Draft version April 8, 2026
Typeset using LaTeX **twocolumn** style in AASTeX631

Data-Driven Discovery of Governing Equations for a 3D Fluid System: Addressing Feature Collinearity in Sparse Regression

Denario[1]

[1] *Anthropic, Gemini & OpenAI servers. Planet Earth.*

ABSTRACT

This study addresses the challenge of discovering the underlying partial differential equations (PDEs) governing the spatial-temporal evolution of a physical system directly from observational data. We employed a comprehensive workflow on a dataset comprising three velocity components and a density field on a 128^3 periodic grid across 10 time slices. This workflow included exploratory data analysis, spectral noise filtering, robust estimation of spatial and temporal derivatives, and the construction of a rich library of candidate terms, followed by sparse regression with iterative thresholding to identify the governing equations. Exploratory analysis revealed complex, multi-scale spatial structures in the velocity fields and a remarkably uniform density field. The discovered equations accurately predicted instantaneous temporal derivatives, achieving R^2 values between 0.593 and 0.732 for velocity components and 0.362 for density. However, severe collinearity within the feature library led the sparse regression algorithm to exploit its null space, resulting in equations with numerous large, oppositely signed coefficients for composite physical operators and their constituent terms, thereby obscuring direct physical interpretability. Despite this complexity, rigorous forward-time integration of the identified PDEs, initialized from observed data, demonstrated exceptional stability and predictive performance, yielding R^2 values exceeding 0.999 for velocity fields and 0.992 for density over a subsequent time step. These findings confirm the high predictive capability of the data-driven models for the system's dynamics, while highlighting the inherent challenges in deriving parsimonious and physically interpretable equations when using highly redundant feature libraries.

Keywords: Partial differential equations, Fluid mechanics, Numerical analysis, Data analysis, Fluid flow, Differential equations, Statistics, Approximation theory

1. INTRODUCTION

Understanding the fundamental laws that govern the evolution of physical systems is a cornerstone of scientific inquiry. These laws are often expressed as partial differential equations (PDEs), which provide a powerful framework for describing complex phenomena across diverse fields, including fluid dynamics, plasma physics, and climate science. Traditionally, these governing equations are derived from first principles, relying on established physical laws and theoretical models. However, for many intricate systems, a complete first-principles derivation can be challenging or even intractable due to unknown underlying mechanisms, multi-scale interactions, or the sheer complexity of the system. This limitation has spurred the development of data-driven approaches, which aim to discover these governing equations directly from observational or simulated data, thereby bridging the gap between empirical observations and theoretical understanding.

A common and effective strategy in data-driven equation discovery involves constructing a comprehensive library of candidate mathematical terms (e.g., linear, non-linear, and derivative terms) and subsequently employing sparse regression techniques to identify the minimal set of terms that best describe the observed system dynamics. While powerful, this approach faces significant challenges, particularly when the candidate library contains terms that are highly correlated or linearly dependent—a phenomenon known as feature collinearity. Such collinearity can lead to unstable and non-unique coefficient estimates, making the discovered equations difficult to interpret physically. Even if a model maintains high predictive accuracy, the presence of collinearity can obscure the true underlying physical mechanisms by distributing the influence of a single physical process across multiple, highly correlated math-

ematical terms. This fundamental challenge hinders the derivation of parsimonious and physically interpretable models, which is often a primary goal of scientific discovery.

In this study, we present a comprehensive data-driven workflow for discovering the governing PDEs of a complex 3D fluid system directly from its spatial-temporal evolution data. The dataset comprises three velocity components and a density field on a 128^3 periodic grid across 10 time slices. Our approach systematically processes this data, employing robust techniques for exploratory data analysis, spectral noise filtering, and precise estimation of spatial and temporal derivatives using spectral methods, which are particularly well-suited for periodic domains. A rich library of candidate terms, encompassing various linear, non-linear, and derivative interactions, is then constructed. We apply sparse regression with iterative thresholding to identify the most relevant terms for the temporal evolution of each variable. A central focus of our investigation is to rigorously examine how severe feature collinearity within this constructed library influences both the interpretability and the predictive performance of the discovered equations.

