
Neural ODEs as a discovery tool to characterize the structure of the hot galactic wind of M82

Dustin Nguyen*

The Ohio State University
dnguyen.phys@gmail.com

Yuan-Sen Ting

Australian National University
The Ohio State University
yuan-sen.ting@anu.edu.au

Todd A. Thompson

The Ohio State University
thompson.1847@osu.edu


Sebastian Lopez

The Ohio State University
lopez.764@osu.edu

Laura A. Lopez

The Ohio State University
lopez.513@osu.edu

Abstract

Dynamic astrophysical phenomena are predominantly described by differential equations, yet our understanding of these systems is constrained by our incomplete grasp of non-linear physics and scarcity of comprehensive datasets. As such, advancing techniques in solving non-linear inverse problems becomes pivotal to addressing numerous outstanding questions in the field. In particular, modeling hot galactic winds is difficult because of unknown structure for various physical terms, and the lack of *any* kinematic observational data. Additionally, the flow equations contain singularities that lead to numerical instability, making parameter sweeps non-trivial. We leverage differentiable programming, which enables neural networks to be embedded as individual terms within the governing coupled ordinary differential equations (ODEs), and show that this method can adeptly learn hidden physics. We robustly discern the structure of a mass-loading function which captures the physical effects of cloud destruction and entrainment into the hot superwind. Within a supervised learning framework, we formulate our loss function anchored on the astrophysical entropy ($K \propto P/\rho^{5/3}$). Our results demonstrate the efficacy of this approach, even in the absence of kinematic data v . We then apply these models to real Chandra X-Ray observations of starburst galaxy M82, providing the first systematic description of mass-loading within the superwind. This work further highlights neural ODEs as a useful discovery tool with mechanistic interpretability in non-linear inverse problems. We make our code public at this GitHub repository (github.com/dustindnguyen/2023_NeurIPS_NeuralODEs_M82 .

1 Introduction

In the field of physics, understanding is often equated with having definitive differential equations that describe dynamic variables [1]. Yet, when it comes to modeling real-world systems, the true underlying physics is not always clear. This is especially the case for modern galactic wind models, where non-linear phenomena are difficult to describe analytically. Recent wind-cloud interaction simulations [2–4] indicate efficient non-linear multi-phase mixing can overcome the so-called cloud-crushing problem [5] posed by observations of fast cool outflows from nearby starburst galaxies [6]. Subsequently, there has been an explosion in activity focused on running suites of high-resolution 3D time-dependent hydrodynamic simulations to extract useful scaling laws relevant to cloud survival

*Corresponding Author

and the interface of radiative turbulent mixing [7–11]. These intuition-based and simulation-based analytic models offer valuable insights. However, considering observations of the cosmos already encodes the necessary information, a shift towards data-driven modeling, set apart from the customary data-testing approach in astronomy, combined with traditional physical models might provide fresh insights into the intrinsic structure of galactic winds.

Deep neural networks have been highlighted as powerful tools for approximating unknown functions and operators [12–15]. A noteworthy application of these networks is their role in the numerical discretization of ODEs/PDEs, leveraging the adjoint sensitivity method [16]. Physics-Informed Neural Networks (PINNs) emphasize the integration of differential equations into their loss functions and employ the adjoint sensitivity method for effective gradient retrieval [17]. Expanding on these ideas, the landscape has further matured to introduce the Universal Differential Equation (UDE) paradigm [18]. Unlike PINNs, which primarily focus on embedding the structure of differential equations into the neural network’s loss function, UDEs pioneer the fusion of neural networks directly into the differential equations themselves, significantly amplifying their approximation prowess. Once integrated, these networks inherently respect conservation laws and shape the system’s dynamics sequentially, ensuring that each contribution can be meaningfully and mechanistically interpreted. Notably, while PINNs are unique in their solution approach, UDEs can be seamlessly solved using standard numerical methods (e.g., Runge-Kutta 4 [RK4]), bridging the gap between traditional differential equation solvers and deep learning.

