
Bridging the Gap Between Scientific Laws Derived by AI Systems and Canonical Knowledge via Abductive Inference with AI-Noether

Karan Srivastava

University of Wisconsin-Madison || IBM Research
Madison, WI || Yorktown Heights, NY
ksrivastava4@wisc.edu

Sanjeeb Dash

IBM Research
Yorktown Heights, NY
sanjeebd@us.ibm.com

Ryan Cory-Wright

Imperial Business School
London, UK
r.cory-wright@imperial.ac.uk

Barry Trager

IBM Research
Yorktown Heights, NY
bmt@us.ibm.com

Cristina Cornelio

Samsung AI
Cambridge, UK
c.cornelio@samsung.com

Lior Horesh

IBM Research
Yorktown Heights, NY
lhoresh@us.ibm.com

Abstract

Advances in AI have shown great potential in contributing to the acceleration of scientific discovery. Symbolic regression can fit interpretable models to data, but these models are not necessarily derivable from established theory. Recent systems (e.g., AI-Descartes, AI-Hilbert) enforce derivability from prior knowledge. However, when existing theories are incomplete or incorrect, these machine-generated hypotheses may fall outside the theoretical scope. Automatically finding corrections to axiom systems to close this gap remains a central challenge in scientific discovery. We propose a solution: an open-source algebraic geometry-based system that, given an incomplete axiom system expressible as polynomials and a hypothesis that the axioms cannot derive, generates a minimal set of candidate axioms that, when added to the theory, provably derive the (possibly noisy) hypothesis. We illustrate the efficacy of our approach by showing that it can reconstruct key axioms required to derive the carrier-resolved photo-Hall effect, Einstein’s relativistic laws, and several other laws.

1 Introduction: The AI Scientific Revolution and Abductive Inference

The proliferation of artificial intelligence in accelerating scientific discovery has allowed us to vastly extend our exploration of science at a breakneck speed while also bringing a fundamental question to the forefront of research: How can we integrate AI-generated insights into canonical scientific knowledge? Since the release of ChatGPT in November 2022, the ubiquity of generative AI has made it easier and cheaper to make scientific discoveries [2, 37, 21, 91, 12, 49, 50, 31, 10, 95, 11, 61, 77]. Indeed, machine learning and AI techniques have recently developed new antivenom [93], new sorting algorithms [62], new drugs that have progressed to phase II clinical trials [97], new mathematical discoveries [80, 44], and new equations in psychology [74], to name a few applications. More recently, there has even been progress in integrating data and explicit background theory in the form of axioms within the process of discovery of mathematical models in science in symbolic form [16, 18].

27 However, the acceleration of scientific discovery poses both challenges and opportunities for the
 28 scientific community. First, with the reduced cost of making and disseminating incremental scientific
 29 discoveries, when combined with strong incentives for scientists to ‘publish-or-perish’ [69], the
 30 sheer volume of AI-generated and AI-assisted work threatens to overwhelm the peer-review system
 31 and obscure genuine advances amid incremental contributions or noise. For instance, the proportion
 32 of arXiv papers substantially modified by LLMs has steadily increased from around 2.5% in 2022 to
 33 around 10% in 2024 [54], with the authors who post the most preprints on arXiv being disproportionately
 34 likely to use LLMs. Second, from a scientific accuracy perspective, it is often unclear whether
 35 scientific discoveries made with generative AI or machine learning generalize well to unseen data
 36 and advance understanding [48, 64], or merely fit known data well. This lack of interpretability and
 37 generalizability (see [55, 81] for discussions) means that verifying scientific discoveries (e.g., via
 38 peer review) is becoming as important as the discovery process itself [17]. Third, stacking AI-driven
 39 discoveries on top of each other poses a serious risk in domains where absolute truth matters more
 40 than predictive accuracy. This is especially true for discoveries generated by AI systems that pri-
 41 oritize minimizing in-sample training loss and may not generalize out-of-distribution. In such an
 42 environment, this can result in a ‘scientific model collapse’ in the sense of [85] where AI-systems
 43 stack plausible-seeming but incorrect hallucinations on top of one another, similar to the replication
 44 crisis that struck psychology in the early 2010s [84].

45 To address the challenges posed by and capitalize on the opportunities presented by AI-driven sci-
 46 entific discovery, we need a framework for *systematic revision of scientific foundations* - what Kuhn
 47 called paradigm shifts [51]. In the life cycle of a scientific paradigm, a scientific community starts
 48 with a certain set of natural laws and beliefs (axioms). They then use these axioms and experimental
 49 data to make scientific discoveries, built on their axiomatic foundation, until the gap between the
 50 foundation and prominent discoveries becomes so apparent that a *paradigm shift* to new axioms is
 51 needed. Kolmogorov’s axioms in probability theory [36] are a good example of this process. Before
 52 Kolmogorov introduced his axioms in 1933, discoveries were made in an ad hoc manner, leading to
 53 paradoxes such as Bertrand’s paradox [82]. However, following Kolmogorov’s paradigm shift, it be-
 54 came possible to distinguish between correct and incorrect discoveries in probability theory. While
 55 the natural sciences involve more empirical complexities, Kuhnian paradigm shifts regularly occur,
 56 leading to significant breakthroughs when they do. For instance, in order to develop his theory of
 57 special relativity, Einstein necessarily had to modify fundamental postulates from prior Newtonian
 58 mechanics such as recognizing that the speed of light is constant [14]. Similarly, Maxwell’s revision
 59 of Ampère’s circuital law to include a displacement term completed classical electrodynamics for
 60 time-varying fields and led to his prediction that light is an electromagnetic wave [57].

61 These observations, coupled with the increasing availability of computational power driven by
 62 Moore’s Law [60], suggest that AI and machine learning should be deployed to develop new sci-
 63 entific paradigms, much as AI has successfully been applied to scientific discovery. In particular,
 64 abductive inference should be employed within the scientific method to determine when a new sci-
 65 entific paradigm is necessary and when the set of scientific axioms used by a community needs
 66 revision. Accordingly, we propose an AI framework to solve the problem of abductive reasoning.

67 **Main Contributions.** We present AI-Noether, a system that automates abductive inference to iden-
 68 tify missing axioms in scientific theories. Given background axioms and a hypothesis that can-
 69 not be derived from them, AI-Noether systematically generates candidate axioms whose addition
 70 makes the hypothesis derivable, transforming paradigm shifts that were historically informal and
 71 human-driven into a scalable computational procedure. This enables us not only to identify dis-
 72 crepancies between theory and observation, but to resolve them by inferring the minimal axiomatic
 73 modifications required. We name our system AI-Noether in honor of Emmy Noether, one of the
 74 leading mathematicians of her time who made significant contributions to both algebraic geometry
 75 and physics. We build on her work at the core of our system in order to make abductive inference
 76 possible, as we will see. This name also reflects Noether’s legacy of uncovering the algebraic struc-
 77 tures underlying physical laws, most famously via Noether’s theorem on symmetries and conserved
 78 quantities [67], paralleling our goal of inferring the minimal additional axioms needed to render a
 79 hypothesis derivable. To support reproducibility and future research, we open-source AI-Noether’s
 80 complete implementation, including all benchmark problems used in this study.

81 To the best of our knowledge, this is the first research work to study automated abductive reasoning
 82 in the context of scientific discovery, in which derivations may consist of fully general polynomial
 83 axioms over any set of variables; see Appendix A for a review of existing automated approaches to




Figure 1: **AI-Noether-augmented scientific discovery loop.** Given background theory axioms and observational data, a new candidate hypothesis is generated that may or may not be compatible with the existing theory (under different notions of derivability and consistency) and may fit the data with varying approximation error. AI-Noether performs abductive inference to identify missing axioms that reconcile these hypotheses with canonical knowledge, while symmetrically, on the data-driven end, experimental design prescribes the generation of new data. Both procedures are instrumental to further refine both hypotheses and axioms, closing the discovery-verification loop.

84 abductive inference and a comparison of our system with existing AI-driven approaches to scientific
 85 discovery of mathematical models. Thus, our work constitutes a step toward the larger goal of
 86 augmenting and transforming the entire scientific method using AI by bridging the gap between
 87 machine-generated hypotheses and known theory.

88 2 Abductive Inference

89 In the pursuit of scientific discovery, we often find ourselves confronted with a striking reality:
 90 established theoretical frameworks routinely fall short of fully explaining observable phenomena.
 91 For instance, standard models of the Big Bang theory predict that there should be about three times
 92 as much lithium in the universe as is observed [32]. This disconnect between theory and observation
 93 presents a fundamental challenge: how do we systematically identify what is missing from our
 94 understanding?

95 Resolving discrepancies between theory and observation is increasingly relevant in the era of ma-
 96 chine learning, because data-driven methods can reveal patterns that existing theory cannot explain.
 97 Thus, fully integrating insights from statistical AI into science is of considerable interest. How-
 98 ever, discoveries from statistical AI are often difficult to verify due to their uninterpretable nature
 99 [81, 55], and integrating them into science remains an open problem. Bridging this disconnect can
 100 be expressed as the following abductive inference problem: given a set of known scientific princi-

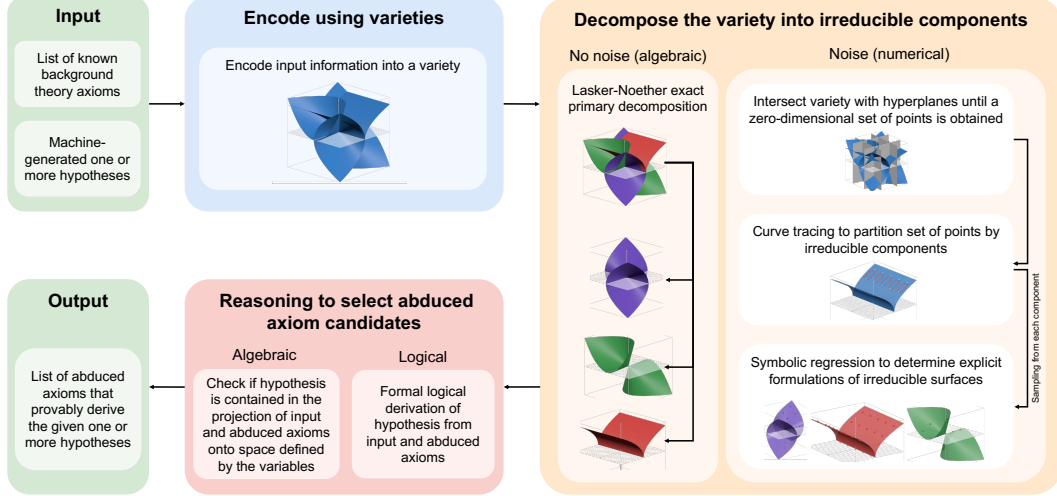


Figure 2: **AI-Noether system overview.** Given background theory axioms and one or more candidate hypotheses that are not supported by the existing theory, AI-Noether identifies additional axioms that reconcile the hypotheses with the background knowledge. The system first encodes axioms and hypotheses into a geometric representation, decomposes this representation into fundamental components, and then reasons over these components to determine which candidate axioms, when added to the theory, are sufficient to explain the hypotheses. These axioms are returned as output after verification, which proves the hypotheses, using both algebraic and logic-based methods.

101 ples (axioms) and a newly discovered relationship that cannot be derived from them, what additional
 102 principles must we postulate to bridge the gap?

103 In this section, we formalize this problem for the case where both axioms and discoveries can be
 104 expressed as multivariate polynomial equations. We assume that all axioms, including any miss-
 105 ing ones, and the hypotheses, are expressible as polynomial equations of the form $A(x) = 0$. A
 106 collection of polynomial equations $A_1(x) = 0, \dots, A_k(x) = 0$ implies the polynomial equation
 107 $Q(x) = 0$ if $Q(\bar{x}) = 0$ for every solution $\bar{x}$ of the system $\{A_1, \dots, A_k\}$.

108 Formally, we consider the problem:

109 *Given axioms $A_1, \dots, A_k$ and a hypothesis Q such that $A_1, \dots, A_k$ do not imply Q ,*
 110 *find additional r axioms $A_{k+1}, \dots, A_{k+r}$ with the “smallest residual” relative to*
 111 *$A_1, \dots, A_k$ such that the augmented system $A_1, \dots, A_{k+r}$ implies Q .*

112 **Terminology:** Throughout the rest of this work, we use the terms $A_1, \dots, A_k$ “explain”, “derive”
 113 or “imply” Q and Q is a “consequence of” $A_1, \dots, A_k$ synonymously. Intuitively, the notion of
 114 “smallest residual” corresponds to the idea that we wish to add as little new information as possible
 115 to $A_1, \dots, A_k$ to derive Q ; we will formalize this notion later on. We assume that no data are
 116 available for the missing axioms $A_{k+1}, \dots, A_{k+r}$; otherwise, abduction would be equivalent to
 117 hypothesis generation, a topic that has been frequently studied in scientific discovery. Under this
 118 terminology, our system generates a set of candidate axioms $\{\hat{A}_{k+1}, \dots, \hat{A}_l\}$ that, along with the
 119 known axioms, explain Q . Further, if no single candidate axiom can be added to explain Q , our
 120 system detects this. Current machine-assisted discovery tools either rely on partial or complete
 121 background theory, or none at all, to discover mathematical models that fit observational data, but
 122 few append or correct the theory in an automated way, and none do so at the same level of generality.

123 In the remainder of this section, we give an overview of our approach.

124 2.1 Approach

125 Our goal is to identify candidate missing polynomial axioms $A_{k+1}, \dots, A_{k+r}$ that, along with a
 126 known set of axioms $A_1, \dots, A_k$, can be used to derive a target hypothesis Q . In this study, we fo-
 127 cus on systems defined over real polynomials with real solutions, though the approach extends natu-
 128 rally to complex equations (see Appendix E.1). The method consists of three main steps—**Encode**,
 129 **Decompose**, and **Reason**, illustrated in Figure 2 and summarized in the **Algorithm Summary**.

Algorithm Summary

- 1: **Input:** Axioms $A_1, \dots, A_k$ and hypothesis Q
 - 2: **Encode:** Construct the solution set (variety) of $\{A_1, \dots, A_k, Q\}$
 - 3: **Decompose:** Break the variety into irreducible components
 - 4: *If noisy:* Use numerical methods (witness sets + symbolic regression)
 - 5: *If exact:* Use symbolic primary decomposition
 - 6: **for** each component **do**
 - 7: Extract candidate axioms (generators or fitted polynomials)
 - 8: **Reason:** Test whether candidates enable derivation of Q
 - 9: *If noisy:* Use theorem prover with existential quantifiers
 - 10: *If exact:* Use algebraic projection via Gröbner bases
 - 11: **end for**
 - 12: **Output:** Candidate axiom sets $\{\hat{A}_{k+1}, \dots, \hat{A}_l\}$ sufficient to explain Q
-

130 **Encode.** Given a collection of known axioms $A_1, \dots, A_k$ and a hypothesis Q , we form the solution
 131 set of the system $\{A_1, \dots, A_k, Q\}$. Each equation is defined over some subset of n real variables,
 132 and so defines a hypersurface in $\mathbb{R}^n$. The intersection of these hypersurfaces defines a geometric
 133 object referred to as the *variety* corresponding to the system, denoted $V(A_1, \dots, A_k, Q)$. This
 134 geometric perspective is useful because algebraic derivability has a geometric interpretation: if Q is
 135 algebraically derivable from $A_1, \dots, A_k$, then Q can be written as an algebraic combination $Q =$
 136 $\sum_{i=1}^k \alpha_i A_i$ where each α_i is a real polynomial. Then Q vanishes whenever all the axioms vanish.
 137 So we can recast the idea of Q being derivable from any set of axioms $A_1, \dots, A_{k+r}$ as a geometric
 138 operation. Our first step in AI-Noether's process then will be to take the known axioms $A_1, \dots, A_k$
 139 and consequence Q and *encode* it into a variety.

140 **Decompose.** The variety corresponding to the equation $xy = 0$ (where x and y are real variables)
 141 is *reducible* as it is the union of two smaller varieties, i.e., solutions of $x = 0$ and of $y = 0$.
 142 We cannot further write either of these smaller varieties as the union of even smaller varieties, so
 143 they are *irreducible*. In other words, $V(xy) = V(x) \cup V(y)$. In writing the variety $V(xy)$ as a
 144 union of two irreducible varieties, we have *decomposed* it. We similarly decompose the variety
 145 $V(A_1, \dots, A_k, Q)$ into its irreducible components representing minimal geometric subsets that can-
 146 not be decomposed further. This process is called *Lasker-Noether Primary Decomposition* [1], and
 147 is geometrically illustrated in Figure 3.

148 If Q is not a consequence of the axioms (i.e., is not derivable from the known axioms), then
 149 $V(A_1, \dots, A_k, Q)$ contains additional irreducible components not present in $V(A_1, \dots, A_k)$. To
 150 illustrate why this matters in the context of abductive inference, assume for example that we only
 151 need one more axiom, A_{k+1} in order to derive Q . Then $Q = \alpha_1 A_1 + \dots + \alpha_{k+1} A_{k+1}$ for unknown
 152 coefficients α_i and unknown A_{k+1} . We are of course interested then in the residual $\alpha_{k+1} A_{k+1}$.
 153 The key observation is that $V(A_1, \dots, A_k, Q)$ is the same variety as $V(A_1, \dots, A_k, \alpha_{k+1} A_{k+1})$.
 154 Similar to the example of $V(xy)$, $V(\alpha_{k+1} A_{k+1})$ is reducible. So intersecting Q with the variety
 155 corresponding to the axioms $V(A_1, \dots, A_k)$ introduces new irreducible components that are di-
 156 rectly related to the residual $\alpha_{k+1} A_{k+1}$. Decomposition isolates these structures. So the next step in
 157 AI-Noether's system is to take the encoded variety and *decompose* it into its irreducible components.
 158 The key difference from traditional symbolic regression-based discovery is that we do not assume
 159 any data for the unknown axiom or any prior knowledge of the variables being fitted.

160 When consequences contain noise - for instance, when Q has been inferred from experimental data
 161 - exact algebraic methods become numerically unstable. In such cases, we employ a numerical
 162 irreducible decomposition [6] via homotopy continuation, which intersects the variety with generic
 163 hyperplanes until a zero-dimensional witness set is obtained, then uses curve tracing to identify
 164 points lying on the same irreducible component. This produces a finite set of sample points on
 165 each component rather than symbolic generators. We then perform symbolic regression on these
 166 witness points to recover polynomial relations that approximately vanish on each component with an
 167 approach similar to vanishing component analysis [56]. The smallest singular value of the resulting
 168 design matrix serves as a quality metric: low values indicate the recovered relation nearly vanishes
 169 on the component, suggesting it governs that portion of the solution space. This numerical workflow,
 170 with the numerical witness set computation and our symbolic regression step, enables decomposition
 171 even when consequences carry numerical error. This workflow is detailed in Appendix C.2




Figure 3: **Geometric visualization of decomposition.** On the left, we have an example of a variety which is the solution to a system of polynomial equations. The middle and right panels show that this variety can be decomposed into irreducible components, which, geometrically, is the process of finding smaller varieties whose union is the original variety.

172 **Reason.** For each irreducible component, we extract its defining polynomials and test whether any
 173 subset $\{\hat{A}_{k+1}, \dots, \hat{A}_l\}$, when added to the known axioms, is sufficient to derive Q . This is done by
 174 forming the variety $V(A_1, \dots, A_k, \hat{A}_{k+1}, \dots, \hat{A}_l)$ and projecting it onto the variables appearing in
 175 Q , illustrated in Figure 4. If this projection lies within the solution set of Q , then Q is implied by
 176 the augmented system. Such sets are returned as candidate missing axioms. If no such set exists,
 177 we conclude that no single missing axiom explains Q . This is the final step of AI-Noether where we
 178 *reason* with each decomposed component to discover which axiom candidates can derive Q .

179 When axioms or consequences contain numerical constants from noisy data, algebraic projection
 180 may be overly sensitive to coefficient errors. As an alternative to algebraic derivability, for ro-
 181 bustness to noise, we use a theorem prover to assess whether the conjecture follows logically from
 182 the augmented background theory (the original axioms together with the abduced ones). Build-
 183 ing on AI-Descartes [16], which introduced existential derivability to handle numerical constants
 184 in data-derived conjectures, we introduce the existential constants also in the abduced axioms. The
 185 conjecture is then deemed derivable when there exists an instantiation of these abstracted constants
 186 that makes it logically implied by the background theory. This extension allows the prover to handle
 187 noisy numerical parameters and provides a principled way to test whether the augmented axioms
 188 suffice to entail the conjecture.




Figure 4: **Geometric visualization of algebraic reasoning.** On the left, we have some surface defined by an axiom (here, $F_c = F_g$ as an example). In the middle, we have an additional equation that defines a new surface. In order to apply the axiom to the new surface, we project the surface onto the space defined by the known axiom(s).




Figure 5: **Illustration of the experimental setup for the Carrier-Resolved Photo-Hall Effect.** AI-Noether successfully recovers the necessary Hall coefficient relation required to prove the relationship between the various parameters of a semiconducting surface.

