Inference of Star Formation and Metallicity Histories from Galaxy Spectra with Score-Based Models

Sacha Perry-Fagant[1,2,3] **Connor Stone**[1,2,3]
Yashar Hezaveh[1,2,3,4,5,6] **Laurence Perreault-Levasseur**[1,2,3,4,5,6]
[1]Université de Montréal [2]Mila [3]Ciela Institute [4]CCA, Flatiron Institute
[5]Perimeter Institute for Theoretical Physics [6]Trottier Space Institute
{sacha.perry-fagant,connor.stone,yashar.hezaveh,
laurence.perreault.levasseur}@umontreal.ca

Abstract

Star formation histories (SFHs) describe when stars in individual galaxies were formed, and are thus a key quantity in understanding galactic evolution. As SFHs themselves are unobservable, information about a galaxy's SFH must be inferred from the galaxy's spectrum, which is an ill-posed inverse problem. In this work, we train a score-based diffusion model to act as a prior over the SFHs. We collect our training data, SFHs and metallicity histories (MHs), from TNG50 simulations, ensuring that the samples include realistic features. We test our model by doing Bayesian posterior inference on mock observations. We apply our model to spectra from the Sloan Digital Sky Survey and find good agreement with previous results. Our work shows it is possible to infer detailed and realistic SFHs and MHs, opening a window to the study of galactic histories at a level of detail never before possible.

1 Introduction

Star formation histories (SFHs) describe when stars in individual galaxies were formed, and are thus a key quantity in understanding galactic evolution. SFHs are necessary for many areas of research in astrophysics, such as studies of dark matter [1, 2], galactic structure and evolution [3, 4], and the expansion of the universe [5]. SFH inference is typically done via spectral energy distribution (SED) fitting of galaxy spectra. A significant challenge in this process is the absence of known, closed-form priors for the SFH derived from first principles. To circumvent this, astrophysical analyses frequently employ simplified parametric models. While simple, these models have proven effective at recovering key galactic properties, such as total stellar mass, with reasonable accuracy [6, 7]. Non-parametric models, in which the SFHs are described by N bins of freely varying star formation, improve over parametric models in terms of flexibility and more accurate inference [7], at the cost of complexity [4]. SFH inference is degenerate as different SFHs can reproduce the same spectra, especially when including metallicity as a result of the age-metallicity degeneracy [7], making it an ill-posed inverse problem. For this reason, it is crucial that only viable SFHs are considered when doing inference. When using non-parametric models for SFH inference, regularization is necessary to ensure physical solutions, but this creates a dilemma. Overly aggressive regularization risks suppressing sharp, bursty features [8], while overly permissive regularization can result in uninformative, broad priors [4].

In this work, we circumvent these issues by using score-based diffusion models trained on the TNG50 cosmological simulations [9, 10] to act as a non-parametric prior over SFHs. In using physically-motivated cosmological simulations, we ensure that our samples include realistic features. To address the age-metallicity degeneracy, we train the model jointly on SFHs as well as metallicity histories (MHs). We demonstrate the capability of our model to do Bayesian posterior inference on both mock spectra and data from the Sloan Digital Sky Survey [11] (SDSS).

Machine Learning and the Physical Sciences Workshop, NeurIPS 2025.

2 Methods

2.1 Data

We generate the SFH and MH data using the TNG50-1 [9, 10] simulations at snapshot 99, where the redshift is effectively zero. Histories are defined using 64 time bins from age 0 to the current age of the universe, calculated using `astropy` [12, 13, 14] assuming a Planck15 [15] cosmology in adherence with TNG50-1, resulting in time bins with a width of approximately 215Myr. Stars in each subhalo are assigned to a bin based on their `GFM_StellarFormationTime`. The sum of all the stars' `GFM_InitialMass` is divided by the width of the bin to store the star formation rate (SFR) in units of $M_\odot$/yr. Metallicities in each bin are calculated by taking the mean of `GFM_Metallicity` weighted by `GFM_InitialMass`. The metallicities Z are converted to $\log_{10} Z/Z_\odot$, assuming solar metallicity $Z_\odot = 0.0142$ in adherence with the stellar population synthesis (SPS) code used in the forward model, described in Section 2.2. Bins with no star formation are assigned a metallicity of 10^{-11} to prevent taking $\log 0$. Because of limitations of the forward model, subhalos with mean metallicity bins having $\log_{10} Z/Z_\odot > 0.5$ are discarded. In addition, only subhalos with `SubhaloFlag=True` and with a total mass $> 10^5 M_\odot$ are used. Before training, the data is z-score normalized and split into a training set (94372 samples) and a test set (23596 samples) to create mock observations with.

