
Evolving Scientific Discovery by Unifying Data and Background Knowledge with AI Hilbert

Ryan Cory-Wright*

Department of Analytics, Marketing and Operations, Imperial College Business School, London, UK
r.cory-wright@imperial.ac.uk

Cristina Cornelio

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

Sanjeeb Dash

IBM Thomas J. Watson Research Center
Yorktown Heights, USA
sanjeebd@us.ibm.com

Bachir El Khadir

IBM Thomas J. Watson Research Center
Yorktown Heights, USA
bachir009@gmail.com

Lior Horesh

IBM Thomas J. Watson Research Center
Yorktown Heights, USA
lhoresh@us.ibm.com

Abstract

The discovery of scientific formulae that parsimoniously explain natural phenomena and align with existing background theory is a key goal in science. Historically, scientists have derived natural laws by manipulating equations based on existing knowledge, forming new equations, and verifying them experimentally. In recent years, data-driven scientific discovery has emerged as a viable competitor in settings with large amounts of experimental data. Unfortunately, data-driven methods often fail to discover valid laws when data is noisy or scarce. Accordingly, recent works combine regression and reasoning to eliminate formulae inconsistent with background theory. However, the problem of searching over the space of formulae consistent with background theory to find one that best fits the data is not well-solved. We propose a solution to this problem when all axioms and scientific laws are expressible via polynomial equalities and inequalities and argue that our approach is widely applicable. We model notions of minimal complexity using binary variables and logical constraints, solve polynomial optimization problems via mixed-integer linear or semidefinite optimization, and prove the validity of our scientific discoveries in a principled manner using Positivstellensatz certificates. The optimization techniques leveraged in this paper allow our approach to run in polynomial time with fully correct background theory (under an assumption that the complexity of our derivation is bounded), or non-deterministic polynomial (NP) time with partially correct background theory. We demonstrate that some famous scientific laws, including Kepler's Third Law of Planetary Motion, the Hagen-Poiseuille Equation, and the Radiated Gravitational Wave Power equation, can be derived in a principled manner from axioms and experimental data.

*Corresponding Author

1 Introduction

A fundamental problem in science involves explaining natural phenomena in a manner consistent with noisy experimental data and a body of potentially inexact and incomplete background knowledge about the universe’s laws [39]. In the past few centuries, The Scientific Method [101] has led to significant progress in discovering new laws. Unfortunately, the rate of emergence of these laws and their contribution to economic growth is stagnating relative to the amount of capital invested in deducing them [27, 21]. Indeed, Dirac [41] noted that it is now more challenging for first-rate physicists to make second-rate discoveries than it was previously for second-rate physicists to make first-rate ones, while Arora et al. [7] found that the marginal value of scientific discoveries to large companies has declined since the fall of the Berlin Wall. Moreover, Bloom et al. [24] have found that research productivity in the United States halves every thirteen years because good scientific ideas are getting harder to find. This phenomenon can be partly explained by analogy to [33]: The Scientific Method has picked most of the “low-hanging fruit” in science, such as natural laws that relate physical quantities using a small number of low-degree polynomials. This calls for more disciplined and principled alternatives to The Scientific Method, which integrate background information and experimental data to generate and verify higher dimensional laws of nature, thereby promoting scientific discovery [c.f. 69, 110]. Accordingly, Figure 1 provides an overview of these alternatives.

Even as the rate of scientific discovery has decreased, the scalability of global optimization methods has significantly improved. Indeed, as we argue in this paper, global optimization methods are now a mature technology capable of searching over the space of scientific laws — owing to Moore’s law and significant theoretical and computational advances by the optimization community [see 22, 58, 15, for reviews]. Indeed, Bertsimas and Dunn [18, Chap. 1] observed that the speedup in raw computing power between 1991 and 2015 is at least six orders of magnitude. Additionally, polynomial optimization has become much more scalable since the works of Lasserre [74] and Parrilo [87], and primal-dual interior-point methods [84, 92, 102] have improved considerably, with excellent implementations now available in, for example, the Mosek solver [6]. Indeed, even methods for non-convex quadratically constrained problems have achieved machine-independent speedups of nearly 200 since their integration within commercial solvers in 2019 [1, 58].

In this paper, we propose a new approach to scientific discovery that leverages these advances by the optimization community. Given a set of background axioms, theorems, and laws expressible as a basic semialgebraic set (i.e., a system of polynomial equalities and inequalities) and observations from experimental data, we derive new laws representable as polynomial expressions that are either exactly or approximately consistent with existing laws and experimental data by solving polynomial optimization problems via linear and semidefinite optimization. By leveraging fundamental results from real algebraic geometry, we obtain formal proofs of the correctness of our laws as a byproduct of the optimization problems. This is notable, because existing automated approaches to scientific discovery, as reviewed in Section 1.1, often rely upon deep learning techniques that do not provide formal proofs and are prone to “hallucinating” incorrect scientific laws that cannot be automatically proven or disproven, analogously to output from state-of-the-art Large Language Models such as GPT-4 [86]. As such, any new laws derived by these systems cannot easily be explained or justified.

Conversely, our approach discovers new scientific laws by solving an optimization problem to minimize a weighted sum of discrepancies between the proposed law and experimental data, plus the distance between the proposed law and its projection onto the set of symbolic laws derivable from background theory. As a result, our approach discovers scientific laws alongside a proof of their consistency with existing background theory *by default*. Moreover, our approach is scalable; it runs in polynomial time with respect to the number of symbolic variables and axioms (when the degree of the polynomial certificates we search over is bounded; see Section 2.2) with a complete and consistent background theory.

We believe our approach could be a first step towards discovering new laws of the universe which involve higher degree polynomials and are impractical for scientists to discover without the aid of modern optimization solvers and high-performance computing environments. Further, our approach is potentially useful for reconciling mutually inconsistent axioms. Indeed, if a system of scientific laws is mutually inconsistent (in the sense that no point satisfies all laws simultaneously), our polynomial optimization problem offers a formal proof of its inconsistency.

1.1 Literature Review

We propose an approach to scientific discovery, which we term **AI-Hilbert**, that uses polynomial optimization to obtain scientific formulae derivable from background theory axioms and consistent with experimental data. This differs from existing works on scientific discovery that express prior knowledge as constraints on the functional form of a learned model (e.g., shape constraints such as monotonicity [38]). Indeed, shape-constrained approaches to scientific discovery have been proposed [71, 72, 43, 38, 19], while discovering scientific laws that are simultaneously derivable from prior knowledge expressed as polynomials and experimental data is, to our knowledge, a new approach.

AI-Hilbert builds upon two areas of the optimization and discovery literature typically considered in isolation: sum-of-squares optimization techniques for solving polynomial optimization problems, and data-driven techniques for symbolic discovery. We now review the relevant literature.

Sum-of-Squares Optimization: Sum-of-squares optimization has been an important component of global optimization methods since the seminal work of Parrilo [87] (see also Lasserre [74]), which combines two key observations. First, sum-of-squares decompositions of multivariate polynomials can be computed via semidefinite optimization, so optimizing over sum-of-squares polynomials is no harder than performing semidefinite optimization. Second, owing to a fundamental result from real algebraic geometry, the Positivstellensatz [70, 103, 90], polynomials of bounded degree defined on basic semialgebraic sets can be certified as non-negative over these sets by representing them




Figure 1: **Comparison of scientific discovery paradigms.** Traditional scientific discovery formulates hypotheses using existing theory and observed phenomena. These hypotheses are validated and tested using data. In contrast, machine learning methods rely on large datasets to identify patterns. AI-Descartes [32] proposes an inversion of the conventional scientific discovery paradigm. It generates hypotheses from observed data and validates them against known theory. However, in AI-Descartes, theory and data remain disjoint and do not mutually enhance one another. In contrast, our work, AI Hilbert, combines data and theory to formulate hypotheses. Unlike conventional methods, insights in data and knowledge embedded in the theory collaboratively reduce the search space. These two components complement each other: theory compensates for noisy or sparse data, while data compensates for inconsistent or incomplete theory. Note that blue denotes components associated with data, purple denotes components linked to theory, and dashed lines represent iterative processes.

as systems of sum-of-squares polynomials (see Section 1.3). Consequently, optimizing over a real polynomial system is equivalent to solving a (larger) sum-of-squares optimization problem, and thus a tractable convex problem. These observations have allowed an entire field of optimization to blossom; see Blekherman et al. [23], Hall [59] for reviews. However, to our knowledge, no works have proposed using sum-of-squares optimization to discover scientific formulae. The closest works are Clegg et al. [31], who propose using Gröbner bases to design proofs of unsatisfiability, Curmei and Hall [38], who propose a sum-of-squares approach to fitting a polynomial to data under very general constraints on the functional form of the polynomial, e.g., non-negativity of the derivative over a box, Ahmadi and El Khadir [3], who propose learning the behavior of noisy dynamical systems via semialgebraic techniques, and Fawzi et al. [45], who propose learning proofs of optimality of stable set problems by combining reinforcement learning with the Positivstellensatz. However, determining whether polynomial optimization is useful for scientific discovery is, to our knowledge, open.

Data-Driven Approaches to Scientific Discovery: The availability of large amounts of scientific data generated and collected over the past few decades has spurred increasing interest in data-driven methods for scientific discovery that aim to identify symbolic equations that accurately explain high-dimensional datasets. Bongard and Lipson [25] and Schmidt and Lipson [98] proposed using heuristics and genetic programming to discover scientifically meaningful formulae, and implement their approach in the *Eureqa* software system [42]. Other proposed approaches are based on mixed-integer global optimization [10, 35], sparse regression [26, 96, 19], Cylindrical Algebraic Decomposition [50], neural networks [65, 73], and Bayesian Markov Chain Monte Carlo approaches [56]. See [68, 13] for reviews of data-driven scientific discovery in fundamental physics and chemistry.

Data-driven approaches have been shown by several authors to perform well in highly overdetermined settings with limited amounts of noise. For instance, Udrescu et al. [107, 106] proposed a method called *AI-Feynman*, which combines neural networks with physics-based techniques to discover symbolic formulae. Moreover, they constructed a benchmark dataset of 100 scientific laws derived from Richard Feynman’s lecture notes [47], with 100,000 noiseless experimental observations of each scientific law, and demonstrated that while the *Eureqa* system could recover an already impressive 71/100 instances from the data, their approach could recover all one hundred; see Cornelio et al. [32] for a review of scientific discovery systems.

Unfortunately, data-driven approaches to scientific discovery have at least three significant drawbacks. First, they are not data efficient [49] and only reliably recover scientific formulae in overdetermined settings with orders of magnitude more data than a human would likely need to make the same discoveries. Indeed, Matsubara et al. [83] recently argued that the sampling regime used by *AI-Feynman* is unrealistic, because it samples values far from those observable in the real world. Moreover, Cornelio et al. [32] recently benchmarked *AI-Feynman* on 81 of the 100 laws, but with 10 (rather than 100,000) observations per law, and where each experimental observation is contaminated with a small amount of noise. In this limited data setting, Cornelio et al. [32] found that *AI-Feynman* recovered 40 of the 81 laws considered, whereas [32] were able to recover 49/81 laws using their symbolic regression solver. This performance degradation is a significant issue in practice because scientific data is typically expensive to obtain and scarce and noisy. Second, purely data-driven methods are agnostic to important background information, such as existing literature, that valid scientific formulae should be consistent with unless there is extraordinary experimental evidence that the literature is incorrect. This implies that data-driven methods search over a larger space of laws than is necessary, require more data than a human would need to derive a valid law, and frequently propose laws that are not scientifically meaningful. Third, data-driven methods typically do not provide interpretable explanations for why their discoveries are valid [c.f. 95], which makes diagnosing whether their discoveries are consistent with existing theory challenging.

To account for background theory in scientific discovery, Cornelio et al. [32] recently proposed an approach called *AI-Descartes*, which iteratively generates plausible scientific formulae using a mixed-integer nonlinear symbolic regression solver [see also 10], and tests whether these formulae are derivable from the background knowledge. If they are not, the method provides a set of reasoning-based measures to compute how distant the formulae induced from the data are from the background theory but is unable to recover the correct formulae. This is because their approach induces potential scientific laws from data and subsequently tests the hypothesis against the background theory, rather than learning from axioms and data simultaneously.

1.2 Contributions and Structure

We propose a new paradigm for scientific discovery that derives polynomial laws simultaneously consistent with experimental data and a body of background knowledge expressible as polynomial equalities and inequalities. We term our approach `AI-Hilbert`, inspired by the work of David Hilbert, one of the first mathematicians to investigate the relationship between sum-of-squares and non-negative polynomials [62].

Our approach automatically provides an axiomatic derivation of the correctness of a discovered scientific law, conditional on the correctness of our background theory. Moreover, in instances with inconsistent background theory, our approach can successfully identify the sources of inconsistency by performing best subset selection to determine the axioms that best explain the data. This is notably different from current data-driven approaches to scientific discovery, which often generate spurious laws in limited data settings and fail to differentiate between valid and invalid discoveries or provide explanations of their derivations. We illustrate our approach by axiomatically deriving some of the most frequently cited natural laws in the scientific literature, including Kepler’s Third Law and Einstein’s Relativistic Time Dilation Law, among other scientific discoveries.

A second contribution of our approach is that it permits fine-grained control of the tractability of the scientific discovery process, by bounding the degree of the coefficients in the Positivstellensatz certificates that are searched over (see Section 1.3, for a formal statement of the Positivstellensatz). This differs from prior works on automated scientific discovery, which offers more limited control over their time complexity. For instance, in the special case of scientific discovery with a complete body of background theory and no experimental data, to our knowledge, the only current alternative to our approach is symbolic regression [see, e.g., 35], which requires genetic programming or mixed-integer nonlinear programming techniques that are not guaranteed to run in polynomial time. On the other hand, our approach searches for polynomial certificates of a bounded degree via a fixed level of the sum-of-squares hierarchy [74, 87], which can be searched over in polynomial time [84, 91].

To contrast our approach with existing approaches to scientific discovery, Figure 2 depicts a stylized version of the scientific method. In this version, new laws of nature are proposed from background theory (which may be written down by humans, automatically extracted from existing literature, or even generated using AI [67]) and experimental data, using classical discovery techniques, data-driven techniques, or `AI-Hilbert`. Observe that data-driven discoveries may be inconsistent with background theory, and discoveries via classical methods may not be consistent with relevant data sources, while discoveries made via `AI-Hilbert` are consistent with background theory and relevant data sources. This suggests that `AI-Hilbert` could be a first step toward discovery frameworks that are less likely to make false discoveries. Moreover, as mentioned in the introduction, `AI-Hilbert` uses background theory to restrict the effective dimension of the set of possible scientific laws, and, therefore, likely requires less data to make scientific discoveries than purely data-driven approaches.

1.3 Background and Notation

The notation is mostly standard to the semidefinite and polynomial optimization literature. We let non-boldface characters such as b denote scalars, lowercase bold-faced characters such as \mathbf{x} denote vectors, uppercase bold-faced characters such as \mathbf{A} denote matrices, and calligraphic uppercase characters such as \mathcal{Z} denote sets. We let $[n]$ denote the set of indices $\{1, \dots, n\}$. We let \mathbf{e} denote the vector of ones, $\mathbf{0}$ denote the vector of all zeros, and \mathbb{I} denote the identity matrix. We let $\|\mathbf{x}\|_p$ denote the p -norm of a vector \mathbf{x} for $p \geq 1$. We let \mathbb{R} denote the real numbers, S^n denote the cone of $n \times n$ symmetric matrices, and S_+^n denote the cone of $n \times n$ positive semidefinite matrices.