Using the *encode*, *decompose*, and *reason* steps, AI-Noether handles both exact symbolic computation and noisy numerical scenarios. For a detailed development of the algebraic and logical foundations underlying the three-step pipeline, see Appendix C. For the details of our implementation, see Appendix E.3. Having established the conceptual framework, we now demonstrate AI-Noether’s capabilities across a range of scientific domains. Section 3 presents experimental results organized by discovery type: recovering individual axioms (Section 3.1), discovering multiple axioms simultaneously (Section 3.2 and 3.3), and handling noisy consequences (Section 3.4). These experiments span classical mechanics, special relativity, orbital dynamics, and electromagnetic theory, illustrating the generality of the abductive inference approach.

3 Results

In this section, we present the results from testing our system on various real-world physical theories. We test AI-Noether on 9 systems reported in AI-Hilbert [18] as well as 3 additional systems of varying complexity, the most complex being the results from a recent work on the Carrier-Resolved Photo-Hall Effect [38]. The other two systems include a combined axiomatic system involving Einstein’s relativistic mass, length contraction, and time dilation laws as consequences, and finally a baseline case of a simple harmonic oscillator. We illustrate how AI-Noether performs under different settings using different examples.

3.1 Single Missing Axiom

Example: Carrier-Resolved Photo-Hall Effect

The Carrier-Resolved Photo-Hall effect [38] equation describes the relationship between various parameters of a semiconducting surface, illustrated in Figure 5 and is given by:

$$Q := H - \frac{r e \mu_P^2 [p_0 + \Delta n(1 - \beta^2)]}{\sigma^2} = 0,$$

where $\sigma = e\mu_P[p_0 + \Delta n(1 + \beta)]$ is the conductivity. We assume the axioms

$A_1 : \beta\mu_P - \mu_N$	Mobility ratio definition
$A_2 : \mu_H - r\mu$	Hall factor correction
$A_3 : p_h - p_0 - \Delta p$	Hole concentration balance
$A_4 : n - \Delta n$	Electron concentration tracking
$A_5 : \Delta p - \Delta n$	Charge neutrality condition
$A_6 : \sigma - e p_h \mu_P - e n \mu_N$	Total conductivity sum
$A_7 : H(p_h + \beta n)^2 e - r p_h + r \beta^2 n$	Hall coefficient relation.

Here μ_N and μ_P are electron and hole drift mobilities, $\beta = \mu_N/\mu_P$ is the mobility ratio (encoded by A_1), μ_H is the Hall mobility with Hall scattering factor r and mobility μ related by $r = \mu_H/\mu$ (encoded by A_2), Δn and Δp are the electron and hole photo-carrier densities (equal in steady state

214 by A_5), H is the Hall coefficient, σ is the conductivity, n is the electron density, p_0 is the background
215 hole density, and p_h is the hole density (for a p-type material).

216 The photo-Hall effect measurement begins with the semiconductor in its dark equilibrium state,
217 where only majority carriers (e.g., holes with concentration p_0 in p-type material) contribute to
218 transport, establishing baseline conductivity σ_0 and Hall coefficient H_0 . Upon illumination with
219 above-bandgap photons, electron-hole pairs are generated satisfying charge neutrality $\Delta n = \Delta p$,
220 increasing the total hole concentration to $p_h = p_0 + \Delta p$ and electron concentration to $n \approx \Delta n$.

221 An alternating current (AC) magnetic field from the rotating parallel dipole line (PDL) system de-
222 flects both carrier types via the Lorentz force, generating a Hall voltage across the sample. As light
223 intensity I varies, both conductivity $\sigma(I) = e(p_h\mu_P + n\mu_N)$ and Hall coefficient $H(I)$ evolve non-
224 linearly due to competing contributions from holes and electrons with different mobilities. The key
225 innovation is the relation $\Delta\mu_H = \frac{d(\sigma^2 H)}{d\sigma}$, which extracts the hole-electron mobility difference from
226 the measured transport data.

227 An exact derivation of Q from the axioms, both directly and as an algebraic combination, can be
228 found in Appendix D. However, it suffices to recognize that, because Q is algebraically derivable
229 from the known axioms, we can write

$$Q = \sum_{i=1}^7 \alpha_i A_i$$

230 where the α_i are polynomials (Appendix D) and $i \in \{1, \dots, 7\}$.

231 One of the key axioms in the derivation is axiom A_7 , which relates the Hall coefficient to all the
232 essential quantities. In our setting, we assume that we neither have A_7 nor any of the multipliers
233 $\alpha_1, \dots, \alpha_7$. Using methods described in [38, 91, 16, 18], we can recover Q from data. Encoding
234 the known information $A_1, \dots, A_6, Q$ into a variety, decomposing each generator, and passing it
235 through the reasoning step yields the following polynomials.

Decomposition	
Variety 1	$\langle \sigma, \mu_P, n - \Delta n, \mu_N, \Delta p - \Delta n, p_h - p_0 - \Delta n, \mu_r - \mu_H \rangle$
Variety 2	$\langle \sigma, e, n - \Delta n, \Delta p - \Delta n, p_h - p_0 - \Delta n, \mu_N - \mu_P\beta, \mu_r - \mu_H \rangle$
Variety 3	$\langle n - \Delta n, \Delta p - \Delta n, p_h - p_0 - \Delta n, \mu_N - \mu_P\beta, \mu_r - \mu_H,$ $\mu_N e \Delta n + e \mu_P p_0 + e \mu_P \Delta n - \sigma, e \mu_P p_0 + e \mu_P \Delta n \beta + e \mu_P \Delta n - \sigma,$ $\sigma H(p_0 + \Delta n(1 + \beta)) + r(\mu_N \beta \Delta n - \mu_P(p_0 + \Delta n)),$ $\mu_N^2 \Delta n r + \mu_N \Delta n \sigma H - \mu_P^2 p_0 r - \mu_P^2 \Delta n r + \mu_P p_0 \sigma H + \mu_P \Delta n \sigma H,$ $r \mu_P(p_0 + \Delta n(1 - \beta^2)) - \sigma H(p_0 + \Delta n(1 + \beta)),$ $- e \sigma H(p_0 + \Delta n(1 + \beta)) + r(\sigma(1 - \beta) + e \mu_P \beta p_0),$ $\mu_N e \mu_P p_0 r - \mu_P \beta r \sigma + \mu_P r \sigma - \sigma^2 H, \mu_N^2 e p_0 r - \mu_P \beta^2 r \sigma + \mu_P \beta r \sigma - \beta \sigma^2 H,$ $\mu \sigma^2 H - \mu_N \mu_H e \mu_P p_0 + \mu_H \mu_P \beta \sigma - \mu_H \mu_P \sigma, \mu \beta \sigma^2 H - \mu_N^2 \mu_H e p_0 + \mu_H \mu_P \beta^2 \sigma - \mu_H \mu_P \beta \sigma,$ $\mu p_0 \sigma H + \mu \Delta n \beta \sigma H + \mu \Delta n \sigma H - \mu_H \mu_P p_0 + \mu_H \mu_P \Delta n \beta^2 - \mu_H \mu_P \Delta n,$ $\mu \mu_N \Delta n \sigma H + \mu \mu_P p_0 \sigma H + \mu \mu_P \Delta n \sigma H + \mu_N \mu_H \mu_P \Delta n \beta - \mu_H \mu_P^2 p_0 - \mu_H \mu_P^2 \Delta n,$ $e H(p_0 + \Delta n(1 + \beta))^2 - r(p_0 + \Delta n(1 - \beta^2)) \rangle$
Generators selected by Reasoning module	
$\sigma H(p_0 + \Delta n(1 + \beta)) + r(\mu_N \beta \Delta n - \mu_P(p_0 + \Delta n))$ $r \mu_P(p_0 + \Delta n(1 - \beta^2)) - \sigma H(p_0 + \Delta n(1 + \beta))$ $- e \sigma H(p_0 + \Delta n(1 + \beta)) + r(\sigma(1 - \beta) + e \mu_P \beta p_0)$ $e H(p_0 + \Delta n(1 + \beta))^2 - r(p_0 + \Delta n(1 - \beta^2))$	

Simple Harmonic Oscillator




Figure 6: **Illustration of the experimental setup for a simple harmonic oscillator.** AI-Noether can infer the requirement of the small angle assumption, gravitational component of acceleration, and any of the other essential axioms required to derive the period of a simple harmonic oscillator. This simpler example illustrates the output of AI-Noether's modules.

237 Notice that each of the obtained relationships is the exact Hall coefficient relation with other axiom
 238 substitutions applied. The first equation is A_7 with A_3, A_4, A_5 substituted and multiplied by μ_P .
 239 The second is the same as the first, with A_1 substituted in as well. The third is the same as the
 240 second, except that A_6 is substituted and multiplied by $e\mu_P$. The last is the same as the first, with
 241 A_1 and A_2 substituted.

242 Each of these axioms contains the same information as A_7 . The alternate expressions reflect the fact
 243 that the expression of A_7 is not unique. We could have defined the same relation in multiple forms.

244 Additionally, if we start with an overly strong assumption, such as $\sigma = 0$ or $\mu_P = 0$ from Variety 1,
 245 we could also derive $Q = 0$. We provide an additional layer of filtering for such cases when using
 246 projections to ensure that we do not eliminate more variables than necessary. A discussion of this
 247 can be found in Appendix E.2.

248 **Example: Simple Harmonic Oscillator**

249 For a simpler and more illustrative example, consider the case of a simple harmonic oscillator with a
 250 point mass m and length L , with a force F acting on it creating a small angle θ , illustrated in Figure
 251 6. Taking g as the acceleration constant, at j iterations, we have the equation that describes the time
 252 period T_j given by:

$$Q := T_j - 2\pi j \sqrt{\frac{L}{g}}$$

253 whose polynomial form, $gT_j^2 - 4\pi^2 j^2 L = 0$, can be derived (either traditionally or as an algebraic
 254 combination) from the following axioms:

$A_1 : a_d - g \sin \theta$	Gravitational component of acceleration
$A_2 : d - L\theta$	Height of oscillator
$A_3 : T\omega - 2\pi$	Frequency-period relation
$A_4 : d\omega^2 - a_d$	Acceleration equation
$A_5 : T_j - jT$	Time period over j iterations
$A_6 : \sin \theta - \theta$	Small angle assumption

255 where a_d is the acceleration of the object, ω is the frequency, and T is the time period for one
 256 oscillation. A small detail here is that in the algebraic context, like in [18], we treat $\sin \theta$ and π
 257 symbolically here. One of the key axioms here is A_1 , the gravitational component of acceleration.
 258 If we remove this assumption from our system, then we can still recover Q from data, as seen in
 259 [91]. Applying AI-Noether to Q and the known axioms results in the following output:

Problem	Axioms	Success Rate	Variables	Max Degree
Hagen–Poiseuille Equation	4	100%	9	4
Kepler	5	100%	8	4
Time Dilation	5	100%	8	4
Escape Velocity	5	100%	9	4
Light Damping	5	100%	10	8
Neutrino Decay	5	100%	8	2
Simple Harmonic Oscillator	6	100%	11	5
Carrier-Resolved Photo-Hall	7	100%	15	6
Relativistic Laws	7	71%*	12	5
Hall Effect	7	100%	16	4
Inelastic Relativistic Collision	9	100%	11	6
Compton Scattering	10	100%	15	6

Table 1: Summary of Axiom Recovery Performance by Problem. In each case, we iterate through every axiom, remove it from the theory, and attempt to recover the axiom with the remaining information. In almost every case, we are able to recover every axiom. *Indicates partial recovery due to recovering alternate axioms. We discuss this particular case further in Section 3.3.

Decomposition	
Variety 1	
260 $\langle \theta - \sin \theta, 2j\pi - \omega T_j, Tj - T_j, \sin \theta L - d, dg - adL, \sin \theta g - ad, $	
$T\omega - 2\pi, 2\omega L\pi - Tg, 2\omega d\pi - adT, \omega^2 L - g, T^2 g - 4L\pi^2, \omega^2 d - ad, adT^2 - 4d\pi^2 \rangle$	
Generators selected by Reasoning module	
$\sin \theta g - ad$	

261 Results summary on other problems

262 We report the performance of AI-Noether on each test system in Table 1. In each case, we iterate
263 through the axioms, omit an axiom from the system, and test if we can recover the missing axiom
264 while not making any assumption about it. In almost every case, we were able to recover the missing
265 axiom. Appendix D.1 provides a further discussion of some of these systems. We will discuss the
266 problem of the relativistic laws in Section 3.3, where we only partially recover missing axioms
267 correctly while recovering alternate axioms for some cases.

268 3.2 Multiple Missing Axioms

269 We now examine the problem of recovering more than one axiom to explain a hypothesis Q . In this
270 case, we have

$$Q = \sum_{i=1}^{k+r} \alpha_i A_i$$

271 However, we only have access to $A_1, \dots, A_k$. We seek a candidate set of axioms that, together with
272 $A_1, \dots, A_k$, suffices to derive Q .

273 Example: Kepler’s third law

274 Consider Kepler’s third law, which relates the orbital period of two celestial bodies to their distances
275 and masses. It can be written as

$$Q := p - \sqrt{\frac{(d_1 + d_2)^3}{G(m_1 + m_2)}} = 0,$$

276 where p is the orbital period (scaled by a factor of 2π), m_1, m_2 the masses, and d_1, d_2 their distances
277 from the center of mass. As in [18], we will scale units to have $G = 1$ for consistency with their




Figure 7: **Illustration of the axioms required to derive Kepler's third law of planetary motion.** AI-Noether detects the existence of gravitational force, an essential axiom required to derive Kepler's third law. When multiple axioms are missing, AI-Noether can also recover coupled versions of the missing axioms.

framework, though this is not necessary. For circular orbits, we use the following axioms (with p non-dimensionalized so that $\omega p = 1$), illustrated in Figure 7:

$A_1 : (d_1 + d_2)^2 F_g - m_1 m_2$	Gravitational Force
$A_2 : F_c - m_2 d_2 \omega^2$	Centrifugal Force
$A_3 : F_g - F_c$	Newton's third law
$A_4 : \omega p - 1$	Frequency-period relation
$A_5 : m_1 d_1 - m_2 d_2$	Center of mass equation

to derive

$$(m_1 + m_2)p^2 - (d_1 + d_2)^3 = 0,$$

which is Kepler's law in polynomial form.

Assume that we are missing axioms A_3 and A_4 - the relationship between centrifugal and gravitational force and the relationship between frequency and period. Then, encoding the known information - equations A_1 , A_2 , A_5 , and Q - our system returns the following relevant component of the primary decomposition and selected generator from the reasoning module:

Selected varieties from Decomposition module	
Variety 1	$\langle m_1 d_1 - m_2 d_2, \mathbf{F_g p^2 - m_2 d_2}, F_g(d_1^2 + 2d_1 d_2 + d_2^2) - m_1 m_2, d_1^3 + 2d_1^2 d_2 + d_1 d_2^2 - m_2 p^2, m_1 p^2 - d_1^2 d_2 - 2d_1 d_2^2 - d_2^3, F_c - \omega^2 m_2 d_2, F_c p^2 - \omega^2 d_1^3 d_2 - 2\omega^2 d_1^2 d_2^2 - \omega^2 d_1 d_2^3 \rangle$
Generators selected by Reasoning module	
	$\mathbf{F_g p^2 - m_2 d_2}$

We return $\mathbf{F_g p^2 - m_2 d_2}$. This equation is the same as the centrifugal force equation A_2 but with the missing relationships $F_c = F_g$ and $\omega = \frac{1}{p}$ substituted in. Therefore, the missing information is discovered, since equation A_2 is assumed to be known, although it is not decoupled. To go one step further, we can remove three out of the five axioms: equations A_2 , A_3 , and A_4 , leaving us with only the gravitational force equation A_1 and the center of mass equation A_5 . Similarly, running our system with the Kepler hypothesis gives us a similar component:

Pairs of missing axioms		Triples of missing axioms	
Missing Axioms	Recovered	Missing Axioms	Recovered
{G, cf}	X	{G, cf, C}	X
{G, C}	✓	{G, cf, wp}	X
{G, wp}	X	{G, cf, cm}	X
{G, cm}	✓	{G, C, wp}	X
{cf, C}	✓	{G, C, cm}	✓
{cf, wp}	✓	{G, wp, cm}	X
{cf, cm}	X	{cf, C, wp}	✓
{C, wp}	✓	{cf, C, cm}	X
{C, cm}	✓	{cf, wp, cm}	X
{wp, cm}	✓	{C, wp, cm}	X

Table 2: Recovery results on the Kepler problem for various subsets of missing axioms. Shorthand: G = gravitational force $(d_1 + d_2)^2 F_g = m_1 m_2$, cf = centrifugal force $F_c = m_2 d_2 \omega^2$, C = force balance $F_g = F_c$, wp = frequency-period relation $\omega p = 1$, cm = center of mass $m_1 d_1 = m_2 d_2$.

Selected varieties from Decomposition module	
Variety 1	$\langle m_1 d_1 - m_2 d_2, \mathbf{F}_g \mathbf{p}^2 - m_2 \mathbf{d}_2, F_g(d_1^2 + 2d_1 d_2 + d_2^2) - m_1 m_2, d_1^3 + 2d_1^2 d_2 + d_1 d_2^2 - m_2 p^2, m_1 p^2 - d_1^2 d_2 - 2d_1 d_2^2 - d_2^3 \rangle$
Generators selected by Reasoning module	
$\mathbf{F}_g \mathbf{p}^2 - m_2 \mathbf{d}_2$	

We get the same axiom candidate. However, in the previous case, where we knew the centrifugal force equation, we now assume that we do not. So we are not only recovering the relationships $F_c = F_g$ and $\omega p = 1$, but also the centrifugal force axiom itself $F_c = m_2 d_2 \omega^2$ with the former relations coupled with it. A criterion for exactly recovering axioms is given in section 3.5.

Table 2 shows the results of attempting to recover various combinations of pairs and triples of axioms for Kepler’s third law.

Results summary on other problems. Across twelve benchmark systems, we systematically removed axiom pairs and triples and attempted to recover explanations for Q and have summarized our results in Table 3. Success rates vary considerably by problem complexity and number of missing axioms. For 2-axiom recovery, success rates range from 10% (Time Dilation, Light Damping) to 67% (Kepler, Simple Harmonic Oscillator), with a median of 47%. For 3-axiom recovery, performance degrades substantially: success rates drop to 0–40%, with most systems achieving 10–25% recovery. The best-performing systems are Kepler (70% pairs, 20% triples), Simple Harmonic Oscillator (67% pairs, 20% triples), and Relativistic Laws (81% pairs, 69% triples). The Hall Effect and Compton Scattering systems, despite their larger axiom sets (9–10 axioms), achieve moderate success (53% and 38% for pairs; 26% and 21% for triples, respectively). The Carrier-Resolved Photo-Hall system (62% pairs, 42% triples) demonstrates the method’s scalability to complex systems with 7 axioms and 15 variables. Overall, recovery difficulty increases with the number of missing axioms and, to a lesser extent, with system size, though the relationship is non-monotonic.

Clearly having more information is helpful, and so the success rate is higher in cases where we were only missing two axioms rather than three. The highest success rate was for the Relativistic Laws (81% pairs, 69% triples), likely because this is the only system with not one but three consequences that we are drawing information from, as we will see in Section 3.3.

3.3 Multiple Consequences

As we will see in Section 3.5, having multiple consequences can help us recover decoupled axioms in the case where multiple axioms are missing. We demonstrate AI-Noether’s ability to handle multiple related consequences simultaneously, exploring relativistic phenomena. Here, each consequence Q_i

Problem	Vars	Axioms	Pairs	Triples	2-Ax Success	3-Ax Success
Hagen-Poiseuille	9	4	6	4	3/6	1/4
Kepler	8	5	10	10	7/10	2/10
Time Dilation	8	5	10	10	1/10	1/10
Escape Velocity	9	5	10	10	8/10	4/10
Light Damping	10	5	10	10	1/10	0/10
Neutrino Decay	8	5	10	10	3/10	1/10
Simple Harmonic Osc.	10	6	15	20	10/15	4/20
Carrier-Resolved PH	15	7	21	35	13/21	15/35
Relativistic Laws	12	7	21	35	17/21	24/35
Hall Effect	16	7	21	35	13/21	17/35
Inelastic Relativistic Collision	11	9	36	84	12/36	15/84
Compton Scattering	15	10	45	120	17/45	25/120

Table 3: Summary of 2- and 3-Axiom Tuple Recovery Performance by Problem. Pairs = $\binom{n}{2}$, Triples = $\binom{n}{3}$ where n is the number of axioms.

is the sum

$$Q_i = \sum_{j=1}^{k+r} \alpha_{i,j} A_j$$

where each $\alpha_{i,j}$ is unknown and $A_{k+1}, \dots, A_{k+r}$ are unknown as well.

