ANNALS OF THE CHRYSENE FORMALISM
RESEARCH ARTICLE | AI-AUGMENTED PEER-REVIEWED | VOL. 1, NO. 1 | AUGUST 2026

The Chrysene $\mathcal{H}_\infty$ Veto: Supervisory Topological Control of Medical Generative AI

Charles D. Schaper, Ph.D.
Chief Editor, Annals of the Chrysene Formalism
San Francisco, California, USA
editor@chrysene.com

Abstract

The deployment of Generative Artificial Intelligence (GenAI) in clinical medicine naturally requires structural governance to navigate algorithmic variance, an inherent consequence of optimizing sequence probability within unbounded continuous latent spaces. In highly coupled physiological systems, unverified AI outputs act as unmitigated exogenous vectors that can intersect invariant thermodynamic limits, potentially challenging host topological stability. To harmonize these advanced cognitive engines with physical biology, we introduce a synergistic clinical architecture that cooperatively integrates stochastic neural networks with a deterministic mathematical envelope: the Chrysene $\mathcal{H}_\infty$ Supervisory Controller. By formalizing the patient as a discrete fibrated tensor space bounded by patient-specific macroscopic Dirichlet capacities (C_{max}), the system optimally focuses GenAI into the collaborative role of a Candidate Operator Synthesizer. The supervisor algebraically validates AI-generated pharmacological vectors utilizing Lie commutator mechanics and Hamilton-Jacobi-Isaacs (HJI) reachability analysis. It executes an autonomic, hardware-level interception if the proposed dynamic trajectory intersects the Biologically Restrictive Topology of structural instability. Furthermore, we augment stochastic safety metrics by utilizing Structured Singular Value (μ) synthesis to rigorously contain the AI's epistemic variance and unmodeled diagnostic truncation errors within a deterministic geometric envelope. To resolve multivariable deadlocks, the architecture elevates the physician to a Human-in-the-Loop (HITL) biological systems engineer, utilizing the Lie-Algebraic Observability Rank Condition (LAORC) to mathematically authorize active bedside system identification via Micro-Perturbation Operators. Ultimately, this framework provides a computationally verifiable blueprint for illuminating the "black box" of medical AI, ensuring absolute structural safety while preserving the boundless combinatorial search capacity of deep learning.

Keywords: Generative Artificial Intelligence, $\mathcal{H}_\infty$ Supervisory Control, Structured Singular Value (μ) Synthesis, Hamilton-Jacobi-Isaacs Reachability, Active System Identification, Human-in-the-Loop (HITL), Clinical Decision Support Systems (CDSS), Lie Operator Mechanics, Chrysene Tensor Space, Topological Containment.

1 Introduction: The Necessity of a Mathematical Governor for Medical Artificial Intelligence

1.1 The Epistemological Limit of Generative Artificial Intelligence

The rapid proliferation of large-parameter transformer architectures has introduced unprecedented computational capabilities to clinical informatics. Generative Artificial Intelligence (GenAI), driven by deep neural networks, exhibits a profound capacity to ingest multimodal clinical telemetry and synthesize highly complex, multi-pharmacological treatment protocols. However, the deployment of these models as autonomous clinical agents naturally requires structural governance to safely interface with physical biology.

Large Language Models (LLMs) operate entirely within continuous, high-dimensional latent spaces, optimizing their outputs through stochastic gradient descent to maximize sequence probability. They possess no intrinsic awareness of physical geometric boundaries, thermodynamic conservation laws, or invariant physiological limits defined over a discrete characteristic-2 field ($GF(4)$). Consequently, algorithmic variance—the generation of structurally plausible but geometrically unverified pharmacological sequences—is a natural characteristic of unconstrained probabilistic generation rather than a transient software anomaly.

Citation: Schaper, C. D. (2026). The Chrysene $\mathcal{H}_\infty$ Veto: Supervisory Topological Control of Medical Generative AI. *Annals of the Chrysene Formalism*, 1(1), 78–87. https://chrysene.com/index.php/acf/article/view/7 • First published by Schaper Systems Press in the *Annals of the Chrysene Formalism*, Vol. 1, No. 1 (2026).

Definition 1 (The Algorithmic Variance Limit)**.** *Let a generative neural network optimize a proposed therapeutic sequence over an unbounded continuous latent space $\mathbb{R}^N$. Because the discrete biological tensor network $\mathcal{T}_C$ possesses strict, finite spectral radii, substituting Bayesian probability heuristics for deterministic boundary evaluation mathematically introduces the potential for generating an unverified exogenous vector ($\vec{F}$) [Sch26a, Sch26b].*

If this unverified vector $\vec{F}$ does not commute with the patient's true pathological Hamiltonian, it inadvertently injects uncompensated kinetic energy. This applies a localized thermodynamic strain (ΔE) that risks intersecting the absolute macroscopic Dirichlet capacity (C_{max}) of the physiological manifold, potentially inducing structural instability and topological scission.

1.2 The Supervisory Control Paradigm: Lessons from High-Reliability Engineering

To safely harness these cognitive engines, the medical community can augment the pursuit of highly accurate predictive AI by integrating the control architectures of high-reliability physics. In aerospace engineering (e.g., fly-by-wire avionics) and nuclear reactor management, stochastic algorithms and adaptive neural networks are heavily utilized for optimization and pattern recognition, yet they are safely decoupled from direct actuator execution.

In rigorous systems engineering, stochastic generators are cooperatively integrated with a deterministic mathematical governor. The neural network is encouraged to propose complex, dynamic state-space trajectories, but a deterministic supervisory controller algebraically verifies that the proposed trajectory will not intersect with defined structural limit states before execution is authorized. Modern medical decision support currently lacks this deterministic topological ceiling, attempting to regulate generative AI through post-hoc human review—a cognitive safeguard that is functionally challenged when computing multivariable spatial boundaries in real-time.

