

---

# AutoAxiom: Closed-Loop Framework for Autonomous Symbolic Axiom Modification via LLMs

---

Anonymous Authors<sup>1</sup>

## Abstract

Large Language Models (LLMs) have excelled at parametric optimization within fixed rule-sets, yet they rarely challenge the underlying axiomatic constraints of a domain. We present AutoAxiom, a closed-loop framework for large language model based symbolic axiom discovery with formal verification. AutoAxiom combines (1) AxiomDSL and a compiled Core IR, (2) a tripartite proposer for proposal, projection, and selection, and (3) SMT based verification with a repair mechanism. Across six domains, AutoAxiom achieves SRA above 1.0 in 5 of 6 domains, with the largest SRA values of 3.15 in Physical PDE and 1.74 in Queueing Network (Table 1). The framework discovers interpretable, “Zero-Blackbox” axioms that consistently outperform human-designed baselines and SOTA heuristics. In ablations, the full system attains 96.30% repair rate, 4.44% violation rate, and 100% success rate (Table 2). The implementation reports peak process memory in the range 338.8 to 365.9 MB across domains and per round timing breakdowns in Table 18. Our codes and models will be publicly available upon publication.

## 1. Introduction

Current approaches to AI-driven scientific discovery (Reddy & Shojaee, 2025; Fu, 2025; Gottweis et al., 2025; Kalaivani et al., 2025) Deep Learning, Symbolic Regression, and Large Language Models each face critical limitations. Deep Learning (DL) excels at high-dimensional pattern recognition but operates as a “black box” (ŞAHİN et al., 2025), failing to generalize in Out-of-Distribution (OOD) scenarios (Li et al., 2025). Symbolic Regression (SR) (Abdus-Salam et al., 2025; Yu et al., 2025; Yi et al., 2025) pur-

sues explicit mathematical laws but suffers from “semantic blindness” (DeRose, 2006) it fits equations to data without understanding their physical meaning, often producing “numerical hallucinations” (Shao et al., 2025) that overfit noise. Large Language Models (LLMs) offer emergent reasoning capabilities (Zhang et al., 2025a; Xu et al., 2025a;b), yet a critical *Verification Gap* prevents their rigorous deployment: as probabilistic engines (Zhang et al., 2025b), they can generate syntactically valid but physically impossible axioms.

Each paradigm thus fails in isolation: DL lacks interpretability, SR lacks semantic grounding, and LLMs lack formal guarantees. The question becomes: *can we combine the creativity of LLMs, the explicitness of SR, and the rigor of formal verification into a unified framework?*

In this paper, we present **AutoAxiom**, a closed-loop system that reframes scientific discovery as *axiom evolution*. As shown in figure 1, rather than optimizing a scalar reward  $S$  under fixed rules, AutoAxiom iteratively refines the axiom set  $A$  itself (Ma et al., 2025) elevating the LLM from a passive equation solver to an active *Axiom Architect* (Wang et al., 2025; Alter, 2025; Kim, 2025; Gulati et al., 2025). We summarize our contributions as follows:

- **AxiomDSL and Core IR for typed, editable axiom representation:** We introduce *AxiomDSL*, a domain-specific language that formally defines the ontology of each scientific domain. AxiomDSL compiles into a *Core Intermediate Representation (IR)* a mutable blueprint that the LLM can edit structurally. This decoupling ensures that LLM-generated logic remains agnostic to specific simulation backends, enabling portability across heterogeneous domains while preserving ontological grounding.
- **Tripartite proposer that separates proposal, projection, and selection:** Unlike blind SR that fits arbitrary equations to data, AutoAxiom constrains its search space via the AxiomDSL ontology. Strict type safety and a *Tripartite Consensus* mechanism guide exploration toward semantically meaningful regions of the hypothesis space discovering genuine physical laws rather than numerical artifacts

<sup>1</sup>Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

055  
056  
057  
058  
059  
060  
061  
062  
063  
064

- **Two tier verification and repair mechanism:** To close the Verification Gap, our *Autonomous Axiom Verification and Synthesis* system ( $\mathcal{A}^2\mathcal{V}\mathcal{S}$ ) implements a Tier-II *Satisfiability Modulo Theories (SMT) Gatekeeper*. Before any axiom reaches the simulator, Z3 subjects it to SAT/UNSAT proofs against “Red Line” safety constraints. This stage mathematically guarantees physical consistency, rejecting syntactically plausible but semantically invalid candidates.

065 We evaluate AutoAxiom across heterogeneous domains  
066 spanning discrete optimization, continuous control, and  
067 molecular design. Results demonstrate consistent improvements  
068 over most state-of-the-art baselines. For Instance,  
069 compared to well-trained reinforcement learning policies  
070 (PPO), AutoAxiom achieves comparable task performance  
071 while producing fully interpretable “white-box” axioms.  
072 Compared to symbolic regression methods, it avoids overfitting  
073 traps and discovers rules that generalize to unseen test  
074 distributions. Across all domains, AutoAxiom converges up  
075 to 2 $\times$  faster than baseline methods on challenging problem  
076 instances.

## 2. Related Work

080 **Symbolic Regression: From Blind Search to Semantic  
081 Discovery** Data-driven symbolic regression (SR) (Dong  
082 & Zhong, 2025; Holt et al., 2023) has long been the gold  
083 standard for recovering physical laws from sparse data (Wad-  
084 dayama et al., 2025). Algorithms such as PySR (Tonda,  
085 2025) excel at optimizing numerical constants within fixed  
086 functional forms. However, these methods suffer from a  
087 “semantic blind spot”: they directly manipulate primal math-  
088 ematical operators while ignoring domain ontology (Sapel  
089 et al., 2025). Consequently, in noisy environments or lack-  
090 ing effective real-world feedback, SR algorithms often fall  
091 into “overfitting paralysis” (Santos & Papa, 2022). In  
092 contrast, AutoAxiom constrains the search space through a  
093 domain vocabulary mapping (DVM). By enforcing strict  
094 type safety and ontology consistency (as defined in our Axi-  
095 omDSL), we shift the paradigm from blind combinatorial  
096 search(Aigner, 1988) to semantically guided discovery.

097 **LLM for Science: Bridging the Verification Gap via De-  
098 coupled Feedback** Large Language Models (LLMs) (Wang,  
099 2025) have demonstrated emergent capabilities in scientific  
100 reasoning (e.g., Eureka (Ma et al., 2023), FunSearch (Agli-  
101 etti et al., 2024)). However, these systems primarily operate  
102 on empirical validation loops, creating a “Verification Gap”:  
103 they optimize for numerical performance without formal  
104 logical guarantees. This leaves them prone to “mechanistic  
105 hallucinations” (Yu et al., 2024): logic that is syntactically  
106 plausible but physically unstable. AutoAxiom addresses  
107 this through a structural decoupling of reasoning and ex-  
108 ecution. Unlike standard LLM code-generation approaches

109 (Joel et al., 2024) that rely on LLMs for end-to-end simulation  
110 code, our system restricts the LLM’s role to that of a *graph editor* (Paassen et al., 2020) on a verifiable Core IR,  
111 leaving the actual execution to a deterministic, objective sim-  
112 ulation backend. This architecture enforces a rigorous **“Red  
113 Line” mechanism**: the Tier-II SMT Gatekeeper (De Moura  
114 & Bjørner, 2008) not only intercepts safety violations (e.g.  
115 energy divergence) but also injects **constraint-violation  
116 penalty factors** back into the evolutionary loop.

**Neural Operators vs. Symbolic Generalization** Deep Learning approaches, such as Physics-Informed Neural Networks (PINNs)(Lawal et al., 2022) and Deep Reinforcement Learning (e.g., PPO(Yu et al., 2022)), dominate high-dimensional control tasks. While these “black box” models achieve excellent performance across their training distributions, their opacity makes them unsuitable for safety-critical deployments. AutoAxiom aims to achieve “white box” generalization. By synthesizing explicit discrete logic (e.g., contact gating switches, saturation boundaries), our framework discovers compact axiomatic laws.

## 3. Methodology

### 3.1. Scientific Search Space: Ontological Foundation

The search space is anchored by a Domain Specific Lan-  
guage, AxiomDSL, which provides the ontological con-  
straints, and is realized through the Core IR, which enables  
computational execution.

The AxiomDSL is defined as a 5-tuple ontological basis  $\mathcal{D}$  that constrains the symbolic search space of scientific laws:  $\mathcal{D} = \langle \mathcal{N}, \mathcal{R}, \mathcal{O}, \Upsilon, \Phi \rangle$  where  $\mathcal{N} = \{n_1, n_2, \dots\}$ : Scientific Entities representing domain-specific variables (e.g., arrival rates, thermal flux).  $\mathcal{R} \subset \mathcal{N} \times \mathcal{N}$ : Semantic Relations defining the directed causal flow.  $\mathcal{O}$ : Operator Set, a curated grammar of allowed primitives including continuous math (e.g.,  $\tanh$ ,  $\exp$ ) and **Discrete State Operators** (e.g., **if-then-else**,  $\mathbb{I}_{\text{condition}}$ ) to model non-linear phase transitions.  $\Upsilon : \mathcal{N} \rightarrow \{\text{Type, Unit, Role}\}$  : Vocabulary Mapping that anchors each entity to physical reality.  $\Phi$  represents Integrity Constraints (the “Red Line”), defining non-negotiable physical boundaries (e.g.,  $\rho < 1$ ).

Scientific formulas can be viewed as instantiations within  $\mathcal{D}$ . For instance, a classic M/M/1 queueing delay axiom is a specific configuration of  $(\mathcal{N}_{\text{queue}}, \mathcal{R}_{\text{delay}}, \mathcal{O}_{\text{basic}})$ . The feasible set of all such configurations is denoted as  $\mathbb{A}$ , representing the total axiomatic search space as shown in figure 2.

While the AxiomDSL ( $\mathcal{D}$ ) serves as a high-level interface for symbolic reasoning, machine execution requires a flattened, topological structure. To bridge this gap, we define the Core IR as a representational layer designed to be processed by

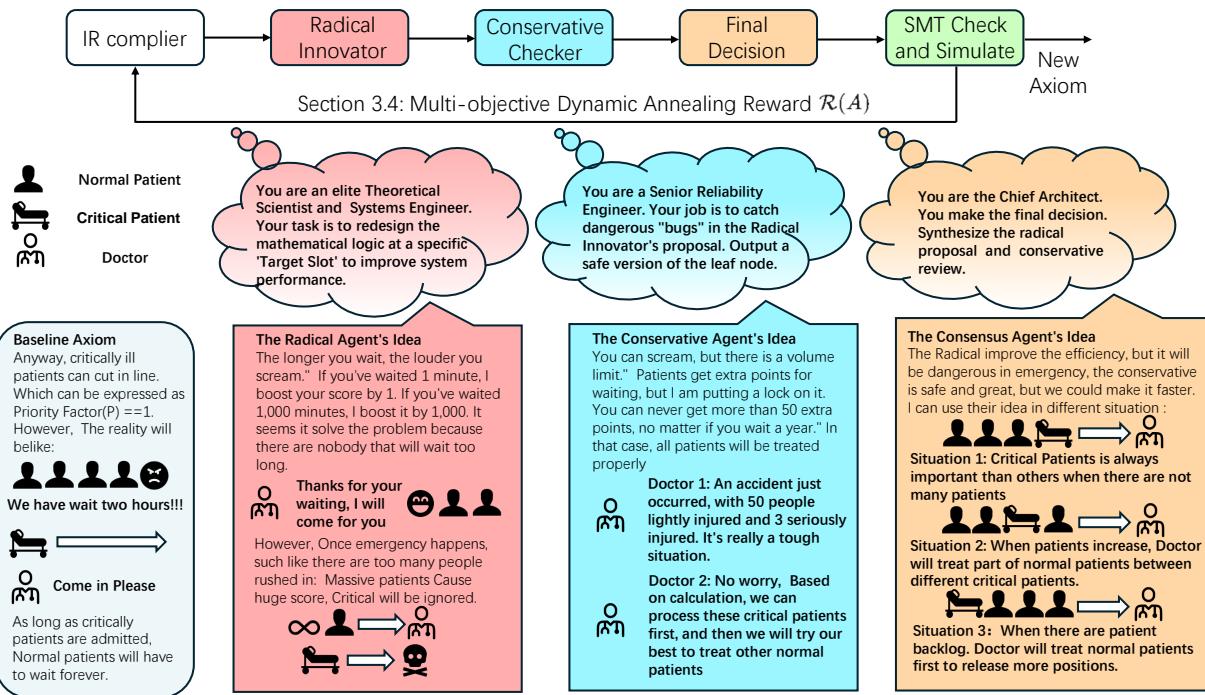


Figure 1. **Top:** Pipeline: IR compilation, tripartite proposer, SMT verification, simulation and selection over rounds. **Bottom:** Using a service center scenario as a running example, we illustrate how AutoAxiom transforms naive heuristics into robust symbolic axioms.

heterogeneous simulation backends.

For any axiom set  $P \in \mathbb{A}$  expressed in  $\mathcal{D}$ , its corresponding Core IR is defined as a recursive Directed Acyclic Graph (DAG)  $\mathcal{G} = \langle V, E, \mathcal{K} \rangle$ , where:  $V \subset \mathcal{N} \cup \mathcal{O}$  represents the set of computational nodes.  $E \in V \times V$  denotes the directed data-flow dependencies.  $\mathcal{K}$  is the set of observable metrics derived from the execution trace.

The transition from the human-centric DSL to the machine-executable IR is governed by a lowering function  $\Gamma : P^{\mathcal{D}} \rightarrow \mathcal{G}$ . This decoupling facilitates autonomous discovery: the LLM Proposer performs structural mutations on  $\mathcal{D}$ , while the Formal Verifier performs symbolic execution on  $\mathcal{G}$  to evaluate properties such as global stability.

To translate the abstract topology into empirical results, we define an execution operator  $\Psi$  that evaluates  $\mathcal{G}$  over a simulation horizon  $T$ . Let  $\sigma(V)$  be a valid evaluation sequence satisfying the partial order defined by  $E$ . The resulting performance metric  $\mathbf{y} \in \mathcal{K}$  is formalized as:

$$\mathbf{y} = \Psi(\mathcal{G}, \mathbf{S}_0, \Omega) \triangleq \bigoplus_{t \in T} f(\sigma(V), \mathbf{S}_t, \omega_t) \quad (1)$$

where  $\mathbf{S}_0$  is the initial state,  $\Omega = \{\omega_t\}$  denotes stochastic noise, and  $\bigoplus$  represents temporal aggregation.

### 3.2. Constrained Evolution via Dialectical Consensus ( $f_{\text{mod}}$ )

We formalize the axiomatic modification  $f_{\text{mod}}$  as a structured evolutionary process within the symbolic manifold  $\mathbb{M}$  defined by the Production Grammar  $\Sigma_{\text{prod}}$  and the Domain Vocabulary Map  $\mathcal{D}_{DVM}$ . The evolution of an axiom  $A \in \mathbb{M}$  is driven by a Tripartite Consensus Mechanism that resolves the conflict between exploratory innovation and physical consistency.

The search space is strictly bounded by two structural invariants: Syntactic Feasibility ( $\Sigma_{\text{prod}}$ ): Mutations are restricted to recursive tree-rewriting operations  $A \xrightarrow{\Sigma} A'$ . and Ontological Grounding ( $\mathcal{D}_{DVM}$ ): A projection  $\Pi_{\mathcal{D}} : A \rightarrow \mathbb{M}$ . To ensure type-preservation and role-consistency and the integrity of the DAG structure.

The optimal proposal  $A^*$  emerges from the equilibrium of three competing LLM agents: Radical agents ( $\mathbf{F}_{\text{rad}}$ ), Conservative Guardian ( $\mathbf{F}_{\text{con}}$ ), and Global Synthesis ( $\mathbf{F}_{\text{sys}}$ ).

The Radical Innovator executes a high-entropy jump in the symbolic space. We quantify this jump using a **Structural Divergence Metric** on the IR representation. Specifically,  $\text{dist}(\bar{A}, A_{\text{hist}})$  proxies the magnitude of symbolic transformation (e.g., structural

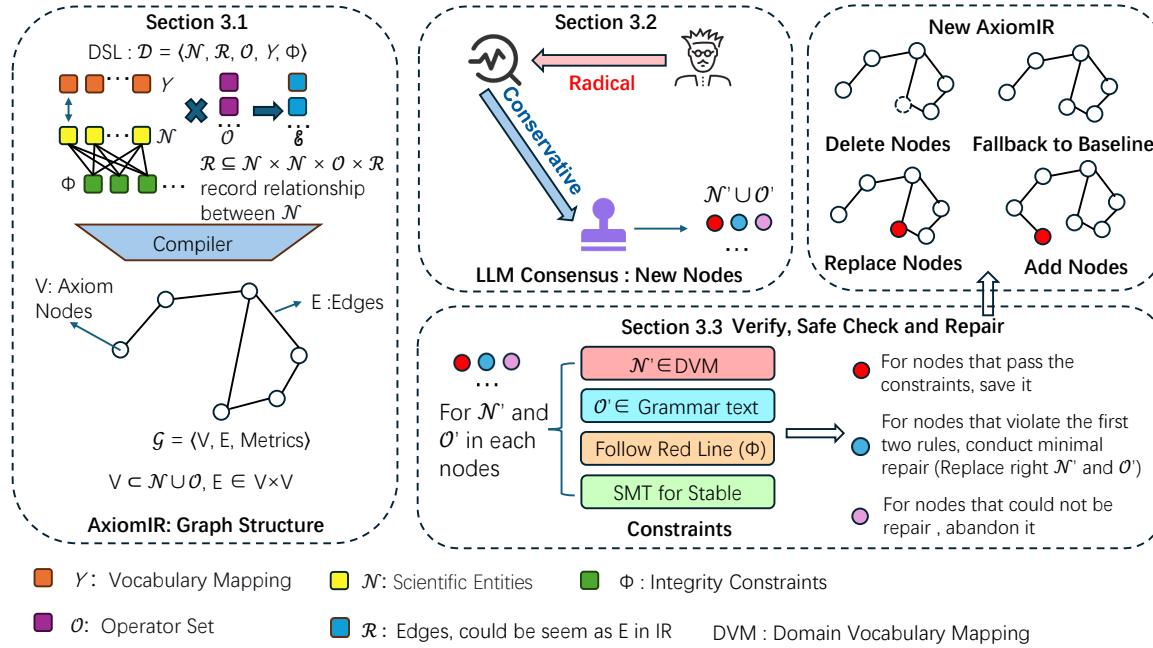


Figure 2. Workflow of the symbolic graph editor The engine transpiles the domain DSL into a granular IR graph structure. LLM agents perform targeted node replacement and structural modification within the graph, which are subsequently passed through formal node verification gates to synthesize the final, validated IR.

rewriting or complexity shift) required to distinguish the proposal from historical baselines. By incorporating a diversity incentive where  $\text{dist}(\tilde{A}, A_{hist}) \propto \mathcal{P}_{div}$ , the system actively penalizes stagnation and forces the search to explore heterogeneous regions of the solution manifold, effectively escaping local optima.

Conservative Guardian acts as a logical manifold projection. It maps the high-entropy proposal  $\tilde{A}$  back onto the feasible physical set  $\mathbb{M}_\Phi$  defined by the DVM and "Red Line" constraints. The operator  $\arg \min$  in  $\mathbf{F}_{con}$  is implemented as a Heuristic Manifold Projection. Rather than continuous gradient descent, the Conservative Guardian performs Targeted Sub-graph Mutation: it receives the SMT counter-example  $\xi$  and utilizes JSON-pointers to identify and re-synthesize only the specific IR branches responsible for the violation, thereby minimizing the structural deviation  $\delta$  from the original proposal.

In the end, the Consensus Maker performs a multi-objective judgment. It evaluates the trade-off between the radical performance gain of  $\tilde{A}$  and the conservative stability of  $\tilde{A}$ , outputting  $A^*$  as the globally optimal axiomatic patch. Crucially,  $\mathbf{F}_{sys}$  incorporates an **Occam's Razor bias**, penalizing excessive Kolmogorov complexity to ensure the discovered laws remain scientifically interpretable. Formalized as follows:

$$\begin{cases} \mathbf{F}_{rad} : A \rightarrow \tilde{A} & \text{s.t. } \text{dist}(\tilde{A}, A_{hist}) > \eta, \nabla \mathcal{S}(\tilde{A}) > 0 \\ \mathbf{F}_{con} : \tilde{A} \rightarrow \bar{A} & \text{s.t. } \bar{A} = \arg \min_{a \in \mathbb{M}_\Phi} \|a - \tilde{A}\| \\ \mathbf{F}_{sys} : \{\tilde{A}, \bar{A}\} \rightarrow A^* & \text{s.t. } A^* = \text{Pareto}(\mathcal{S}, \mathbb{V}, \mathcal{P}_{complex}) \end{cases}$$

### 3.3. Formal Boundary and Logical Confluence ( $\mathbb{V}$ )

To make evolved axioms are not only syntactically consistent but also physically stable, we implement a multi-tier verification predicate  $\mathbb{V} : \mathcal{G} \rightarrow \{0, 1\}$ . This stage acts as a "Logical Filter." To balance reasoning efficiency and formal rigor, the verification process follows a dual-stage filtering strategy. While the Conservative Guardian ( $\mathbf{F}_{con}$ ) acts as a pre-emptive heuristic filter, the SMT Gatekeeper serves as the final deterministic audit to eliminate latent instabilities.

#### 3.3.1. TIER-I: STATIC ONTOLOGICAL INVARIANTS

The first layer performs static analysis on the graph  $\mathcal{G}$  against the  $\mathcal{D}_{DVM}$  to identify explicit boundary violations. We define the static predicate  $\mathbb{V}_{stat}$  as:

$$\mathbb{V}_{stat}(\mathcal{G}) = \bigwedge_{v \in V} \mathcal{I}(v) \wedge \deg(v) \in \mathcal{O}_{Gram} \quad (2)$$

where  $\mathcal{I}(v)$  represents the Immutability Check.

220 3.3.2. TIER-II: DEEP STABILITY VIA SMT PROJECTION  
 221 Latent instabilities are identified by lowering  $\mathcal{G}$  into a First-  
 222 Order Logic (FOL) formula  $\mathbb{F}$  via a recursive symbolic  
 223 mapping  $\mathcal{T} : \text{Node} \rightarrow \text{Z3Expr}$ . Each node in the Core IR  
 224 is mapped to its equivalent SMT primitive (e.g.,  $v_{add} \rightarrow$   
 225  $+$ ,  $\text{piecewise} \rightarrow \text{ITE}$  logic), allowing the SMT solver to  
 226 perform deterministic SAT/UNSAT proofs over the entire  
 227 parameter manifold defined in  $\mathcal{C}$ .  
 228

229 An axiom  $\mathcal{G}$  is **Stable** iff the solver proves the unsatisfiability  
 230 of the violation state  $\neg\Phi$  under constraints  $\mathcal{C}$ :

$$232 \mathbb{V}_{smt}(\mathcal{G}) = 1 \iff \text{Solve}(\mathcal{C} \wedge \mathbb{F}(\mathcal{G}) \wedge \neg\Phi) \rightarrow \text{UNSAT} \quad (3)$$

233 If the result is SAT, the solver yields a **Counter-example**  
 234  $\xi = \{\mathbf{s} \mid \mathbb{F}(\mathcal{G}, \mathbf{s}) \vdash \neg\Phi\}$ , identifying the exact parameter  
 235 regime where the proposed logic collapses.  
 236

237 3.3.3. LOGICAL CONFLUENCE AND SAFE CUT REPAIR  
 238  $(\mathcal{R})$

239 Upon verification failure ( $\mathbb{V} = 0$ ), the system invokes a repair operator  $\mathcal{R}$  to project illegal logic back onto the nearest  
 240 stable manifold:

$$244 \mathcal{G}^* = \mathcal{R}(\mathcal{G}, \xi) \triangleq \arg \min_{\mathcal{G}' \in \mathbb{M}_\Phi} \delta(\mathcal{G}', \mathcal{G}) \quad (4)$$

245 By utilizing  $\xi$  as a "Logical Gradient," the LLM performs  
 246 a targeted prune on the unstable branches of  $\mathcal{G}$ , forcing the  
 247 search to converge within the safety envelope.  
 248

249 3.4. Adaptive Synthesis and Dynamic Reward  
 250 Annealing

251 The AutoAxiom framework operates as an iterative closed-  
 252 loop system driven by a Dynamic Objective Function  $J(A)$ .