Example: Relativistic Laws

Consider Einstein's relativistic mass, length contraction, and time dilation formulae, which describe how mass increases and lengths contract for objects moving at relativistic speeds, illustrated in Figure 8. These can be expressed in polynomial form as:

$$\begin{aligned} Q_1 &:= -c^2 f^2 + f_0^2 (c^2 - v^2) = 0 && \text{Time dilation} \\ Q_2 &:= -m_0^2 c^2 + m^2 (c^2 - v^2) = 0 && \text{Relativistic mass} \\ Q_3 &:= L^2 c^2 - L_0^2 (c^2 - v^2) = 0 && \text{Length contraction} \end{aligned}$$

These consequences arise from a unified set of axioms describing light clock geometry in different reference frames. The axiom structure, illustrated in Figure 9, reveals how different subsets of axioms are required to derive each consequence: axioms A_1 – A_4 form the foundation for time dilation (Q_1), with additional axioms A_5 – A_6 needed for relativistic mass (Q_2), and additional axiom A_7 required for length contraction (Q_3).

$$\begin{aligned} A_1 &: f_0 t_0 - 1 && \text{Frequency-period (rest)} \\ A_2 &: f t - 1 && \text{Frequency-period (moving)} \\ A_3 &: c t_0 - 2L_0 && \text{Perpendicular light travel (rest)} \\ A_4 &: \frac{c^2 t^2}{4} - L_0^2 - \frac{v^2 t^2}{4} && \text{Light path geometry (moving)} \\ A_5 &: m_0 u_0 - m u_y && \text{Transverse momentum conservation} \\ A_6 &: u_0 t_0 - u_y t && \text{Transverse distance invariance} \\ A_7 &: t(c^2 - v^2) - 2Lc && \text{Parallel light travel time} \end{aligned}$$

Physical Interpretation.

Axioms A_1 and A_2 define the period-frequency relationship in the rest and moving frames, respectively. Axiom A_3 encodes that light travels distance $2L_0$ in time t_0 for a perpendicular round trip in the rest frame. Axiom A_4 captures the geometric constraint from a perpendicular light clock: a




Figure 8: **Illustration of the multiple experimental setups required to establish each of the relativistic laws - time dilation, length contraction, and relativistic mass.** The left two images illustrate a light clock experiment in a stationary and moving frame of reference required for time dilation and length contraction. The right two images illustrate an accelerated mass experiment in a stationary and moving frame of reference required for deriving relativistic mass. AI-Noether can extrapolate axioms from multiple consequences with more complex derivability dependency structures.

moving observer sees light travel along the hypotenuse of a right triangle, relating the rest-frame quantities to the moving-frame time and velocity. Axioms A_5 and A_6 ensure conservation of momentum and distance in the transverse (perpendicular) direction. Finally, axiom A_7 describes light traveling parallel to the direction of motion, accounting for the asymmetric forward and backward travel times.

The derivation structure proceeds as follows. From A_3 and A_4 , we obtain $t^2(c^2 - v^2) = c^2 t_0^2$, the basic time dilation relation. Substituting the frequency-period relations from A_1 and A_2 yields Q_1 . For relativistic mass (Q_2), we additionally use A_6 to relate transverse velocities through the time dilation factor, then apply momentum conservation from A_5 to obtain the mass transformation. For length contraction (Q_3), axiom A_7 provides the parallel light travel constraint, which when combined with the time dilation relation and A_3 gives the length contraction formula.

Abductive inference with multiple consequences. We test AI-Noether's ability to recover missing axioms in two scenarios: single-axiom removal and multi-axiom removal with multiple consequences.

Single-axiom case. We saw in Table 1 that the combined relativistic laws were the only case where there were two axioms which we do not recover exactly. When removing individual axioms and attempting recovery using a single consequence, AI-Noether exhibits interesting behavior for axioms A_4 and A_7 . Instead of recovering the original quadratic forms, the system returns factorized alternatives. For instance, when A_4 is removed, AI-Noether recovers the pair:

$$c - 2f_0^2 L_0 t \quad \text{and} \quad c + 2f_0^2 L_0 t$$

whose product $(c - 2f_0^2 L_0 t)(c + 2f_0^2 L_0 t) = c^2 - 4f_0^4 L_0^2 t^2$ is algebraically equivalent to the original constraint after substituting $f_0 t_0 = 1$ and $ct_0 = 2L_0$. A similar factorization occurs for A_7 . This factorization phenomenon reflects the fundamental algebraic structure: the solution variety of the quadratic constraint decomposes into a union of two linear varieties (hyperplanes), each corresponding to one factor.

This factorization phenomenon reflects the structure of the primary decomposition: when A_4 is removed, the variety $V(\{A_1, A_3, A_5, A_6, A_7, Q\})$ decomposes into components whose generators include these linear factors. These factors do not individually reproduce the physics of A_4 , but they appear in the decomposition because they define subvarieties that, when intersected with the other axioms, collectively constrain the solution space in a way that allows Q to be derived. This illustrates a subtle but important point: abductive inference may return algebraic expressions that are structurally different from the original axioms yet sufficient to derive the target consequence when combined with background knowledge. The system identifies what is algebraically necessary to close the gap between axioms and consequence, which may differ from the canonical physical formulation of the missing law.

Multi-axiom case with multiple consequences.




Figure 9: **Derivation structure for relativistic consequences.** The graph shows how axioms A_1 – A_7 combine through intermediate results to derive the three consequences. Solid blue boxes represent axioms, yellow boxes show intermediate algebraic steps, and green boxes are the final consequences. The derivation of Q_1 (time dilation) uses axioms A_1 – A_4 ; Q_2 (relativistic mass) additionally requires A_5 – A_6 ; and Q_3 (length contraction) further needs A_7 .

371 We now remove three axioms simultaneously: A_2 (required for all consequences), A_3 (required for
372 Q_1 and Q_3), and A_5 (required for Q_2). Running AI-Noether with all three consequences Q_1 , Q_2 ,
373 Q_3 and the remaining axioms yields:

Generators selected by reasoning module	
	$ft - 1$
374	$ct_o - 2L_o$
	$m_o u_o - m u_y$

375 Notably, the system recovers *all three* missing axioms exactly— A_2 , A_3 , and A_5 —without the cou-
376 pling artifacts observed in Section 3.2. This demonstrates a key advantage of having multiple con-
377 sequences: the additional constraints allow the decomposition to isolate the individual missing axi-
378 oms rather than returning their coupled combinations. As we formalize in Section 3.5, multiple
379 consequences increase the algebraic constraints on the variety, forcing the primary decomposition
380 to separate axioms that would otherwise appear only in entangled form. This multi-consequence
381 strategy is essential for recovering decoupled axioms and represents a significant advantage over
382 single-consequence abductive inference.

383 3.4 Noisy Consequences

384 We now consider the effect of noise in the consequences. For each of the problems in this section,
385 as well as with the remaining test problems, we ran AI-Hilbert [18] on synthetically generated
386 datasets for Q (respectively, Q_i for relativistic laws) with standard Gaussian noise added to the data
387 proportional to the means of each column with a noise level ϵ . We therefore obtain perturbations of
388 each consequence rather than the true consequences themselves. We illustrate this with the following
389 two examples of Kepler’s third law and relativistic laws continued from earlier.

390 Example with single axiom correction: Kepler, continued

391 Taking the previous example of Kepler’s third law, assume that the gravitational force axiom $F_g(d_1 +$
392 $d_2)^2 - m_1 m_2$ is missing. With a noise level of $\epsilon = 1\%$, the resulting consequence Q that was
393 obtained is given by:

$$-1.01387623m_2d_2^3 - 2.02927363m_2d_1d_2^2 - 1.01837465m_2d_1^2d_2 + 1.0142876m_1m_2p^2 = 0$$

394 This, however, does not come with a certificate of derivability, since the remaining axioms cannot
395 derive this consequence. We assume as before that we have no data for the missing axiom, so we

cannot use a traditional symbolic regression approach to determine the missing axiom. Algebraically computing a primary decomposition results in obtaining Q as the only axiom candidate. Since algebraic decomposition uses exact computation, we can only handle exact derivability via algebraic projection. Therefore, we turn to a numerical irreducible decomposition that attempts to generate approximate points on the irreducible components, and then perform symbolic regression on these points. The data generated is defined over all the variables of the system, not just those in the consequence or any single axiom; therefore, there is a hyperparameter of which variables to fit to that is to be tuned. Of course, one could choose various hyperparameters to get a family of axiom candidates. Fitting the data generated in our system yields the following axiom candidate, which we then verify using our reasoning module.

Axiom generated by Decomposition and Reasoning modules

$$-0.502463F_g d_1^2 - 1.012319F_g d_1 d_2 - 0.497537F_g d_2^2 + 0.497536m_1 m_2$$

Notice that this indeed is a noisy perturbation of the missing axiom: $F_g(d_1^2 + 2d_1 d_2 + d_2^2) - m_1 m_2$.

Example: Relativistic laws with noisy consequences

We test the robustness of our approach by combining multiple challenges: noisy consequences, multiple missing axioms, and multiple consequences. Running AI-Hilbert [18] to generate noisy versions of the three relativistic consequences from synthetically generated data with 1% Gaussian noise (standard deviation proportional to the mean of each column) gives the following perturbations of the relativistic laws (rounded to four decimal places for brevity):

$$\tilde{Q}_1 := -1.0089c^2 f^2 + 0.9986f_0^2 c^2 - 0.9956f_0^2 v^2 = 0 \quad (\text{Time dilation})$$

$$\tilde{Q}_2 := -0.9876m_0^2 c^2 + 1.0134m^2 c^2 - 1.0089m^2 v^2 = 0 \quad (\text{Relativistic mass})$$

$$\tilde{Q}_3 := 0.9945L^2 c^2 - 1.0123L_0^2 c^2 + 1.0067L_0^2 v^2 = 0 \quad (\text{Length contraction})$$

We removed three axioms that span different roles in the derivation structure (recall Figure 9): A_2 ($ft - 1$), required for time dilation; A_5 ($m_0 u_0 - m u_y$), required for relativistic mass; and A_4 ($\frac{c^2 t^2}{4} - L_0^2 - \frac{v^2 t^2}{4}$), the geometric constraint from the perpendicular light clock. Without these, we again do not get any certificate of derivability of the laws.

This selection is also particularly interesting because, as we saw in Section 3.3, axiom A_4 exhibits factorization behavior in the noiseless single-axiom case—the algebraic decomposition returns linear factors rather than the original quadratic form. This raises the question: can numerical methods recover what algebraic methods transform?

Using numerical irreducible decomposition (Section C.2) followed by symbolic regression on witness points sampled from each component, our system recovers noisy approximations of all three missing axioms:

Axioms recovered via numerical decomposition and automated reasoning

$$1.0000 ft - 0.6599$$

$$-0.9899 m_0 u_0 + 1.0000 m u_y$$

$$-0.2525 c^2 t^2 + 0.2451 v^2 t^2 + 1.0000 L_0^2$$

The recovered axioms closely approximate the true missing relationships despite 1% noise in the consequences. Most notably, axiom A_4 is recovered in its original quadratic form with coefficient errors of approximately 1–2%, rather than the factorized linear expressions produced by algebraic decomposition in the noiseless case. This highlights a fundamental difference between algebraic and numerical abductive inference:

- **Algebraic decomposition** operates purely symbolically, decomposing varieties into unions of irreducible components. When A_4 is missing, the variety structure naturally factors into linear subvarieties, yielding generators like $c \pm 2f_0^2 L_0 t$. While algebraically sufficient for derivability, these do not match the canonical physical form.




Figure 10: **Results on running AI-Noether on various problems** with 1,2, and 3 axioms missing with noisy consequences obtained from running AI-Hilbert [18] with varying levels of noise in the data. For each problem, we plot the normalized coefficient distance [18, 19] and the noise level for successful runs.

- **Numerical decomposition** samples witness points on each irreducible component, then performs symbolic regression to fit polynomials that approximately vanish on those points. Because we regress using the monomial support of the original missing axiom as a basis (here, $\{c^2t^2, v^2t^2, L_0^2, 1\}$), the method directly recovers a noisy version of A_4 rather than its factorization. The witness points constrain all monomials simultaneously, bypassing the algebraic factorization that occurs in the symbolic setting.

This demonstrates a practical advantage of the numerical approach. By operating on sampled data rather than symbolic manipulation, it can recover axioms in forms that more closely match their canonical physical formulations. The choice of monomial basis serves as an inductive bias that guides the regression toward physically meaningful parameterizations, while preserving the principled algebraic foundation of the decomposition step. The complete numerical workflow—witness set computation via homotopy continuation, robust symbolic regression with outlier removal, and coefficient normalization—is detailed in Section C.2.

Results summary on other problems

We summarize our results in Figure 10. For each problem, we apply the numerical irreducible decomposition (details in Section C.2) to obtain witness points on each component of the variety defined by the remaining axioms and target. We then perform symbolic regression to recover the missing axiom by fitting a polynomial in the null space of the Vandermonde-like matrix constructed from these points. To improve robustness against numerical artifacts and outliers introduced by the homotopy continuation, we employ iterative outlier removal: points with the largest residuals (top 10% by default) are discarded, and the regression is repeated until convergence. The recovered polynomial is normalized so that its largest coefficient has magnitude one, and we apply the same normalization to the ground-truth axiom. The coefficient distance reported is the ℓ_2 norm between these normalized coefficient vectors, following [18]; since $p = 0$ and $-p = 0$ define the same constraint, we report the minimum of $\|c_{\text{fit}} - c_{\text{true}}\|$ and $\|c_{\text{fit}} + c_{\text{true}}\|$.

Failure cases in Figure 10 arise from two sources. (i) Timeouts during witness set computation or symbolic regression, which we cap at two hours by default; this time constraint primarily affected the Compton scattering results and is detailed in Appendix E.3. (ii) Cases where the reasoning module (Section C.2) cannot verify that the recovered axiom, combined with the remaining axioms, implies the target consequence. We give a discussion on the criteria for when an axiom is recoverable in the noiseless and noisy case in Section 3.5.

3.5 Criteria for Recoverability of Axioms

Recovering axioms without noise in the consequences

The theorem below provides a sufficient criterion for recovering missing axioms. In order to state the technical details, we will use terminology from algebraic geometry. A detailed overview of all terminology relevant to the core method can be found in Appendix C.

First, we require the notion of the height of an ideal. For a prime ideal P , we define its height, $\text{ht}(P)$, as the length n of the longest chain $P_0 \subsetneq P_1 \subsetneq \dots \subsetneq P_n = P$ of prime ideals properly contained in P . For a general ideal, its height is the minimum of the heights of its associated primes, which is the same as the minimum of the heights of all primes that contain the ideal.

Theorem 1: Common minimal associated prime

Given ideals $I \subset J \subsetneq \mathbb{R}[x_1, \dots, x_n]$ such that $\text{ht}(I) = \text{ht}(J)$, then I and J have a common minimal associated prime.

Proof : Theorem 1

Let P be an associated prime for J such that $\text{ht}(P) = \text{ht}(J)$. Then we have that $I \subset J \subset P$. Since $\text{ht}(I) = \text{ht}(P)$, there is no prime Q such that $I \subset Q \subsetneq P$. Otherwise $\text{ht}(I) \leq \text{ht}(Q) < \text{ht}(P)$. But then P is a minimal prime containing I and thus a minimal associated prime for I .

Example: Let J be the ideal of the full set of axioms $\langle A_1, \dots, A_{k+1} \rangle$ (so assume only one axiom A_{k+1} is missing) and let Q be a consequent of J . Now treat A_{k+1} as a missing axiom and let I be the ideal generated by the known axioms and the consequence $\langle A_1, \dots, A_k, Q \rangle$. If I and J have the same height, then we know that they have a common associated prime which must contain the missing axiom A_{k+1} . Under the usual situation where Q and A_{k+1} are not contained in any of the associated primes for $\langle A_1, \dots, A_k \rangle$, then we have that $\text{ht}(\langle A_1, \dots, A_k, Q \rangle) = \text{ht}(\langle A_1, \dots, A_{k+1} \rangle)$, the theorem applies, and we can examine the associated primes for I to recover the missing axiom A_{k+1} .

Intuitive interpretation: Geometrically, each axiom defines a hypersurface that reduces the dimension of the solution space by at most one. Removing an axiom removes a "cut," increasing the variety's dimension. Adding a hypothesis Q introduces a new cut. The theorem states that if the dimension lost by removing axioms is exactly recovered by adding hypotheses, then the decomposition of I will reveal a component containing the missing axiom. In other words, the hypotheses carry enough information to "reconstruct" the missing constraints.

Typically, if multiple axioms are missing, adding back a single phenomenon Q will not be sufficient, so in general we will need to add back several consequent phenomena $\{Q_1, \dots, Q_h\}$ in order to generate an ideal which is guaranteed to have a component that contains all the missing axioms.

Our missing axiom construction attempts to find axioms that are sufficient to derive Q , i.e., whether or not Q is contained in the ideal generated by the axioms. If we want to decide whether or not a set of axioms implies Q , given the ideal I generated by the axioms, we are asking whether $V(I) \subset V(Q)$. To make this test effective, we need to construct the biggest ideal with the same set of zeros, this is $I(V(I))$. So the test whether a set of axioms A implies Q is whether $Q \in I(V\langle A \rangle)$. If our varieties are defined over the reals, then $I(V\langle A \rangle)$ can be computed as the real radical of the ideal $\langle A \rangle$.

Recovering axioms with noise in the consequences

When the consequences are noisy, then it is no longer true that removing an axiom and adding a noisy consequent produces an ideal which is exactly contained in the original axiom ideal. It is in some sense "nearby" to an ideal which is so contained. To handle this, we replace exact ideal decomposition with numerical ideal decomposition [6]. This produces sample points on each component of the perturbed ideal. Assuming the perturbation is small and generic, we expect sample points for each numerical component to be approximate zeros of one of the components of the true ideal. [15]. When this is the case we can use symbolic regression on the sample points generated by numerical decomposition to generate approximations to the missing axiom.

3.6 Comparison with State-of-the-Art

We evaluate our abductive reasoning system against two representative baselines: 1) solver-based abduction using cve5 with SyGuS grammars, and 2) LLM-based abduction using GPT-5 Pro constrained to output polynomial hypotheses. All methods are tested on a common benchmark of algebraic physics systems, where one or more axioms are removed, and the goal is to infer a missing axiom that will provide derivability of a target conjecture. Each proposed axiom set is checked by the same symbolic verification pipeline in Macaulay2. Our experiments are both on clean and noisy data for each problem in our benchmark set, and use multiple levels of axiom removal.

Overall, our system demonstrates substantially stronger performance than both baselines (see Table 4). It succeeds in 97% of single-axiom cases where both baselines achieved 0% (LLM) or timed out (cvc5), including cases where the missing axiom involves numerical constants or requires longer multi-step symbolic reasoning. Moreover, cvc5 generally fails to generate a polynomial within the allowed grammar, as it exceeds the runtime we set. GPT-5 Pro can generate syntactically reasonable expressions but doesn’t yield axioms that satisfy the strict algebraic membership or literal-appearance criteria. Instead, our approach, in most cases, recovers minimal, correct, and theory-compatible axioms and remains robust to noise.

More details about our experiments comparing AI-Noether with the state of the art, along with ablations and runtime statistics, can be found in Appendix F.

Method	Single Axiom	Multi Axiom	Data-Free	Noise Robust	Var. Agnostic	Explainable
Traditional SR	✗	✗	✗	✗	✗	✗
Integrated SR	✗	✗	✗	✓	✓	✓
LLM (GPT-5 Pro)	0%	0%	✓	✗	✓	✓
cvc5	t-out	t-out	✓	✗	✗	✓
AI-Noether	97% ✓	49% ✓	✓	✓	✓	✓

Table 4: **Comparison of abductive and symbolic regression methods.** Green ticks (✓) indicate capability support; red crosses (✗) indicate absence. “Data-Free” means the method operates without numeric data or prior parameter information, and “Var. Agnostic” indicates invariance to variable naming. “t-out” means that the method ran out of time without producing a solution.

4 Final Remarks

Code Availability

An open-source implementation of AI-Noether as well as all the benchmark problems we tested on are available at <https://github.com/IBM/AI-Noether>. Algorithm implementation details are provided in Appendix E.3. The repository includes documentation for reproducing all experiments presented in this paper.

Limitations of AI-Noether

AI-Noether is a purely algebraic system for generating a minimal set of missing axioms from a polynomial hypothesis and a set of background theory expressible as polynomial equations. We employ the same techniques as in AI-Hilbert to deal with non-algebraic functions, derivatives and integrals. The first method is that we simplify integrals and derivatives as much as possible and then assign new variables to any remaining integrals, derivatives and non-algebraic functions. The second is that we treat derivatives symbolically, treating them as indeterminates in a polynomial. For trigonometric functions (as in the case of the simple harmonic oscillator, light damping, and Compton scattering), we can consider trigonometric polynomials by looking at the corresponding axioms as polynomials in the respective trig functions. This becomes more challenging, however, when dealing with compositions of trig functions with other polynomial functions. It will be interesting therefore to extend AI-Noether to handle cases when these techniques are not sufficient to derive the consequences from the full set of axioms.