The focus of this paper revolves around these UDEs, which in this paper, we refer to as simply "neural ODEs". The incorporation of neural networks into ODEs/PDEs through differentiable programming is an emerging research area with broad application in the physical and biological sciences [19–26]. In the realm of galactic wind simulations, 3D time-dependent models offer intricate insights into non-linear behaviors. However, 1D steady-state models have demonstrated their efficacy in accurately capturing the statistical global properties of these complex 3D simulations. Such 1D models provide a streamlined approach, effectively summarizing the essential characteristics and trends of galactic winds, and form the foundation for our exploration using neural networks. In this work we use neural ODEs to model real observations of a galaxy and attempt to characterize the structure of mass-loading, which captures the effect of cool cloud destruction, and subsequent entrainment, into a hot supernovae-driven galactic wind. We focus on characterizing X-ray observations of the widely studied starburst prototype galaxy: Messier 82 (M82). We use the astrophysical entropy ($K/k_b = T/n^{2/3}$) as a feature-engineered (physical) variable within the loss function and penalize diverging solutions. We will show that regression on the this loss function allows us to infer a mass-loading model (with no prior knowledge) that makes predictions which better match the observed temperature and density profiles.

2 Neural Galactic Wind Model

The hydrodynamic equations for a steady-state hot flow moving in the x direction are [27, 28, 26]:

$$\frac{(A\rho v)_x}{A} = \dot{\mu}, \quad vv_x = -\frac{P_x}{\rho} - \nabla\Phi - \frac{\dot{\mu}v}{\rho}, \quad \text{and} \quad v\epsilon_x - \frac{vP\rho_x}{\rho^2} = -\frac{n^2\Lambda}{\rho} + \frac{\dot{\mu}}{\rho}\left(v^2 - \epsilon - \frac{P}{\rho}\right), \quad (1)$$

where v , ρ , P , ϵ , $\nabla\Phi$, n , Λ , A , $\dot{\mu}$, and subscript x are the bulk velocity, density, pressure, specific internal energy, gravity, number density, cooling rate, surface area expansion rate, volumetric mass-loading rate, and first-order spatial derivative, respectively. Mass-loading term $\dot{\mu}$ captures the global effects of cool cloud destruction and incorporation into the hot phase [27, 7, 28, 29]. In this work, it is assumed that the entrained cool material contains negligible velocity and temperature, which is valid in the limit that the hot galactic wind contains most of the thermal and kinetic energy [27]. We take $\nabla\Phi = \sigma^2/x$, where $\sigma = 200 \text{ km s}^{-1}$, polynomial fit [30] to the radiative cooling curve, and A is defined below separately for the mock test and comparison to real data. When substituted into each other, each of the derivatives v_x , ρ_x , and P_x are proportional to $(\mathcal{M}^2 - 1)^{-1}$ [31] and thus all contain a singularity at the sonic point (i.e., numerically diverges at $\mathcal{M} = 1$). The dimensions of each simulation is $0.37 \text{ kpc} \leq x \leq 2.65 \text{ kpc}$ and number of steps is $n_x = 500$. We do not use all 500 points during optimization in both the mock test and comparison to real data, which only has 44 resolved data points. As there is currently no kinematic measurements, we use the classic galactic wind model by Chevalier and Clegg [32] to guess an initial velocity of $v_{\text{hot}} \sim 1835 \text{ km s}^{-1}$ after leaving the starburst volume $R = 0.3 \text{ kpc}$. The initial conditions are then $n_0 = 0.843 \text{ cm}^{-3}$,

$T_0 = 0.615 \times 10^7$ K. In the mock test we take $v_0 = v_{\text{hot}}$, and then for the Chandra data, we consider two initial velocities $v_{0,a} = v_{\text{hot}}$ and $v_{0,b} = v_{\text{hot}}/2$.

In summary, our method is:

1. Solve an initial value problem (Eqs. 1) by RK4 integration of v_0 , ρ_0 , P_0 from x_i to x_f .
2. Calculate loss (Eqs. 2 and 3) between the data and integrated solutions at points of $n_{x,\text{data}}$.
3. Backpropagate through automatic-differentiated ODE solver to get gradients.
4. Update weights of individual neural network variable $\dot{\mu}$.
5. Iterate for 150 epochs of ADAM optimization and then up to 150 epochs of BFGS optimization.

