AI-Newton: A concept-driven physical law discovery system without prior physical knowledge

You-Le Fang,* Dong-Shan Jian,† and Xiang Li‡
School of Physics, Peking University, Beijing 100871, China

Yan-Qing Ma§
School of Physics, Peking University, Beijing 100871, China
and Center for High Energy Physics, Peking University, Beijing 100871, China

(Received 8 December 2025; revised 15 June 2026; accepted 9 October 2026; published 4 November 2026)

While current artificial intelligence (AI)-driven methods excel at deriving empirical models from individual experiments, a significant challenge remains in uncovering the common fundamental physics that underlie these models—a task at which human physicists are adept. To bridge this gap, we introduce AI-Newton, a framework for concept-driven scientific discovery. Our system autonomously derives general physical laws directly from raw, multiexperiment data, operating without supervision or prior physical knowledge. Its core innovations are twofold: (1) proposing interpretable physical concepts to construct laws and (2) progressively generalizing these laws to broader domains. Applied to a large, noisy dataset of mechanics experiments, AI-Newton successfully rediscovers foundational and universal laws, such as Newton's second law, the conservation of energy, and the universal gravitation. This work represents a proof-of-concept step toward autonomous, humanlike scientific discovery.

DOI: 10.1103/dvvc-4rsr

I. INTRODUCTION

For centuries, the ultimate goal of fundamental physics research has been to describe a wide range of phenomena through a small number of discovered laws. Advances in artificial intelligence (AI) have made AI-driven scientific discovery a highly promising new paradigm [1]. Although AI has achieved remarkable results in tackling domain-specific challenges [2,3], the ultimate aspiration from a paradigm-shifting perspective still lies in developing reliable AI systems capable of autonomous scientific discovery directly from a large collection of raw data without supervision [4,5].

Current approaches to automated physics discovery focus on individual experiments, employing either neural network (NN)-based methods [6–25] or symbolic techniques [26–33]. By analyzing data from a single experiment, these methods can construct a specific model capable of predicting future data from the same experiment; if sufficiently simple, such a model may even be expressed in symbolic form [34–36]. Although these methods represent a crucial and successful stage toward automated scientific discovery, they have not yet reached a discovery capacity comparable to that of human physicists.

Human scientists advance further by discerning common patterns across specific models from different experiments and, on that basis, formulating general models that account for data from all such experiments. For instance, Newtonian mechanics provides a unifying and interpretable framework by defining meaningful physical concepts and formulating general laws that are valid across diverse phenomena. Therefore, a central challenge for AI-driven scientific discovery is to evolve beyond descriptive data fitting for isolated experiments toward autonomous systems capable of extracting general theories.

In this article, we present AI-Newton, a concept-driven discovery system, which is designed for the critical question: How to extract concepts and general laws from problem-specific models. AI-Newton integrates an autonomous discovery workflow that is fundamentally built upon plausible reasoning and physical concepts. Given a collection of physical experiments, AI-Newton can gradually formulate a set of general laws applicable across a wide problem scope with neither supervision nor any prior physical knowledge. As a proof-of-concept implementation [37], we apply AI-Newton to classical mechanics. Across 46 diverse experiments, it successfully rediscovers Newton's second law, energy conservation, universal gravitation, and other foundational laws. This controlled setting provides a verifiable initial test of cross-experiment theory formation.

II. KNOWLEDGE BASE AND KNOWLEDGE REPRESENTATION

AI-Newton contains an experiment base and a theory base, as shown in Fig. 1. The experiment base stores physical experiments and corresponding simulated data generators.

*Contact author: eden@stu.pku.edu.cn
†Contact author: dsjian@stu.pku.edu.cn
‡Contact author: lix-PHY@pku.edu.cn
§Contact author: yqma@pku.edu.cn

