
Causal Discovery of Latent Variables in Galactic Archaeology

Zehao Jin*^{1,2} Yuxi(Lucy) Lu*³ Yuan-Sen Ting³ Yujia Zheng⁴ Tobias Buck⁵

Abstract

Galactic archaeology—the study of stellar migration histories—provides insights into galaxy formation and evolution. However, establishing causal relationships between observable stellar properties and their birth conditions remains challenging, as key properties like birth radius are not directly observable. We employ Rank-based Latent Causal Discovery (RLCD) to uncover the causal structure governing the chemodynamics of a simulated Milky Way galaxy. Using only five observable properties (metallicity, age, and orbital parameters), we recover in a purely data-driven manner a causal graph containing two latent nodes that correspond to real physical properties: the birth radius and guiding radius of stars. Our study demonstrates the potential of causal discovery models in astrophysics.

1. Introduction

The advent of large astronomical surveys has propelled much of the current interest in applying deep learning to astronomy. However, the observational nature of astronomy—where controlled experiments are impossible—often requires us to develop a deep understanding of the mechanistic causal relations governing astrophysical systems, beyond merely identifying correlations within observed variables.

Much of the advancement in astrophysics has relied on human intervention through forward modeling to match observations, and the statistical validation of physical models to decipher causal structures between fundamental

drivers and emergent observations. This traditional approach, while successful, is inherently limited by human intuition and the complexity of the systems under study.

This raises the question: would it be possible to discover the underlying causal structure through automated causal graphical inference? Recent advances in causal discovery (e.g., [Spirtes et al., 2001](#); [Pearl, 2009](#); [Pearl et al., 2016](#))—which infers causal relationships from purely observational data—suggest this may indeed be feasible. In principle, latent variable models and directed acyclic graphs provide the statistical framework to distinguish between competing causal structures, potentially revealing relationships that human-guided analysis might overlook.

Causal discovery through probabilistic graphical models has a long history ([Geiger & Heckerman, 1994](#); [Spirtes et al., 1995; 2001](#); [Chickering, 2002](#); [Shimizu et al., 2006](#); [Zhang, 2008](#); [Huang et al., 2018](#); [Deleu et al., 2022](#)), and recently, physical sciences have begun to recognize its potential. ([Runge et al., 2019](#); [Li et al., 2020](#); [Schölkopf et al., 2021](#); [Zhang et al., 2024](#); [Yao et al., 2024](#)). These successes demonstrate that causal discovery can reveal non-obvious relationships in complex systems where traditional approaches may struggle.

While astrophysical studies have been relatively late to adopt these methods, efforts to directly infer causal structures from astronomical data are beginning to emerge, including a preliminary pilot study by [Pasquato et al. \(2023\)](#) and Bayesian analyses of causal structures underlying galaxy–supermassive black hole coevolution ([Jin et al., 2025](#)). However, due to the complexity of these astronomical systems and the indirect nature of observations, many inferred causal graph structures remain tentative and difficult to connect directly to known physical mechanisms.

To bridge this gap between causal discovery and physical understanding, we focus our investigation on galactic archaeology—the study of chemodynamic histories that shape the galaxies today. This field offers an ideal testbed for causal discovery methods because while the underlying physics is relatively well-understood, properties that govern stellar evolution—such as birth radius and historical orbital parameters—are not directly observable. These “missing” latent variables create a suitable scenario for causal discovery: we need methods that can both identify

*Equal contribution ¹Center for Astrophysics and Space Science (CASS), New York University Abu Dhabi, PO Box 129188, Abu Dhabi, UAE ²Center for Astronomy and Astrophysics and Department of Physics, Fudan University, Shanghai 200438, People’s Republic of China ³Center for Cosmology and Astroparticle Physics (CCAPP), The Ohio State University, 191 W. Woodruff Ave., Columbus, OH 43210, USA ⁴Carnegie Mellon University, Pittsburgh, PA, USA ⁵Universität Heidelberg, Interdisziplinäres Zentrum für Wissenschaftliches Rechnen, Im Neuenheimer Feld 205, Heidelberg, Germany. Correspondence to: Zehao Jin <zj448@nyu.edu>.

these hidden factors and establish their causal relationships with observable quantities.