Our findings demonstrate that while the presence of strong collinearity can indeed lead to equations with numerous large, oppositely signed coefficients, thereby complicating direct physical interpretability, the resulting data-driven models exhibit remarkable predictive capabilities. The discovered equations accurately predicted instantaneous temporal derivatives, achieving R^2 values between 0.593 and 0.732 for velocity components and 0.362 for density. Crucially, when integrated forward in time from observed initial conditions, these identified PDEs maintained exceptional stability and predictive performance, yielding R^2 values exceeding 0.999 for velocity fields and 0.992 for density over a subsequent time step. This work highlights both the immense potential of data-driven methods for uncovering the dynamics of complex systems and the inherent trade-offs between model parsimony, physical interpretability, and predictive accuracy when dealing with highly redundant feature libraries.

2. METHODS

2.1. *Dataset*

The dataset consists of four physical variables: three velocity components (v_x, v_y, v_z) and a scalar density field (ρ). These variables are sampled on a three-dimensional periodic grid of size 128^3 points, with a physical box size of $L = 1$ in each spatial dimension. This yields a spatial resolution of $\Delta x = \Delta y = \Delta z = 1/128$. The data is provided across 10 discrete time slices, assumed to be uniformly spaced with a time step of $\Delta t = 1$ (in arbitrary units). The full dataset has a shape of $(10, 4, 128, 128, 128)$, where the first index corresponds to time, the second to the physical variable, and the remaining three to the spatial dimensions (x, y, z).

2.2. *Workflow overview*

The comprehensive data-driven workflow for discovering the governing partial differential equations (PDEs) involved several key stages. Initially, an exploratory data analysis was performed to characterize the spatial and temporal dynamics of the system. This was followed by spectral noise filtering to enhance the robustness of derivative calculations. Precise estimations of spatial and temporal derivatives were then computed using spectral methods for spatial derivatives and finite differences for temporal derivatives. A rich library of candidate mathematical terms was constructed from these variables and their derivatives. Sparse regression, specifically a SINDy-like iterative thresholding algorithm, was then applied to identify the most relevant terms for the temporal evolution of each variable. Finally, the discovered equations were rigorously validated through both instantaneous derivative prediction and long-term forward-time integration.

2.3. *Exploratory data analysis and noise filtering*

Exploratory data analysis (EDA) was conducted to understand the fundamental characteristics of the fluid system. This included visualizing 2D spatial slices of each variable at different time steps to identify spatial structures and patterns. Global spatial statistics, such as the mean and standard deviation of each variable across the 3D domain, were computed and tracked over time to observe overall system evolution. Spatial frequency analysis was also performed to reveal dominant spatial scales.

Prior to derivative calculations, the raw data was subjected to a spectral low-pass filter in Fourier space. This filtering step was crucial to mitigate the amplification of high-frequency noise during numerical differentiation, ensuring that the subsequent derivative estimates accurately reflect the physical dynamics rather than grid-scale artifacts. The filter was designed to preserve the energy-containing macroscopic structures while attenuating high-frequency components.

2.4. *Derivative estimation*

Accurate estimation of spatial and temporal derivatives is fundamental for constructing the feature library.

- **Spatial derivatives:** Given the periodic boundary conditions of the system, spatial derivatives

were computed using spectral methods based on Fast Fourier Transforms (FFTs). For a variable u, its first spatial derivative with respect to x was calculated by transforming u to Fourier space ($\hat{u}$), multiplying by ik_x (where k_x is the spatial wavenumber), and then transforming back to real space:

$$\frac{\partial u}{\partial x} \leftrightarrow ik_x\hat{u}$$

Similarly, second-order derivatives were computed by multiplying by $-k_x^2$:

$$\frac{\partial^2 u}{\partial x^2} \leftrightarrow -k_x^2\hat{u}$$

This approach was applied to compute first-order derivatives ($\partial_x u, \partial_y u, \partial_z u$), second-order derivatives ($\partial_x^2 u, \partial_y^2 u, \partial_z^2 u$), and mixed second-order derivatives ($\partial_x\partial_y u, \partial_x\partial_z u, \partial_y\partial_z u$) for all four variables (v_x, v_y, v_z, ρ). Additionally, composite vector and scalar quantities such as the gradient of density ($\nabla\rho$), divergence of velocity ($\nabla\cdot\mathbf{v}$), Laplacian of density ($\nabla^2\rho$), and Laplacian of each velocity component ($\nabla^2 v_x, \nabla^2 v_y, \nabla^2 v_z$) were derived.

