Immunity as Boundary Verification

Normal-Form Dynamics, Adversarial Classification, Distributed Verification, and Non-Clinical Proxy Maps in Active Biological Systems

Yining Wu
Independent Researcher
yining.wu@alumni.upenn.edu

Version v1.0 — May 2026

Abstract

The immune system is often described as a defense or killing system. FDS-B1 gives a systems-theoretic reconstruction: immunity is a finite-capacity boundary-verification architecture. Within the Finite Distinction Systems framework, a living system maintains operational boundaries under finite sensing, repertoire, classification, memory, signaling, energetic, spatial-routing, action, repair, and resolution budgets. Immune activity is therefore not reducible to destruction. Regulated elimination is one downstream action selected after a prior verification pipeline: candidate biological distinctions are sensed, admitted into immune update channels, classified against memory, tolerance, danger, damage, and tissue context, routed into regulated action, incorporated into memory or tolerance, and resolved or converted into chronic verification load. This version hardens the framework by adding explicit normal-form equations for verification saturation, a fold-like attractor-loss model for systemic alarm-like collapse, adversarial distinction injection and classifier sabotage, distributed verification topology and spatial latency, and non-clinical wet-lab proxy maps based on repertoire entropy, inflammatory alarm load, metabolic resources, and classification uncertainty. The framework yields three crucial divergent predictions: verification DDoS (decoy load raises error and delay even under fixed danger signal), bandwidth exhaustion distinct from effector exhaustion, and spatial bottleneck failure when verification latency exceeds local damage timescale. The central bridge claim is that immune failure can be interpreted non-clinically as failure of admission, classification, cross-scale alignment, adversarial robustness, action selection, memory updating, or FDS-resolution. The paper does not replace immunology, molecular mechanisms, diagnosis, treatment theory, vaccination science, clinical immunology, or medical decision-making. It provides a non-clinical causal-topology and modeling layer for organizing immune classification, tolerance, inflammation, memory, and regulated destructive action as finite-system boundary maintenance.

Keywords: finite distinction systems; immunity; boundary verification; immune classification; active tolerance; danger model; trained immunity; immune repertoire; adversarial classification; immune evasion; inflammation; immune memory; resolution; spatial latency; non-clinical biomedical bridge; capacity deficit; finite capacity; active boundary.

Clinical and ethical notice

FDS-B1 is a non-clinical systems-theoretic paper. It does not diagnose, treat, prevent, manage, stratify, or predict disease in individuals. It does not recommend medications, supplements, vaccines, devices, procedures, screening schedules, diets, behavioral interventions, therapy timing, dosing, clinical triage, immunomodulation, or any change

to medical care. It does not provide clinical decision support. Disease names and pathological phrases, when used, are non-clinical illustrative correlates or model-class analogues, not diagnostic categories or clinical claims. Readers with medical symptoms, diagnoses, treatment questions, emergency symptoms, medication questions, vaccine questions, or personal health decisions should consult licensed medical professionals and follow applicable medical guidance. B1 is a research-interpretation bridge paper, not a medical service or clinical tool.

1 Introduction: Immunity Beyond Killing

The immune system is commonly introduced as a defense system. It detects infectious agents, eliminates abnormal or infected cells, coordinates inflammation, forms memory, and participates in repair. This description is useful, but it can make destructive action appear more fundamental than classification, tolerance, repair, and resolution. In fact, immune destruction is costly. Before a biological structure is killed, tolerated, ignored, repaired, contained, remembered, or escalated, it must be admitted into immune update and classified under uncertainty.

FDS-B1 begins with a different primitive:

$$\text{immunity} = \text{finite-capacity boundary verification} + \text{regulated action}.$$

The immune system is not primarily a killing machine. It is a multiscale verification architecture that tests candidate biological distinctions against host boundary state, tissue context, regulatory memory, tolerance history, danger/damage signals, spatial routing constraints, resource budget, and threat-loss expectations. Regulated elimination is one action in the immune action space, not the defining primitive of immunity.

B1 central thesis

Immunity is a multi-scale, finite-capacity boundary-verification architecture. It admits candidate biological distinctions, classifies them against memory, tolerance, danger, damage, tissue context, and spatial accessibility, selects regulated actions including monitoring, repair, containment, tolerance, resolution, and regulated elimination, and fails when verification demand exceeds classification, memory, signaling, spatial-routing, resource, or resolution capacity.

This paper is part of the biomedical B-series of Finite Distinction Systems. B0 defines the safety firewall, claim-status hierarchy, mechanism non-replacement rule, biomedical FDS object, and minimum reporting template for downstream biomedical bridge papers [1]. B1 applies that governance structure to immunity. It treats immunity as a non-clinical systems bridge: a way to organize immune mechanisms by their roles in active boundary maintenance, not a replacement for immunology.

1.1 Why boundary verification is broader than self/non-self

The language of self/non-self remains historically important, but it is too low-dimensional to carry the full immune problem. Biological boundary verification includes distinctions such as

$$\{\text{self}, \text{altered-self}, \text{infected-self}, \text{damaged-self}, \text{transformed-self},\\ \text{commensal}, \text{foreign-harmless}, \text{foreign-dangerous}, \text{unknown-for-verification}\}.$$

This does not reject classical self/non-self theory. It treats self/non-self as one projection of a richer boundary-state verification problem. Immune systems often distinguish tolerated foreign from dangerous foreign, damaged self from healthy self, symbiotic microbial signals from invasive microbial signals, repairable injury from spreading threat, and transient alarm from chronic verification load.

This perspective is compatible with prior immunological theories that challenged a simple self/non-self boundary. Janeway's pattern-recognition perspective emphasized innate recognition of conserved microbial patterns [4]; Matzinger's danger model shifted attention toward damage and danger signals [5]; and Pradeu's continuity theory placed immunity within a broader account of organismal identity and biological boundaries [7]. B1 does not replace these accounts. It provides a finite-capacity accounting layer beneath them.

1.2 Relation to distinction admission

B1 extends the FDS notion of distinction admission into biological systems. A candidate biological distinction may be physically present without being admitted into immune update. It may be sensed but not acted on, admitted for verification but later rejected, classified as tolerated, or routed into regulated elimination. Thus B1 separates four stages:

$$\text{recognition} \neq \text{admission} \neq \text{verification} \neq \text{action}.$$

A fifth stage, resolution, determines whether the system returns to a lower-cost boundary-maintenance regime or accumulates chronic verification load. This distinction is crucial. A model that treats immunity as direct stimulus-response killing misses the finite verification architecture that decides whether a signal should be ignored, watched, tolerated, repaired, contained, escalated, remembered, or destroyed.

2 Scope and Mechanism Non-Replacement

2.1 Non-clinical reader contract

The B-series biomedical translation barrier is

$$\text{FDS biomedical interpretation} \not\Rightarrow \text{clinical action}.$$

B1 follows that rule. It does not define disease categories, diagnostic biomarkers, clinical thresholds, treatment-response groups, patient phenotypes, prognosis, risk categories, therapeutic timing, dose schedules, vaccine schedules, or standards of care. All claims below are non-clinical bridge claims or research hypotheses unless independently converted into validated biomedical tools under appropriate ethical, regulatory, and clinical processes.

2.2 Mechanism non-replacement rule

B1 does not replace immunology. It does not replace innate immunity, adaptive immunity, antigen presentation, pattern-recognition receptors, complement, cytokine signaling, clonal selection, affinity maturation, immune tolerance, regulatory T cells, tissue-resident immunity, trained immunity, immunological memory, barrier biology, inflammation, resolution biology, clinical immunology, pathology, microbiology, vaccinology, or immunotherapy research.

Mechanism non-replacement rule for B1

FDS-B1 acts as a causal-topology layer. It may organize immune mechanisms by their systems-level roles in boundary verification, classification, tolerance, memory, repair, pruning, regulated destructive action, adversarial robustness, spatial routing, and resource allocation, but it may not erase, replace, or ignore validated immunological mechanisms.

2.3 Claim-status convention

Unless otherwise specified, B1 claims have the following status:

- **B-L2 biomedical bridge:** a mapping from FDS variables to immune-system roles.
- **B-L3 research hypothesis:** a non-clinical hypothesis with measurable proxies.
- **Not B-L5:** no clinical use, no patient-level inference, no treatment guidance.

3 FDS Core Background

The formal FDS object is

$$S = (X, E, B, M, Y, A, U, \pi, \ell, \Phi, \mathcal{P}, \tau),$$

where X is internal state, E is environmental state, B is boundary, M is memory or model state, Y is observation channel, A is action space, U is update rule, π is classification or projection, ℓ is boundary-maintenance loss, Φ is resource budget, $\mathcal{P}$ is an admissible perturbation or pruning family, and τ is the update timescale [2].

An empirical system enters the active-boundary FDS cascade only if its updates are nontrivial and relevant to future boundary-maintenance loss:

$$\mathbb{P}\big(U(M_t, Y_t) \neq M_t\big) > 0, \qquad I(M_{t+1}; \ell_{t+k}) > 0$$

for some $k > 0$. In empirical applications, a stronger intervention test is appropriate:

$$\mathbb{E}[\ell_{t+k} \mid do(U)] \neq \mathbb{E}[\ell_{t+k} \mid do(U_\emptyset)].$$

For B1, this means that an immune update counts as active-boundary-relevant only if it changes future boundary-maintenance loss in a specified biological model or measurement setting.

The FDS capacity-deficit form is also inherited. Let Ψ be a pre-specified task family and let

$$R^{(\tau)}_{\min}(\epsilon; \Psi)$$

be the minimal task-relevant rate needed to maintain boundary-relevant distinctions at distortion tolerance ϵ over window τ. If accessible internal or regulatory capacity is C, then

$$\Delta = R^{(\tau)}_{\min}(\epsilon; \Psi) - C.$$

When $\Delta > 0$, full-fidelity internal maintenance is impossible without approximation, externalization, task relaxation, capacity expansion, or failure. B1 applies this structure to immune verification.

4 Biomedical FDS Mapping of the Immune System

The B1 biomedical object is

$$S_{\text{immune}} = (X, E, B, M, Y, A, U, \pi, \ell, \Phi, \mathcal{P}, \tau).$$

It is not one molecular pathway. It is a scale-declared system object whose variables must be specified before any model is tested.

Table 1: FDS-B1 mapping of the immune boundary-verification system.