2. Methodology

To test whether causal discovery can recover physically meaningful latent variables in galactic archaeology, we require data where both observable properties and ground-truth birth conditions are available. Simulated galaxies provide this unique opportunity, allowing us to validate the causal structures discovered by our method against known physical quantities.

2.1. Simulation Data

The simulated galaxy studied in this work is taken from the NIHAO-UHD project (Buck et al., 2018; 2020), a set of high-resolution cosmological zoomed-in hydrodynamical simulations of Milky Way-mass galaxies. NIHAO (Numerical Investigation of a Hundred Astronomical Objects) comprises 100 simulated galaxies with halo masses ranging from $\sim 10^9 - 10^{12} M_\odot$. The dark matter halos were selected from a large-scale dark matter simulation based on an isolation criterion (no similar mass companion within three virial radii at redshift $z = 0$, Wang et al., 2015).

The NIHAO-UHD simulations employ a modified version of the smoothed particle hydrodynamics (SPH) solver GASOLINE2 (Wadsley et al., 2017), with star formation and feedback modeled following (Stinson et al., 2006; 2013). These galaxies show agreement with observed Milky Way properties and local disk galaxies (Obreja et al., 2019; Buck, 2019), making them suitable analogs for our analysis. The simulations have been extensively validated in studies of MW-mass galaxies (e.g., Buck et al., 2018; 2019; Hilmi et al., 2020; Sestito et al., 2021; Obreja et al., 2022; Lu et al., 2022a;b; Wang et al., 2023).

We focus on the g2.79e12 simulation. We select disk stars with $[\text{Fe}/\text{H}] > -1$ located between 7-10 kpc at present day, a range chosen to ensure adequate sampling while matching typical observational surveys. To incorporate realistic observational uncertainties, we add uncertainties to age (10%), $[\text{Fe}/\text{H}]$ (0.02 dex), and $[\text{O}/\text{Fe}]$ (0.06 dex).

Given the distinct formation histories of the high- and low- α disk populations in the Milky Way (Bensby et al., 2014; Hayden et al., 2015), we separate stars using the criterion $[\text{O}/\text{Fe}] = -0.13[\text{Fe}/\text{H}] + 0.17$. The high- α disk formed rapidly in a turbulent, gas-rich environment, while the low- α disk formed through gradual secular evolution (Conroy et al., 2022; Xiang & Rix, 2022). This pilot study focuses on the low- α disk, where secular processes dominate and causal relationships may be clearer to identify.

For our causal discovery analysis, we use five observable

variables that trace stellar migration: metallicity $[\text{Fe}/\text{H}]$, age, and three orbital parameters—vertical action J_z , angular momentum L_z , and eccentricity e . These quantities are routinely measured in spectroscopic surveys and encode information about both stellar birth conditions and subsequent dynamical evolution.

2.2. Discovery of Causal Structures

Having established our dataset, we now turn to the causal discovery analysis. Inferring causal relationships in astrophysical systems is complicated by unobserved (latent) variables—quantities that influence our observations but cannot be directly measured. These latent variables can confound relationships between observables or mediate their interactions, making causal inference challenging.

To address this challenge, we employ the Rank-based Latent Causal Discovery (RLCD) algorithm (Dong et al., 2024), which can uncover causal structures involving latent factors within linear systems. The key insight behind RLCD is that latent variables leave statistical signatures in the relationships between observed variables—specifically, they create rank deficiencies in the covariance matrix that can be detected and interpreted.