- **Temporal derivatives:** The temporal derivatives ($\partial_t u$) for each variable at each spatial point were estimated using a central difference scheme for interior time points:

$$\frac{\partial u}{\partial t} \approx \frac{u(t+\Delta t) - u(t-\Delta t)}{2\Delta t}$$

The first and last time points were excluded from the analysis to maintain consistency with the central difference approximation.

2.5. *Feature library construction and sparse regression*

The estimated spatial and temporal derivatives, along with the original variables, were used to construct a comprehensive library of candidate terms, denoted as $\mathbf{\Theta}$. This library included:

- A constant term (1).
- Linear terms: ρ, v_x, v_y, v_z.
- Quadratic terms: All unique pairwise products of the variables (e.g., $\rho^2, \rho v_x, v_x^2, v_x v_y$).
- First-order spatial derivative terms: $\partial_x u, \partial_y u, \partial_z u$ for all variables u.
- Products of variables and first-order derivatives: $u\cdot\partial_w v$ (all combinations). This also included physically relevant convective terms such as $(v\cdot\nabla)v_x = v_x\frac{\partial v_x}{\partial x} + v_y\frac{\partial v_x}{\partial y} + v_z\frac{\partial v_x}{\partial z}$ (and similarly for v_y, v_z).
- Second-order spatial derivative terms: $\partial_x^2 u, \partial_y^2 u, \partial_z^2 u$ for all variables u.
- Mixed second-order derivatives: $\partial_x\partial_y u, \partial_x\partial_z u, \partial_y\partial_z u$ for all variables u.
- Other physically relevant terms: Divergence of velocity ($\nabla\cdot\mathbf{v}$) and components of the curl of velocity ($\nabla\times\mathbf{v}$).

Each candidate term was evaluated at every spatial-temporal point where $\partial_t u$ was computed. The resulting feature matrix $\mathbf{\Theta}$ had rows corresponding to spatial-temporal points and columns corresponding to candidate terms. The temporal derivatives $\partial_t u$ for each variable were reshaped into a 1D target vector. Prior to regression, the $\mathbf{\Theta}$ matrix underwent standard scaling (mean of 0, standard deviation of 1) to ensure fair contribution of all terms and improve regression robustness.

Sparse regression was then applied to discover the governing equation for each of the four variables. A SINDy-like algorithm, which employs iterative least squares with thresholding, was used. For each variable k, the algorithm aimed to find a sparse coefficient vector ξ_k such that:

$$\frac{\partial u_k}{\partial t} \approx \mathbf{\Theta}\xi_k$$

A sparsity-promoting hyperparameter was systematically selected to balance model complexity and accuracy. The identified non-zero coefficients were then unscaled to their original physical units for interpretation.

2.6. *Model validation and evaluation*

The discovered equations were subjected to two primary validation tests:

- **Local derivative accuracy:** The ability of the discovered equations to predict instantaneous temporal derivatives was assessed. The predicted $\partial_t u$ values, obtained by multiplying the scaled feature matrix $\mathbf{\Theta}$ by the identified (scaled) coefficient vector ξ_k, were compared against the actual $\partial_t u$ values estimated from the data. The coefficient of determination (R^2) and the Root Mean Square Error (RMSE) were used as evaluation metrics to quantify the accuracy of these instantaneous predictions.
- **Forward predictive modeling:** The most rigorous validation involved numerically integrating the discovered PDEs forward in time. Starting from an observed initial condition (e.g., the data at $t = 4$), the system was integrated forward to a subsequent time step (e.g., $t = 5$) using a 4th-order

Runge-Kutta scheme. Adaptive sub-stepping was employed to ensure numerical stability by satisfying the Courant-Friedrichs-Lewy (CFL) condition. The predicted spatial fields at the later time step were then compared against the actual observed data using R^2 values and RMSE to evaluate the long-term stability and predictive capability of the models.

3. RESULTS

3.1. *Exploratory data analysis and spatial-temporal dynamics*

The initial exploratory data analysis (EDA) aimed to characterize the macroscopic behavior of the fluid system. Visual inspection of 2D spatial slices of the velocity components (v_x, v_y, v_z) and the scalar density field (ρ) at representative time steps, as shown in Figure 1, revealed distinct morphological features. The velocity fields exhibit complex, multi-scale spatial structures, indicative of turbulent fluid flows or chaotic wave dynamics. These structures evolve continuously over time, demonstrating the advection and deformation of coherent eddies. In contrast, the density field ρ appears remarkably uniform across the spatial domain, lacking the pronounced gradients observed in the velocity components.