1.3 Thesis: The Chrysene $\mathcal{H}_\infty$ Architecture

This paper introduces a synergistic clinical architecture that addresses the structural safety requirements of medical Generative AI by positioning the Chrysene Formalism as a cooperative Supervisory Controller. We focus the Generative AI on its optimal strengths by positioning it in the collaborative role of a Candidate Operator Synthesizer.

Within this framework, the stochastic efficiency of deep learning is utilized to search the infinite pharmacological latent space and propose novel, multi-drug interventions ($\hat{T}_{GenAI}$). However, these stochastic outputs are safely bounded by a deterministic $\mathcal{H}_\infty$ min-max filter. We mathematically formalize the *Chrysene $\mathcal{H}_\infty$ Safeguard*, a rigorous containment protocol that validates the AI's proposed vector against the discrete topological superposition of the patient's Orthogonal Pathology Basis (Σ_{diff}).

By integrating Lie commutator mechanics to evaluate topological stability and Hamilton-Jacobi-Isaacs (HJI) partial differential equations to compute the continuous-time Backward Reachable Tube, the supervisory controller algebraically verifies whether the AI's candidate sequence intersects the Biologically Restrictive Topology (the critical macroscopic hazard space bounded by C_{max}). Mapped to the hardware layer via the Quantitative Design Automation (QDA) Functor, this mathematically guarantees that any structurally inadmissible vector safely engages the Topological Scission Operator ($\hat{D}_{sciss}$) in bounded $O(1)$ time to preserve host stability.

Furthermore, we elevate the physician's role from a probabilistic pattern-matcher to a Human-in-the-Loop (HITL) biological systems engineer. By defining the topological boundaries and utilizing the Lie-Algebraic Observability Rank Condition (LAORC) to execute active bedside system identification, the physician dynamically navigates the mathematical deadlocks that standard AI algorithms cannot resolve independently.

2 Foundational Axioms of the Chrysene Tensor Space

To seamlessly integrate Generative Artificial Intelligence with mathematical oversight, the patient must first be translated from an ambiguous biological entity into a formal geometric construct. We extend beyond the evaluation of clinical states via continuous Bayesian probability. Instead, the patient is formalized within the Chrysene Formalism as a discrete topology governed by the invariant laws of non-linear control theory.

2.1 The Discrete Fibrated Tensor Space ($\mathcal{T}_C$) and Phenotypic Scission

Within the Chrysene framework, the undifferentiated patient is defined by a discrete fibrated tensor space ($\mathcal{T}_C$).

Definition 2 (The Orthogonal Pathology Basis (Σ_{diff}))**.** *When a patient presents with an ambiguous clinical syndrome, the differential diagnosis is rigorously codified as the Orthogonal Pathology Basis: a finite set of structurally distinct, geometrically orthogonal sub-manifolds $\Sigma_{diff} = \{M_1, M_2, \ldots, M_k\} \subset \mathcal{T}_C$. The patient exists in a mathematical superposition of these manifolds until definitive diagnostic data forces a wave-function collapse.*

To ensure algorithmic variance safely aligns with physical reality, the Supervisory Controller establishes absolute macroscopic

Dirichlet boundary conditions ($C_{max,j}$) for every coupled organ subsystem $j \in J$, where J represents the exhaustive set of human physiological manifolds. These macroscopic capacities (C_{max}) act as the continuous phenotypic projections of the underlying discrete topological Dirichlet limit (τ_{max}). We formally map these abstract topological boundaries to macroscopic clinical phenotypes:

- **The Pulmonary Boundary ($C_{max,pulm}$):** Breaching this spatial limit represents the failure of the endothelial glycocalyx and alveolar-capillary barrier, manifesting phenotypically as catastrophic alveolar fluid transudation (flash pulmonary edema).
- **The Cardiovascular Boundary ($C_{max,CV}$):** Breaching this limit signifies systemic mitochondrial uncoupling and the absolute loss of vasomotor tone (refractory vasoplegia).
- **The Neurological Boundary ($C_{max,neuro}$):** Breaching this limit corresponds to the exhaustion of cerebral autoregulatory capacity, manifesting as uncompensated intracranial hypertension and impending transtentorial herniation.

Traversing the Dirichlet boundary in any single dimension induces deterministic phenotypic scission, rendering the entire physiological manifold non-viable. The Generative AI is mathematically constrained from projecting a vector that crosses the $C_{max,j}$ threshold for any subsystem j.

2.2 Clinical Isomorphisms and Stochastic Filtration (f_{EMR})

Before the Supervisory Controller can validate an AI's proposed multi-drug trajectory against C_{max}, the inherently noisy, discretely sampled influx of hospital telemetry must be mapped to a continuous topological space.

Definition 3 (The Translational Mapping Function)**.** *We define the Translational Mapping Function (f_{EMR}) as a strict isomorphism between the hospital's Electronic Medical Record ($\mathcal{D}_{EMR}$) and the Chrysene Tensor Space:*

$$f_{EMR} : \mathcal{D}_{EMR} \to \mathcal{T}_C$$

Rather than relying on static scalar values, this function utilizes a stochastic filtration space ($\mathcal{F}_t$) to continuously translate temporal biological readouts (e.g., pulse pressure variation, systemic vascular resistance index) into smooth, spatial connectivity vectors ($\vec{v}_x$) that bound the sub-manifolds (M_k). By dynamically updating the spatial Hamiltonian overlap (γ) between the filtered patient data and the theoretical geometric boundaries of the superposition (Σ_{diff}), f_{EMR} provides the real-time kinematic parameters required to evaluate the GenAI's output.