We also use some notations specific to the sum-of-squares (SOS) optimization literature; see [34] for an introduction to computational algebraic geometry and [23] for a general theory of sum-of-squares and convex algebraic optimization. Specifically, we let $\mathbb{R}[\mathbf{x}]_{n,2d}$ denote the ring of real polynomials in the n -tuple of variables $\mathbf{x} \in \mathbb{R}^n$ of degree $2d$, $P_{n,2d} := \{p \in \mathbb{R}[\mathbf{x}]_{n,2d} : p(\mathbf{x}) \geq 0 \quad \forall \mathbf{x} \in \mathbb{R}^n\}$ denote the convex cone of non-negative polynomials in n variables of degree $2d$, and

$$\Sigma[\mathbf{x}]_{n,2d} := \left\{ p(\mathbf{x}) : \exists q_1, \dots, q_m \in \mathbb{R}[\mathbf{x}]_{n,d}, p(\mathbf{x}) = \sum_{i=1}^m q_i^2(\mathbf{x}) \right\}$$

denote the cone of sum-of-squares polynomials in n variables of degree $2d$, which can be optimized over via $\binom{n+d}{d}$ dimensional semidefinite matrices [c.f. 87] using interior point methods [84]. Note

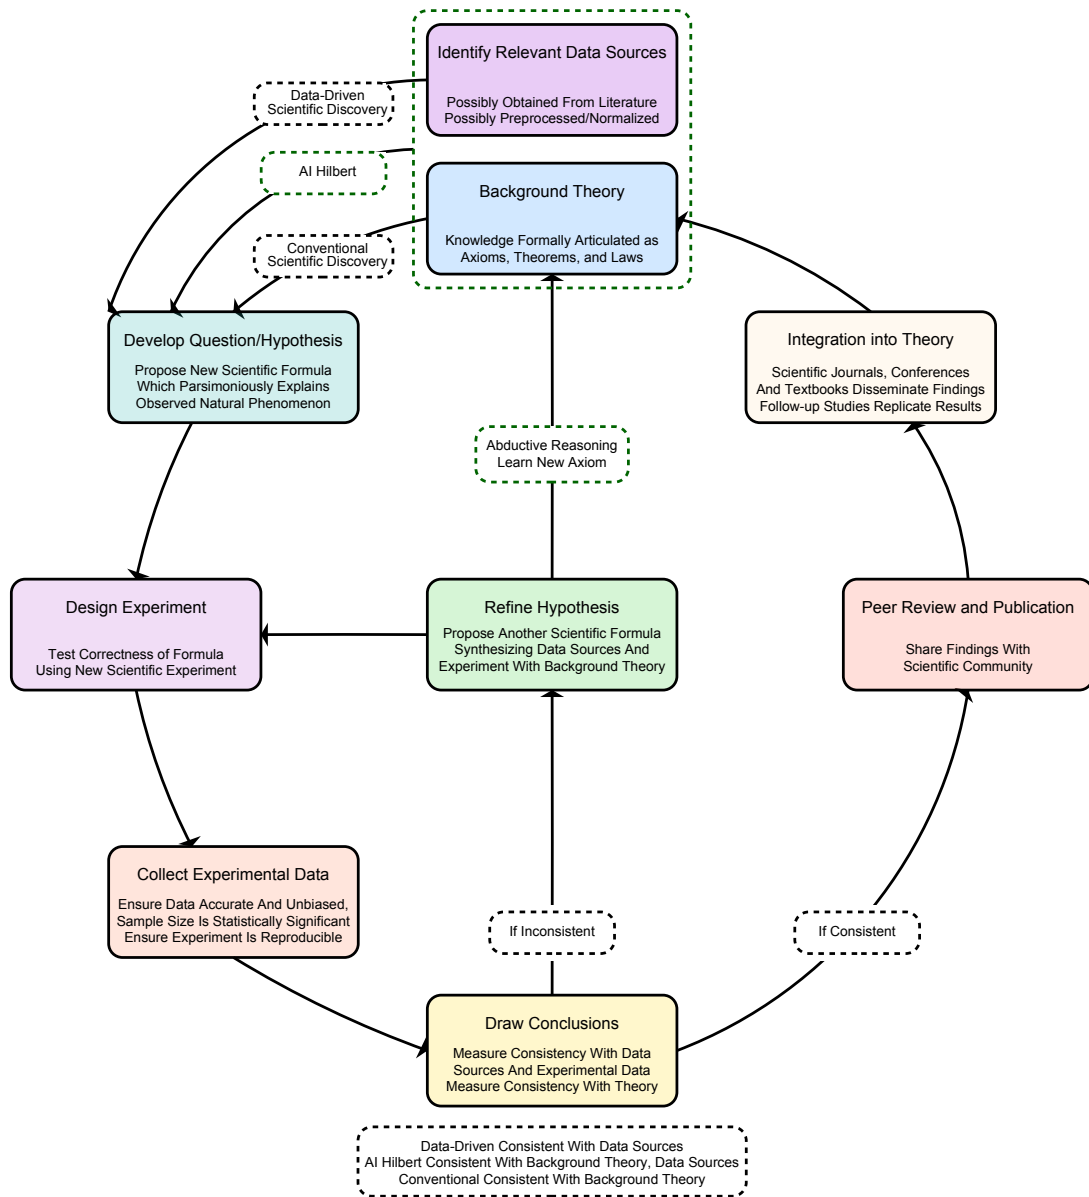


Figure 2: **The scientific method with scientific discoveries made via classical methods, data-driven methods, or AI-Hilbert.** AI-Hilbert proposes scientific laws consistent with a body of background theory formally articulated as polynomial equalities, inequalities, and relevant data sources. This likely allows scientific discoveries to be made using fewer data points than state-of-the-art approaches, and for missing scientific axioms to be deduced via abductive reasoning as part of the scientific discovery process. On the other hand, existing approaches to scientific discovery propose laws that may be inconsistent with either background theory or existing data sources.

that $\Sigma[\mathbf{x}]_{n,2d} \subseteq P_{n,2d}$, and the inclusion is strict unless $n \leq 2, 2d \leq 2$ or $n = 3, 2d = 4$ [61]. Nonetheless, $\Sigma[\mathbf{x}]_{n,2d}$ provides a high-quality approximation of $P_{n,2d}$, since each non-negative polynomial can be approximated (in the ℓ_1 norm of its coefficient vector) to any desired accuracy $\epsilon > 0$ by a sequence of sum-of-squares [75]. If the maximum degree d is unknown, we suppress the dependence on d in our notation.

To define a notion of distance between polynomials, we also use several functional norms. Let $\|\cdot\|_p$ denote the ℓ_p norm of a vector. Let $\boldsymbol{\mu} \in \mathbb{N}^n$ be the vector (μ_1, \dots, μ_n) and $\mathbf{x}^\boldsymbol{\mu}$ stand for the monomial $x_1^{\mu_1} \dots x_n^{\mu_n}$. Then, for a polynomial $q \in \mathbb{R}_{n,2d}[\mathbf{x}]$ with the decomposition $q(\mathbf{x}) = \sum_{\boldsymbol{\mu} \in \mathbb{N}^n: \|\boldsymbol{\mu}\|_1 \leq 2d} a_\boldsymbol{\mu} \mathbf{x}^\boldsymbol{\mu}$, we let the notation $\|q\|_p = \left(\sum_{\boldsymbol{\mu} \in \mathbb{N}^n: \|\boldsymbol{\mu}\|_1 \leq 2d} a_\boldsymbol{\mu}^p \right)^{1/p}$ denote the coefficient norm of the polynomial,

Finally, to derive new laws of nature from existing ones, we repeatedly invoke a fundamental result from real algebraic geometry called the Positivstellensatz [see, e.g., 103]. Various versions of the Positivstellensatz exist, with stronger versions holding under stronger assumptions [see 76, for a review], and any reasonable version being a viable candidate for our approach. For simplicity, we invoke a compact version due to [90], which holds under some relatively mild assumptions but nonetheless lends itself to relatively tractable optimization problems:

Theorem 1 (Putinar’s Positivstellensatz [90], see also Theorem 5.1 of [87]) *Consider the basic (semi)algebraic sets*

$$\mathcal{G} := \{\mathbf{x} \in \mathbb{R}^n : g_1(\mathbf{x}) \geq 0, \dots, g_k(\mathbf{x}) \geq 0\}, \quad (1)$$

$$\mathcal{H} := \{\mathbf{x} \in \mathbb{R}^n : h_1(\mathbf{x}) = 0, \dots, h_l(\mathbf{x}) = 0\}, \quad (2)$$

where $g_i, h_j \in \mathbb{R}[\mathbf{x}]_n$, and \mathcal{G} satisfies the Archimedean property² [see also 23, Chap. 6.4.4], i.e., there exists an $R > 0$ and $\alpha_0, \dots, \alpha_k \in \Sigma[\mathbf{x}]_n$ such that $R - \sum_{i=1}^k \alpha_i x_i^2 = \alpha_0(\mathbf{x}) + \sum_{i=1}^k \alpha_i(\mathbf{x}) g_i(\mathbf{x})$.

Then, for any $f \in \mathbb{R}[\mathbf{x}]_{n,2d}$, the implication

$$\mathbf{x} \in \mathcal{G} \cap \mathcal{H} \implies f(\mathbf{x}) \geq 0$$

holds if and only if there exist SOS polynomials $\alpha_0, \dots, \alpha_k \in \Sigma[\mathbf{x}]_{n,2d}$, and real polynomials $\beta_1, \dots, \beta_l \in \mathbb{R}[\mathbf{x}]_{n,2d}$ such that

$$f(\mathbf{x}) = \alpha_0(\mathbf{x}) + \sum_{i=1}^k \alpha_i(\mathbf{x}) g_i(\mathbf{x}) + \sum_{j=1}^l \beta_j(\mathbf{x}) h_j(\mathbf{x}). \quad (3)$$

Note that strict polynomial inequalities of the form $h_i(\mathbf{x}) > 0$ can be modeled by introducing an auxiliary variable τ and requiring that $h_i(\mathbf{x})\tau^2 - 1 = 0$, and thus our focus on non-strict inequalities in Theorem 1 is without loss of generality [see also 23].

Remarkably, the Positivstellensatz implies that if we set the degree of α_i s to be zero, then polynomial laws consistent with a set of equality-constrained polynomials can be searched over via linear optimization. Indeed, this subset of laws is sufficiently expressive that, as we demonstrate in our numerical results, it allows us to recover Kepler’s third law and Einstein’s dilation law axiomatically. Moreover, the set of polynomial natural laws consistent with polynomial (in)equalities can be searched via semidefinite or sum-of-squares optimization.

We close this section by remarking that one could develop an alternative version of the Positivstellensatz with only inequality constraints, by expressing each equality via two inequalities. However, this increases the number of decision variables in the optimization problems generated by the Positivstellensatz and solved in this paper, and thus decreases the tractability of these optimization problems; see also [23]. Accordingly, we treat equality and inequality constraints separately throughout.

²This assumption is stronger than the compactness assumption on \mathcal{G} found in the Positivstellensatz of [100], but is typically not restrictive in practice, as one could assume that $g_1(\mathbf{x}) = R - \sum_{i=1}^n x_i^2$ for some constant R . Moreover, it is arguably more tractable—it avoids the need to explicitly consider products of the form $g_i g_j$ in the decomposition, although we may require SOS polynomials of a higher degree to generate a valid certificate.

1.4 Structure

The rest of the paper is organized as follows: in Section 2, we describe AI-Hilbert, the scientific discovery system proposed in this paper. In Section 3, we argue that it presents an exciting new approach to scientific discovery, by demonstrating that it can rediscover the Hagen-Poiseuille Equation, Einstein’s Relativistic Time Dilation Law, Kepler’s Third Law, the Radiated Gravitational Wave Power Equation, and the Bell Inequalities (additionally, a detailed comparison of AI-Hilbert with state-of-the-art approaches in Appendix A and examples of scientific discovery via AI-Hilbert in Appendix B). In Section 4, we summarize our conclusions and discuss the limitations of and opportunities arising from this work.

2 Discovering Scientific Formulae Via Polynomial Optimization

In this section, we formally introduce AI-Hilbert. First, in Section 2.1, we provide an overview of our approach. Next, in Section 2.2, we formalize our approach as a polynomial optimization problem. Further, in Section 2.3, we define a new notion of the distance between a polynomial and a (possibly inconsistent or incomplete) set of polynomial background knowledge axioms. Subsequently, in Section 2.4, we discuss the role of background theory on the amount of data required to discover scientific laws, by revisiting some well-studied examples from the machine learning and statistics literature. Finally, in Section 2.5, we specialize our approach to problem settings where a scientist has access to a complete background theory and no experimental data.

2.1 Method Overview

Our scientific discovery method (AI-Hilbert) aims to discover an unknown polynomial formula $q(\cdot) \in \mathbb{R}[x]$ which describes a physical phenomenon, and is both consistent with a background theory of polynomial equalities and inequalities \mathcal{B} (a set of axioms) and a collection of experimental data \mathcal{D} (defined below). We provide a high-level overview of AI-Hilbert in Figure 3 and summarize our procedure in Algorithm 1. The inputs to AI-Hilbert are a four-tuple $(\mathcal{B}, \mathcal{D}, \mathcal{C}(\Lambda), d^c)$, where:

- \mathcal{B} denotes the relevant background theory, expressed as a collection of axioms, in the scientific discovery setting, i.e., the polynomial laws relevant for discovering q . It is the union of the inequalities $\{g_1(\mathbf{x}) \geq 0, \dots, g_k(\mathbf{x}) \geq 0\}$ defining \mathcal{G} and the equalities $\{h_1(\mathbf{x}) = 0, \dots, h_l(\mathbf{x}) = 0\}$ defining \mathcal{H} – where \mathcal{G} and \mathcal{H} are as in (1) and (2), respectively. \mathcal{B} is defined over n variables x_1, \dots, x_n . However, only t of these n variables can be measured and are directly relevant for explaining the observed phenomenon. In particular, we let x_1 denote the dependent variable. The




Figure 3: **Schematic illustration of AI Hilbert and its components.** Using background knowledge encoded as multivariate polynomials, experimental data, and hyperparameters (e.g., a sparsity constraint on the background theory) to control our model’s complexity, we formulate scientific discovery as a polynomial optimization problem, reformulate it as a semidefinite optimization problem, and solve it to obtain both a symbolic model and its formal derivation. Dashed boxes correspond to optional components. An example is introducing an incorrect candidate formula as a new axiom in the background theory.

remaining $n - t$ variables appear in the background theory but are not directly observable³. The background theory \mathcal{B} is defined as *complete* if it contains all the axioms necessary to formally prove the target formula, *incomplete* otherwise. Moreover, \mathcal{B} is called *inconsistent* if it contains axioms that contradict each other, *consistent* otherwise (we define these terms more rigorously in Section 2.3). A special case of inconsistency is when a formula that incorrectly describes the studied phenomenon is added to a consistent background theory (see Section 3.6).

- $\mathcal{D} := \{\bar{x}_i\}_{i \in [m]}$ denotes a collection of data points, or measurements of an observed physical phenomenon, which may be contaminated by noise, e.g., from measurement error. We assume that $\bar{x}_i \in \mathbb{R}^n$ and $\bar{x}_{i,j} = 0$ for $j \geq t + 1$, i.e., the value of $\bar{x}_{i,j}$ – the j th entry of \bar{x}_i – is set to zero for all variables j that cannot or should not be measured.
- \mathcal{C} denotes a set of constraints and bounds which depend on a set of hyper-parameters Λ . Specifically, we consider a global bound on the degree of the polynomial q ; a vector \mathbf{d} restricting individual variable degrees in q ; a hyperparameter λ that models the fidelity to background theory and data; and a bound over the number of axioms that should be included in a formula derivation.
- $d^c(\cdot, \mathcal{G} \cap \mathcal{H})$ denotes a distance function that defines the distance from an arbitrary polynomial to the background theory. We formally define d^c in Section 2.3.

Algorithm 1 provides a high-level description of AI-Hilbert. The procedure first combines the background theory \mathcal{B} and data \mathcal{D} to generate a polynomial optimization problem $\mathbb{P}r$ which targets a specific concept identified by a dependent - or target - variable, included in the set of observable entities that can be measured in the environment (x_1, \dots, x_t) . This is achieved by leveraging the distance d^c (formally defined in Section 2.3) and integrating the bounds and constraints \mathcal{C} (with their hyperparameters Λ) via the `PolyJuMP.jl` Julia package [82]. This corresponds to the `Formulate` step of Algorithm 1, which we formalize in Section 2.2.

Algorithm 1 AI Hilbert for Scientific Discovery

Input: $(\mathcal{B}, \mathcal{D}, \mathcal{C}(\Lambda), d^c)$
1: $\mathbb{P}r \leftarrow \text{Formulate}(\mathcal{B}, \mathcal{D}, \mathcal{C}(\Lambda), d^c)$
2: $\mathbb{P}r^{\text{sd}} \leftarrow \text{Reduce}(\mathbb{P}r)$
3: $q(\mathbf{x}) \leftarrow \text{Solve}(\mathbb{P}r^{\text{sd}})$
Output: $q(\mathbf{x}) = 0$
Output: α, β

AI-Hilbert then reformulates the problem $\mathbb{P}r$ as a semidefinite (or linear if no inequalities are present in the background theory) optimization problem $\mathbb{P}r^{\text{sd}}$, by leveraging standard techniques from sum-of-squares optimization that are now integrated within the `SumOfSquares.jl` and `PolyJuMP.jl` Julia packages, as discussed in Section 1.3. This corresponds to the `Reduce` step of Algorithm 1.

Next, AI-Hilbert solves $\mathbb{P}r^{\text{sd}}$ using a mixed-integer conic optimization solver such as Gurobi [1] or Mosek [6]. This corresponds to the `Solve` step of Algorithm 1.

AI-Hilbert then outputs a candidate formula of the form $q(\mathbf{x}) = 0$ where the only monomials with nonzero coefficients are those that only contain the variables x_1, \dots, x_t , the (independent and dependent) variables that are observed in the environment. The background theory may contain additional variables x_{t+1}, \dots, x_n that are not observed in the environment and that will not appear in the derived law. This is because the axioms in the background theory are not constraints on the functional form of the target polynomial, but rather general scientific laws describing the environment, often not including any of the quantities/variables observed in the data.