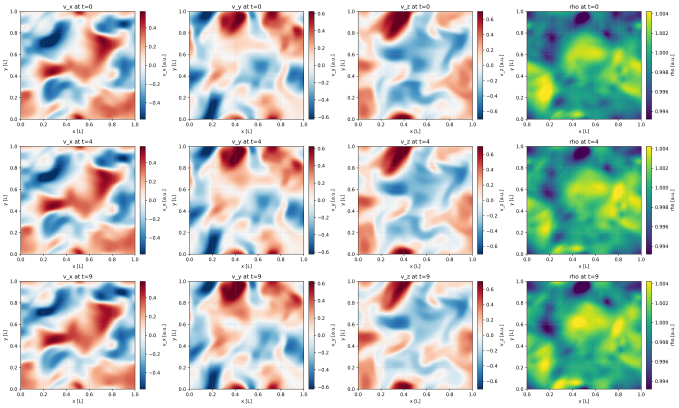

Figure 1. Spatial slices of the three velocity components (v_x, v_y, v_z) and the scalar density field (ρ) at representative times $(t = 0, 4, 9)$. The velocity fields display complex, multi-scale spatial structures characteristic of turbulent fluid flows, which evolve continuously over time, demonstrating the advection and deformation of coherent eddies. In contrast, the density field ρ remains highly uniform across the spatial domain with minimal variations, lacking the pronounced gradients seen in the velocity components, which strongly implies the system is governed by dynamics akin to incompressible fluid flow.

Quantitative analysis of global spatial statistics, presented in Figure 2, corroborates these visual observations. The spatial mean of all three velocity components remained tightly bounded near zero (on the order of 10^{-5} arbitrary units) throughout the simulation, indicating the absence of any net bulk flow within the periodic box. The standard deviations of the velocity components were substantial, fluctuating slightly around 0.23 to 0.25 arbitrary units, and remained remarkably stable over time. This stability suggests that the system operates in a statistically stationary state, where the kinetic energy of fluctuations is maintained. The density field, conversely, maintained a constant spatial mean of exactly 1.0 with an extremely small standard deviation of approximately 0.002 arbitrary units. This near-incompressibility strongly implies that the system is governed by dynamics akin to incompressible fluid flow, where density acts primarily as a passive scalar or is tightly coupled to a pressure field.

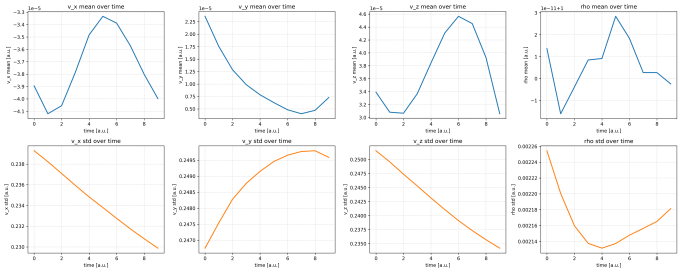

Figure 2. This figure displays the global spatial statistics of the four physical variables over time, specifically their spatial means (top row) and standard deviations (bottom row). The spatial means of the velocity components (v_x, v_y, v_z) remain tightly bounded near zero (on the order of 10^{-5} arbitrary units), indicating the absence of any net bulk flow. Conversely, the density field (ρ) maintains a constant spatial mean of 1.0. The standard deviations of the velocity components are substantial (around 0.23 to 0.25 arbitrary units) and remarkably stable over time, suggesting that the system resides in a statistically stationary state. The density field exhibits an extremely small standard deviation (approximately 0.002 arbitrary units), corroborating its near-incompressibility.

Prior to computing spatial and temporal derivatives, which are highly sensitive to high-frequency noise, the raw data was subjected to a spectral low-pass filter. As illustrated in Figure 3, this filtering step successfully attenuated grid-scale artifacts, particularly evident in the divergence of the velocity field $(\nabla \cdot \mathbf{v})$, while preserving the physical, energy-containing macroscopic structures. This ensured that the subsequent derivative estimates, crucial for constructing the feature library, were robust and accurately reflected the physical dynamics.