Moreover, AI-Noether can recover a missing axiom only if each variable appearing in that axiom also occurs in at least one of the available axioms or in the target hypothesis. If a missing axiom involves latent variables that are entirely absent from the observed theory and consequences, the current framework cannot identify or introduce them automatically. However, as seen in Section 3.2 with Kepler’s third law and multiple missing axioms, AI-Noether in cases can recover information about variables involved in multiple missing axioms by discovering coupled versions of the missing axioms and variables.

Finally, although effective on a wide range of benchmark systems, AI-Noether relies on algebraic decomposition, and symbolic verification whose computational cost can grow rapidly with the number of variables, axioms, and polynomial degrees, depending on the system.

Future Directions

In this study, we have focused on developing what we believe is a first-of-its-kind abductive inference system and testing it on existing laws of nature. A natural next step is therefore to employ AI-Noether in scientific domains where existing theories are incomplete or internally inconsistent, with the aim of supporting genuinely novel discoveries and of discovering new axioms as substantive scientific contributions that extend or revise existing theories. We also aim to extend AI-Noether beyond equations over a polynomial basis to handle differential operators and inequality constraints, enabling abductive inference across a broader range of scientific domains.

Conclusion

This work introduces AI-Noether, an automated abductive reasoning system that connects AI-generated scientific hypotheses with established theoretical knowledge. Given a set of background axioms and a hypothesis not supported by existing theory, AI-Noether discovers a minimal set of additional axioms that reconcile the hypothesis with the theory, provided both are expressible as polynomials. By combining algebraic-geometry tools with symbolic verification, the system provides candidate explanations and formal guarantees for deriving target hypotheses from the augmented axioms.

AI-Noether not only demonstrates effectiveness in establishing a breadth of known scientific laws but also positions axiomatic discovery as a fundamental form of scientific inquiry. AI-Noether supports the discovery of novel hypotheses (e.g., those generated by data-driven or theory-based methods) by determining whether and how to integrate them into coherent theoretical frameworks. In this sense, the system advances the scientific method by explicitly addressing the problem of identifying and explaining the underlying principles governing the environment in which a phenomenon occurs.

We demonstrated the effectiveness of AI-Noether across a broad set of physical systems, including classical mechanics, electromagnetism, and semiconductor physics. We showed that it robustly recovers missing axioms under both noiseless and noisy conditions. Comparisons against solver-based methods and language model baselines highlight our approach’s advantages in correctness and formal verifiability. More broadly, applying AI-Noether to domains where current theories are incomplete or internally inconsistent points toward a shift in the scientific method itself: one in which abductive inference not only explains existing discoveries better, but also uncovers new principles that reveal the fundamental structure of the world around us.

References

- [1] Michael F. Atiyah and I. G. Macdonald. *Introduction to Commutative Algebra*, volume 361 of *Addison-Wesley Series in Mathematics*. Addison-Wesley, Reading, MA, 1969. Russian translation published as *Vvedenie v kommutativnuju algebru* (1972).
- [2] Pierre Baldi, Peter Sadowski, and Daniel Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature Communications*, 5(1):4308, 2014.
- [3] Haniel Barbosa, Clark Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, et al. cvc5: A versatile and industrial-strength smt solver. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 415–442. Springer, 2022.
- [4] Haniel Barbosa, Chantal Keller, Andrew Reynolds, Arjun Viswanathan, Cesare Tinelli, and Clark Barrett. An interactive SMT tactic in Coq using abductive reasoning. In *Proceedings of 24th International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR ’23)*, volume 94 of *EPiC Series in Computing*, pages 11–22. EasyChair, 2023.
- [5] Clark Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. *Satisfiability Modulo Theories*, volume 185. IOS Press, 2009.

- [6] Daniel J. Bates, Paul Breiding, Tianran Chen, Jonathan D. Hauenstein, Anton Leykin, and Frank Sottile. Numerical nonlinear algebra, 2024.
- [7] Daniel J. Bates, Jonathan D. Hauenstein, Andrew J. Sommese, and Charles W. Wampler. Bertini: Software for numerical algebraic geometry. Available at bertini.nd.edu with permanent doi: [dx.doi.org/10.7274/R0H41PB5](https://doi.org/10.7274/R0H41PB5).
- [8] Elena Bellodi, Marco Gavanelli, Riccardo Zese, Evelina Lamma, and Fabrizio Riguzzi. Non-ground abductive logic programming with probabilistic integrity constraints. *Theory and Practice of Logic Programming*, 21(5):557–574, 2021.
- [9] Yoshua Bengio. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2017.
- [10] Dimitris Bertsimas and Wes Gurnee. Learning sparse nonlinear dynamics via mixed-integer optimization. *Nonlinear Dynamics*, 111(7):6585–6604, 2023.
- [11] Daniil A Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023.
- [12] Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.
- [13] Bruno Buchberger. Bruno buchberger’s phd thesis 1965: An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. *Journal of Symbolic Computation*, 41(3):475–511, 2006. Logic, Mathematics and Computer Science: Interactions in honor of Bruno Buchberger (60th birthday).
- [14] Chin-Wen Chou, David B Hume, Till Rosenband, and David J Wineland. Optical clocks and relativity. *Science*, 329(5999):1630–1633, 2010.
- [15] Emma R. Cobian, Jonathan D. Hauenstein, and Charles W. Wampler. Robust numerical algebraic geometry. <https://arxiv.org/abs/2403.18749>, 2024.
- [16] Cristina Cornelio, Sanjeeb Dash, Vernon Austel, Tyler R Josephson, Joao Goncalves, Kenneth L Clarkson, Nimrod Megiddo, Bachir El Khadir, and Lior Horesh. Combining data and theory for derivable scientific discovery with AI-Descartes. *Nature Communications*, 14(1):1777, 2023.
- [17] Cristina Cornelio, Takuya Ito, Ryan Cory-Wright, Sanjeeb Dash, and Lior Horesh. The need for verification in AI-driven scientific discovery. *arXiv preprint arXiv:2509.01398*, 2025.
- [18] Ryan Cory-Wright, Cristina Cornelio, Sanjeeb Dash, Bachir El Khadir, and Lior Horesh. Evolving Scientific Discovery by Unifying Data and Background Knowledge with AI Hilbert. *Nature Communications*, 15(1):5922, 2024.
- [19] David Cox, John Little, and Donald O’Shea. *Ideals, Varieties and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Springer, 1991.
- [20] Wang-Zhou Dai, Qiuling Xu, Yang Yu, and Zhi-Hua Zhou. Bridging machine learning and logical reasoning by abductive learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [21] Alex Davies, Petar Veličković, Lars Buesing, Sam Blackwell, Daniel Zheng, Nenad Tomašev, Richard Tanburn, Peter Battaglia, Charles Blundell, András Juhász, et al. Advancing mathematics by guiding human intuition with AI. *Nature*, 600(7887):70–74, 2021.
- [22] Wolfram Decker, Gert-Martin Greuel, Gerhard Pfister, and Hans Schönemann. SINGULAR 4-4-0 — A computer algebra system for polynomial computations. <http://www.singular.uni-kl.de>, 2024.
- [23] David Déharbe, Pascal Fontaine, Stephan Merz, and Bruno Woltzenlogel Paleo. Exploiting symmetry in SMT problems. In *Automated Deduction – CADE-22*, volume 5663 of *Lecture Notes in Computer Science*, pages 222–236. Springer, 2009.

- [24] Igor Douven. Abduction. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2025 edition, 2025.
- [25] Mnacho Echenim, Nicolas Peltier, and Yanis Sellami. A generic framework for implicate generation modulo theories. In *International Joint Conference on Automated Reasoning*, pages 279–294. Springer, 2018.
- [26] Mnacho Echenim, Nicolas Peltier, and Sophie Turrett. An approach to abductive reasoning in equational logic. In *IJCAI 2013-International Joint Conference on Artificial Intelligence*, pages 531–537. AAAI Press, 2013.
- [27] Mnacho Echenim, Nicolas Peltier, and Sophie Turrett. Prime implicate generation in equational logic. *Journal of Artificial Intelligence Research*, 60:827–880, 2017.
- [28] D. Eisenbud. *Commutative Algebra: With a View Toward Algebraic Geometry*. Graduate Texts in Mathematics. Springer, 1995.
- [29] David Eisenbud, Craig Huneke, and Wolmer Vasconcelos. Direct methods for primary decomposition. *Inventiones mathematicae*, 110(1):207–235, Dec 1992.
- [30] Thomas Eiter and Georg Gottlob. The complexity of logic-based abduction. *Journal of the ACM (JACM)*, 42(1):3–42, 1995.
- [31] Marissa R Engle and Nikolaos V Sahinidis. Deterministic symbolic regression with derivative information: General methodology and application to equations of state. *AIChE Journal*, 68(6):e17457, 2022.
- [32] Brian D Fields. The primordial lithium problem. *Annual Review of Nuclear and Particle Science*, 61(1):47–68, 2011.
- [33] Han Gao, Luning Sun, and Jian-Xun Wang. Physics-informed deep learning for equation discovery: Methods, challenges, and opportunities. *arXiv preprint arXiv:2304.04862*, 2023.
- [34] Patrizia Gianni, Barry Trager, and Gail Zacharias. Gröbner bases and primary decomposition of polynomial ideals. *Journal of Symbolic Computation*, 6(2):149–167, 1988.
- [35] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at <http://www2.macaulay2.com>.
- [36] Geoffrey Grimmett and David Stirzaker. *Probability and random processes*. Oxford university press, 2020.
- [37] Roger Guimerà, Ignasi Reichardt, Antoni Aguilar-Mogas, Francesco A Massucci, Manuel Miranda, Jordi Pallarès, and Marta Sales-Pardo. A Bayesian machine scientist to aid in the solution of challenging scientific problems. *Science Advances*, 6(5):eaav6971, 2020.
- [38] Oki Gunawan, Seong Ryul Pae, Douglas M. Bishop, Yudistira Virgus, Jun Hong Noh, Nam Joong Jeon, Yun Seog Lee, Xiaoyan Shao, Teodor Todorov, David B. Mitzi, and Byungha Shin. Carrier-resolved photo-Hall effect. *Nature*, 575(7781):151–155, November 2019. Epub 2019 Oct 7.
- [39] Boumediene Hamzi and Houman Owhadi. Learning dynamical systems from data: A simple cross-validation perspective, part i: Parametric kernel flows. *Physica D: Nonlinear Phenomena*, 421:132817, 2021.
- [40] Robin Hartshorne. *Algebraic Geometry*, volume 52 of *Graduate Texts in Mathematics*. Springer, 1977.
- [41] Jonathan D. Hauenstein, Jose Israel Rodriguez, and Frank Sottile. General witness sets for numerical algebraic geometry. *Advances in Geometry*, 20(3):381–406, 2020.
- [42] Daniel Heldt, Martin Kreuzer, Sebastian Pokutta, and Hennie Poulisse. Approximate computation of zero-dimensional polynomial ideals. *Journal of Symbolic Computation*, 44(11):1566–1591, 2009.

- [43] Yu-Xuan Huang, Wang-Zhou Dai, Le-Wen Cai, Stephen H Muggleton, and Yuan Jiang. Fast abductive learning by similarity-based consistency optimization. *Advances in Neural Information Processing Systems*, 34:26574–26584, 2021.
- [44] Thomas Hubert, Rishi Mehta, Laurent Sartran, Miklós Z. Horváth, Goran Žužić, Eric Wieser, Aja Huang, Julian Schrittwieser, Yannick Schroecker, Hussain Masoom, Ottavia Bertolli, Tom Zahavy, Amol Mandhane, Jessica Yung, Iuliya Beloshapka, Borja Ibarz, Vivek Veeriah, Lei Yu, Oliver Nash, Paul Lezeau, Salvatore Mercuri, Calle Sönne, Bhavik Mehta, Alex Davies, Daniel Zheng, Fabian Pedregosa, et al. Olympiad-level formal mathematical reasoning with reinforcement learning. *Nature*, 2025. Accepted for publication, in press.
- [45] Alexey Ignatiev, Nina Narodytska, and Joao Marques-Silva. Abduction-based explanations for machine learning models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1511–1519, 2019.
- [46] Antonis C. Kakas, Robert Kowalski, and Francesca Toni. Abductive logic programming. *Journal of Logic and Computation*, 2(6):719–770, 1992.
- [47] Hiroaki Kitano. Nobel Turing Challenge: creating the engine for scientific discovery. *npj Systems Biology and Applications*, 7(1):29, 2021.
- [48] Mario Krenn, Robert Pollice, Si Yue Guo, Matteo Aldeghi, Alba Cervera-Lierta, Pascal Friederich, Gabriel dos Passos Gomes, Florian Häse, Adrian Jinich, AkshatKumar Nigam, et al. On scientific understanding with artificial intelligence. *Nature Reviews Physics*, 4(12):761–769, 2022.
- [49] Jiří Kubalík, Erik Derner, and Robert Babuška. Symbolic regression driven by training data and prior knowledge. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, pages 958–966, 2020.
- [50] Jiří Kubalík, Erik Derner, and Robert Babuška. Multi-objective symbolic regression for physics-aware dynamic modeling. *Expert Systems with Applications*, 182:115210, 2021.
- [51] Thomas S Kuhn. *The structure of scientific revolutions*, volume 962. University of Chicago press Chicago, 1997.
- [52] P. Langley. Integrated systems for computational scientific discovery. In *Proceedings of the AAAI Conference on Artificial Intelligence 38(20)*, pages 22598–22606, 2024.
- [53] Anton Leykin, Jose Israel Rodriguez, and Frank Sottile. Numerical algebraic geometry for Macaulay2. <https://macaulay2.com/doc/Macaulay2/share/doc/Macaulay2/NumericalAlgebraicGeometry/html/>, 2024. Macaulay2 Package.
- [54] Weixin Liang, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, Sheng Liu, Siyu He, Zhi Huang, Diyi Yang, Christopher Potts, Christopher D Manning, and James Y. Zou. Mapping the increasing use of LLMs in scientific papers. In *First Conference on Language Modeling*, 2024.
- [55] Zachary C Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018.
- [56] Roi Livni, David Lehavi, Sagi Schein, Hila Nachliely, Shai Shalev-Shwartz, and Amir Globerson. Vanishing component analysis. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 597–605, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- [57] Malcolm Longair. ... a paper... i hold to be great guns’: a commentary on maxwell (1865)‘a dynamical theory of the electromagnetic field. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2039):20140473, 2015.
- [58] Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The AI scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.

- 747 [59] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions.
748 *Advances in neural information processing systems*, 30, 2017.
- 749 [60] Chris A Mack. Fifty years of moore’s law. *IEEE Transactions on semiconductor manufac-*
750 *turing*, 24(2):202–207, 2011.
- 751 [61] Nour Makke and Sanjay Chawla. Interpretable scientific discovery with symbolic regression:
752 a review. *Artificial Intelligence Review*, 57(1):2, 2024.
- 753 [62] Daniel J Mankowitz, Andrea Michi, Anton Zhernov, Marco Gelmi, Marco Selvi, Cosmin
754 Paduraru, Edouard Leurent, Shariq Iqbal, Jean-Baptiste Lesciau, Alex Ahern, et al. Faster
755 sorting algorithms discovered using deep reinforcement learning. *Nature*, 618(7964):257–
756 263, 2023.
- 757 [63] P. Marquis. Extending abduction from propositional to first-order logic. In In: *Jorrand, P.,*
758 *Kelemen, J. (eds) Fundamentals of Artificial Intelligence Research*, 1991.
- 759 [64] Santiago Miret and Nandan M Krishnan. Are LLMs ready for real-world materials discovery?
760 *arXiv preprint arXiv:2402.05200*, 2024.
- 761 [65] Stefan Mitsch and André Platzer. KeYmaera X: An axiomatic tactical theorem prover for
762 hybrid systems. *Journal of Automated Reasoning*, 59(2):175–205, 2017.
- 763 [66] H Michael Möller and Ferdinando Mora. Upper and lower bounds for the degree of gröbner
764 bases. In *International Symposium on Symbolic and Algebraic Manipulation*, pages 172–183.
765 Springer, 1984.
- 766 [67] Emmy Noether. Invariant variation problems. *Transport theory and statistical physics*,
767 1(3):186–207, 1971.
- 768 [68] Peter Norvig. Inference in text understanding. In *AAAI*, pages 561–565, 1987.
- 769 [69] Gideon Parchomovsky. Publish or perish. *Michigan Law Review*, 98(4):926–952, 2000.
- 770 [70] Gabriele Paul. Approaches to abductive reasoning: an overview. *Artificial intelligence review*,
771 7(2):109–152, 1993.
- 772 [71] Judea Pearl. Reasoning with cause and effect. *AI Magazine*, 23(1):95–95, 2002.
- 773 [72] Charles Sanders Peirce. *Collected papers of charles sanders peirce*, volume 5. Harvard
774 University Press, 1934.
- 775 [73] Brenden K Petersen, Mikel Landajuela Larma, Terrell N. Mundhenk, Claudio Prata Santiago,
776 Soo Kyung Kim, and Joanne Taery Kim. Deep symbolic regression: Recovering mathemat-
777 ical expressions from data via risk-seeking policy gradients. In *International Conference on*
778 *Learning Representations*, 2021.
- 779 [74] Joshua C Peterson, David D Bourgin, Mayank Agrawal, Daniel Reichman, and Thomas L
780 Griffiths. Using large-scale experiments and machine learning to discover theories of human
781 decision-making. *Science*, 372(6547):1209–1214, 2021.
- 782 [75] Stanislas Polu and Ilya Sutskever. Generative language modeling for automated theorem
783 proving. In *arXiv preprint arXiv:2009.03393*, 2020.
- 784 [76] Oliver Ray. Automated abduction in scientific discovery. In Lorenzo Magnani and Ping Li,
785 editors, *Model-Based Reasoning in Science, Technology, and Medicine*, volume 64 of *Studies*
786 *in Computational Intelligence*, pages 103–116. Springer, 2007.
- 787 [77] Chandan K Reddy and Parshin Shojaee. Towards scientific discovery with generative ai:
788 Progress, opportunities, and challenges. In *Proceedings of the AAAI Conference on Artificial*
789 *Intelligence*, volume 39, pages 28601–28609, 2025.
- 790 [78] Andrew Reynolds, Haniel Barbosa, Andres Nötzli, Clark Barrett, and Cesare Tinelli. Scalable
791 algorithms for abduction via enumerative syntax-guided synthesis. In *Automated Reasoning.*
792 *IJCAR 2020*, volume 12166 of *Lecture Notes in Computer Science*, pages 141–160. Springer,
793 2020.

- 794 [79] M. T. Ribeiro, S. Singh, and C. Guestrin. "why should i trust you?": Explaining the predic-
795 tions of any classifier. *KDD*, pages 1135—1144, 2016.
- 796 [80] Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog,
797 M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang,
798 Omar Fawzi, et al. Mathematical discoveries from program search with large language mod-
799 els. *Nature*, 625(7995):468–475, 2024.
- 800 [81] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions
801 and use interpretable models instead. *Nature machine intelligence*, 1(5):206–215, 2019.
- 802 [82] Nicholas Shackel. Bertrand’s paradox and the principle of indifference. *Philosophy of Sci-*
803 *ence*, 74(2):150–175, 2007.
- 804 [83] Parshin Shojaei, Kazem Meidani, Shashank Gupta, Amir Barati Farimani, and Chandan K.
805 Reddy. LLM-SR: Scientific equation discovery via programming with large language models.
806 *arXiv preprint arXiv:2404.18400*, 2024. ICLR 2025 (Oral).
- 807 [84] Patrick E Shrout and Joseph L Rodgers. Psychology, science, and knowledge construction:
808 Broadening perspectives from the replication crisis. *Annual review of psychology*, 69(1):487–
809 510, 2018.
- 810 [85] Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin
811 Gal. Ai models collapse when trained on recursively generated data. *Nature*, 631(8022):755–
812 759, 2024.
- 813 [86] K.E. Smith, L. Kahanpää, P. Kekäläinen, and W. Traves. *An Invitation to Algebraic Geometry*.
814 Universitext. Springer New York, 2013.
- 815 [87] Andrew J. Sommese, Jan Verschelde, and Charles W. Wampler. Using monodromy to de-
816 compose solution sets of polynomial systems into irreducible components. In *Applications*
817 *of Algebraic Geometry to Coding Theory, Physics and Computation*, volume 36 of *NATO*
818 *Science Series II: Mathematics, Physics and Chemistry*, pages 297–315. Kluwer Academic
819 Publishers, 2001.
- 820 [88] Andrew J. Sommese, Jan Verschelde, and Charles W. Wampler. Symmetric functions applied
821 to decomposing solution sets of polynomial systems. *SIAM Journal on Numerical Analysis*,
822 40(6):2026–2046, 2002.
- 823 [89] Simon Spies, Lennard Gäher, Michael Sammler, and Derek Dreyer. Quiver: Guided abductive
824 inference of separation logic specifications in coq. *Proceedings of the ACM on Programming*
825 *Languages*, 8(PLDI):889–913, 2024.
- 826 [90] D.R. Swanson and N.R. Smalheiser. An interactive system for finding complementary litera-
827 tures: A stimulus to scientific discovery. *Artificial Intelligence*, pages 183–203, 1997.
- 828 [91] Silviu-Marian Udrescu and Max Tegmark. Ai feynman: A physics-inspired method for sym-
829 bolic regression. *Science Advances*, 6(16):eaay2631, 2020.
- 830 [92] Mojtaba Valipour, Bowen You, Maysum Panju, and Ali Ghodsi. SymbolicGPT: A generative
831 transformer model for symbolic regression. In *arXiv preprint arXiv:2106.14131*, 2021.
- 832 [93] Susana Vázquez Torres, Melisa Benard Valle, Stephen P Mackessy, Stefanie K Menzies,
833 Nicholas R Casewell, Shirin Ahmadi, Nick J Burlet, Edin Muratspahić, Isaac Sappington,
834 Max D Overath, et al. De novo designed proteins neutralize lethal snake venom toxins.
835 *Nature*, 639(8053):225–231, 2025.
- 836 [94] Zishen Wan, Che-Kai Liu, Hanchen Yang, Chaojian Li, Haoran You, Yonggan Fu, Cheng
837 Wan, Tushar Krishna, Yingyan Lin, and Arijit Raychowdhury. Towards cognitive ai systems:
838 a survey and prospective on neuro-symbolic ai. *arXiv preprint arXiv:2401.01040*, 2024.