In our framework, we model the causal relationships using a Structural Causal Model (SCM). Consider a simple example: if birth radius (latent) influences both current metallicity and angular momentum (observed), then metallicity and angular momentum will be correlated through their shared latent cause. The SCM captures such relationships mathematically through a directed acyclic graph $\mathcal{G} := (\mathbf{V}_G, \mathbf{E}_G)$, where arrows represent causal influences. Each variable V_i in this graph is generated according to a linear equation:

$$V_i = \sum_{V_j \in \text{Pa}_G(V_i)} a_{ij} V_j + \varepsilon_{V_i}, \quad (1)$$

where $\text{Pa}_G(V_i)$ denotes the variables that directly cause V_i (its “parents” in the graph), a_{ij} quantifies the strength of the causal effect from V_j to V_i , and ε_{V_i} represents random noise. The complete set of variables \mathbf{V}_G includes both observed variables \mathbf{X}_G (our five measured quantities) and latent variables \mathbf{L}_G (the hidden factors we aim to discover).

RLCD works by analyzing patterns in how observed variables covary. When latent variables are present, they constrain these covariation patterns in detectable ways. The algorithm identifies how many latent variables exist, which observed variables they influence, and whether latent variables influence each other. This enables reconstruction of a complete causal graph including both observed and hidden factors. The method has been validated on synthetic and real-world datasets (Dong et al., 2024; 2025).

After discovering the causal structure, we quantify the strength of each causal relationship by estimating the coef-

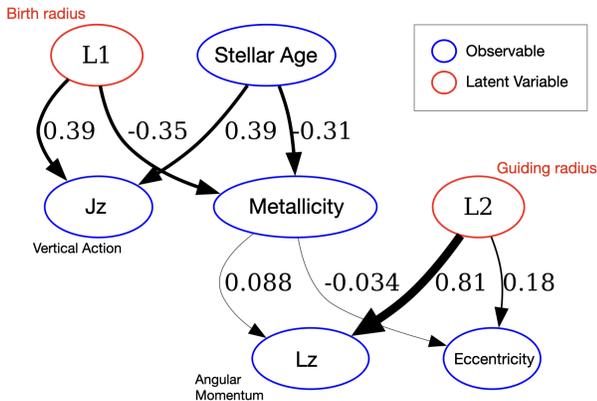


Figure 1. Causal structure discovered by RLCD from five observable stellar properties. Blue ellipses represent observed variables, while red ellipses show the two latent variables identified by the algorithm. The analysis reveals that L_1 corresponds to birth radius and L_2 to guiding radius, as validated by comparison with ground truth. Numbers on arrows indicate the strength of causal relationships (edge coefficients a_{ij}) in the linear SCM.

ficients a_{ij} . As latent variables have no inherent scale, we fix the variance of each latent variable to unity—a standard convention that allows unique parameter estimation (Dong et al., 2025). We use maximum likelihood estimation to find the parameters that best explain our observed data given the discovered causal structure.

Finally, with the causal structure and parameters determined, we estimate the latent variable values for each individual star. This involves finding the latent values that best reconstruct the observed properties according to our linear model, minimizing the prediction error.

3. Results

Applying RLCD to our five observable stellar properties yields the causal structure shown in Figure 1. The algorithm identifies two latent variables, L_1 and L_2 , whose connections to the observed variables provide insights into the underlying physics of stellar migration (Sellwood & Binney, 2002).

The discovered causal structure aligns well with our understanding of galactic chemical evolution. The latent variable L_1 influences both metallicity ($[\text{Fe}/\text{H}]$) and vertical action (J_z), while age shows similar causal connections. This pattern suggests that L_1 encodes information about stellar birth conditions—stars born at different galactic radii have distinct metallicities due to radial abundance gradients, and their vertical motions retain memory of their birth environments through the differential gravitational potential. Meanwhile, L_2 directly influences angular momentum (L_z) and eccentricity (e), the two quantities that together define a star’s guiding radius. Stars can develop eccentric