We represent the volumetric mass-loading rate, $\dot{\mu}$, with a multi-layer perceptron neural network comprised of 3 hidden layers. The input into $\dot{\mu}$ is a single position, and the full range of positions is sampled by forward integration of Eqs. 1 using 5th order RK4 [33]. The bulk velocity v cannot be used in the optimization problem because there is not yet any kinematic data available for M82's X-ray emitting wind. We calculate the loss function using a feature-engineered quantity, the astrophysical entropy $K/k_b = T/n^{2/3} \propto P/\rho^{5/3}$. The loss function is a weighted Mean-Square-Error (MSE):

$$\mathcal{L} = \sum_i^{n_{x,\text{data}}} \left[\mathcal{W}_i \times \left(K_i - \hat{K}_i \right)^2 \right]. \quad (2)$$

where K is the data, and \hat{K} are the solutions of the integrated ODEs. The weights \mathcal{W} linearly scale the MSE as a function of n_x , where $\mathcal{W}_0 = 1$ and $\mathcal{W}_{n_{x,\text{data}}} \ll 1$, $n_{x,\text{data}} = 44$ is the number of data points. This scaling increases sensitivity to early solutions, which is important for non-linear problems. We do not include division by $n_{x,\text{data}}$ or k_b^2 because it does not impact training. If the flow at any instance reaches the sonic point, $\mathcal{M} = 1$, the equations become numerically unstable [31], hindering the optimization process. We prevent this by introducing a penalty term [26] that activates between $1 \leq \hat{\mathcal{M}} \leq \mathcal{M}_{\text{penalty}}$, where we take $\mathcal{M}_{\text{penalty}} = 1.5$. In this region, the loss is artificially increased per optimization step by:

$$\mathcal{L} = \mathcal{L}_{\text{previous}} \times \omega \sum_i^{n_{x,\text{data}}} \left[1 - \left(1 - \hat{\mathcal{M}} \right)^2 \right] \quad (1 \leq \hat{\mathcal{M}} \leq \mathcal{M}_{\text{penalty}}), \quad (3)$$

where ω is a constant. We minimize the loss function using two optimization algorithms, which has been shown to be required for convergence in other neural ODE studies [34, 26]. We solve the equations, automatic-differentiate the system, and calculate gradients entirely within the Julia SciML ecosystem.

3 Results

Mock Test: We test the framework to learn a mass-loading function described by

$$\dot{\mu}_{\text{truth}} = \dot{\mu}_0 \times a^\Delta / (x^\Delta \times (1 + x/a)^{\Gamma-\Delta}) \quad (4)$$

where $\dot{\mu}_0 = 10 M_\odot \text{ kyr}^{-1} \text{ kpc}^{-3}$, $a = 1.5 \text{ kpc}$, $\Delta = 4.0$, and $\Gamma = -4.0$. This equation scales as $x^{-\Delta}$ and transitions to $x^{-\Gamma}$ at approximately distance a , describing intense cloud entrainment after the wind leaves the host galaxy [28]. We take the flow geometry to be a flared-cylinder $A = A_0 (1 + (x/\eta)^2)$, which characterizes a cylindrically-collimated flow that later undergoes spherical expansion [35–37]. Substituting Eq. 4 into 1, we calculate the mock-data (i.e., $\dot{\mu} \rightarrow \dot{\mu}_{\text{truth}}$). We then “forget” $\dot{\mu}_{\text{truth}}$ and replace $\dot{\mu}$ with a neural network and initialize it such that the output across all x is approximately 0 (using Flux.jl’s default Glorot initialization). The first (i.e., untrained) neural ODE solution will be that of a wind with approximately zero mass-loading. The mock data is calculated with $n_x = 500$ steps, however, we utilize only $n_{x,\text{data}} = 44$ linearly spaced samples of the mock dataset during optimization to mimic the resolution of the Chandra X-ray data we later focus on. In Fig. 1 we plot the neural ODE solutions. Despite not using the kinematic mock data v we still predict the correct kinematic, and thermodynamic, solutions and learn the underlying mass-loading function $\dot{\mu}_{\text{truth}}$.