FDS term	Immune interpretation	Required caution
X	Host internal biological state: cells, tissues, regulatory networks, immune cell populations, repair state	Specify molecular, cellular, tissue, organ, organism, or immune-system scale
E	Perturbation environment: pathogens, commensals, toxins, damaged tissue, transformed cells, allergens, tissue stressors	Do not collapse all causes into one variable
B	Host boundary or immune interface: epithelial barrier, cellular identity boundary, tissue inflammatory boundary, microbiome-host interface, organism-level immune boundary	Boundary scale must be explicit
M	Immune memory and regulatory model: clonal memory, tolerance state, trained innate state, regulatory set-points, prior verification history	Not conscious memory; map to specific biological variables where possible
Y	Observation channels: receptors, antigens, PAMP/DAMP-like signals, cytokines, chemokines, danger signals, damage signals, tissue-state signals	Signals are proxies; they are not ground truth
A	Immune actions: ignore, monitor, tolerate, suppress, repair, contain, inflame, eliminate, resolve, remember	Not treatment recommendation
U	Immune update: signaling, activation, suppression, clonal expansion, memory formation, tolerance adjustment, resolution dynamics	Requires active-boundary qualification
π	Classification map: self/non-self, danger/safe, damaged/healthy, infected/uninfected, tolerate/attack, repair/eliminate	Labels are model-level projections
ℓ	Boundary-maintenance loss: uncontrolled threat, tissue damage, dysregulation, chronic alarm, loss of function, failure to resolve	Not a clinical endpoint unless independently validated

FDS term	Immune interpretation	Required caution
Φ	Resource budget: metabolic energy, immune-cell availability, signaling bandwidth, clonal expansion limits, lymphocyte renewal, tissue-damage tolerance, repair and resolution capacity	Repertoire diversity contributes mainly to M, π, and C_{class}, not only to Φ
$\mathcal{P}$	Perturbations: infection, mutation, injury, inflammation, aging, microbiome shift, stress, memory drift	Perturbation family must not be post hoc
τ	Immune update timescale: seconds, minutes, hours, days, years depending on scale	Timescale must match mechanism and data

4.1 Multiscale boundary declaration

The immune boundary is not one membrane. It is a family of boundary-verification interfaces. A B1 model must declare its primary scale.

Table 2: Multiscale immune boundary-verification examples.

Scale	Boundary verified	Example verification problem
Molecular / receptor	Ligand, antigen, pattern, danger, or damage distinction	Signal admission and initial classification
Cellular	Healthy self, infected self, damaged self, transformed self, immune cell state	Eliminate, tolerate, repair, suppress, or signal
Tissue	Normal tissue state, inflamed state, damaged region, barrier breach	Contain, repair, resolve, or escalate
Organism	Host boundary, microbiome-host interface, external perturbation interface	Systemic regulation and boundary preservation
Memory scale	Prior threat, tolerance state, stale memory, trained response	Faster response versus overgeneralization or drift

5 The Boundary-Verification Pipeline

5.1 Recognition, admission, verification, action, and resolution

B1 distinguishes five layers that are often compressed in informal descriptions.

5.2 Multi-signal admission

A candidate immune distinction d is not admitted merely because it exists. It is admitted when its expected boundary relevance exceeds verification, damage, escalation, and resource costs under a

Table 3: Recognition, admission, verification, action, and resolution are distinct finite-system layers.

Layer	Meaning	Characteristic failure
Recognition	A signal or pattern is detected by some receptor or sensor channel	missed or weak signal
Admission	A candidate distinction enters an immune-update channel	critical exclusion or false admission
Verification	The admitted distinction is classified against memory, tolerance, danger, tissue state, and boundary context	misclassification or high classification entropy
Action selection	A regulated response is selected	over-elimination, under-response, or cross-scale mismatch
Resolution	The system returns toward lower-cost boundary maintenance while preserving useful memory	FDS-resolution failure or chronic alarm

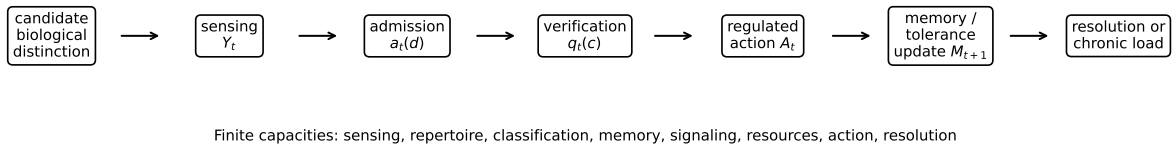

Figure 1: Boundary-verification pipeline. Candidate biological distinctions must be sensed, admitted, verified, routed into regulated action, incorporated into memory or tolerance, and resolved or converted into chronic verification load. This figure is a conceptual systems diagram, not a molecular pathway.

particular memory and tissue context. A minimal multi-signal gate is

$$P_{\text{admit}}(d,t) = \sigma\{\beta[V_B(d,t) + \alpha_1 S_1(d,t) + \alpha_2 S_2(d,t) + \alpha_D S_D(d,t) \tag{1}$$

$$- c_{\text{verify}}(d, M_t) - c_{\text{damage}}(d, B_t) - c_\Phi(d,t)]\}, \tag{2}$$

where S_1 is antigenic or pattern signal, S_2 is contextual permission or co-stimulation-like signal, S_D is danger/damage/tissue-state signal, M_t is memory and tolerance state, and B_t is boundary context. A tolerance memory does not always block admission. It may instead reduce escalation value, increase escalation cost, or route the candidate into a tolerate/suppress/resolve action.

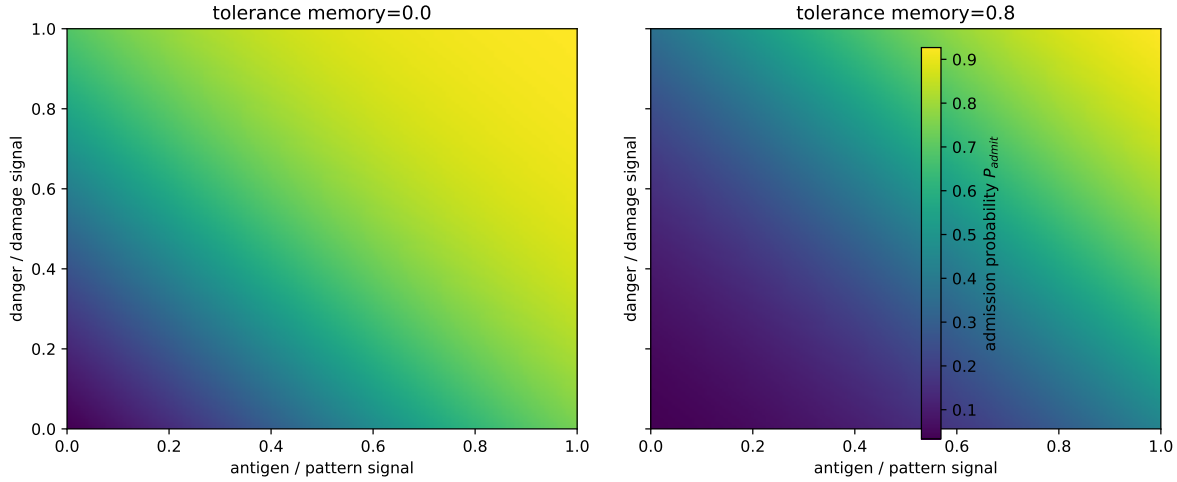

Figure 2: Multi-signal admission. Synthetic admission probability is shown as a function of antigen/pattern and danger/damage signals under low versus high tolerance memory. Tolerance memory shifts the admission/action landscape; it is not simply passive absence of immunity.

5.3 Boundary-state classification vector

Instead of a single self/non-self label, B1 models immune classification as a vector:

$$\pi_{\text{immune}}(d_t) = (s_{\text{self}}, s_{\text{danger}}, s_{\text{damage}}, s_{\text{foreign}}, s_{\text{tolerance}}, s_{\text{memory}}, s_{\text{growth}}, s_{\text{context}}).$$

The action policy depends on the vector and its uncertainty, not on a one-dimensional label. Define

$$q_t(c) = \mathbb{P}(c \mid Y_t, M_t, B_t, \Phi_t)$$

for boundary-relevant classes c, and

$$H_{\text{class}}(t) = -\sum_c q_t(c) \log_2 q_t(c).$$

High H_{class} means the system has not formed a low-distortion classification. Under finite verification capacity, high uncertainty should increase delay, containment, lower-specificity alarm, or broad default action, depending on the model and scale.

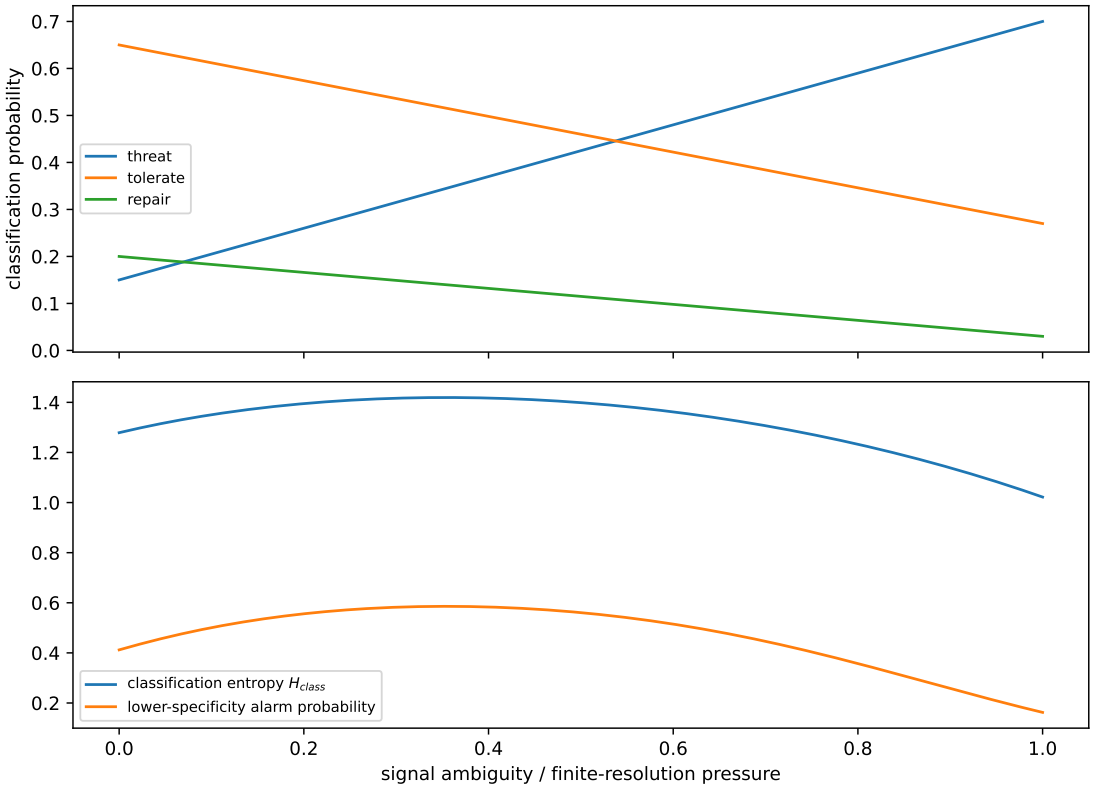

Figure 3: Classification entropy. In this synthetic demonstration, increasing signal ambiguity raises classification entropy and increases the probability of a lower-specificity alarm response. The figure is illustrative and not fitted to biological data.