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

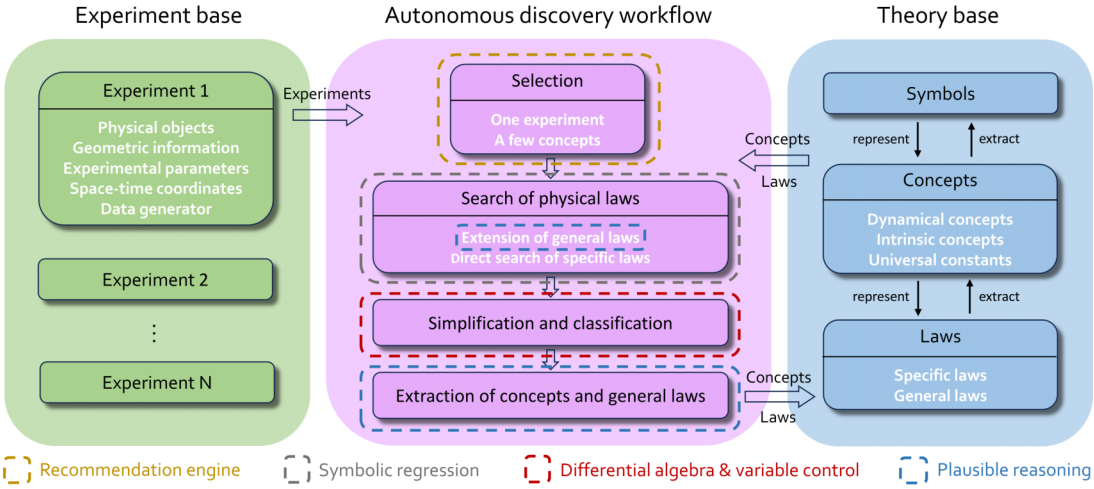


FIG. 1. AI-Newton's experiment base, theory base, and autonomous discovery workflow.

The inputs for each experiment include only the physical objects involved, geometric information, experimental parameters, and space-time coordinates, which define an experiment. To emphasize that no prior physical knowledge is used, all other concepts, such as mass or energy, are autonomously discovered in AI-Newton. The output of each experiment is simulated data with statistical errors.

The theory base stores physical knowledge explicitly in an interconnected library of symbols, concepts, and laws. This design mirrors how human physicists construct concise, universal laws from conceptual building blocks. In contrast to prior work, which interprets latent features in NNs as physical concepts [23,38,39], AI-Newton represents concepts and laws in an explicit, symbolic form. This greatly enhances interpretability and makes the acquired knowledge capable of transferring to new problems. Moreover, the introduction of powerful intermediate concepts allows complex physical laws to be expressed concisely, which in turn makes them more amenable to discovery through techniques like symbolic regression (SR) [26–36]. Initially, the concept layer contains only space-time coordinates; new concepts are autonomously defined and registered using a dedicated physical domain-specific language (DSL). (See Supplemental Material [40] for details.)

A robust knowledge representation is crucial because our goal is for the AI to discover generalizable knowledge across diverse systems, which requires transferring knowledge between different problems. To achieve this, we designed a physical DSL with a well-defined structure. This DSL not only formulates equations but also encodes the properties of physical objects and the relationships between physical quantities. For instance, given the known concepts of coordinate x and time t, the velocity of a ball can be defined in the DSL as

$$C_{\mathrm{id}} := \forall i\text{: Ball}, \quad \mathrm{d}x[i]/\mathrm{d}t, \tag{1}$$

where i indexes the balls and C_{id} denotes the symbol of velocity. In addition to dynamical concepts like velocity, the system also automatically identifies two other types: intrinsic concepts (e.g., mass, spring constant), which depend solely on specific physical objects, and universal constants (e.g., the gravitational constant), which are independent of all other quantities. Both are defined by documenting their measurement procedures. For example, mass of a ball could be defined as

$$C_{\mathrm{id}'} := \forall i\text{: Ball},$$
$$\text{Intrinsic } [\text{ExpName}(o_1 \to i, o_2 \to s), L[s] - L_0[s]], \tag{2}$$

where ExpName is the name of an experiment. In this experiment, the measured ball i is suspended from a fixed spring s, and the spring elongation $L[s] - L_0[s]$ serves as the measurement of the mass. Recording the measurement procedures of intrinsic concepts is essential, since it allows the value of an intrinsic property to be retrieved by invoking its defining experiment, ensuring conceptual consistency across different problems. A summary of representative concepts autonomously discovered and categorized within this framework is provided in Table I.

