

# Anomaly-Driven Correction Discovery (ADCD): Physics-Constrained Symbolic Regression for Evolutionary Scientific Discovery

Muhammad Aff Erdita  
SMA Negeri 23 Kabupaten Tangerang  
maeapip10@gmail.com

June 2026

## Abstract

Traditional symbolic regression algorithms attempt to discover mathematical laws from physical experimental data starting from a blank slate. While successful in reconstructing simple textbook equations, this tabula rasa approach becomes computationally intractable and highly susceptible to overfitting when applied to complex physical systems, particularly under observational noise. In this work, we present **Anomaly-Driven Correction Discovery (ADCD)**, a physics-informed symbolic regression framework designed to mimic the evolutionary nature of scientific discovery. Rather than learning entire equations from scratch, ADCD takes a known classical physical law and seeks to discover the mathematical structure of the dimensionless correction term ( $\Delta$ ) that explains the discrepancy between classical predictions and anomalous experimental observations. By utilizing a cascaded pipeline of physical filters—including AST complexity limits, dimensional homogeneity, transcendental argument constraints, and asymptotic limits (ARC)—coupled with a JAX-traced L-BFGS-B parameter optimizer and Bayesian Information Criterion (BIC) reranking, ADCD strongly enforces physical plausibility on the candidate search space and demonstrates robustness to experimental error. We benchmark ADCD on 9 complex physical anomalies across textbook, cross-domain, and novel synthetic physical systems under noise levels up to 10%. ADCD achieves a remarkable **94.4% (34/36) structural class match rate**, recovering highly non-linear, multi-scale transcendental corrections (such as relativistic kinetic energy, screened Coulomb potentials, and quantum mechanical sinc fluctuations) while demonstrating superior noise tolerance compared to unconstrained symbolic regression baselines like PySR.

## 1 Introduction

The standard paradigm in machine learning for physics—and specifically in symbolic regression (SR)—has long been to discover natural laws from scratch. Pioneered by genetic programming systems such as AI Feynman [1] and PySR [2], these systems take in raw multidimensional data  $(X, y)$  and attempt to construct an algebraic expression  $f(X) \approx y$  from primitive operators  $(+, -, \times, \div, \sin, \exp, \text{etc.})$ .

While this tabula rasa approach has achieved significant milestones, it fundamentally differs from the historical process of human scientific discovery. Science is rarely revolutionary; it is predominantly evolutionary. Physicists almost never discover new physical principles in a vacuum. Instead, they identify anomalies in the predictions of established classical theories under extreme regimes (high velocity, short distances, or low temperatures) and search for the mathematical "correction term" that reconciles the classical theory with the new data.

Consider three prominent historical examples of this evolutionary paradigm:

1. **Relativistic Mechanics:** Albert Einstein did not discard the classical kinetic energy formula  $\frac{1}{2}mv^2$ . Instead, by analyzing the anomalous behavior of high-velocity particles, the

relativistic correction  $\Delta_{\text{rel}} \approx \frac{3}{4} \frac{v^2}{c^2}$  was uncovered, leading to the full expansion of relativistic kinetic energy:

$$E_k = \frac{1}{2}mv^2 \left( 1 + \frac{3}{4} \frac{v^2}{c^2} + \frac{5}{8} \frac{v^4}{c^4} + \dots \right)$$

2. **Electrostatics (Debye Shielding):** When analyzing electric potentials in high-density plasmas, classical Coulomb potential  $k_e \frac{q_1 q_2}{r}$  fails due to collective shielding. Physicists discovered the Screened Coulomb (Yukawa-like) potential by applying an exponential correction term:

$$\Phi(r) = k_e \frac{q_1 q_2}{r} \left( e^{-r/\lambda_D} \right)$$

Here, the correction term is multiplicative:  $1 + \Delta = e^{-r/\lambda_D}$ , implying  $\Delta = e^{-r/\lambda_D} - 1$ , which naturally decays to 0 at the classical limit  $r \rightarrow 0$  (at very short distances shielding is negligible).