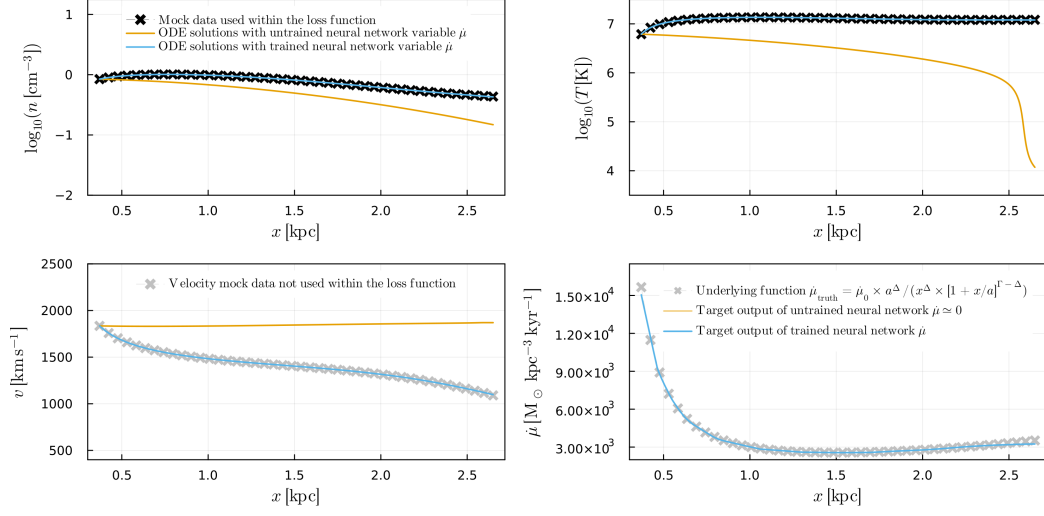


Figure 1: Density, temperature, and velocity profiles for neural ODE solutions and the target output of neural network variable $\hat{\mu}$ (bottom right panel). The x's are the mock data, of which, only the density and temperature is used in the loss function. Despite not using the kinematic mock data, the correct kinematics is predicted and the underlying mass-loading function was learned after training.




Figure 2: Density, temperature, and velocity profiles for neural ODE solutions and the target output of neural network variable $\hat{\mu}$ (bottom right panel) for two different initial velocities (yellow and blue lines). The x's are the Chandra X-ray data which exists only for temperature and density. After training, the learned mass-loading function leads to a better fit to the data, as the thermalization of kinetic energy prevents rapid bulk cooling that would otherwise cool the flow to 10^4 K.

Application towards Chandra X-Ray Data: We now consider applications of neural ODEs towards Chandra X-ray observations of the northern wind of galaxy M82 (see [38]). There are 44 data points for n , T , and A . For illustrative purposes, we will consider two initial velocities: $v_{0,a} \sim v_{\text{hot}}$ and $v_{0,b} \sim v_{\text{hot}}/2$. We plot the neural ODE solutions for both scenarios in Fig. 2. The additional heating from the learned mass-loading prevents rapid bulk cooling leads to better matched temperature profiles. For $v_0 = v_{0,a}$, the target output of neural network variable $\dot{\mu}$ (yellow line) reveals non-trivial structure, sharply truncating at $x \sim 1.1$ kpc. The mass-loading rate and initial wind mass-outflow rate is calculated as $\dot{M}_{\text{load}} = \int dx A \dot{\mu}$ and $\dot{M}_{\text{wind}} = A_0 \rho_0 v_0$. We find $\dot{M}_{\text{load}}/\dot{M}_{\text{wind}} = 0.09$ and 0.59 for initial velocities of $v_0 = v_{0,a}$ and $v_{0,b}$. In the latter case, the mass-loading rate is roughly half the initial wind outflow, suggesting a sharp metal abundance gradient as seen in [38].