These explicit concepts serve as the building blocks for the laws layer, which stores discovered physical laws, such as conserved quantities and dynamical equations. The laws are categorized into specific laws (valid for one experiment with specific forms) and general laws (valid across diverse experiments with general forms). Within this framework, prior research in AI-driven physics discovery has concentrated on identifying specific laws. The introduction of general laws enables AI-Newton to simultaneously describe physics in various complex systems with compact and unified formulations. For instance, consider a system with a ball on an inclined plane connected to a fixed end via a spring. By applying the general law discovered by AI-Newton (Newton's second law

TABLE I. Representative concepts autonomously discovered by AI-Newton, categorized by their numerical dependencies.

Category	Dependence	Examples
Dynamical	Objects and experiments	v, T
Intrinsic	Specific objects only	m, k
Universal	Fully independent	G, g

in the x direction),

$$\forall i : \text{Ball},\ m_i a_{i,x} + (\nabla_i V_g)_x + (\nabla_i V_k)_x = 0, \tag{3}$$

the more complex dynamical equation of the ball can be concretely derived as

$$ma_x - \frac{c_x c_z}{c_x^2 + c_y^2 + c_z^2} mg + \frac{\left[\left(c_y^2 + c_z^2\right)x - c_x\left(c_y y + c_z z\right)\right]}{\left(c_x^2 + c_y^2 + c_z^2\right)L} k\Delta L = 0, \tag{4}$$

where (c_x, c_y, c_z) is the normal vector defining the inclined plane. For multiobject systems, concrete dynamical equations can be much more complex than the general laws, making them hard to be obtained using previous symbolic approaches. These cases highlight the efficacy of our concept-driven hierarchical approach.

III. AUTONOMOUS DISCOVERY WORKFLOW

The autonomous discovery workflow in AI-Newton continuously distills knowledge—expressed as physical concepts and laws—from experimental data, as shown in Fig. 1. Plausible reasoning, a method based on rational inference from partial evidence [41,42], is the key to discovering knowledge. Unlike deductive logic, it produces contextually reasonable rather than universally certain conclusions, mirroring scientific practice where hypotheses precede rigorous verification.

The workflow initiates each trial by selecting an experiment and a few concepts from the knowledge base. This selection is governed by a recommendation engine utilizing a value function inspired by the upper confidence bound (UCB) method [43–47] to balance exploration and exploitation. Both experiments e and concept combinations c are prioritized by a unified score

$$U_t(x) = p_\theta(x) Q_t(x) + \frac{1}{\sqrt{1 + N_t(x)}}, \tag{5}$$

where $x \in \{e, c\}$ and $N_t(x)$ is the selection count. The quality $Q_t(x)$ is exponentially updated by rewards r_t: For experiments, $r_t(e) = r_t^{\text{sl}} + r_t^{\text{gl}} + r_t^{\text{in}}$ measures the count of newly updated specific laws, general laws, and intrinsic properties; for concepts, a binary reward $r_t(c) \in \{0, 1\}$ indicates whether the combination yielded valid relations. The prior probability $p_\theta(x)$ is constantly 1 for experiments, while for concepts, $p_\theta(c)$ is predicted by an NN. The NN is dynamically updated via a time-decaying loss $\mathcal{L}(\theta) = \sum_j w_j (p_\theta(c_j) - \delta_j)^2$, learning from historical binary success labels δ_j. This mechanism enables the emulation of humanlike learning (see Supplemental Material [40] for details).

To ensure that the workflow establishes foundational knowledge before tackling complex experiments, we introduce an era-control strategy. Within a given era, every trial must conclude within a specific wall-clock time limit. If no new knowledge is acquired after a sufficient number of trials, the system advances to a new era with an exponentially increased time limit. Consequently, this strategy keeps the system focused on simpler experiments in the early phases.

The next step of each trial is to explore new laws from the selected experiment and concepts. For specific-law discovery, AI-Newton searches for relations by an SR method combining direct instantiation-verification and differential polynomial regression based on principal component analysis (PCA) [48–51]. The relations are given as

$$\sum_{i=1}^{n_B} c_i B_i - D = 0, \tag{6}$$

where c_i are numerical coefficients, B_i are rational combinations of concepts, and D is a constant. New general laws may also emerge by extending an existing one through plausible reasoning, which mirrors human research practice to heuristically leverage existing general laws. If a previously discovered general law is instantiated in the current experiment as $\mathcal{E}$ but fails to hold, AI-Newton attempts to modify it with a simple correction term:

$$\mathcal{E} + cB - D = 0. \tag{7}$$

For example, kinetic energy conservation is instantiated as $\mathcal{E} = T$, in a single-spring system, where T is kinetic energy, then AI-Newton may obtain

$$T + ckx^2 - D = 0, \tag{8}$$

which restores a conserved quantity. In both routes, candidate expressions are accepted only after passing rigorous statistical hypothesis testing (e.g., a χ^2 test with $\alpha = 0.05$).