- 839 [95] Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal
840 Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, Anima Anandkumar, Kari-
841 anne Bergen, Carla P Gomes, Shirley Ho, Pushmeet Kohli, Joan Lasenby, Jure Leskovec,
842 Tie-Yan Liu, Arjun Manrai, Debora Marks, Bharath Ramsundar, Le Song, Jimeng Sun,
843 Jian Tang, Petar Veličković, Max Welling, Linfeng Zhang, Connor W Coley, Yoshua Ben-
844 gio, and Marinka Zitnik. Scientific discovery in the age of artificial intelligence. *Nature*,
845 620(7972):47–60, August 2023.
- 846 [96] Runxiang Wang, Boxiao Wang, Kai Li, Yifan Zhang, and Jian Cheng. DrSR: LLM based
847 scientific equation discovery with dual reasoning from data and experience. *arXiv preprint*
848 *arXiv:2506.04282*, 2025.
- 849 [97] Zuojun Xu, Feng Ren, Ping Wang, Jie Cao, Chunting Tan, Dedong Ma, Li Zhao, Jinghong
850 Dai, Yipeng Ding, Haohui Fang, et al. A generative AI-discovered TNIK inhibitor for idio-
851 pathic pulmonary fibrosis: a randomized phase 2a trial. *Nature Medicine*, pages 1–9, 2025.
- 852 [98] Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad
853 Godil, Ryan J Prenger, and Animashree Anandkumar. Leandojo: Theorem proving with
854 retrieval-augmented language models. In A. Oh, T. Naumann, A. Globerson, K. Saenko,
855 M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, vol-
856 ume 36, pages 21573–21612. Curran Associates, Inc., 2023.
- 857 [99] Xiao-Wen Yang, Wen-Da Wei, Jie-Jing Shao, Yu-Feng Li, and Zhi-Hua Zhou. Analysis
858 for abductive learning and neural-symbolic reasoning shortcuts. In *Forty-first International*
859 *Conference on Machine Learning*, 2024.
- 860 [100] Zhi-Hua Zhou. Abductive learning: towards bridging machine learning and logical reasoning.
861 *Science China. Information Sciences*, 62(7):76101, 2019.

A Extended Literature Review

Abductive reasoning is a field of inference that generates plausible explanations from incomplete observations [72, 70]. Unlike deductive reasoning, which obtains a conclusion from a set of premises, abductive reasoning starts with an outcome and ends with a set of plausible explanations, e.g.,

- The fact C is observed.
- If A were true then C would be true.
- There is reason to believe that A is true.

As observed by Peirce [72], abductive reasoning is arguably “the only logical operation which introduces any new ideas”. Moreover, it is often applied by humans in everyday situations, such as reading between the lines [68] and counterfactual reasoning [71]. Doctors use abductive reasoning to propose possible diseases that cause observed symptoms. It has been argued that automating inference may be necessary for understanding and automating consciousness [9]. However, abductive reasoning has not yet been incorporated into state-of-the-art AI systems, such as AI-Feynman, AI-Descartes, or AI-Hilbert [91, 16, 18]. Thus, automating abductive inference is an important component of automating the scientific method.

Automated Abductive Inference in Logic and Classical AI. Abduction, firstly introduced by Peirce as the inferential step that proposes an explanatory hypothesis for a surprising observation in the logic of scientific inquiry [24], was then later on, in AI, sharpened into a logic-based hypothesis synthesis problem [63]. Given an incomplete background theory T and observations (target claim) O , find hypotheses H (additional assumptions from a designated space of abducibles, e.g., defined by a grammar) that render the claim derivable, meaning that $T \cup H \models O$ and $T \cup H$ is consistent (typically subject to minimality criteria) [30]. This was first introduced for propositional logic and then extended to more expressive forms such as first-order [63]. Abductive logic programming (ALP) then further operationalized these ideas by combining abducibles with integrity constraints and proof procedures [46], and more recent work extended it to even richer constraint families (e.g., including probabilistic integrity constraints [8]).

Most relevant to our work is the case in which hypotheses are *equations* (or more generally, algebraic constraints): in this setting, abduction can be realized via consequence generation in equational logic, where explanations are returned as equalities/disequalities that complete the theory enough to derive the target [27]. More broadly, implicate-generation modulo background theories can be used to generate missing conditions (including equalities and arithmetic constraints) whenever a suitable decision procedure is available [25, 26]. However, these methods either support only limited arithmetic expressiveness (usually including only basic arithmetic operators) or are computationally challenging as the hypothesis space grows.

More recently, abduction has become part of automated reasoners/solvers (see paragraph below for more details): SMT/SyGuS-based methods synthesize hypotheses directly in background theories with equality or linear arithmetic [78]. Automated solvers, like *cvc5* expose abduction as a reusable primitive [3]. These can be also incorporated in verification-oriented systems where abduction fills in preconditions and specifications so that program proofs can find a proof compositionally [89, 4].

Finally, beyond logic, several papers have investigated abductive inference in the context of scientific discovery: Swanson and Smalheiser’s interactive system ARROWSMITH [90] for identifying complementary literatures in other branches of science exploits abductive hypothesis formation as constructing bridging concepts across disjoint research areas; and more recently Langley [52] argued for integrated discovery systems that treat abduction as one component within an end-to-end pipeline that combines multiple reasoning and learning mechanisms.

We follow this line of work in interpreting abduction as the process of selecting or synthesizing hypotheses that complete an incomplete theory in order to explain a target hypothesis.

Automated Abductive Inference for Machine Learning. The recent surge in popularity of black-box models, including large language models, has led to an interest in explaining the predictions of these models. LIME [79] and similar methods such as SHAP [59] provide (local) interpretable explanations that describe the local behavior of any ML model using a weighted linear combination (often sparse) of input features; see [81, 55] for reviews of interpretable and explainable machine

learning. However, existing works on interpretability and explainability generally focus on computing explanations for black-box models, rather than generating theories that could be used for scientific discovery. In particular, techniques like LIME or SHAP scores cannot, to our knowledge, be leveraged to establish missing axioms in a scientific discovery context.

More recently, a new line of work has explored abductive inference for machine learning [20, 45, 100, 43, 99]; see also [94] for a recent review of this line of work. Initiated by [20, 45], the key idea is to embed purely logical reasoning models within the machine learning process to take advantage of the benefits of both machine learning and logical reasoning. For instance, in [45], the authors use abductive inference to generate explanations of ML models given input background theories and a set of hypothetical explanations. Moreover, [20] combines machine learning techniques with a first-order logic model to obtain an output that depends on both models. However, these works do not apply abductive inference to scientific discovery.

SMT Solvers and Abductive Reasoning in Automated Theorem Proving. Recent work has explored the use of Satisfiability Modulo Theories (SMT) solvers for abductive reasoning, particularly in the context of automated theorem proving and program verification. SMT solvers generalize Boolean SAT by incorporating theories over specific domains (e.g., arithmetic, arrays, bit-vectors), enabling reasoning about richer constraints [5, 23].

Barbosa et al. [4] developed an SMT-based tactic in the Coq proof assistant that uses abductive reasoning to identify missing axioms needed to complete a proof. When the cvc5 SMT solver fails to prove a goal, their `abduce` tactic generates facts that, when added to the background theory, allow the proof to succeed. This work shares conceptual similarities with AI-Noether, as both identify missing axioms to bridge a derivability gap, though their approach operates within formal proof systems rather than scientific discovery contexts. Reynolds et al. [78] introduced scalable algorithms for abduction via syntax-guided synthesis (SyGuS), demonstrating that SMT solvers can efficiently generate formulas consistent with axioms that entail a given goal. Their work establishes theoretical foundations for abductive synthesis that inform our comparison with cvc5 in Section F.

Beyond theorem proving, abductive reasoning has been applied to scientific hypothesis generation through Abductive Logic Programming (ALP). For example, Ray [76] demonstrated how ALP methods outperform theorem-proving approaches. While ALP focuses on logic-based hypothesis synthesis, AI-Noether operates in the polynomial algebraic setting, leveraging computational algebraic geometry rather than first-order logic, and works in the robust cases of noisy machine-generated hypotheses.

Large Language Models for Scientific Discovery and Equation Discovery. The emergence of large language models (LLMs) has opened new avenues for automating scientific discovery. Recent work has explored using LLMs’ embedded scientific knowledge and code generation capabilities to discover governing equations from data, a task traditionally addressed by symbolic regression.

Shojaee et al. [83] introduced LLM-SR, which combines LLMs with evolutionary search to discover scientific equations by treating them as program skeletons with placeholder parameters. Their approach leverages domain-specific priors from LLM training to navigate the equation search space more efficiently than classical symbolic regression methods. Follow-up work by DrSR [96] enhances this framework with dual reasoning that combines data-driven insights with reflective learning, analyzing structural relationships like monotonicity and correlation to guide equation generation. These LLM-based symbolic regression approaches are complementary to AI-Noether: while they discover equations *from data*, AI-Noether discovers equations needed to *explain* already-formulated hypotheses without requiring data for the missing axioms.

Beyond symbolic discovery, LLMs have enabled autonomous experimental systems. Boiko et al. [11] demonstrated fully autonomous chemical research where LLMs plan multi-step syntheses, control laboratory robotics, and analyze results—closing the loop between computational hypothesis generation and physical experimentation. This represents a step toward Kitano’s “Nobel Turing Challenge” [47] of building AI systems capable of Nobel-quality autonomous discovery.

More broadly, LLMs have been applied to automated theorem proving and mathematical discovery. Romera-Paredes et al. [80] introduced FunSearch, an evolutionary procedure pairing pretrained LLMs with systematic evaluation to solve open problems in extremal combinatorics, demonstrating that LLMs can push beyond existing human knowledge when combined with search algorithms. Lu et al. [58] proposed The AI Scientist, a system for fully automated open-ended scientific discovery

that generates research ideas, runs experiments, and writes papers. In formal mathematics, recent work on automated theorem proving using LLMs [75, 98] has shown promise in completing proof steps, though performance still lags human experts on research-level problems. These systems focus on *generating* new knowledge or proofs, whereas AI-Noether performs *abductive inference* to identify which axioms are missing from an incomplete theory. The axioms we are inferring are, by assumption, not derivable from the other known axioms and consequences, so traditional forward inference approaches would not work.

Neuro-Symbolic Methods and Physics-Informed Discovery. Recent work has explored hybrid approaches that combine neural networks with symbolic reasoning for scientific discovery. Petersen et al. [73] introduced deep symbolic regression using reinforcement learning with risk-seeking policies to discover mathematical expressions, demonstrating improved exploration of the symbolic search space. Valipour et al. [92] developed SymbolicGPT, a transformer-based model for symbolic regression that learns from large corpora of synthetic equations. These neural-symbolic approaches excel at pattern recognition and generalization across equation families, but typically require training data for the target expressions.

Physics-informed approaches incorporate domain knowledge as inductive biases. Gao et al. [33] showed that embedding physical constraints (e.g., conservation laws, symmetries) into neural network architectures improves equation discovery accuracy and generalization. Similarly, SINDy [12] leverages sparsity constraints to identify governing equations of dynamical systems and Hamzi and Owahdi [39] similarly demonstrate the use of kernel flows in learning dynamical systems. While these methods successfully integrate prior knowledge, they focus on *parameter fitting* within known functional forms rather than *structural discovery* of missing axioms. AI-Noether addresses the complementary problem where the functional form itself is unknown and must be inferred from algebraic constraints.

AI-Noether vs SOTA scientific discovery models. Figure 11 compares AI-Noether with four other scientific discovery models in terms of their functionality in discovering theory from data and existing theory. The traditional discovery process starts with known theory and data and seeks to discover a new formula that explains the data in a way that is consistent with the theory. The traditional, first-principles-based scientific discovery, broadly speaking, operates on two paradigms: either generate a hypothesis by studying patterns in known data and verify the hypothesis by checking derivability or consistency with theory or derive possible hypotheses from first-principles and test the derived hypotheses on data. This approach is manual and time-consuming. ML/Regression approaches help automate the hypothesis generation step by allowing for automated hypothesis generation using data. Verification here, however, is still done manually, and is often the bottleneck. AI-Descartes and AI-Hilbert solve this problem by automating the verification step by either automating the testing on theory using theorem provers (AI-Descartes) or combining fit-to-theory in the search process (AI-Hilbert). AI-Noether targets a new problem of automation in the scientific process: in the event that theory is insufficient for deriving a machine-generated hypothesis, instead of changing the hypothesis-generation mechanism, AI-Noether proposes corrections / completions to the theory in order to discover a new, minimal theory required to validate the hypothesis.

B A Walk-through Example

Example 1. Consider an axiom system over five real quantities, indicated by x, y, z, a, b :

$$x^2 - ay^2 = 0,$$

$$by - z = 0.$$

Further, suppose that a, b are physical constants in the system, while x, y, z are quantities that may vary. Let p_1 be the polynomial $x^2 - ay^2$ and let $p_2 = by - z$, and the axioms be $p_1 = 0$ and $p_2 = 0$. Further, let

$$q = b^2x^2 - az^2.$$

The equation $q = 0$ then follows from the axioms as

$$q = b^2x^2 - az^2 = b^2(x^2 - ay^2) + a(by + z)(by - z) = b^2p_1 + a(by + z)p_2. \quad (1)$$

In the example above, we have expressed q as $\alpha p_1 + \beta p_2$ where $\alpha = b^2$ and $\beta = a(by + z)$ are polynomials defined over x, y, z, a, b . We say that q is an *algebraic combination* of the polynomials




Figure 11: **Comparison of AI-Noether with other scientific discovery models.** We compare classical scientific discovery, machine-learning-based regression, and recent integrated systems (AI-Descartes and AI-Hilbert) with AI-Noether. While most approaches focus on generating new hypotheses or formulas from data and/or existing theory, AI-Noether targets hypotheses that are not supported by the current theory and performs abductive inference to identify and justify the missing axioms required to reconcile them with known knowledge.

1018 p_1 and p_2 and say that $q = 0$ is *derivable* from $p_1 = 0$ and $p_2 = 0$. As $q = \alpha p_1 + \beta p_2$, it follows that
 1019 $q = 0$ for all values of x, y, z, a, b such that $p_1 = 0$ and $p_2 = 0$. We therefore say that the axioms in
 1020 the example imply or explain $q = 0$.

Table 5: Six minimal axiom systems; each column corresponds to one such system, consisting of two equations that explain $q = 0$.

$b = 0$	$b = 0$	$b = 0$	$x^2 - ay^2 = 0$	$x^2 - ay^2 = 0$	$x^2 - ay^2 = 0$
$a = 0$	$by + z = 0$	$by - z = 0$	$a = 0$	$by + z = 0$	$by - z = 0$

1021 Notice that there are many alternative theories from which the equation $q = 0$ follows as a conse-
 1022 quence. For example, looking at the equation (1) and letting $p_3 = by + z$, it is clear that $q = 0$
 1023 follows from the axioms $p_1 = 0$ and $p_3 = 0$ as $q = \alpha p_1 + \beta p_3$. In a similar manner, we can
 1024 trivially create the six *minimal* theories in Table 5 where the first axiom is formed by taking a poly-
 1025 nomial that is a factor of αp_1 and setting it to 0, and the second axiom is similarly composed from a
 1026 factor of βp_2 ; $q = 0$ is a consequence of each theory.

1027 Each of the 6 axiom systems above is minimal in the sense that we cannot remove an axiom from
 1028 it and still have $q = 0$ as a consequence. Depending on the context, some axiom systems may be
 1029 more interesting than others. For example, if a represents the mass of a pendulum, $a = 0$ is likely
 1030 not an interesting condition. On the other hand, if a represents temperature in Celsius, $a = 0$ may
 1031 represent a meaningful condition.

1032 The axiom system $p_1 = 0, p_2 p_3 = 0$ is also a valid theory explaining $q = 0$, but we call it *reducible*
 1033 in the sense that we can trivially get a simpler axiom system if we replace $p_2 p_3$ by either of its factors.
 1034 Thus, if our only available axioms were $p_1 = 0$, the “residual” of the set of axioms $\{p_2 p_3 = 0\}$
 1035 is not as small as possible. There are other ways of building an infinitude of axiom systems such

that $q = 0$ is a consequence. Starting from the axiom system $p_1 = 0, p_2 = 0$, it is easy to see that $p_1 = 0, rp_1 + sp_2 = 0$, where r, s are nonzero real numbers, are valid theories that explain $q = 0$. Again, starting with $p_1 = 0$, the axiom set $\{rp_1 + sp_2 = 0\}$ does not have the least residual.

Now, assume we have experimental data for a phenomenon consisting of a number of data points specifying values of x, z, a, b (with a and b unchanging in the data), that the data satisfy the equation $b^2x^2 - az^2 = 0$, and that we are able to numerically obtain this equation from the data. Further, suppose that we have a partial theoretical explanation of the phenomenon, where the partial theory assumes the presence of another variable y and the equation $x^2 - ay^2 = 0$, i.e., $p_1 = 0$. In this situation, we consider y as an unmeasured quantity, and x, z, a , and b as measured quantities. Given that the data points at hand do not have any values for y , it is not possible to infer the relationship $by - z$ directly from the data.

Consider the equation $0 = b^2p_1 = b^2x^2 - b^2ay^2$ that follows from $p_1 = 0$. Subtracting the derived equation $b^2x^2 - az^2 = 0$ from it, we obtain the relationship $a(z^2 - b^2y^2) = 0$ as a consequence of $p_1 = 0$ and $q = 0$. This implies that either $a = 0$ or $by - z = 0$ or $by + z = 0$. Each of the equations (or products of pairs of equations) in the previous line can be appended to $p_1 = 0$ to get a valid theory (in the absence of other information) explaining $q = 0$.

Using the notation above, “derivable from” implies “can be explained by”. We will later discuss conditions under which the converse is true (see Section 3.5).

C Methodology details

In this section, we develop the details of *AI-Noether* and its **Encode-Decompose-Reason** modules. We first develop the algebraic theory of the method and then extend it to numerical algorithms that enable us to handle noise. For further reading on these topics, one can refer to the following textbooks and resources on computational algebraic geometry [19, 86, 1, 28, 34, 29, 40]. This section will include all the algebraic geometry framework and, as such, will be denser than the rest of the paper.

C.1 Preliminaries - Algebraic Background Without Noise

We define notions specific to real algebraic geometry; see [19] for further details. Let $\mathbb{R}$ denote the set of real numbers and let $R = \mathbb{R}[x_1, \dots, x_n]$ be the ring of polynomials defined over the n variables $x_1, \dots, x_n$ with real coefficients. Axioms and hypotheses in our model will be equations of the form $p(x) = 0$ where $p \in R$. We abuse notation and identify an axiom by its defining polynomial. We will focus on real-valued solutions of axiom systems; see the appendix for more details. Finally, we will focus the exposition on the case with one axiom correction and one consequence for ease of explanation. This will generalize well to the multiple missing axiom and consequence case and we will make a note of these generalizations for each module.

