This is the accepted manuscript made available via CHORUS. The article has been published as:

Interpretable conservation laws as sparse invariants

Ziming Liu, Patrick Obin Sturm, Saketh Bharadwaj, Sam J. Silva, and Max Tegmark

Phys. Rev. E **109**, L023301 — Published 27 February 2024

DOI: 10.1103/PhysRevE.109.L023301

Discovering New Interpretable Conservation Laws as Sparse Invariants

Ziming Liu,[1] Patrick Obin Sturm,[2] Saketh Bharadwaj,[3] Sam J. Silva,[2] and Max Tegmark[1]

[1] *Department of Physics, Institute of Artificial Intelligence and Fundamental Interactions, Massachusetts Institute of Technology, Cambridge, USA*

[2] *Department of Earth Sciences, University of Southern California, Los Angeles, USA*

[3] *Department of Chemical Engineering, Indian Institute of Technology, Hyderabad, India*

(Dated: February 5, 2024)

Discovering conservation laws for a given dynamical system is important but challenging. In a *theorist* setup (differential equations and basis functions are both known), we propose the **S**parse **I**nvariant **D**etector (SID), an algorithm that auto-discovers conservation laws from differential equations. Its algorithmic simplicity allows robustness and interpretability of the discovered conserved quantities. We show that SID is able to rediscover known and even discover new conservation laws in a variety of systems. For two examples in fluid mechanics and atmospheric chemistry, SID discovers 14 and 3 conserved quantities, respectively, where only 12 and 2 were previously known to domain experts.

I. INTRODUCTION

Conservation laws are important concepts in physics, yet discovering them is challenging. Ideally, the set of discovered conserved quantities should be *complete*, *independent* and *interpretable*. Although several attempts have been made to automate the discovery process with machine learning [1–8], their complicated setups and blackbox nature make it hard to guarantee all these desirable properties. This paper considers a simple yet realistic setup where all these desirable properties can be met.

"Discovering conservation laws" can mean wildly different things for *experimentalists*, *computationalists* and *theorists*, as shown in Table I. Most prior work [1–6] takes on the experimentalist setup, assuming knowledge of neither the differential equations nor the form of conservation laws. [7] takes the computationalist setup, assuming knowledge of differential equations. This work explores the theorist setup, where both differential equations and basis functions of conservation laws are known. Admittedly, this setup is simpler than the other two, but is still realistic when theorists have the differential equations at hand and have educated guesses about the basis functions that may span the conserved quantities.

We propose **S**parse **I**nvariant **D**etector (SID), an algorithm that reveals conservation Laws. SID is incredibly simple in the sense that it only requires linear algorithms (except for sparsification), so the results are much more trustworthy and interpretable than blackbox machine learning methods. Note that SID does not replace us human scientists, but rather acts as a helpful assistant: while humans need to input basis functions (i.e., *formulating* hypotheses) to SID, SID is good at computing conserved quantities (i.e., *testing* hypotheses) based on the given prompt. In this manner, human scientists can focus on the more creative part of the job, while SID does the technical and tedious work. This paper gives two examples where new conserved quantities are successfully discovered by SID: one in fluid mechanics, and another in atmospheric chemistry (see Table II). In the former one, although the new conserved quantities are somewhat expected in hindsight, humans alone may need several more months to find them. In the latter one, a new conserved quantity is found, which was unintended in the design of the model.

TABLE I: Three setups of conservation law discovery

Setup	Experimentalist	Computationalist	Theorist
Model-based	No	Yes	Yes
Known basis	No	No	Yes
Independence	Partial	Yes	Yes
Completeness	No	Partial	Yes
Interpretability	Partial	Partial	Yes
Reference	[1–6]	[7]	This work

TABLE II: The number of conserved quantities known to experts and discovered by SID

	Fluid (2D)	Fluid (3D)	Atmosphere
Known	**8**	12	2
SID	**8** (simpler)	**14**	**3**

II. METHOD

A. Problem setup

We consider a first-order differential equation $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, where $\dot{\mathbf{x}} \equiv \frac{d\mathbf{x}}{dt}$, $\mathbf{x} \equiv (x_1, \cdots, x_d) \in \mathbb{R}^d$ is the state vector and $\mathbf{f} : \mathbb{R}^d \to \mathbb{R}^d$ is a vector field [9]. This ODE formulation is more general than it seems: (1) Hamiltonian systems are subsumed as $\mathbf{x} \equiv (\mathbf{x}', \mathbf{p}')$; (2) Higher-order differential equations (e.g., $\ddot{\mathbf{y}} = \mathbf{f}(\mathbf{y})$) are included as $\mathbf{x} \equiv (\mathbf{y}, \dot{\mathbf{y}}, \cdots)$; (3) Partial differential equations (PDEs) become ODEs once discretized.

