

---

# GalacticFlow: Learning a Generalized Representation of Galaxies with Normalizing Flows

---

**Luca Wolf**

Interdisciplinary Center  
for Scientific Computing,  
University of Heidelberg,  
Im Neuenheimer Feld 205,  
D-69120 Heidelberg

luca.wolf@stud.uni-heidelberg.de

**Tobias Buck**

Interdisciplinary Center  
for Scientific Computing,  
University of Heidelberg,  
Im Neuenheimer Feld 205,  
D-69120 Heidelberg

tobias.buck@iwr.uni-heidelberg.de

## Abstract

State-of-the-art galaxy formation simulations generate data within weeks or months. Their results consist of a random sub-sample of possible galaxies with a fixed number of stars. We propose a ML based method, `GalacticFlow`, that generalizes such results. We use normalizing flows to learn the extended distribution function of galaxies conditioned on global galactic parameters. `GalacticFlow` then provides a continuized and condensed representation of the ensemble of galaxies in the data. Thus, essentially compressing large amounts of explicit simulation data into a small implicit generative model. Our model is able to evaluate any galaxy eDF given by a set of global parameters and allows generating arbitrarily many stars from it. We show that we can learn such a representation, embodying the entire mass range from dwarf to Milky Way mass, from only 90 galaxies in  $\sim 18$  hours on a single RTX 2080Ti and generate a new galaxy of one million stars within a few seconds.

## 1 Motivation

In astrophysics, distribution functions (DFs) play a crucial role in modeling the properties and behavior of galaxies. DFs provide a mathematical description to characterize the statistical properties of the stars in a galaxy. In galactic dynamics, the DF, often denoted as  $f(\boldsymbol{x}, \boldsymbol{v})$ , describes the probability of finding a star with specific phase space coordinates within the galaxy [1]. They have been widely used to describe the dynamics of collisionless systems such as galaxies via Boltzmann's equation [2] and allow to calculate various properties, such as the mass, velocity dispersion, and density profile via integration of the DF over different regions of phase space. Since galaxies typically consist of various populations of stars with different properties, such as age or chemical composition (metallicity) but similar orbit distributions the concept of extended DFs has been introduced [22]. Those eDFs combine the phase space distribution with additional stellar properties in order to model their joint distribution and allow for more complex relations among stars from different populations. In general, those eDFs can become quite complex especially for galaxies with a wide range of stellar populations and complex dynamical histories such that they become analytically intractable. Here we draw on the fact that ML concepts such as normalizing Flows (NFs) are able to describe such complex DFs. Furthermore, automatic differentiation allows to easily compute gradients of the approximated DFs enabling to feed the neural DFs into existing mathematical modelling frameworks such as Boltzmann's equation. Normalizing




Figure 1: Overview of the GalacticFlow network. We use global galaxy properties  $p_i$  as input conditions  $\theta$  to the conditional normalizing flow (cNF),  $p(\varphi, \varpi | \theta)$ , to generate extended distribution functions (eDFs)  $p_\theta(\varphi, \varpi)$  of galaxies. Where  $\varphi$  denotes the phase space coordinates and  $\varpi$  the additional stellar quantities like age and metallicity.