The aforementioned process may generate redundant knowledge causing an explosion in both the theory base and search space that severely hinders continuous discovery under limited resources. To address this, AI-Newton simplifies physical laws into minimal representations in each trial. For the example shown in this article, we employ the Rosenfeld-Gröbner algorithm [52–55] from differential algebra to perform the simplification. During this reduction, expressions are evaluated using a recursive syntax-based complexity metric, which serves strictly as a lexicographical rule to prioritize simpler forms. Furthermore, through controlled-variable analysis, AI-Newton numerically identifies the dependencies of relations on physical objects and experimental parameters, using these dependencies as the basis for classification.

After identifying new laws, AI-Newton extracts new concepts from the processed results through plausible reasoning: A conserved quantity in the current experiment suggests broader utility, triggering its extraction as a new concept. Similarly, it proposes new general laws from directly searched specific laws that also hold in multiple other experiments. Crucially, AI-Newton autonomously determines the domain of applicability for each general law. A specific law is abstracted into a general law only if it mathematically holds across at least $n_{\text{exp}}^{\text{gl}}$ different experiments, where $n_{\text{exp}}^{\text{gl}}$ is a predefined hyperparameter. All accumulated knowledge is then updated to the theory base.

IV. REDISCOVERING LAWS OF NEWTONIAN MECHANICS

To evaluate AI-Newton's performance, we apply it to Newtonian mechanics problems, focusing on a set of 46 predefined experiments. These experiments involve three primary

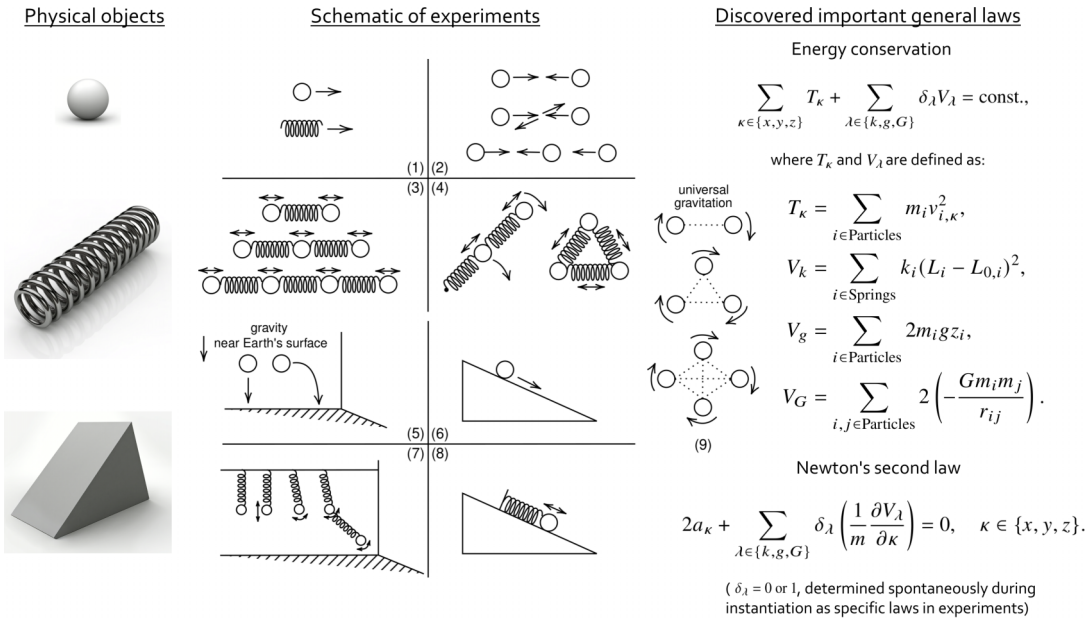


FIG. 2. Schematic of tested experiments and main general laws discovered. Some complex configurations are omitted for clarity. See text for details.

types of physical objects: balls (either small balls or celestial bodies), springs, and inclined planes. The experiments are designed to investigate both isolated and coupled systems, as illustrated in Fig. 2, including

(1) free motion of individual balls and springs;

(2) elastic collision of balls;

(3) coupled systems demonstrating translational vibrations, rotational oscillations, and pendulumlike motions;

(4) gravity-related problems, such as projectile motion and motion on inclined planes, along with complex spring-ball systems;

(5) celestial mechanics problems involving gravitational interactions.