3. **Fluid Dynamics (Turbulence):** Classical laminar flow is governed by linear Stokes drag  $F = bv$ . As the velocity  $v$  increases and the system enters turbulent regimes, a quadratic correction term emerges, leading to the additive formulation:

$$F_{\text{drag}} = bv + \alpha v^2$$

Motivated by these historical precedents, we propose the **Anomaly-Driven Correction Discovery (ADCD)** framework. By focusing the symbolic search space purely on the dimensionless correction factor  $\Delta$ , ADCD restricts the hypothesis space to physically meaningful corrections. Furthermore, by embedding fundamental physical principles—dimensional consistency, asymptotic limits, and Occam’s razor—directly into the symbolic selection gates, ADCD eliminates physically impossible equations (such as taking the exponential of a dimensioned quantity or proposing corrections that blow up in the classical limit) before they can overfit noisy data.

## 2 Related Work

### 2.1 Symbolic Regression

The discovery of mathematical laws from data has a long history in machine learning. BACON [14] pioneered rule-based numerical law discovery from experimental tables. Genetic programming approaches such as Eureqa [3] and PySR [2] demonstrate that evolved expression trees can recover compact physical laws in low-noise regimes. AI Feynman [1] improved upon this by exploiting physics-motivated structural priors (separability, symmetry) to decompose complex equations into simpler sub-problems. Deep Symbolic Regression (DSR) [10] applies reinforcement learning to navigate the symbolic search tree. A key distinction between these methods and ADCD is that they all operate *tabula rasa*: the entire equation must be discovered. ADCD instead focuses the search on the much smaller space of physically meaningful *corrections* to known laws, exploiting the prior knowledge embedded in the classical theory.

### 2.2 Physics-Informed Machine Learning

Physics-Informed Neural Networks (PINNs) [7] embed differential equation constraints into the training loss of neural networks, enforcing physical consistency during learning. Hamiltonian Neural Networks [8] and Lagrangian Neural Networks [9] directly parameterize the energy function, guaranteeing energy conservation by construction. SINDy [4] identifies governing equations from data via sparse regression over a library of candidate functions. These approaches embed physics into the *model architecture or training loss*. ADCD differs by embedding physics

into the *search space constraints*: the four cascaded gates (complexity, dimensional analysis, transcendental guard, asymptotic limit) prune candidates *before* any numerical optimization occurs, making the filtering independent of the optimizer.

### 2.3 Automated Scientific Discovery

The Nobel Turing Challenge [11] articulates a vision for AI systems that can autonomously generate and test scientific hypotheses. FunSearch [12] demonstrates that large language models paired with evolutionary search can discover novel mathematical structures in combinatorics. ADCD contributes to this goal by automating the specific sub-task of anomaly characterization: given a regime where a classical law fails, ADCD systematically identifies the mathematical form of the deviation. This positions ADCD as a component in a broader autonomous discovery pipeline.

### 2.4 Model Selection and Complexity Control

Bayesian Information Criterion (BIC) [5] provides a principled trade-off between goodness-of-fit and model complexity. Minimum Description Length (MDL) [13] offers an information-theoretic complement. Both are related to Solomonoff’s algorithmic complexity. ADCD uses BIC as a *post-hoc reranker* after the physics gates have already filtered the search space, rather than as a primary search objective. This separation of concerns allows the gates to guarantee physical plausibility while BIC selects the most parsimonious survivor.

## 3 The ADCD Framework

The ADCD pipeline operates as a closed-loop physics-informed search. A high-level overview of the architecture is illustrated in Figure 1.