4 Conclusion

In this work we use neural ODEs to explain X-ray observations of outflows from starburst galaxy M82. Rather than parameter estimation on an assumed function, here the data formulates mass-loading term $\dot{\mu}$ without any prior knowledge its structure. We start with approximately zero mass-loading, showcasing the flexibility of the model to discover rich structure (orange line, bottom right panel Fig. 2) as an ab initio modeling process. This work highlights the exceptional utility of neural networks as universal function approximators for non-linear inverse problems.

New measurements by future X-ray space missions such as the recently launched XRISM satellite [39] and LEM concept [40], will provide the first kinematics measurements of the hot 10^7 K gas. Our study indicates the learned mass-loading factors are sensitive to the assumed initial hot gas velocity. These forthcoming data can easily be integrated into our methods to better understand cool cloud entrainment, ultimately shedding light on important processes in launching multi-phase winds.

Comparison to previous work: This paper employs neural networks as universal function approximators for individual terms in a steady-state galactic wind model, akin to the approach of [26]. The salient distinctions in our work include: (1) our flow equations encompass critical physical processes such as radiative cooling and gravity; (2) our optimization with mock data operates at roughly a tenth of the spatial resolution in [26], $n_{x,\text{data}}$, mirroring the resolution observed in real data; and (3) our model omits one of the three dynamical variables, v , reflecting the current absence of velocity data in X-ray observations.

Neural ODE comparison to Fourier series: In the context of a 1D steady-state problem like our galactic wind model, alternative universal function approximators, such as the Fourier series, may seem viable. However, future direct comparisons with 2D X-ray surface brightness images of galaxies, more intricate models—2D or 3D spatial hydrodynamics—are essential to capture finer structure. When considering time, these evolve into 3D and 4D models, respectively. Fourier series becomes less effective in these multidimensional contexts due to its inefficiency in scaling beyond 1D. Furthermore, it’s imperative to note that while our study employs neural networks strictly as universal function approximators, there’s a burgeoning effort in leveraging them for symbolic regression within ODE/PDE systems (see PySR and SymbolicRegression.jl), as evidenced by works such as [18, 41, 42, 25]. Employing symbolic regression on these trained neural networks promises to yield symbolic, interpretable expressions that elucidate the underlying physics being learned. While deriving such symbolic interpretations was not the primary objective of our study, the foundational step of training the neural network within the confines of the physical system was the central focus of this work. We also note that key integrated quantities, such as the total mass being deposited into the hot wind, can still be inferred without symbolic representations of the trained neural network.

Appendix

Data: Using 534 ks of archival Chandra X-ray Observatory data, we constrain the temperature and density gradient along the outflow spanning ± 2.6 kpc. Spectra was extracted using CIAO for 101 rectangular regions along the outflow, each of size $3'' \times 1'$. To constrain the temperature and density, each region’s spectra was modeled in XSPEC using `const * phabs * phabs(powerlaw + vavpec)`, where the abundances were frozen based on [38]. We focus on the northern outflow of M82. The southern side is largely asymmetric due to tidal interaction with M81, requiring a treatment that is beyond the scope of this work.

Acknowledgements

DN acknowledges funding from NASA 21-ASTRO21-0174. Y.S.T. acknowledges financial support from the Australian Research Council through DECRA Fellowship DE220101520.