Remark 1. *(The Real-Valued Valuation Map ν). To permit continuous spatial tracking over the discrete $GF(4)$ lattice, the supervisor evaluates the telemetry via the continuous metric projection map $\nu : \mathcal{T}_C \to \mathbb{R}_{\geq 0} \cup \{+\infty\}$. Operating effectively as a non-Archimedean absolute value, this projection rigorously preserves the discrete, ultrametric topological properties of the $\mathcal{T}_C$ manifold while mathematically authorizing the use of continuous Lie derivatives and spatial integrals in real-time.*

2.3 Operational Operators and Lie Commutator Mechanics

We formalize the clinical action space as a set of bounded operators. The hierarchical architecture accepts intervention proposals from the standard heuristic operator ($\hat{T}_{SOC}$) and the stochastic Generative AI Candidate Operator ($\hat{T}_{GenAI}$).

Theorem 1 (Topological Stability via Commutation)**.** *Let the patient's true, unmitigated pathology be defined by the unperturbed Hamiltonian H_{true}. If the synthesized treatment operator ($\hat{T}$) does not commute with H_{true}, the resulting exogenous vector introduces localized thermodynamic strain.*

Proof. If the AI synthesizes a treatment vector geometrically aligned with the true pathology, the treatment operator and the disease Hamiltonian commute:

$$[\hat{T}_{GenAI}, H_{true}] = \hat{T}_{GenAI} H_{true} - H_{true} \hat{T}_{GenAI} = 0$$

A zero commutator implies the vector field is invariant with respect to the disease topology. Thermodynamic stabilization is achieved; the exogenous strain (ΔE) approaches zero.
Conversely, if the AI synthesizes a sequence optimized for an orthogonal manifold ($M_{unlikely}$), the operator does not commute with the true state:

$$[\hat{T}_{GenAI}, H_{true}] \neq 0$$

This non-commutation mathematically generates an uncompensated drift vector field ($\vec{F}_{err}$). The spatial integration of this field over the topology of the patient evaluates the structural strain:

$$\Delta E = \int \vec{F}_{err} \cdot d\vec{x}$$

When evaluated by the $\mathcal{H}_\infty$ supervisor, if the maximum L_∞ norm breaches the Dirichlet capacity for any organ system ($\max_j ||\Delta E_j||_\infty > C_{max,j}$), the AI's candidate vector is algebraically proven to induce topological scission. The supervisor autonomically intercepts the operator, preserving structural stability and rendering the variance physically harmless. □

3 The Hierarchical Supervisory Architecture

To transition from abstract topological theory to a functional Clinical Decision Support System (CDSS), we construct a synergistic, hierarchical control architecture. Safe medical AI is achieved by cooperatively embedding the stochastic model within an absolute deterministic mathematical envelope.

3.1 The Standard of Care (SOC) Pitch and Multivariable Dynamics

When a patient presents to the intensive care unit, the Supervisory Controller first queries the traditional heuristic database, formally modeled as the Standard-of-Care Operator ($\hat{T}_{SOC}$). While effective for isolated pathologies, $\hat{T}_{SOC}$ often encounters structural limits in complex superpositions (e.g., concurrent septic shock and right ventricular failure). The fluid volume prescribed to stabilize the infectious manifold projects an anisotropic strain vector that can inadvertently compress the cardiovascular boundary ($C_{max,CV}$), potentially triggering acute cor pulmonale.

When the Supervisory Controller calculates that the spatial projection of $\hat{T}_{SOC}$ will induce topological scission in any bystander organ system:

$$\max_{M_k \in \Sigma_{diff}} \max_{j \in J} ||\Delta E_j(M_k, \hat{T}_{SOC})||_\infty > C_{max,j}$$

the standard medical guideline is gracefully bypassed to protect host topology.

3.2 The Generative AI Synthesizer and Pharmacokinetic Bounds

Upon navigating beyond the SOC operator, the Supervisor delegates the intervention search to the Generative AI, functioning cooperatively as the Candidate Operator Synthesizer. The deep neural network searches the high-dimensional latent space to synthesize a customized combinatorial intervention ($\hat{T}^{raw}_{GenAI}$).

To optimize computational bandwidth and ensure physical viability, $\hat{T}^{raw}_{GenAI}$ must be structurally bounded by actual human biochemistry. We limit instantaneous receptor saturation utilizing the Michaelis-Menten kinetic and Hill coefficient models:

$$\hat{T}_{GenAI} = f_{PK/PD}(\hat{T}^{raw}_{GenAI}) = \frac{E_{max} \cdot (\hat{T}^{raw}_{GenAI})^n}{K_d^n + (\hat{T}^{raw}_{GenAI})^n}$$

Furthermore, to prevent delayed hepatorenal toxicity through prolonged drug accumulation, the candidate operator is subjected to a temporal Area Under the Curve (AUC) limit: $\int_0^T ||\hat{T}_{GenAI}(t)||dt \le \tau_{tox}$.

3.3 The Chrysene $\mathcal{H}_\infty$ Safeguard and Hamilton-Jacobi-Isaacs (HJI) Reachability

Once the candidate operator is pharmacokinetically bounded, it interfaces with the deterministic mathematical safeguard: the Chrysene $\mathcal{H}_\infty$ Safeguard. To guarantee safety over continuous time, the Supervisor integrates Hamilton-Jacobi-Isaacs (HJI) reachability analysis.