Finally, AI-Hilbert returns multipliers $\{\alpha_i\}_{i=1}^k, \{\beta_j\}_{j=1}^l$ such that

$$q(\mathbf{x}) = \alpha_0(\mathbf{x}) + \sum_{i=1}^k \alpha_i(\mathbf{x})g_i(\mathbf{x}) + \sum_{j=1}^l \beta_j(\mathbf{x})h_j(\mathbf{x})$$

³One might also exclude measurements for certain variables x_j if they are trivially connected to other measurable variables. For instance, the period of rotation can trivially be obtained from the frequency of rotation, and it is better to avoid having both in the list of measured variables in our system. We treat known quantities such as c - the speed of light - as measurable entities.

if $d^c(q, \mathcal{G} \cap \mathcal{H}) = 0$, which is a certificate of the fact that q is derivable from the background theory. If $d^c > 0$, then `AI-Hilbert` returns a certificate that q is approximately derivable from the background theory, and q is approximately equal to $\alpha_0(\mathbf{x}) + \sum_{i=1}^k \alpha_i(\mathbf{x})g_i(\mathbf{x}) + \sum_{j=1}^l \beta_j(\mathbf{x})h_j(\mathbf{x})$.

2.2 Overall Problem Setting

`AI-Hilbert` aims to discover an unknown polynomial model $q(\mathbf{x}) = 0$, which contains one or more dependent variables raised to some power within the expression (to avoid the trivial solution $q \equiv 0$), is approximately consistent with our axioms \mathcal{G} and \mathcal{H} —meaning d^c is small, and explains our experimental data well, meaning $q(\bar{\mathbf{x}}_i)$ is small for each data point i , and is of low complexity.

Let x_1, \dots, x_t denote the measurable variables, and let x_1 denote the dependent variable which we would like to ensure appears in our scientific law. Let $\Omega = \{\boldsymbol{\mu} \in \mathbb{N}^n : \|\boldsymbol{\mu}\|_1 \leq 2d\}$. Let the discovered polynomial expression be

$$q(x) = \sum_{\boldsymbol{\mu} \in \Omega} a_{\boldsymbol{\mu}} x^{\boldsymbol{\mu}},$$

where $2d$ is a bound on the maximum allowable degree of q . We formulate the following polynomial optimization problem:

$$\begin{aligned} \min_{q \in \mathbb{R}_{n, 2d}} \quad & \sum_{\bar{\mathbf{x}}_i \in \mathcal{D}} |q(\bar{\mathbf{x}}_i)| + \lambda \cdot d^c(q, \mathcal{G} \cap \mathcal{H}) & (4) \\ \text{s.t.} \quad & \sum_{\boldsymbol{\mu} \in \Omega: \mu_1 \geq 1} a_{\boldsymbol{\mu}} = 1, \\ & a_{\boldsymbol{\mu}} = 0 \quad \forall \boldsymbol{\mu} \in \Omega : \sum_{j=t+1}^n \mu_j \geq 1, \end{aligned}$$

where d^c , the distance between q and the background theory, is the optimal value of an inner minimization problem we define in Section 2.3, $\lambda > 0$ is a hyperparameter that balances the relative importance of model fidelity to the data against model fidelity to a set of axioms, the first constraint ensures that x_1 , our dependent variable, appears in q , the second constraint ensures that we do not include any unmeasured variables. In certain problem settings, we constrain $d^c = 0$, rather than penalizing the size of d^c in the objective.

Note that the formulation of the first constraint controls the complexity of the scientific discovery problem via the degree of the Positivstellensatz certificate: a smaller bound on the allowable degree in the certificate yields a more tractable optimization problem but a less expressive family of certificates to search over, which ultimately entails a trade-off that needs to be made by the user. Indeed, this trade-off has been formally characterized by Lasserre [75], who showed that every non-negative polynomial is approximable to any desired accuracy by a sequence of sum-of-squares polynomials, with a trade-off between the degree of the SOS polynomial and the quality of the approximation.

After solving Problem (4), one of two possibilities occurs. Either the distance between q and our background information is 0, or the Positivstellensatz provides a non-zero polynomial

$$r(\mathbf{x}) := q(\mathbf{x}) - \alpha_0(\mathbf{x}) - \sum_{i=1}^k \alpha_i(\mathbf{x})g_i(\mathbf{x}) - \sum_{j=1}^l \beta_j(\mathbf{x})h_j(\mathbf{x})$$

which defines the discrepancy between our derived physical law and its projection onto our background information. In this sense, solving Problem (4) also provides information about the inverse problem of identifying a complete set of axioms that explain q . In either case, it follows from the Positivstellensatz (Theorem 1) that solving Problem (4) for different hyperparameter values and different bounds on the degree of q eventually yields polynomials that explain the experimental data well and are approximately derivable from background theory.

We close this section with two remarks on the generality and complexity of `AI-Hilbert`.

Implicit and Explicit Symbolic Discovery: Most prior work [e.g., 38, 32, 99] aims to identify an unknown symbolic model $f \in \mathbb{R}[x]_{n,2d}$ of the form $y_i = f(\mathbf{x}_i)$ for a set of independent variables of interest $\mathbf{x}_i \in \mathbb{R}^n$ and a dependent variable $y_i \in \mathbb{R}$, while AI-Hilbert searches for an implicit polynomial function q which links the dependent and independent variables. We do this for two reasons. First, many scientific formulae of practical interest admit implicit representations as polynomials, but explicit representations of the dependent variable as a polynomial function of the independent variables are not possible [c.f. 2]. For instance, Kepler’s third law of planetary motion has this property (see Section 3.6). Second, as proven by Artin [9] to partially resolve Hilbert’s 17th problem [c.f. 61], an arbitrary non-negative polynomial can be represented as a sum of squares of rational functions. Therefore, by multiplying by the denominator in Artin’s representation [9], implicit representations of natural laws become a viable and computationally affordable search space.

We remark that the implicit representation of scientific laws as polynomials where $q(\mathbf{x}) = 0$ introduces some degeneracy in the set of optimal polynomials derivable from (4), particularly in the presence of a correct yet overdetermined background theory. For instance, in the derivation of Kepler’s Law of Planetary Motion in Section 3.6, we eventually derive the polynomial $m_1 m_2 G p^2 = m_1 d_1 d_2^2 + m_2 d_1^2 d_2 + 2 m_2 d_1 d_2^2$. Since we have the axiom that $m_1 d_1 = m_2 d_2$, we could instead derive the (equivalent) formula $m_1 m_2 G p^2 = (m_1 + m_2) d_1 d_2 (d_1 + d_2)$. To partly break this degeneracy, we propose to either constrain the degree of the proof certificate and gradually increase it (as is done in (4)) or, (equivalently in a Lagrangian sense) include a term modeling the complexity of our derived polynomial (e.g., $\|q\|_1$, the L_1 -coefficient norm of q) in the objective.

Complexity of Scientific Discovery: Observe that, if the degree of our new scientific law q is fixed and the degree of the polynomial multipliers in the definition in d^c is also fixed, then Problem (4) can be solved in polynomial time⁴ with a consistent set of axioms (resp. nondeterministic polynomial time with an inconsistent set of axioms). This is because solving Problem (4) with a fixed degree and a consistent set of axioms corresponds to solving a semidefinite optimization problem of a polynomial size, which can be solved in polynomial time (assuming that a constraint qualification such as Slater’s condition holds) [84]. Moreover, although solving Problem (4) with a fixed degree and an inconsistent set of axioms corresponds to solving a mixed-integer semidefinite optimization problem, which is NP-hard, recent evidence [40] shows that integer optimization problems can be solved in polynomial time with high probability. This suggests that Problem (4) may also be solvable in polynomial time with high probability. However, if the degree of q is unbounded then, to the best of our knowledge, no existing algorithm solves Problem (4) in polynomial time. This explains why searching for scientific laws of a fixed degree and iteratively increasing the degree of the polynomial laws searched over, in accordance with Occam’s Razor, is a key aspect of our approach.

2.3 Distance to Background Theory and Model Complexity

Scientists often start with experimental measurements and a set of polynomial equalities and inequalities (axioms) which they believe to be true. From these axioms and measurements, they aim to deduce a new law, explaining their data, which includes one or more dependent variables and excludes certain variables. The simplest case of scientific discovery involves a consistent and correct set of axioms that fully characterize the problem. In this case, the Positivstellensatz (Theorem 1) facilitates the discovery of new scientific laws via deductive reasoning, without using any experimental data, as we argue in Section 2.5. Indeed, under an Archimedean assumption, the set of all valid scientific laws corresponds precisely to the preprime (see [34] for a definition) generated by our axioms [90], and searching for the simplest polynomial version of a law which features a given dependent variable corresponds to solving an easy linear or semidefinite feasibility problem.

It is not uncommon to have a set of axioms that is inconsistent (meaning that there are no values of $\mathbf{x} \in \mathbb{R}^n$ that satisfy all laws simultaneously), or incomplete (meaning the axioms do not ‘span’ the space of all derivable polynomials; we provide a formal definition later in this section). Therefore, we require a notion of a distance between a body of background theory (which, in our case, consists

⁴Under the real number complexity model, and under the bit number complexity model under some mild regularity conditions on the semidefinite optimization problems that arise from our sum-of-squares optimization problems. Note that, under the bit complexity model, semidefinite optimization problems cannot always be solved in polynomial time due to the existence of ill-behaved semidefinite problems where all feasible solutions are of doubly exponential size. We refer to Ramana [91] or Laurent and Rendl [77] for a complete characterization of the complexity of semidefinite optimization.

of a set of polynomial equalities and inequalities) and a polynomial. We now establish this definition, treating the inconsistent and incomplete cases separately; note that the inconsistent and incomplete case may be treated via the inconsistent case. We remark that [23, 113] propose related notions of the distance between (a) a point and a variety defined by a set of equality constraints, and (b) the distance between two semialgebraic sets via their Hausdorff distance. However, to our knowledge, the distance metrics in this paper have not previously been proposed.

Incomplete Case: Suppose \mathcal{B} is a background theory (consisting of equalities and inequalities in \mathcal{G} and \mathcal{H}), where \mathcal{G} satisfies the previously defined Archimedean property (see Theorem 1) with constant R , and the axioms are not inconsistent, meaning that $\mathcal{G} \cap \mathcal{H} \neq \emptyset$. Then, a natural notion of distance is the ℓ_2 coefficient distance d^c between q and $\mathcal{G} \cap \mathcal{H}$, which is given by:

$$d^c(q, \mathcal{G} \cap \mathcal{H}) := \min_{\substack{\alpha_0, \dots, \alpha_k \in \Sigma_{n, 2d}[\mathbf{x}], \\ \beta_1, \dots, \beta_l \in \mathbb{R}_{n, 2d}}} \left\| q - \alpha_0 - \sum_{i=1}^k \alpha_i g_i - \sum_{j=1}^l \beta_j h_j \right\|_2.$$

It follows directly from Putinar’s Positivstellensatz that $d(q, \mathcal{G} \cap \mathcal{H}) = 0$ if and only if q is derivable from \mathcal{B} . We remark that this distance has a geometric interpretation as the distance between a polynomial q and its projection onto the algebraic variety generated by $\mathcal{G} \cap \mathcal{H}$. Moreover, by norm equivalence, this is equivalent to the Hausdorff distance [113] between q and $\mathcal{G} \cap \mathcal{H}$.

With the above definition of d^c , and the fact that $\mathcal{G} \cap \mathcal{H} \neq \emptyset$, we say that $\mathcal{G} \cap \mathcal{H}$ is an incomplete set of axioms if there does not exist a polynomial p with a non-zero coefficient on at least one monomial involving a dependent variable, such that $d^c(q, \mathcal{G} \cap \mathcal{H}) = 0$.

Inconsistent Case: Suppose \mathcal{B} is an inconsistent background theory i.e. $\mathcal{G} \cap \mathcal{H} = \emptyset$. Then, a natural approach to scientific discovery is to assume that a subset of the axioms are valid laws, while the remaining axioms are scientifically invalid (or invalid in a specific context, e.g., micro vs. macro-scale). In line with the sparse regression literature [c.f. 20] and related work on discovering nonlinear dynamics [19, 80], we assume that scientific discoveries can be made using at most k correct scientific laws and define the distance between the scientific law and the problem data via a best subset selection problem. Specifically, we introduce binary variables z_i and y_j to denote whether the i th and j -th laws are consistent, and require that $\alpha_i = 0$ if $z_i = 0$ and $\beta_j = 0$ if $y_j = 0$ and $\sum_i z_i + \sum_j y_j \leq \tau$ for a sparsity budget τ . Furthermore, we allow a non-zero ℓ_2 distance between the scientific law f and the reduced background theory, but penalize this distance in the objective. This gives the following notion of distance between a scientific law q and the space $\mathcal{G} \cap \mathcal{H}$:

$$\begin{aligned} d^c(q, \mathcal{G} \cap \mathcal{H}) := \min & \left\| q - \alpha_0 - \sum_{i=1}^k \alpha_i g_i - \sum_{j=1}^l \beta_j h_j \right\|_2, \\ \text{s.t.} & \alpha_i = 0 \text{ if } z_i = 0, \forall i \in \{0, \dots, k\}, \\ & \beta_j = 0 \text{ if } y_j = 0, \forall j \in \{1, \dots, l\}, \\ & \sum_{i=0}^k z_i + \sum_{j=1}^l y_j \leq \tau, \\ & z_0 \dots z_k \in \{0, 1\}, y_1, \dots, y_l \in \{0, 1\}, \\ & \alpha_0, \dots, \alpha_k \in \Sigma_{n, 2d}[\mathbf{x}], \beta_1, \dots, \beta_l \in \mathbb{R}[\mathbf{x}]_{n, 2d}. \end{aligned}$$

It follows directly from the Positivstellensatz that $d = 0$ if and only if q can be derived from \mathcal{B} . If $\tau = m + l$, then we certainly have $d^c = 0$, since the overall system of polynomials is inconsistent and the sum-of-squares proof system can deduce that “ $-1 \geq 0$ ” from inconsistent proof systems, from which it can claim a distance of 0. However, by treating τ as a hyper-parameter and including the quality of the law on experimental data as part of the optimization problem (see Section 2.2), scientific discoveries can be made from inconsistent axioms by incentivizing solvers to set $z_i = 0$ for inconsistent axioms i . Alternatively, a practitioner may wish to explore the Pareto frontier of scientific discoveries that arise as we vary τ , to detect how large the set of correct background knowledge is. Provided there is a sufficiently high penalty cost on poorly explaining scientific data via the derived law, our optimization problem prefers a subset of correct axioms with a non-zero distance d^c to the derived polynomial over a set of inconsistent axioms which gives a distance $d^c = 0$.

2.4 Impact of Background Theory on Amount of Data Needed to Discover Scientific Laws

In the introduction, we suggested that providing a partially complete background theory expressible as polynomial equalities and inequalities may accelerate the scientific discovery process by decreasing the amount of data required to recover a scientific law with high probability. We now justify our claim in the introduction, by reviewing examples from the machine learning literature, which may be viewed as special cases of scientific discovery, where including relevant background theory decreases the number of data points required to recover a scientific law with high probability. The first two examples involve discovery settings where the ground truth is known to be sparse, and imposing a sparsity constraint on the discovered law reduces the amount of data required to recover the law with high probability. Note that these examples are due to [52, 8, 29]. However, the similarities between them strongly suggest that similar relations between the amount of background theory introduced and the amount of data required for scientific discovery also hold in other contexts.