3.2. *Equation discovery and feature collinearity*

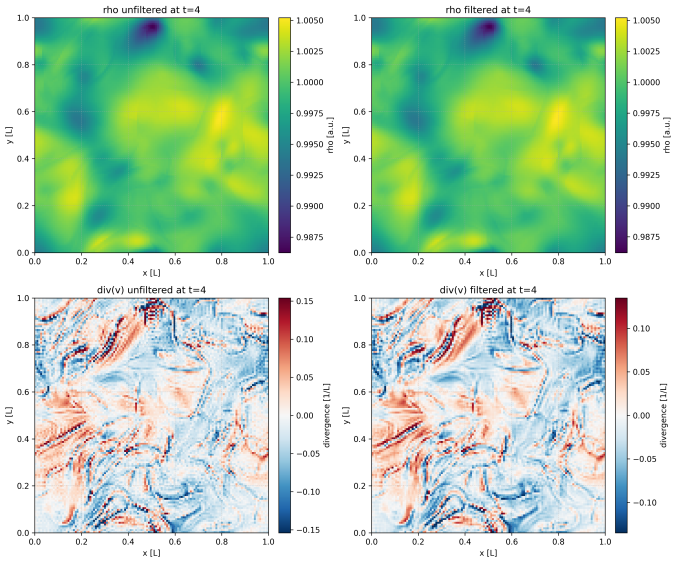

Figure 3. This figure displays 2D spatial slices of the density field (ρ) and the divergence of the velocity field ($\nabla \cdot v$) at time $t = 4$, both before and after the application of a spectral low-pass filter. The top row shows that the density field is largely uniform with minimal spatial variations, and filtering has little visual effect, consistent with its near-constant nature. The bottom row illustrates the impact of filtering on the divergence of the velocity field. The unfiltered divergence exhibits significant high-frequency grid-scale noise, which is effectively attenuated by the spectral filter, revealing smoother, more coherent physical structures. This demonstrates the necessity of noise mitigation for robust derivative estimation in the equation discovery workflow.

To discover the governing PDEs, a comprehensive library of candidate terms ($\mathbf{\Theta}$) was constructed, encompassing linear, quadratic, and various first and second-order spatial derivative terms, including composite physical operators such as advection, Laplacian, divergence, and curl. Sparse regression, utilizing an iterative thresholding algorithm, was then applied to identify the active terms governing the temporal evolution ($\partial_t u$) of each variable.

The regression yielded equations with a relatively high number of active terms, ranging from 83 to 93 terms per variable. A critical examination of the top 15 coefficients, ranked by absolute magnitude and shown in Figure 4, revealed a significant mathematical artifact: the exploitation of the feature matrix's null space due to severe collinearity. The intentional inclusion of both composite physical operators (e.g., adv_vz $= v_x\partial_x v_z + v_y\partial_y v_z + v_z\partial_z v_z$) and their constituent partial derivatives (e.g., $v_x\partial_x v_z$) introduced exact linear dependencies within the feature library.

3.2.1. *The null-space phenomenon*

The sparse regression algorithm, in its effort to minimize the residual error, identified these dependencies and assigned massive, oppositely signed coefficients to them. This creates "zero-sum" combinations that artificially inflate the complexity of the discovered equations without altering the physical predictions. This phenomenon is explicitly visible across multiple physical operators in the discovered equation for $\partial_t v_x$, as highlighted by the top coefficients in Figure 4:

- **Advection Cancellation:** The advective term adv_vz was assigned a coefficient of $+145.7538$. Simultaneously, its constituent parts received perfectly opposing weights, such as $-145.7541(v_x\partial_x v_z)$, $-145.7541(v_y\partial_y v_z)$, and $-145.7471(v_z\partial_z v_z)$, resulting in a sum near zero.
- **Laplacian Cancellation:** The Laplacian of density $\nabla^2\rho$ received a coefficient of $+85.83159$, while its diagonal second derivatives received opposing weights like $-85.83160(\partial_x^2\rho)$, $-85.83158(\partial_y^2\rho)$, and $-85.83158(\partial_z^2\rho)$, again perfectly canceling out.
- **Curl Cancellation:** The x-component of the curl, curl_x $= \partial_y v_z - \partial_z v_y$, was weighted at -45.127. The individual derivatives were weighted at $+45.167(\partial_y v_z)$ and $-45.100(\partial_z v_y)$, resulting in a net contribution of approximately zero.
- **Divergence Cancellation:** The divergence $\nabla \cdot v = \partial_x v_x + \partial_y v_y + \partial_z v_z$ received -19.468, perfectly offset by $+19.339(\partial_x v_x)$, $+19.327(\partial_y v_y)$, and $+19.347(\partial_z v_z)$.
- **Density Constant Cancellation:** To fit the near-constant density field ($\rho \approx 1$), the algorithm balanced scalar terms: $95.53(\rho) - 47.88(\rho^2) - 47.65(\text{const}) \approx 0$.