2.2 Forward model

A galaxy spectrum can be thought of as the sum of spectra of all objects that inhabit it, here assumed to be dominated by stellar light. The spectra of populations of stars can be approximated using single stellar populations (SSPs), which are theoretical spectra for stars born from the same gas cloud (and thus share a similar metallicity) around the same time. Converting (SFH, MH) pairs into mock spectra then consists of summing up SSPs at each age and metallicity, weighted by the total stellar mass. This process results in a composite stellar population (CSP) which acts as our mock spectrum, see Conroy [16] for a review. To create the SSPs, we use python-FSPS[1] [17], a `Python` wrapper for FSPS [18, 19] (Flexible Stellar Population Synthesis). To ensure the forward model is differentiable, we interpolate between a grid of previously extracted SSPs rather than calling FSPS directly. The details of our FSPS configuration and grid generation can be found in Appendix A.1. The max metallicity value allowed by our chosen libraries is $\log_{10} Z/Z_\odot = 0.5$ but we allow our forward model to interpolate spectra if provided with higher values. We find that the interpolated spectra do not differ greatly from $\log_{10} Z/Z_\odot = 0.5$ spectra, so we cap them at this value in our plots and analysis.

The redshift and the dust attenuation parameter A_V are two configurable parameters that must be provided to the model. The redshift encodes how much the spectrum has been stretched by the expansion of the universe since its emission. It is often used interchangeably with age. The (SFH, MH) pairs are always generated up to the current age of the universe, but we only integrate up to the appropriate redshift to compute the CSP; the rest are considered as a prediction. A_V controls dust emission and extinction. We assume Calzetti extinction [20] and use `BAGPIPES`[2] [21] to generate the extinction curve A_λ/A_V and emission spectra. More details can be found in Appendix A.2.

2.3 Score-based models

Score-based models (SBMs) [22] are a type of generative model that produce new samples from a dataset of interest with underlying distribution $p(\mathbf{x}), \mathbf{x} \in \mathbb{R}^n$, through denoising. The model is a neural network that is trained by adding increasing levels of noise $t \in [0, T]$ to the dataset until it can be expressed as $\mathcal{N}(\mathbf{x}; \mathbf{0}, \sigma_T^2 \mathbb{I})$, where σ_T is the maximum value of noise, and learning the score $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ at each level. The score is then used to generate samples from the noised distribution. SBMs can be used for posterior sampling in the case where we have an observation $\mathbf{y}$ produced by a physical process $S(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}^m$, in our case the forward model defined in Section 2.2, with additive noise $\eta \in \mathbb{R}^m$ assumed to be distributed as $\eta \sim \mathcal{N}(0, \Sigma_y)$. Posterior sampling can be done by using the learnt prior score in Bayes' Theorem

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t \mid \mathbf{y}) = \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} \mid \mathbf{x}_t). \tag{1}$$