Example 1 Sparse Linear Regression [52]

Consider a sparse high-dimensional regression model where we aim to recover a τ -sparse regression model $\beta^* \in \mathbb{R}^p$ given access to n noisy linear observations of the form $\mathbf{Y} = \mathbf{X}\beta^* + \mathbf{W} \in \mathbb{R}^n$, where $X_{i,j} \stackrel{\text{iid}}{\sim} \mathcal{N}(0,1)$ and $W_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ for some parameter $\sigma > 0$, it is known that β^* is a τ -sparse vector with binary coefficients, and $n \ll p$ with $p \rightarrow \infty$. Then, it is information-theoretically impossible to recover β if $n \leq \Theta\left(\frac{2\tau \log p}{\log(1 + \frac{2\tau}{\sigma^2})}\right)$ [111]. On the other hand, for $n \geq \Theta\left(\frac{2\tau \log p}{\log(1 + \frac{2\tau}{\sigma^2})}\right)$, the (unique) optimal solution of the polynomial optimization problem (expressible as a binary problem)

$$\min_{\beta \in \mathbb{R}^p} \|\mathbf{Y} - \mathbf{X}\beta\|_2^2 \text{ s.t. } \beta_i^2 = \beta_i \forall i \in [p], \sum_{i \in [p]} \beta_i = \tau$$

is such that

$$\frac{1}{\tau} \|\beta^* - \beta\|_0 \rightarrow 0$$

with high probability as $\tau \rightarrow \infty$. On the other hand, for any $\lambda > 0$, the optimal solution of the popular Lasso method

$$\min_{\beta \in \mathbb{R}^p} \|\mathbf{Y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_1$$

only recovers β^* with high probability when $n \geq \Theta(2\tau + \sigma^2) \log p$. We remind the reader that

$$\Theta(2\tau + \sigma^2) \log p > \Theta\left(\frac{2\tau \log p}{\log(1 + \frac{2\tau}{\sigma^2})}\right).$$

Example 2 Sparse Principal Component Analysis [8]

Consider a sparse principal component analysis setting where we aim to recover a τ -sparse binary vector $\mathbf{x}^* \in \mathbb{R}^n$ given an observed matrix $\mathbf{Y} = \frac{\lambda}{\tau} \mathbf{x}^* \mathbf{x}^{*\top} + \mathbf{W}$, where \mathbf{W} is a $\text{GOE}(n)$ matrix, i.e., is a symmetric matrix with on-diagonal entries taking values $W_{i,i} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \frac{2}{n})$ and off-diagonal entries taking values $W_{i,j} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \frac{1}{n})$, and $\lambda > 0$ is the signal-to-noise ratio. Let λ, τ possibly depend on n , and set $1 \ll \tau \ll n$ as $n \rightarrow \infty$. Then:

- Recovery of \mathbf{x}^* is information-theoretically impossible when $\lambda \ll \frac{\sqrt{\tau}}{\sqrt{n}}$, with high probability.
- The (mixed-integer representable) polynomial optimization problem

$$\max_{\mathbf{x} \in \mathbb{R}^n} \mathbf{x}^\top \mathbf{Y} \mathbf{x} \text{ s.t. } x_i^2 = x_i \forall i \in [n], \sum_{i \in [n]} x_i = \tau$$

achieves exact recovery with high probability when $\lambda \gg \frac{\sqrt{\tau}}{\sqrt{n}}$.

- The diagonal thresholding algorithm of [66] recovers \mathbf{x}^* with high probability if $\lambda \gg \frac{\tau}{\sqrt{n}}$.

- The vanilla PCA method, which disregards background theory encoded via a sparsity constraint and solves the polynomial optimization problem

$$\max_{\mathbf{x} \in \mathbb{R}^n} \mathbf{x}^\top \mathbf{Y} \mathbf{x} \text{ s.t. } \|\mathbf{x}\|_2^2 = \tau$$

fails to recover \mathbf{x}^* with high probability when $\lambda > 1$.

Example 3 *Low-Rank Matrix Completion [29]*

Consider a low-rank matrix completion setting where we aim to recover a fixed rank r $n \times n$ matrix \mathbf{A} given a uniform random sample of its entries $\Omega \subseteq [n] \times [n]$ of size m , where \mathbf{A} satisfies the mutual incoherence property of [29] with constant μ . Then:

- Recovery of \mathbf{A} is information-theoretically impossible when $m \leq \Theta(nr \log n)$, because there are infinitely many rank- r matrices that match all observed entries perfectly [29].
- The polynomial optimization problem [c.f. 17, 15]

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times n}, \mathbf{Y} \in \mathcal{S}^n} \sum_{(i,j) \in \Omega} (X_{i,j} - A_{i,j})^2 \text{ s.t. } \mathbf{Y} \mathbf{X} = \mathbf{X}, \mathbf{Y}^2 = \mathbf{Y}, \text{tr}(\mathbf{Y}) = r$$

achieves exact recovery with high probability provided $m \geq \Theta(nr \log n)$.

- The nuclear norm relaxation of [29] recovers \mathbf{A} with high probability provided $m \geq \Theta(n^{6/5} r \log n)$.
- The naive approach of disregarding the rank constraint and solving

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times n}} \sum_{(i,j) \in \Omega} (X_{i,j} - A_{i,j})^2,$$

which admits the solution $X_{i,j} = 0$ if $(i,j) \notin \Omega$, fails to recover \mathbf{A} with high probability provided $m < n^2$.

The above examples are admittedly more stylized than many scientific discovery settings that arise in practice. Nonetheless, they reveal that in certain circumstances, encoding relevant background provably reduces the amount of data required to recover a scientific law with high probability. This agrees with intuition: if we provide a complete background theory that can be manipulated to recover the scientific law, then, as discussed in the next section, we require no data to recover a scientific law. On the other hand, if we provide no background theory we may require a significant amount of data to recover a law. Therefore, providing relevant background theory that constrains the space of derivable scientific laws should decrease the amount of data needed to recover a scientific law with high confidence. This observation highlights the value of embedding relevant background theory within the scientific discovery process.

2.5 Discovering Scientific Laws From Background Theory Alone

Suppose that the background theory \mathcal{B} constitutes a complete set of axioms that fully describes our physical system. Then, any polynomial which contains our dependent variable x_1 and is derivable from our system of axioms is a valid physical law. Therefore, we need not even collect any experimental data, and we can solve the following feasibility problem to discover a valid law (let $\Omega = \{\boldsymbol{\mu} \in \mathbb{N}^n : \|\boldsymbol{\mu}\|_1 \leq 2d\}$):

$$\begin{aligned}
& \exists \quad q(x) = \sum_{\mu \in \Omega} a_{\mu} x^{\mu} & (5) \\
\text{s.t.} \quad & q(\mathbf{x}) = \alpha_0(\mathbf{x}) + \sum_{j=1}^k \alpha_j(\mathbf{x}) g_j(\mathbf{x}) + \sum_{j=1}^l \beta_j(\mathbf{x}) h_j(\mathbf{x}), \\
& \sum_{\mu \in \Omega: \mu_1 \geq 1} a_{\mu} = 1, \\
& a_{\mu} = 0 \quad \forall \mu \in \Omega : \sum_{j=t+1}^n \mu_j \geq 1, \\
& \alpha_i(\mathbf{x}) \in \Sigma[\mathbf{x}]_{n,2d}, \quad \beta_j(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]_{n,2d},
\end{aligned}$$

where the second and third constraints ensure that we include the dependent variable x_1 in our formula q and rule out the trivial solution $q = 0$, and exclude any solutions q which contain uninteresting symbolic variables respectively.

Note that if we do not have any inequality constraints in either problem, then we may eliminate α_i and obtain a linear optimization problem.

3 Experimental Validation

In this section, we showcase AI-Hilbert’s ability to recover valid scientific laws across various problem settings, including settings with incomplete or inconsistent axioms and experimental data.

The rest of this section is laid out as follows: First, we provide a detailed explanation of AI-Hilbert’s implementation in Section 3.1. Subsequently, we discuss the trade-off between data and theory made by AI-Hilbert in Section 3.2. Next, we validate AI-Hilbert on five selected problems that highlight its different capabilities, as discussed below. Our main findings are as follows:

1. In Sections 3.3–3.4, we demonstrate that AI-Hilbert successfully derives valid scientific laws solely from a complete and consistent background theory. Namely, by deriving the Hagen-Poiseuille Equation (Section 3.3) and deriving the Radiated Gravitational Wave Power Equation (Section 3.4). We further demonstrate this capability on a suite of test instances in Appendix B.
2. In Sections 3.5–3.6, we demonstrate that AI-Hilbert is capable of deriving a valid scientific law from an inconsistent yet complete background theory. We illustrate this capability through two scenarios: Firstly, when two axioms in the background theory contradict each other (Section 3.5), by using measurement data to discern the correct axiom. Secondly, where an incorrect candidate formula is incorporated into a complete and consistent background theory (Section 3.6). In both cases, we use real-life, noisy, experimental data.
3. In Section 3.7, we demonstrate that AI-Hilbert successfully derives a valid scientific law from an incomplete yet consistent background theory and experimental data. Specifically, we demonstrate that AI-Hilbert is data-efficient, in the sense that as the number of axioms supplied increases, less data is required to successfully derive a scientific law.
4. In Section 3.8, we showcase AI-Hilbert’s ability to handle and discover inequalities, namely the Bell Inequalities. This is notable because existing scientific discovery methods, to our knowledge, cannot handle inequalities.
5. We compare AI-Hilbert with some widely used methods from the literature on a suite of test instances in Appendix A. We demonstrate that AI-Hilbert outperforms these methods by rediscovering more of the scientific laws in this test bed. This occurs because AI-Hilbert integrates data and theory systematically via sum-of-squares optimization.

3.1 Implementation Details

We now illustrate the lower-level implementation of AI-Hilbert via a synthetic example where the axioms are assumed to be consistent and complete. Consider a semialgebraic system in two real

variables x and y which comprises the axioms:

$$x^2 + y^2 - 2 = 0, \quad (6)$$

$$y - x^3 = 0, \quad (7)$$

where we let $h_1(x, y) = 0$ and $h_2(x, y) = 0$ denote Equations (6)–(7), respectively. These axioms can be viewed as being true in a subdomain of \mathbb{R}^2 , i.e., when $|x| = 1$.

Let the set $S = \{(1, 1), (-1, -1)\}$ contain all points satisfying equations (6)–(7). If

$$q(x, y) = \beta_1(x, y)h_1(x, y) + \beta_2(x, y)h_2(x, y), \quad (8)$$

where β_1, β_2 are polynomials in x, y , then q is a polynomial that vanishes on S . On the other hand, to certify that a polynomial $q(x, y)$ vanishes on S , we search for a polynomial $q(x, y)$ such that $q(x, y) = g(x, y)p(x, y)$ and q satisfies the expression in (8). For example, $q(x, y) = x - y$ vanishes on S ; setting $\beta_1 = -\frac{1}{2}x, \beta_2 = -1$, we have

$$q = \beta_1 h_1 + \beta_2 h_2 = x - y + \frac{1}{2}(x^3 - xy^2) = (x - y)(1 + \frac{1}{2}x(x + y)). \quad (9)$$

Suppose we wish to find a polynomial function $q(x, y)$ that vanishes on S and explains the dataset

$$\bar{\mathbf{x}} = \begin{bmatrix} .5 & .5 \\ -2 & -2 \\ 3 & 3 \\ -3 & -3 \end{bmatrix}$$

where the columns are observed values of x and y , respectively, and each row represents a datapoint. Note that these data observations are noiseless observations from the polynomial $f(x, y) = x - y$, which we aim to recover. Then we search for q satisfying (8) that fits $\bar{\mathbf{x}}$.

We assume that β_1, β_2, q are unknown polynomials that are comprised of monomials of degree at most 3, i.e., the ten monomials in the vector

$$\text{mon} = (1, y, y^2, y^3, x, xy, xy^2, x^2, x^2y, x^3),$$

where $\text{mon}_1 = 1, \text{mon}_2 = y, \dots, \text{mon}_{10} = x^3$. Let

$$\beta_1(x, y) = \sum_{j=1}^{10} a_j \text{mon}_j, \quad \beta_2(x, y) = \sum_{j=1}^{10} b_j \text{mon}_j, \quad \text{and} \quad q(x, y) = \sum_{j=1}^{10} c_j \text{mon}_j. \quad (10)$$

Further, let \mathbf{v} denote the vector $(a_1, \dots, a_{10}, b_1, \dots, b_{10}, c_1, \dots, c_{10})$, $\bar{\mathbf{x}}_i$ denote the i th row of $\bar{\mathbf{x}}$, and $q(\bar{\mathbf{x}}_i)$ be the value of the polynomial q evaluated at the point $(\bar{x}, \bar{y}) = \bar{\mathbf{x}}_i$. Note that $q(\bar{\mathbf{x}}_i) = \sum_{j=1}^{10} c_j \text{mon}_j(\bar{\mathbf{x}}_i)$ is a linear function of the unknowns c_1, \dots, c_{10} .

If $\bar{\mathbf{x}}_i$ is a noiseless experimental observation, we should have $q(\bar{\mathbf{x}}_i) = 0$. Accordingly, we interpret any nonzero value of $|q(\bar{\mathbf{x}}_i)|$ as the error when q is evaluated at $\bar{\mathbf{x}}_i$, and aim to minimize this error when selecting q . Equation (8) implies the following linear equations in the variables a_i, b_i, c_i , which are obtained by equating the coefficients of the monomials in Equation (8):

$$c_1 = -2a_1 \quad [\text{coef. of } 1] \quad (11)$$

$$c_2 = -2a_2 + b_1 \quad [\text{coef. of } y] \quad (12)$$

$$c_3 = a_1 - 2a_3 + b_2 \quad [\text{coef. of } y^2] \quad (13)$$

\vdots

$$0 = a_3 + b_4 \quad [\text{coef. of } xy^3] \quad (14)$$

Let these constraints be denoted by $Av = \mathbf{0}$, and let m , the number of data points, be an integer between 1 and 4. Then, we solve the linear optimization problem

$$\begin{aligned}
\min \quad & 100 \left(\sum_{i=1}^m t_i \right) + \sum_{j=1}^{10} w_j & (15) \\
\text{s.t.} \quad & t_i \geq q(\bar{\mathbf{x}}_i) & i = 1, \dots, m \\
& t_i \geq -q(\bar{\mathbf{x}}_i) & i = 1, \dots, m \\
& w_j \geq c_j & j = 1, \dots, 10 \\
& w_j \geq -c_j & j = 1, \dots, 10 \\
& Av = \mathbf{0} \\
& \sum_{j=4}^{10} c_j = 1 \\
& v \in \mathbb{R}^{30}, t_i, w_j \geq 0 & i = 1, \dots, m, j = 1, \dots, 10.
\end{aligned}$$

In Problem (15), the first two constraints imply that $t_i \geq |q(\bar{\mathbf{x}}_i)|$; the third and fourth force $w_i \geq |c_i|$. The second-to-last constraint forces a monomial containing x to be present in q to avoid the trivial solution $q \equiv 0$. We sometimes use a different right-hand-side value to get non-fractional c_i values. Thus, the optimization problem above searches for a polynomial q that minimizes a weighted combination of the L_1 -error of q and the L_1 -coefficient norm of q . The latter term is a regularization term incentivizing a sparse q .

If we set $m = 1$ and solve (15), we obtain the (correct) function $q(x, y) = x - y$ in (15). In other words, a single data point from $\bar{\mathbf{x}}$ suffices to “recover” $x - y$ with the choice of objective in (15).

However, let LPF be the linear optimization problem obtained by dropping the second term in the objective in (15) (and not incentivizing sparsity). If we solve LPF with $m = 1$, we get $q' = 4x^3 - x^2 - y^2 - 4y + 2$ as a solution. This is not a multiple of $(x - y)$ but vanishes at the points in S and the first datapoint in \mathcal{D} . We need $m = 2$ (and use two datapoints) before we get

$$q' = y^3 + x^2y - 2xy^2 + 4x - 4y = (x - y)((x - y)^2 + 4).$$

This illustrates the role of sparsity and regularization in reducing the amount of data required to recover a scientific law. Note that $(x - y)^2 + 4$ is strictly positive for all real x, y , and therefore $q' = 0$ on S indicates that $x - y = 0$ on S .

Further, if we drop the constraints $Av = \mathbf{0}$ and allow q to be an arbitrary degree-3 polynomial in x, y (and not equal to $\beta_1 h_1 + \beta_2 h_2$), then we cannot recover a multiple of $x - y$ until we set $m = 4$ and use all datapoints in $\bar{\mathbf{x}}$. In the latter case, we get $q = x^3 - y^3$. This highlights the value of background theory in restricting the space of feasible scientific laws and reducing the amount of data required to recover a scientific law.

If the only measured variable is x , then we would try to eliminate y in the final formula. Multiplying (7) by $y + x^3$ and subtracting the result from (6), we get the expression $x^2 + x^6 - 2 = 0$. This satisfies the x -components of points in S . Sometimes, there may be a unique way of eliminating variables and getting a formula on the measured variables.

3.2 Trade off Between Data and Theory

There is a fundamental trade-off between the amount of background theory available and the amount of data needed for scientific discovery. Indeed, with a complete and consistent set of background theories, it is possible to perform scientific discovery without any experimental data via the Positivstellensatz (see Section 1.3). On the other hand, the purely data-driven approaches to scientific discovery reviewed in Section 1.1 often require many noiseless or low-noise experimental observations to successfully perform scientific discovery. More generally, as the amount of relevant background theory increases, the amount of experimental data required by AI-Hilbert to perform scientific discovery cannot increase, because increasing the amount of background theory decreases the effective VC-dimension of our scientific discovery problem [see, e.g., 60]. This can be seen in the machine learning examples discussed in Section 2.4, where imposing a sparsity constraint (i.e., providing a relevant axiom) reduces the number of data observations needed to discover a ground truth model; see also [79, 57] for an analysis of the asymptotic consistency of shape-constrained regression. We provide further evidence of this experimentally in Section 3.7.