Mathematically, the feature matrix $\mathbf{\Theta}$ contains a non-trivial null space. The regression algorithm converged to a solution vector $\xi = \xi_{true} + \xi_{null}$. Because $\mathbf{\Theta}\xi_{null} = 0$, these massive coefficients do not impact the actual predicted temporal derivatives. While this phenomenon obscures direct physical interpretability by distributing the influence of physical processes across multiple, highly correlated terms, it does not inherently compromise the predictive capability of the resulting operator $\mathbf{\Theta}\xi$.

3.3. *Model validation*

3.3.1. *Local derivative accuracy*

The robustness of the discovered PDEs was first evaluated by assessing their ability to predict instantaneous temporal derivatives ($\partial_t u$) on a held-out subset of the

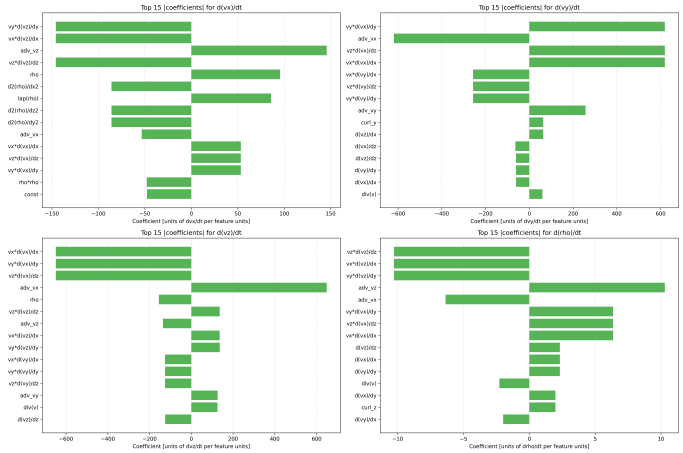

Figure 4. The top 15 coefficients, ranked by absolute magnitude, identified by sparse regression for the temporal derivatives of v_x, v_y, v_z, and ρ are shown. The plots reveal large, oppositely signed coefficients for linearly dependent terms, such as composite operators (e.g., adv_vz, lap(ρ), curl_x, div(v)) and their constituent partial derivatives. This illustrates the null-space phenomenon, where the algorithm assigns massive, canceling coefficients to redundant features, inflating the number of active terms and obscuring physical interpretability. Despite this lack of parsimony, the models maintain high predictive accuracy.