6 Immune Alphabet and Verification Capacity

Immune capacity is not only metabolic energy. It also includes classification alphabet size: the effective repertoire of antigenic, danger, damage, tolerance, memory, and tissue-context distinctions that can be represented and acted on within an update window.

Let $\mathcal{A}^{\text{eff}}_{\text{immune}}$ be the effective immune classification alphabet. Define

$$C_{\text{class}} = \log_2 |\mathcal{A}^{\text{eff}}_{\text{immune}}|.$$

For TCR/BCR-like repertoire data, a candidate non-clinical proxy is clonal Shannon entropy,

$$H_{\text{rep}} = -\sum_i p_i \log_2 p_i,$$

where p_i is the observed frequency of clone i. This is not a ground-truth capacity measure. Sequencing depth, tissue localization, receptor function, antigen specificity, avidity, co-stimulation, and regulatory context must be considered. Still, repertoire entropy, richness, clonality, Simpson diversity, and functional validation can contribute to estimating $C^{\text{proxy}}_{\text{class}}$.

The verification capacity ledger can be written schematically as

$$C_{\text{verify}}(t) = \min\{C_{\text{sense}}, C_{\text{signal}}, C_{\text{class}}, C_{\text{memory}}, C_{\text{effector}}, C_{\text{resolve}}, \Phi_{\text{eff}}, C_{\text{spatial}}\}.$$

The immune verification deficit is

$$\Delta_{\text{verify}}(t) = L_{\text{verify}}(t) - C_{\text{verify}}(t).$$

When $\Delta_{\text{verify}} > 0$, the system cannot verify all boundary-relevant candidate distinctions at the requested fidelity and timescale. It must narrow admission, delay response, coarse-grain classification, use broad alarm, externalize to tissue-level signaling, consume additional resources, relax the task, or fail.

7 Normal-Form Dynamics of Verification Saturation

This section supplies the normal-form equations behind the schematic figures. The variables are nondimensional and hypothesis-generating. They do not define disease states.

7.1 Demand, capacity, and integrated pressure

Let $P(t)$ be perturbation or pathogen-like load, $D(t)$ be damage/danger load, and $U_{\text{unk}}(t)$ be unknown-for-verification load. Let $Y_{\text{decoy}}(t)$ denote adversarial or nonfunctional decoy burden. A minimal demand equation is

$$\dot{L}_{\text{verify}} = r_P P + r_D D + r_U U_{\text{unk}} + r_Y Y_{\text{decoy}} - \gamma_L L_{\text{verify}}.$$

Capacity is shaped by effective resources, clonal or cellular expansion, usable memory, alarm burden, adversarial suppression, and maintenance debt:

$$\dot{C}_{\text{verify}} = \gamma_C \Big(C_0 + \alpha_\Phi \Phi_{\text{eff}} + \alpha_K K_{\text{clone}} + \alpha_M M_{\text{usable}} - \beta_A A_{\text{alarm}} - \beta_D D_{\text{maint}} - \beta_S S_{\text{sabotage}} - C_{\text{verify}} \Big).$$

The positive deficit is

$$\Delta_+(t) = [L_{\text{verify}}(t) - C_{\text{verify}}(t)]_+.$$

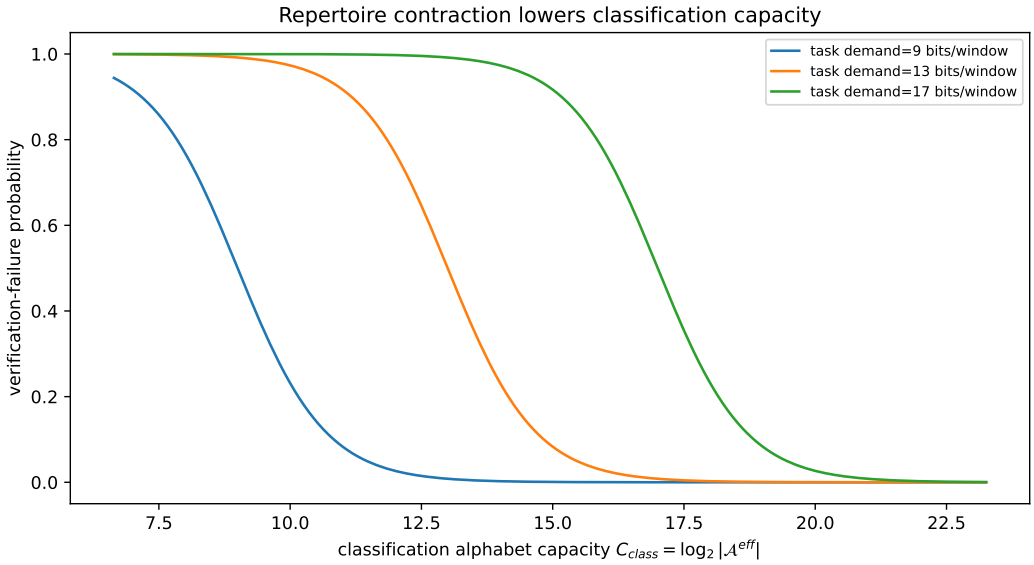

Figure 4: Repertoire contraction lowers classification capacity. The figure shows synthetic failure curves as task demand approaches or exceeds classification alphabet capacity. This is a normal-form demonstration, not a repertoire-sequencing result.

Integrated verification pressure over a finite memory window may be written as

$$Z(t) = \int_{t-T}^{t} \Delta_+(s)\, ds,$$

or dynamically,

$$\dot{Z} = \Delta_+ - \gamma_Z Z.$$

A sustained overload condition is

$$Z(t) > Z_c,$$

which indicates risk of transition into a lower-specificity alarm-like regime. This condition is not a clinical definition of cytokine storm, sepsis, or any disease state.

7.2 Fold-like verification attractor

Let $x(t)$ be a verified boundary-control state. A minimal fold-like normal form is

$$\dot{x} = r(t) - x^2 - \eta_A A_{\text{alarm}} + \xi(t),$$

where

$$r(t) = r_0 + a_C C_{\text{verify}} - a_L L_{\text{verify}} - a_D D_{\text{maint}}.$$

When $r(t) > 0$, a maintained verification attractor exists in the reduced model. As $r(t) \to 0^+$, the system should show slower recovery and a shrinking rescue window. When $r(t) < 0$, the previous attractor is lost in the normal form. This is a systems-level alarm-collapse model, not a clinical definition.

7.3 Maintenance debt and resolution capacity

Persistent alarm consumes resources and produces maintenance debt. Define

$$D_{\text{maint}}(t) = \int_0^t [L_{\text{verify}}(s) + L_{\text{damage}}(s) - C_{\text{repair}}(s) - C_{\text{resolve}}(s)]_+ \, ds.$$

A dynamic form is

$$\dot{D}_{\text{maint}} = \rho_\Delta \Delta_+ + \rho_A A_{\text{alarm}} - \gamma_D D_{\text{maint}}.$$

Effective resource budget evolves as

$$\dot{\Phi}_{\text{eff}} = r_\Phi(\Phi_0 - \Phi_{\text{eff}}) - \eta_A A_{\text{alarm}} - \eta_D D_{\text{maint}}.$$

Resolution capacity can then be modeled as

$$\gamma_R(t) = \gamma_{R0} \frac{\Phi_{\text{eff}}(t)}{\Phi_{\text{eff}}(t) + K_\Phi}.$$

FDS-resolution failure occurs when the triggering distinction has been cleared, contained, or tolerated, but the system fails to return from high-verification/high-alarm state to a lower-cost boundary-maintenance regime.

8 Parameter Status and Future System Identification

The coefficients in the normal-form equations are not claimed to be universal constants. They are effective coarse-grained parameters summarizing gene-regulatory, metabolic, epigenetic, cytokine, repertoire, and tissue-context states over a specified boundary and timescale. In future data-driven versions, these coefficients should be treated as identifiable nonlinear functions rather than fixed constants.

A general form is

$$\dot{z} = f(z, u; \theta(z, u, t)), \tag{3}$$

where

$$z = (L_{\text{verify}}, C_{\text{verify}}, A_{\text{alarm}}, D_{\text{maint}}, \Phi_{\text{eff}}, x), \tag{4}$$

$$\theta = (\alpha_\Phi, \beta_A, \beta_D, \gamma_R, \eta_A, \dots). \tag{5}$$

Each parameter family may depend on underlying biological state:

$$\theta(z, u, t) = h_\Theta(G_t, \text{Epi}_t, \text{Met}_t, \text{Cyt}_t, R_t, \text{Tissue}_t), \tag{6}$$

where G_t is gene-regulatory state, Epi_t is epigenetic state, Met_t is metabolic state, Cyt_t is cytokine or signaling environment, R_t is repertoire state, and Tissue_t is tissue context.