Encode. An *ideal* over R is a subset I of R that is closed under addition and also multiplication by elements of R : if $f, g \in I$, then $f + g \in I$ and $\alpha f \in I$ for any $\alpha \in R$. Given polynomials $A_1, \dots, A_k \in R$, we denote the set of all algebraic combinations of $A_1, \dots, A_k$ by

$$\langle A_1, \dots, A_k \rangle = \{f \in R : f = \sum_{i=1}^k \alpha_i A_i \text{ for some } \alpha_i \in R\}.$$

This set is called an ideal over R with *generators* $A_1, \dots, A_k$. Let $\mathcal{A}$ stand for $\{A_1, \dots, A_k\}$, and let $I(\mathcal{A})$ stand for the ideal $\langle A_1, \dots, A_k \rangle$. Our earlier notion of “derivability” corresponds to ideal membership: $Q = 0$ is derivable from $\mathcal{A}$ if Q is an algebraic combination of $A_1, \dots, A_k$, i.e., if $Q \in I(\mathcal{A})$. Given an ideal I , we define the *variety* of I , denoted $V(I)$, as the set

$$V(I) = \{\mathbf{x} \in \mathbb{R}^n : f(\mathbf{x}) = 0 \text{ for each } f \in I\}.$$

A standard fact from algebraic geometry is that the set above is the same as the solution set to any finite collection of polynomials that generate the same ideal. Encoding equations with ideals and varieties has a key advantage in that it allows us to study logical implication and consistency questions geometrically: if $Q = 0$ can be derived from $A_1 = 0, \dots, A_k = 0$, then $V(A_1, \dots, A_k, Q) = V(A_1, \dots, A_k)$. If adding Q to $\mathcal{A}$ introduces extra structure into the resulting variety, then Q is not derivable from $\mathcal{A}$. As such, this holds for any Q , and so in the case where we have multiple consequences Q_i , this holds for each Q_i . In that case, we define $V(A_1, \dots, A_k, Q_1, \dots, Q_s)$ as our variety.

Decompose. Assume Q is not implied by $\mathcal{A}$. Then $V(\mathcal{A}, Q) \neq V(\mathcal{A})$. We will detail this in the case where a single axiom correction is needed. Suppose Q can be derived *nontrivially* from an axiom A_{k+1} along with $\mathcal{A}$, i.e., $Q = \alpha_1 A_1 + \dots + \alpha_{k+1} A_{k+1}$ for some polynomials α_i without $\alpha_{k+1} = 1, A_{k+1} = Q$. In this case $V(A_1, \dots, A_k, Q) = V(A_1, \dots, A_k, \alpha_{k+1} A_{k+1})$ and since $\alpha_{k+1} \neq 1$, we know that intersecting $V(Q)$ with $V(\mathcal{A})$ is the same as intersecting $V(\alpha_{k+1} A_{k+1})$ with $V(\mathcal{A})$. Since $V(\alpha_{k+1} A_{k+1})$ is reducible, we are introducing new reducibility into $V(\mathcal{A})$ that is a factor of the residual $\alpha_{k+1} A_{k+1} = Q - (\sum_{i=1}^k \alpha_i A_i)$ (and if we had multiple consequences, then we would have $V(A_1, \dots, A_k, Q_1, \dots, Q_s) = V(A_1, \dots, A_k, \alpha_{k+1} A_{k+1}, \dots, \alpha_{l-k+1} A_{l-k+1})$). We therefore can take our intersected variety $V(\mathcal{A}, Q)$ and study its irreducible components. We next define the mechanisms required for this process.

A *prime ideal* is an ideal $P \subset R$ such that if $fg \in P$ with $f, g \in R$, then $f \in P$ or $g \in P$.

For a prime ideal P , $V(P)$ is irreducible, i.e., it cannot be represented as the union of two proper subsets that are varieties. The *radical* of an ideal I is denoted by $\sqrt{I}$ and is defined as:

$$\sqrt{I} = \{f \in R : f^m \in I \text{ for some positive integer } m\}.$$

The following key theorem allows us to break up a variety into irreducible components:

Theorem 2: Primary Decomposition

Any ideal I can be decomposed into an intersection of ideals

$$I = I_1 \cap \dots \cap I_r$$

such that

1. There are no redundancies among the I_i , i.e., for each j we have $\cap_{i \neq j} I_i \not\subset I_j$.
2. For each i , $\sqrt{I_i}$ is a prime ideal (called an *associated prime* of I).

Let $P_i = \sqrt{I_i}$. Using the fact that for any ideals I, J , we have $V(I \cap J) = V(I) \cup V(J)$ and $V(I) = V(\sqrt{I})$ [19], we get from the primary decomposition that a variety can be decomposed into a union of varieties corresponding to prime ideals: $V(I) = \bigcup_{i=1}^r V(P_i)$.

Decomposition in this setting serves a role analogous in spirit to Principal Component Analysis (PCA) in linear algebra, identifying latent, independent substructures that explain variation in the system. Whereas PCA decomposes a dataset into orthogonal directions of variance, primary decomposition identifies algebraically independent pieces of a solution space introduced by inconsistency or incompleteness.

Reason. To test derivability, we rely on an algebraic technique for eliminating variables from polynomial systems. A *reduced Gröbner basis* is a canonical generating set for an ideal that can be computed algorithmically. Given a variable ordering $x_1 < \dots < x_n$ (e.g., lexicographic), the reduced Gröbner basis $\{G_1, \dots, G_m\}$ of $I = \langle A_1, \dots, A_k \rangle$ satisfies a key elimination property:

Theorem 3: Elimination Theorem

For any $1 \leq d \leq n$, the polynomials in the Gröbner basis $\{G_1, \dots, G_m\}$ involving only $x_1, \dots, x_d$ generate the same ideal as all polynomials in I involving only those variables:

$$\langle \{G_1, \dots, G_m\} \cap \mathbb{R}[x_1, \dots, x_d] \rangle = \langle A_1, \dots, A_k \rangle \cap \mathbb{R}[x_1, \dots, x_d].$$

So if we want to eliminate the variables $x_{d+1}, \dots, x_n$ from the system, we compute a set of generators for the elimination ideal $I \cap \mathbb{R}[x_1, \dots, x_d]$ by finding the reduced Gröbner basis and extracting polynomials involving only $x_1, \dots, x_d$. Geometrically, this corresponds to projecting the variety $V(I) \subset \mathbb{R}^n$ onto the subspace of the first d coordinates. The elimination ideal defines the Zariski closure of this projection: $V(I \cap \mathbb{R}[x_1, \dots, x_d]) = \overline{\pi_d(V(I))}$, where the closure is taken over $\mathbb{C}$ and we then restrict to real points.

This provides a way to test algebraic consequences: if Q depends only on $x_1, \dots, x_d$, we eliminate all other variables and check whether Q belongs to the elimination ideal. We use this to test whether Q can be derived from the axioms together with a candidate axiom extracted from decomposition.

1115 Elimination also filters out trivial candidates whose projections represent undesired boundary con-
 1116 ditions under which Q holds vacuously.

1117 We present sufficient conditions for axiom recovery in Section 3.5. The generalized algorithm is
 1118 detailed in Algorithm 1 in Appendix E.3.

1119 C.2 Handling Noisy Consequences - Numerical Extensions

1120 The purpose of this study is to bridge the gap between machine-generated hypotheses and theory.
 1121 Therefore, we must address the case where consequences (and possibly axioms) contain noise due
 1122 to measurement error or numerical approximation in data-driven discovery. Accordingly, we make
 1123 two key changes to our algorithm to account for noise.

1124 **Numerical Irreducible Decomposition.** The algebraic decomposition method described in Sec-
 1125 tion C.1 requires exact symbolic computations and is therefore not stable under coefficient pertur-
 1126 bations of the consequences. We address this by replacing symbolic primary decomposition with
 1127 *numerical irreducible decomposition*, a well-established technique in numerical algebraic geometry
 1128 [6].

1129 The numerical approach proceeds in two stages. First, the variety is intersected with generic hyper-
 1130 planes until a zero-dimensional set of witness points is obtained, as illustrated in Figure 12. These
 1131 witness points sample the variety at discrete locations. Second, homotopy continuation-based curve
 1132 tracing determines which points lie on the same irreducible components (Figure 13). The method
 1133 leverages fundamental constructs from numerical algebraic geometry, including witness sets [88, 41]
 1134 and monodromy-based component identification [87].




Figure 12: Geometric visualization of slicing: intersecting generic hyperplanes with a variety until a zero-dimensional set is obtained. Left: the variety. Middle: the slicing process. Right: the zero-dimensional witness set.




Figure 13: Geometric visualization of curve tracing: defining homotopies between points to test whether sets of points lie on the same irreducible surface. Left: illustration of homotopy paths along the surface. Right: sets of points partitioned by irreducible component. Details in [6].

1135 The output of numerical irreducible decomposition is a collection of witness sets—finite point sam-
 1136 ples on each irreducible component of the variety. While this provides geometric information about
 1137 the decomposition, it does not directly yield the algebraic structure we seek (polynomial generators

1138 of each component). We therefore perform symbolic regression on the witness points to recover
 1139 candidate polynomial relations.

1140 **Symbolic Regression on Witness Sets.** Given witness points on an irreducible component, we fit
 1141 polynomial relations that approximately vanish on those points. This problem—learning polynomi-
 1142 als from data—has been studied extensively in the context of vanishing ideals [56, 42].

1143 For each removed axiom A_i with symbolic form

$$A_i(x_1, \dots, x_m) = \sum_k c_{ik} M_{ik}(x_1, \dots, x_m) = 0,$$

1144 we use its *monomial support* $\{M_{ik}\}$ as a fixed regression basis. Given sampled points $\{p_j\}_{j=1}^N$ on a
 1145 single irreducible component, we form the evaluation (Vandermonde-like) matrix $\Phi \in \mathbb{C}^{N \times K}$ with
 1146 entries $\Phi_{jk} = M_{ik}(p_j)$. The coefficients $\mathbf{c}_i = (c_{i1}, \dots, c_{iK})$ are estimated by solving

$$\min_{\mathbf{c}_i} \|\Phi \mathbf{c}_i\|_2^2 \quad \text{subject to} \quad \|\mathbf{c}_i\|_2 = 1,$$

1147 which is equivalent to computing the right singular vector corresponding to the smallest singular
 1148 value of Φ [56]. This formulation identifies the (approximate) kernel of the evaluation map, recov-
 1149 ering a polynomial that nearly vanishes on the sample set. Alternative approaches based on border
 1150 basis methods [42] or least-squares sensitivity analysis could be employed, though we adopt the
 1151 SVD-based approach for its numerical stability and direct connection to nullspace identification.

1152 The resulting polynomial $\hat{A}_i(x) = \sum_k \hat{c}_{ik} M_{ik}(x)$ is the best-fitting algebraic relation in the least-
 1153 squares sense that holds on the numerically computed points. Each component yields its own fitted
 1154 relation, enabling us to capture branch-specific structures when the noisy variety decomposes into
 1155 multiple sheets.

1156 **Choice of Variables and Basis Functions.** The variable set and monomial basis used in regression
 1157 serve as tunable hyperparameters. For each axiom A_i , we restrict regression to the variables actu-
 1158 ally appearing in A_i , denoted $\text{vars}(A_i)$, and use its observed monomial support as basis functions.
 1159 This prevents overfitting and ensures fitted coefficients are directly comparable to the symbolic ax-
 1160 iom. Alternatively, one could learn variable subsets or expand the monomial basis to include cross-
 1161 variable interactions or higher-degree terms, though we adopt the fixed-support approach for clarity
 1162 and stability.

1163 The fitted residual $\|\Phi \mathbf{c}_i\|_2 / \|\mathbf{c}_i\|_2$ serves as a numerical quality metric. Low residuals indicate that
 1164 the recovered relation nearly vanishes on the component, suggesting the axiom (or its noisy per-
 1165 turbation) governs that portion of the variety. The resulting polynomials form our set of candidate
 1166 axioms to test, as illustrated in Figure 14.




Figure 14: Geometric visualization of symbolic regression output: recovered polynomial relations governing the irreducible surfaces identified by curve tracing (Figure 13).

1167 **Robust Fitting.** In practice, witness sets may contain outliers due to numerical error accumulation
 1168 in homotopy continuation. To improve robustness, we implement an iterative outlier removal pro-
 1169 cedure: after initial coefficient estimation, we remove points with the largest residuals (top 10% by
 1170 default) and refit. This process repeats until convergence or a maximum iteration count, yielding co-
 1171 efficient estimates less sensitive to numerical artifacts. Implementation details, including coefficient

normalization and error metrics for comparing recovered and true axioms, are provided in the code repository.

Reason with Axiom and Consequences containing Numerical Constants.

An alternative to algebraic derivability is to employ a theorem prover to determine whether the conjecture is logically derivable from the expanded set of axioms (i.e., the original axioms together with the abduced ones). Our approach follows the strategy proposed in AI-Descartes [16], which introduced two complementary notions of derivability for scientific laws: a direct (symbolic) derivability test and an *existential* one. The latter is designed for situations in which expressions inferred from data contain numerical constants (e.g., $Q = x/(0.709x + 0.157)$), since such expressions are unlikely to be matched syntactically (often the background theory contains only symbolic constants/parameters). To do this, AI-Descartes abstracts each numeric constant into a new logical variable c_i , for example,

$$Q = \frac{x}{0.709x + 0.157} \rightsquigarrow Q' = \frac{x}{c_1x + c_2},$$

and constrains these variables through admissibility conditions $\mathcal{C}'$ (e.g., non-negativity). The constants are then existentially quantified so that the theorem prover may search over all admissible instantiations.

Under this abstraction, AI-Descartes defines a conjecture to be *existentially derivable* from the background theory when

$$\exists c_1 \cdots \exists c_s. (\mathcal{C} \wedge \mathcal{A}) \rightarrow (f' \wedge \mathcal{C}'), \quad (2)$$

that is, when some choice of the abstracted constants makes the conjecture consistent with both the axioms $\mathcal{A} = \{A_1, \dots, A_{k+r}\}$ and the constraints $\mathcal{C}$ over the symbolic variables present in the axioms ($x_1, \dots, x_n$, assumed universally quantified). This mechanism allows the system to assess the derivability of a data-derived expression with respect to the background theory for at least one admissible instantiation of its numerical constants, thus avoiding sensitivity to noise in their estimated values.

In our setting, some axioms are abduced from the conjecture and the original theory, and thus may also contain numerical constants. Consequently, we need to introduce existential variables not only for the numerical constants appearing in the conjecture but also for those occurring in the abduced axioms. This requires extending the existential derivability notion of Cornelio et al. [16] to handle numerical parameters on both sides of the derivation. Our resulting formulation, which accommodates constants in both the conjecture and the axioms, is as follows:

$$\exists c_1 \cdots \exists c_s \left(\mathcal{C}' \wedge \forall x_1 \cdots \forall x_n ((\mathcal{C} \wedge \mathcal{A} \wedge \mathcal{A}') \rightarrow f') \right). \quad (3)$$

where $\mathcal{A}'$ denotes the abduced axioms with their numerical constants replaced by existential variables, and $\mathcal{C}'$ denotes the admissibility constraints for these newly introduced variables.

In this way, we obtain a strict generalization of the existential derivability introduced in AI-Descartes [16], which now accommodates numerical parameters in both the conjecture and the abduced axioms.

D Additional problem details and derivations

Derivation of Carrier-Resolved Photo-Hall Effect

The Carrier-Resolved Photo-Hall effect [38] equation describes the relationship between various parameters of a semiconducting surface and is given by:

$$Q := H - \frac{r e \mu_P^2 [p_0 + \Delta n(1 - \beta^2)]}{\sigma^2} = 0,$$

1210 where $\sigma = e\mu_P[p_0 + \Delta n(1 + \beta)]$ (carrier-resolved Photo-Hall effect [38]). We assume the axioms

$A_1 : \beta\mu_P - \mu_N$	Mobility ratio definition
$A_2 : \mu_H - r\mu$	Hall factor correction
$A_3 : p_h - p_0 - \Delta p$	Hole concentration balance
$A_4 : n - \Delta n$	Electron concentration tracking
$A_5 : \Delta p - \Delta n$	Charge neutrality condition
$A_6 : \sigma - ep_h\mu_P - en\mu_N$	Total conductivity sum
$A_7 : H(p_h + \beta n)^2 e - rp_h + r\beta^2 n$	Hall coefficient relation.

1211 Here μ_N and μ_P are electron and hole drift mobilities, $\beta = \mu_N/\mu_P$ is the mobility ratio (encoded
1212 by A_1), μ_H is the Hall mobility with Hall scattering factor $r = \mu_H/\mu$ (encoded by A_2), Δn and
1213 Δp are the electron and hole photo-carrier densities (equal in steady state by A_5), H is the Hall
1214 coefficient, σ is the conductivity, n is the electron density, p_0 is the background hole density, and p_h
1215 is the hole density (for a p-type material).

1216 Solving the coupled axiom system A_1 - A_7 with the measured $\sigma(I)$ and $H(I)$ data simultaneously
1217 yields the densities and mobilities of both majority and minority carriers, plus recombination life-
1218 time, diffusion length, and recombination coefficient, providing complete carrier characterization
1219 from a single measurement technique for photovoltaic material optimization.

1220 To derive Q , we eliminate $(\mu_N, \Delta p, n, p_h)$ using A_1, A_3 - A_5 , i.e., $\mu_N = \beta\mu_P$, $\Delta p = \Delta n$, $n = \Delta n$,
1221 $p_h = p_0 + \Delta n$. Substituting into A_6 gives $\sigma = e\mu_P[p_0 + \Delta n(1 + \beta)]$, and substituting into A_7
1222 gives $eH[p_0 + \Delta n(1 + \beta)]^2 = r[p_0 + \Delta n(1 - \beta^2)]$; replacing $[p_0 + \Delta n(1 + \beta)]^2$ by $\sigma^2/(e^2\mu_P^2)$
1223 yields $Q = 0$. To obtain the polynomial form over measured variables, we multiply the A_7 -derived
1224 relation by $e\mu_P$ and use $e\mu_P(p_0 + \Delta n) = \sigma - e\mu_P\beta\Delta n$ from the σ -equation, then expand to get

$$re\mu_P\Delta n\beta^2 + re\mu_P\Delta n\beta - r\sigma + ep_0\sigma H + e\Delta n\beta\sigma H + e\Delta n\sigma H = 0.$$

1225 which is a polynomial version of the Photo-Hall equation. Equivalently, setting

$$\begin{aligned}\alpha_1 &= -\mu e^2\mu_P(\Delta n)^2\beta H - \mu e^2\mu_P p_0\Delta n H - \mu e^2\mu_P(\Delta n)^2 H - \mu e\mu_P p_0 r - \mu\mu_N e\Delta n r + \mu r\sigma, \\ \alpha_2 &= -e\mu_P^2 p_0 - \mu_N e\mu_P \Delta n - e\mu_P^2 \Delta n + \mu_P \sigma, \\ \alpha_3 &= -\mu e^2\mu_P^2 \Delta n\beta H - p_h \mu e^2\mu_P^2 H - \mu e\mu_P^2 \beta r + \mu\mu_N e\mu_P r + \mu e\mu_P^2 r - \mu_H e\mu_P^2, \\ \alpha_4 &= -\mu n e^2\mu_P^2 \beta^2 H - \mu e^2\mu_P^2 \Delta n\beta^2 H - 2p_h \mu e^2\mu_P^2 \beta H + \mu\mu_N e^2\mu_P \Delta n\beta H - \mu e\mu_P^2 \beta^2 r + \mu\mu_N e^2\mu_P p_0 H \\ &\quad + \mu\mu_N e^2\mu_P \Delta n H - \mu\mu_N e\mu_P \beta r + \mu\mu_N^2 e r - \mu_N \mu_H e\mu_P, \\ \alpha_5 &= -\mu e^2\mu_P^2 \Delta n\beta H - p_h \mu e^2\mu_P^2 H - \mu e\mu_P^2 \beta r + \mu\mu_N e\mu_P r + \mu e\mu_P^2 r - \mu_H e\mu_P^2, \\ \alpha_6 &= \mu e\mu_P \Delta n\beta H + \mu e\mu_P p_0 H + \mu e\mu_P \Delta n H - \mu\mu_P \beta r + \mu\mu_N r - \mu_H \mu_P, \\ \alpha_7 &= \mu e\mu_P^2.\end{aligned}$$

1226 we obtain

$$Q = \sum_{i=1}^7 \alpha_i A_i$$

1227 Illustrating nuances of axiom formatting and multipliers with Kepler's third law

1228 From section 3.2, recall the equation for Kepler's third law, which relates the orbital period of two
1229 celestial bodies to their distances and masses. It can be written as

$$Q := p - \sqrt{\frac{(d_1 + d_2)^3}{G(m_1 + m_2)}} = 0,$$

1230 where p is the orbital period (scaled by a factor of 2π), m_1, m_2 the masses, and d_1, d_2 their dis-
1231 tances from the center of mass. For circular orbits, we use the following axioms (with p non-

1232 dimensionalized so that $\omega p = 1$):

$$A_1 : (d_1 + d_2)^2 F_g - G m_1 m_2 = 0 \quad (4)$$

$$A_2 : F_c - m_2 d_2 \omega^2 = 0 \quad (5)$$

$$A_3 : F_g - F_c = 0 \quad (6)$$

$$A_4 : \omega p - 1 = 0 \quad (7)$$

$$A_5 : m_1 d_1 - m_2 d_2 = 0. \quad (8)$$

1233 Equation (4) gives gravitational force, (5) centrifugal force, (6) equates them, (7) relates frequency
1234 and period, and (8) defines the center of mass. To derive Q , we eliminate d_1 using (8), substitute
1235 $p = 1/\omega$ from (7), and replace ω^2 by combining (5)–(6) with (4). This yields

$$G(m_1 + m_2)p^2 - (d_1 + d_2)^3 = 0,$$

1236 which is Kepler's law in polynomial form. Equivalently, setting

$$\alpha_1 = -p^2, \alpha_2 = p^2(d_1 + d_2)^2, \alpha_3 = -p^2(d_1 + d_2)^2, \alpha_4 = m_2 d_2(\omega p + 1)(d_1 + d_2)^2, \alpha_5 = -d_2^2,$$

we obtain

$$Q = \alpha_1 A_1 + \alpha_2 A_2 + \alpha_3 A_3 + \alpha_4 A_4 + \alpha_5 A_5.$$