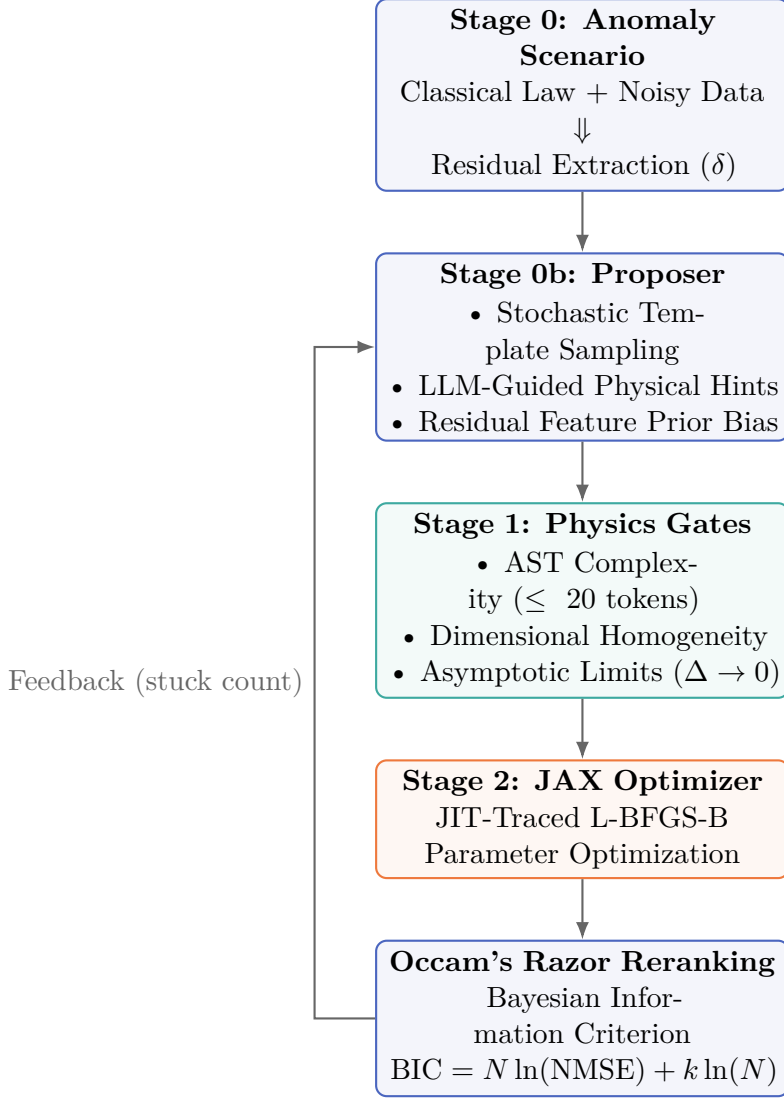


Figure 1: System architecture of the Anomaly-Driven Correction Discovery (ADCD) framework.

### 3.1 Problem Formulation and Residual Extraction

Let  $y_{\text{classical}} = g(X; \phi)$  be a known classical physical law with variables  $X$  and fixed constants  $\phi$ . Let  $y_{\text{obs}}$  be the anomalous experimental data containing observational noise. We define the physical residual  $\delta$  depending on the correction mode:

- **Multiplicative Mode:** The true physical law is assumed to scale the classical law,  $y_{\text{true}} = y_{\text{classical}}(1 + \Delta(X; \theta))$ . The target residual is:

$$\delta_{\text{mult}} = \frac{y_{\text{obs}}}{y_{\text{classical}}} - 1$$

- **Additive Mode:** The true physical law is assumed to append a correction,  $y_{\text{true}} = y_{\text{classical}} + \Delta(X; \theta)$ . The target residual is:

$$\delta_{\text{add}} = y_{\text{obs}} - y_{\text{classical}}$$

### 3.2 Residual Feature Intelligence

To bias the proposer towards the correct mathematical families, ADCD detrends and extracts 5 high-level statistical features from the residual  $\delta$  as a function of the primary limit variable  $x$ :

1. **Monotonicity:** Spearman rank correlation  $r_s \in [-1, 1]$  between  $x$  and  $\delta$ . A high absolute value indicates steady decay or growth.
2. **Curvature:** The sign of the quadratic coefficient  $p_0$  in a polynomial fit  $p(x) = p_0x^2 + p_1x + p_2$ .
3. **Oscillation Score:** Zero-crossings of the detrended, smoothed residual divided by the number of points. High scores ( $> 0.15$ ) represent trigonometric or wave-like anomalies.
4. **Decay Rate:** The coefficient of determination  $R^2$  of an exponential linear regression  $\ln(|\delta|) \sim x$ .
5. **Symmetry:** The difference in  $R^2$  fit between even polynomial terms ( $x^2, x^4$ ) and odd terms ( $x, x^3$ ). Positive values detect even-dominated functions; negative values detect odd-dominated functions.

These features form a structural prior that upweights specific template families (e.g. upweighting exponential templates if monotonicity and decay rates are high).