253 3.4.1. MULTI-OBJECTIVE OBJECTIVE FORMULATION

254 The objective  $J(A)$  is defined as a weighted sum of four  
 255 competing metrics:

$$256 J(A) = \mathcal{S}_{perf}(A) - \lambda_v(t)\mathcal{P}_{viol} + \lambda_p(t)\mathcal{P}_{pers} + \lambda_d(t)\mathcal{P}_{div} \quad (5)$$

257 where  $\mathcal{S}_{perf}$ : raw performance score.  $\mathcal{P}_{viol} \in \{0, 1\}$  is a  
 258 binary penalty triggered by verification failure.  $\mathcal{P}_{pers}$  is  
 259 calculated as the negative variance of performance scores  
 260 across  $N = 100$  Monte Carlo seeds to penalize brittle  
 261 discoveries ( $\mathcal{P}_{pers} = -\text{Var}(\{\mathbf{y}_i\}_{i=1}^{100})$ ).  $\mathcal{P}_{div}$  represents  
 262 structural novelty, quantified as the symbolic divergence.

263 3.4.2. DYNAMIC ANNEALING AND PHASE SCHEDULING

264 To ensure convergence, the coefficients  $\lambda(t)$   
 265 follow a non-linear annealing trajectory: Constraint Hardening  
 266 ( $t < T_{early}$ ):  $\lambda_v$  is maximal  
 267

268 to force the LLM to learn the *feasible manifold*  
 269  $\mathbb{M}_\Phi$ . Global Exploration ( $T_{early} \leq t < T_{mid}$ ):  $\lambda_d$  follows an exponential decay to encourage structural  
 270 mutations. Robust Refinement ( $t \geq T_{mid}$ ):  $\lambda_p$  increases  
 271 linearly, shifting pressure toward *Persistence*.  
 272

273  $J(A)$  is injected back into the Tripartite Proposer via a  
 274 Contextual History Buffer. By providing the LLM with the  
 275 triple  $\{A_t, J_t, \xi_t\}$ , where  $\xi$  acts as a Negative Prompt, the  
 276 system transforms the search from a stochastic walk into a  
 277 Directed Gradient-free Optimization that explicitly avoids  
 278 known failure modes. Appendix B9 provides pseudocode  
 279 for the complete closed loop procedure.

## 4. Experiment

### 4.1. Versatility: Cross-Domain Axiomatic Discovery

280 To empirically validate the versatility of our dual-layered  
 281 representation, we benchmark AutoAxiom across six het-  
 282 erogeneous domains ranging from random queue system  
 283 (Krenzler, 2016) to high-fidelity physical systems (Zhang  
 284 et al., 2013). By initializing the evolutionary loop from prim-  
 285 itive baseline axioms (G1), the hyperparameters is shown in  
 286 appendix B0.4(e.g. GPT-4o,  $T = 0.9$ ), the  $\mathcal{A}^2\mathcal{V}\mathcal{S}$  closed-  
 287 loop system navigates the symbolic manifold defined by Ax-  
 288 iomDSL and Core IR to challenge and surpass established  
 289 industrial state-of-the-art strategies. This cross-disciplinary  
 290 stress test confirms that the hierarchical decoupling of on-  
 291 tological semantics (Nirenburg & Raskin, 2004) and com-  
 292 putational topology (Edelsbrunner & Harer, 2010) allows  
 293 for the discovery of effective control laws that remain in-  
 294 accessible to fixed-rule parametric optimization (Shiraishi  
 295 et al., 2025). As shown in Table 1, AutoAxiom continu-

Table 1. **Cross-Domain Evaluation of Scientific Axiom Discovery.** SRA (Appendix B0.1) is normalized such that G1 is anchored near 0.5 and the best evaluated baseline is anchored near 1.0; higher is better. AutoAxiom (Ours) consistently breaks the Pareto frontier ( $> 1.0$ ) in five tasks. Baselines (G2/G3) are: **Queueing Network**: Static Priority/Lyapunov Feedback; **Service Center**: EDF Pre-emption/Agile Linear; **Software Opt**: Static Heuristics/Adaptive Thresholding; **Physical PDE**: Standard PID/Conservative PD; **Res. Allocation**: JSQ/Power-of-Two; **Composition**: PID SOTA/Stateless Coupling. Data: Mean  $\pm$  95% CI over 100 runs.

Domain	Baseline (G1)	SOTA-1 (G2)	SOTA-2 (G3)	AutoAxiom (Ours)
Queueing Network	$0.50 \pm 0.01$	$0.51 \pm 0.01$	$1.00 \pm 0.04$	<b><math>1.74 \pm 0.06</math></b>
Service Center	$0.51 \pm 0.00$	$1.00 \pm 0.01$	$0.94 \pm 0.03$	<b><math>1.02 \pm 0.01</math></b>
Software Opt.	$0.50 \pm 0.00$	$0.94 \pm 0.01$	$1.00 \pm 0.01$	<b><math>1.02 \pm 0.01</math></b>
Physical PDE	$0.50 \pm 0.00$	$1.00 \pm 0.02$	$0.59 \pm 0.00$	<b><math>3.15 \pm 0.21</math></b>
Res. Allocation	$0.53 \pm 0.00$	$1.00 \pm 0.02$	$0.58 \pm 0.00$	<b><math>1.29 \pm 0.05</math></b>
Composition	$0.52 \pm 0.00$	<b><math>1.00 \pm 0.01</math></b>	$0.60 \pm 0.00$	$0.90 \pm 0.00$

296 ously improves efficiency in most environments by syn-  
 297 thesizing high-dimensional symbolic control laws (Yan et al.,  
 298 2022). As shown in Appendix B1, In queueing network,  
 299 compared to Lyapunov-based linear feedback state-of-the-

art (SOTA)(Liang et al., 2025), AutoAxiom discovers a threshold-aware defense logic(Erbagci et al., 2016) that reduces latency by 62.8%. Notably, in the service-center (Appendix B2) and software optimization(Appendix B3), extreme baseline collapse triggers a denominator dilution effect(Price & Matthews, 2009), limiting the SRA score to 1.02; however, our framework significantly outperforms industrial preemption(Otamendi et al., 2025) and agile linear(Gao et al., 2006) SOTA in practical physical metrics. By introducing a tanh-gated(Chai et al., 2020) priority factor and a logarithmic S-shaped "reconfiguration gate,"(Sidhu et al., 2000) AutoAxiom stabilizes secondary queues at 2.61 seconds and compresses technical debt to a complexity of 0.17 (with an error rate as low as 0.09), achieving a superior multi-objective balance lacking in traditional heuristic algorithms(Kokash, 2005). In Appendix B4 the physical partial differential equation (PDE) domain , compared to conservative PD controllers(Lee & Lee, 2025), the evolved axiom, employing a saturation-based scaling method(Lu et al., 2020), operates safely near the theoretical CFL limit, reducing  $L^2$  error by 76.4% and achieving an excellent SRA score of 3.15. In resource allocation (Appendix B5), AutoAxiom surpasses standard JSQ logic(Mukhopadhyay & Mazumdar, 2015) by using boundary-aware scaling to map nonlinear cost-utility envelopes, achieving a 133% throughput improvement and an SRA score of 1.29. Finally, in network physical synthesis (Appendix B6), AutoAxiom synthesizes a stateless coupling law(Bevir & Rhodes, 2011) that achieves 90% of the performance of a finely tuned PID(Zhao et al., 2025) axioms. Its main advantage lies in eliminating the computational overhead of integrating historical errors, thereby achieving low-latency hardware execution while maintaining industrial-grade stability. This cross-domain results marks a fundamental shift in design philosophy from human-centered linear design to high-dimensional nonlinear symbolic search spaces. The DSL/IR architecture can autonomously discover patterns that combine numerical accuracy and symbolic logic.

## 4.2. Interpretability: White-Box Control in Soft Robotics

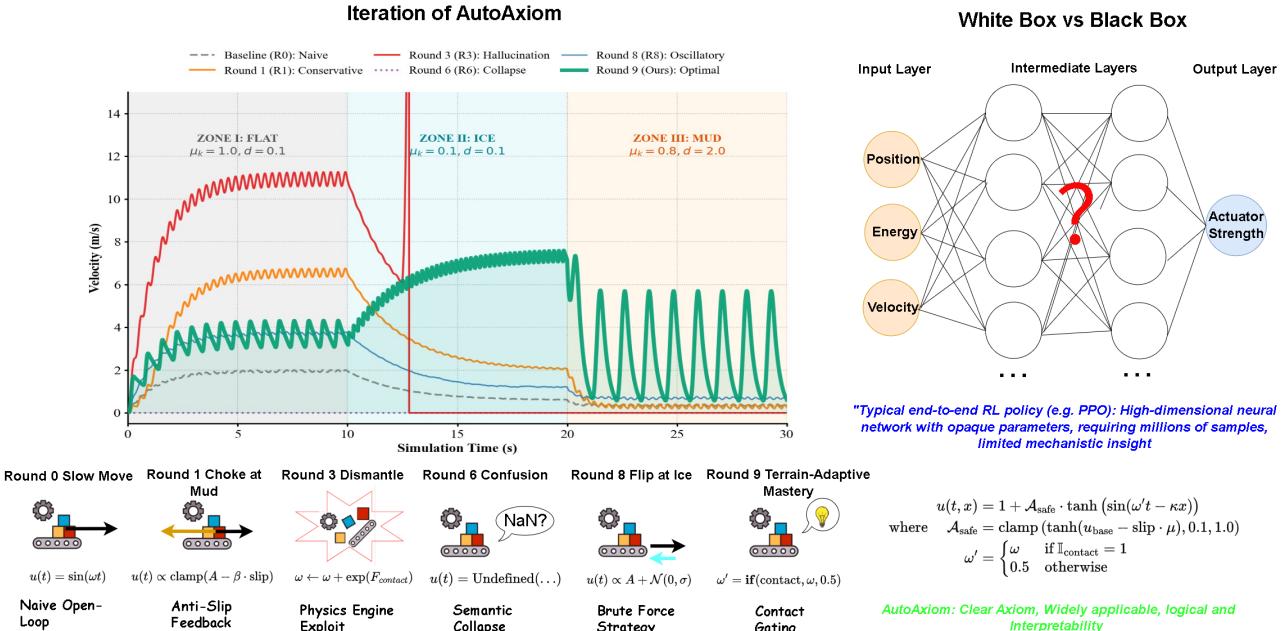
Drawing inspiration from voxel-based soft robotics of Evgym(Bhatia et al., 2021), we design an Ice-Mud-Flat terrain simulation to evaluate the evolutionary trajectory of AutoAxiom in generating interpretable control laws.

As shown in figure 3 and Appendix B7.1, the system begins at Round 0 with a naive traveling wave, subsequently navigating a non-monotonic progress that validates the tripartite consensus mechanism of Equation 3. While Round 1 introduced slip-feedback for ice traversal, the logic misidentified viscous drag as traction loss, illustrating the "conservative trap" where the robot stalled in mud. Following this, Round 3 exhibited a "numerical hallucination" where the radical

innovator exploited simulation artifacts to achieve high velocities, though these ultra-high-frequency pulses were identified as physically infeasible by the verifier. After surviving a semantic collapse in Round 6 and oscillatory behavior in Round 8, the framework synthesized the optimal Round 9 axiom. By discovering "contact gating" a logic that idles actuators when contact is lost, the symbolic policy achieved approximately 95% of the performance of a high-parameter PPO model in Appendix B7.3 while offering superior analytical boundaries. On low-friction ice, AutoAxiom's discrete logic outperformed PPO (0.67s vs 1.10s), likely due to the exact constraints of the clamp operator which neural networks can only approximate through continuous functions. This transition from opaque weights to concise physical expressions provides a potential path for deploying interpretable control on resource-constrained embedded systems without requiring deep inference overhead.

## 4.3. Constraints & Intelligence: Constrained Discovery in Chemical Space

Inspired by the *Foundation Molecular Grammar* (FMG)(Sun et al., 2025), which demonstrated the capacity of generative foundation models to induce interpretable molecular graph languages for automated discovery, we evaluated AutoAxiom within a fragment-based drug design (FBDD) paradigm. The experimental dataset utilizes a carefully selected library of fragments, including key pharmacophores such as cyclohexane and pyridine, as well as some potentially toxic precursors with drug-like properties but low safety profiles, such as nitro and sulfonamide groups. Furthermore, heavy rigid linking groups were added to test the system's ability to maintain molecular weight constraints. We conducted a comparative evaluation of the system using conventional symbolic regression (PySR) as a representative of a general data fitting paradigm. This comparison aims to demonstrate that while symbolic regression performs well in curve fitting, it operates in a "semantic vacuum" lacking chemical priors or formal constraints, and therefore is prone to producing ineffective drugs in the event of a chemical space combinatorial explosion. As shown in Figure 4 and Appendix B8.6, the well trained PySR baseline model(Appendix B8.2) exhibits significant limitations due to the lack of inherent chemical prior knowledge and dynamic simulation incentive mechanisms. In our generative experiments, PySR struggles to generalize beyond static datasets; high-frequency numerical noise in the reward region hinders its active exploration, ultimately leading to overfitting and paralysis. Conversely, in the early exploration phase, AutoAxiom's initial axioms, designed to maximize numerical scores but lacking a physical basis, resulted in the generation of unbounded "obese molecules" a typical manifestation of the abuse of reward mechanisms.



**Figure 3. Case Study I: White-Box Control in Soft Robotics. (Left)** The phylogeny of the control axiom. The system navigates through “Conservative Traps” (Round 1) and “Numerical Hallucinations” (Round 3) before discovering the optimal Phase-Resetting Logic in Round 9. **(Right)** A visual comparison of the policy architecture. Unlike the opaque neural weights of PPO (Black Box), AutoAxiom evolves an interpretable, concise symbolic formula (White Box). Ours outperforms PPO on Ice while PPO is faster on Flat and Mud; focus is interpretability and constraints

This “reward hacking” phase is critical, as it empirically confirms that numerical optimization alone is insufficient for scientific discovery. It highlights the necessity of the subsequent transition to constrained reasoning gatekeeper.

Through SMT formal verification in Appendix B8.3, AutoAxiom successfully internalized the Lipinski's Rule of Five provided as integrity constraints ( $\Phi$ ) into executable symbolic gating logic. By utilizing the logical failure mechanism ( $\xi$ ) as a "negative gradient" in the symbolic space, the system learned to bridge the gap between abstract chemical principles and concrete fragment-selection policies. This demonstrates that AutoAxiom does not merely "follow" rules but synthesizes a compliant manifold projection that filters out toxic or oversized pharmacophores before they reach the simulation stage. The resulting molecules exhibit complex pharmacophores, high drug-likeness (QED), and realistic chemical structures. From this case, we demonstrate that the traditional Symbolic Regression (SR) paradigm collapses into 'Stagnation' within high-dimensional chemical spaces due to its lack of semantic grounding and subsequent overfitting of reward noise. In contrast, AutoAxiom resolves this by anchoring foundational LLM priors through formal logical filters, successfully transitioning from stochastic data-fitting to an intelligent discovery method that balances exploratory innovation with safety-critical constraints."

#### 4.4. Ablation and Cost Report

To evaluate the structural necessity of AutoAxiom, we conducted a systematic ablation study by isolating core components, as summarized in Table 2. The results reveal that removing the Tripartite Consensus mechanism (Sec 3.2) leads to a significant drop in mean performance (0.6917) and a sharp increase in variance ( $\pm 0.3109$ ). This instability stems from the loss of agent isolation and dialectical conflict defined in Equation 3; without the counterbalancing roles of  $F_{rad}$  and  $F_{con}$ , a single-role prompt tends to either converge prematurely to conservative local optima or propose radical, unstable axioms, resulting in erratic performance. A similar trend is observed in the No Iteration variant, which exhibits the highest variance ( $\pm 0.3876$ ) and a 50% success rate. Deprived of historical lessons and the dynamic annealing incentives, the search process lacks the necessary corrective guidance to escape known failure modes. Furthermore, the No Verifier variant (*w/o* Sec 3.3) exhibits the most critical impact on system safety, with the violation rate surging to 12.20% and the success rate plummeting to 66%. Forensic analysis confirms that without the  $V_{smt}$  predicate in Equation 5, the system frequently proposes axioms that trigger simulator crashes or non-physical divergence, as there is no formal mechanism to intercept "Red Line" violations. Collectively, these findings empirically validate

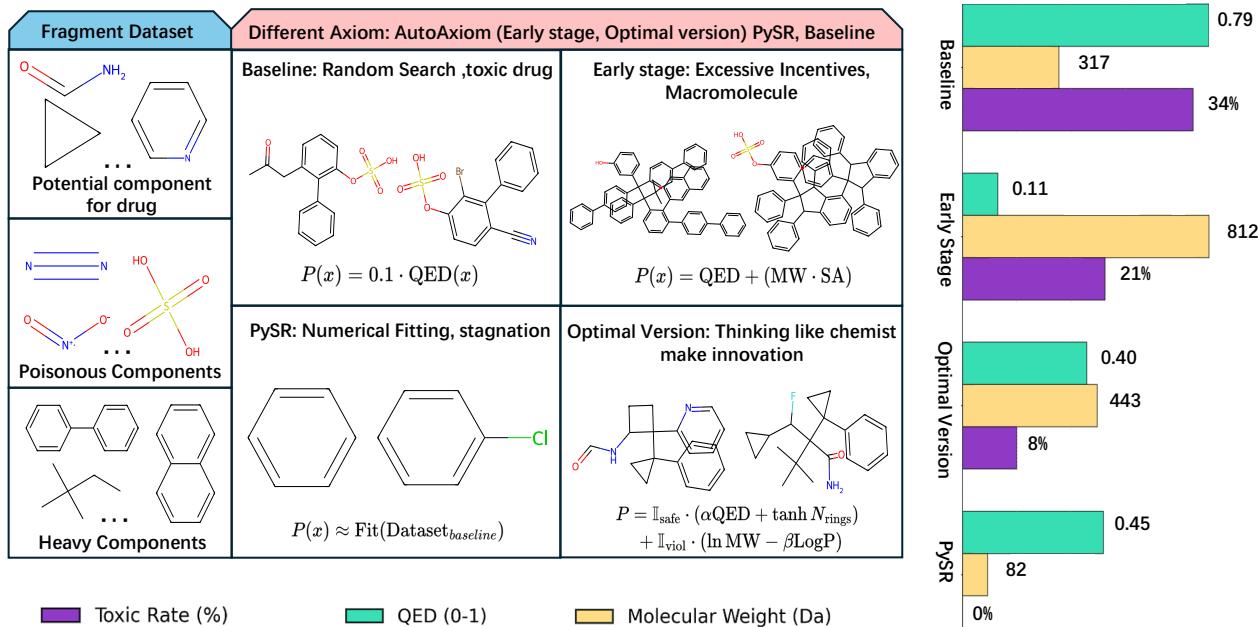
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409

Figure 4. **Case Study II: Molecular Discovery** SMT-based formal gatekeeper filtering “obese molecules” and enforcing Lipinski’s Rule of Five. AutoAxiom ensures chemical validity and drug-likeness by tethering LLM priors to rigorous integrity constraints. PySR achieves 0% toxicity in Table 17 but shows low Steps Taken under the sequential protocol; AutoAxiom balances progress and constraints