The above observations can also be explained via real algebraic geometry: requiring that a scientific law is consistent with an axiom is equivalent to restricting the space of valid scientific laws to a subset of the space of discoverable scientific laws. As such, an axiom is equivalent to infinitely many data observations that penalize scientific laws outside a subspace but provide no information that discriminates between scientific laws within the subspace.

3.3 Deriving the Hagen-Poiseuille Equation

We consider the problem of deriving the velocity of laminar fluid flow through a circular pipe, from a simplified version of the Navier-Stokes equations, an assumption that the velocity can be modeled by a degree-two polynomial in the radius of the pipe, and a no-slip boundary condition. Letting $u(r)$ denote the velocity in the pipe where r is the distance from the center of the pipe, R denote the radius of the pipe, Δp denote the pressure differential throughout the pipe, L denote the length of the pipe, and μ denote the viscosity of the fluid, we have the following velocity profile for $r \in [0, R]$:

$$u(r) = \frac{-\Delta p}{4L\mu}(r^2 - R^2). \quad (16)$$

We now derive this law axiomatically, by assuming that the velocity profile can be modeled by a symmetric polynomial, and iteratively increasing the degree of the polynomial until we obtain a polynomial solution, consistent with Occam's Razor. Accordingly, we initially set the degree of u to be two and add together the following terms with appropriate polynomial multipliers:

$$u = c_0 + c_2 r^2, \quad (17)$$

$$\mu \frac{\partial}{\partial r} \left(r \frac{\partial}{\partial r} u \right) - r \frac{dp}{dx} = 0, \quad (18)$$

$$c_0 + c_2 R^2 = 0, \quad (19)$$

$$L \frac{dp}{dx} = -\Delta p. \quad (20)$$

Here Equation (17) posits a quadratic velocity profile in r , Equation (18) imposes a simplified version of the Navier-Stokes equations in spherical coordinates, Equation (19) imposes a no-slip boundary condition on the velocity profile of the form $u(R) = 0$, and Equation (20) posits that the pressure gradient throughout the pipe is constant. The variables in this axiom system are $u, r, R, L, \mu, \Delta p, c_0, c_2$, and $\frac{dp}{dx}$. We treat $c_0, c_2, \frac{dp}{dx}$ as variables that cannot be measured, and use the `differentiate` function in Julia to symbolically differentiate $u = c_0 + c_2 r^2$ with respect to r in Equation (18) before solving the problem, giving the equivalent expression $4c_2 \mu r - r \frac{dp}{dx}$. Solving Problem (5) with u as the dependent variable, and searching for polynomial multipliers (and polynomial q) of degree at most 3 in each variable and an overall degree of at most 6, we get the expression:

$$4rL\mu u - r\Delta p(R^2 - r^2) = 0,$$

which confirms the result. The associated polynomial multipliers for Equations (17)–(20) are:

$$\begin{aligned} &4rL\mu, \\ &r^2L - LR^2, \\ &4rL\mu, \\ &r^3 - rR^2. \end{aligned}$$

3.4 Radiated Gravitational Wave Power

We now consider the problem of deriving the power radiated from gravitational waves emitted by two-point masses orbiting their common center of gravity in a Keplerian orbit, as originally derived by Peters and Mathews [89] and verified for binary star systems by Hulse and Taylor [64]. Specifically, [89] showed that the average power generated by such a system is:

$$P = -\frac{32G^4}{5c^5 r^5} (m_1 m_2)^2 (m_1 + m_2),$$

where P is the (average) power of the wave, $G = 6.6743 \times 10^{-11} m^3 kg^{-1} s^{-2}$ is the universal gravitational constant, c is the speed of light, m_1, m_2 are the masses of the objects, and we assume that the two objects orbit a constant distance of r away from each other. Note that this equation is one of the twenty so-called bonus laws considered in the work introducing AI-Feynman [107], and notably, is one of only two such laws that neither AI-Feynman nor Eureqa [42] were able to derive. We now derive this law axiomatically, by combining the following axioms with appropriate multipliers:

$$\omega^2 r^3 - G(m_1 + m_2) = 0, \quad (21)$$

$$x5(m_1 + m_2)^2 c^5 P + G \text{Tr} \left(\frac{d^3}{dt^3} \left(m_1 m_2 r^2 \begin{pmatrix} x^2 - \frac{1}{3} & xy & 0 \\ xy & y^2 - \frac{1}{3} & 0 \\ 0 & 0 & -\frac{1}{3} \end{pmatrix} \right) \right) = 0, \quad (22)$$

$$x^2 + y^2 = 1, \quad (23)$$

where we make the variable substitution $x = \cos \phi, y = \sin \phi$, Tr stands for the trace function, and we manually define the derivative of a bivariate degree-two trigonometric polynomial in $\phi = \phi_0 + \omega t$ in (x, y) in terms of (x, y, ω) as the following linear operator:

$$\frac{d}{dt} \left(\begin{pmatrix} \sin \phi \\ \cos \phi \end{pmatrix}^\top \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \begin{pmatrix} \sin \phi \\ \cos \phi \end{pmatrix} \right) = \omega \begin{pmatrix} \sin \phi \\ \cos \phi \end{pmatrix}^\top \begin{pmatrix} a_{1,2} + a_{2,1} & a_{1,1} - a_{2,2} \\ a_{1,1} - a_{2,2} & -a_{1,2} - a_{2,1} \end{pmatrix} \begin{pmatrix} \sin \phi \\ \cos \phi \end{pmatrix}.$$

Note that Equation (21) is a restatement of Kepler's previously derived third law of planetary motion, Equation (22) provides the gravitational power of a wave when the wavelength is large compared to the source dimensions, by linearizing the equations of general relativity, the third equation defines the quadruple moment tensor, and Equation (23) (which we state as $x^2 + y^2 = 1$ within our axioms) is a standard trigonometric identity. Solving Problem (5) with P as the dependent variable, and searching for a formula involving P, G, r, c, m_1, m_2 with polynomial multipliers of degree at most 20, and allowing each variable to be raised to a power for the variables $(P, x, y, \omega, G, r, c, m_1, m_2)$ of at most $(1, 4, 4, 4, 3, 6, 1, 5, 5)$ respectively, then yields the following equation:

$$\frac{1}{4} P r^5 c^5 (m_1 + m_2)^2 = \frac{-8}{5} G^4 m_1^2 m_2^2 (m_1 + m_2)^3, \quad (24)$$

which verifies the result. Note that this equation is somewhat expensive to derive, owing to fact that searching over the set of degree 20 polynomial multipliers necessitates generating a large number of linear equalities, and writing these equalities to memory is both time and memory intensive. Accordingly, we solved Problem (5) using the MIT SuperCloud environment [93] with 640 GB RAM. The resulting system of linear inequalities involves 416392 candidate monomials, and takes 14368s to write the problem to memory and 6.58s to be solved by Mosek. This shows that the correctness of the universal gravitational wave equation can be confirmed via the following multipliers:

$$\frac{-8}{5} G m_1^2 m_2^2 (\omega^4 r^6 (x^2 + y^2)^2 + \omega^2 r^3 G (m_1 + m_2) + G^2 (m_1 + m_2)^2), \quad (25)$$

$$\frac{1}{20} r^5, \quad (26)$$

$$\frac{-8}{5} \omega^4 r^6 G^2 m_1^2 m_2^2 (m_1 + m_2) (x^2 + y^2 + 1). \quad (27)$$

Finally, Figure 4 illustrates how the Positivstellensatz derives this equation, by demonstrating that (setting $m_1 = m_2 = c = G = 1$), the gravitational wave equation is precisely the set of points (ω, r, P) where our axioms hold with equality.

3.5 Deriving Einstein's Relativistic Time Dilation Law

Next, we consider the problem of deriving Einstein's relativistic time dilation formula from a complete set of background knowledge axioms plus an inconsistent "Newtonian" axiom, which posits that light behaves like a mechanical object. We distinguish between these axioms using data on the relationship between the velocity of a light clock and the relative passage of time, as measured experimentally by Chou et al. [30] and stated explicitly in the work of Cornelio et al. [32, Tab. 6].




Figure 4: Illustration of the Positivstellensatz and its ability to recover the Radiation Gravitational Wave Power Equation in the special case where $m_1 = m_2 = c = G = 1$. Keeping other variables constant, the points that obey the power equation are the intersection of the points that obey Kepler's Third Law and the points of a linearized equation from general relativity, and the wave equation is recoverable by adding these other equations with appropriate polynomial multipliers.

Einstein's law describes the relationship between how two observers in relative motion to each other observe time, and demonstrates that observers moving at different speeds experience time differently. Indeed, letting the constant c denote the speed of light, the frequency f of a clock moving at a speed v is related to the frequency f_0 of a stationary clock via

$$\frac{f}{f_0} = \sqrt{1 - \frac{v^2}{c^2}}. \quad (28)$$

We now derive this law axiomatically, by adding together the following five axioms with appropriate polynomial multipliers:

$$cdt_0 - 2d = 0, \quad (29)$$

$$cdt - 2L = 0, \quad (30)$$

$$4L^2 + 4d^2 - v^2dt^2 = 0, \quad (31)$$

$$f dt_0 = 1, \quad (32)$$

$$f dt = 1, \quad (33)$$

plus the following (inconsistent) Newtonian axiom:

$$dt^2(v^2 + c^2) - 4L^2 = 0, \quad (34)$$

where dt_0 denotes the time required for a light to travel between two stationary mirrors separated by a distance d , and dt denotes the time required for light to travel between two similar mirrors moving at velocity v , giving a distance between the mirrors of L .

These axioms have the following meaning: Equation (29) relates the time required for light to travel between two stationary mirrors to their distance, Equation (30) similarly relates the time required for light to travel between two mirrors in motion to the effective distance L , Equation (31) relates the physical distance between the mirrors d to their effective distance L induced by the motion v via the Pythagorean theorem, and Equations (32)-(33) relate frequencies and periods. Finally, Equation (34) assumes (incorrectly) that light behaves like other mechanical objects, meaning if it is emitted orthogonally from an object traveling at velocity v , then it has velocity $\sqrt{v^2 + c^2}$.

By solving Problem (4) with a cardinality constraint that we include at most $\tau = 5$ axioms (corresponding to the exclusion of one axiom), a constraint that we must exclude either Equation (30) or Equation (34), f as the dependent variable, experimental data in f, f_0, v, c to separate the valid and invalid axioms (obtained from [32, Tab. 6] by setting $f_0 = 1$ to transform the data in $(f - f_0)/f_0$ into data in f, f_0, v, c as variables that we would like to appear in our polynomial formula $q(\mathbf{x}) = 0 \forall \mathbf{x} \in \mathcal{G} \cap \mathcal{H}$, and searching the set of polynomial multipliers of degree at most 2 in each term, we obtain the law:

$$-c^2 f_0^2 + c^2 f^2 + f_0^2 v^2 = 0, \quad (35)$$

in 6.04 seconds using Gurobi version 9.5.1. Moreover, we immediately recognize this as a restatement of Einstein's law 28. This shows that the correctness of Einstein's law can be verified by multiplying the (consistent relativistic set of) axioms by the following polynomials:

$$2df_0^2 f^2 + cf_0 f^2, \quad (36)$$

$$-cf_0^2 f - 2f_0^2 f^2 L, \quad (37)$$

$$-f_0^2 f^2, \quad (38)$$

$$-2cdf_0 f^2 - c^2 f^2, \quad (39)$$

$$c^2 dt f_0^2 f - dt f_0^2 f v^2 + c^2 f_0^2 - f_0^2 v^2. \quad (40)$$

Moreover, it verifies that relativistic axioms, particularly the axiom $cdt = 2L$, fit the light clock data of [30] better than Newtonian axioms, because, by the definition of Problem (4), AI-Hilbert selects the combination of $\tau = 5$ axioms with the lowest discrepancy between the discovered scientific formula and the experimental data.

3.6 Deriving Kepler's Third Law of Planetary Motion

In this section, we consider the problem of deriving Kepler's third law of planetary motion from a complete set of background knowledge axioms plus an incorrect candidate formula as an additional axiom, which is to be screened out using experimental data. To our knowledge, this paper is the first work that addresses this particularly challenging problem setting. Indeed, none of the approaches to scientific discovery reviewed in the introduction successfully distinguish between correct and incorrect axioms via experimental data by solving a single optimization problem. The primary motivation for this experiment is to demonstrate that AI-Hilbert provides a system for determining whether, given a background theory and experimental data, it is possible to improve upon a state-of-the-art scientific formula using background theory and experimental data.

Kepler's law describes the relationship between the distance between two bodies, e.g., the sun and a planet, and their orbital periods and takes the form:

$$p = \sqrt{\frac{4\pi^2(d_1 + d_2)^3}{G(m_1 + m_2)}}, \quad (41)$$

where $G = 6.6743 \times 10^{-11} m^3 kg^{-1} s^{-2}$ is the universal gravitational constant, m_1 and m_2 are the masses of the two bodies, d_1 and d_2 are the respective distances between m_1, m_2 and their common center of mass, and p is the orbital period. We now derive this law axiomatically by adding together the following five axioms with appropriate polynomial multipliers:

$$d_1 m_1 - d_2 m_2 = 0, \quad (42)$$

$$(d_1 + d_2)^2 F_g - G m_1 m_2 = 0, \quad (43)$$

$$F_c - m_2 d_2 w^2 = 0, \quad (44)$$

$$F_c - F_g = 0, \quad (45)$$

$$wp = 1, \quad (46)$$

plus the following (incorrect) candidate formula (as an additional axiom) proposed by Cornelio et al. [32] for the exoplanet dataset (where the mass of the planets can be discarded as negligible when added to the much bigger mass of the star):

$$p^2 m_1 - 0.1319(d_1 + d_2)^3 = 0. \quad (47)$$

Here F_g and F_c denote the gravitational and centrifugal forces in the system, and w denotes the frequency of revolution. Note that we replaced p with $2\pi p$ in our definition of revolution period in order that π does not feature in our equations; we divide p by 2π after deriving Kepler’s law.

The above axioms have the following meaning: Equation (42) defines the center of mass of the dynamical system, Equation (43) defines the gravitational force of the system, Equation (44) defines the centrifugal force of the system, Equation (45) matches the centrifugal and dynamical forces, and Equation (46) relates the frequency and the period of revolution.

Accordingly, we solve our polynomial optimization problem under a sparsity constraint that at most $\tau = 5$ axioms can be used to derive our model, a constraint that $d^c = 0$ (meaning we need not specify the hyperparameter λ in (4)), by minimizing the objective

$$\sum_{i=1}^n |q(\bar{x}_i)|,$$

where q is our implicit polynomial and $\{\bar{x}_i\}_{i=1}^4$ is a set of observations of the revolution period of binary stars stated in [32, Tab. 5]. Searching over the set of degree-five polynomials q derivable using degree six certificates then yields a mixed-integer linear optimization problem in 18958 continuous and 6 discrete variables, with the solution:

$$m_1 m_2 G p^2 - m_1 d_1 d_2^2 - m_2 d_1^2 d_2 - 2 m_2 d_1 d_2^2 = 0, \quad (48)$$

which is precisely Kepler’s third law. The validity of this equation can be verified by adding together our axioms with the weights:

$$-d_2^2 p^2 w^2, \quad (49)$$

$$-p^2, \quad (50)$$

$$d_1^2 p^2 + 2 d_1 d_2 p^2 + d_2^2 p^2, \quad (51)$$

$$d_1^2 p^2 + 2 d_1 d_2 p^2 + d_2^2 p^2, \quad (52)$$

$$m_1 d_1 d_2^2 p w + m_2 d_1^2 d_2 p w + 2 m_2 d_1 d_2^2 p w + m_1 d_1 d_2^2 + m_2 d_1^2 d_2 + 2 m_2 d_1 d_2^2, \quad (53)$$

as previously summarized in Figure 3. This is significant, because existing works on symbolic regression and scientific discovery [107, 56] often struggle to derive Kepler’s law, even given observational data. Indeed, our approach is also more scalable than deriving Kepler’s law manually; Johannes Kepler spent four years laboriously analyzing stellar data to arrive at his law [97].

3.7 Kepler Revisited With Missing Axioms

In this section, we revisit the problem of deriving Kepler’s third law of planetary motion considered in Section 3.6, with a view to verify AI-Hilbert’s ability to discover scientific laws from a combination of theory and data. Specifically, rather than providing a complete (albeit inconsistent) set of background theory as in Section 3.6, we suppress a subset of the axioms (42)–(46) and investigate the number of noiseless data points required to recover Equation (48). To simplify our analysis, we set $G = 1$ and generate noiseless data observations by sampling the values of the independent variables (the masses of the two bodies and the distance between them) uniformly at random in the ranges observed in real data (i.e., exoplanet dataset in AI-Descartes [32]) and computing the value of the dependent variable (the revolution period) using the ground truth formula.