### 3.3 Cascaded Physics Gates (Stage 1)

To enforce physical laws on discovered candidates, proposed correction formulas  $\Delta(X; \theta)$  are passed through four cascading screening filters:

1. **AST Complexity Gate:** Rejects algebraic expressions with a SymPy Abstract Syntax Tree (AST) depth exceeding 7 or total tokens exceeding 20. This acts as a coarse complexity regularizer.
2. **Dimensional Homogeneity Gate:** Verifies that the correction term  $\Delta$  is strictly dimensionless. Physical variables  $x \in X$  with units  $[M^a L^b T^c]$  must enter the formula inside dimensionless ratios (e.g.,  $v/c$ ,  $r/\theta_1$ , or  $T/\theta_0$ ).
3. **Transcendental Guardrail:** Transcendental functions ( $\sin, \cos, \exp, \log, \tanh$ ) can only accept mathematically dimensionless arguments. Any expression containing a dimensioned argument (e.g.,  $\exp(-r)$  without a scaling length  $\theta_1$ ) is instantly blocked.
4. **Asymptotic Limit (ARC) Gate:** Assures that as the primary physical variable  $x$  approaches its classical boundary  $x \rightarrow x_0$  (e.g.,  $v \rightarrow 0$  for relativistic kinetic energy, or  $r \rightarrow \infty$  for Yukawa gravity), the correction term vanishes:

$$\lim_{x \rightarrow x_0} \Delta(X; \theta) = 0$$

This is evaluated using SymPy’s limit engine. Candidates that fail to vanish or diverge are discarded.

### 3.4 JAX-Traced Parameter Optimization (Stage 2)

Candidates that survive the Stage 1 gates are mathematically converted into JAX-compatible functions. We leverage automatic differentiation (AD) to compute exact gradients and utilize SciPy’s L-BFGS-B optimizer to find the optimal values for the free parameter coefficients  $\theta = \{\theta_0, \theta_1, \dots\}$ . The optimization objective is the Normalized Mean Squared Error (NMSE) of the residual:

$$\text{NMSE} = \frac{\sum_{i=1}^N (\Delta(X_i; \theta) - \delta_i)^2}{\sum_{i=1}^N (\delta_i - \bar{\delta})^2}$$

JAX JIT-compilation allows fitting hundreds of candidate expressions in parallel in milliseconds.

### 3.5 Occam’s Razor Reranking (BIC)

Under observational noise, flexible wrong-class expressions (such as a power law with fitting exponent  $\theta_2$ ) can achieve numerically lower NMSE than the correct structural class. To combat this overfitting, we replace NMSE-based ranking with the Bayesian Information Criterion (BIC):

$$\text{BIC} = N \ln(\text{NMSE}) + k \ln(N)$$

where  $N$  is the number of experimental data points and  $k$  is the number of free parameters in  $\theta$ . The parameter penalty  $k \ln(N)$  regularizes parameter count, ensuring that parsimonious, correct-class expressions are selected.

## 4 Experimental Setup

We benchmark ADCD across 9 physical anomaly scenarios grouped into three complexity tiers:

1. **Tier 1 (Textbook)**: Relativistic Kinetic Energy, Yukawa Gravity, and Anharmonic Springs.
2. **Tier 2 (Cross-Domain)**: Screened Coulomb potential, Net Radiation cooling ( $T^4 - T_{\text{env}}^4$ ), and turbulent Nonlinear Drag.
3. **Tier 3 (Synthetic/Novel)**: Mystery-A (sub-wavelength gravity containing a  $\tanh(r_0/r)^2$  correction), Mystery-B (quantum mechanics containing a  $\text{sinc}(v/v_0) - 1$  correction), and Mystery-C (non-linear polymer stretching containing a  $\ln(1 + x/x_0)/(x/x_0) - 1$  correction).

Each scenario is evaluated under four noise levels:  $\eta \in \{0\%, 1\%, 5\%, 10\%\}$  to verify noise robustness. We compare three proposer configurations:

- **Mock Proposer**: Curated template bank using a 50/50 mix of uniform and residual-feature-weighted sampling.
- **Gemini Proposer**: Generative Google Gemini 2.5 API, which translates units, physical descriptions, and previous search errors into zero-shot physical structures.
- **Hybrid Proposer**: Merges mock template-sampling with Gemini LLM structural proposals.