The complexities of experiments are systematically increased by varying the number of physical objects and spatial dimensions, encompassing high-degree-of-freedom problems such as coupled oscillations of chained two-ball-two-spring systems on inclined planes, rotational dynamics of four-ball-four-spring systems, and other complex configurations. To evaluate the workflow under controlled statistical uncertainty, all test data are generated by solving differential equations and adding independent Gaussian errors. Under rigorous validation criteria, the relative noise level in these tests is set up to ∼1% (or ∼0.01% for experiments involving universal gravitation). Since the accumulation loop operates on validated laws rather than on raw data, these data-facing conditions are separable from the cross-experiment accumulation mechanism under study.

We evaluated the performance of our proof-of-concept implementation on an Intel Xeon Platinum 8370C (128 threads @ 3.500 GHz) platform with NVIDIA A40 GPU, configured with 64 cores for parallel processing. With max trials set to 1200 and an average run-time of 48 h, the system identified approximately 90 physical concepts and 50 general laws on average across the test cases. A catalog and classification are provided in Supplemental Material [40]. The discoveries include significant general laws such as energy conservation and Newton's second law along with their relevant concepts, as shown in Fig. 2, and together describe the tested systems from simple to high-degree-of-freedom configurations.

Statistical discovery progression on ten test cases is illustrated in Fig. 3, showing the timing distribution of important concept discoveries. This discovery progression exhibits an incremental pattern, where AI-Newton first explores simple concepts (e.g., mass) before advancing to more complex ones (e.g., force). For example, kinetic energy is first discovered in elastic collisions, after which plausible reasoning introduces elastic potential V_k, gravitational potential near Earth's surface

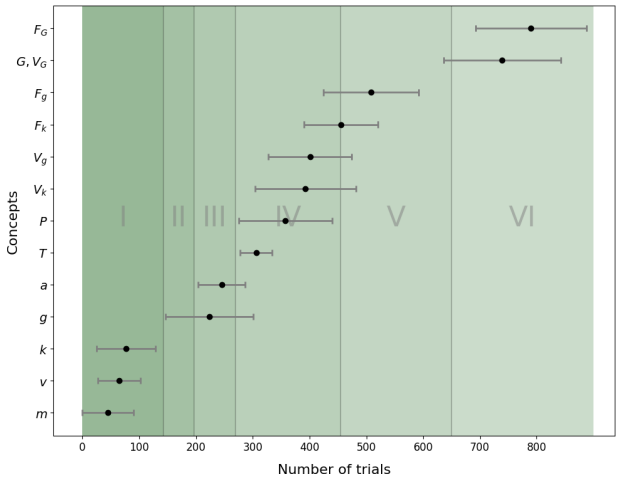


FIG. 3. Statistical analysis of concept discovery timing on ten test cases, recording the mean and standard deviation of discovery timings for key concepts. Number of trials means the number of analysis trial attempt has been done, not distinguishing which experiment. Roman numerals (I, II, ...) in the background indicate the eras defined by the era-control strategy.

V_g, and universal gravitational potential V_G when this conservation law fails in broader experiments, ultimately yielding energy conservation. The discovery of Newton's second law follows an analogous progression: It is first proposed in a simple experimental context and then generalized through plausible reasoning.

It is important to emphasize that the system is able to independently discover and unify fundamental concepts from disparate physical contexts. For instance, it can derive the concept of "mass" through two distinct experimental routes: from the static elongation of a spring under gravity (defining gravitational mass, m_g) and from the experiment of a horizontal spring-mass oscillation system (defining inertial mass, m_i). Critically, the system then autonomously verifies the numerical equivalence of m_g and m_i, effectively indicating a cornerstone of general relativity—the weak equivalence principle—from raw data alone. Furthermore, although symmetry is not presented as a concept in the discovery process, the discovered invariants and laws implicitly respect the underlying physical symmetries of the experimental data, like spatiotemporal translation invariance in isolated systems. Explicitly representing symmetry as concepts and incorporating it into general laws will be left for future work.