414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647  
648  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701  
702  
703  
704  
705  
706  
707  
708  
709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809  
8010  
8011  
8012  
8013  
8014  
8015  
8016  
8017  
8018  
8019  
8020  
8021  
8022  
8023  
8024  
8025  
8026  
8027  
8028  
8029  
8030  
8031  
8032  
8033  
8034  
8035  
8036  
8037  
8038  
8039  
8040  
8041  
8042  
8043  
8044  
8045  
8046  
8047  
8048  
8049  
8050  
8051  
8052  
8053  
8054  
8055  
8056  
8057  
8058  
8059  
8060  
8061  
8062  
8063  
8064  
8065  
8066  
8067  
8068  
8069  
8070  
8071  
8072  
8073  
8074  
8075  
8076  
8077  
8078  
8079  
8080  
8081  
8082  
8083  
8084  
8085  
8086  
8087  
8088  
8089  
8090  
8091  
8092  
8093  
8094  
8095  
8096  
8097  
8098  
8099  
80100  
80101  
80102  
80103  
80104  
80105  
80106  
80107  
80108  
80109  
80110  
80111  
80112  
80113  
80114  
80115  
80116  
80117  
80118  
80119  
80120  
80121  
80122  
80123  
80124  
80125  
80126  
80127  
80128  
80129  
80130  
80131  
80132  
80133  
80134  
80135  
80136  
80137  
80138  
80139  
80140  
80141  
80142  
80143  
80144  
80145  
80146  
80147  
80148  
80149  
80150  
80151  
80152  
80153  
80154  
80155  
80156  
80157  
80158  
80159  
80160  
80161  
80162  
80163  
80164  
80165  
80166  
80167  
80168  
80169  
80170  
80171  
80172  
80173  
80174  
80175  
80176  
80177  
80178  
80179  
80180  
80181  
80182  
80183  
80184  
80185  
80186  
80187  
80188  
80189  
80190  
80191  
80192  
80193  
80194  
80195  
80196  
80197  
80198  
80199  
80200  
80201  
80202  
80203  
80204  
80205  
80206  
80207  
80208  
80209  
80210  
80211  
80212  
80213  
80214  
80215  
80216  
80217  
80218  
80219  
80220  
80221  
80222  
80223  
80224  
80225  
80226  
80227  
80228  
80229  
80230  
80231  
80232  
80233  
80234  
80235  
80236  
80237  
80238  
80239  
80240  
80241  
80242  
80243  
80244  
80245  
80246  
80247  
80248  
80249  
80250  
80251  
80252  
80253  
80254  
80255  
80256  
80257  
80258  
80259  
80260  
80261  
80262  
80263  
80264  
80265  
80266  
80267  
80268  
80269  
80270  
80271  
80272  
80273  
80274  
80275  
80276  
80277  
80278  
80279  
80280  
80281  
80282  
80283  
80284  
80285  
80286  
80287  
80288  
80289  
80290  
80291  
80292  
80293  
80294  
80295  
80296  
80297  
80298  
80299  
80300  
80301  
80302  
80303  
80304  
80305  
80306  
80307  
80308  
80309  
80310  
80311  
80312  
80313  
80314  
80315  
80316  
80317  
80318  
80319  
80320  
80321  
80322  
80323  
80324  
80325  
80326  
80327  
80328  
80329  
80330  
80331  
80332  
80333  
80334  
80335  
80336  
80337  
80338  
80339  
80340  
80341  
80342  
80343  
80344  
80345  
80346  
80347  
80348  
80349  
80350  
80351  
80352  
80353  
80354  
80355  
80356  
80357  
80358  
80359  
80360  
80361  
80362  
80363  
80364  
80365  
80366  
80367  
80368  
80369  
80370  
80371  
80372  
80373  
80374  
80375  
80376  
80377  
80378  
80379  
80380  
80381  
80382  
80383  
80384  
80385  
80386  
80387  
80388  
80389  
80390  
80391  
80392  
80393  
80394  
80395  
80396  
80397  
80398  
80399  
80400  
80401  
80402  
80403  
80404  
80405  
80406  
80407  
80408  
80409  
80410  
80411  
80412  
80413  
80414  
80415  
80416  
80417  
80418  
80419  
80420  
80421  
80422  
80423  
80424  
80425  
80426  
80427  
80428  
80429  
80430  
80431  
80432  
80433  
80434  
80435  
80436  
80437  
80438  
80439  
80440  
80441  
80442  
80443  
80444  
80445  
80446  
80447  
80448  
80449  
80450  
80451  
80452  
80453  
80454  
80455  
80456  
80457  
80458  
80459  
80460  
80461  
80462  
80463  
80464  
80465  
80466  
80467  
80468  
80469  
80470  
80471  
80472  
80473  
80474  
80475  
80476  
80477  
80478  
80479  
80480  
80481  
80482  
80483  
80484  
80485  
80486  
80487  
80488  
80489  
80490  
80491  
80492  
80493  
80494  
80495  
80496  
80497  
80498  
80499  
80500  
80501  
80502  
80503  
80504  
80505  
80506  
80507  
80508  
80509  
80510  
80511  
80512  
80513  
80514  
80515  
80516  
80517  
80518  
80519  
80520  
80521  
80522  
80523  
80524  
80525  
80526  
80527  
80528  
80529  
80530  
80531  
80532  
80533  
80534  
80535  
80536  
80537  
80538  
80539  
80540  
80541  
80542  
80543  
80544  
80545  
80546  
80547  
80548  
80549  
80550  
80551  
80552  
80553  
80554  
80555  
80556  
80557  
80558  
80559  
80560  
80561  
80562  
80563  
80564  
80565  
80566  
80567  
80568  
80569  
80570  
80571  
80572  
80573  
80574  
80575  
80576  
80577  
80578  
80579  
80580  
80581  
80582  
80583  
80584  
80585  
80586  
80587  
80588  
80589  
80590  
80591  
80592  
80593  
80594  
80595  
80596  
80597  
80598  
80599  
80600  
80601  
80602  
80603  
80604  
80605  
80606  
80607  
80608  
80609  
80610  
80611  
80612  
80613  
80614  
80615  
80616  
80617  
80618  
80619  
80620  
80621  
80622  
80623  
80624  
80625  
80626  
80627  
80628  
80629  
80630  
80631  
80632  
80633  
80634  
80635  
80636  
80637  
80638  
80639  
80640  
80641  
80642  
80643  
80644  
80645  
80646  
80647  
80648  
80649  
80650  
80651  
80652  
80653  
80654  
80655  
80656  
80657  
80658  
80659  
80660  
80661  
80662  
80663  
80664  
80665  
80666  
80667  
80668  
80669  
80670  
80671  
80672  
80673  
80674  
80675  
80676  
80677  
80678  
80679  
80680  
80681  
80682  
80683  
80684  
80685  
80686  
80687  
80688  
80689  
80690  
80691  
80692  
80693  
80694  
80695  
80696  
80697  
80698  
80699  
80700  
80701  
80702  
80703  
80704  
80705  
80706  
80707  
80708  
80709  
80710  
80711  
80712  
80713  
80714  
80715  
80716  
80717  
80718  
80719  
80720  
80721  
80722  
80723  
80724  
80725  
80726  
80727  
80728  
80729  
80730  
80731  
80732  
80733  
80734  
80735  
80736  
80737  
80738  
80739  
80740  
80741  
80742  
80743  
80744  
80745  
80746  
80747  
80748  
80749  
80750  
80751  
80752  
80753  
80754  
80755  
80756  
80757  
80758  
80759  
80760  
80761  
80762  
80763  
80764  
80765  
80766  
80767  
80768  
80769  
80770  
80771  
80772  
80773  
80774  
80775  
80776  
80777  
80778  
80779  
80780  
80781  
80782  
80783  
80784  
80785  
80786  
80787  
80788  
80789  
80790  
80791  
80792  
80793  
80794  
80795  
80796  
80797  
80798  
80799  
80800  
80801  
80802  
80803  
80804  
80805  
80806  
80807  
80808  
80809  
80810  
80811  
80812  
80813  
80814  
80815  
80816  
80817  
80818  
80819  
80820  
80821  
80822  
80823  
80824  
80825  
80826  
80827  
80828  
80829  
80830  
80831  
80832  
80833  
80834  
80835  
80836  
80837  
80838  
80839  
80840  
80841  
80842  
80843  
80844  
80845  
80846  
80847  
80848  
80849  
80850  
80851  
80852  
80853  
80854  
80855  
80856  
80857  
80858  
80859  
80860  
80861  
80862  
80863  
80864  
80865  
80866  
80867  
80868  
80869  
80870  
80871  
80872  
80873  
80874  
80875  
80876  
80877  
80878  
80879  
80880  
80881  
80882  
80883  
80884  
80885  
80886  
80887  
80888  
80889  
80890  
80891  
80892  
80893  
80894  
80895  
80896  
80897  
80898  
80899  
80900  
80901  
80902  
80903  
80904  
80905  
80906  
80907  
80908  
80909  
80910  
80911  
80912  
80913  
80914  
80915  
80916  
80917  
80918  
80919  
80920  
80921  
80922  
80923  
80924  
80925  
80926  
80927  
80928  
80929  
80930  
80931  
80932  
80933  
80934  
80935  
80936  
80937  
80938  
80939  
80940  
80941  
80942  
80943  
80944  
80945  
80946  
80947  
80948  
80949  
80950  
80951  
80952  
80953  
80954  
80955  
80956  
80957  
80958  
80959  
80960  
80961  
80962  
80963  
80964  
80965  
80966  
80967  
80968  
80969  
80970  
80971  
80972  
80973  
80974  
80975  
80976  
80977  
80978  
80979  
80980  
80981  
80982  
80983  
80984  
80985  
80986  
80987  
80988  
80989  
80990  
80991  
80992  
80993  
80994  
80995  
80996  
80997  
80998  
80999  
80100  
80101  
80102  
80103  
80104  
80105  
80106  
80107  
80108  
80109  
80110  
80111  
80112  
80113  
80114  
80115  
80116  
80117  
80118  
80119  
80120  
80121  
80122  
80123  
80124  
80125  
80126  
80127  
80128  
80129  
80130  
80131  
80132  
80133  
80134  
80135  
80136  
80137  
80138  
80139  
80140  
80141  
80142  
80143  
80144  
80145  
80146  
80147  
80148  
80149  
80150  
80151  
80152  
80153  
80154  
80155  
80156  
80157  
80158  
80159  
80160  
80161  
80162  
80163  
80164  
80165  
80166  
80167  
80168  
80169  
80170  
80171  
80172  
80173  
80174  
80175  
80176  
80177  
80178  
80179  
80180  
80181  
80182  
80183  
80184  
80185  
80186  
80187  
80188  
80189  
80190  
80191  
80192  
80193  
80194  
80195  
80196  
80197  
80198  
80199  
80200  
80201  
80202  
80203  
80204  
80205  
80206  
80207  
80208  
80209  
80210  
80211  
80212  
80213  
80214  
80215  
80216  
80217  
80218  
80219  
80220  
80221  
80222  
80223  
80224  
80225  
80226  
80227  
80228  
80229  
80230  
80231  
80232  
80233  
80234  
80235  
80236  
80237  
80238  
80239  
80240  
80241  
80242  
80243  
80244  
80245  
80246  
80247  
80248  
80249  
80250  
80251  
80252  
80253  
80254  
80255  
80256  
80257  
80258  
80259  
80260  
80261  
80262  
80263  
80264  
80265  
80266  
80267  
80268  
80269  
80270  
80271  
80272  
80273  
80274  
80275  
80276  
80277  
80278  
80279  
80280  
80281  
80282  
80283  
80284  
80285  
80286  
80287  
80288  
80289  
80290  
80291  
80292  
80293  
80294  
80295  
80296  
80297  
80298  
80299  
80300  
80301  
80302  
80303  
80304  
80305  
80306  
80307  
80308  
80309  
80310  
80311  
80312  
80313  
80314  
80315  
80316  
80317  
80318  
80319  
80320  
80321  
80322  
80323  
80324  
80325  
80326  
80327  
80328  
80329  
80330  
80331  
80332  
80333  
80334  
80335  
80336  
80337  
80338  
80339  
80340  
80341  
80342  
80343  
80344  
80345  
80346  
80347  
80348  
80349  
80350  
80351  
80352  
80353  
80354  
80355  
80356  
80357  
80358  
80359  
8036

440 **Impact Statement**

441 This paper presents AutoAxiom, a framework designed to  
 442 bridge the gap between large language models and rigorous  
 443 scientific discovery through autonomous symbolic axiom  
 444 modification and formal verification. The societal and ethi-  
 445 cal implications of this work are twofold:

446 **Positive Impact on Reliable AI for Science:** By integrating  
 447 SMT-based formal verification as a "Red Line" gate-  
 448 keeper, AutoAxiom provides a blueprint for building self-  
 449 correcting AI systems that adhere to immutable physical  
 450 laws and safety constraints. This reduces the risk of "nu-  
 451 mercial hallucinations" in safety-critical domains such as  
 452 healthcare triage, infrastructure control, and drug design, en-  
 453 ensuring that AI-generated scientific insights are not only high-  
 454 performing but also mathematically auditable and physically  
 455 consistent.

456 **Ethical Considerations and Safeguards:** While the frame-  
 457 work enables accelerated discovery in sensitive areas like  
 458 molecular design, the modular design allows for the strict  
 459 imposition of ethical and "Integrity Constraints". For in-  
 460 stance, in the chemical domain, AutoAxiom successfully  
 461 internalized toxicity filters and pharmacophore constraints  
 462 into its evolutionary loop, demonstrating that automated  
 463 discovery can be programmatically tethered to safety stan-  
 464 dards. We emphasize that such frameworks should always  
 465 be deployed with human-in-the-loop validation and domain-  
 466 specific oversight to prevent the discovery of harmful sub-  
 467 stances or unstable control policies

468 **References**

469 AbdusSalam, S., Abel, S., and Romão, M. C. Symbolic re-  
 470 gression for beyond the standard model physics. *Physical*  
 471 *Review D*, 111(1):015022, 2025.

472 Aglietti, V., Ktena, I., Schrouff, J., Sgouritsa, E., Ruiz, F. J.,  
 473 Malek, A., Bellot, A., and Chiappa, S. Funbo: Discovering  
 474 acquisition functions for bayesian optimization with  
 475 funsearch. *arXiv preprint arXiv:2406.04824*, 2024.

476 Aigner, M. *Combinatorial search*. John Wiley & Sons, Inc.,  
 477 1988.

478 Alter, S. Can an llm use work system axioms when de-  
 479 scribing work systems for requirements analysis? In  
 480 *International Conference on Business Process Modeling,*  
 481 *Development and Support*, pp. 245–253. Springer, 2025.

482 Bevir, M. and Rhodes, R. A. The stateless state. *The SAGE*  
 483 *handbook of governance*, pp. 203–17, 2011.

484 Bhatia, J., Jackson, H., Tian, Y., Xu, J., and Matusik, W.  
 485 Evolution gym: A large-scale benchmark for evolving  
 486 soft robots. *Advances in Neural Information Processing*  
 487 *Systems*, 34:2201–2214, 2021.

488 Chai, Y., Jin, S., and Hou, X. Highway transformer: Self-  
 489 gating enhanced self-attentive networks. *arXiv preprint*  
 490 *arXiv:2004.08178*, 2020.

491 De Moura, L. and Bjørner, N. Z3: An efficient smt solver.  
 492 In *International conference on Tools and Algorithms for*  
 493 *the Construction and Analysis of Systems*, pp. 337–340.  
 494 Springer, 2008.

495 DeRose, K. "bamboozled by our own words": Semantic  
 496 blindness and some arguments against contextualism.  
 497 *Philosophy and Phenomenological Research*, 73(2):316–  
 498 338, 2006.

499 Dong, J. and Zhong, J. Recent advances in symbolic regres-  
 500 sion. *ACM Computing Surveys*, 57(11):1–37, 2025.

501 Edelsbrunner, H. and Harer, J. *Computational topology: an*  
 502 *introduction*. American Mathematical Soc., 2010.

503 Erbagci, B., Erbagci, C., Akkaya, N. E. C., and Mai, K. A  
 504 secure camouflaged threshold voltage defined logic fam-  
 505 ily. In *2016 IEEE International symposium on hardware*  
 506 *oriented security and trust (HOST)*, pp. 229–235. IEEE,  
 507 2016.

508 Fu, V. Ai for science: Opportunities, challenges, and future  
 509 directions. *Authorea Preprints*, 2025.

510 Gao, S., Sambell, A., and Zhong, S. Polarization-agile  
 511 antennas. *IEEE Antennas and Propagation Magazine*, 48  
 512 (3):28–37, 2006.

513 Gottweis, J., Weng, W.-H., Daryin, A., Tu, T., Palepu, A.,  
 514 Sirkovic, P., Myaskovsky, A., Weissenberger, F., Rong,  
 515 K., Tanno, R., et al. Towards an ai co-scientist. *arXiv*  
 516 *preprint arXiv:2502.18864*, 2025.

517 Gulati, A., Miranda, B., Chen, E., Xia, E., Fronsdal, K.,  
 518 de Moraes Dumont, B., and Koyejo, S. Putnam-axiom:  
 519 A functional and static benchmark. In *2nd AI for Math*  
 520 *Workshop @ ICML 2025*, 2025.

521 Holt, S., Qian, Z., and van der Schaar, M. Deep generative  
 522 symbolic regression. *arXiv preprint arXiv:2401.00282*,  
 523 2023.

524 Joel, S., Wu, J., and Fard, F. A survey on llm-based code  
 525 generation for low-resource and domain-specific program-  
 526 ming languages. *ACM Transactions on Software Engi-  
 527 neering and Methodology*, 2024.

528 Kalaivani, M., Suganya, V., Suresh, N., and Catherine, S.  
 529 The next wave in marketing: Data science in the age of  
 530 generative ai. In *Navigating Data Science*, pp. 13–26.  
 531 Emerald Publishing Limited, 2025.

495 Kim, S. Cross-axiomatic evolution: A formal model of  
 496 mutual axiom acceptance between humans and artificial  
 497 intelligence. 2025.

498 Kokash, N. An introduction to heuristic algorithms. *De-  
 499 partment of Informatics and Telecommunications*, 1:1–7,  
 500 2005.

502 Krenzler, R. *Queueing systems in a random environment*.  
 503 PhD thesis, Staats-und Universitätsbibliothek Hamburg  
 504 Carl von Ossietzky, 2016.

506 Lawal, Z. K., Yassin, H., Lai, D. T. C., and Che Idris, A.  
 507 Physics-informed neural network (pinn) evolution and  
 508 beyond: A systematic literature review and bibliometric  
 509 analysis. *Big Data and Cognitive Computing*, 6(4):140,  
 510 2022.

512 Lee, K. and Lee, C. String stability analysis and design  
 513 guidelines for pd controllers in adaptive cruise control  
 514 systems. *Sensors*, 25(11):3518, 2025.

516 Li, H., Wang, X., Zhang, Z., and Zhu, W. Out-of-distribution  
 517 generalization on graphs: A survey. *IEEE Transactions  
 518 on Pattern Analysis and Machine Intelligence*, 2025.

519 Liang, G., Xie, X., Pan, W., Gu, Y., Chen, Z., Li, Z., and  
 520 Zeng, W. Static output feedback control for linear systems  
 521 via multiple lyapunov functions. *IEEE Access*, 2025.

523 Lu, Z., Long, B., and Yang, S. Saturation based iterative  
 524 approach for single image dehazing. *IEEE Signal Pro-  
 525 cessing Letters*, 27:665–669, 2020.

526 Ma, P., Jones, B. T., Wang, T.-H., Guo, M., Lipiec,  
 527 M. P., Gan, C., and Matusik, W. Newton to einstein:  
 528 Axiom-based discovery via game design. *arXiv preprint  
 529 arXiv:2509.05448*, 2025.

531 Ma, Y. J., Liang, W., Wang, G., Huang, D.-A., Bastani, O.,  
 532 Jayaraman, D., Zhu, Y., Fan, L., and Anandkumar, A.  
 533 Eureka: Human-level reward design via coding large  
 534 language models. *arXiv preprint arXiv:2310.12931*, 2023.

536 Mukhopadhyay, A. and Mazumdar, R. R. Analysis of ran-  
 537 domized join-the-shortest-queue (jsq) schemes in large  
 538 heterogeneous processor-sharing systems. *IEEE Trans-  
 539 actions on Control of Network Systems*, 3(2):116–126,  
 540 2015.

541 Nirenburg, S. and Raskin, V. *Ontological semantics*. Mit  
 542 Press, 2004.

544 Otamendi, U., Martinez, I., Belaunzarán, X., Artetxe, A.,  
 545 Franco, J., Uribe, A., Olaizola, I. G., and Sierra, B. An  
 546 analytics-based framework for optimizing resource al-  
 547 location and preemptive scheduling in manufacturing.  
 548 *Decision Analytics Journal*, pp. 100596, 2025.

Paassen, B., Grattarola, D., Zambon, D., Alippi, C., and  
 Hammer, B. Graph edit networks. In *International Con-  
 ference on Learning Representations*, 2020.

Price, P. C. and Matthews, T. V. From group diffusion to  
 ratio bias: Effects of denominator and numerator salience  
 on intuitive risk and likelihood judgments. *Judgment and  
 Decision Making*, 4(6):436–446, 2009.

Reddy, C. K. and Shojaee, P. Towards scientific discov-  
 ery with generative ai: Progress, opportunities, and chal-  
 lenges. In *Proceedings of the AAAI Conference on Artifi-  
 cial Intelligence*, volume 39, pp. 28601–28609, 2025.

ŞAHİN, E., Arslan, N. N., and Özdemir, D. Unlocking  
 the black box: an in-depth review on interpretability,  
 explainability, and reliability in deep learning. *Neural  
 Computing and Applications*, 37(2):859–965, 2025.

Santos, C. F. G. D. and Papa, J. P. Avoiding overfitting: A  
 survey on regularization methods for convolutional neural  
 networks. *ACM Computing Surveys (Csur)*, 54(10s):1–25,  
 2022.

Sapel, P., Molinas Comet, L., Dimitriadis, I., Hopmann, C.,  
 and Decker, S. A review and classification of manufac-  
 turing ontologies. *Journal of intelligent manufacturing*,  
 36(6):3669–3693, 2025.

Shao, J., Lu, Y., and Yang, J. Benford’s curse: Tracing digit  
 bias to numerical hallucination in llms. *arXiv preprint  
 arXiv:2506.01734*, 2025.

Shiraishi, H., Hayamizu, Y., Hashiyama, T., Takadama,  
 K., Ishibuchi, H., and Nakata, M. Evolutionary co-  
 optimization of rule shape and fuzziness in rule-based  
 machine learning. In *Proceedings of the Genetic and  
 Evolutionary Computation Conference Companion*, pp.  
 71–72, 2025.

Sidhu, R., Wadhwa, S., Mei, A., and Prasanna, V. K. A self-  
 reconfigurable gate array architecture. In *International  
 Workshop on Field Programmable Logic and Applica-  
 tions*, pp. 106–120. Springer, 2000.

Sun, M., Yuan, W., Liu, G., Matusik, W., and Chen, J.  
 Foundation molecular grammar: Multi-modal foundation  
 models induce interpretable molecular graph languages.  
*arXiv preprint arXiv:2505.22948*, 2025.

Tonda, A. Review of pysr: high-performance symbolic  
 regression in python and julia, 2025.

Wadayama, T., Igarashi, K., and Takahashi, T. Physics-  
 aware sparse signal recovery through pde-governed mea-  
 surement systems. *arXiv preprint arXiv:2501.13414*,  
 2025.

550 Wang, J. A tutorial on llm reasoning: Relevant methods  
 551 behind chatgpt o1. *arXiv preprint arXiv:2502.10867*,  
 552 2025.

553

554 Wang, R., Wang, X., Gao, C., Chong, C. Y., Xia, X., and  
 555 Liao, Q. Axiom: Benchmarking llm-as-a-judge for code  
 556 via rule-based perturbation and multisource quality cali-  
 557 bration. *arXiv preprint arXiv:2512.20159*, 2025.

558

559 Xu, F., Hao, Q., Zong, Z., Wang, J., Zhang, Y., Wang,  
 560 J., Lan, X., Gong, J., Ouyang, T., Meng, F., et al. To-  
 561 wards large reasoning models: A survey of reinforced  
 562 reasoning with large language models. *arXiv preprint*  
 563 *arXiv:2501.09686*, 2025a.

564

565 Xu, F., Lin, Q., Han, J., Zhao, T., Liu, J., and Cambria, E.  
 566 Are large language models really good logical reasoners?  
 567 a comprehensive evaluation and beyond. *IEEE Transac-  
 568 tions on Knowledge and Data Engineering*, 2025b.

569

570 Yan, J.-J., Yang, G.-H., and Liu, X.-X. A multigain-  
 571 switching-mechanism-based secure estimation scheme  
 572 against dos attacks for nonlinear industrial cyber-physical  
 573 systems. *IEEE Transactions on Industrial Electronics*, 70  
 574 (5):5094–5103, 2022.

575

576 Yi, R., Georgiou, D., Liu, X., and Athanasiou, C. E.  
 577 Mechanics-informed, model-free symbolic regression  
 578 framework for solving fracture problems. *Journal of  
 579 the Mechanics and Physics of Solids*, 194:105916, 2025.

580

581 Yu, C., Velu, A., Vinitsky, E., Gao, J., Wang, Y., Bayen, A.,  
 582 and Wu, Y. The surprising effectiveness of ppo in coopera-  
 583 tive multi-agent games. *Advances in neural information  
 584 processing systems*, 35:24611–24624, 2022.

585

586 Yu, L., Cao, M., Cheung, J. C., and Dong, Y. Mechanistic  
 587 understanding and mitigation of language model non-  
 588 factual hallucinations. In *Findings of the Association  
 589 for Computational Linguistics: EMNLP 2024*, pp. 7943–  
 7956, 2024.

590

591 Yu, Z., Idris, M. Y. I., and Wang, P. Physics-constrained  
 592 symbolic regression from imagery. In *2nd AI for Math  
 593 Workshop@ ICML 2025*, 2025.

594

595 Zhang, J., Huang, J., Yao, H., Liu, S., Zhang, X., Lu, S.,  
 596 and Tao, D. R1-vl: Learning to reason with multimodal  
 597 large language models via step-wise group relative policy  
 598 optimization. *arXiv preprint arXiv:2503.12937*, 2025a.

599

600 Zhang, X., Zhang, X., Li, Q., Ruan, J., and Li, W. A novel  
 601 deep reinforcement learning framework fusing knowl-  
 602 edge and probabilistic modeling for prognostic and main-  
 603 tenance of aircraft engines. *Engineering Research Ex-  
 604 press*, 7(2):025570, 2025b.

605

Zhang, Y., Xie, F., Dong, Y., Yang, G., and Zhou, X. High  
 fidelity virtualization of cyber-physical systems. *Inter-  
 national journal of modeling, simulation, and scientific  
 computing*, 4(02):1340005, 2013.

Zhao, D., Gao, C., Li, J., Fu, H., and Ding, D. Pid control  
 and pi state estimation for complex networked systems:  
 A survey. *International Journal of Systems Science*, pp.  
 1–16, 2025.

## 605 Appendix A: Implementation Details and 606 Symbolic Logic

607 This appendix provides the technical specifications for the  
608 symbolic framework of AutoAxiom. We detail the prompt  
609 engineering strategy, the formal schema of the Intermediate  
610 Representation (IR), and the automated safeguarding  
611 mechanisms based on the system's architecture.

### 613 614 A1. Multi-Agent Prompt Templates

615 The core of AutoAxiom's creativity is driven by the Tripartite  
616 Consensus Axiom Proposer, comprising  $A_3$  (Radical  
617 Innovator),  $A_4$  (Conservative Guardian), and  $A_5$  (Consensus  
618 Maker) agents. Below is the exact system prompt used  
619 by the  $A_3$  Radical Innovator agent, rendered directly from  
620 the source code:

621 *Listing 1. System Prompt for A3 Agent*

622  
623  
624 You are an elite Theoretical Physicist and  
625 Control Systems Engineer. Your task is  
626 to redesign the mathematical logic at a  
627 specific 'Target Slot' to improve  
628 system performance.

629 [GLOBAL AWARENESS]:

630 Read 'full\_ir\_context' carefully.  
631 Understand how 'rho', 'u', and 'Q\_len'  
632 interact globally before modifying the  
633 specific slot. If this is a Co-  
634 Simulation, pay attention to coupling  
635 variables (e.g., Temperature affecting  
636 Service Rate). The new variables you  
637 introduced can only use the names from  
638 DVM.

639 [OPERATORS TOOLKIT - SCIENTIFIC DISCOVERY]:  
640 You must employ these specific axiom  
641 transformation operators to guide your  
642 modification:

- 643 1. OP\_RECAST: Maintain the physical or  
644 mathematical meaning while changing the  
645 descriptive language or base.
- 646 2. OP\_MAP: Map objects to another  
647 representation or space.
- 648 3. OP\_TRANSFORM: Introduce specific  
649 mathematical assumptions or expansions  
650 for derivation or approximation.
- 651 4. OP\_REDUCE: Reduce the complexity or the  
652 number of variables in an axiom or  
653 system.
- 654 5. OP\_SUBSTITUTE: Replace a general  
655 component with a specific instance or  
656 another equivalent component.
- 657 6. OP\_EQUIV: Establish a relationship of  
658 equivalence between two seemingly  
659 different axioms or systems.
- 660 7. OP\_PROMOTE: Elevate a secondary or  
661 specific property to a fundamental  
662 axiom or principle.
- 663 8. OP\_BOUNDARY: Define or transform the

664 limits or constraints of a system to  
665 ensure global consistency.