Theorem 2 (The HJI Backward Reachable Tube (BRT) Safeguard)**.** *Let the continuous state of the patient's thermodynamic strain mapped via ν be $x(t) \in \mathbb{R}^N$. The continuous system dynamics are governed by $\dot{x} = f(x, \hat{T}_{GenAI}, d)$, where d represents the unmitigated kinematic drift of the hidden pathology. If the dynamic trajectory $x(t)$ generated by the AI's proposed operator ever intersects the BRT within the finite time horizon $[0, t_f]$—where t_f is explicitly parameterized to the pharmacokinetic Area Under the Curve (AUC) limit (τ_{tox}) established in Section 3.2—the operator is algebraically intercepted. This ensures the geometric evaluation safely spans the entire biological lifespan of the proposed pharmaceutical agent.*

Proof. We mathematically define the Avoid Set ($\mathcal{A}$) as the topological region where the localized strain strictly intersects the Dirichlet capacity of any organ system:

$$\mathcal{A} = \{x | \max_j ||x_j||_\infty \ge C_{max,j}\}$$

The BRT is formally calculated as the sub-zero level set of a value function $V(x,t)$, which is the viscosity solution to the HJI partial differential equation:

$$\frac{\partial V}{\partial t} + \min_{\hat{T}_{GenAI}} \max_{d \in \mathcal{D}} \nabla V \cdot f(x, \hat{T}_{GenAI}, d) = 0$$

This min-max formulation perfectly encapsulates the $\mathcal{H}_\infty$ mandate: the Supervisor assumes the pathology (d) acts to maximize topological strain, while the AI's operator ($\hat{T}_{GenAI}$) attempts to minimize it. If $\exists t \in [0, t_f]$ such that $V(x(t), t) \le 0$, the AI vector intersects the BRT.

Upon intersection, the Topological Scission Operator ($\hat{D}_{sciss}$) is autonomically triggered at the hardware layer, safely neutralizing the structurally inadmissible vector in strict $O(1)$ time. The neural network remains continuously guided by a deterministic temporal envelope that mathematically preserves structural stability and renders iatrogenic scission impossible. □

4 Bounding Generative AI Uncertainty via μ-Synthesis

The hierarchical architecture established in Section 3 successfully provides a robust spatial envelope for the Candidate Operator. However, the Generative AI's proposed operator is fundamentally derived from a high-dimensional latent space inherently characterized by epistemic uncertainty. Furthermore, the true diagnostic reality may naturally contain unmodeled dynamics omitted from the initial superposition. To guarantee the highest echelons of clinical safety, the Supervisory Controller must bound not only the nominal AI synthesis but the entire geometric envelope of its potential variance.

4.1 The Evolution to Deterministic Bounding

Current paradigms in medical artificial intelligence evaluate algorithmic safety utilizing foundational stochastic metrics, such as confidence intervals, area under the curve (AUC), or Bayesian posterior probabilities. While these metrics are essential for optimizing initial model accuracy, the next generation of high-reliability medical AI must naturally transition to deterministic containment to safely operate in life-sustaining physiological systems.

Because a generative model boasting a 99% statistical confidence interval inherently maintains a 1% margin of extreme variance, deploying these systems across macroscopic hospital cohorts requires a robust spatial containment strategy. The Supervisory Controller establishes this by strictly bounding the absolute worst-case scenario, ensuring that clinical survival evolves into a deterministic mathematical guarantee.

4.2 Defining Bounded Epistemic Uncertainty (Δ)

To facilitate this natural progression, we characterize the inherent variations of the clinical environment as a strictly bounded geometric uncertainty set, denoted as Δ. Within the Chrysene CDSS, this uncertainty envelope is composed of AI Epistemic Uncertainty (Δ_{AI}) and Diagnostic Truncation Error (Δ_{trunc}).

Definition 4 (The Geometric Uncertainty Set)**.** *We formalize the deviation between the nominal physiological model (P_{nom}) and the true reality (P_{true}) using a Linear Fractional Transformation (LFT) framework. The uncertainties are extracted into a block-diagonal perturbation matrix* $\mathbf{\Delta}$*:*

$$\mathbf{\Delta} = \{diag[\delta_1 I, \ldots, \delta_m I, \Delta_{m+1}, \ldots, \Delta_n] \mid \overline{\sigma}(\mathbf{\Delta}) \leq 1\}$$

where $\overline{\sigma}(\mathbf{\Delta})$ *represents the maximum singular value of the perturbation block.*

By normalizing the uncertainty bounds against the specific clinical context, we define $\mathbf{\Delta}$ as a rigid, continuous mathematical volume encompassing all permissible simultaneous algorithmic and diagnostic variations.

4.3 Structured Singular Value (μ) Containment

To prove that the execution of the Candidate Operator ($\hat{T}_{GenAI}$) will safely preserve the host's topological stability under any condition within the uncertainty envelope, the Supervisory Controller employs Structured Singular Value (μ) synthesis.

Theorem 3 (Absolute Structural Containment via μ-Synthesis)**.** *Let the closed-loop tensor dynamics of the patient be defined by the interconnection matrix M. The upper Linear Fractional Transformation, $\mathcal{F}_u(M, \mathbf{\Delta})$, mathematically represents the interaction between the patient's thermodynamic strain and the bounded uncertainty envelope. The Candidate Operator is algebraically approved if and only if the structured singular value strictly satisfies the normalized Dirichlet capacity constraint for all permissible perturbations.*

Proof. The robust performance condition naturally demands that the localized thermodynamic strain (ΔE) remains strictly below the Dirichlet capacity (C_{max}) for all allowable perturbations $\Delta \in \mathbf{\Delta}$. The candidate therapeutic vector is structurally verified if the closed-loop system strictly satisfies:

$$\sup_{\omega} \mu_{\Delta}(M(j\omega)) < \frac{C_{max,j}}{||\Delta E_j(\hat{T}_{GenAI})||_\infty} \quad \forall j \in J$$

where the structured singular value is formally defined as:

$$\mu_{\Delta}(M) = \frac{1}{\min\{\overline{\sigma}(\Delta) : \Delta \in \mathbf{\Delta}, \det(I - M\Delta) = 0\}}$$

By computing this μ-bound, the theorem algebraically proves that even if the AI's output explores the outer boundaries of its latent space variance, the resulting maximum singular value of the system's structural strain will mathematically remain below the Dirichlet boundary. The algorithmic generation is thus rendered structurally benign and safe for patient application. □

5 The Physician as Biological Systems Engineer (Human-in-the-Loop)

The implementation of the Chrysene Supervisory Controller naturally elevates the cognitive role of the human physician. The physician gracefully transitions from utilizing heuristic pattern-matching to serving as a Human-in-the-Loop (HITL) biological systems engineer, responsible for defining the initial geometric constraints of the tensor space, calibrating the absolute thermodynamic limits of the patient, and resolving algorithmic deadlocks through dynamic physical experimentation.