A *conserved quantity* (short as CQ) is a scalar function

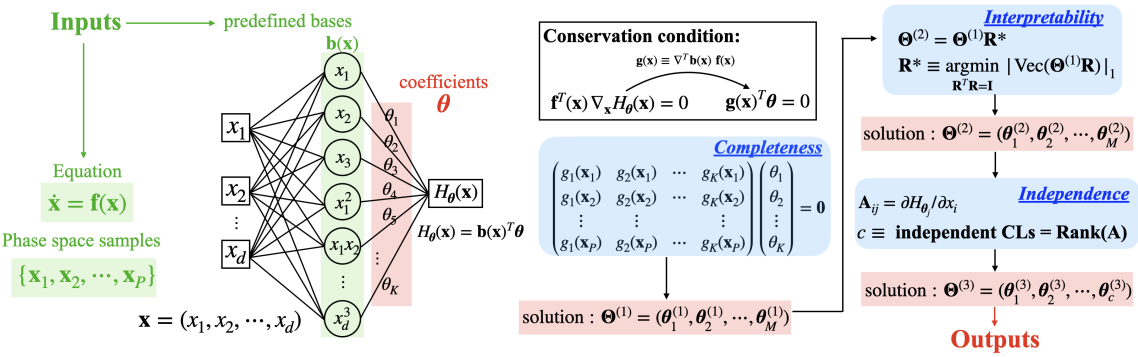


FIG. 1: SID workflow: **Inputs** are differential equations, basis functions and sample points; **Outputs** are a set of conserved quantities which are complete, independent and interpretable.

$H(\mathbf{x}) : \mathbb{R}^d \to \mathbb{R}$, such that its value remains constant along any trajectory obeying $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ [10]. As proved in [7], a necessary and sufficient condition for $H(\mathbf{x})$ being a conserved quantity is $\nabla H(\mathbf{x}) \cdot \mathbf{f}(\mathbf{x}) = 0$, since

$$0 = \dot{H} = \nabla H(\mathbf{x}) \cdot \dot{\mathbf{x}} = \nabla H(\mathbf{x}) \cdot \mathbf{f}(\mathbf{x}). \tag{1}$$

Given the differential equation $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, we hope to find a set of conserved quantities $\{H_1, \cdots, H_c\}$ which satisfies these three properties:

- *Independence*: they are functionally independent. None of them can be written as (possibly) nonlinear combinations of others, i.e., $g(H_1, \cdots, H_c) = 0 \Rightarrow g = 0$.
- *Completeness*: any conserved quantity H (in the function space spanned by basis functions) can be expressed by them, i.e., there exists g such that $H = g(H_1, \cdots, H_c)$.
- *Interpretability*: conserved quantities can be written as (hopefully simple) symbolic formulas.

B. Solving the linear equation and completeness

The prior work [7] parametrizes the conserved quantities $H_{\boldsymbol{\theta}}(\mathbf{x})$ as neural networks and learns the parameters $\boldsymbol{\theta}$ to make $|\nabla H_{\boldsymbol{\theta}}(\mathbf{x}) \cdot \mathbf{f}(\mathbf{x})|^2$ close to zero. However, neural network training may get stuck at local minima, so the results are not reliable. Moreover, the parameterized conserved quantities are not immediately interpretable.

We consider a simpler setup. Assume that we know $H_{\boldsymbol{\theta}}(\mathbf{x})$ to be a linear combination of K predefined basis functions $b_i(\mathbf{x})$ $(1 \le i \le K)$ such that

$$H_{\boldsymbol{\theta}}(\mathbf{x}) = \sum_{i=1}^{K} \theta_i b_i(\mathbf{x}) \equiv \boldsymbol{\theta} \cdot \mathbf{b}(\mathbf{x}), \tag{2}$$