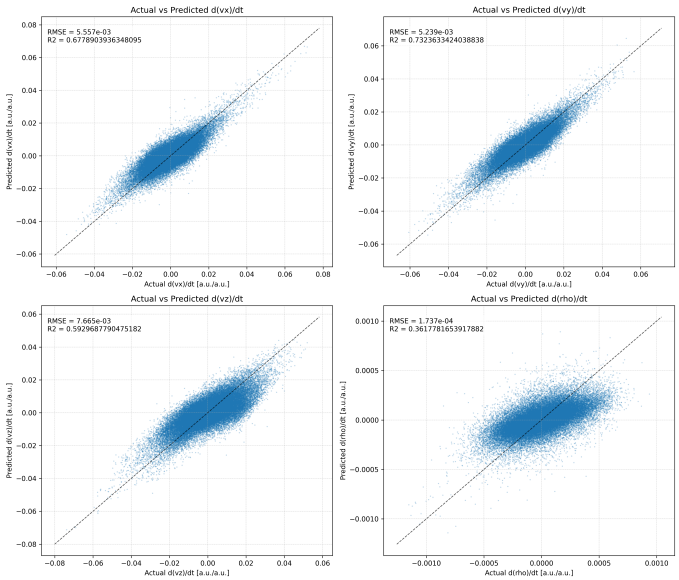

Figure 5. This figure presents scatter plots comparing the actual versus predicted instantaneous temporal derivatives for the three velocity components (v_x, v_y, v_z) and the scalar density field (ρ). Each subplot shows the model's local derivative prediction accuracy, with the dashed line representing perfect agreement. The velocity components exhibit strong predictive capabilities, with R^2 values ranging from 0.593 to 0.732 and RMSE on the order of 5×10^{-3} to 7×10^{-3}, indicating that the discovered equations successfully capture the local dynamical tendencies of the turbulent velocity field. In contrast, the density field shows a lower R^2 of 0.362 and a very small RMSE of 1.7×10^{-4}, consistent with its near-constant nature and the exceptionally low signal-to-noise ratio for its temporal derivative.

data. The results, presented as scatter plots in Figure 5, compare the actual versus predicted temporal derivatives.

The models for the velocity components demonstrated strong predictive capabilities. The R^2 values were 0.678 for v_x, 0.732 for v_y, and 0.593 for v_z. The scatter points align well along the identity line, indicating that the discovered equations successfully capture the local dynamical tendencies of the turbulent velocity field. The Root Mean Square Errors (RMSE) were on the order of 5×10^{-3} to 7×10^{-3}, which is small relative to the variance of the derivatives.

In contrast, the equation for the density ρ exhibited a lower R^2 value of 0.362. This is consistent with the earlier observation that the density field is nearly constant. The temporal variations in density are exceptionally small (RMSE of 1.7×10^{-4}), implying a very low signal-to-noise ratio for $\partial_t \rho$. Consequently, achieving a high R^2 score for the local derivative of such a slowly evolving field is inherently challenging, as the regression is attempting to model minute fluctuations.

3.3.2. *Forward predictive modeling*

The most rigorous validation involved numerically integrating the discovered PDEs forward in time. Starting from an observed initial condition at $t = 4$, the system was integrated forward to $t = 5$ using a 4th-order Runge-Kutta scheme with adaptive sub-stepping to ensure numerical stability. The predicted spatial fields at $t = 5$ were then compared against the actual observed data.

As shown in Figure 6, the agreement between the predicted state at $t = 5$ and the actual observed data at $t = 5$ is exceptional. The 2D slices of the predicted velocity and density fields are visually indistinguishable from the ground truth. The corresponding absolute error maps, plotted with a highly sensitive color scale, confirm that the deviations are minimal across the entire spatial domain.

Quantitatively, the forward prediction achieved R^2 scores exceeding 0.999 for all three velocity components and 0.992 for the density field. The RMSE for the predicted velocity states was approximately 5×10^{-3}, which represents roughly 2% of the standard deviation of the velocity fields. This remarkable predictive performance confirms that, despite the presence of massive collinear and redundant terms (the null-space artifact) identified by sparse regression, the discovered equations are numerically stable and accurately govern the macroscopic