## 2 Related Work: Normalizing flows in astrophysics

NFs are a machine learning (ML) tool, that allows to learn complex, high-dimensional, multi-modal probability distribution functions (pdf) from given data points  $x_i \in \mathbb{R}^n$ . Besides their study in probabilistic ML and application to many general tasks, such as image generation, they have recently been found useful in a wide range of physics problems. Particularly in astrophysics, they are used to model the phase space density of stars and infer the Galactic dark matter density and gravitational potential via Boltzmann’s equation [6, 14]. [8] applied NFs to describe the color-magnitude diagram of the Gaia data, and [21] used conditional NF models (cNFs) to model quasar continua while [23] used them to model the element distribution of stars. Here we go one step further and not only learn the 6d phase space density of stars, but their 10d eDF that further covers age, metallicity, and oxygen/iron abundance. We describe galactic DFs via NFs that are trained on state-of-the-art simulations of galaxy formation to directly learn the stellar eDFs from high quality data. In particular, we use cNFs to describe the eDFs conditioned on global galaxy properties such as total stellar mass, average metallicity or star formation history, that provide additional contextual information when sampling. As such, we describe a population of galactic DFs with a single cNF model, illustrated in Fig. 1. The idea behind GalacticFlow is based on the fact that galaxy morphology is linked to underlying physical parameters via various scaling relations, most famously, the Mass-Metallicity [13, 24], Tully-Fisher [25] and Faber-Jackson relation [12]. Furthermore, galaxies grow hierarchically, i.e. by merging with others, resulting in older galaxies being heavier and more elliptical, while younger galaxies are lighter and more spiral-like giving rise to the famous Hubble sequence [16]. Learning these different DFs with a NF promises a compressed model, that allows generating arbitrarily many stars from exact DFs of arbitrary galaxies in very short times. In comparison to other generative models, this further allows direct evaluations of those DFs.

## 3 Method

A normalizing flow constructs a diffeomorphism  $f^{-1}$ , such that  $y_i = f^{-1}(x_i)$  follows a simple base distribution [e.g. a gaussian 18, 20]. The change of variable theorem relates  $f^{-1}$  to the data’s modeled pdf via:  $\log p(x) = \log p'(f^{-1}(x)) + \log |J_{f^{-1}}(x)|$ , where  $p'$  is the prior distribution in  $y$ -space. To find  $f^{-1}$ , that yields Gaussian distributed  $y$ -values, one sets  $p' = \mathcal{N}(0, 1)$  and maximizes  $\mathcal{L} = \log p(x)$ , the mean log likelihood of the data. This is equivalent to minimizing the Kullback–Leibler divergence of  $p_{\text{true}}$  from  $p$  [20]. After training, we obtain an exact modeled pdf and are able to generate arbitrarily large samples by drawing  $y \sim \mathcal{N}(0, 1)$  and applying  $f$ .  $f^{-1}$  is composed of several functions  $f_1^{-1}, \dots, f_m^{-1}$  that use the chain rule to compute  $\log p(x)$ . These functions can be chosen as simple parametric functions, e.g. splines, with neural networks (NNs) determining their parameters individually per data point. Coupling layers [e.g. 9, 10], maintain the first half of dimensions and input them into NNs for the transformation’s parameters of the second




Figure 2: Corner plot of the 10D stellar parameter space. The diagonal shows a Gaussian KDE of the respective marginal distribution with data from the validation set in blue and model in orange. The off-diagonal shows 2D histograms of ground truth data in the bottom left corner and the corresponding GalacticFlow model the top right corner. Note that the images for the samples we have flipped x and y axis to have the same orientation as the data, for better comparison. They then share the same axis and the same histogram grid with the respective data plot.

half. This is followed by a permutation of the dimensions and repeated several times and satisfies the requirements of a diffeomorphism with efficiently computeable inverse and Jacobian determinant.

### 3.1 Model architecture

Our NF architecture uses very flexible rational-quadratic spline coupling layers [11] alternating with invertible  $1 \times 1$  convolutions as generalized permutations [17] as presented by [11]. We also adapt "double coupling layers", as used by [23] which transforms the second half of dimensions with parameters determined from the first half followed by a transformation of the first half with the new second half. This choice enables both forward ( $f^{-1}$ ) and backward ( $f$ ) transformations in a single pass through the NNs, contrasting with auto-regressive methods. This allows training and sampling to be fast at the same time. This base NF can easily be expanded to learn conditional pdfs. To achieve this, we follow [23] and input conditional values into a second NN and multiply the outputs of both networks to determine the transformation parameters. These conditional values must be provided for every point in the training set and as necessary additional information for each sample drawn. We choose 14 flow layers (convolution layer followed by coupling layer),  $K = 10$  spline bins per coupling layer and 3 hidden layers with each 128 neurons for the coupling layer NNs. In our tests we found this architecture to be working best, as smaller models struggled to fit the data while larger models were over fitting. However, we did not perform extensive cross-testing of all hyper-parameters. Our final cNF model operates on 10D data and employs a 4D condition space that we describe further below. With this architecture the model has about  $3 \cdot 10^6$  trainable parameters, taking about 11MB of disk space.