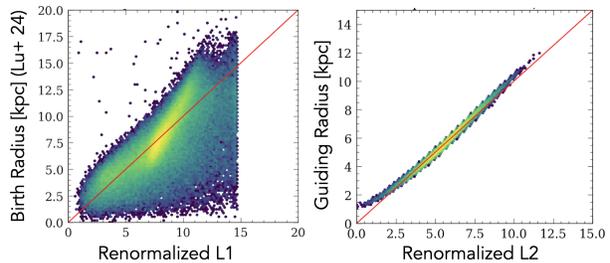


Figure 2. Validation of discovered latent variables. *Left*: Comparison between renormalized L_1 and birth radius inferred using the Lu et al. (2024) method. *Right*: True guiding radius from simulation versus renormalized L_2 . The discovered L_1 achieves comparable performance to (Lu et al., 2024) for birth radius inference, while L_2 directly recovers the true guiding radius.

orbits through gravitational scattering with giant molecular clouds—a process known as “blurring” that preserves angular momentum while increasing eccentricity.

Based on these causal relationships, we hypothesize that L_1 corresponds to birth radius (R_b) and L_2 to guiding radius (R_g). Both quantities play fundamental roles in galactic dynamics: radial migration displaces stars from their birth radii over time (Minchev et al., 2013; Sharma et al., 2021; Lu et al., 2024), while orbital heating increases random motions around the guiding radius (Sellwood & Binney, 2002). Birth radius serves as the primary indicator of radial migration because it represents the initial condition from which stars have migrated, while guiding radius tracks the secular evolution of orbital eccentricity. Traditional methods estimate birth radii from age and metallicity under the assumption that the metallicity gradient is linear (Minchev et al., 2018; Lu et al., 2024), while guiding radii can be calculated directly from angular momentum. The causal structure discovered by RLCD independently recovers these known relationships.

To test our hypothesis, we leverage the simulation’s complete stellar histories. After estimating individual stellar values for L_1 and L_2 using the identified parameters, we map these dimensionless quantities to physical units via polynomial transformations, accounting for the centering and standardization in our analysis. For L_1 , we compare against birth radii inferred using the established method of Lu et al. (2024), which estimates birth radii from $[\text{Fe}/\text{H}]$ and age based on assumptions from (Lu et al., 2022a). For L_2 , we can directly compare against the true guiding radius from the simulation.

Figure 2 presents the validation results. For guiding radius (right panel), the agreement between L_2 and the ground truth is notable, with minimal scatter about the one-to-one relation. This tight correlation confirms that L_2 captures guiding radius—unsurprising given that guiding radius is directly encoded in the angular momentum that L_2 influences. For birth radius (left panel), we compare L_1 against the Lu

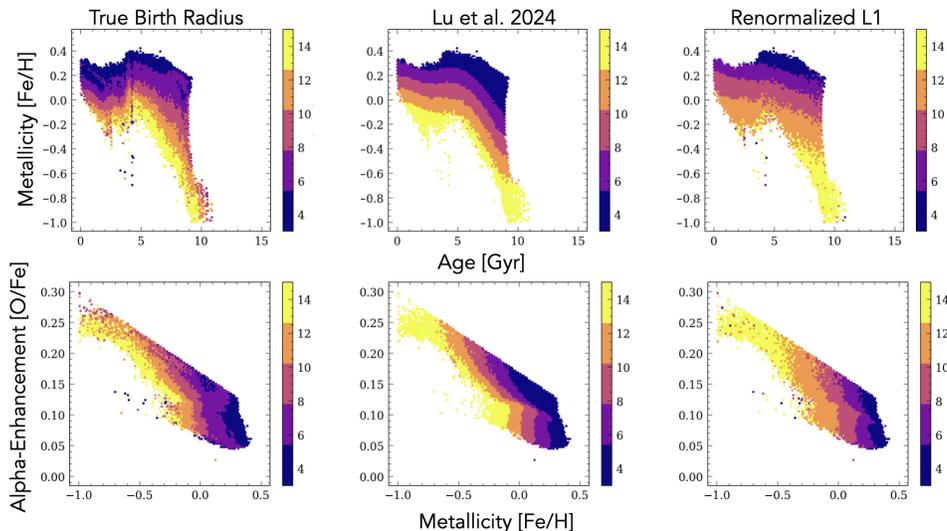


Figure 3. Distribution of birth radii in chemical abundance space. *Left*: True birth radius from simulation. *Middle*: Birth radius inferred using the Lu et al. (2024) method. *Right*: Renormalized L_1 from causal discovery. Top row shows the age-metallicity plane, bottom row shows the alpha-abundance plane. All three methods exhibit consistent patterns: younger, metal-rich stars originate from smaller galactic radii (yellow) while older, metal-poor stars come from larger radii (blue), reflecting inside-out galaxy formation. The similarity across columns validates that the discovered latent variable L_1 captures birth radius information despite using no prior knowledge.