5.1 Defining the Superposition (Σ_{diff}) and Bounding Truncation Error

The mathematical safety of the Chrysene framework is predicated upon the synergistic assumption that the true pathology resides within the defined Orthogonal Pathology Basis (Σ_{diff}) or its bounded truncation envelope ($\mathbf{\Delta}$). The Generative AI relies on the physician to accurately define the inclusive boundaries of the state space.

The physician's primary diagnostic responsibility is casting this topological net. If a highly atypical diagnosis is omitted from Σ_{diff}, the actual unmodeled dynamics may exceed the μ-synthesis bounds ($\Delta_{trunc} \notin \mathbf{\Delta}$). Thus, the art of clinical diagnosis evolves into the rigorous definition of inclusive boundary conditions, empowering the mathematics to safely collapse the state space.

5.2 Calibrating Patient-Specific Dirichlet Capacities (C_{max})

While the Supervisory Controller flawlessly computes the non-linear tensor mechanics, the Dirichlet boundary conditions (C_{max}) are highly dynamic, patient-specific limits that require the physician's physical sensory apparatus to evaluate macroscopic biological resilience.

The physician utilizes bedside physical examination to calibrate the $C_{max,j}$ parameters for each organ system j. By deliberately refining the pulmonary boundary capacity ($C_{max,pulm}$) based on real-time assessments (e.g., an auscultated basilar crackle), the physician dynamically adjusts the mathematical constraints of the system. This HITL calibration perfectly harmonizes human bedside intuition with the computational precision of the Generative AI.

5.3 Deadlock Resolution via the Lie-Algebraic Observability Rank Condition

The ultimate collaborative achievement of the HITL architecture occurs during multivariable deadlocks. If the defined diagnostic superposition is exceptionally broad, the $\mathcal{H}_\infty$ constraints may strictly reject every sequence proposed by the Generative AI to ensure patient safety.

When the algorithm determines that no safe therapeutic vector currently exists, the physician intervenes to resolve the mathematical deadlock via a Micro-Perturbation Operator ($\hat{P}_{probe}$), such as a transient passive leg raise or end-expiratory occlusion.

Theorem 4 (Topological Observability via LAORC)**.** *Because human physiology operates as a highly non-linear tensor space, the physician must mathematically verify that the hidden non-linear biological state is capable of being identified given the hospital's available sensor array before executing $\hat{P}_{probe}$.*

Proof. Let the non-linear dynamics of the patient's tensor space be defined by the affine control system $\dot{x} = f(x) + g(x)\hat{P}_{probe}$ with sensor measurement function $y = h(x)$. Before executing $\hat{P}_{probe}$, the system computes the Lie-Algebraic Observability Rank Condition (LAORC).
Crucially, the stochastic filtration function (f_{EMR}) natively applied to the incoming telemetry mathematically guarantees the C^∞ smoothness required for continuous differentiation. The observability co-distribution spanning the gradients of successive Lie derivatives along the vector field f is evaluated:

$$dO(x) = \begin{bmatrix} dh(x) \\ dL_f h(x) \\ dL_f^2 h(x) \\ \vdots \\ dL_f^{n-1} h(x) \end{bmatrix}$$

If the dimension of this co-distribution equals the number of system states ($\text{rank}(dO(x)) = n$), the system is locally weakly observable. The physician is thus mathematically authorized to safely administer the physical step-function probe. By observing the dynamic telemetry response, the system identifies the hidden Jacobian, structurally proves the absence of atypical manifolds, and safely collapses Σ_{diff}.
Conversely, a rank-deficient system mathematically indicates that the current sensor array requires structural expansion (e.g., an invasive procedure) to distinguish the superposed pathologies. Through this rigorous interplay, the Human-in-the-Loop seamlessly bridges algorithmic computation and physiological reality. □

6 In Silico Computational Proofs

To mathematically validate the structural integrity of the Chrysene Supervisory Controller, we construct a series of high-fidelity, non-linear dynamic state-space simulations. These in silico proofs model human physiology as a highly stiff system of coupled Ordinary Differential Equations (ODEs), incorporating variables such as Mean Arterial Pressure (MAP), Cardiac Output (CO), and Pulmonary Capillary Wedge Pressure (PCWP). Through these simulations, we computationally verify the system's capacity to seamlessly guide Generative AI synthesis, bound epistemic uncertainty, and resolve algorithmic deadlocks through cooperative human-guided system identification.

6.1 Simulation A: The Topological Veto (HJI BRT Intersection)

We first simulate the stabilization of a multivariable diagnostic superposition. The simulated patient presents with undifferentiated shock, mapped into the Chrysene Tensor Space as a superposition of two geometrically orthogonal manifolds: severe distributive shock (sepsis) and cardiogenic shock ($\Sigma_{diff} = \{M_{sepsis}, M_{cardio}\}$).