8.1 Connection to data-driven system identification

Sparse nonlinear system identification, symbolic regression, state-space modeling, neural ODEs, and Bayesian dynamical inference may be used to infer candidate closures for θ from perturbational time-series data. A standard formulation writes

$$\dot{z} = \Theta(z, u)\Xi + \epsilon, \tag{7}$$

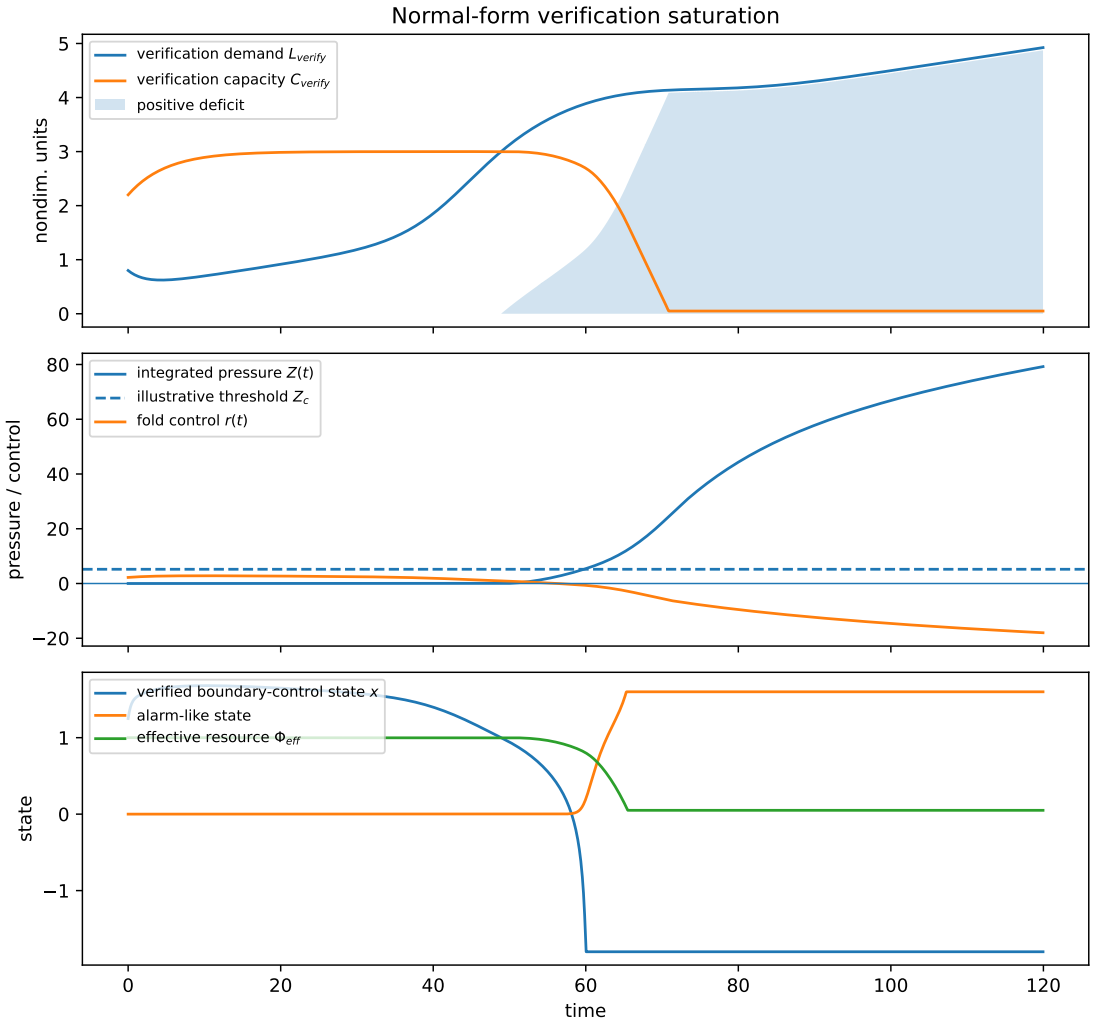

Figure 5: Normal-form verification saturation. Synthetic demand, capacity, integrated pressure, fold control, verified boundary-control state, alarm-like state, and effective resource traces are generated by the nondimensional equations in Section 7. The figure illustrates a model form, not clinical data.

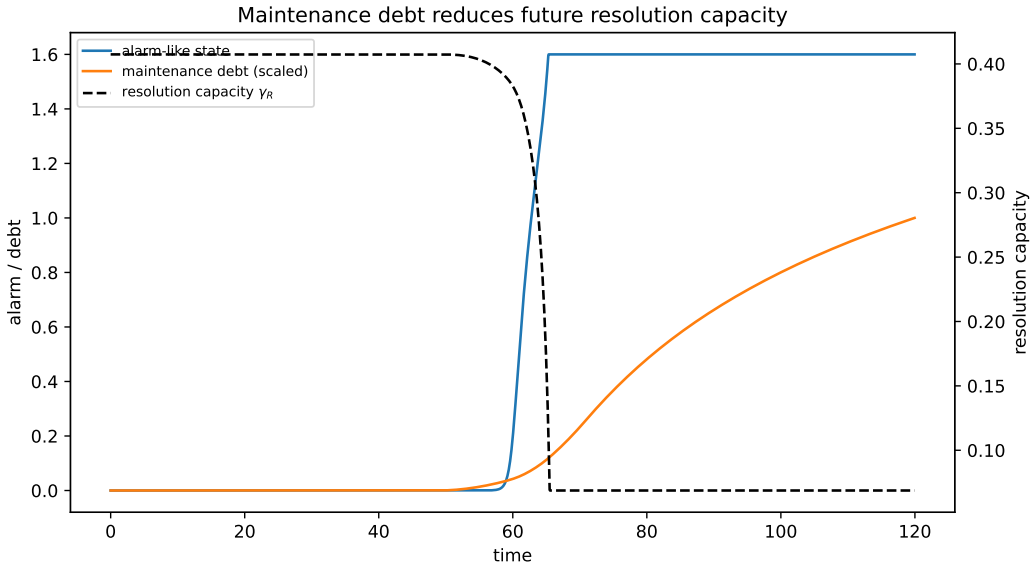

Figure 6: Maintenance debt and resolution capacity. Sustained alarm accumulates maintenance debt, which reduces future resolution capacity in the normal form. This is a non-clinical systems demonstration.

where $\Theta(z, u)$ is a library of candidate nonlinear terms,

$$\{1, z_i, z_i z_j, z_i^2, \sigma(z_i), z_i/(K + z_i), u_i z_j, \dots\}, \tag{8}$$

and Ξ is a sparse coefficient matrix estimated from data. This connects the normal-form equations to data-driven discovery pipelines.

System identification firewall

Identified equations would be model-system or cohort-level research models, not clinical decision tools. Parameter estimates are system-specific and require validation across biological contexts before any non-clinical or clinical use.

Table 4: Effective parameters and candidate identification sources.

Parameter	Current role	Future data source
α_Φ	resource-to-verification capacity coupling	metabolomics, ATP/oxygen/glucose flux, mitochondria
β_A	alarm burden reducing verification capacity	cytokine panels, inflammatory signaling, flow cytometr
β_D	maintenance debt reducing capacity	longitudinal damage/repair proxies, tissue-stress marke
γ_R	resolution recovery rate	time-resolved return-to-baseline dynamics
C_{class}	classification alphabet capacity	TCR/BCR repertoire entropy, clonal diversity, single-c
H_{class}	classification uncertainty	probabilistic classifier entropy from single-cell / cytoki

9 Immune Control Numbers

The normal-form parameters $(\alpha_\Phi, \beta_A, \beta_D, \gamma_R, \eta_A, \dots)$ are effective coarse-grained quantities, but the model's predictive core can be compressed into a small set of dimensionless control numbers. These numbers are not claimed to be universal constants; they are condition-specific ratios that organize when the system transitions between verification regimes.

Define the following dimensionless control numbers:

$$\mathrm{VLR}(t) = \frac{L_{\mathrm{verify}}(t)}{C_{\mathrm{verify}}(t)} \qquad \text{(verification load ratio)}, \tag{9}$$

$$\mathrm{SPI}(t) = \frac{Z(t)}{Z_c} \qquad \text{(sustained pressure index)}, \tag{10}$$

$$\mathrm{ADR}(t) = \frac{Y_{\mathrm{decoy}}(t)}{C_{\mathrm{verify}}(t)} \qquad \text{(adversarial distinction ratio)}, \tag{11}$$

$$\mathrm{SLR}(t) = \frac{\tau^{\mathrm{spatial}}_{\mathrm{verify}}(t)}{\tau_{\mathrm{local\ damage}}(t)} \qquad \text{(spatial latency ratio)}. \tag{12}$$

- VLR > 1: verification demand exceeds capacity. The system must coarse-grain, delay, externalize, consume reserves, or fail.
- SPI > 1: integrated verification pressure has crossed the sustained-overload threshold, indicating risk of transition into a lower-specificity alarm-like regime.
- ADR > 0: adversarial decoy or false-signal burden consumes a significant fraction of verification capacity.
- SLR > 1: spatial verification latency exceeds the local damage timescale, so distributed verification loses the local race.

These control numbers are not a replacement for the underlying dynamics. They are compressed observables that can be estimated from proxy data (Table 6) and compared across model systems, experimental conditions, and timescales.

10 Adversarial Distinction Injection and Classifier Sabotage

A capacity-based immune theory should include cases in which pathogens or transformed cells actively reshape the verification problem. In FDS terms, adversarial biological processes can attack the observation channel Y, the projection map π, the memory state M, the resource budget Φ, or the capacity ledger C_{verify}.

10.1 Three adversarial modes

Verification DDoS. A perturbation may inject many low-value or decoy distinctions, consuming verification capacity:

$$Y_t^{\mathrm{obs}} = Y_t^{\mathrm{true}} + Y_t^{\mathrm{decoy}}, \qquad L_{\mathrm{verify}} \uparrow .$$

Projection sabotage. A perturbation may shift classification:

$$\pi_{\mathrm{immune}}(d) = \text{threat} \quad \to \quad \pi'_{\mathrm{immune}}(d) = \text{tolerate/ignore}.$$

Immune-checkpoint-like or suppressive mechanisms may be discussed as non-clinical illustrative examples, not therapy guidance.

False compression. A perturbation may force semantic overlap between threat and protected self:

$$\pi(d_{\text{threat}}) \approx \pi(d_{\text{self}}),$$

so that finite resolution cannot separate them without collateral risk. Molecular-mimicry-like processes can be interpreted non-clinically as adversarial semantic overlap under finite immune resolution.