### 3.2 Training

To demonstrate our model, we train a GalacticFlow model on simulation data from the NIHAO project [3–5, 26] that models 95 galaxies from dwarf galaxy to the Milky Way mass scale. From this simulation suite we extract for each galaxy the 10 dimensional joint distribution of stellar position  $\mathbf{x}$ , velocity  $\mathbf{v}$ , total metallicity  $Z$ , iron and oxygen abundance  $[\text{Fe}/\text{H}]$ ,  $[\text{O}/\text{Fe}]$  and stellar age  $\tau$ . Additionally, we record easy to observe galactic properties such as total stellar mass  $M_*$ , median stellar age  $\tau_{50}$ , the 10-th percentile stellar age  $\tau_{10}$  and the average stellar metallicity  $Z$  for each galaxy to be used as input conditions to the cNF. We normalize the data to zero mean and unity standard deviation. Total stellar mass is transformed to log space before normalizing, for a more natural spacing. The inverse of this transformation needs to be applied after sampling and the Jacobian needs to be respected in case of pdf evaluation. For training we use the RAdam optimizer [19] and an exponential learning rate schedule with an initial learning rate of  $\text{lr} = 9 \cdot 10^{-5}$  and a decrease of a factor  $\gamma = 0.998$  every 10 and once the learning rate falls below  $3 \cdot 10^{-6}$ , every 120 optimisation steps. We train 20 different models, where in each model 5 randomly chosen galaxies are left out from the training data and remain reserved for validation only. Using a batch size of 1024, we train for 10 epochs on a single RTX 2080Ti, which takes about 18 hours of wall clock time. Drawing a sample from this flow takes about 3.5s per GPU per million stars, in the limit of large sample sizes.

## 4 GalacticFlow: Learning a generalized generative model for galaxies

We evaluate the the quality of the GalacticFlow model by comparing 100 generated galaxy samples with their ground truth data from the validation set. In this way we can judge how the model performs in providing a generalized representation that interpolates in galactic parameter space. A "generated galaxy" in this context refers to a sample from the flow with conditions matching those of a ground truth galaxy, including the same number of stars. It is vital to note, that we neither expect nor intend a perfect agreement between a GalacticFlow model and a real galaxy. After all, each galaxy possesses a unique merger history and distinctive characteristics that cannot be generalized with only 4 conditions. Thus, our objective with GalacticFlow is to capture general trends in galactic eDFs rather than memorizing individual formation histories. An overly close agreement would only be possible on an over-fitted training galaxy. In Fig. 2 we show a qualitative

comparison of one example galaxy from the validation set with its GalacticFlow model in a corner plot. The agreement of model and data can be deemed quite good. GalacticFlow demonstrates the capacity to interpolate and generate realistic galaxies for parameter sets not seen during training. Across the entire mass range, it successfully captures various shapes of stellar eDFs. However, some challenges are evident in reproducing sharp peaks, which may relate to the naive sampling procedure [7]. We anticipate that model quality and its ability to recognize trends will improve considerably with a more diverse training set. Finding a metric to quantitatively evaluate the performance of our model is challenging, as we need a sample-based difference measure between two 10D distributions. In the spirit of our model we also require the metric to measure 'physical closeness', at least in a sense that, e.g., galaxies of similar mass are deemed close. We choose to employ linear maximum mean discrepancy [MMD 15] as we found it to fulfill these requests. We generate a sample of an unseen galaxy and compute the linear MMD to all 90 data galaxies, recording the difference in log stellar mass  $\Delta \log M_{\text{stars}}$  between model and data galaxy. We repeat this process for all 5 unseen galaxies of all 20 models, to calculate a relation of statistical distance as a function of  $\Delta \log M_{\text{stars}}$  as shown in Fig. 3. This relation enables us to evaluate the quality of generalization and the recognition of general, known trends by the model. We show here MMD vs.  $\Delta \log M_{\text{stars}}$  as stellar mass is a primary condition and a key predictor of galaxy morphology. The clear trend of increasing MMD with increasing magnitude of  $\Delta \log M_{\text{stars}}$  and the statistical similarity between model and data