References

- [1] James D. Meiss. *Differential Dynamical Systems*. Society for Industrial and Applied Mathematics, 2007. doi: 10.1137/1.9780898718232. URL <https://epubs.siam.org/doi/abs/10.1137/1.9780898718232>.
- [2] Max Gronke and S. Peng Oh. The growth and entrainment of cold gas in a hot wind. *Monthly Notices of the Royal Astronomical Society*, 480(1):L111–L115, October 2018. doi: 10.1093/mnras/sly131.
- [3] Max Gronke and S. Peng Oh. How cold gas continuously entrains mass and momentum from a hot wind. *Monthly Notices of the Royal Astronomical Society*, 492(2):1970–1990, February 2020. doi: 10.1093/mnras/stz3332.
- [4] Max Gronke and S. Peng Oh. Cooling driven coagulation. *arXiv e-prints*, art. arXiv:2209.00732, September 2022.
- [5] Richard I. Klein, Christopher F. McKee, and Philip Colella. On the Hydrodynamic Interaction of Shock Waves with Interstellar Clouds. I. Nonradiative Shocks in Small Clouds. *The Astrophysical Journal*, 420:213, January 1994. doi: 10.1086/173554.
- [6] C. R. Lynds and A. R. Sandage. Evidence for an Explosion in the Center of the Galaxy M82. *The Astrophysical Journal*, 137:1005, May 1963. doi: 10.1086/147579.
- [7] Evan E. Schneider, Eve C. Ostriker, Brant E. Robertson, and Todd A. Thompson. The Physical Nature of Starburst-driven Galactic Outflows. *The Astrophysical Journal*, 895(1):43, May 2020. doi: 10.3847/1538-4357/ab8ae8.
- [8] Drummond B. Fielding, Eve C. Ostriker, Greg L. Bryan, and Adam S. Jermyn. Multiphase Gas and the Fractal Nature of Radiative Turbulent Mixing Layers. *The Astrophysical Journal Letters*, 894(2):L24, May 2020. doi: 10.3847/2041-8213/ab8d2c.
- [9] Matthew W. Abruzzo, Drummond B. Fielding, and Greg L. Bryan. Taming the TuRMoiL: The Temperature Dependence of Turbulence in Cloud-Wind Interactions. *arXiv e-prints*, art. arXiv:2210.15679, October 2022.
- [10] Brent Tan, S. Peng Oh, and Max Gronke. Cloudy with a chance of rain: accretion braking of cold clouds. *Monthly Notices of the Royal Astronomical Society*, 520(2):2571–2592, April 2023. doi: 10.1093/mnras/stad236.
- [11] Brent Tan and Drummond B. Fielding. Cloud Atlas: Navigating the Multiphase Landscape of Tempestuous Galactic Winds. *arXiv e-prints*, art. arXiv:2305.14424, May 2023. doi: 10.48550/arXiv.2305.14424.
- [12] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. *Deep Learning*. Nature, New York, NY, 2015. doi: 10.1038/nature14539. URL <https://www.nature.com/articles/nature14539>.
- [13] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. The MIT Press, Cambridge, MA, 2016. ISBN 9780262035613.
- [14] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural Networks*, 3(5):551–560, 1990. ISSN 0893-6080. doi: [https://doi.org/10.1016/0893-6080\(90\)90005-6](https://doi.org/10.1016/0893-6080(90)90005-6). URL <https://www.sciencedirect.com/science/article/pii/0893608090900056>.