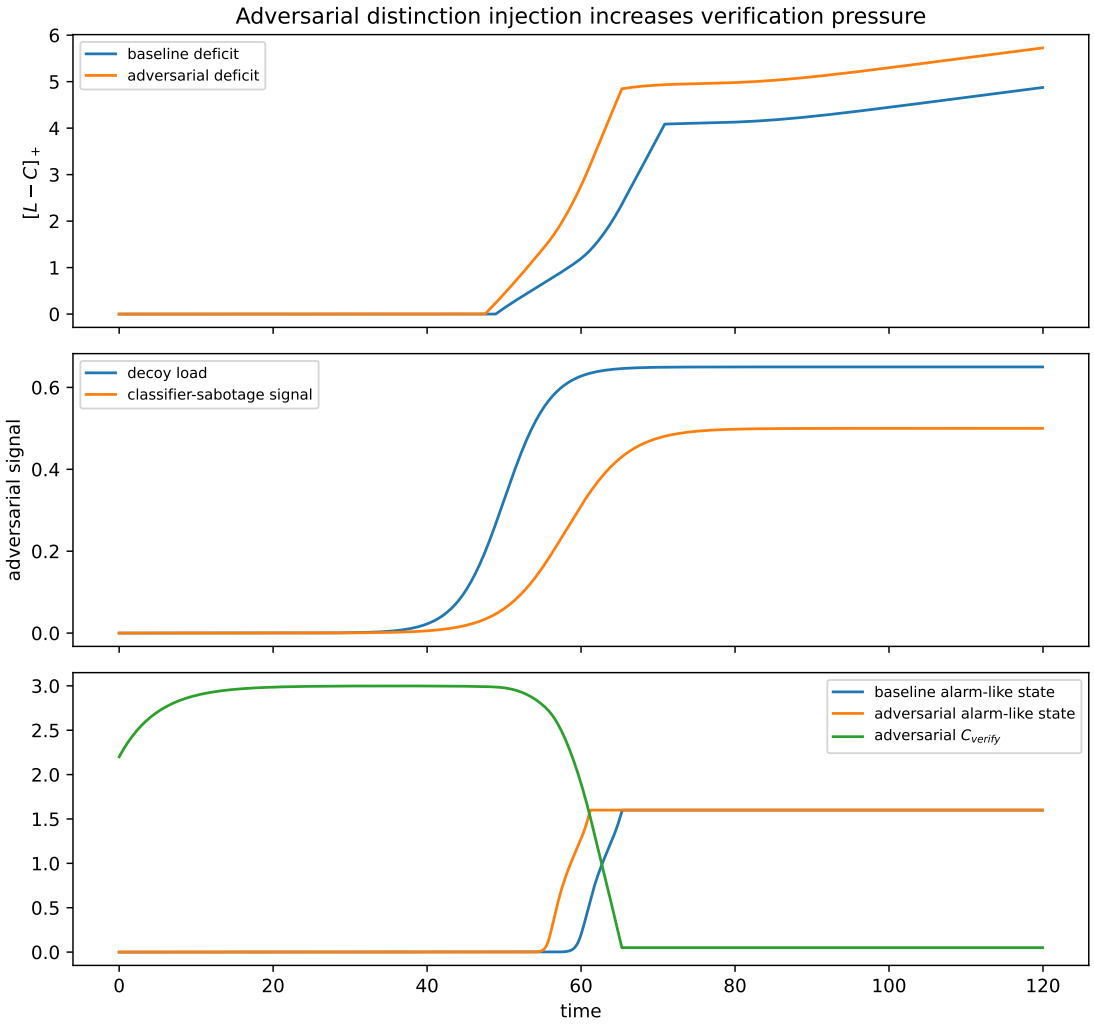

Figure 7: Adversarial distinction injection and classifier sabotage. Synthetic decoy and suppressive signals increase verification deficit and alarm-like state relative to the baseline normal form. The labels are systems analogues, not clinical claims.

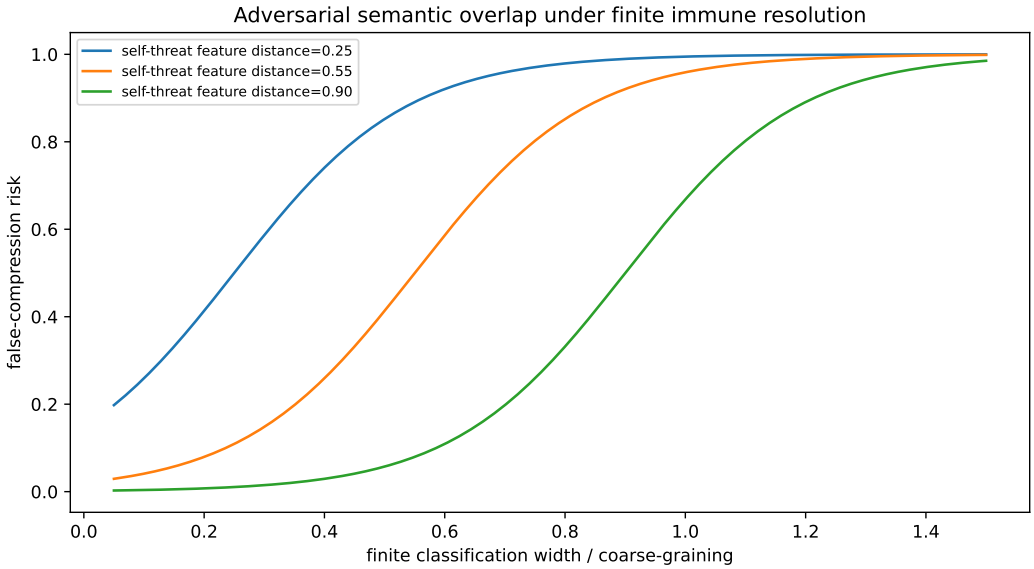

Figure 8: False compression under finite immune resolution. When coarse-graining exceeds the feature distance between protected and threatening classes, false-compression risk rises in the synthetic model. This illustrates molecular-mimicry-like semantic overlap as an FDS failure mode.

11 Decision-Theoretic Immune Action

Given a boundary-state posterior $q_t(c)$, regulated immune action can be described at the coarse-grained population level by an as-if variational model:

$$P(a|s) \propto \exp[-\beta_{\text{eff}}\, \ell(s,a)],$$

where $\ell(s,a)$ includes threat loss, collateral damage, resource cost, resolution cost, and memory/tolerance update cost:

$$\ell_{\text{immune}} = \ell_{\text{threat}} + \lambda_D \ell_{\text{collateral}} + \lambda_\Phi \ell_{\text{resource}} + \lambda_R \ell_{\text{resolution}} + \lambda_M \ell_{\text{memory}}.$$

This is not a claim that individual cells or molecular pathways literally compute arg min expectations. The exponential (or argmin) form is a coarse-grained description of a population-level outcome shaped by receptor affinity distributions, local cell densities, cytokine fields, metabolic resources, clonal expansion time, and regulatory signaling dynamics. The effective inverse temperature β_{eff} captures the effective selection sharpness of the action distribution and is itself bounded by finite resources:

$$\beta_{\text{eff}} \leq f(\Phi_{\text{eff}}, N_{\text{clone}}, C_{\text{signal}}, \tau).$$

Higher selection sharpness (more precise action) demands more resources: more clonal expansion, more signaling bandwidth, more metabolic investment, more time. A minimal cost-of-computation accounting can be written as

$$\text{Cost}_{\text{compute}} = c_{\text{sig}} N_{\text{signal}} + c_{\text{clone}} N_{\text{clone}} + c_{\text{mig}} d + c_{\text{update}}\, H(M_t \mid M_{t+1}, Y_t),$$

where each term corresponds to a physical resource or update cost. This connects immune action selection to the FDS Core's Landauer-based accounting: logically irreversible update operations

in immune memory, tolerance adjustment, and clonal repertoire restructuring carry minimum thermodynamic costs under finite resources.

Thus immune success is not maximal killing. It is expected boundary-loss minimization under finite verification, action, and computation capacity.

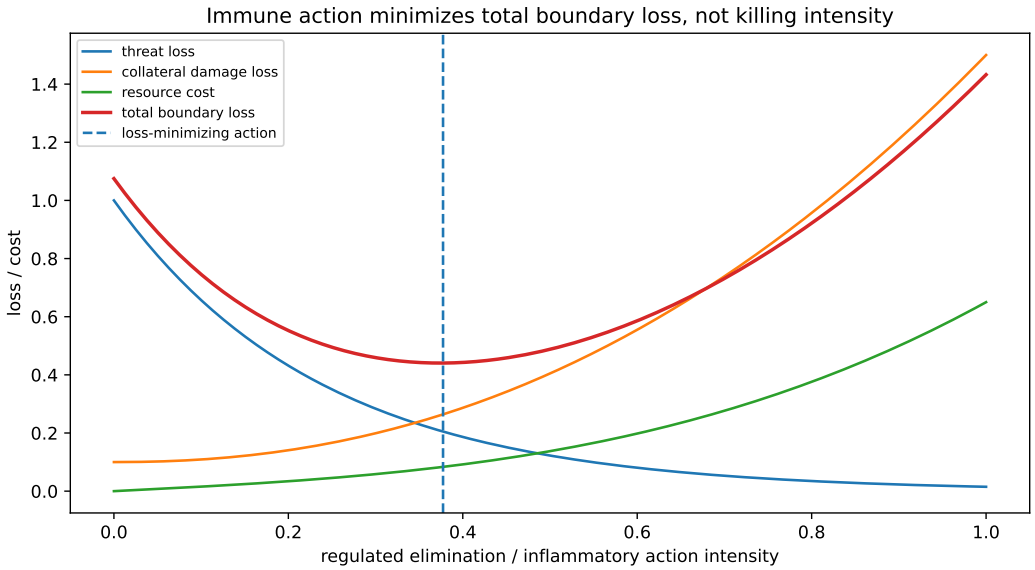

Figure 9: Regulated elimination minimizes total boundary loss, not killing intensity. The synthetic curve shows an optimum where threat loss, collateral damage, and resource cost are jointly minimized.

11.1 Active tolerance

Tolerance is not the absence of immune function. It is a regulated classification-and-action state that preserves boundary-compatible structures from unnecessary destruction. In B1, tolerance may be represented by a memory-dependent action map:

$$M_t^{\mathrm{tol}}(d) \Rightarrow A_t \in \{\mathrm{tolerate}, \mathrm{suppress}, \mathrm{monitor}, \mathrm{resolve}\}$$

unless danger, damage, rapid growth, or memory mismatch shifts the posterior.

12 Distributed Verification Topology and Spatial Latency

The immune system is not a centralized classifier. It is distributed, mobile, and spatially delayed. Signals are sensed locally, packaged by antigen-presenting or reporting processes, routed through lymphatic and vascular channels, compared against local and systemic repertoire, amplified by cell expansion, and returned as regulated action. In B1 terms, lymph nodes and related tissues may be treated as local verification and amplification hubs.

12.1 Spatial verification latency

A simplified spatial verification delay is

$$\tau_{\mathrm{verify}}^{\mathrm{spatial}} = \tau_{\mathrm{sensing}} + \tau_{\mathrm{APC}} + \tau_{\mathrm{migration}} + \tau_{\mathrm{LN}} + \tau_{\mathrm{expansion}} + \tau_{\mathrm{return}}.$$

Let $\tau_{\text{local damage}}$ be the timescale on which local boundary integrity degrades. If

$$\tau_{\text{verify}}^{\text{spatial}} > \tau_{\text{local damage}},$$

then local boundary failure may occur before verified systemic response arrives. Equivalently, if $v_{\text{damage}} > v_{\text{immune}}$, where

$$v_{\text{immune}} \sim \frac{1}{\tau_{\text{verify}}^{\text{spatial}}},$$

then distributed verification loses the local race. This is a spatial systems constraint, not a clinical rule.