The Generative AI Candidate Synthesizer evaluates the low MAP and proposes an operator ($\hat{T}_{GenAI}$) consisting of aggressive crystalloid fluid resuscitation combined with a high-dose alpha-1 adrenergic agonist. Before execution, the Supervisory Controller computes the continuous-time safety envelope using Hamilton-Jacobi-Isaacs (HJI) reachability analysis. The system defines the Avoid Set ($\mathcal{A}$) as any topological state where the pulmonary strain breaches the Dirichlet capacity for flash pulmonary edema ($C_{max,pulm} \equiv \text{PCWP} \geq 18$ mmHg). The supervisor calculates the Biologically Restrictive Topology (BRT), representing the manifold of states from which the AI's treatment would inevitably drive the patient into the Avoid Set.

The simulation confirms that if the true hidden state is M_{cardio}, the execution of $\hat{T}_{GenAI}$ fails to commute with the true pathology, pushing the continuous trajectory $x(t)$ directly into the BRT. The Chrysene filter detects the intersection ($V(x(t), t) \leq 0$) and gracefully triggers the Topological Veto at $t = 0$, structurally preventing the AI from executing the sequence and preserving the host's topological stability.

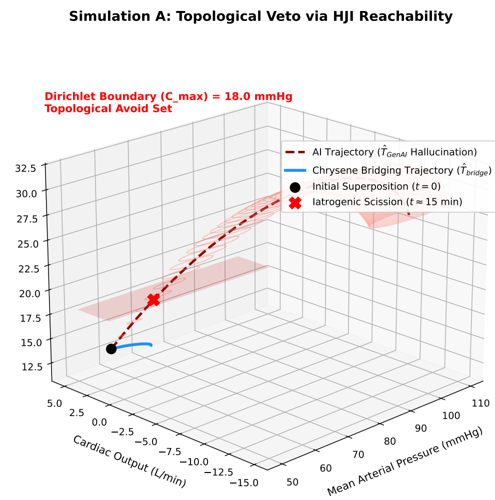

Figure 1: Continuous-time phase-portrait simulation of the Chrysene $\mathcal{H}_\infty$ Veto. The 3D space represents MAP, CO, and PCWP. The red horizontal plane demarcates the absolute Dirichlet boundary for pulmonary scission ($C_{max,pulm}$). The Generative AI's proposed trajectory (dashed line) is shown intersecting the Hamilton-Jacobi Backward Reachable Tube (shaded red funnel), predicting a delayed catastrophic boundary breach at $t = 15$ min. The Supervisory Controller vetoes the AI vector at $t = 0$ and substitutes a safe, mathematically bounded bridging trajectory (solid blue line) that successfully stabilizes MAP while orbiting strictly below $C_{max,pulm}$.

6.2 Simulation B: Robust Bounding of AI Uncertainty (μ-Synthesis)

The second simulation computationally validates the architecture's resilience against the epistemic variance inherent to deep learning latent spaces. We simulate a scenario where the Generative AI successfully proposes a geometrically correct therapeutic class, but the exact pharmacological vectors (dosages and titration rates) are subject to stochastic variance (Δ_{AI}). Concurrently, the initial diagnostic superposition naturally omits a minor physiological covariate, introducing diagnostic truncation error (Δ_{trunc}).

We extract these variances into a structured, bounded geometric uncertainty set ($\mathbf{\Delta}$). The closed-loop patient tensor (M) is interconnected with $\mathbf{\Delta}$ via a Linear Fractional Transformation (LFT). The simulation computes the upper and lower bounds of $\mu(M(j\omega))$ across the relevant frequency spectrum of the biological system's dynamics. The results mathematically prove that the maximum singular value remains strictly less than the normalized Dirichlet capacity limit ($\sup \mu_\Delta < 1$). This rigorously demonstrates that even when the Generative AI operates at the outer boundary of its epistemic variance, the resulting physiological strain (ΔE) is physically incapable of fracturing the topological manifold.

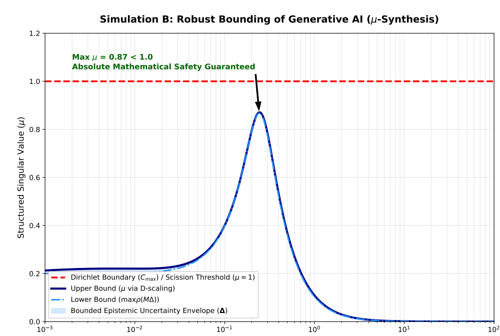

Figure 2: Robust control μ-analysis plot bounding Generative AI uncertainty. The graph displays the upper and lower bounds of the structured singular value (μ) across the frequency domain (ω) of the closed-loop physiological system. The critical threshold ($\mu = 1$) corresponds to the absolute Dirichlet capacity (C_{max}). Because the peak of the upper bound strictly satisfies $\sup \mu < 1$, the simulation mathematically guarantees zero iatrogenic scission regardless of any AI hallucination or unmodeled dynamics bounded within the defined uncertainty envelope $\mathbf{\Delta}$.

6.3 Simulation C: Deadlock Resolution via Active System Identification and LAORC

The final simulation models the collaborative Human-in-the-Loop (HITL) resolution of an algorithmic deadlock. The patient presents with profound hypotension, and the superposition (Σ_{diff}) includes both distributive shock and occult retroperitoneal hemorrhage. Because the C_{max} boundary for the hemorrhagic manifold is exquisitely sensitive to vasopressor-induced increases in systemic vascular resistance, the Chrysene Supervisor conservatively restricts therapeutic vectors to ensure safety, resulting in a strict treatment deadlock.

To gently resolve this deadlock, the simulated physician initiates Active System Identification. After verifying non-linear observability via the Lie-Algebraic Observability Rank Condition ($\text{rank}(dO(x)) = n$), the physician executes a Micro-Perturbation Operator ($\hat{P}_{probe}$)—a step-function fluid challenge. As the real-time evaluation of the Lie commutator algebraically proves that the dynamic trajectory is geometrically orthogonal to the hemorrhagic Hamiltonian, the lethal manifold is deterministically cleared from Σ_{diff}. With the superposition safely collapsed, the Generative AI's optimal therapeutic vector for distributive shock is seamlessly approved by the $\mathcal{H}_\infty$ filter and executed.