et al. method. This consistency between two independent approaches—one theory-driven, one data-driven—suggests both methods capture similar physical information about stellar birth conditions.

Figure 3 provides further validation by examining birth radius distributions across chemical abundance space. All three approaches—ground truth, Lu et al. inference, and discovered L_1 —show consistent patterns. In the age-[Fe/H] plane (top row), younger, metal-rich stars consistently appear at smaller birth radii, while older, metal-poor stars occupy larger radii, reflecting the inside-out growth of the galactic disk. The [O/Fe]-[Fe/H] plane (bottom row) reveals how radial abundance gradients at different epochs imprint on stellar chemistry. The discovered L_1 reproduces these chemical evolution patterns without incorporating any prior knowledge, demonstrating that RLCD extracts physically meaningful information from the data.

These results demonstrate that causal discovery can recover physically meaningful latent variables from observational data alone. The method not only identifies the correct number of hidden factors but also determines their proper causal relationships with observables. For guiding radius, the near-perfect recovery validates our approach. For birth radius, the agreement with established methods and consistency across chemical spaces confirms that causal discovery captures the same underlying physics that astronomers have uncovered through theoretical work. While we have assumed linear relationships for this proof-of-concept, future work with nonlinear causal discovery methods may further improve the reconstruction of stellar birth properties.

4. Broader Impact

This work demonstrates that causal discovery can uncover physically meaningful relationships in astrophysical systems from observational data alone. By successfully identifying birth radius and guiding radius as the latent variables governing stellar migration—without incorporating prior astrophysical knowledge—we validate the potential of automated causal inference in astronomy. The discovered structure not only aligns with decades of theoretical understanding but achieves predictive performance comparable to traditional domain-specific methods. This suggests that causal discovery can complement human-guided theoretical approaches, potentially revealing relationships that intuition might overlook in increasingly complex astronomical datasets.

Looking forward, our results establish a foundation for applying causal discovery across astronomy, where many phenomena involve unobservable latent variables—from dark matter shaping galaxy dynamics to stellar interiors driving evolution. As astronomical surveys grow in scale and dimensionality, such data-driven methods become essential for extracting physical insights. Future work extending to broader stellar populations and real observations will face additional challenges from selection effects and measurement uncertainties. The combination of causal discovery with flexible machine learning architectures promises interpretable AI systems that both predict and explain, guided by discovered causal structures rather than opaque correlations—moving toward a new paradigm where data-driven methods and physical understanding mutually reinforce our exploration of the physical processes that engender the universe.