where only $\boldsymbol{\theta} \in \mathbb{R}^K$ are learnable parameters to be determined and the vector $\mathbf{b} : \mathbb{R}^d \to \mathbb{R}^K$ defines the basis functions. Since the number of conserved quantities can exceed one, we define a set of parameters $\boldsymbol{\Theta} \equiv \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \cdots\}$ and their corresponding functions $H_{\boldsymbol{\Theta}} \equiv \{H_{\boldsymbol{\theta}} | \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$. As shown in FIG. 1, Eq. (2) is equivalent to a neural network whose last linear layer contains the only trainable parameters. With Eq. (2), the conservation condition Eq. (1) becomes:

$$\mathbf{g}(\mathbf{x})^T\boldsymbol{\theta} = 0, \quad \mathbf{g}(\mathbf{x}) \equiv (\nabla \mathbf{b}(\mathbf{x}))\mathbf{f}(\mathbf{x}), \tag{3}$$

which is a linear equation of $\boldsymbol{\theta}$. Remember that in our setup, both $\mathbf{b}(\mathbf{x})$ and $\mathbf{f}(\mathbf{x})$ are known, so $\mathbf{g}(\mathbf{x}) \equiv (\nabla \mathbf{b}(\mathbf{x}))\mathbf{f}(\mathbf{x})$ is known as well. In practice, we draw P random points $\mathbf{x}_i$ $(1 \le i \le P)$ from phase space. A solution $\boldsymbol{\theta}$ should make Eq. (3) hold for all $\mathbf{x}_i$, or more explicitly,

$$\underbrace{\begin{pmatrix} g_1(\mathbf{x}_1) & g_2(\mathbf{x}_1) & \cdots & g_K(\mathbf{x}_1) \\ g_1(\mathbf{x}_2) & g_2(\mathbf{x}_2) & \cdots & g_K(\mathbf{x}_2) \\ \vdots & \vdots & & \vdots \\ g_1(\mathbf{x}_P) & g_2(\mathbf{x}_P) & \cdots & g_K(\mathbf{x}_P) \end{pmatrix}}_{\mathbf{G}} \underbrace{\begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_K \end{pmatrix}}_{\boldsymbol{\theta}} = \mathbf{0}, \tag{4}$$

which is simply linear regression. In practice, we apply singular value decomposition to $\mathbf{G} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T$, where $\mathbf{U} \in \mathbb{R}^{P\times P}$ and $\mathbf{V} \in \mathbb{R}^{K\times K}$ are orthogonal matrices, $\boldsymbol{\Sigma} \in \mathbb{R}^{P\times K}$ is diagonal with singular values $0 \le \sigma_1 \le \sigma_2 \le \cdots$. We count σ_i as effectively zero if $\sigma_i < \epsilon \equiv 10^{-8}$. The number of vanishing singular values, denoted M, is equal to the dimensionality of the solution space (null space), which is spanned by the first M columns of $\mathbf{V}^T$, denoted $\boldsymbol{\Theta}^{(1)} \equiv (\boldsymbol{\theta}_1^{(1)}, \boldsymbol{\theta}_2^{(1)}, \cdots, \boldsymbol{\theta}_M^{(1)}) \in \mathbb{R}^{K\times M}$. The linear structure obviously gives **completeness** (in the space spanned by basis functions), since any solution $\boldsymbol{\theta}$ can be expressed as a linear combination of columns of $\boldsymbol{\Theta}^{(1)}$.

C. Interpretability

In order to gain more interpretability, we want $\boldsymbol{\Theta}^{(1)}$ to be sparse. Note that if $\mathbf{R} \in \mathbb{R}^{M\times M}$ is an orthogonal matrix, the columns of $\boldsymbol{\Theta}^{(2)} = \boldsymbol{\Theta}^{(1)}\mathbf{R}$ also form a set of complete and orthogonal solutions. Therefore we can encourage sparsity by finding and applying the orthogonal

matrix that minimizes the following:

$$\mathbf{R}^* = \underset{\mathbf{R}^T\mathbf{R}=\mathbf{I}}{\operatorname{argmin}} ||\mathbf{\Theta}^{(1)}\mathbf{R}||_1, \quad \mathbf{\Theta}^{(2)} = \mathbf{\Theta}^{(1)}\mathbf{R}^*, \tag{5}$$

where $||\mathbf{M}||_1 \equiv \sum_{ij} |M_{ij}|$ denotes the L_1-norm of a matrix $\mathbf{M}$, encouraging sparsity.

