain change has already occurred. A reader who wants to resist the conclusion has a clear and correctly-located target: dispute that refusal is properly formalized as a domain-level contraction rather than as an extreme case of within-domain reweighting, in which case the incompatibility may soften or disappear. What such a reader cannot do is show the argument wrong by pointing out that active-inference models can represent zero-probability policies, since the claim was never that they cannot; the claim is conditional on Section 4’s definition throughout, concerns the constituting event rather than its settled result, and was never intended to hold independently of that definition.

Whether latent variables of the kind the source theory proposes can be added to an active-inference model without quietly reintroducing something equivalent to the admissibility account’s own domain-contraction primitive is, so far as this paper is aware, an open question the source theory raises but does not resolve. What is not open, given the definition Section 4 adopts, is the narrower claim: ordinary free-energy minimization, over a fixed generative model and domain with no commitment variables, cannot produce the constituting event of refusal as that definition specifies it, whatever it can do with the resulting zero once constituted by some other means. That is a specific, falsifiable point of disagreement between two formal frameworks, not a difference in vocabulary papered over two compatible pictures, and its force rises or falls with the definition it depends on rather than with any general comparison of the two frameworks’ scope.

7 A Category, Not a Metaphor

The temptation, on noticing that three independent literatures have each concluded that boundaries compute, is to treat the convergence itself as the discovery: boundaries are active, not passive, and three fields have each said so in their own language. This paper has argued that the more useful discovery lies one level down, in the fact that the three fields say three different things once the vocabulary is actually translated rather than merely juxtaposed. Screening is a property of a joint distribution. Pursuit is a property of a regulatory process and a system’s distance from what it corrects toward. Deletion is a property of an accumulated history of irreversible commitment. A system could in principle exhibit any one without either of the others, and in the refusal case specifically, exhibiting deletion is formally shown to be incompatible with deriving it from free-energy minimization alone.

The present comparison does not support treating these three mechanisms as instances of a single underlying computation, though this is narrower than a claim that no such unification is possible; a deeper formalism deriving all three as special cases may well exist. What the comparison does establish is a constraint any such formalism would have to meet: it would need to preserve, not erase, the distinctions argued for here, and it would need to reproduce, not dissolve, the specific incompatibility of Section 6. A unification that merely renamed the three mechanisms without accounting for why one resists derivation from another would not have unified anything; it would have hidden the finding this paper reports. Short of that standard being met, the responsible position is to treat boundary-formation as

a genuine category of computational object with at least three known, non-interchangeable implementations, resisting the flattening move, common in interdisciplinary synthesis and equally available to this paper's own account as to either of the other two, in which noticing that several fields use boundary-like language is mistaken for having shown they mean the same thing by it. What the three accounts actually share is narrower and more defensible than a single mechanism: each has found that treating a boundary as inert geometry fails to predict what the system does, and each has built a specific, checkable alternative in its place. That shared failure of the passive picture, not a shared positive mechanism, is what unifies the literature this paper has reviewed.

References

- Friston, Karl. 2019. "A Free Energy Principle for a Particular Physics." *arXiv preprint* arXiv:1906.10184.
- Kirchhoff, Michael, Thomas Parr, Ensor Palacios, Karl Friston, and Julian Kiverstein. 2018. "The Markov Blankets of Life: Autonomy, Active Inference and the Free Energy Principle." *Journal of the Royal Society Interface* 15(138): 20170792.
- Levin, Michael. 2021. "Bioelectric Signaling: Reprogrammable Circuits Underlying Embryogenesis, Regeneration, and Cancer." *Cell* 184(8): 1971–1989.
- Levin, Michael. 2023. "Bioelectric Networks: The Cognitive Glue Enabling Evolutionary Scaling from Physiology to Mind." *Animal Cognition* 26: 1–27.
- Flyxion. 2026a. *RSVP: Relativistic Scalar–Vector Plenum, A Field Theory of Admissible Historical Structure*. Unpublished monograph.
- Flyxion. 2026b. *Structured Irreversibility: Discrete Rewriting and Coherence Realization*. Unpublished monograph.
- Flyxion. 2026c. *The Autonomy of Refusal: Constraint, Residue, and the Geometry of Persistence*. Unpublished monograph.
- Flyxion. 2026d. *Throwing the Game: Refusal, Event-Driven Cognition, and the Survival of Value Beyond Autoregressive Intelligence*. Unpublished manuscript, revised edition.