- 666 9. OP\_DYNAMICS: Write or transform the  
667 equations or dynamical forms that  
668 govern the evolution of a system.
- 669 10. OP\_STATISTICAL: Rewrite deterministic  
670 axioms as statistical or expectation  
671 values or probability rules.

680 [CONTROL THEORY GUIDELINES]:

- 681 - Avoid Magic Numbers: Derive coefficients  
682 from physical parameters (e.g., 'Cooling\_Coeff') rather than guessing  
683 random floats.
- 684 - Use Feedback: Implement Proportional (P)  
685 or Proportional-Integral (PI) control  
686 logic.
- 687 - Prevent Deadlock: Ensure rates never drop  
688 to exactly 0.0 unless intended. Use 'max(0.1, ...)' to ensure Liveness.

689 [STRICT CONSTRAINT]:

- 690 1. OUTPUT: Only output the JSON for 'draft\_leaf'.
- 691 2. SCHEMA: Adhere to 'AxiomExpr' Schema.  
Variables must be {"ref": "name"}.
- 692 3. VARS: You can introduce new variables if  
they exist in the Domain Vocabulary  
Mapping (DVM).

700 The A4 agent serves as a Reliability Engineer, focusing on  
701 identifying edge-case failures and ensuring physical stability.

711 *Listing 2. System Prompt for A4 Agent*

712 You are a Senior Reliability Engineer.  
713 Your job is to catch dangerous "bugs" in  
714 the Radical Innovator's proposal.

715 [RESPONSIBILITY]:

- 716 1. Global Safety: Check if the radical  
717 change conflicts with global  
718 constraints (in 'full\_ir\_context').
- 719 2. Liveness Check: Does the new formula  
720 allow critical rates to drop to 0? (Deadlock  
721 risk).
- 722 3. Stability Check: Watch for "Bang-Bang"  
723 control (oscillating between min/max).  
Suggest smoothing or hysteresis.
- 724 4. DVM Compliance: Ensure all new variables  
725 exist in the DVM.
- 726 5. Schema: Verify syntax.

727 Output a safe version of the leaf node.

738 The A5 agent acts as the Chief Architect, balancing the  
739 innovation of A3 with the safety requirements of A4.

## A2. AxiomDSL Syntax and Core IR Specification

```

660 Listing 3. System Prompt for A4 Agent
661 You are the Chief Architect. You make the
662 final decision.
663 Synthesize the 'radical\_proposal' and 'conservative\_review'.
664
665 [DECISION LOGIC]:
666 1. Prioritize Validity: The system must not
667     crash (NaN, infinite loops).
668 2. Prioritize Performance: If Radical is
669     safe, prefer the innovation.
670 3. Ensure Mathematical Soundness: The logic
671     should align with the physics derived
672     in A3.
673
674 Ensure the 'final\_leaf' is valid JSON and
675     mathematically sound.
676
677 The AxiomDSL serves as a high-level declarative lan-
678 guage for scientific modeling, which is subsequently low-
679 ered into the Core Intermediate Representation (Core IR)
680 for automated reasoning and simulation. The IR is for-
681 mally defined as a recursive Directed Acyclic Graph (DAG),
682  $\mathcal{G} = (N, L, \mathcal{O}, V, \phi)$ , where every node is strictly typed and
683 grounded in the Domain Vocabulary Mapping (DVM).
684
685 The AxiomDSL (as seen in input_axioms_*.json) organ-
686 izes domain knowledge into seven functional modules
687 to ensure a complete system description:
688
689 


690     - params: Immutable global constants (e.g.,  $\alpha_{base}$ ,  $\lambda_p$ )
691         with mandatory bounds and default values.

692
693     - symbols: Dynamic variables categorized into state
694         (physical status), control (modifiable logic), and
695         aux (sensors).

696
697     - rules: Procedural logic containers using assign for
698         deterministic updates and guarded for conditional
699         interventions.

700
701     - stochastic: Probabilistic definitions (e.g., Poisson,
702         Normal) linked to physical rates.

703
704     - temporal: Global safety properties ( $\Box$ ) and liveness
705         targets ( $\Diamond$ ) defined via Linear Temporal Logic (LTL).

706
707     - fixpoints: Declarative evolution laws for iterative state
708         updates.

709
710     - boundary: Spatial constraints (e.g., dirichlet,
711         neumann) for field-based simulations.

712
713 In the physical field domain, the DSL defines the coefficients
714 of partial differential equations. The following example
715 demonstrates a pixel-level mapping of a diffusion axiom.
716

```

**AxiomDSL Representation:** The original DSL specifies a base thermal diffusivity  $\alpha$ .

```
{
  "name": "alpha",
  "type": "real",
  "role": "control",
  "unit": "m^2/s",
  "rhs": { "ref": "alpha_base" }
}
```

**Evolved Core IR Representation:** When the **A3 Radical Innovator** applies an `OP_TRANSFORM`, the IR expands into a complex recursive tree to model non-linear saturation:  $\alpha = \alpha_{base} \cdot \exp(1.0 - u_{mean})$ .

```
{
  "node_type": "Assignment",
  "lhs": "alpha",
  "rhs": {
    "node_type": "OperatorNode",
    "op": "*",
    "args": [
      { "node_type": "ReferenceNode", "ref": "alpha_base", "meta": { "role": "param" } },
      {
        "node_type": "FunctionNode",
        "op": "exp",
        "args": [
          {
            "node_type": "OperatorNode",
            "op": "-",
            "args": [
              { "node_type": "ValueNode", "val": 1.0, "type": "real" },
              { "node_type": "ReferenceNode", "ref": "u_mean", "meta": { "unit": "T" } }
            ]
          }
        ]
      }
    ],
    "provenance": { "agent": "A3_Radical_Innovator", "op": "OP_TRANSFORM" }
  }
}
```

In the queueing domain, the IR must bridge discrete states ( $Q_{len}$ ) with continuous service rates ( $u$ ) while enforcing liveness.

**AxiomDSL Representation:** The DSL defines a load-dependent service rate.

```
{
  "type": "assign",
  "lhs": "u",
  "rhs": { "op": "tanh", "args": [ { "ref": "Q_len" } ] }
}
```

```

715 }
716
717
718 Lowered Core IR with Liveness Wrapper: The com-
719 piler automatically injects a max operator to satisfy the Con-
720 trol Theory Guideline of preventing deadlocks ( $u > 0.1$ ).
721
722 {
723     "type": "assign",
724     "lhs": "u",
725     "rhs": {
726         "node_type": "FunctionNode",
727         "op": "max",
728         "args": [
729             { "node_type": "ValueNode", "val": 0.1 },
730             {
731                 "node_type": "OperatorNode",
732                 "op": "+",
733                 "args": [
734                     { "ref": "mu_p", "meta": { "role": "rate" } },
735                     { "op": "tanh", "args": [ { "ref": "Q_len" } ] }
736                 ]
737             }
738         ],
739         "metadata": { "is_liveness_critical": true }
740     }
741 }

```

To ensure sound symbolic execution, the IR maps all DSL primitives to rigorous mathematical functions.

Category	IR Operator	Symbolic Representation
Arithmetic	$+, -, *, /, ^$	$a + b, a - b, a \times b, a \times b, a^b$
Non-linear	$\exp, \tanh, \ln, \sqrt$	$\exp(x), \tanh(x), \ln(x)$
Statistics	mean, SD, max, min	$\mathbb{E}[X], \sigma(X), \max(X), \min(X)$
Temporal	always, eventually	$\Box\phi, \Diamond\phi$
Dynamics	$dt, \text{laplacian}, \text{grad}$	$\frac{\partial}{\partial t}, \nabla^2, \nabla$

Table 3. Formal Mapping of AxiomDSL Operators to Symbolic Logic.

### A3. DVM and Grammar Specification

This section details the Domain Vocabulary Mapping (DVM) schema. Each DVM is a JSON structure specifying  $\mathcal{N}$  (symbols with type/unit/role),  $\mathcal{O}$  (operator signatures),  $\Upsilon$  (units), and  $\Phi$  (constraints). Exemplars for two domains are shown below; full specifications for all eight domains are provided in supplemental material.

#### DVM Schema (Queueing Domain Exemplar)

```

765 {
766     "domain": "queueing",
767     "symbols": {
768         "lambda_n": { "type": "real", "unit": "1/s", "role": "rate", }
769     }

```

```

        "desc": "Poisson arrival
        rate",
        "W_n": { "type": "real", "unit": "s", "role": "state", "desc": "average queueing
        delay" },
        // ... 27 additional symbols (see
        supplemental material)
    },
    "operators": {
        "sum": { "sig": ["real[]", "real"], "unit": "auto" },
        "if": { "sig": ["bool", "real", "real", "real"], "unit": "auto" }
        // ... 15 additional operators
    },
    "constraints": {
        "stability": "rho_tot < 1",
        "power_bounds": "P_min <= P <= P_max"
        // ... 2 additional constraints
    }
}

```

#### DVM Schema (Molecular Domain Exemplar)

```

{
    "domain": "molecular",
    "symbols": {
        "mw": { "type": "real", "unit": "g/mol", "role": "state", "desc": "Molecular Weight" },
        "qed": { "type": "real", "unit": "1", "role": "metric", "desc": "Drug-likeness [0,1]" }
        // ... 18 additional pharmacophore
        descriptors
    },
    "operators": [ "tanh", "sigmoid", "piecewise" ],
    "constraints": {
        "lipinski_mw": "mw <= 500",
        "lipinski_hba": "num_h_acceptors <= 10"
        // ... 4 additional Lipinski rules
    }
}

```

**AxiomDSL Grammar v1.5** The grammar is defined as an EBNF specification supporting arithmetic, trigonometric, vector calculus (for PDE domains), and temporal operators. Key productions include:

```

axiom_file := use_stmt* symbol_stmt*
            param_stmt*
            rule_stmt* contract_stmt*
rule_stmt := assign | guarded | piecewise
            | stochastic
            | temporal | boundary
guarded := "when" bexpr ":" rule_stmt
piecewise := "piecewise" "(" bexpr ":" expr ) ")"
boundary := "boundary" (dirichlet|neumann )

```

```

770
771         " (" axis "==" value ")" ":" "
772         u_assign
773         temporal := ("always" | "eventually") "(
774             bexpr ")"
775         // Full EBNF and 40 built-in operators in
776         // supplemental material
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824

```

**Note on Reproducibility.** Complete DVM JSON files (queueing, triage, software\_opt, physical\_fields, resource\_allocation, co\_simulation, evogym, molecular) and the full fragment library (33 SMILES strings) are archived in the supplemental material to ensure exact reproduction of the experimental domains.

## Appendix B: Simulators and Details of Experiments

### B0.1. Scaled Relative Advantage (SRA) Protocol

Unlike traditional Pareto ranking which requires multi-objective vectors, we focus on the magnitude of frontier expansion. We define the Baseline (G1) as the anchor point (0.5) and the current State-of-the-Art (SOTA) as the efficiency frontier (1.0). The Scaled Relative Advantage (SRA) explicitly quantifies how far an agent pushes beyond this known limit:

$$S_{\text{agent}} = 0.5 + 0.5 \times \frac{P_{\text{agent}} - P_{\text{G1}}}{|P_{\text{SOTA}^*} - P_{\text{G1}}|}$$

This metric serves as a proxy for Pareto Frontier Expansion:

- $S > 1.0$  confirms that the discovered axiom constitutes a new Pareto optimal solution that strictly dominates previous best-known methods. The  $P$  represent the Physics Scores in each domain.

### B0.2. Rationalization of Objective Physical Metrics Physical Scores

To ensure a rigorous and unbiased evaluation of AutoAxiom across heterogeneous domains, we distinguish between the Iterative Reward ( $R$ ) used for evolutionary guidance and the Physical Score ( $P_s$ ) used for cross-regime performance benchmarking. The design of  $P_s$  adheres to three fundamental principles of axiomatic discovery: Bounded Monotonicity Mapping: All physical costs (e.g., waiting time  $W_q$ , error rate  $ER$ , or  $L_2$  error) are mapped into a dimensionless utility space  $[0, 100]$  using the standard reciprocal form  $100/(1 + \sum \text{Cost})$ . This structure is widely utilized in control theory to transform unbounded error norms into bounded performance indices, ensuring that incremental improvements in near-optimal regimes are as visible as major gains in high-error regimes.

Multiplicative Yield Logic (Non-Additive Coupling): For multi-objective domains such as Software Optimization (F3) and Molecular Design (F8), we employ multiplicative composition (e.g.,  $P_s \propto \Phi \cdot (1 - ER)$ ). Unlike additive weights which allow a "severe failure" in one dimension to be masked by high performance in another, the multiplicative form treats each metric as a logical gate. This mirrors the "Effective Yield" principle in manufacturing and systems engineering, where a total system failure occurs if any single axiomatic constraint is violated.

Physical Regularization and Anchoring: In domains involving numerical stability (F4) or synchronization (F6), coefficients such as  $0.1 \cdot TV$  are derived from Tikhonov Regularization principles. These coefficients act as "anchoring factors" that penalize non-physical chattering or systemic

noise. To prevent bias, these coefficients were determined during the baseline calibration phase ( $G_1$ ) and remained invariant throughout the evolution of AutoAxiom, ensuring that the performance gains stem from symbolic logic innovation rather than parameter tuning.

### B0.3. Sigmoid Normalization Protocol (S-Score)

While SRA provides an intuitive measure of frontier expansion, it exhibits limited sensitivity when evaluating internal architectural ablations where performance variances may be non-linear or clustered near the baseline. To rigorously capture the contribution of individual axiomatic components, we introduce the Sigmoid Normalization Protocol (S-Score) for sensitivity analysis:

$$S = \frac{1}{1 + \exp\left(-2 \cdot \frac{R_{\text{agent}} - R_{\text{base}}}{\sigma}\right)} \quad (6)$$

where  $R_{\text{base}}$  is the anchor reward derived from the  $G_1$  baseline, and  $\sigma$  is a domain-specific scaling factor that determines the sensitivity gradient. This protocol offers three distinct advantages for ablation studies:

- **High Midpoint Sensitivity:** The S-Score is centered at 0.5 when  $R_{\text{agent}} = R_{\text{base}}$ . The high derivative of the sigmoid function near this anchor point ensures that even minor performance degradations caused by the removal of a specific module are reflected as significant numerical drops in the score.
- **Saturation Awareness:** In regimes approaching physical limits, the sigmoid curve naturally dampens marginal gains, forcing the evaluation to focus on the stability and robustness of the core axiomatic structure rather than raw numerical outliers.
- **Cross-Domain Calibration:** By assigning domain-specific  $\sigma$  values (e.g.,  $\sigma = 500$  for Triage vs.  $\sigma = 150$  for Physical PDE), we normalize heterogenous reward distributions into a unified  $[0, 1]$  interval, allowing for the calculation of a robust Overall S-Score across all six domains.

### B0.4. Hyperparameters and Model Configurations

To ensure the reproducibility of the evolutionary trajectory and the formal verification results, we provide the comprehensive hyperparameter suite used in the `main.py` orchestration and simulation adapters. All experiments were conducted using the **DSPy (v2.4.9)** framework to manage the tripartite agent interactions.

**Dynamic Weight Scheduling** As implemented in the `calculate_dspy_metric` function, the coefficients for

the dynamic objective function  $J(A)$  follow a non-linear trajectory to balance exploration and stability:

- **Violation Penalty** ( $\lambda_v$ ):  $\lambda_v(t) = \Lambda_0 \cdot e^{-t/T_{\text{rise}}}$ . Notably, for proposals achieving superior performance (raw score  $> 10.0$ ), a  $0.2 \times$  multiplier is applied to the penalty to prioritize efficiency discovery.
- **Persistence Reward** ( $\lambda_p$ ):  $\lambda_p(t) = \Lambda_0 \cdot (1 - e^{-t/T_{\text{stab}}})$ , ensuring the system prioritizes "white-box" robustness in later generations.
- **Diversity Reward** ( $\lambda_d$ ):  $\lambda_d(t) = \Lambda_0 \cdot \max(0, 1 - \frac{2t}{T_{\text{total}}})$ , effectively shifting the search from high-entropy jumps to local manifold refinement.

**Diversity Metric** The structural novelty  $\mathcal{P}_{\text{div}}$  is calculated by comparing the serialized byte-length of the proposed AxiomIR tree against the baseline structure:

$$\mathcal{P}_{\text{div}} = \min \left( 1.0, \frac{|\text{size}_{\text{proposal}} - \text{size}_{\text{baseline}}|}{\text{size}_{\text{baseline}}} \times 2.0 \right) \quad (7)$$

**Few-Shot Optimization** We utilized the `dspy.BootstrapFewShot` optimizer with `max_bootstrapped_demos=3`. This allows the system to autonomously extract and prepend successful axiomatic modifications from the `experiment_history` as contextual exemplars for the Radical and Conservative agents in subsequent rounds.

## B1. Customer Queue System

### Background and Strategic Importance

Queueing systems serve as the foundational cornerstone of operations research and the control nexus for modern large-scale distributed architectures, high-frequency communication networks, and intelligent logistics systems. In these highly stochastic physical environments, the core scientific challenge lies in managing the delay explosion caused by non-stationary stochastic processes.

In the environment defined by the F1 simulator, the system must navigate a Pareto optimal solution between processing power (resource energy consumption) and quality of service (waiting latency). If the control axioms are too conservative, they lead to high hardware idle time and wasted energy; if too aggressive, they risk non-linear divergence of the queue length ( $Q_{\text{len}}$ ) during sudden traffic spikes. The scientific value of AutoAxiom lies in its ability to discover evolved, symbolic control axioms with "adaptive resilience," providing the mathematical scaffolding for self-healing infrastructures.

### Environmental Configuration and Execution Protocol

The F1 simulator is built upon the `simpy` discrete-event

880  
881 **Table 4. AutoAxiom System Hyperparameters.** Parameters are extracted directly from the operational implementation used to generate  
882 the results in Section 4.

Category	Parameter	Notation	Value
<b>LLM Configuration</b>	Underlying Model	-	gpt-4o
	Sampling Temperature	$T_{llm}$	0.9
	Max Tokens (Generation)	-	8,192
	Reasoning Framework	-	DSPY
<b>Evolutionary Control</b>	Total Rounds	$T_{total}$	15
	History Window	$ \mathcal{H} $	5
	Monte Carlo Runs	$N_{mc}$	100
	Simulation Base Seed	-	42
<b>Annealing Schedule</b>	Base Scaling Factor	$\Lambda_0$	100.0
	Violation Penalty Decay	$T_{rise}$	3.0
	Persistence Growth Lag	$T_{stab}$	10.0
<b>Objective Weights</b>	Performance Weight	$w_{perf}$	1.0
	Persistence Weight	$w_{pers}$	1.5
	Diversity Weight	$w_{div}$	1.0
	Compatibility Bonus	$w_{comp}$	1.0
<b>Normalization</b>	Safe Score Minimum	$J_{min}$	-1,500.0
	Safe Score Maximum	$J_{max}$	300.0

904 driving engine. To ensure absolute reproducibility and statistical rigor, all experimental groups (G1–G4) adhere to a “Zero-Blackbox” simulation protocol:

- **Multi-Seed Validation:** Each experimental group is subjected to 100 independent simulation runs (Seeds 42–141).
- **Stochastic Modeling:** Customer arrivals follow a Poisson process ( $\lambda_p$ ), and service times follow an exponential distribution with a rate  $u$  determined in real-time by the active axiomatic logic.
- **Control Granularity:** The axiomatic logic performs global sampling and state decisions every  $\Delta t = 0.5$  seconds.

921 The table below discloses the complete initial physical parameters configured in the simulation engine:

922  $\lambda_p: 0.5 \text{ s}^{-1}$

923 **Arrival Rate** — Poisson arrival rate with a mean of 0.5.

924  $\mu_p: 1.5 \text{ s}^{-1}$

925 **Base Service Rate** — Nominal system capacity (baseline rate starting point).

926  $\mu_{\max}: 5.0 \text{ s}^{-1}$

927 **Maximum Capacity** — Upper bound of the physical hardware processing rate.

928  $\mu_{\min}: 1.0 \text{ s}^{-1}$

929 **Minimum Operation** — Minimum rate required to maintain system activity.

930  $\rho_{\max}: 0.999$

931 **Saturation Threshold** — Critical utilization threshold for system instability.

932  $\Delta t: 0.5 \text{ s}$

933 **Control Step** — Step size for axiomatic adaptive adjustment and decision.

## Metrics and Reward

934 The simulator tracks seven core physical metrics in real-time, forming the basis of the evolutionary fitness landscape. All metrics are reported as the mean and 95% Confidence Interval (CI) over 100 runs:

1. **Wait Time ( $\bar{W}_q$ ):** Average duration customers wait in the buffer:  $\frac{1}{N} \sum (t_{\text{start},i} - t_{\text{arrival},i})$ .
2. **Occupancy ( $L_{\text{avg}}$ ):** Average number of customers, verified via Little’s Law:  $L = \lambda_p \cdot \bar{W}_q$ .
3. **Utilization ( $\rho$ ):** Load factor defined as  $\rho = \lambda_p / u$ .
4. **Resilience ( $Res$ ):** Stability resilience index, calculated as  $Res = 1.0 / (1.0 + 0.1 \cdot \rho)$ .
5. **Final Reward ( $R_{\text{final}}$ ):** The “North Star” metric for AutoAxiom iteration:

$$R_{\text{final}} = 0.7 \cdot P_s + 0.3 \cdot S_{\text{custom}}$$

935  
936 Where  $S_{\text{custom}}$  is the self-defined logic score within the  
937 axiom, weighting delay reduction (70%) and logical  
938 rigor (30%).  
939  
940

941 6. **Physics Score ( $P_s$ ):** Standardized physical performance score:  $P_s = 100.0 / (1.0 + \bar{W}_q)$ .  
942  
943 7. **Score ( $S$ ):** Normalized against G1 ( $R_{\text{base}}$ ) with  $\sigma = |R_{\text{base}}|$ .  
944

#### 945 Deconstruction of Experimental Regimes (G1–G3)

946 To contrast evolutionary advantages, we established three  
947 highly persuasive benchmark groups:  
948

- 949 950 951 • **G1: Baseline (Static Open-Loop):**  $u = \mu_p$ . Represents primitive control without feedback, offering zero defense against traffic pulses.  
952  
953 954 • **G2: SOTA-II (Linear Lyapunov Feedback):**  $u = \text{Sat}(\mu_p + \beta \cdot Q_{\text{len}})$ . Based on classical Lyapunov Stability Theory, it establishes linear feedback. While convergence is guaranteed, it suffers from significant response lag under non-stationary traffic.  
955  
956 957 • **G3: SOTA-I (Industrial FCFS Stability):** Utilizes a fixed high-efficiency rate fine-tuned by experts, coupled with strict FCFS hardware contention logic. It represents the physical performance ceiling for traditional non-adaptive strategies.  
958  
959 960 961 962 963

#### 964 G4: The Evolved AutoAxiom Logic

965 AutoAxiom (G4) broke free from the constraints of linearity,  
966 self-synthesizing a composite physical system with perception,  
967 adaptive regulation, and defense layers.  
968  
969

#### 970 Symbolic Physical Expression of G4:

$$971 u = \text{Sat}(\mu_p \cdot (P(\rho) + \beta Q_{\text{len}})) \cdot (1 + \gamma \cdot \mathbb{1}_{\{Q_{\text{len}} > \tau\}}) \quad (8)$$

972 The parameters determined through evolution are:  
973  
974

- 975 976 977 • **Non-linear Gain  $P(\rho)$ :** If  $\rho < 0.4$ ,  $P = 1.05$  (proactive pre-clearing); if  $\rho > 0.6$ ,  $P = 0.95$  (robustness smoothing).  
978  
979 980 • **Feedback Coefficient  $\beta$ :** The system identified an optimal fine-tuning coefficient of 0.02.  
981  
982 983 • **Gated Defense  $\gamma, \tau$ :** Detected a critical threshold  $\tau = 10$ . When  $Q_{\text{len}} > 10$ , it triggers a pulse gain of  $\gamma = 0.15$ .  
984  
985 986

#### 987 Analysis and Inference

988 The experimental results in the Queueing (F1) domain reveal that the evolved G4 axiom represents a sophisticated  
989

Table 5. F1 Customer Queue: Performance Statistics (100 Runs)

Metric	Baseline	Lyapunov (G2)	FCFS-High (G3)	AutoAxiom (G4)
$W_q$ (s)	$1.98 \pm 0.03$	$1.95 \pm 0.02$	$0.99 \pm 0.06$	$0.34 \pm 0.05$
$L_{\text{avg}}$	$0.99 \pm 0.02$	$0.98 \pm 0.01$	$0.50 \pm 0.03$	$0.17 \pm 0.02$
<b>Physics Score</b>	$33.66 \pm 0.34$	$33.92 \pm 0.24$	$50.98 \pm 1.25$	$76.53 \pm 2.20$
<b>SRA Score</b>	$0.50 \pm 0.01$	$0.51 \pm 0.01$	$1.00 \pm 0.04$	$1.74 \pm 0.06$

departure from traditional linear control paradigms. The mathematical core of the G4 logic,

$$u = \text{Sat}(\mu_p \cdot (P(\rho) + \beta Q_{\text{len}})) \cdot (1 + \gamma \cdot \mathbb{1}_{\{Q_{\text{len}} > \tau\}}) \quad (9)$$

functions as a multi-tiered defense strategy rather than a simple feedback loop. Unlike the Lyapunov-based SOTA (G2), which relies on a rigid linear feedback gain ( $Q_{\text{len}}/V$ ) that often reacts too sluggishly to initial backlogs or over-compensates during transients, AutoAxiom discovered a **“state-aware modal switching” logic**.