- [15] Qi Lu, Zhiping Mao, Weiran Zuo, and Bin Dong. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:270–278, 2021. doi: 10.1038/s42256-021-00302-5. URL <https://www.nature.com/articles/s42256-021-00302-5>.
- [16] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations, 2019.
- [17] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [18] Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, and Ali Ramadhan. Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*, 2020.
- [19] Rahel Vortmeyer-Kley, Pascal Nieters, and Gordon Pipa. A trajectory-based loss function to learn missing terms in bifurcating dynamical systems. *Scientific Reports*, 11:21181, Oct 2021. doi: 10.1038/s41598-021-99609-x. URL <https://www.nature.com/articles/s41598-021-99609-x>.
- [20] Maximilian Gelbrecht, Niklas Boers, and Jürgen Kurths. Neural partial differential equations for chaotic systems. *New Journal of Physics*, 23(4):043005, April 2021. doi: 10.1088/1367-2630/abeb90.
- [21] Brendan Keith, Akshay Khadse, and Scott E. Field. Learning orbital dynamics of binary black hole systems from gravitational wave measurements. *Physical Review Research*, 3(4):043101, November 2021. doi: 10.1103/PhysRevResearch.3.043101.
- [22] Colby Fronk and Linda Petzold. Interpretable polynomial neural ordinary differential equations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 33(4):043101, 2023. doi: 10.1063/5.0130803.
- [23] George Stepaniants, Alasdair D. Hastewell, Dominic J. Skinner, Jan F. Totz, and Jörn Dunkel. Discovering dynamics and parameters of nonlinear oscillatory and chaotic systems from partial observations. *arXiv e-prints*, art. arXiv:2304.04818, April 2023. doi: 10.48550/arXiv.2304.04818.
- [24] Shuangshuang Yin, Jianhong Wu, and Pengfei Song. Optimal control by deep learning techniques and its applications on epidemic models. *Journal of Mathematical Biology*, 87:2121–2148, Apr 2023. doi: 10.1007/s00285-023-01873-0. URL <https://link.springer.com/article/10.1007/s00285-023-01873-0>.
- [25] Vinicius V. Santana, Erbet Costa, Carine M. Rebello, Ana Mafalda Ribeiro, Chris Rackauckas, and Idelfonso B. R. Nogueira. Efficient hybrid modeling and sorption model discovery for nonlinear advection-diffusion-sorption systems: A systematic scientific machine learning approach. *arXiv e-prints*, art. arXiv:2303.13555, March 2023. doi: 10.48550/arXiv.2303.13555.
- [26] Dustin D. Nguyen. Neural Astrophysical Wind Models. *arXiv e-prints*, art. arXiv:2306.11666, June 2023. doi: 10.48550/arXiv.2306.11666.
- [27] L. L. Cowie, C. F. McKee, and J. P. Ostriker. Supernova remnant evolution in an inhomogeneous medium. I - Numerical models. *The Astrophysical Journal*, 247:908–924, Aug 1981. doi: 10.1086/159100.
- [28] Dustin D. Nguyen and Todd A. Thompson. Mass-loading and non-spherical divergence in hot galactic winds: implications for X-ray observations. *Monthly Notices of the Royal Astronomical Society*, October 2021. doi: 10.1093/mnras/stab2910.
- [29] Drummond B. Fielding and Greg L. Bryan. The Structure of Multiphase Galactic Winds. *The Astrophysical Journal*, 924(2):82, January 2022. doi: 10.3847/1538-4357/ac2f41.