D. Independence

Although columns of $\mathbf{\Theta}^{(2)}$ are linearly independent, $H_{\mathbf{\Theta}^{(2)}}$ are not guaranteed to be functionally independent. Take the 1D harmonic oscillator $\mathbf{x} = (x, p)$, for example. Restricting basis functions to be polynomials in x and p up to the 4th order, there are two solutions:

$$H_{\boldsymbol{\theta}_1} = x^2 + p^2, H_{\boldsymbol{\theta}_2} = H^2_{\boldsymbol{\theta}_1} = x^4 + 2x^2p^2 + p^2, \tag{6}$$

where $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are orthogonal (hence independent), but $H_{\boldsymbol{\theta}_2} = H^2_{\boldsymbol{\theta}_1}$, so they are not functionally independent. Consequently, we want a subset of $\mathbf{\Theta}^{(2)}$, denoted $\mathbf{\Theta}^{(3)}$, such that $H_{\mathbf{\Theta}^{(3)}}$ is both independent and complete (i.e., can generate $H_{\mathbf{\Theta}^{(2)}}$). The first question is: how many elements, denoted c, does $\mathbf{\Theta}^{(3)}$ have? As shown in [7], c is equal to the rank of the following matrix:

$$\mathbf{A} = \begin{pmatrix} \frac{\partial H_{\boldsymbol{\theta}_1}}{\partial x_1} & \frac{\partial H_{\boldsymbol{\theta}_2}}{\partial x_1} & \cdots & \frac{\partial H_{\boldsymbol{\theta}_M}}{\partial x_1} \\ \frac{\partial H_{\boldsymbol{\theta}_1}}{\partial x_2} & \frac{\partial H_{\boldsymbol{\theta}_2}}{\partial x_2} & \cdots & \frac{\partial H_{\boldsymbol{\theta}_M}}{\partial x_2} \\ \vdots & \vdots & & \vdots \\ \frac{\partial H_{\boldsymbol{\theta}_1}}{\partial x_d} & \frac{\partial H_{\boldsymbol{\theta}_2}}{\partial x_d} & \cdots & \frac{\partial H_{\boldsymbol{\theta}_M}}{\partial x_d} \end{pmatrix} \tag{7}$$

which hinges on the fact that **gradients of functionally dependent functions are linearly dependent** [11]. In practice, applying singular value decomposition to $\mathbf{A}$ gives $\mathbf{A} = \mathbf{U}'\mathbf{\Sigma}'\mathbf{V}'^T$, where $\mathbf{U}' \in \mathbb{R}^{d\times d}$ and $\mathbf{V}' \in \mathbb{R}^{M\times M}$ are orthogonal matrices, and $\mathbf{\Sigma}'$ is a diagonal matrix with singular values $s_1 \geq s_2 \geq \cdots \geq 0$. We count s_i as effectively non-zero if $s_i > \epsilon = 10^{-8}$. The number of non-zero singular values is equal to rank($\mathbf{A}$), which is in turn equal to c. After determining c, we aim to obtain $\mathbf{\Theta}^{(3)}$ by selecting c elements from $\mathbf{\Theta}^{(2)}$. The selection process is as follows: (1) We assign each conserved quantity a complexity score (based on entropy [12]) and sort them from the simplest to the most complex. (2) Starting from an empty set $\mathbf{\Theta}^{(3)}$, looping over element $\boldsymbol{\theta} \in \mathbf{\Theta}^{(2)}$, we add $\boldsymbol{\theta}$ to $\mathbf{\Theta}^{(3)}$ if $H_{\boldsymbol{\theta}}$ is independent of $H_{\mathbf{\Theta}^{(3)}}$ (functions already added), until $\mathbf{\Theta}^{(3)}$ contains c elements.

III. RESULTS

To better illustrate SID, we apply it to three dynamical systems.

A. Systems biology

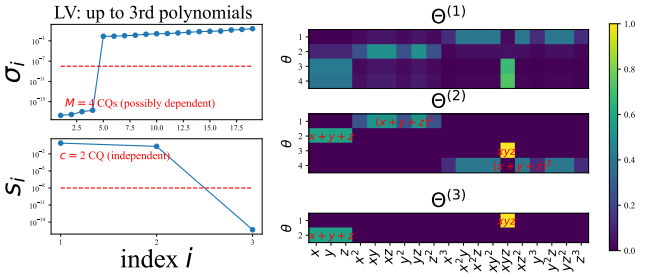
Our first application is from systems biology. The Lotka–Volterra equations (LV hereafter) describe how population of many species evolve in time via inter-species interactions. We study this particular equation:

$$\dot{x} = x(y - z), \dot{y} = y(z - x), \dot{z} = z(x - y) \tag{8}$$

with two known conserved quantities $H_1 = x + y + z$ and $H_2 = xyz$. We define basis functions to be all polynomials up to 3rd order (including $K = 19$ terms, shown in FIG. 2b). We draw $P = 100$ data points from the standard Gaussian distribution, i.e., $\mathbf{x}_i \equiv (x_i, y_i, z_i) \sim \mathcal{N}(0, \mathbf{I}_{3\times3})$ $(1 \leq i \leq P)$.

Within the function space spanned by the basis functions, $M = 4$ conserved quantities are found, since FIG. 2a top shows that there are 4 vanishing singular values σ_i. The coefficients of four CQs are somewhat mixed, shown in FIG. 2b top. After sparsification (Eq. (5)), the coefficients become less entangled, shown in FIG. 2b middle, although their represented conserved quantities are still dependent. Among the 4 CQs, only $c = 2$ are independent, since FIG. 2a bottom shows that there are 2 non-vanishing singular values s_i. FIG. 2b bottom shows the final outputs: the two conserved quantities agree with our prior knowledge.

What will happen if we choose the set of basis functions to be smaller or larger? (1) smaller: if we include polynomials up to the 1st or the 2nd order ($K = 3$ or $K = 9$) only, then only H_1 can be discovered by SID, while H_2 is not discovered. (2) larger: if we instead include polynomials up to order 4,5,6 ($K = 34, 55, 83$), then both H_1 and H_2 are still discovered, although the sparsification and independence process may take longer than with only 3rd order polynomials. Detailed results are included in Appendix B. The take-home message is that SID does not replace human scientists since it requires the input of basis functions (formulate hypothesis) from human scientists. SID is good at testing hypotheses (which could be technical and tedious), however, it is human scientists who formulate hypotheses (which requires creativity).



(a) (b)

FIG. 2: Three-species Lokta-Volterra equation. SID correctly discovers that: (a) there are 4 CQs (top) in polynomials up to 3rd order, yet only 2 of them (bottom) are independent. (b) coefficients of conserved quantities $\mathbf{\Theta}^{(i)}(i = 1, 2, 3)$, with more interpretability.

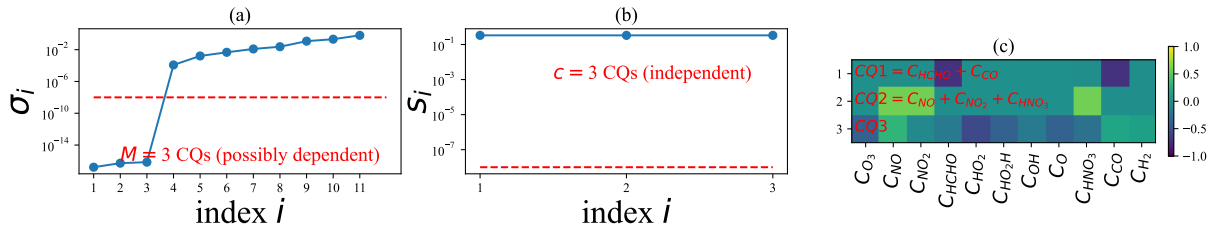


FIG. 3: (a)(b) SID discovers 3 independent conserved quantities in the ozone photochemical production model. (c) The coefficients of these linear conserved quantities. The first two correspond to the known carbon and nitrogen conservation, while CQ_3 is identified for the first time.

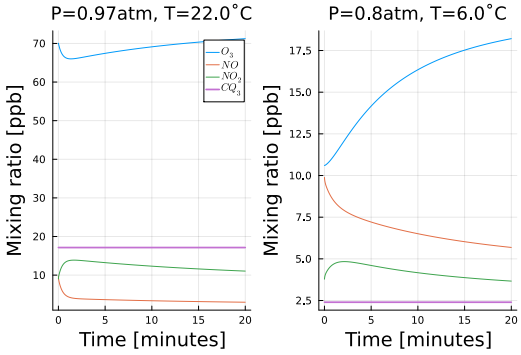


FIG. 4: The evolution of concentrations in simulation. Under both conditions, CQ_3 is well conserved.