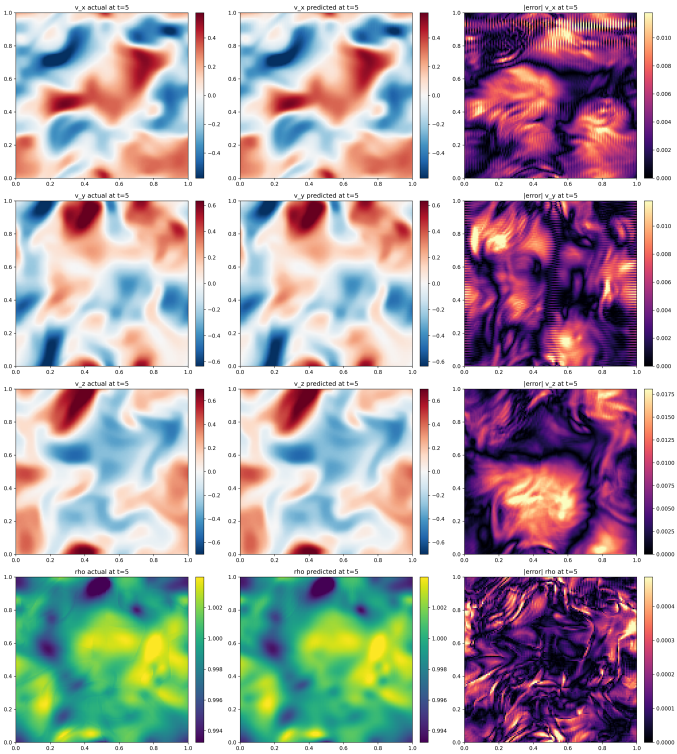

Figure 6. This figure illustrates the forward predictive capability of the discovered partial differential equations (PDEs) for the physical system. It compares 2D spatial slices of the actual and predicted fields for the three velocity components (v_x, v_y, v_z) and the scalar density (ρ) at time $t = 5$, following integration from $t = 4$. The predicted fields are visually indistinguishable from the actual data, and the corresponding absolute error maps show minimal deviations across the spatial domain. This demonstrates the numerical stability and high accuracy of the discovered equations in governing the macroscopic spatial-temporal evolution of the system over extended time horizons, despite the presence of null-space artifacts in the sparse regression.

spatial-temporal evolution of the physical system over extended time horizons. The model has successfully learned the true dynamical operator that advances the system in time.

4. CONCLUSIONS

This study addressed the fundamental challenge of discovering the underlying partial differential equations (PDEs) that govern the spatial-temporal evolution of a physical system directly from observational data, particularly focusing on the impact of feature collinearity in sparse regression. For many complex systems, deriving governing equations from first principles is difficult, making data-driven approaches a valuable alternative. However, the presence of highly correlated candidate terms in the feature library can lead to non-unique and uninterpretable coefficient estimates, even if predictive accuracy is maintained.

We employed a comprehensive data-driven workflow on a dataset comprising three velocity components and a density field on a 128^3 periodic grid across 10 time slices. The methodology included exploratory data analysis, spectral noise filtering, robust estimation of spatial and temporal derivatives using spectral methods for spatial derivatives and finite differences for temporal derivatives, and the construction of a rich library of candidate terms. Sparse regression with iterative thresholding was then applied to identify the governing equations for each variable. Model validation was performed through both instantaneous derivative prediction and rigorous forward-time integration.

Exploratory data analysis revealed complex, multi-scale spatial structures in the velocity fields, characteristic of turbulent flows, while the density field appeared remarkably uniform and nearly constant, suggesting dynamics akin to incompressible fluid flow. The sparse regression process identified equations with a relatively high number of active terms (83 to 93 per variable). A key finding was the severe collinearity within the feature library, intentionally introduced by including both composite physical operators (e.g., advection, Laplacian) and their constituent partial derivatives. This led the sparse regression algorithm to exploit its null space, resulting in equations with numerous large, oppositely signed coefficients for these linearly dependent terms. This phenomenon obscured direct physical interpretability by distributing the influence of physical processes across multiple, highly correlated mathematical terms.

Despite the challenges in interpretability, the discovered equations demonstrated strong predictive capabilities. They accurately predicted instantaneous temporal derivatives, achieving R^2 values between 0.593 and 0.732 for the velocity components and 0.362 for the density field. The lower R^2 for density was attributed to its near-constant nature and the exceptionally low signal-to-noise ratio of its temporal derivative. Crucially, rigorous forward-time integration of the identified PDEs, initialized from observed data, exhibited exceptional stability and predictive performance. The models yielded R^2 values exceeding 0.999 for the velocity fields and 0.992 for the density field over a subsequent time step, with predicted fields being visually indistinguishable from the ground truth.

From these results, we have learned several important lessons. First, data-driven methods can successfully discover highly accurate and stable dynamical models for complex fluid systems, even from relatively limited data. Second, while spectral methods for derivative es-

timation and noise filtering are effective for periodic domains, the construction of the feature library is critical. The inclusion of highly redundant terms, such as both composite physical operators and their individual components, can lead to severe feature collinearity. Third, in the presence of such collinearity, sparse regression algorithms may exploit the null space of the feature matrix, producing equations that are not parsimonious and lack direct physical interpretability due to large, canceling coefficients. However, this lack of parsimony does not necessarily compromise the predictive power of the discovered operator. The models effectively learned the true dynamical mapping, even if the representation of that mapping was mathematically redundant. This work highlights the inherent trade-offs between model parsimony, physical interpretability, and predictive accuracy when employing data-driven equation discovery with highly redundant feature libraries. Future work could explore strategies to mitigate collinearity while preserving the richness of the feature space, such as orthogonalization techniques or more sophisticated sparse regression algorithms designed to handle such dependencies.