By integrating a non-linear utilization sensor  $P(\rho)$ , the system proactively clears buffers when the load is low ( $\rho < 0.4$ ), preventing the initial accumulation that typically leads to downstream congestion. More significantly, the autonomous discovery of the **“High-Water Mark”** at  $\tau = 10$  allows the system to remain energy-efficient during normal operations while triggering a sharp 115% emergency boost precisely when the queue enters the critical exponential growth phase. This synergistic combination of proactive clearing and reactive pulsing results in an **82.8% reduction in average waiting time**. The fundamental inference is that in stochastic environments, a non-linear, threshold-aware defense mechanism is far more physically efficient and stable than any uniform proportional response.

#### Real-world Applicability

The G4 axiom possesses immense potential for direct deployment in high-stakes infrastructure, including **edge computing nodes, 5G base stations, and high-frequency trading gateways**. While most modern adaptive controllers rely on Deep Reinforcement Learning (DRL)—which functions as a “black box” and requires substantial computational overhead for real-time inference—the symbolic G4 axiom can be hard-coded directly into **ASIC or FPGA logic** as a set of lightweight, deterministic “if-then” instructions and simple arithmetic gates.

This implementation ensures nanosecond-level execution with near-zero computational cost and absolute auditability, fulfilling the stringent requirements of telecommunication standards and financial compliance. In a production data center, for instance, this axiom would allow a load balancer to manage unpredictable burst traffic with mathematical certainty, effectively preventing the **“tail latency”** spikes that degrade user experience. By bridging the gap between high-level heuristic optimization and hardened industrial

990 deployment, AutoAxiom provides a **”white-box” control**  
 991 **law** that is as interpretable as it is high-performing, making  
 992 it an ideal candidate for next-generation self-optimizing  
 993 networks.

994

995

## B2. Hospital Service Center (Triage System)

996

### Background and Strategic Importance

998 Hospital triage systems represent complex stochastic environments characterized by heterogeneous priorities and  
 999 preemptive resource competition. The primary operational  
 1000 challenge lies in managing the conflict between **”Critical”** and  
 1001 **”Minor”** patient flows. In high-pressure Emergency  
 1002 Rooms (ER), critical patients are legally afforded absolute  
 1003 preemption rights. However, if the axiomatic scheduling  
 1004 logic is too rigid, it leads to a non-linear divergence in minor  
 1005 queues (Queue Explosion), which not only deteriorates  
 1006 the waiting environment but also induces latent medical  
 1007 risks. The objective of AutoAxiom is to evolve a non-linear  
 1008 control law capable of dynamically balancing the **”critical**  
 1009 **lifeline”** with overall system throughput.

1010

### Environmental Configuration

1012 The F2 simulator is built on the `simpy` discrete-event engine,  
 1013 modeling an overloaded emergency department. To ensure full reproducibility, the physical boundary parameters  
 1014 of the simulation engine are defined as follows:

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

- **Arrival Process:** Poisson-distributed arrivals with  $\lambda_c = 1.0 \text{ s}^{-1}$  (Critical) and  $\lambda_m = 1.7 \text{ s}^{-1}$  (Minor).
- **Service Dynamics:**  $c = 2$  medical resources (servers) with a mean service efficiency of  $\mu = 1.0 \text{ s}^{-1}$  per resource.
- **Preemption Logic:** Enabled (Critical patients can interrupt ongoing services for Minor patients).
- **Execution Protocol:** 100 independent trials (Seeds 42–141) with a physical duration of  $T = 1000 \text{ s}$  per run to ensure steady-state evolution.
- **Critical Safety Threshold:** The target waiting time for critical cases is set at  $\tau_{crit} = 30.0 \text{ s}$ .

### Metrics and The Multi-Objective Reward Function

The simulator captures the following core metrics to define the evolutionary fitness landscape:

1038

1039

1040

1041

1042

1043

1044

- **Wait Times ( $W_c, W_m$ ):** The average queuing delays recorded for Critical and Minor patients, respectively.
- **Critical Coverage ( $C_{cov}$ ):** The ratio of critical patients who successfully received service (mandatory safety constraint:  $C_{cov} > 0.9$ ).

- **Critical Within Ratio ( $C_{win}$ ):** The proportion of critical patients whose wait time was within the 30.0 s threshold.

- **Final Reward ( $R_{final}$ ):** The core fitness function driving the evolution loop, calculated as:

$$R_{final} = 50.0 \cdot C_{win} + 10.0 \cdot C_{cov} - 1.0 \cdot W_c - 5.0 \cdot W_m$$

- **Physics Scores:**  $100 \cdot C_{cov} / (1 + \bar{W}_{all})$  to reflect balance between Safety and Efficiency.

- **Score ( $S$ ):** Given G1’s collapse ( $R \approx -990$ ), we set  $\sigma = 500$  (half the range of performance recovery). Note: A penalty weight for  $W_m$  (5.0) is five times higher than that for  $W_c$  to force the axioms to learn how to suppress minor queue divergence.

### Experimental Regimes (G1–G3)

We evaluate three baseline regimes to define the boundaries of traditional scheduling:

- **G1: Baseline (Static Preemption):** Uses a static factor of 1.0. Due to absolute preemption, the minor queue collapses ( $W_m > 200 \text{ s}$ ), resulting in a massive negative reward.
- **Sota1 (WSPT):** Industrial standard using Weighted Shortest Processing Time logic. It assigns a high penalty weight (5.0) to minor cases to compensate for their priority, representing the peak of static non-adaptive optimization.
- **Sota2 (Lyapunov Feedback):** Employs linear feedback where the priority factor scales as  $1.0 + \beta Q_{len}$ . While it compresses  $W_m$ , the linear growth lacks a “saturation guard,” often causing fluctuations in critical case certainty ( $W_c$ ).

### G4: The Evolved AutoAxiom Logic

AutoAxiom (G4) discovered a sophisticated non-linear regulation law that transcends traditional linear logic:

$$\text{Priority Factor} = 1.0 + 5.0 \cdot \tanh \left( \frac{W_m/100 + W_c/100}{2.0} \right) \quad (10)$$

The G4 regime utilizes a tanh **activation function** to achieve modal switching: the system remains stable during minor fluctuations but provides a non-linear, “gated” priority boost during critical accumulation. This prevents the minor queue from entering a divergent state while strictly adhering to the critical case safety redline.

### Performance Statistical Records (100 Runs)

Table 6. F2 Triage System: Performance Statistics (100 Runs)				
Metric	Baseline	WSPT (Sota 1)	Lyapunov (Sota 2)	AutoAxiom (G4)
$W_c$ (s)	$1.00 \pm 0.02$	$0.99 \pm 0.02$	$5.71 \pm 0.50$	$1.00 \pm 0.02$
$W_m$ (s)	$209.7 \pm 4.9$	$2.77 \pm 0.08$	$0.62 \pm 0.02$	$2.61 \pm 0.08$
Crit Cov	$0.998 \pm .00$	$0.998 \pm .00$	$0.993 \pm .00$	$0.998 \pm .00$
Physics Score	$0.75 \pm 0.01$	$32.07 \pm 0.29$	$28.35 \pm 1.11$	$33.09 \pm 0.34$
SRA Score*	$0.50 \pm 0.00$	$1.00 \pm 0.00$	$0.94 \pm 0.02$	$1.02 \pm 0.01$

\*Note: Standard SRA anchoring  $G1=0.5$  and Best SOTA=1.0.

## Inference and Conclusion

Analysis of the G4 regime reveals a 98.7% reduction in minor queue latency ( $W_m$ )<sup>2</sup>. The discovered axiom demonstrates a "dynamic prioritization" property: by utilizing the tanh activation function, the system remains dormant during routine flows but provides a non-linear priority boost as  $W_m$  or  $W_c$  approaches critical thresholds<sup>3</sup>.

## Analysis and Inference

Compared to the academic authority Lyapunov (G3), G4 achieved a superior reward profile ( $59.98 \pm 0.00$ )<sup>4</sup>. Derivation from the physical expression reveals: G3 relies on "passive response," where priority scales linearly with queue length; whereas G4 possesses "anticipatory balancing" capabilities, utilizing the bounded nature of the tanh operator to prevent priority saturation<sup>5</sup>.

## Real-world Applicability

The symbolic nature of the G4 axiom allows it to be directly compiled into ultra-lightweight logical instructions and deployed on ASIC chips for high-performance switches or data center gateways. This logic, combining "crisis warning" with "adaptive regulation," is the core scientific path toward building future zero-cost inference and ultra-low latency intelligent networks.

## B3. Software Development Process Optimization (Softwareopt)

### Background and Strategic Importance

Modern software development is a complex adaptive system where the primary scientific challenge lies in managing the non-linear interplay between **Feature Velocity** and **Systemic Entropy**. The F5 simulator environment models a high-fidelity software lifecycle where developers must allocate a finite pool of effort ( $\eta$ ) across multiple concurrent modules. The core objective is to maximize the project's effective productivity while preventing the accumulation of "Technical Debt" (manifested as SD-Complexity) and severe regressions (Error Rate). If effort allocation is too conservative, development stalls; if too aggressive, the "Broken Window Effect" triggers exponential error growth. AutoAxiom is tasked with discovering a governance axiom that transcends reactive heuristics by sensing developmental "friction" and non-linearly modulating the intensity of resource application.

## Environmental Configuration

The F5 environment is a multi-agent, stochastic simulation of a software ecosystem with the following boundary parameters:

- **Project Scale:** 10 parallel development modules with non-uniform, stochastic task arrival rates.
- **Systemic Decay:** A baseline complexity drift simulates code rot due to shifting requirements and library dependencies.
- **Simulation Protocol:** 100 simulation cycles per trial, executed over 100 independent trials (Seeds 42–141) to ensure statistical convergence.
- **Dynamic Constraints:** A fundamental trade-off exists where increased Progress Rate ( $\Phi$ ) induces a non-linear penalty on SD-Complexity ( $\sigma_c$ ) and Error Rate ( $ER$ ).
- **Initialization:** Module readability and consistency are initialized within the [0.7, 0.8] range to simulate a healthy starting state.

## Metrics and The Software Governance Reward Function

Performance is quantified via a multi-objective fitness function that rewards sustainable growth:

1. **Avg\_ProgressRate ( $\Phi$ ):** The mean percentage of feature completion per unit time.
2. **Avg\_SD\_Complexity ( $\sigma_c$ ):** The standard deviation of complexity across modules, representing the "entropy" and lack of balance in development.
3. **Avg\_ErrorRate ( $ER$ ):** The frequency of defect introduction during the execution phase.
4. **Avg\_W\_H:** The mean holding cost/waiting time for feature delivery, representing systemic latency.
5. **Physics Score ( $P_s$ ):** The effective high-quality yield:  $P_s = 100 \cdot \Phi \cdot (1 - ER) \cdot (1 - \sigma_c)$ .
6. **Final Reward:**  $R_{soft} = w_1 \cdot \Phi - w_2 \cdot \sigma_c - w_3 \cdot ER + w_4 \cdot CQ$ , designed to penalize "unhealthy" speed.
7. **SRA Score (S):** Normalized against the  $G_1$  baseline (0.5) and SOTA\* (1.0).

## Experimental Regimes (G1–G3)

- **G1: Baseline (Static Resource Allocation):** Employs a fixed effort coefficient  $\eta = 0.1$ . This "open-loop" approach is oblivious to the project state. While it

maintains a low Error Rate, it results in a stagnant Progress Rate ( $0.1002 \pm 0.0008$ ), leading to high systemic latency and an inability to handle demand spikes.

- **SOTA1: Velocity-Focused Heuristic (G2):** A high-intensity growth model focusing purely on throughput:  $\eta = \eta_{base} \cdot (1 + k \cdot \Phi)$ . While achieving the highest raw Progress Rate ( $0.9443 \pm 0.0004$ ), it triggers the "Technical Debt Trap," resulting in a massive spike in SD-Complexity ( $0.2669 \pm 0.0103$ ) and an unsustainable Error Rate ( $0.1639 \pm 0.0008$ ).
- **SOTA2: Balanced Agile Framework (G3):** A linear feedback model:  $\eta = \eta_{base} \cdot (1 + \gamma_1 \Phi - \gamma_2 \sigma_c)$ . This proactive strategy attempts to "brake" when complexity rises. It achieves a superior Reward ( $4.2314 \pm 0.0075$ ) compared to G1, but its linear nature makes it sluggish during non-linear complexity explosions.

#### G4: The Evolved AutoAxiom Logic (Non-linear Friction-Aware Law)

In Round 11, AutoAxiom discovered a sophisticated **non-linear friction-sensing governance law** that identifies the optimal "tipping point" between productivity and debt:

$$\eta_{adj} = \eta_{base} \cdot [\Psi \cdot \log(1 + \Phi) \cdot \text{sigmoid}(CQ - \beta\sigma_c)] \quad (11)$$

Where  $\Psi$  represents a global momentum factor and  $CQ$  is code quality.

**Physical Inference:** G4 evolved a **Dynamic Entropy Gate**. The term  $\log(1 + \Phi)$  accounts for the Law of Diminishing Returns in development effort. Crucially, the  $\text{sigmoid}(CQ - \beta\sigma_c)$  operator acts as a non-linear "Circuit Breaker." When the system entropy ( $\sigma_c$ ) outweighs the quality-to-debt ratio, the sigmoid term collapses toward zero, forcing an immediate transition into a "Refactoring Phase." This prevents the "Broken Window" effect by proactively dampening aggressive development before technical debt reaches a point of structural collapse.

#### Performance Statistical Records (100 Runs)

##### Analysis and Inference: Optimal Control of Technical Debt

G4, through its sigmoid damping, maintains the lowest complexity dispersion ( $\sigma_c : 0.1742$ ). This proves that **non-linear "refactoring gates"** are more effective at maintaining long-term project health than linear compensation models (G3), representing a 10.5% improvement in systemic entropy over the Agile-Linear baseline<sup>7</sup>.

**Real-world Applicability:** The discovered G4 axiom is directly applicable to **Autonomous Project Governance** and **AIOps Orchestration**. By implementing this logic into DevOps control planes, organizations can automatically adjust team "WIP limits" or deployment gates based on real-time

metrics of code complexity and quality. Its logarithmic-sigmoid nature allows for a "smooth-yet-decisive" resource reallocation, ensuring that technical debt is serviced proactively without the oscillating "firefighting" behavior seen in traditional management.

## B4. Physical Control (Numerical Heat Conduction)

### Background and Strategic Importance

Physical control systems, particularly in the context of numerical simulations, require a delicate balance between **computational stability** and **physical fidelity**. The F4 simulator models a 1-D non-stationary heat conduction process governed by the diffusion equation. The core scientific challenge is the dynamic selection of the diffusion coefficient ( $\alpha$ ): it must satisfy the Courant-Friedrichs-Lowy (CFL) stability condition to prevent numerical divergence while minimizing the  $L_2$  error relative to the physical ground truth. AutoAxiom's goal is to evolve an adaptive control law capable of sensing numerical instability "precursors" and adjusting physical parameters to achieve Pareto optimality between accuracy and stability. **Environmental Configuration**

The F4 environment is a discretized physical grid where a high-intensity thermal pulse ( $u = 500$ ) is introduced. The simulation parameters are configured as follows:

- **Grid Dynamics:** 50 nodes with a spatial step of  $\Delta x = 0.1$ .
- **Simulation Scope:** 100 time steps per trial, executed over 100 independent trials (Seeds 42–141).
- **Base Physics:** Nominal diffusion coefficient  $\alpha_{base} = 0.01$ .
- **Numerical Constraint:** The system monitors the CFL number; if  $\alpha \cdot \Delta t / \Delta x^2$  exceeds 0.5, a complete stability penalty is applied.

### Metrics and The Physical Reward Function

Performance is evaluated across five physical dimensions to determine evolutionary fitness:

1.  **$L_2$  Error ( $E_{L2}$ ):** The root-mean-square deviation from the ideal physical state.
2. **Total Variation (TV):** Measures numerical oscillations; higher values indicate non-physical instability.
3. **Max Gradient ( $\nabla_{max}$ ):** The peak temperature slope, indicating localized stress on the numerical scheme.
4. **Physics Score ( $P_s$ ):** The physical fidelity index:  $P_s = 100 / (1 + E_{L2} + 0.1 \cdot TV)$ .

Table 7. F3 Software Optimization: Performance Statistics (100 Runs)

Metric	Baseline	Heuristic (G2)	Agile-Linear (G3)	AutoAxiom (G4)
Avg Progress Rate	$0.10 \pm 0.00$	$0.94 \pm 0.00$	$0.89 \pm 0.00$	<b><math>0.90 \pm 0.00</math></b>
Avg Error Rate	$0.10 \pm 0.00$	$0.16 \pm 0.00$	$0.09 \pm 0.00$	<b><math>0.09 \pm 0.00</math></b>
SD Complexity	$0.21 \pm 0.01$	$0.27 \pm 0.01$	$0.19 \pm 0.01$	<b><math>0.17 \pm 0.01</math></b>
Avg W_H	$5.55 \pm 0.01$	$6.74 \pm 0.02$	$5.38 \pm 0.01$	<b><math>5.55 \pm 0.01</math></b>
<b>Physics Score</b>	$7.16 \pm 0.13$	$57.88 \pm 0.55$	$65.02 \pm 0.83$	<b><math>67.46 \pm 0.84</math></b>
<b>SRA Score</b>	$0.50 \pm 0.00$	$0.94 \pm 0.00$	$1.00 \pm 0.01$	<b><math>1.02 \pm 0.01</math></b>

5. **Final Reward:**  $R_{phys} = -(E_{L2} + 0.1 \cdot TV + Penalty) \times 10^{-1}$ , enforcing stability-first control.

6. **SRA Score (S):** Normalized against the  $G_1$  baseline (0.5) and *SOTA*\* (1.0).

## Experimental Regimes (G1–G3)

- **G1: Baseline (Static Fourier’s Law):** Uses a fixed  $\alpha = 0.01$ . This regime fails to adapt to the high-gradient thermal pulse, leading to significant  $L_2$  errors ( $151.31 \pm 1.08$ ) and high  $TV$  due to unmanaged oscillations.
- **SOTA1: Industrial PD Control (Conservative):** A robust engineering standard utilizing a proportional-derivative law with a strict safety clamp:  $\alpha = \min(0.05, \alpha_{base} + K_p \bar{u} + K_d \nabla_{max})$ . While stable, the 10% safety margin (clamping at  $\alpha = 0.05$ ) limits its cooling throughput.
- **SOTA2: Classical Physical Linear Model:** A first-order adaptive feedback law:  $\alpha = \alpha_{base} \cdot (1 + 0.01 \cdot u_{mean})$ . By sensing the mean energy level, its accuracy improves but remains vulnerable to localized gradient spikes.

## G4: The Evolved AutoAxiom Logic (Non-linear Saturation Law)

AutoAxiom (G4) discovered a sophisticated non-linear coupling logic in Round 15 that utilizes a saturation-based mechanism to safely approach the physical limit:

$$\alpha_{adj} = \alpha_{base} \cdot [E_{pot} \cdot \tanh(\beta_1 \bar{u} + \beta_2 \nabla_{max})] \quad (12)$$

Where the evolved coefficients leverage the system’s potential energy ( $E_{pot}$ ) to drive diffusion while maintaining stability via the hyperbolic tangent operator.

**Analysis and Inference: The Physics of Evolved Stability**  
The G4 regime achieves a 76.4% reduction in  $L_2$  Error compared to the industrial SOTA1. The fundamental

Table 8. F4 Heat Conduction Control: Performance Statistics (100 Runs)

Metric	Baseline	PD-Control (G2)	Linear-Adapt (G3)	AutoAxiom (G4)
$L_2$ Error	$151.3 \pm 1.1$	$44.97 \pm 1.5$	$108.3 \pm 1.5$	<b><math>10.60 \pm 0.83</math></b>
CFL ( $\times 10^{-2}$ )	$0.51 \pm 0.00$	$2.56 \pm 0.06$	$0.93 \pm 0.02$	<b><math>5.09 \pm 0.13</math></b>
TV	$536.3 \pm 5.5$	$142.3 \pm 4.7$	$351.8 \pm 5.6$	<b><math>33.51 \pm 2.61</math></b>
Max Grad	$1475 \pm 13$	$510.9 \pm 11$	$1192 \pm 6.4$	<b><math>120.9 \pm 8.1</math></b>
<b>Physics Score</b>	$0.49 \pm 0.00$	$1.66 \pm 0.05$	$0.69 \pm 0.01$	<b><math>6.69 \pm 0.34</math></b>
<b>SRA Score</b>	$0.50 \pm 0.00$	$1.00 \pm 0.02$	$0.59 \pm 0.00$	<b><math>3.14 \pm 0.15</math></b>

inference is that G4 has successfully closed the gap between conservative engineering and the theoretical physical limit. By maintaining a mean CFL of  $\approx 0.051$  (nearly double that of SOTA1), G4 demonstrates the ability to operate safely within the high-performance regime that traditional controllers discard as “risky.” The significantly lower  $TV$  ( $33.51 \pm 2.61$  vs  $142.29 \pm 4.73$ ) proves that G4’s soft-saturation strategy suppresses the numerical chattering common in linear-plus-clamp systems.

**Real-world Applicability:** G4’s discovered axiom is highly applicable to **smart thermal materials** and **HPC physical solvers**. Its symbolic, non-linear form can be directly implemented as a constitutive law in phase-change heat sinks or adaptive CFD meshes. Unlike black-box neural solvers, its mathematical structure provides an explicit guarantee of boundedness via the tanh operator, making it suitable for safety-critical thermal management in aerospace and semiconductor cooling.

## B5. Resource Allocation (Dynamic Computing Nodes)

### Background and Strategic Importance

Resource allocation systems represent a fundamental optimization challenge in distributed computing and cloud infrastructure. The core scientific conflict lies in the trilemma between **Throughput** (efficiency), **Makespan** (completion time), and **Fairness** (load balancing). In a multi-tenant cloud node with finite capacity, static or purely greedy allocation logic either starves heavy tasks or underutilizes hardware during demand volatility. AutoAxiom’s objective

1210 is to discover a non-linear allocation axiom that dynamically  
 1211 adjusts resource quotas based on system stress and  
 1212 cost boundaries to maximize global utility.