B. Fluid mechanics

Arguably the biggest puzzle in fluid mechanics is turbulence [13, 14]. Turbulence, and chaos in general, are due to lack of sufficient conserved quantities. For example, a two-body problem is regular by having 8 degrees of freedom (dof) and 7 CQs, while a three-body problem is chaotic with 12 dofs but only 6 CQs. Therefore, studying conserved quantities of fluid systems is relevant to understanding turbulence. As a preliminary step, we study conserved quantities of a fluid element in ideal fluid (zero viscosity and incompressible). In 2D (3D), The fluid element is a triangle (tetrahedron), which is represented by its 3 (4) vertices [15]. Effectively, we can view the system as 3 (4) "free" particles, with the only constraint being that the area (volume) of the triangle (tetrahedron) should remain unchanged. The equations of motion are included in Appendix D, which appear a bit intimidating (especially for 3D).

Fluid dynamics experts (including some authors of this paper) have attempted to find the conserved quantities with pencil and paper. Given the complexity of the calculations, it is impressive that they found 8 (12) conserved quantities for 2D (3D). However, they were unsure whether there were more undiscovered conserved quantities, and whether the discovered ones were in their simplest. So we turn to SID for help. The results below are for the basis function set selected to be polynomials up to 2nd (3rd) order for 2D (3D), but more polynomial orders are also tried in Appendix D.

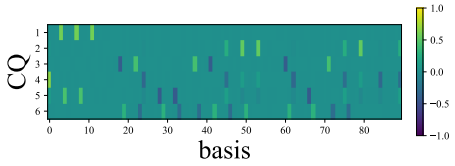


FIG. 5: SID's discovered conserved quantities (coefficients) for the simplified fluid system (2D) with the basis set chosen to be polynomials up to order 2.

For the 2D case, SID finds 8 conserved quantities, agreeing with experts' expectation. Interestingly, the conserved quantities found by SID appear to be simpler. In fact, all the conserved quantities discovered by SID are 1st or 2nd order polynomials (see FIG. 5), while experts found a 4th order polynomial, which later turned out to be a combination of two 2nd order are conserved quantities discovered by SID (see Appendix D for details).

For the 3D case, SID finds 14 conserved quantities, while experts found only 12. The two new conserved quantities can be interpreted as the angular momentum in the center of mass (COM) frame. They are non-trivial because it is easy to (falsely) think that the COM angular momentum is dependent on the angular momentum and the linear momentum [16]. Although humans alone will probably get the results right at the end of the day without SID, SID can take care of subtle details automatically, thus saving human experts' mental labor to a great extent.

C. Atmospheric chemistry

We next apply SID to a truncated atmospheric chemistry model of photochemical ozone production [17], where an exotic new conserved quantity is found. This simplified dynamical system contains 11 species and 10 reactions involved in ozone formation, including NOx, organic, and radical chemistry [18]. A key characteristic of this system is conservation of carbon and nitrogen atoms, H_C and H_N, respectively. Though species in this

model contain two other elements, hydrogen and oxygen, neither are conserved, as H_2O molecules are not one of the 11 species whose concentrations are tracked, and diatomic oxygen O_2 is treated as an infinite source and sink due to its abundance. H_C and H_N are implied in the coefficients of a stoichiometric matrix $\mathbf{B} \in \mathbb{Z}^{11,10}$ used in prior work to enforce conservation of atoms in machine learning surrogate models [17]. H_C and H_N can be represented by linear combinations of species concentrations, the coefficients of which form a basis for the null space of $\mathbf{B}^T$. Further details are provided in Appendix C.

We applied SID to simulation trajectories expecting to discover up to 2 conserved quantities which are linear combinations of concentrations. The training data are points on simulation trajectories at pressure $P = 0.95$ atm and temperature $T = 20.0°\text{C}$. As shown in FIG. 3, besides H_C and H_N, SID surprisingly discovers a third conserved quantity CQ_3 that is a linear combination of species concentrations (C_X means the concentration of X):

$$\begin{aligned} CQ_3 \approx\ & 6C_{O_3} - 5C_{NO} + C_{NO_2} + 3C_{HCHO} \\ & + 9C_{HO_2} + 6C_{HO_2H} + 2C_{OH} + 6C_O \\ & + 4C_{HNO_3} - 3C_{CO} - 2.21C_{H_2} \end{aligned} \tag{9}$$