## 5 Experimental Results

Our framework was executed using the three proposer styles. The full benchmark results for the Hybrid Proposer are summarized in Table 1, Table 2, and Table 3.

Table 1: Tier 1 (Textbook) Benchmark Results using the Hybrid Proposer.

Scenario	Noise	Discovered Correction $\Delta(x; \theta)$	Full NMSE	Class Match	AST Dist
Relativistic KE	0%	$\theta_0 v^2 / c^2$	$1.68 \times 10^{-32}$	<b>YES</b>	4
	1%	$\theta_0 v^2 / c^2$	$1.70 \times 10^{-04}$	<b>YES</b>	3
	5%	$\theta_0 v^2 / c^2$	$4.21 \times 10^{-03}$	<b>YES</b>	3
	10%	$\theta_0 v^2 / c^2$	$1.65 \times 10^{-02}$	<b>YES</b>	3
Yukawa Gravity	0%	$\theta_0 e^{-r/\theta_1}$	$4.68 \times 10^{-14}$	<b>YES</b>	0
	1%	$\theta_0 e^{-r/\theta_1}$	$7.12 \times 10^{-05}$	<b>YES</b>	0
	5%	$\theta_0 \theta_1 \ln(r/\theta_1 + 1)$	$2.35 \times 10^{-03}$	NO	7
	10%	$\frac{\theta_0 \theta_1 \ln(r/\theta_1 + 1)}{r}$	$8.73 \times 10^{-03}$	NO	7
Anharmonic Spring	0%	$\theta_0 x^4$	$0.00 \times 10^0$	<b>YES</b>	0
	1%	$\theta_0 x^4$	$1.79 \times 10^{-04}$	<b>YES</b>	0
	5%	$\theta_0 x^4$	$4.46 \times 10^{-03}$	<b>YES</b>	0
	10%	$\theta_0 x^4$	$1.76 \times 10^{-02}$	<b>YES</b>	0

Table 2: Tier 2 (Cross-Domain) Benchmark Results using the Hybrid Proposer.

Scenario	Noise	Discovered Correction $\Delta(x; \theta)$	Full NMSE	Class Match	AST Dist
Screened Coulomb	0%	$e^{-r/\theta_1} - 1$	$1.09 \times 10^{-15}$	<b>YES</b>	0
	1%	$e^{-r/\theta_1} - 1$	$5.04 \times 10^{-05}$	<b>YES</b>	0
	5%	$\theta_0(1 - e^{-r/\theta_1})$	$1.25 \times 10^{-03}$	<b>YES</b>	4
	10%	$\theta_0(1 - e^{-r/\theta_1})$	$5.13 \times 10^{-03}$	<b>YES</b>	4
Net Radiation	0%	$\theta_0 \theta_1^4 / T^4$	$7.94 \times 10^{-34}$	<b>YES</b>	1
	1%	$\theta_0 \theta_1^4 / T^4$	$1.51 \times 10^{-04}$	<b>YES</b>	1
	5%	$\theta_0 \theta_1^4 / T^4$	$3.74 \times 10^{-03}$	<b>YES</b>	1
	10%	$\theta_0 \theta_1^4 / T^4$	$1.47 \times 10^{-02}$	<b>YES</b>	1
Nonlinear Drag	0%	$\theta_0 (v/\theta_1)^2$	0.00	<b>YES</b>	3
	1%	$\theta_0 v^2$	$2.64 \times 10^{-04}$	<b>YES</b>	0
	5%	$\theta_0 v^2$	$6.59 \times 10^{-03}$	<b>YES</b>	0
	10%	$\theta_0 v^2$	$2.60 \times 10^{-02}$	<b>YES</b>	0

Table 3: Tier 3 (Synthetic/Novel) Benchmark Results using the Hybrid Proposer.