To exploit the fact that our data observations are noiseless, we solve the following variant of (4):

$$\begin{aligned}
& \min_{q \in \mathbb{R}_{n,2d}} \frac{\lambda_1}{\sqrt{|\mathcal{D}|}} \sum_{\bar{x}_i \in \mathcal{D}} |q(\bar{x}_i)| + \lambda_2 \cdot d^c(q, \mathcal{G} \cap \mathcal{H}) + (1 - \lambda_1 - \lambda_2) \|q\|_1 \\
& \text{s.t.} \quad \sum_{\mu \in \Omega: \mu_1 \geq 1} a_\mu = 1, \\
& \quad a_\mu = 0 \quad \forall \mu \in \Omega : \sum_{j=t+1}^n \mu_j \geq 1, \\
& \quad \sum_{\bar{x}_i \in \mathcal{D}} |q(\bar{x}_i)| \leq |\mathcal{D}| \epsilon, \\
& \quad \sum_{\mu \in \Omega: \mu_1 = 0} a_\mu \leq -1/10 \quad \bigvee \quad \sum_{\mu \in \Omega: \mu_1 = 0} a_\mu \geq 1/10
\end{aligned}$$

where we set $\lambda_1 = 0.9, \lambda_2 = 0.01, \epsilon = 10^{-7}$ for all experiments, seek a degree 4 polynomial q using a proof certificate of degree at most 6, and use L_1 -coefficient norm of q as a regularization term analogously to Lasso regression [104]. Note that the second-to-last constraint ensures that the derived polynomial q explains all (noiseless) observations up to a small tolerance. Further, the last constraint is imposed as a linear inequality constraint with an auxiliary binary variable via the big- M technique [53, 16], to ensure that the derived formula includes at least one term not involving the rotation period.

Figure 5 depicts the ℓ_2 coefficient distance between the scientific formula derived by AI-Hilbert and Equation (48) mod $d_1 m_1 - d_2 m_2 = 0$ as we increase the number of data points, where we suppress the axiom $d_1 m_1 - d_2 m_2 = 0$ (left), where we suppress all axioms (right). In both cases, there is an all-or-nothing phase transition [see also 51, for a discussion of this phenomenon throughout machine learning] in AI-Hilbert’s ability to recover the scientific law, where before a threshold number of data points, AI-Hilbert cannot recover Kepler’s law, and beyond the threshold, AI-Hilbert recovers the law exactly.

Note that the increase in the coefficient distance before $m = 71$ data points reflects that solutions near $q = 0$ (and thus closer in coefficient norm to the ground truth) are optimal with $m = 0$ data points but do not fit a small number of data points perfectly, while polynomials q further from the ground truth in coefficient norm fit a small number of data points perfectly. Indeed, the total error with respect to the training data is less than 10^{-5} for all values of m in both problem settings.




Figure 5: Coefficient distance between scientific formula derived by AI-Hilbert and ground truth vs. the number of data points when we omit some axioms (left) or all axioms (right).

Figure 5 reveals that when only the axiom $d_1 m_1 - d_2 m_2 = 0$ is missing, it is possible to perform scientific discovery with as few as 46 data points, while at least 71 data points are needed when all axioms are missing. This is because the axiom $d_1 m_1 - d_2 m_2 = 0$ multiplied by the term in the proof certificate $-d_2^2 p^2 w^2$ (Section 3.6) is of a similar complexity as Kepler’s Third Law. Thus, the value of $d_1 m_1 - d_2 m_2 = 0$ is 46 data points, while the value of all axioms is 71 data points. The value of data compared to background theory depends, in general, on the problem setting and the data quality, as well as how well dispersed the data samples are.

3.8 Bell Inequalities

We now consider the problem of deriving Bell Inequalities in quantum mechanics. Bell Inequalities [14] are well-known in physics because they provide bounds on the correlation of measurements in any multi-particle system which obeys local realism (i.e., for which a joint probability distribution exists), that are violated experimentally, thus demonstrating that the natural world does not obey local realism. For ease of exposition, we prove a version called the GHZ inequality [54]. Namely, given random variables A, B, C which take values on $\{\pm 1\}$, for any joint probability distribution describing A, B, C , it follows that

$$\mathbb{P}(A = B) + \mathbb{P}(A = C) + \mathbb{P}(B = C) \geq 1, \quad (54)$$

but this bound is violated experimentally [44].

We derive this result axiomatically, using Kolmogorov's probability axioms and the specialization of our sum-of-squares framework to linear optimization proposed in Section 2.5, which is a valid specialization because the entire problem is linear. In particular, letting $p_{-1,1,-1} = \mathbb{P}(A = -1, B = 1, C = -1)$, deriving the largest lower bound for which this inequality holds is equivalent to solving the following linear optimization problem:

$$\min p_{AB} + p_{BC} + p_{AC} \quad \text{s.t. } p \in \mathcal{S},$$

where $\mathcal{S} := \{p \geq 0, e^\top p = 1\}$, $p_{AB} := p_{-1,-1,-1} + p_{-1,-1,1} + p_{1,1,-1} + p_{1,1,1}$ and p_{AC}, p_{BC} are defined similarly.

We solve this problem using Gurobi and Julia, which verifies that $\gamma = 1$ is the largest value for which this inequality holds, and obtains the desired inequality. Moreover, the solution to its dual problem yields the certificate $2p_{-1,-1,-1} + 2p_{1,1,1} \geq 0$ which verifies that 1 is indeed a valid lower bound for $p_{AB} + p_{BC} + p_{AC}$, by adding $e^\top p$ to the left-hand side of this certificate and 1 to the right-hand side.

To further demonstrate the generality and utility of our approach, we now derive a more challenging Bell inequality, namely the so-called I3322 inequality (c.f. [48]). Given particles $A_1, A_2, A_3, B_1, B_2, B_3$ which take values on $\{\pm 1\}$, the inequality reveals that for any joint probability distribution, we have:

$$\begin{aligned} \mathbb{E}[A_1] - \mathbb{E}[A_2] + \mathbb{E}[B_1] - \mathbb{E}[B_2] - \mathbb{E}[(A_1 - A_2)(B_1 - B_2)] \\ + \mathbb{E}[(A_1 + A_2)B_3] + \mathbb{E}[A_3(B_1 + B_2)] \leq 4. \end{aligned}$$

Using the same approach as previously, and defining p to be an arbitrary discrete probability measure on $\{\pm 1\}^6$, we verify that the smallest such upper bound which holds for each joint probability measure is 4, with the following polynomial certificate modulo $e^\top p = 1$:

$$\begin{aligned} &4p_{2,1,1,1,1,1} + 4p_{1,2,1,1,1,1} + 8p_{2,2,1,1,1,1} + 4p_{2,1,2,1,1,1} + 4p_{1,2,2,1,1,1} + 8p_{2,2,2,1,1,1} \\ &+ 4p_{1,1,1,2,1,1} + 8p_{2,1,1,2,1,1} + 4p_{1,2,1,2,1,1} + 8p_{2,2,1,2,1,1} + 4p_{2,1,2,2,1,1} + 4p_{2,2,2,2,1,1} \\ &+ 4p_{1,1,1,1,2,1} + 4p_{2,1,1,1,2,1} + 12p_{1,2,1,1,2,1} + 12p_{2,2,1,1,2,1} + 8p_{1,2,2,1,2,1} + 8p_{2,2,2,1,2,1} \\ &+ 8p_{1,1,1,2,2,1} + 8p_{2,1,1,2,2,1} + 12p_{1,2,1,2,2,1} + 12p_{2,2,1,2,2,1} + 4p_{1,2,2,2,2,1} + 4p_{2,2,2,2,2,1} \\ &+ 4p_{1,1,2,1,1,2} + 4p_{2,1,2,1,1,2} + 4p_{1,2,2,1,1,2} + 4p_{2,2,2,1,1,2} + 4p_{1,1,1,2,1,2} + 4p_{2,1,1,2,1,2} \\ &+ 4p_{1,1,2,2,1,2} + 4p_{2,1,2,2,1,2} + 4p_{1,1,1,1,2,2} + 8p_{1,2,1,1,2,2} + 4p_{2,2,1,1,2,2} + 4p_{1,1,2,1,2,2} \\ &+ 8p_{1,2,2,1,2,2} + 4p_{2,2,2,1,2,2} + 8p_{1,1,1,2,2,2} + 4p_{2,1,1,2,2,2} + 8p_{1,2,1,2,2,2} + 4p_{2,2,1,2,2,2} \\ &+ 4p_{1,1,2,2,2,2} + 4p_{1,2,2,2,2,2} \geq 0 \end{aligned}$$

where an index of 1 denotes that a random variable took the value -1 and an index of 2 denotes that a random variable took the value 1, and the random variables are indexed in the order $A_1, A_2, A_3, B_1, B_2, B_3$.

4 Discussion and Future Developments

In this work, we proposed a new approach to scientific discovery that leverages ideas from real algebraic geometry and mixed-integer optimization to discover new scientific laws from a possibly inconsistent or incomplete set of scientific axioms and noisy experimental data. This improves

existing approaches to scientific discovery that typically propose plausible scientific laws from either background theory alone or data alone. Indeed, by combining data and background theory in the discovery process, we potentially allow scientific discoveries to be made in previously inhospitable regimes where there is limited data and/or background theory, and gathering data is expensive. We hope our approach serves as an exciting tool that assists the scientific community in efficiently and accurately explaining the natural world.

Inspired by the success of `AI-Hilbert` in rediscovering existing scientific laws, we conclude by discussing some exciting research directions that are natural extensions of this work.

Improving the Generality of AI-Hilbert: This work proposes a symbolic discovery framework that combines background theory expressible as a system of polynomial equalities and inequalities, or that can be reformulated as such a system (e.g., in a Polar coordinate system, by substituting $x = r \cos \theta, y = r \sin \theta$ and requiring that $x^2 + y^2 = r^2$). However, many scientific discovery contexts involve background theory that cannot easily be expressed via polynomial equalities and inequalities, including differential operators, integrals, and limits, among other operators. Therefore, extending `AI-Hilbert` to encompass these non-polynomial settings would be of interest.

We point out that several authors have already proposed extensions of the sum-of-squares paradigm beyond polynomial basis functions, and these works offer a promising starting point for performing such an extension. Namely, Löfberg and Parrilo [81] (see also Bach [11] and Bach and Rudi [12]) propose an extension to trigonometric basis functions, and Fawzi et al. [46] propose approximating univariate non-polynomial functions via their Gaussian quadrature and Padé approximants. Moreover, Huchette and Vielma [63] advocate modeling non-convex functions via piecewise linear approximations with strong dual bounds. Using such polynomial approximations of non-polynomial operators offers one promising path for extending `AI-Hilbert` to the non-polynomial setting.

Automatically Parameter-Tuning AI-Hilbert: `AI-Hilbert` requires hyperparameter optimization by the user to trade-off the importance of fidelity to a model, fidelity to experimental data, and complexity of the symbolic model. Therefore, one extension of this work could be to automate this hyperparameter optimization process, by automatically solving mixed-integer and semidefinite optimization problems with different bounds on the degree of the proof certificates and different weights on the relative importance of fidelity to a model and fidelity to data, and using machine learning techniques to select solutions most likely to satisfy a scientist using `AI-Hilbert`; see also [112] for a review of automated hyperparameter optimization.

Improving the Scalability of AI-Hilbert: One limitation of our implementation of `AI-Hilbert` is that it relies on reformulating sum-of-squares optimization problems as semidefinite problems and solving them via primal-dual interior point methods (IPMs) [84, 85]. This arguably presents a limitation, because the Newton step in IPMs [see, e.g., 5] requires performing a memory-intensive matrix inversion operation. Indeed, this matrix inversion operation is sufficiently expensive that, in our experience, `AI-Hilbert` was unable to perform scientific discovery tasks with more than $n = 15$ variables and a constraint on the degree of the certificates searched over of $d = 20$ or greater (in general, runtime and memory usage is a function of both the number of symbolic variables and the degree of the proof certificates searched over).

To address this limitation and enhance the scalability of `AI-Hilbert`, there are at least three future directions to explore. First, one could exploit ideas related to the Newton polytope (or convex hull of the exponent vectors of a polynomial) [94] to reduce the number of monomials in the sum-of-squares decompositions developed in this paper, as discussed in detail in [23, Chap 3.3.4]. Second, one could use presolving techniques such as chordal sparsity [55, 109] or partial facial reduction [88, 114] to reduce the number of variables in the semidefinite optimization problems that arise from sum-of-squares optimization problems. Third, one could attempt to solve sum-of-squares problems without using computationally expensive interior point methods for semidefinite programs, e.g., by using a Burer-Monteiro factorization approach [28, 78] or by optimizing over a second-order cone inner approximation of the positive semidefinite cone [4].

Acknowledgements

We are grateful to Ken Clarkson, Joao Goncalves, Phokion Kolatis, Jon Lenchner, and Nimrod Megiddo (IBM) for valuable discussions on scientific discovery and two anonymous referees and the senior editor for valuable comments that improved the manuscript. Ryan Cory-Wright and Bachir El Khadir both gratefully acknowledge IBM Research for hosting them as IBM Goldstine postdoctoral fellows while most of this work was conducted.

References

- [1] T. Achterberg. What’s new in Gurobi 9.0, 2019.
- [2] A. A. Ahmadi, E. De Klerk, and G. Hall. Polynomial norms. *SIAM Journal on Optimization*, 29(1):399–422, 2019.
- [3] A. A. Ahmadi and B. E. Khadir. Learning dynamical systems with side information. *SIAM Review*, 65(1):183–223, 2023.
- [4] A. A. Ahmadi and A. Majumdar. DSOS and SDSOS optimization: More tractable alternatives to sum of squares and semidefinite optimization. *SIAM Journal on Applied Algebra and Geometry*, 3(2):193–230, 2019.
- [5] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton. Primal-dual interior-point methods for semidefinite programming: Convergence rates, stability and numerical results. *SIAM Journal on Optimization*, 8(3):746–768, 1998.
- [6] E. D. Andersen and K. D. Andersen. The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. *High Performance Optimization*, pages 197–232, 2000.
- [7] A. Arora, S. Belenzon, and A. Pataconi. The decline of science in corporate R&D. *Strategic Management Journal*, 39(1):3–32, 2018.
- [8] G. B. Arous, A. S. Wein, and I. Zadik. Free energy wells and overlap gap property in sparse PCA. In *Conference on Learning Theory*, pages 479–482. PMLR, 2020.
- [9] E. Artin. Über die zerlegung definiter funktionen in quadrate. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 5, pages 100–115. Springer, 1927.
- [10] V. Austel, S. Dash, O. Gunluk, L. Horesh, L. Liberti, G. Nannicini, and B. Schieber. Globally optimal symbolic regression. *arXiv preprint arXiv:1710.10720*, 2017.
- [11] F. Bach. Sum-of-squares relaxations for information theory and variational inference. *arXiv preprint arXiv:2206.13285*, 2022.
- [12] F. Bach and A. Rudi. Exponential convergence of sum-of-squares hierarchies for trigonometric polynomials. *SIAM Journal on Optimization*, 33(3):2137–2159, 2023.
- [13] Z. J. Baum, X. Yu, P. Y. Ayala, Y. Zhao, S. P. Watkins, and Q. Zhou. Artificial intelligence in chemistry: current trends and future directions. *Journal of Chemical Information and Modeling*, 61(7):3197–3212, 2021.
- [14] J. S. Bell. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3):195, 1964.
- [15] D. Bertsimas, R. Cory-Wright, S. Lo, and J. Pauphilet. Optimal low-rank matrix completion: Semidefinite relaxations and eigenvector disjunctions. *arXiv preprint arXiv:2305.12292*, 2023.
- [16] D. Bertsimas, R. Cory-Wright, and J. Pauphilet. A unified approach to mixed-integer optimization problems with logical constraints. *SIAM Journal on Optimization*, 31(3):2340–2367, 2021.
- [17] D. Bertsimas, R. Cory-Wright, and J. Pauphilet. Mixed-projection conic optimization: A new paradigm for modeling rank constraints. *Operations Research*, 70(6):3321–3344, 2022.