This additional quantity is linearly independent of H_C and H_N and is not in the null space of $\mathbf{B}^T$. CQ_3 has a relative variation of less than 0.1% in 995 of 1000 simulated cases. Two representative simulation trajectories are shown in Figure 4, where CQ_3 holds under different chemical and meteorological conditions. The evolving concentrations of O_3, NO, and NO_2 are included as contrasts to the invariance of CQ_3. We have not yet identified the underlying cause of CQ_3, and whether it is physically exact or numerically approximate. We have ruled out symmetry corresponding to hydrogen conservation: when explicitly incorporating production of H_2O as an additional buildup species, SID identifies approximate hydrogen conservation as well as a fourth conserved quantity (see Appendix C). This implies that CQ_3 might be a non-trivial conserved quantity that is worth thorough study in future work.

IV. CONCLUSIONS

We have presented an algorithm SID to automatically discover conserved quantities from dynamical equations. In constrast to previous blackbox models, SID is guaranteed to be robust and interpretable thanks to its algorithmic simplicity. We demonstrate the power of SID on two examples in atmospheric chemistry and fluid mechanics, revealing new conserved quantities hitherto unknown to human experts. Although SID does not replace human scientists, it is a helpful assistant that can facilitate the discovery process. Promising future directions include applying SID to a broader range of applications, e.g., large-scale molecule dynamics. We discussed the possiblities in Appendix A.

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

[2] Y.-i. Mototake, Interpretable conservation law estimation by deriving the symmetries of dynamics from trained deep neural networks, Physical Review E **103**, 033303 (2021).

[3] S. J. Wetzel, R. G. Melko, J. Scott, M. Panju, and V. Ganesh, Discovering symmetry invariants and conserved quantities by interpreting siamese neural networks, Physical Review Research **2**, 033499 (2020).

[4] S. Ha and H. Jeong, Discovering invariants via machine learning, Phys. Rev. Res. **3**, L042035 (2021).

[5] S. Arora, A. Bihlo, R. Brecht, and P. Holba, Model-free machine learning of conservation laws from data, arXiv preprint arXiv:2301.07503 (2023).

[6] M. F. Kasim and Y. H. Lim, Constants of motion network, arXiv preprint arXiv:2208.10387 (2022).

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

[8] E. Kaiser, J. N. Kutz, and S. L. Brunton, Discovering conservation laws from data for control, in *2018 IEEE Conference on Decision and Control (CDC)* (IEEE, 2018) pp. 6415–6421.

[9] Our method can also handle non-autonomous systems, i.e., $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t)$, as long as basis functions also depend on time t.

[10] The constant itself can be different for different trajectories.

[11] Suppose H_1 is dependent on H_2, H_3, then there exists a function g such that $H_1 = g(H_2, H_3)$. Taking gradients on both es gives $\nabla H_1, \frac{\partial g}{\partial H_2}\nabla H_2 + \frac{\partial g}{\partial H_3}\nabla H_3$, which means ∇H_1 is linear combination of ∇H_2 and ∇H_3. So $\mathbf{A} = [\nabla H_1, \nabla H_2, \nabla H_3]$ has rank two (if H_2 and H_3 are independent).

[12] For a vector $\boldsymbol{\theta}$, we can associate it to a probability distribution $p_i \equiv |\theta_i|/(\sum_i |\theta_j|)$, whose entropy is $S = -\sum_i p_i \log p_i$.

[13] G. Falkovich and K. R. Sreenivasan, *Lessons from hydrodynamic turbulence*, Tech. Rep. (Abdus Salam International Centre for Theoretical Physics, 2006).

[14] P. Constantin, On the euler equations of incompressible fluids, Bulletin of the American Mathematical Society **44**, 603 (2007).

[15] A. Pumir, E. Bodenschatz, and H. Xu, Tetrahedron deformation and alignment of perceived vorticity and strain in a turbulent flow, Physics of Fluids **25**, 035101 (2013).

[16] Also, the COM angular momentum has three components in 3D, so why are there only two new conserved quantities, not three? It turns out that the third component is not independent of the 14 conserved quantities.

[17] P. O. Sturm and A. S. Wexler, Conservation laws in a neural network architecture: enforcing the atom balance

of a julia-based photochemical model (v0.2.0), Geoscientific Model Development **15**, 3417 (2022).

[18] P. O. Sturm and A. S. Wexler, A mass- and energy-conserving framework for using machine learning to speed computations: a photochemistry example, Geoscientific Model Development **13**, 4435 (2020).