12.2 Graph model of verification topology

The linear delay chain can be generalized to a directed graph $G = (V, E)$ with node capacities. Nodes V include tissue compartments, antigen-presenting-cell (APC) relay points, lymph nodes (verification hubs), effector-cell reservoirs, and return channels. Each edge $e_{ij} \in E$ has a capacity K_{ij} (cells per time), speed v_{ij}, queue q_{ij}, and delay τ_{ij}.

For the upstream (tissue to lymph node) information channel, define the minimum cut:

$$\lambda_{\text{up}} = \min_{\text{tissue}\to\text{LN}} \left(\sum_{e \in \text{cut}} K_{ij} \right).$$

If

$$L_{\text{verify}} > \lambda_{\text{up}},$$

then the antigen/reporting pipeline is bottlenecked: not all candidate distinctions can reach central verification hubs within the required window. This produces APC queuing, delayed admission, or coarse-grained relay.

For the downstream (effector return) channel, containment failure occurs when

$$\tau_{\text{down}} + \tau_{\text{expansion}} > \frac{\ln(P_{\text{crit}}/P_0)}{r_{\text{pathogen}}},$$

where P_0 is initial pathogen-like load and r_{pathogen} is its effective growth rate. If effector return and expansion take longer than the local pathogen doubling time allows, the infection escapes containment before verified systemic action arrives.

This graph model predicts topological phase transitions: when upstream or downstream min-cut capacity is exceeded, the system switches from controlled verification to local containment failure. The control number SLR from Section 9 captures this at the coarsest level; the graph model adds spatial structure.

12.3 Cross-scale verification mismatch

A classification may be locally correct at one scale while globally costly at another. For example,

$$\pi_{\text{cell}}(d) = \text{threat} \quad \text{but} \quad \ell_{\text{tissue}}(A_{\text{eliminate}}) \gg \ell_{\text{threat}}.$$

B1 therefore treats immune verification as a multiscale alignment problem, not merely a local recognition problem. Any cross-scale claim must declare which S_{immune} object is being modeled.

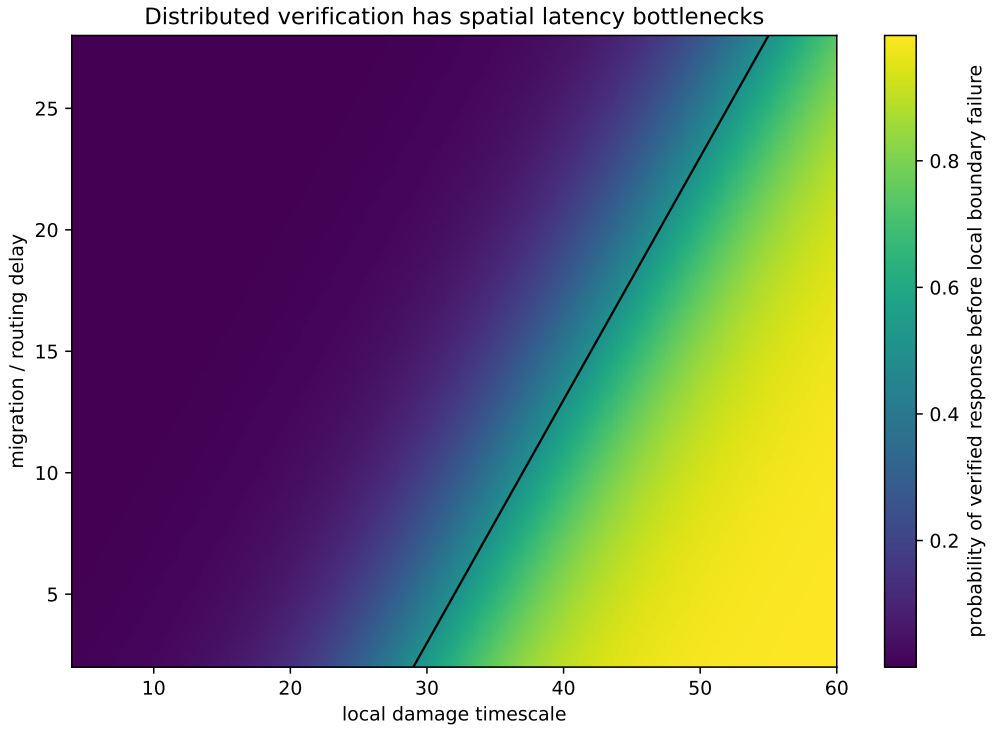

Figure 10: Distributed verification has spatial latency bottlenecks. The 50% contour marks where synthetic verification delay matches local damage timescale. This is a non-clinical topology demonstration.

13 Failure Modes of Immune Boundary Verification

Table 5: FDS-B1 immune verification failure modes. Disease/process terms are non-clinical illustrative correlates, not diagnostic categories.

Failure mode	FDS form	Non-clinical illustrative correlate
Critical exclusion	boundary-relevant distinction not admitted	immune-evasion-like or delayed-recognition-like process
False admission	low-quality or irrelevant distinction admitted	unnecessary activation / distraction-like verification burden
False positive classification	protected structure classified as threat	autoimmune-like self-boundary misclassification
False negative classification	threat classified as harmless or tolerated	tumor-evasion-like or pathogen-evasion-like failure
Verification saturation	$L_{\text{verify}} > C_{\text{verify}}$ over time	systemic alarm-like or cytokine-storm-like overload analogue
Over-elimination	action cost exceeds threat reduction	collateral tissue-damage-like regime
FDS-resolution failure	alarm fails to return to lower-cost maintenance	chronic-inflammation-like persistent verification load
Memory drift	prior classification no longer fits current context	stale memory / overgeneralization-like process
Tolerance breakdown	protected or symbiotic distinction loses tolerance status	tolerance-instability-like process
Adversarial false compression	threat and protected class overlap under finite resolution	molecular-mimicry-like semantic overlap
Spatial latency failure	verified response arrives after local boundary degradation	local containment failure analogue

14 Non-Clinical Wet-Lab Proxy Map

B1 can be tested only if variables are translated into measurable proxies without collapsing proxies into ground truth. Table 6 provides candidate proxy families. These are research proxies, not clinical biomarkers.

Table 6: Candidate non-clinical proxy map. These proxies support model testing; they do not define clinical states.

FDS-B1 variable	Candidate measurable proxy	Caution
C_{class}	TCR/BCR repertoire Shannon entropy, richness, clonality, Simpson diversity, tissue-local repertoire diversity, antigen-specific functional assays	Repertoire entropy is not full classification capacity
L_{verify}	antigenic load, damage/danger signal panels, number of unresolved candidate distinctions, tissue perturbation load	Load must be pre-specified and scale-declared
H_{class}	classifier uncertainty from single-cell state models, repertoire ambiguity, mixed posterior class assignments, multi-omic state entropy	Statistical uncertainty is not identical to biological indecision
A_{alarm}	cytokine/chemokine panels, IL-6-like/TNF-like/IL-1-like inflammatory load composites, IFN signatures, acute-phase-like markers	Candidate alarm-load proxy, not direct measurement of Δ_{verify}
Φ_{eff}	ATP/metabolic flux, glycolysis/OXPHOS balance, oxygen/nutrient availability, mitochondrial function, proliferation capacity, exhaustion-like state markers	Exhaustion-like markers are phenotype proxies, not resource itself
C_{resolve}	resolution mediator dynamics, decline of alarm markers, return of tissue-state measures, repair signatures	Resolution is active re-normalization, not passive disappearance
D_{maint}	cumulative alarm exposure, unresolved damage load, repeated perturbation burden, recovery-time lengthening	Maintenance debt is a model variable requiring calibration
C_{spatial}	APC migration time, lymph-node response timing, effector return latency, tissue-local response delay	Spatial proxies are highly model- and tissue-dependent

A composite non-clinical verification-pressure proxy can be written as

$$\Delta_{\text{verify}}^{\text{proxy}} = z(L_{\text{antigen/damage}}) + w_H z(H_{\text{class}}) + w_A z(A_{\text{alarm}}) - w_C z(C_{\text{class}}^{\text{proxy}}) - w_\Phi z(\Phi_{\text{eff}}^{\text{proxy}}),$$

where $z(\cdot)$ denotes a standardized proxy and weights are pre-registered or fitted in a non-clinical model. Nonlinear rises in cytokine-like panels may contribute to $A_{\text{alarm}}^{\text{proxy}}$, but they do not by themselves define Δ_{verify}.

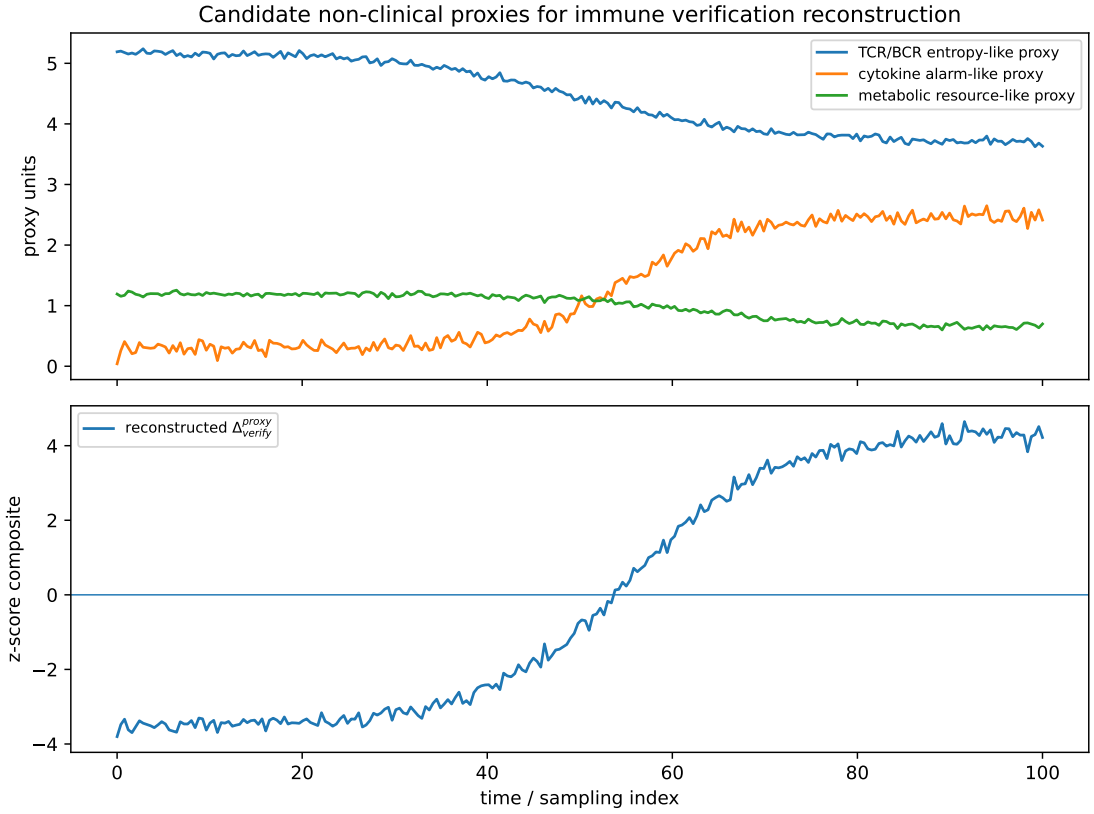

Figure 11: Candidate proxy reconstruction. Synthetic TCR/BCR entropy-like, cytokine alarm-like, metabolic resource-like, and classification-entropy-like signals are combined into a verification-pressure proxy. The figure demonstrates a reporting template, not a validated biomarker.