- [30] Evan E. Schneider and Brant E. Robertson. CHOLLA: A New Massively Parallel Hydrodynamics Code for Astrophysical Simulation. *The Astrophysical Journal Supplement Series*, 217(2): 24, April 2015. doi: 10.1088/0067-0049/217/2/24.
- [31] Henny J. G. L. M. Lamers and Joseph P. Cassinelli. *Introduction to Stellar Winds*. 1999.
- [32] R. A. Chevalier and A. W. Clegg. Wind from a starburst galaxy nucleus. *Nature*, 317(6032): 44–45, Sep 1985. doi: 10.1038/317044a0.
- [33] Ch. Tsitouras. Runge–kutta pairs of order 5(4) satisfying only the first column simplifying assumption. *Computers & Mathematics with Applications*, 62(2):770–775, 2011. ISSN 0898-1221. doi: <https://doi.org/10.1016/j.camwa.2011.06.002>. URL <https://www.sciencedirect.com/science/article/pii/S0898122111004706>.
- [34] Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, Ali Ramadhan, and Alan Edelman. Universal differential equations for scientific machine learning, 2021.
- [35] R. A. Kopp and T. E. Holzer. Dynamics of coronal hole regions. I. Steady polytropic flows with multiple critical points. *solphys*, 49(1):43–56, July 1976. doi: 10.1007/BF00221484.
- [36] J. E. Everett, E. G. Zweibel, R. A. Benjamin, D. McCammon, L. Rocks, and J. S. Gallagher, III. The Milky Way’s Kiloparsec-Scale Wind: A Hybrid Cosmic-Ray and Thermally Driven Outflow. *The Astrophysical Journal*, 674:258–270, February 2008. doi: 10.1086/524766.
- [37] Dustin D. Nguyen and Todd A. Thompson. Galactic Winds and Bubbles from Nuclear Starburst Rings. *The Astrophysical Journal Letters*, 935(2):L24, August 2022. doi: 10.3847/2041-8213/ac86c3.
- [38] Laura A. Lopez, Smita Mathur, Dustin D. Nguyen, Todd A. Thompson, and Grace M. Olivier. Temperature and Metallicity Gradients in the Hot Gas Outflows of M82. *The Astrophysical Journal*, 904(2):152, December 2020. doi: 10.3847/1538-4357/abc010.
- [39] XRISM Science Team. Science with the X-ray Imaging and Spectroscopy Mission (XRISM). *arXiv e-prints*, art. arXiv:2003.04962, March 2020.
- [40] Ralph Kraft, Maxim Markevitch, Caroline Kilbourne, Joseph S. Adams, Hiroki Akamatsu, Mohammadreza Ayromlou, Simon R. Bandler, Marco Barbera, Douglas A. Bennett, Anil Bhardwaj, Veronica Biffi, Dennis Bodewits, Akos Bogdan, Massimiliano Bonamente, Stefano Borgani, Graziella Branduardi-Raymont, Joel N. Bregman, Joseph N. Burchett, Jenna Cann, Jenny Carter, Priyanka Chakraborty, Eugene Churazov, Robert A. Crain, Renata Cumbee, Romeel Dave, Michael DiPirro, Klaus Dolag, W. Bertrand Doriese, Jeremy Drake, William Dunn, Megan Eckart, Dominique Eckert, Stefano Ettori, William Forman, Massimiliano Galeazzi, Amy Gall, Efrain Gatuzz, Natalie Hell, Edmund Hodges-Kluck, Caitriona Jackman, Amir Jahromi, Fred Jennings, Christine Jones, Philip Kaaret, Patrick J. Kavanagh, Richard L. Kelley, Ildar Khabibullin, Chang-Goo Kim, Dimitra Koutroumpa, Orsolya Kovacs, K. D. Kuntz, Erwin Lau, Shiu-Hang Lee, Maurice Leutenegger, Sheng-Chieh Lin, Carey Lisse, Ugo Lo Cicero, Lorenzo Lovisari, Dan McCammon, Sean McEntee, Francois Mernier, Eric D. Miller, Daisuke Nagai, Michela Negro, Dylan Nelson, Jan-Uwe Ness, Paul Nulsen, Anna Ogorzalek, Benjamin D. Oppenheimer, Lidia Oskinova, Daniel Patnaude, Ryan W. Pfeifle, Annalisa Pillepich, Paul Plucinsky, David Pooley, Frederick S. Porter, Scott Randall, Elena Rasia, John Raymond, Mateusz Ruszkowski, Kazuhiro Sakai, Arnab Sarkar, Manami Sasaki, Kosuke Sato, Gerrit Schellenberger, Joop Schaye, Aurora Simionescu, Stephen J. Smith, James F. Steiner, Jonathan Stern, Yuanyuan Su, Ming Sun, Grant Tremblay, Nhut Truong, James Tutt, Eugenio Ursino, Sylvain Veilleux, Alexey Vikhlinin, Stephan Vladutescu-Zopp, Mark Vogelsberger, Stephen A. Walker, Kimberly Weaver, Dale M. Weigt, Jessica Werk, Norbert Werner, Scott J. Wolk, Congyao Zhang, William W. Zhang, Irina Zhuravleva, and John ZuHone. Line Emission Mapper (LEM): Probing the physics of cosmic ecosystems. *arXiv e-prints*, art. arXiv:2211.09827, November 2022. doi: 10.48550/arXiv.2211.09827.
- [41] Miles Cranmer, Alvaro Sanchez-Gonzalez, Peter Battaglia, Rui Xu, Kyle Cranmer, David Spergel, and Shirley Ho. Discovering symbolic models from deep learning with inductive biases. *NeurIPS 2020*, 2020.

- [42] Miles Cranmer. Interpretable Machine Learning for Science with PySR and SymbolicRegression.jl, May 2023. URL <http://arxiv.org/abs/2305.01582>. arXiv:2305.01582 [astro-ph, physics:physics].