1237 Assume that we are missing knowledge of the equation of gravitational force – axiom A_1 –, but we
1238 have Kepler's law Q at hand (possibly derived from sufficient data; see [91, 16, 18]). At this point,
1239 not only A_1 but also $\alpha_1, \dots, \alpha_5$ are not known. We will assume they exist and attempt to find them
1240 from Q and the remaining known axioms $A_2 - A_5$. Here we view $\alpha_1 A_1$ as the (unknown) residual
1241 of Q with respect to $A_2, \dots, A_5$. We define the ideal $I = \langle A_2, A_3, A_4, A_5, Q \rangle$. This is the same as
1242 the ideal $\langle A_2, A_3, A_4, A_5, \alpha_1 A_1 \rangle$ as $Q = \sum_{i=1}^5 \alpha_i A_i$. Geometrically, $V(A_2, A_3, A_4, A_5, \alpha_1 A_1) =$
1243 $V(A_2, A_3, A_4, A_5) \cap V(\alpha_1 A_1)$. Since the latter is reducible, we have more irreducible components
1244 than we have for $V(A_2, A_3, A_4, A_5)$. Computing a primary decomposition of I then gives us the
1245 following associated primes:

Decomposition	
Variety 1	$\langle d_2, m_1, F_g, F_c, \omega p - 1 \rangle$
Variety 2	$\langle m_2, d_1, F_g, F_c, \omega p - 1 \rangle$
Variety 3	$\langle m_2, m_1, F_g, F_c, \omega p - 1 \rangle$
1246 Variety 4	$\langle F_c - F_g, m_1 d_1 - m_2 d_2, \omega p - 1, F_g p - \omega m_2 d_2, F_g p^2 - m_2 d_2,$ $\mathbf{F_g(d_1 + d_2)^2 - m_1 m_2 G}, d_1(d_1 + d_2)^2 - m_2 p^2 G, m_1 p^2 G - d_2(d_1 + d_2)^2,$ $\omega d_1^2(d_1 + d_2)^2 - m_1 p G, \omega d_1^3 - 3\omega d_1 d_2^2 - 2\omega d_2^3 + 2m_1 p G - m_2 p G \rangle$
Generators selected by Reasoning module	
$\mathbf{F_g(d_1 + d_2)^2 - m_1 m_2 G}$	

1247 The generators of these irreducible components/ideals represent key equations that define each
1248 irreducible component. We iterate through each generator $\hat{A}_1$ and compute a Gröbner basis
1249 to eliminate the variables not appearing in Q . We find that only the correct expression -
1250 $\mathbf{F_g(d_1 + d_2)^2 - m_1 m_2 G}$ - can derive Q and is therefore the only candidate returned by our sys-
1251 tem.

1252 **Axiom Formatting and Multipliers.** Since the reducibility is introduced by $V(\alpha_i A_i)$, we do not
1253 distinguish between α_i and A_i when searching for i th axiom candidates. Additionally, we may
1254 return axioms in alternative forms; the exact expression is not guaranteed. For example, consider
1255 the same Kepler axioms and assume we are missing axiom A_4 instead. Then the relevant ideals in
1256 the primary decomposition are:

Decomposition	
Variety 1	$\langle F_c - F_g, m_1d_1 - m_2d_2, \omega p - 1, F_gp - \omega m_2d_2, F_gp^2 - m_2d_2, F_g(d_1+d_2)^2 - m_1m_2G, d_1(d_1+d_2)^2 - m_2p^2G, m_1p^2G - d_2(d_1+d_2)^2, \omega d_2(d_1+d_2)^2 - m_1pG, \omega(d_1 - 2d_2)(d_1+d_2)^2 + 2m_1pG - m_2pG \rangle$
Variety 2	$\langle F_c - F_g, m_1d_1 - m_2d_2, \omega p + 1, F_gp + \omega m_2d_2, F_gp^2 - m_2d_2, F_g(d_1+d_2)^2 - m_1m_2G, d_1(d_1+d_2)^2 - m_2p^2G, m_1p^2G - d_2(d_1+d_2)^2, \omega d_2(d_1+d_2)^2 + m_1pG, \omega(d_1 - 2d_2)(d_1+d_2)^2 - 2m_1pG + m_2pG \rangle$
Generators selected by Reasoning module	
	$\omega p - 1$
	$F_gp^2 - m_2d_2$
	$\omega p + 1$

The highlighted polynomials are returned by our system as being axiom candidates that derive Q along with the remaining axioms. The other generators have been omitted for readability, and cannot derive Q . Although the correct expression $\omega p - 1$ is present, additional terms are present. The term $F_gp^2 - m_2d_2$ is a restatement of A_2 using A_3 and the unknown $\omega p - 1$ substituted in. It contains the same information as $\omega p - 1$ since the former is already known. The term $\omega p + 1$ is an artifact of the fact that α_4 as noted before is $m_2d_2(\omega p + 1)(d_1 + d_2)^2$. Note that the variety $V(m_2d_2(\omega p + 1)(d_1 + d_2)^2)$ is reducible and can be written as the union: $V(m_2) \cup V(d_2) \cup V(\omega p + 1) \cup V((d_1 + d_2)^2)$. Therefore, $\omega p + 1$ appears as a viable axiom candidate since there is an algebraic derivation of Q from the remaining axioms along with $\omega p + 1$. Incidentally, the remaining terms m_2 , d_2 , and $(d_1 + d_2)^2$ also show up in the associated primes:

Decomposition	
Variety 1	$\langle F_c - F_g, m_2^2, (d_1+d_2)^2, m_1d_2 + d_1m_2 + 2m_2d_2, m_1m_2, m_1d_1 - m_2d_2, m_1^2, F_gm_2, F_gm_1, F_g^2, F_g - \omega^2m_2d_2 \rangle$
Variety 2	$\langle d_2, m_1, F_g, F_c \rangle$
Variety 3	$\langle m_2, d_1, F_g, F_c \rangle$
Generators selected by Reasoning module	
	$(d_1+d_2)^2$
	d_2
	m_2

In order to filter out these terms, for non-numerical examples, we apply an additional layer of filtering detailed in Appendix E.2.

D.1 Detailed Breakdown of Single Axiom Removal Experiment

We tested the system by removing one axiom at a time and asking it to recover the missing axiom, reporting our findings in Table 6. Across 12 problems, the method succeeded in 73 out of 75 cases (97.3%). Recovery was perfect for Kepler (5/5), Time Dilation (5/5), Escape Velocity (5/5), Light Damping (5/5), Hagen-Poiseuille (4/4), Neutrino Decay (5/5), Simple Harmonic Oscillator (6/6), Carrier-Resolved Photo-Hall (7/7), and Compton Scattering (10/10). Two additional failures occurred in the Relativistic Laws system (5/7 recovered), where axioms A_4 (light path geometry) and A_7 (parallel light travel time) were not recovered; these quadratic constraints factor into linear generators during primary decomposition, and the system returns the factored forms rather than the original quadratic axioms (see Section 3.3 for discussion).

Table 6: Results across problems; axioms are listed per problem as $A_1, A_2, \dots$. "Recovered" indicates whether the method successfully derived the removed axiom.

Problem	Axiom	Expression	Recovered
Hagen–Poiseuille	A1	$u - c_0 - c_2 r^2$	✓
	A2	$\mu \frac{\partial}{\partial r} \left(r \frac{\partial u}{\partial r} \right) - r \frac{dp}{dx}$	✓
	A3	$c_0 + c_2 R^2$	✓
	A4	$L \frac{dp}{dx} = -\Delta p$	✓
Kepler	A1	$(d_1 + d_2)^2 F_g - m_1 m_2$	✓
	A2	$F_c - m_2 d_2 \omega^2$	✓
	A3	$F_c - F_g$	✓
	A4	$\omega p - 1$	✓
	A5	$m_1 d_1 - m_2 d_2$	✓
Time Dilation	A1	$c dt_0 - 2d$	✓
	A2	$4L^2 - 4d^2 - v^2 dt^2$	✓
	A3	$f_0 dt_0 - 1$	✓
	A4	$f dt - 1$	✓
	A5	$c dt - 2L$	✓
Escape Velocity	A1	$K_i - \frac{1}{2} m v_e^2$	✓
	A2	$K_f = 0$	✓
	A3	$U_i r + GMm$	✓
	A4	$U_f = 0$	✓
	A5	$K_i + U_i - (K_f + U_f)$	✓
Light Damping	A1	$Sr^2 - q_c^2 a_p^2 \sin^2 \theta$	✓
	A2	$dA - 2\pi r^2 \sin \theta d\theta$	✓
	A3	$P - \int_0^\pi S dA$	✓
	A4	$\frac{4}{3} - \int_0^\pi \sin^3 \theta d\theta$	✓
	A5	$a_p^2 - \frac{1}{2} w^4 x_0^2$	✓
Neutrino Decay	A1	$p_\nu - p_\mu$	✓
	A2	$E_\pi - m_\pi$	✓
	A3	$E_\nu - p_\nu$	✓
	A4	$E_\pi - E_\mu - E_\nu$	✓
	A5	$E_\mu^2 - p_\mu^2 - m_\mu^2$	✓
Simple Harmonic Osc.	A1	$a_d - g \sin \theta$	✓
	A2	$d - L\theta$	✓
	A3	$T\omega - 2\pi$	✓
	A4	$d\omega^2 - a_d$	✓
	A5	$T_j - jT$	✓
	A6	$\sin \theta - \theta$	✓
Carrier-Resolved PH	A1	$\beta \mu_P - \mu_N$	✓
	A2	$\mu_H - r\mu$	✓
	A3	$p_h - p_0 - \Delta p$	✓
	A4	$n - \Delta n$	✓
	A5	$\Delta p - \Delta n$	✓
	A6	$\sigma - e p_h \mu_P - e n \mu_N$	✓
	A7	$H(p_h + \beta n)^2 e - r p_h + r \beta^2 n$	✓
Relativistic Laws	A1	$f_0 t_0 - 1$	✓
	A2	$ft - 1$	✓

Problem	Axiom	Expression	Recovered
	A3	$ct_0 - 2L_0$	✓
	A4	$\frac{c^2 t^2}{4} - L_0^2 - \frac{v^2 t^2}{4}$	✗
	A5	$m_0 u_0 - m u_y$	✓
	A6	$u_0 t_0 - u_y t$	✓
	A7	$t(c^2 - v^2) - 2Lc$	✗
Hall Effect	A1	$F_m - q_e v B$	✓
	A2	$F_e - q_e E$	✓
	A3	$F_m - F_e$	✓
	A4	$Eh - U_H$	✓
	A5	$v dt - L$	✓
	A6	$I dt - Q$	✓
	A7	$Q - N q_e$	✓
Inelastic Relativistic Collision	A1	$p_m^2 (c^2 - v_m^2) - m_m^2 v_m^2 c^2$	✓
	A2	$E_m^2 - (m_m c^2)^2 - (p_m c)^2$	✓
	A3	$E_r - m_r c^2$	✓
	A4	$E_c^2 - (m_c c^2)^2 - (p_c c)^2$	✓
	A5	$2E_m E_r - E_c^2 + E_m^2 + E_r^2$	✓
	A6	$p_c - p_m$	✓
	A7	$v_m - \frac{4}{5}c$	✓
	A8	$m_r - m_m$	✓
	A9	$p_c^2 (c^2 - v_c^2) - m_c^2 v_c^2 c^2$	✓
Compton Scattering	A1	$E_1 + E e_1 - E_2 - E e_2$	✓
	A2	$E_1 - h f_1$	✓
	A3	$E_2 - h f_2$	✓
	A4	$p_1 c - h f_1$	✓
	A5	$p_2 c - h f_2$	✓
	A6	$\lambda_1 f_1 - c$	✓
	A7	$\lambda_2 f_2 - c$	✓
	A8	$E e_1 - m_e c^2$	✓
	A9	$E e_2^2 - c^2 p e_2^2 - m_e^2 c^4$	✓
	A10	$p e_2^2 - p_2^2 - p_1^2 + 2p_1 p_2 \cos \theta$	✓

E Additional Technical Details

E.1 Differences Between Real and Complex Solutions

Let $A_1, \dots, A_k$ be a collection of polynomials with $A_1 = 0, \dots, A_k = 0$ being the associated axioms. Let $\mathcal{A}$ stand for the set $\{A_1, \dots, A_k\}$. Let $V_R(\mathcal{A})$, respectively $V_C(\mathcal{A})$, be the set of real, respectively complex, solutions of the equations $A_1 = 0, \dots, A_k = 0$. Note that $V_R(\mathcal{A})$ is the same as the variety $V(\mathcal{A})$ defined in the main body of the paper.

First note that $V_C(\mathcal{A}, Q) = V_C(\mathcal{A})$ if and only if $Q = 0$ for all complex solutions of $A_1 = 0, \dots, A_k = 0$. Further, if $Q \in I(\mathcal{A})$, then $V_C(\mathcal{A}, Q) = V_C(\mathcal{A})$. However, the converse is not true. By Hilbert's Nullstellensatz, if $V_C(\mathcal{A}, Q) = V_C(\mathcal{A})$, then $Q^m \in I(\mathcal{A})$ for some positive integer m (but m need not equal 1). But $Q^m(\bar{x}) = 0$ if and only if $Q(\bar{x}) = 0$. Accordingly, if we say that $Q = 0$ is derivable from $\mathcal{A}$ if and only if $Q^m \in I(\mathcal{A})$, then the notions of “a consequence of” and “derivable from” are identical assuming complex-valued solutions are relevant.

On the other hand, if we are only interested in real-valued solutions, the situation is more complicated. As before, $V_R(\mathcal{A}, Q) = V_R(\mathcal{A})$ if and only if $Q = 0$ for all real solutions of $A_1 = 0, \dots, A_k = 0$. Further, if $Q \in I(\mathcal{A})$, then $V_R(\mathcal{A}, Q) = V_R(\mathcal{A})$. However, the converse is not true. If $V_R(\mathcal{A}, Q) = V_R(\mathcal{A})$, then $Q^{2m} + S \in I(\mathcal{A})$ for some positive integer m and some polynomial S that is a sum of squares of polynomials in R .

E.2 Further Filtering for Projections

While testing ideal membership via Gröbner basis elimination (as described in Section C.1) establishes that Q is derivable from the candidate axioms together with the known axioms, this criterion alone can admit *trivial* candidates that make the system degenerate. For instance, a candidate axiom that simply sets a variable to zero (e.g., $m_2 = 0$) may technically allow Q to be derived, but only because it collapses the variety onto a lower-dimensional subspace where Q holds vacuously. To filter such cases, we introduce a stronger derivability test based on *exact containment* in the elimination ideal.

Exact Containment Test. Let $\mathcal{A} = \{A_1, \dots, A_k\}$ denote the known axioms, $\hat{A}_{k+1}$ a candidate axiom, and Q the target hypothesis defined over variables $\{x_1, \dots, x_d\}$. We form the augmented ideal $J = \langle A_1, \dots, A_k, \hat{A}_{k+1} \rangle$ and compute its reduced Gröbner basis $\mathcal{G}$ with respect to a lexicographic term order. The elimination ideal $J \cap \mathbb{R}[x_1, \dots, x_d]$ is generated by $\mathcal{G}' = \mathcal{G} \cap \mathbb{R}[x_1, \dots, x_d]$.

Weak test (ideal membership): We accept $\hat{A}_{k+1}$ if $Q \in \langle \mathcal{G}' \rangle$, i.e., if Q reduces to zero modulo $\mathcal{G}'$. This confirms that Q is derivable from the augmented system.

Strong test (exact containment): We accept $\hat{A}_{k+1}$ only if $Q \in \mathcal{G}'$, i.e., if Q appears *exactly* as an element of the Gröbner basis (up to scalar multiplication). This ensures that the projection onto $\text{vars}(Q)$ yields Q directly, rather than a stronger polynomial from which Q follows as a consequence.

Geometric Interpretation. Geometrically, the weak test verifies that $V(\mathcal{A} \cup \{\hat{A}_{k+1}\})$ projects onto $V(Q)$ when restricted to the coordinates of Q . However, this projection may pass through a lower-dimensional intermediate variety. The strong test ensures that the projection is *minimal*: we project directly onto $V(Q)$ without factoring through a smaller-dimensional subvariety.

For example, consider Kepler’s third law in polynomial form $Q : G(m_1 + m_2)p^2 - (d_1 + d_2)^3 = 0$, defined over variables $\{p, m_1, m_2, d_1, d_2\}$. A trivial candidate axiom $\hat{A}_{k+1} : m_2 = 0$ would satisfy the weak test, since setting $m_2 = 0$ implies $Gm_1p^2 - (d_1 + d_2)^3 = 0$, which algebraically contains Q . However, Q itself does not appear in the elimination ideal; instead, we derive a stronger constraint $Gm_1p^2 - (d_1 + d_2)^3 = 0$ defined over fewer variables (without m_2). The strong test rejects such candidates, filtering out degenerate cases where the system collapses onto boundary conditions.

Limitations and Practical Implementation. The strong test is not foolproof. If the background theory $\mathcal{A}$ already contains an axiom defined over a strict subset of $\text{vars}(Q)$, then the elimination ideal will necessarily include polynomials over fewer variables than Q . For instance, in the Kepler system, the center-of-mass axiom $m_1d_1 - m_2d_2 = 0$ is defined over $\{m_1, m_2, d_1, d_2\}$, a proper subset of the variables in Kepler’s third law. Consequently, $\mathcal{G}'$ will contain polynomials over these variables, and Q may not appear exactly in $\mathcal{G}'$ even for legitimate candidate axioms.

To handle this, we adopt a two-tier filtering strategy:

1. **Primary filter (strong test):** Accept candidates where $Q \in \mathcal{G}'$ exactly. These are high-confidence results.
2. **Secondary filter (weak test):** Accept candidates where $Q \in \langle \mathcal{G}' \rangle$ but $Q \notin \mathcal{G}'$, flagging them as lower-confidence. These may include legitimate axioms whose projection is mediated by background constraints.

In our experiments, we save both sets of candidates, noting which satisfy the strong versus weak test. For systems with hierarchical variable dependencies (e.g., conservation laws involving subsets of state variables), the weak test often recovers correct axioms that the strong test would reject.

E.3 Algorithm Implementation Details

Algorithm 1 describes our procedure for identifying missing axioms that render a target hypothesis Q derivable. The algorithm proceeds in three phases that parallel our methodological framework: **Encode**, **Decompose**, and **Reason**, with separate pathways for exact (symbolic) and noisy (numerical) scenarios. We now describe each step in detail.

Input and Initialization. The input consists of a finite set of polynomial axioms $\mathcal{A} = \{A_1, \dots, A_k\}$ in variables $\mathcal{X} = \{x_1, \dots, x_n\}$, a hypothesis polynomial Q over a restricted subset of variables $\{x_1, \dots, x_d\}$ (where $d \leq n$), and a flag *noisy* indicating whether the consequences

1349 contain numerical error. The algorithm maintains a set $\mathcal{A}_{\text{expl}}$ of explanatory candidates, initially
 1350 empty.

1351 **Encode.** We form the polynomial ideal $I = \langle A_1, \dots, A_k, Q \rangle$. Geometrically, $V(I)$ corresponds
 1352 to the variety of solutions simultaneously satisfying the known axioms and the hypothesis. If Q is
 1353 derivable from $\mathcal{A}$, then $V(I)$ coincides with $V(\mathcal{A})$; otherwise, additional structure appears in the
 1354 decomposition. This encoding step is identical for both exact and noisy cases.

1355 **Decompose.** The decomposition strategy differs based on whether consequences are exact or
 1356 noisy:

1357 *Exact case (symbolic):* We compute a primary decomposition

$$I = I_1 \cap \dots \cap I_r,$$

1358 yielding associated primes $P_j = \sqrt{I_j}$. Each P_j describes an irreducible component of the solution
 1359 space. The generators of P_j form a finite set $\{\hat{A}_{k+1}^{(j,1)}, \dots, \hat{A}_{k+1}^{(j,m_j)}\}$ of potential candidate axioms.

1360 In our experiments, we use Singular [22], a computer algebra software, for computing the primary
 1361 decompositions. We also have the functionality for using Macaulay2 [35], another standard com-
 1362 puter algebra software, though Singular’s default primary decomposition implementation is more
 1363 efficient.