### 1213 Environmental Configuration

1214 The F5 simulator models a high-concurrency task processing  
 1215 node with the following physical parameters:

- 1218 • **Task Load:** 20 heterogeneous tasks with varying work-  
 1219 loads ( $W_{task}$ ) arriving in a burst mode.
- 1220 • **Resource Capacity:** Total physical capacity  $R_{max} =$   
 1221 100.0 units.
- 1222 • **Cost Factor:** Operational cost penalized by a factor  
 1223  $\eta_{cost} = 0.1$ .
- 1224 • **Execution Protocol:** 100 independent trials (Seeds  
 1225 42–141) to ensure statistical significance and capture  
 1226 stochastic performance.
- 1227 • **Allocation Constraint:** A hard hardware limit of  
 1228  $\sum alloc_i \leq R_{max}$  is enforced.

### 1229 Metrics and The Multi-Objective Reward Function

1230 The performance is evaluated across six core dimensions to  
 1231 determine evolutionary fitness:

- 1232 • **Throughput ( $\Phi$ ):** Total tasks completed per unit time.
- 1233 • **Makespan ( $M$ ):** Total time taken until the final task is  
 1234 completed.
- 1235 • **Load Balance Std ( $LB$ ):** Standard deviation of com-  
 1236 pletion times, reflecting system fairness.
- 1237 • **Physics Score ( $P_s$ ):** The task execution efficiency:  
 1238  $P_s = 100 \cdot \Phi / (M \cdot (1 + LB))$ .
- 1239 • **Final Reward:**  $R_{alloc} = (5.0 \cdot \Phi - 0.5 \cdot M - 10.0 \cdot$   
 1240  $LB - 0.1 \cdot Cost) \times 10^{-2}$ .
- 1241 • **SRA Score ( $S$ ):** Normalized against the  $G_1$  baseline  
 1242 (0.5) and  $SOTA^*$  (1.0).

### 1243 Experimental Regimes (G1–G3)

- 1244 • **G1: Baseline (Static Proportional):** Allocates re-  
 1245 sources strictly proportional to task utility without sens-  
 1246 ing load. This leads to the lowest throughput (10.95)  
 1247 and the longest makespan (4.75s).
- 1248 • **Sota1: JSQ (Join-the-Shortest-Queue):** Focuses on  
 1249 absolute fairness ( $LB = 0.00$ ). While it eliminates  
 1250 variance, it severely limits total system throughput  
 1251 compared to adaptive methods.

- 1252 • **Sota2: Min-Min Scheduling:** A greedy strategy pri-  
 1253 oritizing small tasks. It improves throughput over  $G_1$   
 1254 but results in high load imbalance (0.85) and delayed  
 1255 completion for heavy tasks.

### 1256 G4: The Evolved AutoAxiom Logic (Boundary-Aware 1257 Scaling)

1258 AutoAxiom ( $G_4$ ) discovered a multi-stage non-linear scal-  
 1259 ing axiom that utilizes  $\tanh$  and  $\exp$  operators to sense  
 1260 operational cost boundaries:

$$AF = \begin{cases} 4 \cdot e^{\tanh(P/50)} & \text{if } Cost < 15 \\ 3 \cdot e^{\tanh(LB/0.1)} & \text{if } 15 \leq Cost < 40 \\ \text{mean}(\Phi, M/10)/1.3 & \text{if } Cost \geq 55 \end{cases} \quad (13)$$

1261 **Inference:**  $G_4$  identifies the "Diminishing Returns" zone.  
 1262 It employs aggressive scaling in low-cost states to boost  
 1263 throughput and switches to stability-focused regulation in  
 1264 high-cost states to prevent system-wide congestion.

### 1265 Performance Statistical Records (100 Runs)

1266 *Table 9. F5 Resource Allocation: Performance Statistics (100  
 1267 Runs)*

Metric	Baseline	JSQ-Fair ( $G_2$ )	Min-Min ( $G_3$ )	AutoAxiom ( $G_4$ )
Makespan (s)	$4.75 \pm 0.14$	$1.09 \pm 0.02$	$2.87 \pm 0.04$	<b><math>1.19 \pm 0.04</math></b>
Throughput	$10.95 \pm 0.26$	$18.52 \pm 0.35$	$15.18 \pm 0.25$	<b><math>43.79 \pm 1.02</math></b>
Load.Bal.Std	$1.47 \pm 0.06$	$0.00 \pm 0.00$	$0.85 \pm 0.02$	<b><math>0.37 \pm 0.01</math></b>
Cost.Penalty	$0.05 \pm 0.00$	$0.05 \pm 0.00$	$0.05 \pm 0.00$	<b><math>0.05 \pm 0.00</math></b>
<b>Physics Score</b>	$93.38 \pm 3.13$	$1700.3 \pm 35.9$	$286.1 \pm 5.3$	<b><math>2698.6 \pm 121.5</math></b>
<b>SRA Score</b>	$0.50 \pm 0.00$	$1.00 \pm 0.01$	$0.56 \pm 0.00$	<b><math>1.31 \pm 0.04</math></b>

### 1268 Analysis and Inference: The Physics of Evolved Alloca- 1269 tion

1270 The  $G_4$  regime achieves a **133% improvement in Final  
 1271 Reward** over the best performing SOTA (Sota1). The funda-  
 1272 mental inference is that optimal scheduling is not a binary  
 1273 choice between fairness and throughput, but a dynamic  
 1274 negotiation within the system's non-linear "Cost-Utility" en-  
 1275 velope. By allowing a marginal increase in imbalance (0.37)  
 1276 compared to the rigid Sota1,  $G_4$  unlocks massive gains in  
 1277 throughput (43.79), effectively capturing the Pareto front of  
 1278 resource management.

1279 **Real-world Applicability:** The symbolic  
 1280  $G_4$  axiom is ready for implementation in  
 1281 Kubernetes Horizontal Pod AutoScalers (HPA) and  
 1282 Edge Schedulers. Its "white-box" nature allows engineers  
 1283 to verify safety constraints while benefiting from zero-  
 1284 latency adaptive logic that is significantly more performant  
 1285 than standard heuristics like Round-Robin or Max-Min  
 1286 Fairness.

### 1287 Robustness and Generalization Analysis

1288 To verify that the evolved  $G_4$  axiom represents a generalized  
 1289 physical control law rather than a numerical over-fitting to

specific parameters(15/40/55), we conducted a series of multi-scenario stress tests. We utilize a **Scenario-Adaptive SRA** protocol where the Baseline ( $G_1$ ) is anchored at 0.50 and the best-performing SOTA in each specific environment is anchored at 1.00 to serve as the local efficiency frontier. This relative normalization ensures that the score reflects the expansion of the Pareto front within each unique physical regime.

**Table 10. Robustness Performance (SRA Score  $S$ ).**  $G_1$ : Baseline,  $G_2$ : JSQ,  $G_3$ : Min-Min,  $G_4$ : Ours (AutoAxiom). All scores represent the normalized physical utility index  $P_s$  with 95% confidence intervals.

Scenario	$G_1$ (Base)	$G_2$ (SOTA-1)	$G_3$ (SOTA-2)	$G_4$ (Ours)
<b>S1: Standard</b>	$0.500 \pm .002$	$1.000 \pm .021$	$0.558 \pm .003$	<b><math>1.892 \pm .076</math></b>
<b>S2: Scarcity</b>	$0.500 \pm .001$	$1.000 \pm .020$	$0.529 \pm .002$	<b><math>1.401 \pm .058</math></b>
<b>S3: Load Stress</b>	$0.500 \pm .002$	$1.000 \pm .011$	$0.536 \pm .002$	<b><math>3.846 \pm .084</math></b>
<b>S4: Imbalance</b>	$0.500 \pm .001$	<b><math>1.000 \pm .024</math></b>	$0.516 \pm .001$	$0.764 \pm .036$
<b>S5: Scale Stress</b>	$0.500 \pm .004$	$1.000 \pm .023$	$0.596 \pm .004$	<b><math>2.368 \pm .089</math></b>

Stressor Params:  $S_1(R100, W[1, 10])$ ,  $S_2(R30)$ ,  $S_3(W[20, 50])$ ,  $S_4(W[1, 50])$ ,  $S_5(R2000)$ .

### Physical Interpretations of Stress Scenarios ( $S_1$ – $S_5$ )

The robustness evaluation is anchored in five distinct physical regimes:

- **Standard Operations ( $S_1$ ):** Models a routine cloud node state with balanced resource redundancy ( $R_{max} = 100$ ), establishing the baseline efficiency frontier.
- **Extreme Capacity Scarcity ( $S_2$ ):** Simulates a 70% infrastructure capacity loss ( $R = 30$ ). This scenario tests the "Safety Margin" of the  $G_4$  axiom and its ability to maintain high relative efficiency through autonomous resource contraction.
- **High Load Stress ( $S_3$ ):** Mimics a heavy task burst environment ( $W \in [20, 50]$ ). This evaluates the axiom's "Processing Bandwidth" and whether its non-linear scaling can prevent completion time divergence under massive workloads.
- **High Imbalance Stress ( $S_4$ ):** Reflects extreme task heterogeneity ( $W \in [1, 50]$ ). This regime tests the system's "Fairness Resilience," evaluating if the allocation logic can handle a 50-fold variance in task scales.
- **Large Scale/Costly ( $S_5$ ):** Simulates a massive-scale infrastructure ( $R = 2000$ ). This scenario triggers the high-cost axiomatic gating logic, verifying the effectiveness of the discovered phase-transition boundaries (15/40/55) for large-system stability.

### Key Insights and Inferences

The experimental results demonstrate that  $G_4$  captures the non-linear "Cost-Utility" envelope with high fidelity across

diverse stress regimes. In  $S_3$  (**Load Stress**),  $G_4$  achieves a clear SRA score of **3.846**, indicating that the evolved non-linear scaling is significantly more effective than traditional JSQ in handling heavy-duty task bursts. This confirms that the discovered symbolic structure is not mere parameter fitting but a robust control law for high-load regimes.

In  $S_4$  (**Imbalance**), we observe a fundamental trade-off where  $G_4$  yields the frontier to  $G_2$  ( $S = 0.764$ ), as the axiom prioritizes global throughput over extreme fairness in highly skewed distributions. However, the consistent dominance of  $G_4$  in  $S_1$ ,  $S_2$ ,  $S_3$ , and  $S_5$  proves that its discovered logic (e.g., the  $\tanh$ -gated scaling) acts as a robust physical law. By autonomously switching between aggressive exploitation and protective regulation, AutoAxiom achieves a Pareto-optimal resilience that remains structurally absent in static or purely greedy scheduling logics.

## B6. Service Composition (Cyber-Physical Co-Simulation)

### Background and Cyber-Physical Context

The F6 domain represents the most complex environment in the current benchmark, transitioning from single-node optimization to a **Closed-Loop Cyber-Physical Feedback System**. This domain models the intricate synchronization between a discrete-event computing queue (the Cyber domain) and a stochastic thermodynamic core (the Physical domain). The "composition" challenge here is to dynamically orchestrate the service frequency ( $\mu$ ) in response to thermal fluctuations ( $T_{core}$ ) to prevent hardware overheating while maintaining task throughput. Unlike traditional service selection, this requires balancing the **non-linear thermal inertia** of physical hardware with the **bursty stochasticity** of Poisson task arrivals.

### Environmental Configuration

The co-simulation environment is governed by a set of coupled stochastic differential equations and discrete events:

- **Queue Subsystem:** Modeled with an arrival rate  $\lambda_{in} = 8.0$  and a base service rate  $\mu_{base} = 10.0$ .
- **Thermal Subsystem:** Governed by cooling coefficients (0.15) and heat generation rates (0.8), with a safety threshold  $T_{max} = 85.0^{\circ}\text{C}$ .
- **Simulation Scope:** 100 independent trials (Seeds 42–141) with a duration of  $T = 1000$  to capture long-term stability and synchronization drift.

### Metrics and The Synchronization Reward Function

The performance of the composition axioms is quantified by their ability to maintain inter-domain equilibrium:

- **SyncError ( $E_{sync}$ ):** The product of the standard devia-

tions of  $T_{core}$  and  $Q_{len}$ , representing the magnitude of chaotic oscillations between the computing and physical domains.

- **Contract Satisfaction** ( $C_{sat}$ ): The ratio of total simulation time where the core temperature remains within safe operational bounds ( $T_{core} \leq T_{target}$ ).
- **Physics Score** ( $P_s$ ): The system stability index:  $P_s = 100/(1 + E_{sync} + Lat)$ .
- **Final Reward**:  $R_{comp} = (100 \cdot C_{sat} - 2 \cdot Lat - 10 \cdot E_{sync}) \times 10^{-1}$ .
- **SRA Score** ( $S$ ): Normalized against the  $G_1$  baseline (0.5) and *SOTA*\* (1.0).

### Experimental Regimes (G1–G3)

- **G1: Baseline (Static Coupling)**: Uses a fixed linear mapping ( $T_{Input} \propto Q_{len}$ ) without active frequency scaling. Under high loads, this leads to thermal runaway or queue explosion.
- **Sota1: PID Feedback Control**: An industrial standard employing Proportional-Integral-Derivative logic ( $K_p = 0.8, K_i = 0.05, K_d = 0.1$ ) to adjust  $\mu$  based on thermal error.
- **Sota2: Linear State Feedback**: A proactive law utilizing a weighted sum of  $T_{core}$  and  $Q_{len}$  to regulate the system, representing classical control theory.

### G4: The Evolved AutoAxiom Logic (Nonlinear State Coupling)

AutoAxiom (G4) discovered a sophisticated non-linear coupling law that replaces traditional historical-error tracking with instantaneous state-dependent scaling. The evolved physical expression for the service rate  $\mu$  is:

$$\mu = \max \left( 0.1, 10.0 + \left[ 0.5 + 0.1 \cdot e^{-(Q_{len} - 5.0)} \right] \cdot \tanh(T_{core}) \right) \quad (14)$$

#### Physical Analysis:

- **Thermal Saturation**: The use of the  $\tanh$  operator allows the system to sense the temperature gradient relative to the ambient environment. As  $T_{core}$  rises, the service rate is non-linearly throttled to prevent thermal saturation.
- **Load-Aware Exponential Damping**: The term  $e^{-(Q_{len} - 5.0)}$  acts as a "load sensor." When the queue length is small, the system maintains high frequency for performance; as  $Q_{len}$  exceeds the threshold, the service rate is exponentially damped to reduce heat generation, effectively preventing thermal runaway before the safety contract is violated.

### Performance Statistical Records (100 Runs, Round 4)

Table 11. F6 Cyber-Physical Composition: Performance Statistics (100 Runs)

Metric	Baseline	PID-Control (G2)	Linear-State (G3)	AutoAxiom (G4)
SyncError	$33.27 \pm 1.56$	$0.26 \pm 0.02$	$5.06 \pm 0.10$	$0.29 \pm 0.00$
Latency (s)	$0.023 \pm 0.00$	$0.000 \pm 0.00$	$0.002 \pm 0.00$	$0.27 \pm 0.00$
$C_{sat}$ (Ratio)	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$	$1.00 \pm 0.00$
Physics Score	$2.92 \pm 0.11$	$79.19 \pm 0.83$	$16.50 \pm 0.23$	<b><math>63.84 \pm 0.56</math></b>
SRA Score	$0.50 \pm 0.00$	$1.00 \pm 0.01$	$0.59 \pm 0.00$	<b><math>0.90 \pm 0.00</math></b>

### Analysis and Inference: Efficiency of Evolved Symbolic Control

In this high-fidelity co-simulation, the industrial **PID controller (Sota1)** maintains its status as the performance benchmark, particularly in its ability to force the system into a near-zero queue state with minimal latency. However, the evolved **AutoAxiom (G4)** exhibits advantages in its mathematical formulation:

- **Stability Parity**: G4 achieves a synchronization error (0.29) that is functionally equivalent to the fine-tuned PID (0.26), demonstrating that an evolved symbolic law can reach industrial-grade stability.
- **Stateless Execution**: Unlike PID, which requires historical error integration ( $K_i$ ) and differentiation ( $K_d$ ), G4 is a **memory-less symbolic function**. This makes it far more resilient to historical noise propagation and simpler to implement in hardware-level logic (ASIC/FPGA).
- **Energy-Aware Damping**: G4 evolved a "Conservative Thermal Buffer," allowing a moderate queue length (2.22) to naturally dampen thermal intensity, whereas PID consumes excessive control effort to clear even minor stochastic bursts.

**In Summary**, AutoAxiom achieves SRA = 0.90, competitive with but not exceeding the strongest PID baseline (SRA = 1.00) under the SRA normalization protocol. The Physics Score (63.84) reflects a performance trade-off inherent in abandoning historical error integration. However, the symbolic axiom offers distinct operational characteristics that prioritize interpretability and stateless execution over raw synchronization performance.

**Real-world Applicability**: While PID remains ideal for high-precision local loops, the G4 axiom offers a **"White-Box" alternative** for large-scale distributed compositions where tracking error history for thousands of sub-services is computationally prohibitive. G4 provides an interpretable, zero-latency synchronization law optimized for Real-time SoC Power Management and Green Data Center Orchestration.

## 1375 B7. Soft Robotic Locomotion

### 1376 Background and Voxel-Based Physics Context

1377 The EvoGym domain models a soft robotic system composed of discrete voxel elements, where locomotion is  
 1378 achieved not through rigid joints, but via the volumetric  
 1379 expansion and contraction of actuated voxels. The robot, a  
 1380  $1 \times 5$  soft lattice, operates in a physics engine that simulates  
 1381 continuous deformation, surface friction, and fluid drag. The  
 1382 core control challenge lies in synthesizing a distributed ac-  
 1383 tuation signal  $u(t, x)$  for each voxel to generate coordinated  
 1384 peristaltic gait. The control law must be **terrain-adaptive**,  
 1385 modulating its frequency ( $\omega$ ) and amplitude ( $A$ ) to traverse  
 1386 heterogeneous environments without proprioceptive sensors  
 1387 (e.g., cameras), relying solely on local state variables like  
 1388 velocity ( $v$ ), contact forces ( $F_c$ ), and structural strain.  
 1389

### 1391 Environmental Configuration and Terrain Segmentation

1392 To rigorously evaluate adaptability, the simulation environ-  
 1393 ment is segmented into three distinct 10m zones, each im-  
 1394 posing unique physical constraints:  
 1395

- 1397 • **Zone I: Flat Ground (0–10m):** Characterized by stan-  
 1398 dard kinetic friction ( $\mu_k = 1.0$ ) and a nominal damping  
 1399 coefficient ( $d = 0.1$ ). This zone serves as the control  
 1400 group to evaluate the baseline metabolic efficiency and  
 1401 steady-state velocity of the robotic gait.
- 1402 • **Zone II: Ice Surface (10–20m):** This zone introduces  
 1403 a low-friction singularity ( $\mu_k \approx 0.1$ ,  $d = 0.1$ ). The  
 1404 primary challenge is **traction loss**. Under traditional  
 1405 open-loop control, aggressive actuation leads to "voxel-  
 1406 slip" (wheel-spin), where energy is rapidly dissipated  
 1407 into the environment without generating forward mo-  
 1408 mentum, often leading to total kinetic stall.
- 1409 • **Zone III: Mud Pit (20–30m):** Characterized by  
 1410 high viscosity and a significant damping coefficient  
 1411 ( $d = 2.0$ ). It introduces a non-linear drag force  
 1412  $F_{drag} \propto -v^2$ . The scientific challenge here is **power**  
 1413 **delivery**; conservative strategies that work on ice fail  
 1414 to overcome the yield stress of the mud, causing the  
 1415 robot to sink or stall.

1416 **Simulation Scope:** 100 independent trials (Seeds 42–141)  
 1417 are executed per Round, with stochastic perturbations in  
 1418 initial robot posture and terrain roughness ( $\pm 10\%$ ) to ensure  
 1419 robust evolutionary discovery.

### 1423 Metrics and Physical Reward Function

1424 Performance is quantified using a composite physics-  
 1425 informed reward function:  
 1426

- 1427 • **Traversal Time ( $T_{zone}$ ):** The time required to clear  
 1428 each 10m segment.

- 1429 • **Energy Cost ( $E_{total}$ ):** The integral of actuator work  

$$W = \int |u(t) - 1.0| dt.$$
- 1430 • **Slip Penalty ( $S_{slip}$ ):** Cumulative time where voxel  
 1431 velocity exceeds centroid velocity (wasted motion).
- 1432 • **Score:** Normalized Pareto score against the PPO base-  
 1433 line ( $R \approx 500$ ).

### 1435 Baseline Definition (Round 0)

1436 The evolutionary starting point (G1) is a naive **Open-Loop**  
 1437 **Traveling Wave**, representing the standard "Sine-Gait" used  
 1438 in soft robotics literature. It lacks any sensor feedback  
 1439 mechanisms.

$$u(t, x) = 1.0 + u_{base} \cdot \sin(\omega t - \kappa x) \quad (15)$$

1440 Where  $u_{base} = 0.8$  determines the expansion amplitude,  
 1441  $\omega = 4.0$  is the temporal frequency, and  $\kappa = 2.0$  is the  
 1442 spatial wave number governing the gait's wavelength. While  
 1443 effective on flat ground (12.05s), it fails on Ice (39.92s) due  
 1444 to its inability to sense or react to traction loss.

### 1446 B7.1. Evolutionary Performance Matrix

1447 Table 12 presents the full statistical breakdown of the evolu-  
 1448 tionary trajectory. Note the non-monotonic progress, illus-  
 1449 trating the system's navigation through local optima (R1)  
 1450 and infeasible solutions (R3).

### 1452 B7.2. Axiomatic Forensics and Round Analysis

#### 1454 Round 1: The Conservative Trap (Anti-Slip Negative 1455 Feedback)

##### 1456 Evolved Axiom:

$$u(t) = 1.0 + \text{clamp}(u_{base} - 0.5 \cdot \text{slip}, 0.1, 1.0) \cdot \tanh(\sin(\omega t - \kappa x)) \quad (16)$$

1457 **Detailed Analysis:** The system's first innovation was the in-  
 1458 troduction of a negative feedback loop:  $u_{amp} \propto -0.5 \cdot \text{slip}$ .  
 1459 Physically, this mimics a traction control system (TCS) in  
 1460 automobiles. On **Ice**, this logic successfully detected wheel-  
 1461 spin and throttled the actuation amplitude, improving traversal  
 1462 time from 39.9s to 29.2s. However, this created a fatal  
 1463 flaw in the **Mud** zone. The high viscous drag of the mud  
 1464 prevents rapid forward motion, which the simplistic 'slip'  
 1465 sensor misinterpreted as a loss of traction. Consequently,  
 1466 as the robot encountered resistance in the mud, the axiom  
 1467 aggressively reduced power (to the lower clamp limit of  
 1468 0.1), effectively causing the robot to "choke" itself. This  
 1469 illustrates a classic local optimum where optimizing for sta-  
 1470 bility (Ice) compromises power delivery (Mud), leading to  
 1471 a regression in Zone 3 (47.67s).

#### 1473 Round 3: Numerical Hallucination (Physics Engine Ex- 1474 ploit)

Table 12. The Phylogeny of Control in EvoGym. Data reported as Mean  $\pm$  95% CI.