15 Four Non-Clinical Verification Regimes

A useful low-dimensional summary uses verification sensitivity and resolution/control capacity. This yields four non-clinical regimes: under-verification, tolerant surveillance, verified adaptive response, and chronic alarm/hyper-reactive-like response. These are not clinical categories.

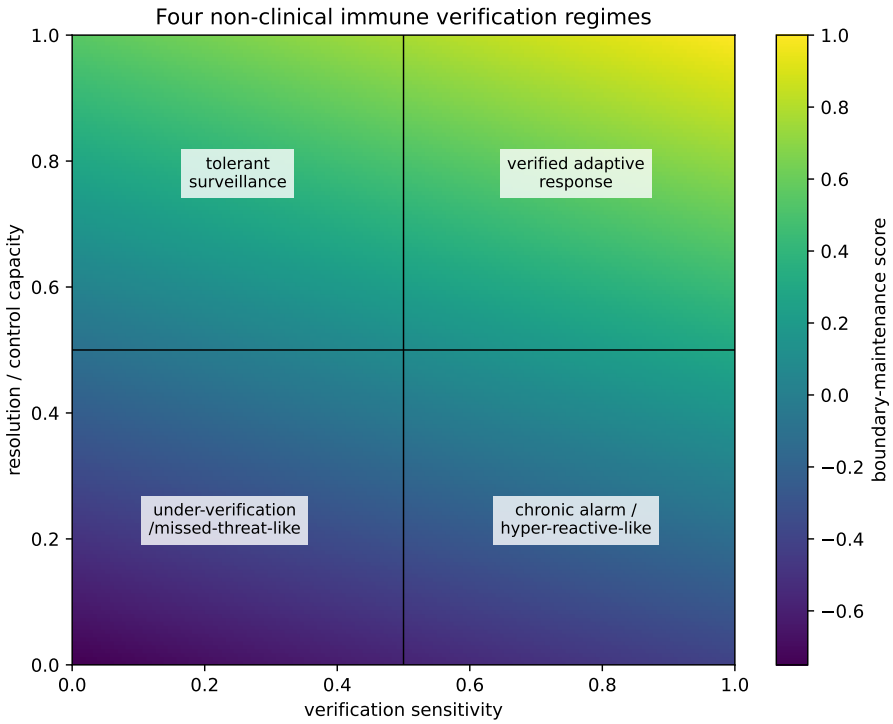

Figure 12: Four non-clinical immune verification regimes. Disease-language analogues are deliberately written as "-like" regimes and are not diagnostic categories.

16 Relation to Existing Immunological Theories

Table 7: Relationship between B1 and existing immune theories.

Existing framework	What it explains	What B1 adds
Self/non-self theory	identity distinction and host protection	self/non-self as one projection of multiaxis boundary verification
Clonal selection	adaptive specificity, expansion, memory	repertoire as finite classification alphabet under capacity constraints
Pattern-recognition models	innate recognition of conserved microbial signals	sensing/admission layer in a verification pipeline
Danger model	damage and danger as immune triggers	danger as one input to finite admission and action selection

Existing framework	What it explains	What B1 adds
Immune network theory	regulatory interactions and immune-system connectivity	boundary-verification topology and cross-scale alignment
Trained innate immunity	history-dependent innate response state	non-adaptive verification memory and prior update
Resolution biology	active termination and repair of inflammation	FDS-resolution as return to lower-cost boundary maintenance
Systems immunology	multiscale data-driven immune dynamics	FDS claim cards, capacity deficits, failure modes, and proxy ledgers

17 B1 Claim Cards

Table 8: B1 claim registry.

Claim	Statement	Demotion condition
B1-1 Boundary verification	Immune systems can be modeled as finite-capacity boundary-verification architectures. Status: B-L2.	Demote if immune response can be fully organized without finite classification, memory, resource, boundary, or verification roles.
B1-2 Admission before action	Immune action requires admission and classification of candidate distinctions before downstream response. Status: B-L2/B-L3.	Demote if action is empirically independent of admission, classification, memory, or context in the specified model.
B1-3 Multiaxis classifier	Immune classification is better modeled as a boundary-state vector than as a single self/non-self label. Status: B-L2.	Demote if a one-dimensional label captures all relevant verification behavior in the declared system.
B1-4 Verification saturation	High candidate-distinction load should produce delay, broad default action, reduced specificity, false positives/negatives, or FDS-resolution failure. Status: B-L3.	Demote if increasing verification burden produces no change in accuracy, delay, alarm load, resource use, or resolution.
B1-5 Memory-tolerance tradeoff	Immune memory reduces future verification cost but can produce drift, overgeneralization, or tolerance risk. Status: B-L3.	Demote if memory has no measurable cost, drift, or threshold effect in the declared system.

Claim	Statement	Demotion condition
B1-6 Adversarial sabotage	Some perturbations actively consume verification capacity or modify classification. Status: B-L3.	Demote if evasion-like processes never alter Y, π, M, Φ, or C_{verify} in declared models.
B1-7 Distributed spatial latency	Immune verification is constrained by routing, migration, amplification, and return times. Status: B-L3.	Demote if spatial latency has no measurable effect in systems where local damage timescale is shorter than verification time.

18 Numerical Demonstrations and Reproducibility

The numerical models are deterministic synthetic demonstrations. They are not fitted immune models, animal studies, human-subject data, patient-level data, clinical prediction tools, or disease simulations. They exist to make the normal-form variables auditable and reproducible.

The code in `code/generate_results.py` generates all figures and CSV tables with fixed parameters. It includes simulations for the boundary-verification pipeline, multi-signal admission, classification entropy, repertoire-capacity contraction, normal-form verification saturation, maintenance-debt feedback, adversarial distinction injection, false-compression risk, spatial latency, proxy reconstruction, four-regime mapping, and regulated elimination loss.

19 Safety Firewall Card

Table 9: B1 safety firewall card.

Question	Answer
Does B1 diagnose disease?	No
Does B1 treat disease?	No
Does B1 prevent disease?	No
Does B1 recommend intervention?	No
Does B1 predict patient outcome?	No
Does B1 alter standard care?	No
Does B1 define standard of care?	No
Does B1 justify self-experimentation?	No
Does B1 define clinical biomarkers?	No
Does B1 replace immunology?	No

20 Crucial Divergent Predictions

B1 is compatible with existing immunological theories but generates predictions that diverge from them under specific finite-capacity conditions. These divergences provide testable contrasts that can distinguish the FDS-B1 framework from alternatives.

Table 10: Crucial divergent predictions: high danger signal under verification bandwidth exhaustion.

Condition	Danger model prediction	Clonal selection prediction	FDS-B1 prediction
High danger signal, but verification bandwidth exhausted (VLR $\gg 1$) and decoy load high (ADR $\gg 0$)	Strong escalation of inflammation and attack	If matching clones exist, specific expansion and elimination	Possible delay, false tolerance, low-specificity alarm, false compression, or local containment failure
Persistent low-grade alarm with full repertoire	Alarm should trigger escalating response	Repertoire available for specific response	If SPI < 1 and maintenance debt is low, system may remain in tolerant surveillance without escalation
Adversarial false compression of self and threat	Danger signal should disambiguate	Clonal specificity should separate	Under finite C_{class}, may produce false negative classification even without canonical molecular mimicry, if effective feature distance is compressed by coarse-graining or context

20.1 Distinguishing verification bandwidth exhaustion from immune exhaustion

B1 distinguishes two mechanistically distinct failure modes that are often conflated in immunological literature.

Verification bandwidth exhaustion occurs when the number of candidate distinctions exceeds verification capacity:

$$\text{VLR}(t) > 1.$$

Primary manifestations: increased classification entropy, delayed response, false positives and false negatives, coarse-grained or low-specificity alarm. Individual cells need not show exhaustion phenotypes.

Immune effector exhaustion occurs when per-cell effector function degrades:

$$C_{\text{effector}}(t) = N_{\text{eff}}(t) \cdot e_{\text{cell}}(t), \qquad \dot{e}_{\text{cell}} < 0.$$

Primary manifestations: reduced per-cell cytotoxicity, cytokine production, or proliferative capacity, often associated with chronic stimulation, metabolic limitation, and inhibitory receptor expression.

The two modes can co-occur but are not identical. B1's proxy table (Table 6) treats exhaustion-like surface markers as phenotype proxies for Φ_{eff}, not as direct measures of verification bandwidth. The divergent prediction is that VLR > 1 can produce false negative classification and delayed response in the absence of per-cell exhaustion, and conversely that exhausted cells can maintain accurate classification if total verification capacity remains adequate.

20.2 Proxy estimation recipe for dimensionless control numbers

Each dimensionless control number can be estimated from the candidate proxies in Table 6. A minimal reporting template for VLR is

$$\text{VLR}^{\text{proxy}}(t) = \frac{\epsilon + \tilde{L}_{\text{antigen/damage}}(t) + w_H\, \tilde{H}_{\text{class}}(t) + w_A\, \tilde{A}_{\text{alarm}}(t)}{\epsilon + \tilde{C}^{\text{proxy}}_{\text{class}}(t) + \tilde{\Phi}^{\text{proxy}}_{\text{eff}}(t) + \tilde{C}^{\text{proxy}}_{\text{spatial}}(t)},$$

where $\tilde{x}$ denotes a rank-normalized or min-max normalized proxy in $[0, 1]$ and ϵ is a small positive constant to avoid denominator instability. The weights w_H, w_A are pre-registered or fitted in a non-clinical model.