To further clarify the distinction from conventional symbolic regression, we performed task-aligned comparisons with representative single-dataset methods, including direct trajectory fitting (PySR and Operon) and derivative-based equation discovery (SINDy variants). In representative many-degree-of-freedom systems, these methods produce either experiment-specific fits with poor extrapolation or differential equations with unstable rollout. In contrast, the law instantiated from AI-Newton remains compact, interpretable, and predictive. Detailed comparative results are presented in Supplemental Material [40].

V. SUMMARY

We introduce AI-Newton, a framework designed for the autonomous construction of a transferable theoretical framework directly from raw, multiexperiment data, operating without supervision or prior physical knowledge. By extracting unifying concepts and general laws that span diverse systems, this approach moves beyond descriptive data fitting on isolated experiments. Our main contributions are based on plausible reasoning, enabling us to (1) propose physical concepts from the extracted laws and (2) extend an existing general law by adding new terms, thereby adapting it to describe a wider range of experiments. Introducing interpretable physical concepts allows discovered laws to remain concise, making them more tractable for SR to identify. Furthermore, iteratively constructing general laws from existing ones enables a gradual, scalable discovery process. Applied to a large, noisy dataset of mechanics experiments, this proof-of-concept implementation of AI-Newton successfully rediscovers foundational laws, including Newton's second law, the conservation of energy, and the law of universal gravitation. This work thus constitutes an initial step toward AI systems capable of contributing to frontier scientific research.

As a first step, we employ AI-Newton to rediscover known physical laws—a task where direct reliance on large language models (LLMs) can make it difficult to distinguish autonomous rediscovery from recall of existing knowledge. The symbolic-regression search interacts with the discovery loop through a defined interface—within each trial, it receives the processed data of the selected experiment together with the concepts recommended by the theory base and returns candidate relations expressed in the DSL, with acceptance decided by the validation rules external to the search—so it can be replaced by other symbolic-regression methods without altering the remaining workflow. In future applications to frontier science, however, the DSL, the recommendation engine, and the plausible reasoning components of the framework could be replaced or augmented by LLMs. This integration could provide the system with access to existing scientific knowledge, enabling a more informed and efficient discovery process.

ACKNOWLEDGMENTS

The funding agency is the National Natural Science Foundation of China, Grant No. 12325503 (Crossref Funder ID: 10.13039/501100001809). The High-performance Computing Platform of Peking University provided computing resources.

DATA AVAILABILITY

The data and source code supporting this study are openly available in the AI-Newton repository [37].

[1] Y. Xu, X. Liu, X. Cao, C. Huang, E. Liu, S. Qian, X. Liu, Y. Wu, F. Dong, C.-W. Qiu *et al.*, Artificial intelligence: A powerful paradigm for scientific research, Innovation **2**, 100179 (2021).

[2] 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 (London) **620**, 47 (2023).

[3] X. Zhang, L. Wang, J. Helwig, Y. Luo, C. Fu, Y. Xie, M. Liu, Y. Lin, Z. Xu, K. Yan, *et al.*, Artificial intelligence for science in quantum, atomistic, and continuum systems, Found. Trends Mach. Learn. **18**, 385 (2025).

[4] C. Lu, C. Lu, R. T. Lange, J. Foerster, J. Clune, and D. Ha, The AI Scientist: Towards fully automated open-ended scientific discovery, arXiv:2408.06292.

[5] C. K. Reddy and P. Shojaee, Towards scientific discovery with generative AI: Progress, opportunities, and challenges, in *Proceedings of the AAAI Conference on Artificial Intelligence* (AAAI Press, Washington, DC, 2025), Vol. 39, pp. 28601–28609.

[6] M. Schmidt and H. Lipson, Distilling free-form natural laws from experimental data, Science **324**, 81 (2009).

[7] S. Brunton, J. Proctor, and J. Kutz, Discovering governing equations from data: Sparse identification of nonlinear dynamical systems, Proc. Natl. Acad. Sci. USA **113**, 3932 (2015).

[8] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, Data-driven discovery of coordinates and governing equations, Proc. Natl. Acad. Sci. USA **116**, 22445 (2019).

[9] T. Wu and M. Tegmark, Toward an artificial intelligence physicist for unsupervised learning, Phys. Rev. E **100**, 033311 (2019).