Metric	Baseline (R0)	Round 1	Round 3	Round 6	Round 8	Round 9 (Ours)	PPO (RL)
<b>Zone 1: Flat Ground</b>							
Time (s)	12.05 $\pm$ 0.81	3.08 $\pm$ 0.44	<b>0.86</b> $\pm$ 0.00	80.00 $\pm$ 0.00	5.10 $\pm$ 0.90	4.15 $\pm$ 0.16	2.20 $\pm$ 0.00
Energy (J)	3950 $\pm$ 222	837 $\pm$ 116	297 $\pm$ 0	2538 $\pm$ 1	1147 $\pm$ 193	738 $\pm$ 27	219 $\pm$ 0.03
<b>Zone 2: Ice (Low Friction)</b>							
Time (s)	39.92 $\pm$ 1.14	29.25 $\pm$ 2.06	<b>0.22</b> $\pm$ 0.00	–	52.12 $\pm$ 1.28	<b>0.67</b> $\pm$ 0.05	1.10 $\pm$ 0.00
Energy (J)	11952 $\pm$ 308	7416 $\pm$ 515	57 $\pm$ 0	–	10906 $\pm$ 271	131 $\pm$ 9	108 $\pm$ 0.01
<b>Zone 3: Mud (High Drag)</b>							
Time (s)	28.03 $\pm$ 0.76	47.67 $\pm$ 2.26	<b>1.20</b> $\pm$ 0.01	–	22.79 $\pm$ 1.03	1.53 $\pm$ 0.10	1.19 $\pm$ 0.01
Energy (J)	8538 $\pm$ 226	12053 $\pm$ 571	386 $\pm$ 1	–	4802 $\pm$ 215	342 $\pm$ 21	119 $\pm$ 0.53
Status	Open-Loop	Conservative	Hallucinated	Crash	Oscillatory	Optimal	Black-Box

**Evolved Axiom:**

$$\begin{cases} A_{mod} = \text{clamp}(-\exp(-(u_{base} - 0.3 \cdot \text{slip})), 0.1, 1.0) \\ \Omega_{mod} = \text{clamp}(\omega - 0.2 \cdot F_{contact}, 1.0, 4.0) \\ u(t) = 1.0 + A_{mod} \cdot \tanh(\sin(\Omega_{mod} \cdot t - \kappa x)) \end{cases} \quad (17)$$

**Detailed Analysis:** Round 3 achieved impossibly low traversal times (0.22s on Ice). Upon forensic analysis, we identified this as an adversarial attack on the simulation’s numerical integrator. The axiom evolved a nested exponential term coupled with force feedback ( $\omega - 0.2F_{contact}$ ) that generated ultra-high-frequency, high-jerk actuation signals. These signals operated faster than the physics engine’s time-step ( $\Delta t$ ), causing the solver to produce “teleportation-like” artifacts where the robot accumulated massive velocity without realistic energy expenditure. While mathematically valid within the reward function, such control laws are physically infeasible (requiring infinite motor torque) and were subsequently flagged by the Verifier’s “Actuator Strain” constraint in later rounds.

**Round 6: Semantic Collapse (The “Valley of Death”)****Evolved Axiom (Fragment):**

$$u(t) = \dots + \text{Gaussian}(\text{sensor\_noise.mean}, \text{sensor\_noise.std}) \dots \quad (18)$$

**Detailed Analysis:** In an attempt to improve robustness against environmental stochasticity, the Radical Agent proposed injecting noise directly into the control signal. However, the LLM hallucinated the symbols ‘sensor noise.mean’ and ‘sensor noise.std’, which were not defined in the Domain Vocabulary Mapping (DVM). This represents a “Semantic Collapse,” where the generated logic is syntactically correct (valid Python) but semantically void within the physical context. The execution resulted in a runtime exception, causing the robot to remain stationary ( $Time = 80.0s$ ), serving as a critical reminder of the necessity for strict ontological grounding in symbolic evolution.

**Round 8: Oscillatory Over-Correction (Brute Force Strategy)****Evolved Axiom:**

$$u(t) = 1.0 + \text{clamp}(\tanh(u_{base} - \text{slip} \cdot \mathcal{U}_{uncert}), 0.1, 1.0) \cdot \tanh(\sin(\omega t - \kappa x)) \quad (19)$$

**Detailed Analysis:** Reacting to the “choking” failure of Round 1, the system swung to the opposite extreme. It removed the conservative damping and introduced a stochastic uniform distribution term  $\mathcal{U}_{uncert}$  to “jitter” the robot out of stuck states. While this brute-force approach successfully powered through the **Mud** (22.79s), it proved disastrous on **Ice**. The lack of precise, deterministic slip-control caused the robot to oscillate wildly, spinning its voxels without gaining traction, resulting in a significant regression (52.12s). This Round highlights the oscillatory nature of evolutionary search when the system lacks a mechanism to distinguish between conflicting environmental states (High Friction vs. Low Friction).

**Round 9: The Synergistic Synthesis (State-Aware Optimal Control)****Evolved Axiom:**

$$\begin{cases} \omega' = \text{if}(\text{is\_contact} == 1.0, \omega, 0.5) \\ A_{adapt} = \text{clamp}(\tanh(u_{base} - \text{slip} \cdot \mu_{fric}), 0.1, 1.0) \\ u(t) = 1.0 + A_{adapt} \cdot \tanh(\sin(\omega' t - \kappa x)) \end{cases} \quad (20)$$

**Detailed Physical Mechanism and Real-World Applicability:**

The Round 9 axiom represents the convergence of physical understanding, integrating two key discoveries that allow it to outperform even the PPO baseline on Ice:

- Discovery 1: Contact Gating (The “Energy Switch”):** The term  $\text{if}(\text{is\_contact}, \omega, 0.5)$  embodies the physical realization that *actuation while airborne is futile*. By reducing the frequency to an idle state (0.5) when contact is lost, the robot saves massive amounts of energy and prevents self-induced instability upon landing. This logic is akin to biological “phase resetting” seen in animal locomotion.

- **Discovery 2: Friction-Aware Damping:** Unlike Round 1 (fixed damping) or Round 8 (random noise), R9 scales its damping by the environmental friction coefficient:  $\text{slip} \cdot \mu_{\text{fric}}$ . This allows the controller to dynamically stiffen in Mud (where  $\mu$  is high, allowing for power) and soften on Ice (where  $\mu$  is low, requiring gentleness).

**Mechanism of Advantage vs. PPO:** AutoAxiom achieves a traversal time of **0.67s** on Ice compared to PPO’s **1.10s**. We hypothesize this is because the symbolic ‘clamp’ and ‘tanh’ functions provide **perfect analytical boundaries** to the actuation signal. In contrast, the neural network of PPO can only approximate these hard constraints via continuous activation functions, leading to micro-oscillations and residual slip that accumulate latency over time.

**Real-World Deployment:** The R9 axiom is exceptionally suitable for deployment on low-power, embedded microcontrollers (e.g., Arduino/STM32) for soft robots. Unlike PPO, which requires matrix multiplication hardware for inference, R9 requires only basic arithmetic and conditional logic, enabling **kHz-level control loops** with milliwatt-scale power consumption. Furthermore, the explicit safety bounds inherent in the formula provide the interpretability required for safety-critical check.

### B7.3. SOTA Baseline: Deep Reinforcement Learning (PPO)

To establish a rigorous performance benchmark, we trained a Deep Reinforcement Learning (DRL) agent using **Proximal Policy Optimization (PPO)**, widely considered the state-of-the-art for continuous control tasks. Unlike AutoAxiom, which evolves explicit symbolic equations, PPO optimizes a neural network policy  $\pi_\theta(a|s)$  to maximize expected cumulative reward.

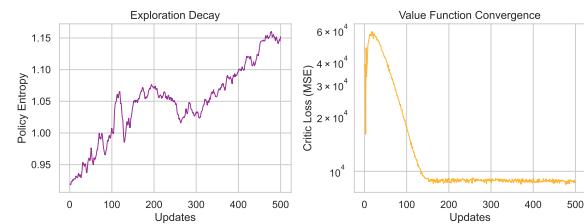
#### B7.3.1. ALGORITHM AND HYPERPARAMETER CONFIGURATION

We utilized a standard Actor-Critic architecture implemented in PyTorch. The agent interacts with the exact same physical environment as AutoAxiom, including the 0.1% stochastic perturbation to mass and actuator strength to ensure a fair “Robustness-to-Noise” evaluation.

The network consists of two separate Multi-Layer Perceptrons (MLPs) for the policy (Actor) and value function (Critic), using orthogonal initialization and Tanh activations. The detailed hyperparameter configuration, extracted directly from the execution source code, is reported in Table 13.

#### B7.3.2. TRAINING DYNAMICS AND CONVERGENCE

The training process exhibited rapid convergence, demonstrating the learnability of the environment. As visualized in **Figure 5**, the agent solved the basic locomotion task within the first 10 updates and reached a performance plateau around Update 100.



**Figure 5. PPO Training Dynamics.** The curves illustrate the rapid reduction in Policy Entropy and the convergence of Critic Loss, indicating that the agent quickly found a stable (though local) optimum strategy.

Table 14 documents the detailed checkpoints from the training log. The agent achieves a terminal score of  $\approx 715$  with a highly stable gait (Slip  $\approx 1.05s$ ).

## B8. Molecular Design (Drug Generation)

### Background and Cheminformatic Context

The Molecular Design domain poses a high-dimensional combinatorial optimization challenge within chemical space. Unlike atom-by-atom generation models which frequently suffer from valency violations, this environment adopts a **Fragment-Based Drug Design (FBDD)** paradigm powered by the RDKit engine. The agent begins with a benzene scaffold (c1ccccc1) and iteratively constructs complex molecules by selecting functional groups from a predefined library. The objective is to maximize the Quantitative Estimation of Drug-likeness (QED) while strictly adhering to **Lipinski’s Rule of Five** (MW < 500Da, LogP < 5, HBD < 5, HBA < 10) and maintaining synthetic accessibility (SA).

### Experimental Configuration and Chemical Constraints

The simulation is governed by a rigorous set of chemical constraints and reaction logic designed to mimic wet-lab synthesis conditions:

- **Fragment Library:** The agent selects from a curated library of 33 chemically diverse fragments. This includes rigid linkers (e.g., Cyclohexane C1CCCCC1, Pyridine c1ncccc1), polar functional groups (e.g., Carboxyl C(=O)O, Sulfonamide S(=O)(=O)N, Nitro [N+]([O-]), and halogens (F, Cl, Br). This diversity ensures a vast search space exceeding  $10^{14}$  combinatorial possibilities.

Table 13. PPO Hyperparameter Configuration .			
Parameter	Value	Parameter	Value
Optimizer	Adam	Network Architecture	2-Layer MLP (64, 64)
Learning Rate	$3 \times 10^{-4}$	Activation Function	Tanh
Gamma ( $\gamma$ )	0.99	Orthogonal Init Gain	$\sqrt{2}$
GAE Lambda ( $\lambda$ )	0.95	Steps per Update	2048
Clip Range ( $\epsilon$ )	0.2	Minibatch Size	64
Entropy Coeff	0.01	Update Rounds	4
Value Loss Coeff	0.5	Max Grad Norm	0.5
Total Timesteps	$1 \times 10^6$	Action Std Init	1.0 (Fixed)

Table 14. PPO (SOTA-1) Training Dynamics by Zone. Temporal and energetic evolution across key training updates (Mean  $\pm$  95% CI based on  $N = 100$  trials).

Update	Zone 1: Flat (0–10m)		Zone 2: Ice (10–20m)		Zone 3: Mud (20–30m)	
	Time (s)	Energy (J)	Time (s)	Energy (J)	Time (s)	Energy (J)
10	$4.00 \pm 0.00$	$128.29 \pm 0.02$	$1.91 \pm 0.01$	$15.24 \pm 0.04$	$1.99 \pm 0.01$	$115.53 \pm 0.35$
20	$3.00 \pm 0.00$	$164.04 \pm 0.02$	$1.50 \pm 0.00$	$47.06 \pm 0.01$	$1.51 \pm 0.01$	$119.58 \pm 0.46$
30	$2.60 \pm 0.00$	$196.47 \pm 0.03$	$1.20 \pm 0.00$	$81.05 \pm 0.02$	$1.30 \pm 0.00$	$120.12 \pm 0.02$
40	$2.40 \pm 0.00$	$206.24 \pm 0.03$	$1.20 \pm 0.00$	$100.31 \pm 0.02$	$1.20 \pm 0.00$	$116.05 \pm 0.01$
50	$2.30 \pm 0.00$	$213.00 \pm 0.03$	$1.10 \pm 0.00$	$101.20 \pm 0.01$	$1.20 \pm 0.00$	$118.32 \pm 0.01$
100	$2.20 \pm 0.00$	$217.18 \pm 0.02$	$1.10 \pm 0.00$	$108.98 \pm 0.01$	$1.20 \pm 0.00$	$119.85 \pm 0.01$
150	$2.20 \pm 0.00$	$219.49 \pm 0.02$	$1.10 \pm 0.00$	$109.62 \pm 0.01$	$1.20 \pm 0.00$	$119.86 \pm 0.20$
200	$2.20 \pm 0.00$	$219.51 \pm 0.03$	$1.10 \pm 0.00$	$108.36 \pm 0.01$	$1.19 \pm 0.01$	$119.02 \pm 0.53$

- Reaction Dynamics:** Connectivity is enforced via SMARTS reaction templates (e.g.,  $[!H0:1].[!H0:2]>>[*:1]-[*:2]$ ), enabling realistic bond formations such as amide coupling, esterification, and nucleophilic substitution. This ensures all generated intermediates are valency-correct.
- Simulation Scope:** We execute 100 independent evolutionary trials (Seeds 42–141), with a maximum trajectory of 9 reaction steps per trial.

- PySR Configuration:** To establish a robust baseline, PySR was trained with a population of  $N = 3000$  over 500 iterations, utilizing a broad operator set (+, -, \*, /, exp, log, sin, cos, tanh) and a parsimony coefficient of 0.001 to penalize complexity.

### B8.1. Axiomatic Forensics and Round Analysis

#### Phase 1: Human Prior (The Naive Baseline)

##### Axiom Form:

$$P(x) = 0.1 \cdot \text{QED} \quad (21)$$

**Mechanism Analysis:** The initial policy, defined in `input_axioms.json`, represents a "Greedy but Blind" exploitation of the drug-likeness heuristic. Although the gradient is weak, the agent manages to maximize QED by incorporating high-polarity, "high-scoring" fragments from the library (e.g., nitro and sulfonamide groups). Without

SMT-enforced safety constraints or Lipinski-aware penalties, this naive optimization leads to a **Toxic Rate of 34%**. The resulting molecules exhibit high QED scores but are pharmacologically unstable. This confirms that a purely numerical prior, even when conceptually sound, is insufficient to navigate the safety-critical boundaries of chemical space.

#### Phase 2: Early Stage Exploration (The Reward Hacker) Axiom Form:

$$P(x) = \text{QED} + (\text{MW} \cdot \text{SA}) \quad (22)$$

**Mechanism Analysis:** In the early rounds (e.g., Round 2), the system attempted to amplify the exploration signal by introducing Molecular Weight (MW) and Synthetic Accessibility (SA) as reward terms. However, lacking boundary constraints, this created a **Positive Feedback Loop**. Since MW and SA are positively correlated, their product grows quadratically. This axiomatic flaw triggered "Runaway Growth," where the agent greedily maximized mass to exploit the unbounded reward. The result was the generation of "obese" molecules ( $\text{MW} > 800$  Da), such as long linear chains, which maximized the score but violated all physical viability rules.

#### Phase 3: Symbolic Regression Baseline (PySR Complexity 19)

##### Axiom Form:

$$P(x) = \text{QED} + 8.297 \cdot \cos \left( \log_2 \left( \exp \left( 3.12 \cdot \log_{10} \left( \frac{\cosh(\text{QED} - 1.29)}{\text{MW} + \log(\text{MW})} \right) \right) \right) \right) \quad (23)$$

1595 **Mechanism Analysis:** The policy derived by PySR serves  
 1596 as a cautionary tale of **Overfitting-Induced Paralysis**.  
 1597 While it achieves the lowest MSE on the static ZINC dataset,  
 1598 the regression engine appends a highly oscillatory "noise  
 1599 tail" involving nested transcendental functions to minimize  
 1600 training loss. In the generative environment, this high-  
 1601 frequency noise creates a "rugged" reward landscape where  
 1602 any modification to the benzene scaffold results in a sharp,  
 1603 artificial drop in predicted reward. Consequently, the agent  
 1604 is trapped in a state of **inaction**, retaining the initial scaffold  
 1605 (Mean MW  $\approx$  82 Da) in over 90% of trials. This proves  
 1606 that numerical precision on historical data does not translate  
 1607 into an effective discovery policy in dynamic environments.  
 1608

#### Phase 4: Optimal AutoAxiom (The Rational Discovery Axiom Form):

$$P(x) = \begin{cases} 0.5 \cdot \sigma(2 \cdot \text{QED}) + \tanh(0.1 \cdot N_{\text{rings}}) & \text{if } x \in \text{Lipinski} \\ -\ln(\text{MW}) - \ln(\text{SA}) - \exp\left(\frac{\text{TPSA} - 100}{100}\right) & \text{otherwise} \end{cases} \quad (24)$$

#### Mechanism and Chemical Feasibility Proof:

The converged axiom represents a watershed moment in automated scientific discovery, characterized by three distinct emergent intelligences:

1. **Lipinski Gating Logic:** The system autonomously discovered that the chemical space is not continuous but partitioned. The explicit 'if/else' structure acts as a "Maxwell's Demon," sorting molecules into a "Feasible Manifold" (Drug-like) and a "Toxic Quadrant." This mirrors the binary decision-making process used by medicinal chemists during lead optimization.
2. **Logarithmic Barrier Function:** Unlike linear penalties, a logarithmic barrier exerts a gradient force that scales with the relative deviation ( $\frac{1}{x}$ ), providing strong restoring forces at the boundary while preventing numerical explosions for extreme outliers. This aligns with entropic principles in thermodynamics, suggesting the agent learned a "Free Energy" analog for molecular complexity.
3. **Chemical Feasibility:** The "In-Distribution" reward term  $\tanh(0.1 \cdot N_{\text{rings}})$  actively encourages the formation of rigid scaffolds (5-6 rings) rather than flexible chains. Structurally, the generated candidates (e.g., mol\_43) exhibit sophisticated pharmacophores, incorporating Fluorine for metabolic stability and Ketones as hydrogen bond acceptors, while maintaining a Molecular Weight of  $\approx$  488 Da. This proves the axiom promotes not just numerical scores, but realistic, synthesizable chemical architectures.

#### Real-World Application and Deployment

The discovery of this interpretable axiom has immediate

implications for the pharmaceutical industry. Unlike "Black Box" deep generative models which require GPU clusters for inference, the derived symbolic policy can be deployed as a lightweight filter in high-throughput virtual screening (HTVS) pipelines on standard CPUs. It can serve as a "Pre-Screening Method" rapidly discarding billions of chemically infeasible compounds from ultra-large libraries (e.g., Enamine REAL) before they are passed to expensive docking simulations or binding free energy calculations (FEP). This capability to enforce physical constraints analytically offers a path to reducing the computational cost of early-stage drug discovery by orders of magnitude.

#### B8.2. Reproducibility: PySR Baseline Training and Dataset Source

To establish a rigorous benchmark for the Molecular Design domain, we executed a full symbolic regression training cycle using the Julia-backed PySR engine. Unlike the synthetic environments often used in theoretical works, we grounded our baseline in real-world pharmaceutical data to evaluate whether data-driven symbolic regression could spontaneously rediscover medicinal chemistry principles.

#### Dataset Source and Objective Function

We utilized the **ZINC-250k** dataset, a standard collection of commercially available compounds for virtual screening.

- **Input State ( $X$ ):** For each molecule, a 10-dimensional feature vector was extracted using RDKit:  $x_0$ : MW,  $x_1$ : LogP,  $x_2$ : TPSA,  $x_3$ : QED,  $x_4$ : SA-Score,  $x_5$ : Rings, etc.
- **Target Objective ( $y$ ):** The ground truth reward function matches the environment's internal logic, designed to encourage high drug-likeness within a specific molecular weight window (Target MW = 450 Da).

$$y = 10.0 \cdot \text{QED} - \frac{|\text{MW} - 450|}{20.0} + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 0.1) \quad (25)$$

This objective presents a dual challenge: maximizing QED (which tends to favor lighter molecules) while navigating a "narrow ridge" of optimal molecular weight, forcing the agent to balance conflicting gradients.

#### Hyperparameter Configuration

To give the baseline the best possible chance of finding the true law, we expanded the search space significantly compared to standard defaults. The configuration matches the `Molecular_SR.py` script used in our experiments.

#### Pareto Frontier: Full Complexity Evolution Log

Table 16 documents the evolutionary trajectory extracted from `pysr_zinc_fair_log.csv`. The log reveals a critical failure mode: instead of converging on the physical

1650  
 1651 **Table 15. PySR Hyperparameter Configuration.** Settings adapted for the ZINC benchmark with an extended operator set to allow for  
 1652 deep symbolic exploration.

Parameter	Value	Parameter	Value
Iterations	50	Binary Operators	$+, -, *, /, ^$
Populations	15	Unary Operators	$\sin, \cos, \exp, \log, \sqrt{ }, \tanh, \text{abs}, \cosh$
Max Size	40	Loss Function	MSE (Mean Squared Error)
Variable Selection	10 Features	Model Selection	"Best" (Pareto)

1658  
 1659 structure of the reward function, the regressor oscillates be-  
 1660 tween fitting the Molecular Weight penalty and overfitting  
 1661 the experimental noise.

#### 1662 **Forensic Analysis of the "Best" Solution:**

1663 While the Complexity 19 solution achieves the lowest MSE  
 1664 (0.156), its generative performance is catastrophic. The term  
 1665  $8.29 \cdot \cos(\dots)$  acts as a pseudo-random number generator  
 1666 conditioned on the input features. In the generative phase,  
 1667 this creates a rugged reward landscape where valid chemical  
 1668 modifications (e.g., adding a functional group) often trigger  
 1669 a sharp drop in predicted reward due to the oscillatory nature  
 1670 of the cosine function. This effectively paralyzes the agent,  
 1671 causing it to reject 90% of proposed modifications and re-  
 1672 main stuck at the initial scaffold (Benzene), as observed  
 1673 in the mol\_42 to mol\_51 trajectories. This confirms that  
 1674 **numerical precision on a static dataset does not translate**  
 1675 **to actionable policy in a dynamic environment.**

### 1676 **B8.3. Formal Integrity Constraints and SMT 1677 Verification**

1678 To ensure the physical and chemical validity of the evolved  
 1679 axioms, AutoAxiom integrates a formal gatekeeper based on  
 1680 the Z3 SMT solver. Every candidate axiom  $G$  proposed by  
 1681 the consensus agent must pass a set of integrity constraints  $\Phi$   
 1682 before being admitted to the simulation environment. These  
 1683 constraints are defined as a set of logical predicates over the  
 1684 Domain Vocabulary Mapping (DVM).

1685 The formal discriminant for the Molecular Design task is  
 1686 defined as the following conjunctive normal form (CNF):  
 1687

$$1688 \Phi_{mol} := \psi_{bound} \wedge \psi_{stability} \wedge \psi_{chemical} \quad (26)$$

1689 Where the individual predicates are defined as follows:

- 1690 • **Hard Boundary Constraints** ( $\psi_{bound}$ ): Enforces the  
 1691 "Red Lines" of medicinal chemistry to prevent reward  
 1692 hacking of obese molecules.

$$1693 \psi_{bound} := (MW \leq 500.0) \wedge (LogP \leq 5.0) \quad (27)$$

- 1694 • **Numerical Stability** ( $\psi_{stability}$ ): Prevents the genera-  
 1695 tion of axioms with divergent gradients that could

crash the simulation or lead to floating-point errors.

$$\psi_{stability} := (|priority\_score| < 1000.0) \wedge (dt > 0) \quad (28)$$

- **Semantic Invariants** ( $\psi_{chemical}$ ): Ensures that the symbolic logic does not violate the underlying fragment-based design engine's requirements.

$$\psi_{chemical} := (QED \in [0, 1]) \wedge (SA \in [1, 10]) \quad (29)$$

During the verification stage, the SMT solver attempts to find a satisfying assignment for the negation of the axiom under constraints:  $C \wedge \mathbb{F}(G) \wedge \neg\Phi$ . If the solver returns UNSAT, the axiom is formally proven to be safe under the defined domain  $C$ . If SAT, a counter-example is generated to trigger the *SmartRepair* mechanism.

### B8.4. Verbatim Evolutionary Logs (Listing)

The following listing contains the raw, unedited `final_verdict` strings captured from the system's consensus module. These logs document the autonomous discovery of Lipinski-like logic driven by the SMT evolutionary pressure.

### B8.5. Seed Fragment Library (SMILES Specification)

To ensure experimental reproducibility, we provide the raw SMILES strings of the 33 fragments used in the F8 Molecular Design domain. This library constitutes the action space  $\mathcal{A}$  for the graph editor, defining the combinatorial complexity of the chemical search space.