References

- Bensby, T., Feltzing, S., and Oey, M. S. Exploring the Milky Way stellar disk. A detailed elemental abundance study of 714 F and G dwarf stars in the solar neighbourhood. *Astronomy & Astrophysics*, 562:A71, February 2014. doi: 10.1051/0004-6361/201322631.
- Buck, T. On the origin of the chemical bimodality of disc stars: a tale of merger and migration. *Monthly Notices of the Royal Astronomical Society*, 491(4): 5435–5446, 11 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz3289. URL <https://doi.org/10.1093/mnras/stz3289>.
- Buck, T., Ness, M. K., Macciò, A. V., Obreja, A., and Dutton, A. A. Stars Behind Bars. I. The Milky Way’s Central Stellar Populations. *Astrophysical Journal*, 861(2):88, July 2018. doi: 10.3847/1538-4357/aac890.
- Buck, T., Ness, M., Obreja, A., Macciò, A. V., and Dutton, A. A. Stars behind Bars II: A Cosmological Formation Scenario for the Milky Way’s Central Stellar Structure. *Astrophysical Journal*, 874(1):67, March 2019. doi: 10.3847/1538-4357/aaffd0.
- Buck, T., Obreja, A., Macciò, A. V., Minchev, I., Dutton, A. A., and Ostriker, J. P. NIHAO-UHD: the properties of MW-like stellar discs in high-resolution cosmological simulations. *Monthly Notices of the Royal Astronomical Society*, 491(3):3461–3478, January 2020. doi: 10.1093/mnras/stz3241.
- Chickering, D. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, 01 2002. doi: 10.1162/153244303321897717.
- Conroy, C., Weinberg, D. H., Naidu, R. P., Buck, T., Johnson, J. W., Cargile, P., Bonaca, A., Caldwell, N., Chandra, V., Han, J. J., Johnson, B. D., Speagle, J. S., Ting, Y.-S., Woody, T., and Zaritsky, D. Birth of the Galactic Disk Revealed by the H3 Survey. *arXiv e-prints*, art. arXiv:2204.02989, April 2022. doi: 10.48550/arXiv.2204.02989.
- Deleu, T., Góis, A., Emezue, C. C., Rankawat, M., Lacoste-Julien, S., Bauer, S., and Bengio, Y. Bayesian structure learning with generative flow networks. In *The 38th Conference on Uncertainty in Artificial Intelligence*, 2022. URL <https://openreview.net/forum?id=HElfed8j9g9>.
- Dong, X., Huang, B., Ng, I., Song, X., Zheng, Y., Jin, S., Legaspi, R., Spirtes, P., and Zhang, K. A versatile causal discovery framework to allow causally-related hidden variables. In *The Twelfth International Conference on Learning Representations*, 2024.
- Dong, X., Ng, I., Huang, B., Sun, Y., Jin, S., Legaspi, R., Spirtes, P., and Zhang, K. On the parameter identifiability of partially observed linear causal models. *Advances in Neural Information Processing Systems*, 37:30740–30771, 2025.
- Geiger, D. and Heckerman, D. Learning gaussian networks. In *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, UAI’94, pp. 235–243, San Francisco, CA, USA, 1994. Morgan Kaufmann Publishers Inc. ISBN 1558603328.
- Hayden, M. R., Bovy, J., Holtzman, J. A., Nidever, D. L., Bird, J. C., Weinberg, D. H., Andrews, B. H., Majewski, S. R., Allende Prieto, C., Anders, F., Beers, T. C., Bizyaev, D., Chiappini, C., Cunha, K., Frinchaboy, P., García-Hernández, D. A., García Pérez, A. E., Girardi, L., Harding, P., Hearty, F. R., Johnson, J. A., Mészáros, S., Minchev, I., O’Connell, R., Pan, K., Robin, A. C., Schiavon, R. P., Schneider, D. P., Schultheis, M., Shetrone, M., Skrutskie, M., Steinmetz, M., Smith, V., Wilson, J. C., Zamora, O., and Zasowski, G. Chemical Cartography with APOGEE: Metallicity Distribution Functions and the Chemical Structure of the Milky Way Disk. *Astrophysical Journal*, 808(2):132, August 2015. doi: 10.1088/0004-637X/808/2/132.
- Hilmi, T., Minchev, I., Buck, T., Martig, M., Quillen, A. C., Monari, G., Famaey, B., de Jong, R. S., Laporte, C. F. P., Read, J., Sanders, J. L., Steinmetz, M., and Wegg, C. Fluctuations in galactic bar parameters due to bar-spiral interaction. *Monthly Notices of the Royal Astronomical Society*, 497(1):933–955, September 2020. doi: 10.1093/mnras/staa1934.
- Huang, B., Zhang, K., Lin, Y., Schölkopf, B., and Glymour, C. Generalized score functions for causal discovery. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD ’18, pp. 1551–1560, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355520. doi: 10.1145/3219819.3220104. URL <https://doi.org/10.1145/3219819.3220104>.
- Jin, Z., Pasquato, M., Davis, B. L., Deleu, T., Luo, Y., Cho, C., Lemos, P., Perreault-Levasseur, L., Bengio, Y., Kang, X., Macciò, A. V., and Hezaveh, Y. Causal discovery in astrophysics: Unraveling supermassive black hole and galaxy coevolution. *The Astrophysical Journal*, 979(2):212, jan 2025. doi: 10.3847/1538-4357/ad9ded. URL <https://dx.doi.org/10.3847/1538-4357/ad9ded>.
- Li, Y., Torralba, A., Anandkumar, A., Fox, D., and Garg, A. Causal discovery in physical systems from videos. *Advances in Neural Information Processing Systems*, 33: 9180–9192, 2020.