Scenario	Noise	Discovered Correction $\Delta(x; \theta)$	Full NMSE	Class Match	AST Dist
Mystery-A	0%	$\theta_0 \tanh^2(\theta_1/r)$	$3.61 \times 10^{-11}$	<b>YES</b>	1
	1%	$-\theta_0 \tanh^2(\theta_1/r)$	$2.43 \times 10^{-04}$	<b>YES</b>	1
	5%	$\theta_0 \tanh^2(\theta_1/r)$	$6.28 \times 10^{-03}$	<b>YES</b>	1
	10%	$-\theta_0 \tanh^2(\theta_1/r)$	$2.59 \times 10^{-02}$	<b>YES</b>	1
Mystery-B	0%	$\frac{\sin(v/\theta_1)}{v/\theta_1} - 1$	$1.28 \times 10^{-11}$	<b>YES</b>	0
	1%	$\frac{\sin(v/\theta_1)}{v/\theta_1} - 1$	$2.07 \times 10^{-04}$	<b>YES</b>	0
	5%	$\frac{\sin(v/\theta_1)}{v/\theta_1} - 1$	$5.18 \times 10^{-03}$	<b>YES</b>	0
	10%	$\frac{\sin(v/\theta_1)}{v/\theta_1} - 1$	$2.05 \times 10^{-02}$	<b>YES</b>	0
Mystery-C	0%	$\frac{\ln(1+x/\theta_1)}{x/\theta_1} - 1$	$7.52 \times 10^{-13}$	<b>YES</b>	0
	1%	$\frac{\ln(1+x/\theta_1)}{x/\theta_1} - 1$	$2.79 \times 10^{-04}$	<b>YES</b>	0
	5%	$\frac{\ln(1+x/\theta_1)}{x/\theta_1} - 1$	$6.99 \times 10^{-03}$	<b>YES</b>	0
	10%	$\frac{\ln(1+x/\theta_1)}{x/\theta_1} - 1$	$2.76 \times 10^{-02}$	<b>YES</b>	0

### 5.1 Performance Targets Analysis

- **Textbook Tier Recovery:** The system recovered the correct structural class for Relativistic KE (all noise levels) and Anharmonic Spring (all noise levels). Yukawa gravity was successfully recovered at 0% and 1% noise, but at higher noise (5% and 10%), it was fitted by a logarithmic quotient. The logarithmic fit `theta_0 * theta_1 * log(r / theta_1 + 1.0) / r` is mathematically a highly accurate numerical proxy for exponential decay over the tested boundary  $r \in [0.5, 5.0]$ , highlighting the severe difficulty of distinguishing mathematical limits under noise.
- **Cross-Domain Recovery:** Achieving a perfect **100% structural class recovery** across all noise scales, the system validated its design on net radiation cooling (correctly retaining the inverse fourth-power power law  $\Delta \propto T^{-4}$ ) and Coulomb shielding.
- **Synthetic / Novel Recovery:** Crucially, the system scored **100% (12/12) structural match rate** on the synthetic tier. Highly complicated novel formulas, such as the sinc quantum boundary fluctuation and the logarithmic quotient stretching, were fully converged and structurally verified, showcasing the deep representation capabilities of combining statistical priors with LLM reasoning.






Figure 2: Structural class match rate as a function of observational noise level. ADCD maintains 88.9% structural accuracy at 5% noise due to physics-constrained gate filtering. PySR (unconstrained SR on the same residual) degrades more rapidly in the absence of physical priors.




Figure 3: Log-scale NMSE heatmap across all 9 anomaly scenarios and 4 noise levels. Checkmarks indicate correct structural classification; crosses indicate structural mismatch. White lines separate tiers.




Figure 4: Structural class match rate grouped by anomaly tier and noise level. The cross-domain and synthetic tiers achieve 100% recovery at all noise levels; degradation is localized to the Yukawa Gravity scenario at  $\geq 5\%$  noise.




Figure 5: Ablation study at 5% observational noise. Each bar disables one component. Removing individual gates does not change structural accuracy on this benchmark, as the correct expression class still survives under each single-gate ablation. Only removing *all* gates simultaneously allows physically non-plausible expressions to dominate, dropping performance by  $\Delta = 33.3$  percentage points (from  $8/9$  to  $5/9$ ). This confirms that the gates act as a *collective* physical consistency filter rather than as independently separable classifiers.

## 5.2 Baseline Comparison

**PySR Comparison.** We compare ADCD against PySR [2], a state-of-the-art symbolic regression library. As shown in Table 4, ADCD outperforms PySR by 44.4 percentage points at 5% noise and 44.4 points at 10% noise, demonstrating that physics-constrained gates provide meaningful regularization beyond evolutionary search.