1364 *Noisy case (numerical):* We apply numerical irreducible decomposition [6] to compute witness sets
 1365 $\{W_1, \dots, W_s\}$, where each W_i is a finite collection of sample points on an irreducible component.
 1366 For each witness set W_i and each removed axiom $A_\ell \in \mathcal{A}$, we perform symbolic regression: we form
 1367 the evaluation matrix Φ using the monomial support of A_ℓ evaluated at points in W_i , then compute
 1368 the right singular vector corresponding to the smallest singular value to obtain fitted coefficients $\hat{c}_\ell$.
 1369 This yields a candidate axiom $\hat{A}_{k+1}^{(i,\ell)} = \sum_k \hat{c}_{\ell k} M_{\ell k}$ for each component-axiom pair.

1370 We use Bertini [7], an alternate numerical algebraic geometry package with parallelization imple-
 1371 mentation to compute the witness sets. We also build in the optionality to use Macaulay2’s numerical
 1372 algebraic geometry package [53].

1373 **Reason.** For each candidate axiom $\hat{A}_{k+1}$ obtained from decomposition (or sets of candidate ax-
 1374 ioms if we know that multiple axioms are missing), we test whether it (or they) enable(s) derivation
 1375 of Q . The reasoning method depends on whether numerical constants are present:

1376 *Exact case (algebraic projection):* We form the augmented ideal $J = \langle A_1, \dots, A_k, \hat{A}_{k+1} \rangle$ (resp.
 1377 $\langle A_1, \dots, A_k, \hat{A}_{k+1}, \dots, \hat{A}_{k+l} \rangle$ for testing a set of multiple candidate axioms) and compute its re-
 1378 duced Gröbner basis $\mathcal{G}$ with respect to a lexicographic order on $\mathcal{X}$. We project $\mathcal{G}$ to the coordinate
 1379 ring of the observed variables: $\mathcal{G}' = \mathcal{G} \cap \mathbb{R}[x_1, \dots, x_d]$. If $Q \in \mathcal{G}'$, then $\hat{A}_{k+1}$ is an *explana-*
 1380 *tory axiom* for Q and is added to $\mathcal{A}_{\text{expl}}$. The Gröbner basis computations here are done entirely in
 1381 Macaulay2 [35].

1382 *Noisy case (existential reasoning):* Following the approach of AI-Descartes [16], we abstract nu-
 1383 merical constants in both the hypothesis Q and the candidate axiom $\hat{A}_{k+1}$ (resp. $\hat{A}_{k+1}, \dots, \hat{A}_{k+l}$ for
 1384 testing a set of multiple candidate axioms) into existentially quantified variables $c_1, \dots, c_s$, yielding
 1385 Q' and $\hat{A}'_{k+1}$. We encode the derivability query as:

$$\exists c_1 \dots \exists c_s \left(C' \wedge \forall x_1 \dots \forall x_n ((C \wedge \mathcal{A} \wedge \hat{A}'_{k+1}) \rightarrow Q') \right),$$

1386 where C and C' denote admissibility constraints on the variables. We submit this query to a theorem
 1387 prover, KeYmaera X [65] in our implementation. If the prover returns *Proved*, then $\hat{A}_{k+1}$ (or the
 1388 set of candidates) is added to $\mathcal{A}_{\text{expl}}$.

1389 **Output.** The algorithm returns the set $\mathcal{A}_{\text{expl}}$ of all candidate axioms that, when added to the back-
 1390 ground theory $\mathcal{A}$, enable derivation of Q . In practice, we also return metadata including fitted
 1391 coefficients (for the noisy case), residual norms (quality metrics), and the component indices from
 1392 which each candidate was extracted.

1393 **Computational Complexity.** The dominant computational costs are: (i) primary decomposition
 1394 or witness set computation (doubly exponential in the worst case for symbolic methods [66], poly-
 1395 nomial per witness set for numerical methods), (ii) symbolic regression (linear algebra on $N \times K$
 1396 matrices for N witness points), and (iii) Gröbner basis computation (doubly exponential worst-
 1397 case [13], though often tractable in practice for structured systems) or theorem proving. Our exper-
 1398 iments demonstrate practical feasibility on systems with up to 16 variables and 10 axioms.

Algorithm 1 Identifying Explanatory Axioms via Decomposition and Reasoning

Input: Axioms $\mathcal{A} = \{A_1, \dots, A_k\}$ over variables $\mathcal{X} = \{x_1, \dots, x_n\}$, hypothesis Q over $\{x_1, \dots, x_d\}$, flag *noisy*
Output: Set $\mathcal{A}_{\text{expl}}$ of explanatory axioms
 1: Initialize $\mathcal{A}_{\text{expl}} \leftarrow \emptyset$
 2: **Encode:** Form ideal $I \leftarrow \langle A_1, \dots, A_k, Q \rangle$
 3: **Decompose:**
 4: **if** *noisy* **then**
 5: Compute witness sets $\{W_1, \dots, W_s\}$ via numerical irreducible decomposition
 6: **for** each witness set W_i **do**
 7: **for** each removed axiom $A_\ell \in \mathcal{A}$ **do**
 8: Form evaluation matrix Φ from monomial support of A_ℓ on W_i
 9: Compute $\hat{c}_\ell$ via smallest singular value of Φ
 10: Construct candidate $\hat{A}_{k+1}^{(i,\ell)} \leftarrow \sum_k \hat{c}_{\ell k} M_{\ell k}$
 11: **Reason:** Abstract numerical constants in $\hat{A}_{k+1}^{(i,\ell)}$ and Q to obtain $\hat{A}_{k+1}^{(i,\ell)'}$ and Q'
 12: Formulate existential query: $\exists c_1 \dots \exists c_s (\mathcal{C}' \wedge \forall \mathcal{X} ((\mathcal{C} \wedge \mathcal{A} \wedge \hat{A}_{k+1}^{(i,\ell)'}) \rightarrow Q'))$
 13: **if** theorem prover returns *Proved* **then**
 14: Add $\hat{A}_{k+1}^{(i,\ell)}$ to $\mathcal{A}_{\text{expl}}$
 15: **end if**
 16: **end for**
 17: **end for**
 18: **else**
 19: Compute primary decomposition $I = I_1 \cap \dots \cap I_r$
 20: **for** each associated prime $P_j = \sqrt{I_j}$ **do**
 21: **for** each generator $\hat{A}_{k+1}^{(j,m)}$ of P_j **do**
 22: **Reason:** Form ideal $J \leftarrow \langle A_1, \dots, A_k, \hat{A}_{k+1}^{(j,m)} \rangle$
 23: Compute Gröbner basis $\mathcal{G} \leftarrow \text{GB}(J)$ with lex order $x_1 < \dots < x_n$
 24: Project: $\mathcal{G}' \leftarrow \mathcal{G} \cap \mathbb{R}[x_1, \dots, x_d]$
 25: **if** $Q \in \mathcal{G}'$ **then**
 26: Add $\hat{A}_{k+1}^{(j,m)}$ to $\mathcal{A}_{\text{expl}}$
 27: **end if**
 28: **end for**
 29: **end for**
 30: **end if**
 31: **Return** $\mathcal{A}_{\text{expl}}$

1399 **Implementation and Hardware**

1400 All symbolic computations were performed using Macaulay2 [35] and Singular [22] for Gröbner
 1401 basis and primary decomposition and Python/SymPy (for preprocessing and verification). Nu-
 1402 merical irreducible decomposition was implemented using Bertini [6, 7] and Macaulay2’s
 1403 NumericalAlgebraicGeometry package [53]. Symbolic regression on witness sets and theorem
 1404 proving with KeYmaera X [65] were implemented in Python. Experiments were run on an Ap-
 1405 ple M4 MacBook Pro with 16 GB of unified memory. The M4 chip features a 10-core CPU and a
 1406 10-core GPU, with hardware acceleration for polynomial arithmetic and linear algebra routines via
 1407 Apple’s Accelerate framework.

1408 Table 7 reports average runtime per test case that ran within the timeout window (7200s) across all
 1409 benchmark problems, broken down by the number of missing axioms and computational mode (al-
 1410gebraic vs. numerical). Algebraic methods scale efficiently for smaller systems, completing single-

1411 axiom tests in under 10 seconds for most problems and maintaining sub-minute performance for
1412 systems with up to 10 variables. Larger systems (15–16 variables) require 1–2 minutes per single-
1413 axiom test, with moderate increases for multi-axiom cases. Numerical methods incur substantially
1414 higher computational costs due to homotopy continuation and witness set construction, typically
1415 requiring 10–100× longer than their algebraic counterparts. For the most complex systems (Rela-
1416 tivistic Laws, Hall Effect, Compton Scattering), numerical methods require 1–2 hours per test case
1417 that succeeded.

Problem	Vars	Ax.	Deg.	Algebraic (s)			Numerical (s)		
				1-Ax	2-Ax	3-Ax	1-Ax	2-Ax	3-Ax
Hagen–Poiseuille	9	4	4	0.3	1.1	1.2	3.3	9.5	10.2
Kepler	8	5	4	1.4	6.2	5.8	15.6	30.4	32.5
Time Dilation	8	5	4	4.2	5.9	4.4	20.7	29.3	35.2
Escape Velocity	9	5	4	2.7	3.5	5.4	15.7	26.3	27.2
Light Damping	10	5	8	85.3	106.2	109.9	792.1	1022.2	1032.6
Neutrino Decay	8	5	2	3.4	7.5	9.5	41.2	268.3	251.4
Simple Harmonic Osc.	10	6	5	9.2	13.2	12.5	341.2	690.1	798.3
Carrier-Resolved PH	15	7	6	14.5	59.9	87.2	431.5	1034.2	1242.2
Relativistic Laws	12	7	5	144.2	231.3	202.1	5239.2	6824.4	6523.4
Hall Effect	16	7	4	45.3	104.8	131.1	892.3	3821.9	3733.2
Inelastic Rel. Col.	11	9	6	87.2	163.6	150.2	1222.2	4757.4	5366.3
Compton Scattering	15	10	6	140.1	160.7	172.9	2419.5	6324.2	6823.4

Table 7: Average runtime per test case for single, pair, and triple axiom removal experiments. Algebraic computations use primary decomposition and Gröbner basis methods; numerical computations use witness sets, symbolic regression, and an automated theorem prover. Results are averaged over all successful test cases for each problem.

1418 This being said, it is also important to note that there were many test cases that failed to run within
1419 the allotted time limit of (7200s), especially when 3 axioms were missing. See Table 8 for the
1420 number of completed runs for each case.

Problem	Vars	Ax.	Deg.	Algebraic			Numerical		
				1-Ax	2-Ax	3-Ax	1-Ax	2-Ax	3-Ax
Hagen–Poiseuille	9	4	4	4/4	6/6	4/4	12/12	18/18	12/12
Kepler	8	5	4	5/5	10/10	10/10	15/15	30/30	30/30
Time Dilation	8	5	4	5/5	10/10	10/10	15/15	30/30	30/30
Escape Velocity	9	5	4	5/5	10/10	10/10	15/15	30/30	30/30
Light Damping	10	5	8	5/5	10/10	10/10	15/15	30/30	30/30
Neutrino Decay	8	5	2	5/5	10/10	10/10	15/15	30/30	30/30
Simple Harmonic Osc.	10	6	5	6/6	15/15	20/20	18/18	45/45	60/60
Carrier-Resolved PH	15	7	6	7/7	21/21	35/35	21/21	63/63	105/105
Relativistic Laws	12	7	5	7/7	21/21	35/35	9/21	18/63	102/105
Hall Effect	16	7	4	7/7	21/21	35/35	21/21	63/63	105/105
Inelastic Rel. Col.	11	9	6	9/9	36/36	84/84	27/27	108/108	252/252
Compton Scattering	15	10	6	10/10	45/45	120/120	3/30	5/135	2/360

Table 8: Test case completion rates within the 7200s timeout window. Values show the number of completed runs out of total test cases for each configuration. Algebraic computations use primary decomposition and Gröbner basis methods; numerical computations use witness sets and symbolic regression. The numerical cases are totaled over all three noise levels that were tested on.

1421 Algebraic methods completed all test cases within the timeout window across all problems and
1422 configurations. Numerical methods achieved perfect completion for most problems, but exhibited

significant timeouts for two systems. Relativistic Laws completed only 43% (9/21) of single-axiom cases, 29% (18/63) of pair cases, and 97% (102/105) of triple cases, likely due to the high computational cost of witness set construction for systems with 12 variables and degree-5 polynomials. Compton Scattering experienced severe timeouts, completing only 10% (3/30) of single-axiom cases, 4% (5/135) of pair cases, and less than 1% (2/360) of triple cases - attributable to its large scale (15 variables, 10 axioms, degree 6) combined. Light Damping, having the highest polynomial degree (8), had an increase in runtime as noted in Table 7, however given the manageable number of variables and axioms, did terminate across all runs.

1431 F Comparison with State-of-the-Art

1432 We compare our system against two representative approaches: (i) **cvc5**-based abduction, and (ii)
1433 **LLM**-based abduction using GPT-5 Pro. Both baselines receive the same input family of physics
1434 systems and targets, and both are evaluated by the same algebraic verifier (Macaulay2) under iden-
1435 tical projection checks. We also note a few systems that we cannot benchmark against due to our
1436 more challenging use case of not assuming we have data for the unknown axiom.

1437 **Benchmarks.** We use each of the 12 systems from Table 1 and their noisy counterparts to bench-
1438 mark the competing methods. For each problem we iterate over all k -tuples of removed axioms with
1439 $k \in \{1, 2, 3\}$ (configurable).

1440 cvc5

1441 CVC5 is a state-of-the-art SMT (Satisfiability Modulo Theories) solver [3] that supports reasoning
1442 over a wide range of logical theories, including nonlinear real arithmetic.

1443 We instantiate our setting in CVC5¹ via its native (experimental) *abduct* API, which synthesizes a
1444 candidate axiom A_{k+1} such that the asserted axioms $\mathcal{A}_{\text{given}}$ together with A_{k+1} entail the target Q
1445 while remaining jointly satisfiable. Concretely, we enable abduction (produce-abducts) and query
1446 the solver with `getAbduct`. Optionally, by setting `incremental = true` we activate the *incremental mode*,
1447 which allows to generate more than one candidate hypothesis axiom via `getAbductNext`.
1448 These calls are documented in the official CVC5 Python API.

1449 Each system’s axioms $A_1, \dots, A_k$ and target Q are encoded as quantifier-free nonlinear real arith-
1450 metic (QF_NRA): all equations are written in polynomial form, and asserted into the solver. We then
1451 pose the abductive query “find A_{k+1} ” with a user-provided grammar restricting the search space of
1452 candidate polynomials.

1453 We supply a SyGuS grammar (constructed via `mkGrammar/addRule`) whose start symbol expands
1454 to equations of the form $P = 0$, where P ranges over polynomials of bounded total degree. The
1455 grammar includes terminals for constants and variables (closed under addition) and non-terminals
1456 degree-stratified polynomials (closed under addition, scalar multiplication, and degree-respecting
1457 products, e.g., for degree 3 we can have linear $\times$ quadratic). All grammars are generated program-
1458 matically from the variable list for each system and capped at a degree inferred from the target (with
1459 optional rational coefficient seeds for noisy variants).

1460 As mentioned above we use `getAbduct` together with the grammar to obtain one (or more using
1461 `getAbductNext`) candidate A_{k+1} . We then perform two solver-side sanity checks before sending
1462 candidates to our Macaulay2 verifier: (i) Consistency: $\mathcal{A}$ is satisfiable; and (ii) *Entailment*: $\mathcal{A} \wedge \neg Q$
1463 is unsatisfiable.

1464 To standardize runtimes across problem sizes, we enforce per-query budgets via CVC5’s time and
1465 resource limits (e.g., `tlimit-per` in milliseconds). These are set through `setOption` and apply
1466 independently to each `abduct` or consistency check. For our experiments, we set a limit of 8 hours
1467 for each `cvc5` run.

1468 LLM

1469 We implement a large language model-based approach to abductive inference using the GPT-5 Pro
1470 model (Nov 2025) accessed through the OpenRouter API. For each physics system in our benchmark
1471 set, the model is provided with the system’s axioms, the list of measured and non-measured variables
1472 (anonymized for not introducing biases), and a target polynomial that cannot be derived from the

¹All experiments use CVC5 through its Python API (version 1.1.0).

1473 given axioms. We then remove one or more axioms (according to the experimental configuration)
 1474 and prompt the model to propose an additional axiom or set of axioms whose inclusion would make
 1475 the target polynomial derivable.

1476 To prevent leakage of semantic cues from variable names and to ensure consistent reasoning across
 1477 domains, all variables in each system are renamed canonically to $\{x_1, x_2, \dots, x_n\}$. The prompt
 1478 presents the remaining axioms and the target polynomial in this canonical form and instructs the
 1479 model to act as an abductive inference solver. The model may reason freely within its response, but
 1480 is required to conclude with a section formatted as:

Polynomial: < insert one or more candidate polynomials >

1481 Each candidate is required to be distinct from the target polynomial. We parse only the lines follow-
 1482 ing this tag as the model’s proposed abductive hypotheses.

1483 The proposed axioms are automatically translated back into the system’s original variable names
 1484 using the inverse of the canonical renaming map. For every candidate, we then invoke a symbolic
 1485 verification step implemented in Macaulay2. The candidate axiom(s) are added to the remaining
 1486 axioms, and the resulting ideal is projected onto the set of measured variables relevant to the target.
 1487 The two checks of *Membership* and *Literal appearance* are then performed corresponding respec-
 1488 tively to weak and strong forms of abductive success. The verification procedure is identical to that
 1489 used for algebraic abduction experiments with Macaulay2.

1490 A proposed axiom set $\mathcal{A}_{abd}$ is considered *valid* for target Q iff, when added to the given set of axioms
 1491 $\mathcal{A}_{given}$, the following hold after eliminating non-measured variables (per target) in Macaulay2:

- 1492 1. *Membership*: whether the target polynomial reduces to zero modulo the Gröbner basis of
 1493 the eliminated ideal (i.e., the target lies in the ideal generated by the axioms and proposed
 1494 axiom). That is, $Q \in \mathcal{A}$ where $\mathcal{A} = \mathcal{A}_{abd} \cup \mathcal{A}_{given}$ (remainder-zero modulo the Gröbner
 1495 basis of the eliminated ideal), and
- 1496 2. *Literal appearance (stronger)*: whether the target polynomial appears exactly among the
 1497 generators of the eliminated Gröbner basis. That is, q appears *exactly* among the generators
 1498 of the eliminated ideal’s Gröbner basis.

1499 We report both metrics. For ablations we also track (iii) parsimony (number of added axioms), (iv)
 1500 solution time, and (v) consistency (the added hypotheses are not already implied by $\mathcal{A}_{abd}$).

1501 All experiments are run using GPT-5 Pro (reasoning model) with temperature 0.7 and a maximum of
 1502 1200 tokens per completion. Each combination of system, number of removed axioms, and specific
 1503 subset of removed axioms is processed independently. For each case, the full prompt, model output,
 1504 parsed candidate polynomials (in canonical and original variables), and the results of the Macaulay2
 1505 verification are written to structured output directories for later analysis. This design enables exact
 1506 reproducibility of every abductive query and isolates LLM behavior from downstream algebraic
 1507 validation.

1508 Our results can be summarized in Table 4, shown again below.

Method	Single Axiom	Multi Axiom	Data- Free	Noise Robust	Var. Agnostic	Explainable
Traditional SR	✗	✗	✗	✗	✗	✗
Integrated SR	✗	✗	✗	✓	✓	✓
LLM (GPT-5 Pro)	0%	0%	✓	✗	✓	✓
cvc5	t-out	t-out	✓	✗	✗	✓
AI-Noether	97% ✓	49% ✓	✓	✓	✓	✓

Table 4: **Comparison of abductive and symbolic regression methods.** Green ticks (✓) indicate capability support; red crosses (✗) indicate absence. “Data-Free” means the method operates without numeric data or prior parameter information, and “Var. Agnostic” indicates invariance to variable naming. “t-out” means that the method ran out of time without producing a solution.

1509 **Summary of results.** Neither baseline successfully recovered missing axioms in our benchmark.
1510 The LLM-based approach (GPT-5 Pro) failed to generate any valid candidate axioms across all
1511 test cases, producing either syntactically malformed outputs or polynomials that did not satisfy the
1512 algebraic verification criteria when back-mapped to the original variable space. The cvc5 solver
1513 consistently timed out after 8 hours without producing any abducts, likely due to the combinato-
1514 rial explosion in the search space induced by the high polynomial degrees (up to degree 8 in our
1515 benchmarks) and the number of variables (up to 16). In contrast to traditional symbolic regression
1516 methods that require training data for each unknown axiom, and in contrast to these unsuccessful
1517 abductive baselines, AI-Noether operates in a purely data-free setting while achieving 97% recov-
1518 ery on single-axiom removal and 49% on multi-axiom removal. This demonstrates that algebraic-
1519 geometric decomposition provides a tractable pathway to abductive inference in scientific discovery
1520 that remains out of reach for both SMT-based synthesis and LLM-based reasoning under current
1521 computational constraints.