[10] S. Greydanus, M. Dzamba, and J. Yosinski, *Hamiltonian Neural Networks* (Curran Associates Inc., Red Hook, NY, 2019).

[11] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, Lagrangian neural networks, in *ICLR 2020 Deep Differential Equations Workshop*, arXiv:2003.04630.

[12] B. M. De Silva, D. M. Higdon, S. L. Brunton, and J. N. Kutz, Discovery of physics from data: Universal laws and discrepancies, Front. Artif. Intell. **3**, 25 (2020).

[13] Z. Liu and M. Tegmark, Machine learning conservation laws from trajectories, Phys. Rev. Lett. **126**, 180604 (2021).

[14] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, Physics-informed machine learning, Nat. Rev. Phys. **3**, 422 (2021).

[15] Z. Liu, V. Madhavan, and M. Tegmark, Machine learning conservation laws from differential equations, Phys. Rev. E **106**, 045307 (2022).

[16] G. Camps-Valls, A. Gerhardus, U. Ninad, G. Varando, G. Martius, E. Balaguer-Ballester, R. Vinuesa, E. Diaz, L. Zanna, and J. Runge, Discovering causal relations and equations from data, Phys. Rep. **1044**, 1 (2023).

[17] C. Cornelio, S. Dash, V. Austel, T. R. Josephson, J. Goncalves, K. L. Clarkson, N. Megiddo, B. El Khadir, and L. Horesh, Combining data and theory for derivable scientific discovery with AI-Descartes, Nat. Commun. **14**, 1777 (2023).

[18] P. Lemos, N. Jeffrey, M. Cranmer, S. Ho, and P. Battaglia, Rediscovering orbital mechanics with machine learning, Mach. Learn.: Sci. Technol. **4**, 045002 (2023).

[19] Z. Liu, P. O. Sturm, S. Bharadwaj, S. J. Silva, and M. Tegmark, Interpretable conservation laws as sparse invariants, Phys. Rev. E **109**, L023301 (2024).

[20] R. Cory-Wright, C. Cornelio, S. Dash, B. E. Khadir, and L. Horesh, Evolving scientific discovery by unifying data and background knowledge with AI Hilbert, Nat. Commun. **15**, 5922 (2024).

[21] D. Zheng, V. Luo, J. Wu, and J. B. Tenenbaum, Unsupervised learning of latent physical properties using perception-prediction networks, in *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, Proceedings of Machine Learning Research Vol. R16 (PMLR, 2018), pp. 496–506.

[22] M. Tegmark, Latent representations of dynamical systems: When two is better than one, arXiv:1902.03364.

[23] R. Iten, T. Metger, H. Wilming, L. del Rio, and R. Renner, Discovering physical concepts with neural networks, Phys. Rev. Lett. **124**, 010508 (2020).

[24] B. Chen, K. Huang, S. Raghupathi, I. Chandratreya, Q. Du, and H. Lipson, Automated discovery of fundamental variables hidden in experimental data, Nat. Comput. Sci. **2**, 433 (2022).

[25] Q. Li, T. Wang, V. Roychowdhury, and M. K. Jawed, Metalearning generalizable dynamics from trajectories, Phys. Rev. Lett. **131**, 067301 (2023).

[26] S.-M. Udrescu and M. Tegmark, AI Feynman: A physics-inspired method for symbolic regression, Sci. Adv. **6**, eaay2631 (2020).

[27] 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, Adv. Neural Inf. Process. Syst. **33**, 4860 (2020).

[28] T. Bendinelli, L. Biggio, and P.-A. Kamienny, Controllable neural symbolic regression, in *International Conference on Machine Learning* (PMLR, 2023), pp. 2063–2077.

[29] W. Tenachi, R. Ibata, and F. I. Diakogiannis, Deep symbolic regression for physics guided by units constraints: Toward the automated discovery of physical laws, Astrophys. J. **959**, 99 (2023).

[30] Y. Tian, W. Zhou, M. Viscione, H. Dong, D. Kammer, and O. Fink, Interactive symbolic regression with co-design mechanism through offline reinforcement learning, Nat. Commun. **16**, 3930 (2025).

[31] M. Cranmer, Interpretable machine learning for science with PySR and SymbolicRegression.jl, arXiv:2305.01582.