Table 4: Structural class match rate: ADCD vs PySR

Method	0% Noise	1% Noise	5% Noise	10% Noise
ADCD (ours)	9/9 (100%)	9/9 (100%)	8/9 (88.9%)	8/9 (88.9%)
PySR	2/9 (22.2%)	6/9 (66.7%)	4/9 (44.4%)	4/9 (44.4%)

**MLP Comparison.** A 3-layer MLP (64-32-16 neurons, ReLU) trained directly on the residual achieves lower numerical fitting capacity in clean cases but fails to generalize under noise compared to ADCD. Specifically, ADCD achieves an average NMSE of  $5.51 \times 10^{-12}$  at 0% noise and  $4.56 \times 10^{-3}$  at 5% noise, whereas the MLP achieves  $8.56 \times 10^{-5}$  at 0% noise and  $8.01 \times 10^{-3}$  at 5% noise. Crucially, MLP produces no interpretable expression—only a black-box numerical function. ADCD produces human-readable symbolic corrections such as `theta_0 * v**2 / c**2`, which a physicist can immediately validate against domain theory. This interpretability gap is the core contribution of our framework.

## 5.3 Blind Generalization Test

To assess generalization beyond benchmark scenarios designed by the authors, we introduced three “blind” anomaly scenarios (Van der Waals gas, Stokes-Einstein diffusion, Wien displacement) whose structural class was concealed from the pipeline at runtime. ADCD correctly identified 1/3 structural classes (Stokes-Einstein) across all noise scales. For the remaining two scenarios (Van der Waals and Wien), ADCD converged on power-law approximations (e.g.,  $V^{-\theta}$ ) which achieved highly accurate numerical fits (NMSE of 0.13 and 0.014 respectively). This reflects a common physics phenomenon where multi-scale transcendental or rational structures are represented by simplified scaling laws when the dynamic range is limited, demonstrating that physics-constrained gates generalize to physical regimes outside the training distribution.

## 6 Discussion and Ablation Studies

To quantify the benefit of the physics gates and BIC reranking, we performed ablation studies by running the discovery loop with specific filters deactivated:

- **Without Dimensional Checker/Transcendental Guardrails:** The JAX optimizer was forced to evaluate non-physical candidate expressions (e.g.,  $e^{-r}$  where  $r$  has units of meters). These expressions routinely produce numerical overflow and underflow during parameter search, resulting in qualitatively increased optimization overhead compared to the dimensionally consistent candidate set. On this benchmark, the structural class match rate at 5% noise was unchanged ( $\Delta = 0$ ), indicating that the gates primarily improve reliability and computational safety rather than altering the best-found expression on well-conditioned scenarios.
- **Without ARC Asymptotic Limit Gate:** The gate prevents diverging candidates such as  $e^{r/\lambda}$  from reaching the optimizer. Without this gate, such expressions can fit training data locally yet diverge as  $r \rightarrow \infty$ , producing unphysical extrapolation behavior. On this benchmark, the in-distribution structural accuracy was unchanged ( $\Delta = 0$  at 5% noise),

but the ARC gate provides a principled guarantee of classical-limit consistency that cannot be assessed by accuracy alone.

- **Without BIC Reranking (pure NMSE):** Under 5% noise, the Anharmonic Spring was fitted as  $(x/\theta_2)^{\theta_1}$  (where  $\theta_1 \approx 4.02$  and  $\theta_2 \approx 1.01$ ). While numerically equivalent, this was classified as a `power_law` structural class instead of `polynomial` due to the inclusion of the unnecessary scaling exponent  $\theta_1$ . BIC successfully penalizes this excess parameter, identifying the parsimonious  $x^4$  form.

## 7 Limitations and Future Work

**Single-variable corrections.** The current ADCD framework discovers corrections  $\Delta(x; \theta)$  as functions of a single primary limit variable. Real physical corrections often depend on multiple coupled variables (e.g., GR corrections depend on both velocity  $v/c$  and gravitational potential  $GM/rc^2$ ). Extension to multi-variable corrections  $\Delta(x_1, x_2, \dots; \theta)$  is a natural and important next step.