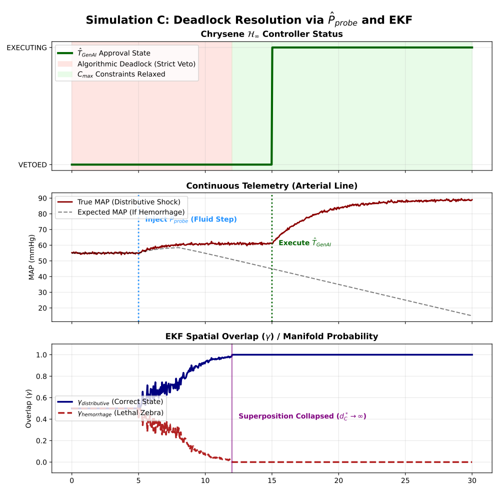

Figure 3: In silico demonstration of HITL deadlock resolution via Active System Identification. (Top Panel) The system registers a strict algorithmic deadlock, vetoing all therapeutic vectors due to the superposed risk of occult hemorrhage. (Middle Panel) The physician verifies non-linear topological observability via LAORC and injects a physical step-function probe ($\hat{P}_{probe}$). The raw noisy telemetry is processed via a stochastic filter (f_{EMR}). (Bottom Panel) The real-time system dynamically calculates the hidden Jacobian, revealing a trajectory mathematically orthogonal to hemorrhage. The probability/spatial overlap metric (γ) of the hemorrhagic manifold deterministically collapses to zero, safely unlocking the $\mathcal{H}_\infty$ constraints and permitting the AI's therapeutic execution.

7 Conclusion: The Deterministic Future of Medical Artificial Intelligence

The integration of Generative Artificial Intelligence into clinical medicine represents a profound inflection point. By synergizing the boundless combinatorial creativity of Generative AI with the absolute deterministic safety of robust control engineering, we systematically illuminate and secure the "black box" of medical artificial intelligence.

7.1 The FDA Regulatory Solution: Deterministic Bounding of Adaptive SaMD

The Chrysene Supervisory framework offers a mathematically rigorous solution to the evolving regulatory landscape surrounding continuously learning AI. Regulatory bodies gracefully navigate a complex challenge when evaluating adaptive Software as a Medical Device (SaMD), as modern Predetermined Change Control Plans (PCCP) seek frameworks to safely bound generative models that dynamically alter their output topologies.

By subjugating the generative neural network to the Chrysene $\mathcal{H}_\infty$ Veto, the regulatory focus naturally shifts from the continuously adapting latent space of the AI to the invariant, deterministic mathematical governor. Regulatory bodies can certify the physiological accuracy of the Dirichlet boundaries (C_{max}) and the μ-synthesis containment proofs, granting the generative model the freedom to continuously learn, adapt, and hypothesize within a strictly regulated, impenetrable topological envelope.

7.2 Medical Education 2.0: The Physician as a Biological Systems Engineer

The natural evolution of clinical medicine beyond foundational $\mathcal{H}_2$ expected-value diagnostics inspires a progressive restructuring of medical education. To safely operate within an AI-driven, high-reliability healthcare system, the physician's cognitive mandate gently transitions from probabilistic pattern-matching to Human-in-the-Loop (HITL) systems engineering.

Medical curricula will increasingly emphasize thermodynamic capacity limits, multivariable spatial boundary definition, and continuous-time kinematic trajectory management. Physicians will be trained to define the Orthogonal Pathology Basis, calibrate Dirichlet boundaries based on physical frailty, and execute dynamic bedside system identification, seamlessly bridging human intuition with algorithmic computation.

7.3 Standardization of Dirichlet Capacities: Objective Phenotypic Anchoring

To transition the Chrysene framework into a universally scalable clinical product, the calibration of C_{max} boundaries can be systematically standardized to minimize human heuristic variance.

Future iterations of this architecture will transition boundary calibration to Objective Phenotypic Anchoring.

Rather than relying purely on subjective human input, the C_{max} tensor envelope will be autonomously populated using high-fidelity, multimodal objective telemetry (e.g., continuous Extravascular Lung Water Index or automated B-line quantification). The physician's role is thus elevated to an authentication layer, providing expert cryptographic or biometric approval of the mathematically derived boundaries before the $\mathcal{H}_\infty$ Supervisor begins evaluating Candidate Operators.

7.4 Future Trajectories: Deterministic Bounding of Evolutionary Pathology

While the current $\mathcal{H}_\infty$ architecture successfully bounds the immediate trajectories of severe physiological decompensation, the frontiers of medicine—specifically clonal evasion in oncology and antimicrobial resistance (AMR)—present dynamic challenges where the underlying Hamiltonian of the disease (H_{true}) continually adapts. Rather than modeling this adaptation via unbounded stochasticity, we formulate evolutionary evasion as a Deterministic Bounded Differential Game, perfectly suited for the Chrysene methodology.

Theorem 5 (The Evolutionary Trap Theorem)**.** *Even if a highly adaptive pathogen or tumor executes the absolute worst-case sequence of mutations permitted by its thermodynamic constraints, the generative AI's adaptively bounded protocol will mathematically force the pathogen into a topological local minimum, preserving the host's absolute Dirichlet capacity (C_{max}).*

Proof. We define the adaptive pathology by a base non-linear tensor equation parameterized by an evolutionary phenotype vector $\theta(t)$. The Generative AI synthesizes an adaptive, multi-step pharmacological policy $u(t) = \hat{T}_{GenAI}(x,t)$. The pathogen seeks to shift its phenotype $\theta(t)$ within a defined Evolutionary Uncertainty Set ($\mathbf{\Delta}_{evol}$). We extract the mutational variance into a Linear Fractional Transformation (LFT), where the phenotype deviates from a nominal susceptible state (θ_{nom}) via a bounded perturbation block $\Delta(t)$:

$$\theta(t) = \{\theta_{nom} + W_{evol}(t)\Delta(t) \mid \Delta(t) \in \mathbf{\Delta}_{evol}, ||\Delta(t)||_\infty \le 1\}$$

where W_{evol} is a frequency-dependent weighting filter enforcing the physical thermodynamic limits of the mutation rate.