Figure 3: Linear maximum mean discrepancy vs. difference in log stellar mass  $\Delta \log M_{\text{stars}}$  between GalacticFlow models and data galaxies (orange points) and corresponding ground truth galaxies and data galaxies (contours).

(points vs. contours) indicates that GalacticFlow eDFs align well with those of real galaxies of similar mass. Thus, the model effectively learns a meaningful generalization of galactic parameter space.

## 5 Conclusion

We find that GalacticFlow can learn a generalized representation of galaxies with conditional normalizing flows, that captures and respects the underlying galactic parameter space very well and interpolates it nicely. After a moderate training time ( $\sim 18$ h), we obtain a compact generative model ( $3 \cdot 10^6$  trainable parameters,  $\sim 11$ MB), that allows to generate realistic galaxies in any desired size for any desired set of galactic parameters in short times (3.5s per GPU per million stars).

Our code is publicly available on GitHub<sup>1</sup> and the galaxy data can be found on Zenodo<sup>2</sup>. This includes our algorithms for pre-processing the data as well as the source code for the cNF. The code is documented and can easily adapted for reuse. Furthermore we provide a user friendly API for loading our pre-trained models, training your own models and generating/sampling galaxies as desired.

## Broader impact statement

The authors are not aware of any immediate ethical or societal implications of this work. This work purely aims to aid scientific modelling and proposes to apply normalizing flows to learn the distribution function of stars in galaxies. Thinking more broadly, the concept of GalacticFlow is to extract and interpolate high-dimensional data from spares simulation results to facilitate faster and more versatile modelling. This might help to recognize how much information is already contained in legacy data that can be made accessible for better modelling and as such help to save energy by running less expensive simulations.

## References

- [1] James Binney. Distribution functions for the Milky Way. *Monthly Notices of the Royal Astronomical Society*, 401(4):2318–2330, February 2010. doi: 10.1111/j.1365-2966.2009.15845.x.
- [2] James Binney and Scott Tremaine. *Galactic Dynamics: Second Edition*. 2008.
- [3] Tobias Buck. 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, February 2020. doi: 10.1093/mnras/stz3289.
- [4] Tobias Buck, Melissa Ness, Aura Obreja, Andrea V. Macciò, and Aaron A. Dutton. Stars behind Bars II: A Cosmological Formation Scenario for the Milky Way’s Central Stellar Structure. *ApJ*, 874(1):67, March 2019. doi: 10.3847/1538-4357/aaffd0.
- [5] Tobias Buck, Aura Obreja, Andrea V. Macciò, Ivan Minchev, Aaron A. Dutton, and Jeremiah P. Ostriker. NIHAO-UHD: the properties of MW-like stellar discs in high-resolution cosmological simulations. *MNRAS*, 491(3):3461–3478, January 2020. doi: 10.1093/mnras/stz3241.
- [6] Matthew R Buckley, Sung Hak Lim, Eric Putney, and David Shih. Measuring galactic dark matter through unsupervised machine learning. *Monthly Notices of the Royal Astronomical Society*, 521(4):5100–5119, 3 2023. doi: 10.1093/mnras/stad843. URL <https://doi.org/10.1093/mnras/stad843>.
- [7] Florentin Coeurdoux, Nicolas Dobigeon, and Pierre Chainais. Normalizing flow sampling with langevin dynamics in the latent space, 2023.
- [8] Miles D. Cranmer, Richard Galvez, Lauren Anderson, David N. Spergel, and Shirley Ho. Modeling the Gaia Color-Magnitude Diagram with Bayesian Neural Flows to Constrain Distance Estimates. *arXiv e-prints*, art. arXiv:1908.08045, August 2019. doi: 10.48550/arXiv.1908.08045.