- Lu, Y., Buck, T., Minchev, I., and Ness, M. K. Reliability and limitations of inferring birth radii in the Milky Way disc. *Monthly Notices of the Royal Astronomical Society*, 515(1):L34–L38, September 2022a. doi: 10.1093/mnras/ slac065.
- Lu, Y. L., Ness, M. K., Buck, T., and Carr, C. Turning points in the age-metallicity relations - created by late satellite infall and enhanced by radial migration. *Monthly Notices of the Royal Astronomical Society*, 512(4):4697–4714, June 2022b. doi: 10.1093/mnras/stac780.
- Lu, Y. L., Minchev, I., Buck, T., Khoperskov, S., Steinmetz, M., Libeskind, N., Cescutti, G., Freeman, K. C., and Ratcliffe, B. There is no place like home - finding birth radii of stars in the Milky Way. *Monthly Notices of the Royal Astronomical Society*, 535(1):392–405, November 2024. doi: 10.1093/mnras/stae2364.
- Minchev, I., Chiappini, C., and Martig, M. Chemodynamical evolution of the Milky Way disk. I. The solar vicinity. *Astronomy & Astrophysics*, 558:A9, October 2013. doi: 10.1051/0004-6361/201220189.
- Minchev, I., Anders, F., Recio-Blanco, A., Chiappini, C., de Laverny, P., Queiroz, A., Steinmetz, M., Adibekyan, V., Carrillo, I., Cescutti, G., Guiglion, G., Hayden, M., de Jong, R. S., Kordopatis, G., Majewski, S. R., Martig, M., and Santiago, B. X. Estimating stellar birth radii and the time evolution of Milky Way’s ISM metallicity gradient. *Monthly Notices of the Royal Astronomical Society*, 481(2):1645–1657, December 2018. doi: 10.1093/mnras/ sty2033.
- Obreja, A., Dutton, A. A., Macciò, A. V., Moster, B., Buck, T., van de Ven, G., Wang, L., Stinson, G. S., and Zhu, L. NIHAO XVI: the properties and evolution of kinematically selected discs, bulges, and stellar haloes. *Monthly Notices of the Royal Astronomical Society*, 487(3):4424–4456, August 2019. doi: 10.1093/mnras/stz1563.
- Obreja, A., Buck, T., and Macciò, A. V. A first estimate of the Milky Way dark matter halo spin. *Astronomy & Astrophysics*, 657:A15, January 2022. doi: 10.1051/0004-6361/202140983.
- Pasquato, M., Jin, Z., Lemos, P., Davis, B. L., and Macciò, A. V. Causa prima: cosmology meets causal discovery for the first time. *arXiv e-prints*, art. arXiv:2311.15160, November 2023. doi: 10.48550/arXiv.2311.15160.
- Pearl, J. *Causality*. Cambridge university press, 2009.
- Pearl, J., Glymour, M., and Jewell, N. P. *Causal inference in statistics : a primer*. Wiley, Chichester, West Sussex, 2016. ISBN 9781119186854.
- Runge, J., Bathiany, S., Bollt, E., Camps-Valls, G., Coumou, D., Deyle, E., Glymour, C., Kretschmer, M., Mahecha, M. D., Muñoz-Marí, J., et al. Inferring causation from time series in earth system sciences. *Nature communications*, 10(1):2553, 2019.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. Toward causal representation learning. *Proceedings of the IEEE*, 109(5): 612–634, 2021.
- Sellwood, J. A. and Binney, J. J. Radial mixing in galactic discs. *Monthly Notices of the Royal Astronomical Society*, 336(3):785–796, November 2002. doi: 10.1046/j.1365-8711.2002.05806.x.
- Sestito, F., Buck, T., Starkenburg, E., Martin, N. F., Navarro, J. F., Venn, K. A., Obreja, A., Jablonka, P., and Macciò, A. V. Exploring the origin of low-metallicity stars in Milky-Way-like galaxies with the NIHAO-UHD simulations. *Monthly Notices of the Royal Astronomical Society*, 500(3):3750–3762, January 2021. doi: 10.1093/mnras/ staa3479.
- Sharma, S., Hayden, M. R., and Bland-Hawthorn, J. Chemical enrichment and radial migration in the Galactic disc - the origin of the $[\alpha\text{Fe}]$ double sequence. *Monthly Notices of the Royal Astronomical Society*, 507(4):5882–5901, November 2021. doi: 10.1093/mnras/stab2015.
- Shimizu, S., Hoyer, P. O., Hyvrinen, A., and Kerminen, A. J. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(4):2003–2030, 2006.
- Spirtes, P., Meek, C., and Richardson, T. Causal inference in the presence of latent variables and selection bias. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 499–506, August 1995.
- Spirtes, P., Glymour, C., and Scheines, R. *Causation, Prediction, and Search*. The MIT Press, 01 2001. ISBN 9780262284158. doi: 10.7551/mitpress/1754.001.0001. URL <https://doi.org/10.7551/mitpress/1754.001.0001>.
- Stinson, G., Seth, A., Katz, N., Wadsley, J., Governato, F., and Quinn, T. Star formation and feedback in smoothed particle hydrodynamic simulations - I. Isolated galaxies. *Monthly Notices of the Royal Astronomical Society*, 373(3):1074–1090, December 2006. doi: 10.1111/j.1365-2966.2006.11097.x.
- Stinson, G. S., Bovy, J., Rix, H. W., Brook, C., Roškar, R., Dalcanton, J. J., Macciò, A. V., Wadsley, J., Couchman, H. M. P., and Quinn, T. R. MaGICC thick disc - I. Comparing a simulated disc formed with stellar feedback