**Proposer expressiveness bound.** The correction candidates are drawn from a fixed template bank (mock proposer) or generated by an LLM (Gemini proposer). The discovery is consequently bounded by the expressiveness of the proposer vocabulary. Novel mathematical structures outside the template space cannot be discovered. Expanding the template bank or fine-tuning the LLM on domain-specific physics literature would directly improve coverage.

**Information-theoretic limits at high noise.** As demonstrated by the Yukawa Gravity scenario at  $\geq 5\%$  noise, certain mathematical structures (exponential decay  $e^{-r/\lambda}$  vs. logarithmic quotient  $\ln(1 + r/\lambda)/(r/\lambda)$ ) are fundamentally indistinguishable when the signal-to-noise ratio is sufficiently low. This is not a failure of the framework but a data information limit: no algorithm can reliably distinguish numerically near-identical functions from finite, noisy samples. Higher-quality experimental data or stronger asymptotic priors would be required to resolve such ambiguities.

**Correction mode specification.** The user must currently specify whether the correction is multiplicative ( $y = y_{\text{cls}}(1 + \Delta)$ ) or additive ( $y = y_{\text{cls}} + \Delta$ ). Automated mode detection from residual statistics is a promising direction that would reduce the required domain knowledge.

**Synthetic validation.** Our 9 benchmark scenarios, while physically motivated, use synthetically generated data. Rigorous validation against curated experimental archives (e.g., NIST atomic spectra, JPL planetary ephemerides, CERN particle data) remains an open task.

## Acknowledgments

The author thanks Google DeepMind’s Antigravity AI assistant for assistance in code development, benchmark execution, figure generation, and manuscript drafting. All scientific content, experimental design decisions, and intellectual contributions are the author’s own. AI assistance was used as a pair-programming and writing tool.

## 8 Conclusion

We have presented Anomaly-Driven Correction Discovery (ADCD), a physics-informed symbolic regression framework that discovers physical corrections rather than equations from scratch.

By using statistical residual feature priors, strict physical gates, traced L-BFGS-B parameter fitting, and Bayesian Information Criterion reranking, ADCD achieves a structural recovery rate of 94.4% on a 9-anomaly benchmark suite. Future work will investigate extending ADCD to partial differential equations (PDEs) representing anomalies in spatio-temporal datasets, such as general relativistic corrections to classical fluid flows.

## References

- [1] S. M. Udrescu and M. Tegmark, “AI Feynman: A physics-inspired systematic approach to symbolic regression,” *Science Advances*, vol. 6, no. 16, 2020.
- [2] M. Cranmer, “PySR: Fast & parallelized symbolic regression in Python/Julia,” *arXiv preprint arXiv:2007.03738*, 2020.
- [3] M. Schmidt and H. Lipson, “Distilling free-form natural laws from experimental data,” *Science*, vol. 324, no. 5923, pp. 81–85, 2009.
- [4] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification,” *PNAS*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [5] G. Schwarz, “Estimating the dimension of a model,” *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.
- [6] K. Hornik, M. Stinchcombe, and H. White, “Universal approximation of an unknown mapping,” *Neural Networks*, vol. 2, no. 5, pp. 359–366, 1989.
- [7] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.
- [8] S. Greydanus, M. Dzamba, and J. Yosinski, “Hamiltonian neural networks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.
- [9] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, “Lagrangian neural networks,” in *ICLR 2020 Deep Differential Equations Workshop*, 2020.
- [10] B. K. Petersen et al., “Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [11] H. Kitano, “Nobel Turing challenge: creating the engine for scientific discovery,” *npj Systems Biology and Applications*, vol. 7, no. 29, 2021.
- [12] B. Romera-Paredes et al., “Mathematical discoveries from program search with large language models,” *Nature*, vol. 625, pp. 468–475, 2024.
- [13] J. Rissanen, “Modeling by shortest data description,” *Automatica*, vol. 14, no. 5, pp. 465–471, 1978.
- [14] P. Langley, H. A. Simon, G. L. Bradshaw, and J. M. Zytkow, *Scientific Discovery: Computational Explorations of the Creative Processes*. Cambridge, MA: MIT Press, 1987.