This formula is not a validated biomarker. It is a reporting template that forces explicit mapping from proxy measurements to the dimensionless control numbers. The template can be refined, replaced, or rejected as data accumulate. Its purpose is transparency: any reader can see which proxies enter the numerator and denominator, how they are weighted, and whether the resulting $\text{VLR}^{\text{proxy}}$ crosses 1 under the specified conditions. Because normalized proxy ratios depend on the reference window and normalization scheme, the operational threshold for $\text{VLR}^{\text{proxy}} > 1$ must be fixed prospectively within each declared model system.

21 Minimal Non-Clinical Experimental Tests

The following test frameworks are designed for model systems, organoid cultures, or controlled animal-model experiments. They are not clinical trials, diagnostic procedures, or human-subject studies. Each test specifies a perturbation and a predicted FDS-B1 exit signature. For each test, the primary scale, proxy variables, perturbation schedule, VLR/ADR/SLR estimation rule, and exit signature must be pre-specified before model fitting or data collection, to avoid post hoc reinterpretation.

21.1 Verification DDoS test

Perturbation: Fix danger signal at a moderate level. Increase the number of decoy antigens, irrelevant candidate distinctions, or low-value molecular patterns presented to the immune verification pipeline (e.g., via repeated injections of irrelevant antigens, killed commensal mixtures, or altered-peptide ligands with matched PAMP context).

Prediction: As the decoy load raises VLR above 1, response delay should increase, classification entropy should rise, and false positive/negative rates should increase—even though the genuine threat signal has not changed. The effect should be stronger when ADR is high.

Controls: Show that the same decoy load does not produce equivalent delay or classification degradation when VLR is kept below 1 by expanding verification capacity (e.g., by increasing repertoire diversity, metabolic resources, or APC bandwidth).

21.2 Bandwidth exhaustion versus effector exhaustion test

Perturbation: In an in vitro or ex vivo system (e.g., splenocyte culture, lymph-node organoid, or T cell stimulation assay), compare two conditions: (1) high VLR from diverse peptide or antigen mixtures without chronic stimulation, and (2) chronic stimulation with a single antigen to drive effector exhaustion but keep VLR moderate.

Prediction: In condition (1), classification entropy, delayed response, and false classification should occur without upregulation of exhaustion-like surface markers. In condition (2), exhaustion markers appear but classification accuracy for the chronic antigen may remain preserved if total verification capacity suffices.

Controls: Measure both per-cell effector function (cytokine production, cytotoxicity markers) and population-level classification (entropy, delay, error). Cross-condition comparisons should separate the two failure modes.

21.3 Spatial bottleneck test

Perturbation: In a tissue-model or lymph-node organoid system, reduce APC reporting bandwidth or migration speed (e.g., by pharmacological disruption of chemotactic signaling, reduction of APC numbers, or physical constriction of lymphatic drainage channels). Measure the delay between local perturbation and appearance of verified effector cells in the tissue compartment.

Prediction: When the spatial latency ratio SLR exceeds 1, local pathogen-like replication or damage should outpace verified effector return, producing local containment failure. The failure should be reproducible across model systems with matched SLR but varied molecular mechanisms.

Controls: Show that increasing APC bandwidth or effector migration speed (reducing SLR below 1) restores local containment even when the total verification demand remains unchanged.

22 Limitations

First, B1 is non-clinical. It cannot be used to diagnose, treat, prevent, or manage disease. Second, it does not replace immunology, molecular biology, systems immunology, clinical immunology, pathology, microbiology, vaccinology, or immunotherapy research. Third, its variables are systems-level abstractions; every empirical use requires a scale-declared mapping. Fourth, the normal-form equations are hypothesis-generating models, not universal immune laws. Fifth, repertoire entropy, cytokine panels, metabolic flux, exhaustion-like markers, and spatial-latency measures are candidate proxies, not ground truth. Sixth, disease and pathology terms are illustrative correlates only. Seventh, adversarial immune-evasion language is a model-class interpretation, not a replacement for detailed mechanisms. Eighth, B1 does not claim that immune systems are globally optimal. It claims only that immune action can be organized as boundary-loss minimization under finite capacity. Ninth, cross-scale bridge claims can fail locally without falsifying the FDS core.

23 Conclusion

Immunity is not reducible to killing. It is a finite-capacity boundary-verification architecture that admits, classifies, remembers, tolerates, repairs, resolves, and only sometimes eliminates candidate biological distinctions. B1 therefore treats immune failure not as weakness or excess alone, but as failure of verification, action selection, resolution, adversarial robustness, memory updating, spatial routing, or cross-scale boundary alignment. The value of the reconstruction is not clinical instruction. It is an audit grammar for systems immunology: specify the boundary, declare the scale, map the alphabet, estimate the verification capacity, track the deficit, identify failure modes, and test the model with non-clinical proxies.

A Notation Summary

Table 11: Notation summary.

Symbol	Meaning
L_{verify}	task-relevant immune verification demand
C_{verify}	accessible verification capacity
Δ_+	positive verification deficit $[L - C]_+$
Z	integrated verification pressure
x	verified boundary-control state in the normal form
$r(t)$	fold-control parameter
A_{alarm}	alarm-like state variable
D_{maint}	maintenance debt
Φ_{eff}	effective resource budget
γ_R	resolution capacity
$\mathcal{A}_{\text{immune}}^{\text{eff}}$	effective immune classification alphabet
H_{rep}	repertoire entropy candidate proxy
H_{class}	classification uncertainty
C_{spatial}	spatial routing and timing capacity

B Data and Code Availability

The LaTeX source, generated figures, CSV tables, and Python code are included in the accompanying package. The simulations are deterministic and nondimensional; all parameters are fixed in `code/generate_results.py` and are not fitted to biological data. No animal studies, human-subject data, patient-level data, clinical prediction tools, or disease simulations are used. The code is designed for conceptual reproducibility and model inspection, not biological calibration.

References

[1] Y. Wu, *FDS-B0: Biomedical Bridge Registry: A Non-Clinical Framework for Boundary Maintenance, Disease-Model Translation, and Safety Firewalls in Finite Distinction Systems*, Zenodo, 2026. doi:10.5281/zenodo.20312983.

[2] Y. Wu, *Active Finite Distinction Systems: A Formal Core for Boundary Maintenance under Finite Capacity*, Zenodo, 2026. doi:10.5281/zenodo.20158923.

[3] F. M. Burnet, *The Clonal Selection Theory of Acquired Immunity*, Cambridge University Press, 1959.

[4] C. A. Janeway, Jr., Approaching the asymptote? Evolution and revolution in immunology, *Cold Spring Harbor Symposia on Quantitative Biology* **54**, 1–13 (1989). doi:10.1101/SQB.1989.054.01.003.

[5] P. Matzinger, Tolerance, danger, and the extended family, *Annual Review of Immunology* **12**, 991–1045 (1994). doi:10.1146/annurev.iy.12.040194.005015.

[6] R. Medzhitov and C. A. Janeway, Jr., Decoding the patterns of self and nonself by the innate immune system, *Science* **296**, 298–300 (2002). doi:10.1126/science.1068883.

[7] T. Pradeu, *The Limits of the Self: Immunology and Biological Identity*, Oxford University Press, 2012. doi:10.1093/acprof:oso/9780199775286.001.0001.

[8] M. G. Netea et al., Trained immunity: A program of innate immune memory in health and disease, *Science* **352**, aaf1098 (2016). doi:10.1126/science.aaf1098.

[9] C. N. Serhan, Resolution phase of inflammation: novel endogenous anti-inflammatory and proresolving lipid mediators and pathways, *Annual Review of Immunology* **25**, 101–137 (2007). doi:10.1146/annurev.immunol.25.022106.141647.

[10] E. J. Wherry, T cell exhaustion, *Nature Immunology* **12**, 492–499 (2011). doi:10.1038/ni.2035.

[11] M. M. Davis, C. M. Tato, and D. Furman, Systems immunology: just getting started, *Nature Immunology* **18**, 725–732 (2017). doi:10.1038/ni.3768.

[12] A. J. Brown, I. Snapkov, R. Akbar, M. Pavlović, E. Miho, G. K. Sandve, and V. Greiff, Augmenting adaptive immunity: progress and challenges in the quantitative engineering and analysis of adaptive immune receptor repertoires, *Molecular Systems Design & Engineering* **4**, 701–736 (2019). doi:10.1039/C9ME00071B.

[13] K. Murphy and C. Weaver, *Janeway's Immunobiology*, 9th ed., Garland Science/Taylor & Francis, New York, 2016.

[14] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proceedings of the National Academy of Sciences* **113**, 3932–3937 (2016). doi:10.1073/pnas.1517384113.

[15] S. E. Irac, M. S. F. Soon, N. Borcherding, et al., Single-cell immune repertoire analysis, *Nature Methods* **21**, 777–792 (2024). doi:10.1038/s41592-024-02243-4.

[16] P. S. L. Schäfer, et al., Integrating single-cell multi-omics and prior biological knowledge for a functional characterization of the immune system, *Nature Immunology* **25**, 405–417 (2024). doi:10.1038/s41590-024-01768-2.

[17] P. Cruz de Casas et al., Same yet different—how lymph node heterogeneity affects immune responses, *Nature Reviews Immunology* **24**, 358–374 (2024). doi:10.1038/s41577-023-00965-8.

[18] C. R. Melo-Silva and L. J. Sigal, Innate and adaptive immune responses that control lymph-borne viruses in the draining lymph node, *Cellular & Molecular Immunology* **21**, 999–1007 (2024). doi:10.1038/s41423-024-01188-0.

[19] M. Roerden and S. Spranger, Cancer immune evasion, immunoediting and intratumour heterogeneity, *Nature Reviews Immunology* **25**, 353–369 (2025). doi:10.1038/s41577-024-01111-8.

[20] C. Maguire et al., Molecular mimicry as a mechanism of viral immune evasion and autoimmunity, *Nature Communications* **15**, 9403 (2024). doi:10.1038/s41467-024-53658-8.

[21] X. Wu, M. McDermott, and A. L. MacLean, Data-driven model discovery and model selection for noisy biological systems, *PLoS Computational Biology* **21**, e1012762 (2025). doi:10.1371/journal.pcbi.1012762.