- [18] D. Bertsimas and J. Dunn. *Machine learning under a modern optimization lens*. Dynamic Ideas Press, 2019.
- [19] D. Bertsimas and W. Gurnee. Learning sparse nonlinear dynamics via mixed-integer optimization. *Nonlinear Dynamics*, pages 1–20, 2023.
- [20] D. Bertsimas, A. King, and R. Mazumder. Best subset selection via a modern optimization lens. *The Annals of Statistics*, 44(2):813 – 852, 2016.
- [21] J. Bhattacharya and M. Packalen. Stagnation and scientific incentives. Technical report, National Bureau of Economic Research, 2020.
- [22] R. Bixby and E. Rothberg. Progress in computational mixed integer programming—a look back from the other side of the tipping point. *Annals of Operations Research*, 149(1):37, 2007.
- [23] G. Blekherman, P. A. Parrilo, and R. R. Thomas. *Semidefinite optimization and convex algebraic geometry*. SIAM, 2012.
- [24] N. Bloom, C. I. Jones, J. Van Reenen, and M. Webb. Are ideas getting harder to find? *American Economic Review*, 110(4):1104–1144, 2020.
- [25] J. Bongard and H. Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948, 2007.
- [26] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- [27] E. Brynjolfsson, D. Rock, and C. Syverson. Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *The economics of artificial intelligence: An agenda*, pages 23–57. University of Chicago Press, 2018.
- [28] S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- [29] E. J. Candes and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [30] C.-W. Chou, D. B. Hume, T. Rosenband, and D. J. Wineland. Optical clocks and relativity. *Science*, 329(5999):1630–1633, 2010.
- [31] M. Clegg, J. Edmonds, and R. Impagliazzo. Using the Gröebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 174–183, 1996.
- [32] C. Cornelio, S. Dash, V. Austel, T. Josephson, J. Goncalves, K. Clarkson, N. Megiddo, B. E. Khadir, and L. Horesh. Combining data and theory for derivable scientific discovery with AI-Descartes. *Nature Communications*, 14(1777), 2023.
- [33] T. Cowen. *The great stagnation: How America ate all the low-hanging fruit of modern history, got sick, and will (eventually) feel better: A Penguin eSpecial from Dutton*. Penguin, 2011.
- [34] D. Cox, J. Little, and D. O’Shea. *Ideals, varieties, and algorithms: An introduction to computational algebraic geometry and commutative algebra*. Springer Science & Business Media, 2013.
- [35] A. Cozad and N. V. Sahinidis. A global MINLP approach to symbolic regression. *Mathematical Programming*, 170:97–119, 2018.
- [36] M. Cranmer. PySR: Fast & parallelized symbolic regression in Python/Julia. doi.org/10.5281/zenodo.4041459, Sept. 2020.
- [37] M. Cranmer, A. Sanchez-Gonzalez, P. Battaglia, R. Xu, K. Cranmer, D. Spergel, and S. Ho. Discovering symbolic models from deep learning with inductive biases. *NeurIPS 2020*, 2020.

- [38] M. Curmei and G. Hall. Shape-constrained regression using sum of squares polynomials. *Operations Research*, 2023.
- [39] H. W. De Regt. Understanding, values, and the aims of science. *Philosophy of Science*, 87(5):921–932, 2020.
- [40] S. S. Dey, Y. Dubey, and M. Molinaro. Branch-and-bound solves random binary ips in polytime. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 579–591. SIAM, 2021.
- [41] P. A. Dirac. Directions in physics. Lectures delivered during a visit to Australia and New Zealand, August/September 1975. 1978.
- [42] R. Dubčáková. Eureka: software review, 2011.
- [43] M. R. Engle and N. V. Sahinidis. Deterministic symbolic regression with derivative information: General methodology and application to equations of state. *AIChE Journal*, 68(6):e17457, 2022.
- [44] A. Fahmi. Locality, Bell’s inequality and the GHZ theorem. *Physics Letters A*, 303(1):1–6, 2002.
- [45] A. Fawzi, M. Malinowski, H. Fawzi, and O. Fawzi. Learning dynamic polynomial proofs. *Advances in Neural Information Processing Systems*, 32, 2019.
- [46] H. Fawzi, J. Saunderson, and P. A. Parrilo. Semidefinite approximations of the matrix logarithm. *Foundations of Computational Mathematics*, 19:259–296, 2019.
- [47] R. P. Feynman, R. B. Leighton, and M. Sands. The Feynman lectures on physics; vol. i. *American Journal of Physics*, 33(9):750–752, 1965.
- [48] M. Froissart. Constructive generalization of Bell’s inequalities. *Nuovo Cimento B;(Italy)*, 64(2), 1981.
- [49] N. Fujinuma, B. DeCost, J. Hattrick-Simpers, and S. E. Lofland. Why big data and compute are not necessarily the path to big materials science. *Communications Materials*, 3(1):59, 2022.
- [50] N. Fulton, S. Mitsch, J.-D. Quesel, M. Völp, and A. Platzer. KeYmaera X: An axiomatic tactical theorem prover for hybrid systems. In *Automated Deduction-CADE-25: 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings 25*, pages 527–538. Springer, 2015.
- [51] D. Gamarnik. The overlap gap property: A topological barrier to optimizing over random structures. *Proceedings of the National Academy of Sciences*, 118(41):e2108492118, 2021.
- [52] D. Gamarnik and I. Zadik. High dimensional regression with binary coefficients. estimating squared error and a phase transition. In *Conference on Learning Theory*, pages 948–953. PMLR, 2017.
- [53] F. Glover. Improved linear integer programming formulations of nonlinear integer problems. *Management Science*, 22(4):455–460, 1975.
- [54] D. M. Greenberger, M. A. Horne, A. Shimony, and A. Zeilinger. Bell’s theorem without inequalities. *American Journal of Physics*, 58(12):1131–1143, 1990.
- [55] A. Griewank and P. L. Toint. On the existence of convex decompositions of partially separable functions. *Mathematical Programming*, 28:25–49, 1984.
- [56] R. Guimerà, I. Reichardt, A. Aguilar-Mogas, F. A. Massucci, M. Miranda, J. Pallarès, and M. Sales-Pardo. A Bayesian machine scientist to aid in the solution of challenging scientific problems. *Science Advances*, 6(5):eaav6971, 2020.
- [57] A. Guntuboyina and B. Sen. Nonparametric shape-restricted regression. *Statistical Science*, 33(4):568–594, 2018.

- [58] S. D. Gupta, B. P. Van Parys, and E. K. Ryu. Branch-and-bound performance estimation programming: A unified methodology for constructing optimal optimization methods. *Mathematical Programming*, 2023.
- [59] G. Hall. Applications of sums of squares polynomials. In P. Parrilo and R. Thomas, editors, *Sum of Squares: Theory and Applications*, volume 77. Proceedings of Symposia in Applied Mathematics, 2020.
- [60] T. Hastie, R. Tibshirani, J. H. Friedman, and J. H. Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.
- [61] D. Hilbert. Über die darstellung definitiver formen als summe von formenquadraten. *Mathematische Annalen*, 32(3):342–350, 1888.
- [62] D. Hilbert. Mathematical problems. In *Mathematics*, pages 273–278. Chapman and Hall/CRC, 2019.
- [63] J. Huchette and J. P. Vielma. Nonconvex piecewise linear functions: Advanced formulations and simple modeling tools. *Operations Research*, 2022.
- [64] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *The Astrophysical Journal*, 195:L51–L53, 1975.
- [65] R. Iten, T. Metger, H. Wilming, L. Del Rio, and R. Renner. Discovering physical concepts with neural networks. *Physical Review Letters*, 124(1):010508, 2020.
- [66] I. M. Johnstone and A. Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.
- [67] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [68] G. Karagiorgi, G. Kasieczka, S. Kravitz, B. Nachman, and D. Shih. Machine learning in the search for new fundamental physics. *Nature Reviews Physics*, 4(6):399–412, 2022.
- [69] H. Kitano. Nobel turing challenge: creating the engine for scientific discovery. *npj Systems Biology and Applications*, 7(1):29, 2021.
- [70] J.-L. Krivine. Anneaux préordonnés. *Journal d’analyse mathématique*, 12:p–307, 1964.
- [71] J. Kubalík, E. Derner, and R. Babuška. Symbolic regression driven by training data and prior knowledge. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, pages 958–966, 2020.
- [72] J. Kubalík, E. Derner, and R. Babuška. Multi-objective symbolic regression for physics-aware dynamic modeling. *Expert Systems with Applications*, 182:115210, 2021.
- [73] M. Landajuela, C. S. Lee, J. Yang, R. Glatt, C. P. Santiago, I. Aravena, T. Mundhenk, G. Mulcahy, and B. K. Petersen. A unified framework for deep symbolic regression. *Advances in Neural Information Processing Systems*, 35:33985–33998, 2022.
- [74] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- [75] J. B. Lasserre. A sum of squares approximation of nonnegative polynomials. *SIAM Review*, 49(4):651–669, 2007.
- [76] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, pages 157–270. Springer, 2009.
- [77] M. Laurent and F. Rendl. Semidefinite programming and integer programming. *Handbooks in Operations Research and Management Science*, 12:393–514, 2005.

- [78] B. Legat, C. Yuan, and P. Parrilo. Low-rank univariate sum of squares has no spurious local minima. *SIAM Journal on Optimization*, 33(3):2041–2061, 2023.
- [79] E. Lim and P. W. Glynn. Consistency of multidimensional convex regression. *Operations Research*, 60(1):196–208, 2012.
- [80] J. Liu, S. Rosen, C. Zhong, and C. Rudin. Okridge: Scalable optimal k-sparse ridge regression. *Advances in Neural Information Processing Systems*, 36, 2024.
- [81] J. Lofberg and P. A. Parrilo. From coefficients to samples: A new approach to SOS optimization. In *2004 43rd IEEE Conference on Decision and Control (CDC)(IEEE Cat. No. 04CH37601)*, volume 3, pages 3154–3159. IEEE, 2004.
- [82] M. Lubin, O. Dowson, J. D. Garcia, J. Huchette, B. Legat, and J. P. Vielma. Jump 1.0: recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*, pages 1–9, 2023.
- [83] Y. Matsubara, N. Chiba, R. Igarashi, and Y. Ushiku. SRSD: Rethinking datasets of symbolic regression for scientific discovery. In *NeurIPS 2022 AI for Science: Progress and Promises*, 2022.
- [84] Y. Nesterov and A. Nemirovskii. *Interior-point polynomial algorithms in convex programming*. SIAM, 1994.
- [85] Y. E. Nesterov and M. J. Todd. Self-scaled barriers and interior-point methods for convex programming. *Mathematics of Operations Research*, 22(1):1–42, 1997.
- [86] OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [87] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2):293–320, 2003.
- [88] F. Permenter and P. Parrilo. Partial facial reduction: Simplified, equivalent SDPs via approximations of the PSD cone. *Mathematical Programming*, 171:1–54, 2018.
- [89] P. C. Peters and J. Mathews. Gravitational radiation from point masses in a Keplerian orbit. *Physical Review*, 131(1):435, 1963.
- [90] M. Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3):969–984, 1993.
- [91] M. V. Ramana. An exact duality theory for semidefinite programming and its complexity implications. *Mathematical Programming*, 77:129–162, 1997.
- [92] J. Renegar. *A mathematical view of interior-point methods in convex optimization*. SIAM, 2001.
- [93] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, M. Jones, A. Klein, L. Milechin, J. Mullen, A. Prout, A. Rosa, C. Yee, and P. Michaleas. Interactive supercomputing on 40,000 cores for machine learning and data analysis. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2018.
- [94] B. Reznick. Extremal PSD forms with few terms. *Duke Mathematical Journal*, 45(2):363–374, 1978.
- [95] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- [96] S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz. Data-driven discovery of partial differential equations. *Science Advances*, 3(4):e1602614, 2017.
- [97] J. L. Russell. Kepler’s laws of planetary motion: 1609–1666. *The British Journal for the History of Science*, 2(1):1–24, 1964.

- [98] M. Schmidt and H. Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, 2009.
- [99] M. Schmidt and H. Lipson. Symbolic regression of implicit equations. In *Genetic Programming Theory and Practice VII*, pages 73–85. Springer, 2009.
- [100] K. Schmüdgen. The moment problem on compact semi-algebraic sets. *Mathematische Annalen*, 289:283–313, 1991.
- [101] H. A. Simon. Does scientific discovery have a logic? *Philosophy of Science*, 40(4):471–480, 1973.
- [102] A. Skajaa and Y. Ye. A homogeneous interior-point algorithm for nonsymmetric convex conic optimization. *Mathematical Programming*, 150(2):391–422, 2015.
- [103] G. Stengle. A Nullstellensatz and a Positivstellensatz in semialgebraic geometry. *Mathematische Annalen*, 207(2):87–97, 1974.
- [104] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.
- [105] S. Udrescu, A. Tan, J. Feng, O. Neto, T. Wu, and M. Tegmark. AI Feynman 2.0: Pareto-optimal symbolic regression exploiting graph modularity. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [106] S.-M. Udrescu, A. Tan, J. Feng, O. Neto, T. Wu, and M. Tegmark. AI Feynman 2.0: Pareto-optimal symbolic regression exploiting graph modularity. *Advances in Neural Information Processing Systems*, 33:4860–4871, 2020.
- [107] S.-M. Udrescu and M. Tegmark. AI Feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16), 2020.
- [108] S.-M. Udrescu and M. Tegmark. AI Feynman: A physics-inspired method for symbolic regression. *Science Advances*, 2020.
- [109] L. Vandenberghe, M. S. Andersen, et al. Chordal graphs and semidefinite optimization. *Foundations and Trends® in Optimization*, 1(4):241–433, 2015.
- [110] H. Wang, T. Fu, Y. Du, W. Gao, K. Huang, Z. Liu, P. Chandak, S. Liu, P. Van Katwyk, A. Deac, et al. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023.
- [111] W. Wang, M. J. Wainwright, and K. Ramchandran. Information-theoretic limits on sparse signal recovery: Dense versus sparse measurement matrices. *IEEE Transactions on Information Theory*, 56(6):2967–2979, 2010.
- [112] T. Yu and H. Zhu. Hyper-parameter optimization: A review of algorithms and applications. *arXiv preprint arXiv:2003.05689*, 2020.
- [113] W. Zhao and G. Zhou. Hausdorff distance between convex semialgebraic sets. *Journal of Global Optimization*, pages 1–21, 2023.
- [114] Y. Zhu, G. Pataki, and Q. Tran-Dinh. Sieve-SDP: A simple facial reduction algorithm to preprocess semidefinite programs. *Mathematical Programming Computation*, 11:503–586, 2019.

A Comparison with Methods from the Literature

To illustrate our system’s capabilities, Table 1 compares AI-Hilbert with four approaches from the literature in terms of their ability to recover various scientific laws from background theory and experimental data.

The data for the first two problems comes from real measurements (denoted with “R” in the table) while for the other problems we rely on simulated data (denoted with “S” in the table). The data is created similarly to the approach used in the literature (see [32] or [108] for more details): we sample ten data points uniformly at random for each problem, independently for each variable in a given continuous range of values. The range is either defined following ranges in real measurements or set artificially (e.g., $(0, 1]$). We then add 1% Gaussian noise to the dependent variable (following [108]), thereby creating small, noisy datasets to resemble real-life data.

We used the background theory provided by Cornelio et al. [32] for the first eight problems and manually extracted the axioms for the remaining ones.

The systems we consider are:

- **AI-Descartes** [32] is a system that combines logical reasoning with symbolic regression to generate meaningful formulae explaining a given phenomenon. It starts by generating a set of candidate models from data alone using symbolic regression and then validates and ranks them using a background theory. In this way the system ensure that the chosen hypothesis is not only close to the data but also respects the laws already known about the environment, often providing a logical proof of its correctness. However, in the experiments below, we only run the symbolic regression module of AI-Descartes.
- **AI Feynman** [107, 105] is a symbolic regression approach integrating deep learning methodologies with the exploitation of some features commonly found in physics functions including unit presence, utilization of low-order polynomials, and symmetry. The algorithm comprises several modules handling polynomial fitting, exhaustive enumeration, neural network fitting, among other tasks, ultimately generating a list of potential formulas.
- **PySR** [36, 37] employs a combination of regularized evolution, simulated annealing, and gradient-free optimization techniques to search for equations that best match the provided input data.
- **Bayesian Machine Scientist (BMS)** [56] utilizes a Markov chain Monte Carlo approach to traverse the potential model space, leveraging prior expectations learned from a large empirical corpus of mathematical expressions.

The above systems return a score (in some cases this corresponds to the error) and the complexity for each output formula. The systems PySR and Bayesian Machine Scientist return a best candidate. For AI-Descartes and AI-Feynman, we declare success if the correct formula is among the candidates returned, whereas for PySR and Bayesian Machine Scientist, we check if the best candidate formula is correct.

A.1 Experimental results for comparison with prior literature

We now compare AI-Hilbert with four methods from the literature in Table 1. Note that all four existing methods that we compare AI-Hilbert against require experimental data to successfully discover scientific laws; further, we take results for these methods for the FSRD problems, the Relativistic time dilation problem, and Kepler’s third law from [32]. On the other hand, AI-Hilbert combines (potentially limited and noisy) data and (potentially corrupted) theory, including the two limiting cases where either only data or only theory are supplied.

AI-Hilbert outperforms the existing systems, accurately recovering formulas in all studied problems. Moreover it achieves this with less data, by leveraging relevant background theory. Finally, our method is the only one capable of recovering symbolic constants (while the other systems combine them into single or multiple real values) thereby preserving the meaning and semantics of each individual constant.