---

<sup>1</sup><https://github.com/luwo9/GalacticFlow>

<sup>2</sup><https://doi.org/10.5281/zenodo.8389554>

- [9] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation, 2015.
- [10] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp, 2017.
- [11] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows, 2019.
- [12] S. M. Faber and R. E. Jackson. Velocity dispersions and mass-to-light ratios for elliptical galaxies. *ApJ*, 204:668–683, March 1976. doi: 10.1086/154215.
- [13] Anna Gallazzi, Stéphane Charlot, Jarle Brinchmann, Simon D. M. White, and Christy A. Tremonti. The ages and metallicities of galaxies in the local universe. *Monthly Notices of the Royal Astronomical Society*, 362(1):41–58, September 2005. doi: 10.1111/j.1365-2966.2005.09321.x.
- [14] Gregory M. Green, Yuan-Sen Ting, and Harshil Kamdar. Deep potential: Recovering the gravitational potential from a snapshot of phase space. *The Astrophysical Journal*, 942(1):26, jan 2023. doi: 10.3847/1538-4357/aca3a7. URL <https://doi.org/10.3847/1538-4357/aca3a7>.
- [15] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012. URL <http://jmlr.org/papers/v13/gretton12a.html>.
- [16] E. P. Hubble. Extragalactic nebulae. *ApJ*, 64:321–369, December 1926. doi: 10.1086/143018.
- [17] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions, 2018.
- [18] Ivan Kobyzev, Simon J. D. Prince, and Marcus A. Brubaker. Normalizing Flows: An Introduction and Review of Current Methods. *arXiv e-prints*, art. arXiv:1908.09257, August 2019. doi: 10.48550/arXiv.1908.09257.
- [19] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond, 2021.
- [20] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference, 2021.
- [21] David M. Reiman, John Tamanas, J. Xavier Prochaska, and Dominika Āurovčíková. Fully probabilistic quasar continua predictions near Lyman- $\alpha$  with conditional neural spline flows. *arXiv e-prints*, art. arXiv:2006.00615, May 2020. doi: 10.48550/arXiv.2006.00615.
- [22] Jason L. Sanders and James Binney. Extended distribution functions for our Galaxy. *Monthly Notices of the Royal Astronomical Society*, 449(4):3479–3502, June 2015. doi: 10.1093/mnras/stv578.
- [23] Yuan-Sen Ting and David H. Weinberg. How many elements matter? *The Astrophysical Journal*, 927(2):209, 3 2022. doi: 10.3847/1538-4357/ac5023. URL <https://doi.org/10.3847/1538-4357/ac5023>.
- [24] Christy A. Tremonti, Timothy M. Heckman, Guinevere Kauffmann, Jarle Brinchmann, Stéphane Charlot, Simon D. M. White, Mark Seibert, Eric W. Peng, David J. Schlegel, Alan Uomoto, Masataka Fukugita, and Jon Brinkmann. The Origin of the Mass-Metallicity Relation: Insights from 53,000 Star-forming Galaxies in the Sloan Digital Sky Survey. *ApJ*, 613(2):898–913, October 2004. doi: 10.1086/423264.
- [25] R. B. Tully and J. R. Fisher. A new method of determining distances to galaxies. *A&A*, 54: 661–673, February 1977.
- [26] Liang Wang, Aaron A. Dutton, Gregory S. Stinson, Andrea V. Macciò, Camilla Penzo, Xi Kang, Ben W. Keller, and James Wadsley. NIHAO project - I. Reproducing the inefficiency of galaxy formation across cosmic time with a large sample of cosmological hydrodynamical simulations. *MNRAS*, 454(1):83–94, November 2015. doi: 10.1093/mnras/stv1937.