to the Milky Way. *Monthly Notices of the Royal Astronomical Society*, 436(1):625–634, November 2013. doi: 10.1093/mnras/stt1600.

Wadsley, J. W., Keller, B. W., and Quinn, T. R. Gasoline2: a modern smoothed particle hydrodynamics code. *Monthly Notices of the Royal Astronomical Society*, 471(2):2357–2369, October 2017. doi: 10.1093/mnras/stx1643.

Wang, K., Carrillo, A., Ness, M. K., and Buck, T. The individual abundance distributions of disc stars across birth radii in GALAH. *arXiv e-prints*, art. arXiv:2307.04724, July 2023. doi: 10.48550/arXiv.2307.04724.

Wang, L., Dutton, A. A., Stinson, G. S., Macciò, A. V., Penzo, C., Kang, X., Keller, B. W., and Wadsley, J. NIHAO project - I. Reproducing the inefficiency of galaxy formation across cosmic time with a large sample of cosmological hydrodynamical simulations. *Monthly Notices of the Royal Astronomical Society*, 454(1):83–94, November 2015. doi: 10.1093/mnras/stv1937.

Xiang, M. and Rix, H.-W. A time-resolved picture of our Milky Way’s early formation history. *Nature*, 603(7902):599–603, March 2022. doi: 10.1038/s41586-022-04496-5.

Yao, D., Muller, C., and Locatello, F. Marrying causal representation learning with dynamical systems for science. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Zhang, J. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16-17): 1873–1896, 2008.

Zhang, K., Xie, S., Ng, I., and Zheng, Y. Causal representation learning from multiple distributions: A general setting. In *International Conference on Machine Learning*, pp. 60057–60075. PMLR, 2024.