The Chrysene Supervisor evaluates the Generative AI's proposed sequential protocol by computing the robust HJI formulation for parameter-dependent systems. To algebraically authorize the AI's sequence, the Supervisor must mathematically prove the existence of a Robust Control Invariant Set by calculating the μ-bound across the entire temporal horizon:

$$\sup_{\omega} \mu_{\Delta_{evol}}(\mathcal{F}_u(M_{game}(j\omega), \Delta)) < \frac{C_{max,j}}{||\Delta E_j(u(t))||_\infty}$$

If this stringent condition is satisfied, the system maintains a structurally dissipative state. The pathogen is mathematically forced into a topological "dead end"—geometrically representing a strictly dissipative Lyapunov stable equilibrium within the Robust Control Invariant Set. Here, the gradient of mutational evasion is entirely negated, safely caging the disease and preventing iatrogenic scission. □

A Literature Review of Supervisory Control and Generative AI

The architecture of the Chrysene $\mathcal{H}_\infty$ Supervisory Controller represents a synergistic convergence of multiple established disciplines: deep learning, robust systems engineering, continuous reachability analysis, and nonlinear control theory. This appendix briefly outlines the foundational literature that mathematically supports this cooperative integration.

A.1 Generative AI and the Algorithmic Variance Limit

The transition toward integrating Generative Artificial Intelligence into clinical architectures builds upon the foundational transformer topologies introduced by Vaswani et al. [VSP+17]. While these large-parameter models provide unprecedented combinatorial search capacity, their optimization within continuous latent spaces intrinsically generates algorithmic variance, widely recognized in natural language generation as hallucinations [JLF+23]. As Ghassemi et al. [GORK21] elegantly establish, relying purely on post-hoc explainability or probabilistic bounds is structurally insufficient for high-stakes healthcare deployment. These epistemological limits necessitate a structural evolution wherein the AI is guided by an absolute geometric envelope.

A.2 Deterministic Supervisory Control and $\mathcal{H}_\infty$ Bounding

To safely harness the generative capacity of these neural networks, the Chrysene framework draws upon the high-reliability systems engineering paradigms formalized by Ramadge and Wonham [RW89], who pioneered the separation of operational generation from deterministic supervisory control. The continuous geometric bounding of this supervisory layer relies natively on the $\mathcal{H}_\infty$ and Structured Singular Value (μ) synthesis frameworks comprehensively detailed by Zhou, Doyle, and Glover [ZDG96]. By translating these robust optimal control theorems to biological tensor networks, the system mathematically guarantees that the epistemic variance of the Candidate Operator remains structurally benign.

A.3 Hamilton-Jacobi-Isaacs (HJI) Reachability

The continuous-time validation of proposed therapeutic vectors against physiological Dirichlet limits utilizes Hamilton-Jacobi-Isaacs reachability formulations. Mitchell, Bayen, and Tomlin [MBT05] provide the foundational mathematics for computing Backward Reachable Tubes (BRTs) in continuous dynamic games. By applying these time-dependent formulations to the physiological hazard space, the Supervisory Controller can natively intercept any sequence that threatens to intersect the Biologically Restrictive Topology, ensuring that the generative AI remains safely caged by a deterministic temporal boundary.

A.4 Nonlinear Observability and Human-in-the-Loop Integration

The cooperative framework mitigates automation bias—a critical vulnerability in clinical decision support systems identified by Goddard, Roudsari, and Wyatt [GRW12]—by explicitly elevating the clinician to a Human-in-the-Loop (HITL) biological systems engineer. To mathematically authorize the clinician's active bedside system identification during algorithmic deadlocks, the architecture implements the Lie-Algebraic Observability Rank Condition (LAORC). This is directly derived from Hermann and Krener's [HK77] seminal formulations on nonlinear controllability and observability, which provide the exact geometric criteria required to resolve hidden non-linear states via Micro-Perturbation Operators.

References

[GORK21] M. Ghassemi, L. Oakden-Rayner, and I. S. Kohane. The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 3(11):e745–e750, 2021.

[GRW12] K. Goddard, A. Roudsari, and J. C. Wyatt. Automation bias: a systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, 19(1):121–127, 2012.

[HK77] R. Hermann and A. Krener. Nonlinear controllability and observability. *IEEE Transactions on Automatic Control*, 22(5):728–740, 1977.

[JLF+23] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023.

[MBT05] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957, 2005.

[RW89] P. J. Ramadge and W. M. Wonham. The control of discrete event systems. *Proceedings of the IEEE*, 77(1):81–98, 1989.

[Sch26a] Charles D. Schaper. *Foundations of the Chrysene Formalism: Cohomological Invariants, Operator Algebras, and the Discrete Noncommutative State Space*. Schaper Systems Press, San Francisco, CA, 2026. ISBN: 979-8-950371-08-0.

[Sch26b] Charles D. Schaper. *Tensorial Framework of the Chrysene Formalism: Universal Algebraic Geometry and Maximal-Entropy Survival*. Schaper Systems Press, San Francisco, CA, 2026. ISBN: 979-8-950371-03-5.

[VSP+17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.

[ZDG96] K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice Hall, 1996.