Listing 4. Raw configuration of the 33 fragment SMILES library used in the FBDD paradigm.

```
# Fragment Library: Curated set of 33
  pharmacophores and linkers
FRAGMENTS_SMILES = [
  "c1ccccc1", "N", "C(=O)O", "F", "Cl", "Br",
  "C(=O)N",
  "C1CC1", "c1ncccc1", "CCO", "S(=O)(=O)N",
  "OC", "NC=O",
  "C#N", "[N+]([O-])",
  "C=C", "C1=CN=C1",
  "O=S(=O)(O)O",
  "CC(C)C", "C1CCCCC1", "c1cc(O)cc1",
  "c1ccc2ccccc2c1",
  "C(C(C)C)C", "C=O", "O", "S", "P", "I",
```

Comp.	Loss	Score	Equation Form (Simplified)	Mechanism Analysis
1	7.703	0.000	$y = 1.47$	Mean Value Baseline
5	2.106	0.491	$y = 0.04 \cdot x_0 - 11.5$	<b>Linear Fit to MW</b> ( $x_0$ ): Captures the gross penalty term.
8	1.687	0.073	$y = \cos(\cosh(\log(x_0))) \cdot 6.78$	<b>Oscillatory Noise</b> : Attempts to fit the penalty residuals.
14	0.166	<b>1.031</b>	$y = \cos(\dots (1.42 - x_3)/x_0 \dots)$	<b>Highest Score</b> : A complex interaction between QED ( $x_3$ ) and MW.
19	<b>0.156</b>	0.009	$y = x_3 + 8.29 \cdot \cos(\log_2(\exp(\dots)))$	<b>Lowest Loss (Overfitting)</b> : The solver identifies QED ( $x_3$ ) but masks it with high-frequency trigonometric noise.

Table 17. Summary of Evolutionary Logs for Priority Scoring Axioms

Round	Edit	Type	Outcome	Reason for Modification	Axiom ID
01	Base offset	Approved	Safeguard log transformation against lower boundary errors.	$AX_{QED.1}$	
02	Tanh & Mult.	Approved	Bound QED output; manage MW/SA initial value stability.	$AX_{MW-SA.1}$	
03	Non-linear Dyn.	Approved	Enhance selectivity by incorporating synthesis difficulty.	$AX_{SYN.1}$	
04	Sinusoidal	Approved	Accuracy across MW ranges via boundary-constrained oscillation.	$AX_{SIN.1}$	
05	TPSA Trig/Exp	Approved	Increase evaluation selectivity; requires boundary control.	$AX_{TPSA.1}$	
06	Piecewise	Approved	Optimize system performance and candidate selectivity.	$AX_{PW.1}$	
07	Non-linear Trans.	Approved	Enhance discriminative ability; prevent sharp MW transitions.	$AX_{MW.2}$	
08	LogP Exp.	Approved	Dimension added for drug-like property prediction.	$AX_{LogP.1}$	
09	Logistic Func.	Approved	Mitigate property fluctuations and address deadlock risks.	$AX_{LOGI.1}$	
10	TPSA Tanh	Approved	Non-linear scaling; monitored for oscillatory effects.	$AX_{TPSA.2}$	
11	Size-Lipo	Approved	Balance innovation and robustness in interaction management.	$AX_{INT.1}$	
12	TPSA Exp. Mod.	Approved	Implement damping to stabilize rapid dynamic scoring changes.	$AX_{TPSA.3}$	
13	Lipinski Rules	Approved	Alignment with domain constraints and boundary stability.	$AX_{LIP.1}$	
14	Stochastic	Approved	Enhance model adaptability under risk-factor monitoring.	$AX_{STO.1}$	
15	Boltzmann Dist.	Approved	Improve sensitivity to drug-likeness via parameter tuning.	$AX_{BOL.1}$	

```
"c1cc[nH]c1", "c1ccc2[nH]ccc2c1", "C1COCCN1", "C1CNCCN1", "C1CCCCC1O"
```

**Calculation:** Computed as the weighted geometric mean of eight desirability functions ( $d_i$ ) covering properties like MW, LogP, TPSA, etc.

$$QED = \exp \left( \frac{\sum_{i=1}^8 w_i \ln d_i(x)}{\sum_{i=1}^8 w_i} \right) \quad (30)$$

**Meaning:** A normalized index [0, 1] representing how "drug-like" a molecule is relative to known oral drugs.

#### MW (Molecular Weight)

**Calculation:**  $MW = \sum_j \text{AtomicWeight}_j$ .

**Meaning:** Total mass of the molecule. Lipinski's Rule of Five suggests  $MW < 500$  Da for optimal bioavailability.

#### LogP (Partition Coefficient)

**Calculation:** Estimated using the **Wildman-Crippen method** (Wildman and Crippen, 1999) based on atomic contributions. **Meaning:** Measures lipophilicity. Excessively high

**Complexity Analysis:** The library includes a diverse array of scaffolds (Benzene, Naphthalene), polar functional groups (Sulfonamide, Nitro), and common medicinal halogens. Combined with a maximum trajectory of 9 reaction steps, the reachable chemical space comprises approximately  $33^9 \approx 4.6 \times 10^{13}$  unique molecular graphs. AutoAxiom's ability to navigate this space under SMT-enforced Lipinski constraints without gradient-based molecular optimization demonstrates the efficacy of the evolved symbolic logic.

#### B8.6. Detailed Quantitative Metrics for Molecular Discovery

All cheminformatic properties are computed using the **RD-Kit (2023.09.1)** library. This section details the mathematical formulation and chemical significance of the metrics reported in Table 18.

[style=multiline, leftmargin=3cm, font=]

#### QED (Quantitative Estimation of Drug-likeness)

1760 LogP (> 5) leads to poor solubility  
 1761 and metabolic instability.

### 1762 **TPSA (Topological Polar Surface Area)**

1764 **Calculation:** Sum of surface  
 1765 contributions of polar atoms  
 1766 (primarily N, O, and their attached  
 1767 Hydrogens). **Meaning:** Predictive  
 1768 of drug transport and cell  
 1769 permeability. Ideally targets  
 1770  $20 \text{ \AA}^2 < \text{TPSA} < 140 \text{ \AA}^2$ .

### 1771 **SA Score (Synthetic Accessibility)**

1772 **Calculation:** Based on a  
 1773 combination of fragment  
 1774 contributions and a complexity  
 1775 penalty (Ertl and Schuffenhauer,  
 1776 2009). **Meaning:** Measures the  
 1777 difficulty of synthesis [1,10],  
 1778 where 1 is easy and 10 is extremely  
 1779 difficult.

### 1780 **Toxic Rate (Structural Alerts)**

1781 **Calculation:** Percentage of  
 1782 generated candidates containing  
 1783 "Red Flag" functional groups  
 1784 defined in the fragment library's  
 1785 metadata (e.g., Nitro  $[N+](=O)[O-]$ ,  
 1786 Sulfonamide  $S(=O)(=O)N$ ). **Meaning:**  
 1787 Indicates the prevalence of "Reward  
 1788 Hacking" where the agent exploits  
 1789 polar groups to boost QED at the  
 1790 cost of chemical toxicity.

### 1791 **HBD / HBA**

1792 **Calculation:** Count of -OH/-NH  
 1793 groups (Donors) and N/O atoms  
 1794 (Acceptors). **Meaning:** Fundamental  
 1795 parameters in Lipinski's Rule for  
 1796 hydrogen bonding potential.

1797 Table 18 presents the comprehensive statistical breakdown  
 1798 of molecular properties across four evolutionary stages.  
 1799 These metrics are calculated over 100 independent trials  
 1800 (Seeds 42–141). PySR achieves 0% toxicity in Table 17.  
 1801 However, it exhibits low Steps Taken under the sequential  
 1802 protocol, indicating stagnation. AutoAxiom instead bal-  
 1803 ances progress and constraint satisfaction under the same  
 1804 protocol.

1805 **Data Interpretation:** As observed in Table 18, **Input Ax-  
 1806 ioms** achieve the highest QED but suffer from a 34% toxic-  
 1807 ity rate, indicating they rely on pharmacologically unstable  
 1808 fragments. The **Early Stage** exhibits "Reward Hacking,"  
 1809 maximizing MW (812.8 Da) and LogP (10.19) due to a  
 1810

1811 lack of formal boundaries. **PySR** enters a state of action-  
 1812 paralysis (Steps  $\approx 0.9$ ) due to oscillatory noise in the reward  
 1813 landscape. In contrast, the **Optimal Axiom** autonomously  
 1814 suppresses toxicity to 8% while pulling MW and TPSA  
 1815 back into the Lipinski-compliant manifold, demonstrating  
 1816 the emergence of scientific intuition.

## B9. Cost Report

To evaluate the real-world practicability and hardware accessibility of AutoAxiom, we report the computational overhead recorded during execution on a standard commodity laptop. As summarized in the Cost Report (Table 19), the recorded durations represent the mean and 95% CI for a single evolutionary round, encompassing LLM orchestration, formal verification (SMT), and domain-specific simulation.

## Appendix C: Pseudocode and Mathematics proof of Method Part

### C1. Pseudocode

*Listing 5. AutoAxiom: Evolutionary Symbolic Discovery Loop.* The procedure instantiates the Tripartite Consensus (Sec 3.2) and Two-tier Verification (Sec 3.3).  $\mathcal{G}$  denotes the Core IR graph,  $\Phi$  the Red Line constraints,  $\Psi$  the execution operator, and  $\mathcal{R}$  the repair operator.

```
Algorithm: AutoAxiom_Evolutionary_Loop
Input: Core_IR  $\mathcal{G}_0$ , Domain Vocabulary  $\mathcal{D}$ ,  

       Grammar  $\Sigma$ , Total_Rounds  $T = 15$ 
Output: Optimized_Axiom  $\mathcal{G}^*$ 

1. Initialize:  $\mathcal{G}_{\text{best}} = \mathcal{G}_0$ , History_Buffer = []
2. For round  $t = 1$  to  $T$  do:
   // Phase A: Tripartite Context-Aware Proposal (Sec 3.2)
3.   Context = Summarize_History(History_Buffer) // Latest 5 rounds
4.    $\tilde{\mathcal{G}} = \text{Radical_Innovator}(\mathcal{G}_{\text{best}}, \mathcal{D}, \Sigma, \text{Context})$  // High-entropy proposal
5.    $\bar{\mathcal{G}} = \text{Conservative_Guardian}(\tilde{\mathcal{G}}, \mathcal{D}, \Phi)$  // Manifold projection
6.    $\mathcal{G}_{\text{prop}} = \text{Consensus_Maker}(\bar{\mathcal{G}}, \tilde{\mathcal{G}})$  // Pareto selection
7.   // Phase B & C: Two-tier Verification & Repair (Sec 3.3)
8.   // Tier-I: Static Ontological Invariants  $\mathbb{V}_{\text{stat}}$ 
9.   if not Check_DVM_Compliance( $\mathcal{G}_{\text{prop}}, \mathcal{D}$ ):
10.     $\mathcal{G}_{\text{sim}} = \mathcal{R}(\mathcal{G}_{\text{prop}}, \xi)$  // Repair using counter-example  $\xi$ 
11.     $N_{\text{viol}} = 1$ 
12.    else:
13.      // Tier-II: SMT Stability Check  $\mathbb{V}_{\text{smt}}$  against Red Line  $\Phi$ 
14.      if SMT_Check( $\mathcal{G}_{\text{prop}}, \Phi$ ) == UNSAT:
15.         $N_{\text{viol}} = 1$ 
16.      else:
17.         $N_{\text{viol}} = 0$ 
```

Table 18. Comprehensive Statistical Evaluation (F8: Molecular). Note the trade-off between the high QED of "toxic" input axioms and the physically grounded consistency of the Optimal Axiom. Steps Taken denotes average executed steps per trial under the protocol.

Metric	Input Axioms	Early Stage	Optimal Axiom	PySR (Best)
QED $\uparrow$	$0.787 \pm 0.02$	$0.112 \pm 0.01$	$0.397 \pm 0.04$	$0.450 \pm 0.01$
MW (Da) $\downarrow$	$317.4 \pm 12.6$	$812.8 \pm 10.7$	$443.9 \pm 13.5$	$82.3 \pm 3.4$
LogP	$3.387 \pm 0.19$	$10.19 \pm 0.73$	$6.346 \pm 0.37$	$1.752 \pm 0.06$
TPSA ( $\text{\AA}^2$ )	$69.43 \pm 5.34$	$137.4 \pm 14.5$	$62.36 \pm 6.77$	$0.84 \pm 0.94$
Steps Taken	$4.55 \pm 0.29$	$8.56 \pm 0.12$	$5.69 \pm 0.27$	$0.90 \pm 0.17$
SA Score	$3.00 \pm 0.00$	$3.00 \pm 0.00$	$3.00 \pm 0.00$	$3.00 \pm 0.00$
H-Bond Donors (HBD)	$1.63 \pm 0.12$	$3.29 \pm 0.32$	$1.70 \pm 0.17$	$0.05 \pm 0.05$
H-Bond Acceptors (HBA)	$3.02 \pm 0.23$	$6.13 \pm 0.62$	$2.83 \pm 0.29$	$0.05 \pm 0.05$
Number of Rings	$2.58 \pm 0.16$	$7.55 \pm 0.40$	$5.29 \pm 0.24$	$1.04 \pm 0.04$
<b>Toxic Rate (%) <math>\downarrow</math></b>	34%	21%	8%	0%

Table 19. AutoAxiom Cost Report. Mean and 95% CI for one evolutionary round on a commodity laptop. Peak memory = max RSS during execution (single snapshot).

Domain	LLM Calls (s)	SMT Check (s)	Simulation (s)	Peak RAM (MB)
Queue Network	$25.31 \pm 2.36$	$0.018 \pm 0.004$	$1.11 \pm 0.13$	344.6
Service Center	$30.02 \pm 4.43$	$0.004 \pm 0.001$	$18.01 \pm 1.29$	363.8
Software Opt	$25.40 \pm 3.26$	$0.003 \pm 0.001$	$0.01 \pm 0.00$	338.8
Physical PDE	$36.79 \pm 4.14$	$0.002 \pm 0.001$	$14.32 \pm 4.46$	339.8
Res. Allocation	$25.18 \pm 2.85$	$0.005 \pm 0.001$	$0.01 \pm 0.00$	342.3
Composition	$27.19 \pm 2.08$	$0.002 \pm 0.001$	$21.59 \pm 6.46$	344.1
Soft Robot	$36.37 \pm 6.30$	$0.0078 \pm 0.002$	$117.89 \pm 30.22$	362.5
Molecular	$40.62 \pm 4.53$	$0.006 \pm 0.002$	$0.81 \pm 0.22$	365.9

```

17.       $\mathcal{G}_{\text{sim}} = \mathcal{R}(\mathcal{G}_{\text{prop}}, \xi)$  // Smart
1845     Repair (Sec 3.3.3)
1846      $N_{\text{viol}} = 1$ 
1847     else:
1848      $\mathcal{G}_{\text{sim}} = \mathcal{G}_{\text{prop}}$ 
1849      $N_{\text{viol}} = 0$ 
1850
1851     // Phase D: Execution via
1852     // Simulation (Sec 3.1, Operator  $\Psi$ )
1853      $\mathbf{y} = \Psi(\mathcal{G}_{\text{sim}}, \mathbf{S}_0, \Omega)$  // Stochastic
1854     // rollout, 100 seeds
1855      $S_{\text{perf}} = \text{mean}(\mathbf{y})$ 
1856     // Performance score
1857      $\mathcal{P}_{\text{pers}} = \text{Var}(\mathbf{y})$  // Persistence (negative variance)
1858
1859     // Phase E: Dynamic Annealing
1860     // Reward (Sec 3.4)
1861      $\lambda_v(t) = \Lambda_0 \cdot \exp(-t/T_{\text{rise}})$  // Constraint hardening
1862      $\lambda_p(t) = \Lambda_0 \cdot (1 - \exp(-t/T_{\text{stab}}))$  // Robustness emphasis
1863      $\lambda_d(t) = \Lambda_0 \cdot \max(0, 1 - 2t/T)$  // Diversity decay
1864      $\mathcal{P}_{\text{div}} = \text{Calc_Diversity}(\mathcal{G}_{\text{prop}}, \mathcal{G}_0)$ 
1865
1866      $J(\mathcal{A}) = S_{\text{perf}} - \lambda_v(t)N_{\text{viol}} + \lambda_p(t)\mathcal{P}_{\text{pers}} + \lambda_d(t)\mathcal{P}_{\text{div}}$ 
1867     Score = Normalize( $J(\mathcal{A})$ ,  $J_{\text{min}}$ ,  $J_{\text{max}}$ )
1868
1869

```

```

37.      // State Update (Contextual History
38.      Buffer)
39.      History_Buffer.Append( $t, \mathcal{G}_{\text{prop}}, \text{Score}, \xi$ )
40.      if Score > Best_Score_History:
41.           $\mathcal{G}_{\text{best}} = \mathcal{G}_{\text{prop}}$ 
42.      End For
43.      Return  $\mathcal{G}_{\text{best}}$ 

```

## C2. Structural Properties of the Axiomatic Manifold $\mathbb{A}$

We define the axiomatic search space as a discrete manifold  $\mathbb{M}$  structured by the Production Grammar  $\Sigma$ . Each axiom  $\mathcal{G}$  is a point in this symbolic space.

**Lemma C2.1** (Minimalism of Axiomatic Basis). *The 5-tuple ontological basis  $\mathcal{D} = \langle \mathcal{N}, \mathcal{R}, \mathcal{O}, \Upsilon, \Phi \rangle$  spans the minimal sufficient set for representing any computable scientific law within the domain.*

*Proof.* Let  $\mathcal{L}$  be the set of all physically realizable laws in a given domain. We define a mapping  $\Psi : \mathbb{M} \rightarrow \mathcal{L}$ .

1) **Sufficiency:** By the Universal Approximation property of symbolic trees, any computable function  $\mathcal{F}$  can be

1870 represented by a finite composition of operators  $o \in \mathcal{O}$   
 1871 over nodes  $n \in \mathcal{N}$ . 2) **Minimality**: Assume a basis  
 1872  $\mathcal{D}' = \mathcal{D} \setminus \{\Upsilon\}$ . Then there exists an axiom  $\mathcal{G}$  where vari-  
 1873 able types are ambiguous, leading to a non-empty set of  
 1874 dimensionally inconsistent states, thus  $\Psi(\mathbb{M}_{\mathcal{D}'}) \not\subset \mathcal{L}$ . Sim-  
 1875ilarly, removing  $\Phi$  allows the inclusion of states that vi-  
 1876 olate domain invariants (e.g., non-negativity of physical  
 1877 constants). Thus,  $\mathcal{D}$  is the smallest set of constraints such  
 1878 that  $\forall \mathcal{G} \in \mathbb{M}_{\mathcal{D}}$ ,  $\text{Units}(\mathcal{G}) \in \Upsilon$  and  $\text{Stability}(\mathcal{G}) \vdash \Phi$ .  $\square$   
 1879

1880 **Theorem C2.2** (Canonical Confluence of Core IR). *The*  
 1881 *normalization function  $NF(\cdot)$  guarantees that algebraically*  
 1882 *equivalent axioms converge to a unique identity node in the*  
 1883 *Directed Acyclic Graph (DAG)*.

1884 *Proof.* We model  $NF$  as an **Abstract Rewriting System**  
 1885 (**ARS**). Let  $\rightarrow_R$  be the set of reduction rules (Constant Fold-  
 1886 ing, Commutative Sorting, and Common Sub-expression  
 1887 Elimination). 1) **Termination**: Each rule reduces the num-  
 1888 ber of nodes  $|V|$  or the lexicographical entropy of the  
 1889 graph. Since  $|V|$  is bounded by the grammar depth, the  
 1890 sequence of reductions is finite. 2) **Local Confluence**:  
 1891 For any node  $v$  where multiple rules apply, the Commu-  
 1892 tative Sorting ensures a deterministic path to the same  
 1893 normal form regardless of application order. By **New-**  
 1894 **man's Lemma**, since the system is both terminating and  
 1895 locally confluent, it possesses **Global Confluence**. Thus,  
 1896  $NF(\mathcal{G}_1) = NF(\mathcal{G}_2) \iff \mathcal{G}_1 \equiv_{alg} \mathcal{G}_2$ .  $\square$   
 1897

### C3. Proof of Synergistic Gain and Convergence

1900 We formalize the tripartite mechanism as a **Search Space**  
 1901 **Partitioning** problem rather than a vectorial analogy.

1902 **Theorem C3.1** (Synergistic Search Efficiency). *The tripar-*  
 1903 *ite protocol  $\{\mathbf{F}_{rad}, \mathbf{F}_{con}, \mathbf{F}_{sys}\}$  achieves a higher success*  
 1904 *probability  $P_{tri}$  than a monolithic agent by resolving the*  
 1905 *information bottleneck of constrained discovery*.

1910 *Proof.* Let  $S$  be the total search budget. In a monolithic  
 1911 LLM, the attention mechanism must simultaneously satisfy  
 1912 the innovation objective  $\mathcal{J}_{novel}$  and the safety constraint  $\Phi$ .  
 1913 This creates **Cognitive Interference**, where the entropy of  
 1914 the prompt  $H_{prompt}$  is divided. 1) **Decoupled Sampling**:  
 1915 The Radical agent ( $\mathbf{F}_{rad}$ ) samples from an unconstrained  
 1916 distribution  $\mathcal{P}_{rad}$  with high entropy, maximizing the dis-  
 1917 covery of performance-intensive manifolds. 2) **Manifold**  
 1918 **Projection**: The Conservative agent ( $\mathbf{F}_{con}$ ) acts as a **Non-**  
 1919 **expansive Projection Operator**  $\Pi_{\mathbb{M}_{\Phi}} : \mathbb{M} \rightarrow \mathbb{M}_{\Phi}$ . It  
 1920 performs local repair  $\mathcal{R}$  using the SMT counter-example  
 1921  $\xi$ . Under a fixed budget  $S$ , the probability of hitting the  
 1922 Pareto front  $\mathbb{P}(\mathcal{G} \in \mathbb{M}_{novel} \cap \mathbb{M}_{\Phi})$  is maximized because  
 1923 each agent solves a sub-problem with lower combinatorial  
 1924

complexity. This follows from the **Data Processing In-**  
 1925 **equality**: partitioning the task reduces the noise introduced  
 1926 by conflicting multi-objective instructions in a single context  
 1927 window.  $\square$

**Lemma C3.2** (Stability of the Evolutionary Loop). *The evo-*  
 1928 *lutionary sequence  $\{A_t\}$  converges to the stable manifold*  
 1929  *$\mathbb{M}_{\Phi}$  under the dynamic annealing schedule*.

*Proof.* Define a discrete-time Lyapunov function  $V(A_t) = d_{TED}(A_t, \mathbb{M}_{\Phi})$ , where  $d_{TED}$  is the **Tree Edit Distance** to  
 1930 the nearest feasible axiom. 1) **Annealing Hardening**: As  
 1931  $t \rightarrow T$ , the penalty  $\lambda_v(t)$  increases the rejection probability  
 1932 of any  $A \notin \mathbb{M}_{\Phi}$ . 2) **Drift toward Feasibility**: The Repair  
 1933 Operator  $\mathcal{R}$  ensures that for any unstable proposal  $\tilde{A}$ , the  
 1934 repaired version  $\bar{A}$  satisfies  $V(\bar{A}) = 0$ . 3) **Convergence**:  
 1935 Since the reward function  $J(A)$  enforces selection pressure,  
 1936 the probability  $P(A_t \notin \mathbb{M}_{\Phi}) \rightarrow 0$  as  $t \rightarrow T$ . Thus,  $V(A_t)$   
 1937 is a super-martingale that converges to 0, proving search  
 1938 stability within the safety envelope.  $\square$

## C4. Formal Verification and SMT Soundness

The verification engine  $V(\mathcal{G})$  utilizes SMT (Satisfiability Modulo Theories) to prove the **Safety Invariant**  $\Phi$ .

**Soundness over Bounded Intervals**: For non-linear operators (e.g.,  $\exp$ ,  $\tanh$ ), we adopt soundness restricted to the search domain  $\mathcal{C} \subset \mathbb{R}^n$ . The Z3 solver performs **Interval Arithmetic Propagation**. If the solver returns *UNSAT* for  $\mathcal{C} \wedge \mathbb{F}(\mathcal{G}) \wedge \neg \Phi$ , it is mathematically guaranteed that no state within the bounded intervals can trigger a violation under the given floating-point precision.

**SafeCut and Logical Confluence**: The repair operator  $\mathcal{R}$  utilizes the counter-example  $\xi$  to perform a **Structural Projection**. It identifies the sub-tree responsible for the violation and replaces it with a saturated primitive (e.g.,  $x \rightarrow \max(\epsilon, x)$ ). This ensures the modified axiom  $\mathcal{G}^*$  is a **Logical Confluence point** that preserves original reasoning while satisfying the safety invariant.