We remark that for some methods from the literature, inexact recovery (denoted with a \checkmark^*) may occur because these approaches are data-driven, and thus usually cannot derive symbolic constants

associated with scientific laws. Instead, they group information concerning multiple constants into one that is less interpretable. For instance, consider the Compton Scattering problem (B.7), the correct formula is $l_2 = l_1 + \frac{h}{m_e c}(1 - \cos(\theta))$ while data-driven methods usually only derive $l_2 = l_1 + k(1 - \cos(\theta))$ (where k is a real number). Moreover, sometimes integers are reported with a small numerical error (e.g., 1, 001 or 0.999 instead of 1).

Table 1: Comparison of AI-Hilbert with state-of-the-art discovery methods in terms of ability to recover scientific laws from background theory and experimental data.

Data	AI-Hilbert	AI-Descartes [32]	AI-Feynman [107]	PySR [36]	BMS [56]
S [3.3] Hagen Poiseuille	✓	✗	✗	✗	✗
S [3.4] Gravitational Wave Power	✓	✗	✗	✗	✗
R [3.5] Relat. Time Dilat.	✓	✗	✗	✗	✗
R [3.6] Kepler's 3 Law	✓	✓*	✗	✓	✗
S [3.8] Bell inequalities [†]	✓	✗	✗	✗	✗
S [B.1] I.15.10 FSRD	✓	✓*	✗	✗	✗
S [B.1] I.27.6 FSRD	✓▷	✓*	✓*	✓	✗
S [B.1] I.34.8 FSRD	✓▷	✓	✓*	✓	✓*
S [B.1] I.43.16 FSRD	✓	✓*	✓*	✓	✓*
S [B.1] II.10.9 FSRD	✓	✓*	✓*	✓*	✓*
S [B.1] II.34.2 FSRD	✓	✓*	✓*	✓*	✓*
S [B.2] Inelastic Relativ. Collision	✓	✓*	✗	✗	✗
S [B.3] Decay of Pion [◊]	✓	✓*	✗	✗	✓*
S [B.4] Radiation Damping	✓	✗	✗	✓*	✓*
S [B.5] Escape Velocity	✓	✓	✓*	✓*	✓*
S [B.6] Hall Effect	✓	✓*	✗	✓*	✓*
S [B.7] Compton Scattering	✓	✓*	✓*	✓*	✓*

Notation used in table:

- ✓ denotes the successful recovery of a scientific law.
- ✗ denotes the failure to recover a scientific law.
- ✓* denotes recovery up to constants, but not exact recovery.
- ✓▷ denotes successful recovery when some variables that are potentially not observable are still declared to be observable.

[†] Inequalities are not supported by the methods from the literature that we benchmarked against.

[◊] To generate data for this problem, we varied the mass values of both the pion and the muon. These variations were intended to mimic small differences in the observed values of the masses, rather than altering the actual values across a broad range of values.

B Discovery of Additional Problems From Background Theory

We validate `AI-Hilbert` on six problems from Feynman’s lectures, considered earlier in `AI-Descartes` [32]. We also examined six additional problems to demonstrate `AI-Hilbert`’s capacity to rediscover scientific laws from background theory alone. Specifically, in Appendix B.2, we derive the inelastic relativistic collision law, in Appendix B.3, we derive the kinetic energy and momentum when a pion at rest decays into a muon and a neutrino, in Appendix B.4 we derive the light scattering law, in Appendix B.5 we derive the minimal velocity that enables an object to overcome the gravitational pull of a planet, in Appendix B.6 we derive the Hall potential of an electrical conductor, and in Appendix B.7 we derive the Compton scattering formula.

B.1 Problems from FSRD

We consider the background theory provided in `AI-Descartes` [32] for the following six problems from Feynman’s Lectures: I.15.10, I.27.6, I.34.8, I.43.16, II.10.9 and II.34.2. As shown in the previous table (Table 1), `AI-Hilbert` obtained the correct symbolic expression for each problem. For the second and third problems, we declared some potentially nonmeasurable variables measurable and permitted `AI-Hilbert` to search for polynomials incorporating these variables. However, the final polynomial found by `AI-Hilbert` did not include these nonmeasurable variables.

B.2 Inelastic Relativistic Collision

We consider the problem of computing the momentum and speed of a particle created by colliding a fast-moving particle with a stationary particle. This problem appears within the exercises in the Feynman Lectures on Physics series [47].

We first introduce relevant notation: let m denote the mass of the moving particle, c be the speed of light, $v_m := 4/5c$ denote the speed of the moving particle pre-collision, m_c and v_c denote the mass and speed, respectively, of the composite particle after the collision, p_m and p_c denote the momentum of the moving particle before and after collision, respectively, and E_m, E_r, E_c denote, respectively, the energy of the moving particle before collision, the resting particle before collision and the composite particle after collision.

Then, given that a particle of mass m moving at a speed of $v_m = 4c/5$ collides in elastically with a similar particle at rest, we have the relations

$$p_c = \frac{4m_c v_c c}{\sqrt{3(c^2 - v_c^2)}}, m_c = \frac{4m}{\sqrt{3}}$$

We now derive these relations axiomatically using `AI-Hilbert`, from the following axioms (via appropriate polynomial multipliers):

$$p_m^2(c^2 - v_m^2) - m^2 v_m^2 c^2 = 0 \tag{55}$$

$$E_m^2 - (m c^2)^2 - p_m^2 c^2 = 0 \tag{56}$$

$$E_r - m c^2 = 0 \tag{57}$$

$$E_c^2 - m_c^2 c^4 - p_c^2 c^2 = 0 \tag{58}$$

$$2E_m E_r - E_c^2 + E_m^2 + E_r^2 = 0 \tag{59}$$

$$p_c - p_m = 0 \tag{60}$$

$$v_m^2 - 16/25 c^2 = 0 \tag{61}$$

$$p_c^2(c^2 - v_c^2) - m_c^2 v_c^2 c^2 = 0. \tag{62}$$

Equation (55) gives the momentum of the moving particle pre-collision (by rewriting the expression $p_m = \frac{mv}{\sqrt{1-v^2/c^2}}$ as a polynomial equation). Equation (56) models the energy of the moving particle pre-collision, Equation (57) models the energy of the particle at rest pre-collision, Equation (58) models the energy of the combined particle post-collision, Equation (59) models the conservation of energy (squaring all terms to reduce the degree of the proof certificate; we could instead introduce $E_c - E_m - E_r = 0$ and obtain essentially the same result), Equation (60) models the conservation of momentum, Equation (61) defines the velocity of the moving particle pre-collision (again squaring

both sides to reduce the degree of the proof certificate), and Equation (62) defines the momentum of the composite particle.

By requiring that only the terms m, m_c, c appear in our final formula, AI-Hilbert provides the following polynomial equality as an output

$$-1.33m_c^4(c^2)^3 + 5.33m^2m_c^2(c^2)^3 + 9.48m^4(c^2)^3 = 0$$

Moreover, this equality can be manipulated to rediscover the two relations at the start of this section. In particular, it follows directly from the equality that $m_c^2 = 16/3m^2$, and hence $m_c = \frac{4m}{\sqrt{3}}$

Moreover, substituting this expression for m_c into Equation (62) and re-arranging yields:

$$p_c = \frac{16mv_c c}{3\sqrt{c^2 - v_c^2}}$$

B.3 Decay of Pion into Muon and Neutrino

Next, we consider the problem of deriving the kinetic energy and momentum generated when a pion at rest decays into a muon and a neutrino.

We first introduce relevant notation: let m_π, m_μ and m_ν stand for, respectively, the mass of the pion, muon, and neutrino. Further, let p_μ be the momentum of the muon, p_ν be the momentum of the neutrino, and let E_π, E_μ, E_ν be the respective total energy of the pion, the muon, and the neutrino. Given these quantities, the momentum of the neutrino satisfies the relation

$$p_\nu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi}. \quad (63)$$

Accordingly, we now derive this relation by adding together the following background theory axioms with appropriate polynomial multipliers (normalizing in order that $c = 1$ for simplicity):

$$p_\nu - p_\mu = 0 \quad (64)$$

$$E_\pi - m_\pi = 0 \quad (65)$$

$$E_\nu - p_\nu = 0 \quad (66)$$

$$E_\pi - E_\mu - E_\nu = 0 \quad (67)$$

$$E_\mu^2 - p_\mu^2 - m_\mu^2 = 0 \quad (68)$$

where Equation (64) ensures that the decay obeys conservation of momentum, Equation (65) defines the energy of the pion (recall that we set $c = 1$ here), Equation (66) defines the energy of the neutrino (which equals its momentum as it is massless), Equation (67) enforces conservation of energy, and Equation (68) defines the kinematic equation of the muon particle.

We assume that the only measured quantities are p_ν, m_π, m_μ . Then AI-Hilbert derives the equality

$$-2m_\pi^2 + 2m_\mu^2 + 4p_\nu m_\pi = 0$$

which is equivalent to the original relation (63).

B.4 Radiation Damping and Light Scattering

We now consider the problem of computing the total amount of energy radiated by a non-relativistic acceleration of a charge.

Denote by S the amount of energy that passes per square meter per second through a surface normal to the radiation, let q_c represent an electric charge accelerating at rate a_p , let r and θ represent, respectively, the distance and angle at which the radiation is observed, and suppose that the charge is oscillating with displacement x_0 at frequency w . Given this notation, we have the natural law:

$$P = \frac{4}{3}\pi q_c^2 x_0^2 w^4 \quad (69)$$

For this problem, we use the following background theory:

$$Sr^2 - q_c^2 a_p^2 \sin(\theta)^2 = 0 \quad (70)$$

$$dA - 2\pi r^2 \sin(\theta) d\theta = 0 \quad (71)$$

$$P - \int_0^\pi S dA = 0 \quad (72)$$

$$\frac{4}{3} - \int_0^\pi \sin(\theta)^3 d\theta = 0 \quad (73)$$

$$a_p^2 - \frac{1}{2} w^4 x_0^2 = 0 \quad (74)$$

where Equation (70) models the rate of radiation of energy, Equation (71) represents a differential of the area of a spherical segment, Equation (72) is the power to rate of radiation relation, Equation (73) is a definite integral from 0 to π , such an integral can be provided as background knowledge, or can be solved by symbolic/numerical integrator, and lastly Equation (74) refers to the average of the acceleration of an oscillating charge over a cycle (squared).

AI-Hilbert provides the relation

$$4.0P - 16.76q_c^2 w^4 x_0^2 = 0$$

which can easily be manipulated to recover Equation (69).

B.5 Escape Velocity

We now consider the problem of deriving the escape velocity of a sphere-shaped planet, that is, the minimal velocity needed for any object attempting to exit a planet to overcome the planet's gravitational pull. Let M denote the mass of the planet, r denote its radius, m denote the mass of the object, v_e denote the escape velocity of the object, K_i and K_f denote the initial and final kinetic energy of the object, U_i and U_f denote the initial and final potential energy of the system, and G denote the universal gravitational constant. We have the relation:

$$v_e = \sqrt{\frac{2GM}{r}} \quad (75)$$

We now recover this scientific law via AI-Hilbert from the background theory consisting of the following polynomial equalities:

$$K_i - \frac{1}{2} m v_e^2 = 0 \quad (76)$$

$$K_f - 0 = 0 \quad (77)$$

$$U_i r + GMm = 0 \quad (78)$$

$$U_f - 0 = 0 \quad (79)$$

$$K_i + U_i - (K_f + U_f) = 0, \quad (80)$$

where Equation (76) models the initial kinetic energy of the object, Equation (77) models the final kinetic energy of the object, Equations (78)–(79) reflect the initial and final potential energy of the body, and Equation (80) models conservation of energy.

Assuming that the only measured quantities are G, M, m, r and v_e . AI-Hilbert uncovers the following polynomial relation that is derivable from the background theory:

$$2GMm - m v_e^2 r = 0.$$

This can easily be rearranged to recover Equation (75).

B.6 Hall Effect

The Hall effect produces a potential difference (the Hall voltage, U_H) across an electrical conductor transverse to an electric current in the conductor and to an applied magnetic field perpendicular to the

current. Let a metal plate of length L , width h , and depth d be subjected to a potential field across its length dimension. The plate is placed in a homogeneous magnetic field B and perpendicular to the electric field E (oriented with the depth dimension). The electron charges q_e moving at velocity v across the length dimension, will be deflected by the magnetic field (due to a Lorentz force F_m). This will entail an excess of electrons on one side of the plate, and a deficiency on the other side, resulting in an electric field E and respective electric force F_e opposite to the magnetic force. Let the potential difference across the width dimension of the plate be U_H . Finally, let I denote the current, which is defined by the number of electron charges N crossing per time interval dt . The Hall potential is given by the expression:

$$U_H = \frac{hLIB}{Nq_e}. \quad (81)$$

Equation (81) is recoverable from the following background theory via AI-Hilbert:

$$F_m - q_e v B = 0 \quad (82)$$

$$F_e - q_e E = 0 \quad (83)$$

$$F_m - F_e = 0 \quad (84)$$

$$Eh - U_H = 0 \quad (85)$$

$$v dt - L = 0 \quad (86)$$

$$I dt - N q_e = 0 \quad (87)$$

where Equation (82) models the Lorentz magnetic force, Equation (83) represents the electric force (E is the homogeneous electric field between the upper and lower metal plate), Equation (84) models Newton's third law, Equation (85) relates the electric potential U_H to field E per plate of width h , Equation (86) models the traversal velocity $v = \frac{L}{dt}$ across the plate length L , and Equation (87) models the amount of charge Q , is the current I times dt (can be measured by an amp meter).

We assume the measured variables are L, h, I, B, N, q_e and U_H . AI-Hilbert provides the relation

$$q_e^2 U_H L N^2 - q_e B h L^2 I N = 0,$$

which can easily be rearranged to recover Equation (81).

B.7 Compton Scattering

Compton scattering refers to the behavior of high frequency photons as they scatter after collision with a charged particle, usually an electron e in an atom. Specifically, when a photon knocks out a loosely bound electron from the outer valence shells of of an atom (or molecule), a new photon is emitted from the atom traveling at an angle θ to the incoming photon's path. Compton related the shift in photon wavelengths to the scattering angle: $\lambda_2 - \lambda_1 = \frac{h}{m_e c} (1 - \cos \theta)$ where λ_1, λ_2 are the initial and final photon wavelengths, m_e is the electron rest mass, c is the speed of light, and h is Planck's constant.

The following equations give a complete set of axioms needed to derive Compton's formula. Let E_r be the energy of the electron at rest (i.e., an electron bound to an atom), let E_m be the energy of the moving electron (after it is knocked out from the atom). Let E_1, E_2 be the initial and final photon energies, f_1, f_2 be the initial and final photon frequencies, and p_1, p_2 be the initial and final photon momentum values. The first constraint is just the expression for conservation of energy. The second and third constraints give the photon energy in terms of frequency. The subsequent two constraints give photon momentum in terms of frequency, and the two constraints after that relate wavelength and frequency. The next constraint is the mass-energy equivalence of the electron at rest, while the constraint after that gives the energy of the moving electron via the relativistic energy-momentum relation. The last equation is a restatement of conservation of momentum (after squaring).

$$E_1 + E_r - E_2 - E_m = 0 \quad (88)$$

$$E_1 - hf_1 = 0 \quad (89)$$

$$E_2 - hf_2 = 0 \quad (90)$$

$$p_1c - hf_1 = 0 \quad (91)$$

$$p_2c - hf_2 = 0 \quad (92)$$

$$\lambda_1 f_1 - c = 0 \quad (93)$$

$$\lambda_2 f_2 - c = 0 \quad (94)$$

$$E_r - m_e c^2 = 0 \quad (95)$$

$$E_m^2 - pe_2^2 c^2 - m_e^2 c^4 = 0 \quad (96)$$

$$pe_2^2 - p_1^2 - p_2^2 + 2p_1 p_2 \cos \theta = 0 \quad (97)$$

If we assume that the only measured variables are $h, c, m_e, \lambda_1, \lambda_2$ and $\cos \theta$, AI Hilbert provides the relationship

$$h^2 c^2 - h^2 c^2 \cos^2 \theta - \lambda_2 h m_e c^3 + \lambda_1 h m_e c^3 - \lambda_2 h m_e c^3 \cos \theta + \lambda_1 h m_e c^3 \cos \theta = 0,$$

which becomes, after rearranging terms

$$(\lambda_2 - \lambda_1) h m_e c^3 (1 + \cos \theta) = h^2 c^2 (1 - \cos^2 \theta).$$

Dividing both sides by $h m_e c^3 (1 + \cos \theta)$, we get the desired relationship. An important assumption we made above is that measurements of $\cos \theta$ are available and not just of θ . This is reasonable in many areas of physics as we often take the component of velocity, force etc. in a direction at an angle θ to the direction of the velocity, force etc. It is also reasonable to assume that h, c, m_e are measured quantities, as these are known quantities.