[32] M. Du, Y. Chen, Z. Wang, L. Nie, and D. Zhang, Large language models for automatic equation discovery of nonlinear dynamics, Phys. Fluids **36**, 097121 (2024).

[33] B. Romera-Paredes, M. Barekatain, A. Novikov, M. Balog, M. P. Kumar, E. Dupont, F. J. R. Ruiz, J. S. Ellenberg, P. Wang, O. Fawzi, P. Kohli, and A. Fawzi, Mathematical discoveries from program search with large language models, Nature (London) **625**, 468 (2024).

[34] M. Valipour, B. You, M. Panju, and A. Ghodsi, SymbolicGPT: A generative transformer model for symbolic regression, in *Efficient Natural Language and Speech Processing (NeurIPS 2022 Workshop)* (NeurIPS, 2022).

[35] X. Chu, H. Zhao, E. Xu, H. Qi, M. Chen, and H. Shao, Scalable neural symbolic regression using control variables, arXiv:2306.04718.

[36] S. Mežnar, S. Džeroski, and L. Todorovski, Efficient generator of mathematical expressions for symbolic regression, Mach. Learn. **112**, 4563 (2023).

[37] Y.-L. Fang, D.-S. Jian, X. Li, and Y.-Q. Ma, AI-Newton [software and data], GitHub, 2025, https://github.com/Science-Discovery/AI-Newton.

[38] C. Wang, H. Zhai, and Y.-Z. You, Emergent Schrödinger equation in an introspective machine learning architecture, Sci. Bull. **64**, 1228 (2019).

[39] B.-B. Li, Y. Gu, and S.-F. Wu, Discover physical concepts and equations with machine learning, Phys. Scr. **100**, 086011 (2025).

[40] See Supplemental Material at http://link.aps.org/supplemental/10.1103/dvvc-4rsr for additional details on physical domain-specific language, recommendation engine, and computational engine.

[41] G. Pólya, *Mathematics and Plausible Reasoning: Induction and Analogy in Mathematics* (Princeton University Press, Princeton, NJ, 1990), Vol. 1.

[42] G. Pólya, *Mathematics and Plausible Reasoning: Patterns of Plausible Inference* (Princeton University Press, Princeton, NJ, 1990), Vol. 2.

[43] T. Lai and H. Robbins, Asymptotically efficient adaptive allocation rules, Adv. Appl. Math. **6**, 4 (1985).

[44] T. L. Lai, Adaptive treatment allocation and the multi-armed bandit problem, Ann. Stat. **15**, 1091 (1987).

[45] R. Agrawal, Sample mean based index policies by $O(\log n)$ regret for the multi-armed bandit problem, Adv. Appl. Probab. **27**, 1054 (1995).

[46] M. N. Katehakis and H. Robbins, Sequential choice from several populations, Proc. Natl. Acad. Sci. USA **92**, 8584 (1995).

[47] P. Auer, Using confidence bounds for exploitation-exploration trade-offs, J. Mach. Learn. Res. **3**, 397 (2002).

[48] K. Pearson, LIII. On lines and planes of closest fit to systems of points in space, London, Edinburgh Dublin Philos. Mag. J. Sci. **2**, 559 (1901).

[49] L. Wang, Discovering phase transitions with unsupervised learning, Phys. Rev. B **94**, 195105 (2016).

[50] H. Kiwata, Deriving the order parameters of a spin-glass model using principal component analysis, Phys. Rev. E **99**, 063304 (2019).

[51] D. Yevick, Conservation laws and spin system modeling through principal component analysis, Comput. Phys. Commun. **262**, 107832 (2021).

[52] F. Boulier, D. Lazard, F. Ollivier, and M. Petitot, Representation for the radical of a finitely generated differential ideal, in *Proceedings of the 1995 International Symposium on Symbolic and Algebraic Computation*, ISSAC '95 (Association for Computing Machinery, New York, 1995), pp. 158–166.

[53] F. Boulier, D. Lazard, F. Ollivier, and M. Petitot, Computing representations for radicals of finitely generated differential ideals, applicable algebra in engineering, Commun. Comput. **20**, 73 (2009).

[54] Maplesoft, Different algebra in Maple, Maplesoft Help Center, 2024, https://cn.maplesoft.com/support/help/Maple/view.aspx?path=DifferentialAlgebra.

[55] Maplesoft, Maple 2024, https://www.maplesoft.com/products/maple/.