[1] github.com/dfm/python-fsps

[2] github.com/ACCarnall/bagpipes

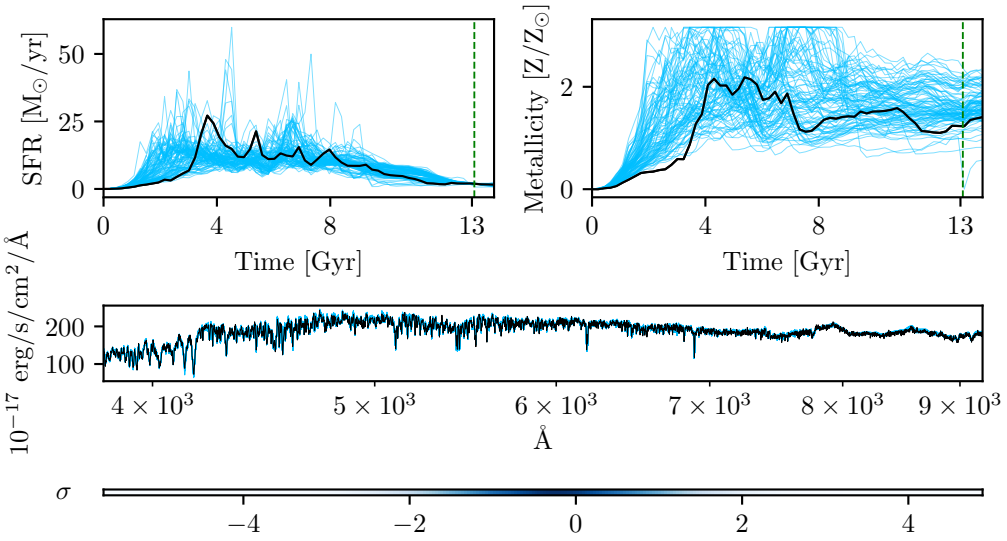

Figure 1: Top: 100 SFH and MH samples (blue) for the mock TNG observation 535410 with $A_V = 0.2$ and redshift$= 0.05$, compared against ground truth (black). The time axis shows age of the universe and the dashed green line shows the time at redshift 0.05. Middle: observed spectrum (black) and posterior spectra (blue). Bottom: density of residuals.

The last term in equation 1 is the likelihood, which we approximate at each noise level t using the convolved likelihood approximation (CLA) [23, 24]

$$p_t(\mathbf{y} \mid \mathbf{x}_t) \approx \mathcal{N}(\mathbf{y} \mid S(\mathbf{x}_t), \Sigma_\mathbf{y} + \sigma_t^2 A A^\top)\,, \tag{2}$$

where A is the Jacobian of the forward model. As our forward model S is non-linear, we approximate $AA^\top$ with a diagonal matrix as described in Appendix B.2, which we have found to work well for our experiments. We use the `Python` package `score_models` [3] [25, 26, 27] for training and sampling. We employ the variance exploding (VE) SDE and use an NCSNpp [22] (Noise Conditional Score Network++). Other model choices and hyperparameters are described in Appendix B.1. Gradient clipping is applied when calculating the score of the likelihood. As getting sufficient samples for one test point currently takes several hours, work is being done to speed up sampling.

3 Experiments

We test our model on mock data and SDSS data. We obtain between 100-120 samples for each experiment and use the top 100 samples based on mean residuals for our plots and analysis.

Mock data Mock observations are created by selecting a random sample from the test set and feeding it through the forward model. The noise is sampled assuming an SNR$= 50$: $\Sigma_\mathbf{y} = \frac{1}{50}\mathbb{E}[\tilde{\mathbf{y}}]\mathbb{I}$ where $\tilde{\mathbf{y}}$ is the observation without noise. The target wavelengths are chosen to match an arbitrary SDSS observation, but can be configured to any range within the wavelengths provided by FSPS: $[90\text{Å}, 10^8\text{Å}]$. We show an example of posterior SFHs and MHs in Figure 1, demonstrating that our posteriors have coverage of the ground truth. We are able to recover the ground truth cumulative mass of $1.056 \times 10^{11} M_\odot$ with our samples $1.067 \pm 0.058 \times 10^{11} M_\odot$. Due to computational limitations, we showcase only one experiment. Future work will test coverage on a number of test samples.

SDSS data We apply our model to SDSS data and compare our results to estimates from the MaNGA FIREFLY Value-Added-Catalog [28], which is an application of the FIREFLY SED fitting code [29] to the SDSS MaNGA IFU survey data [30]. We convert the spatially resolved estimates to a 2D history by summing the values in each bin. This catalogue is also used to find the redshift and

[3] github.com/AlexandreAdam/score_models

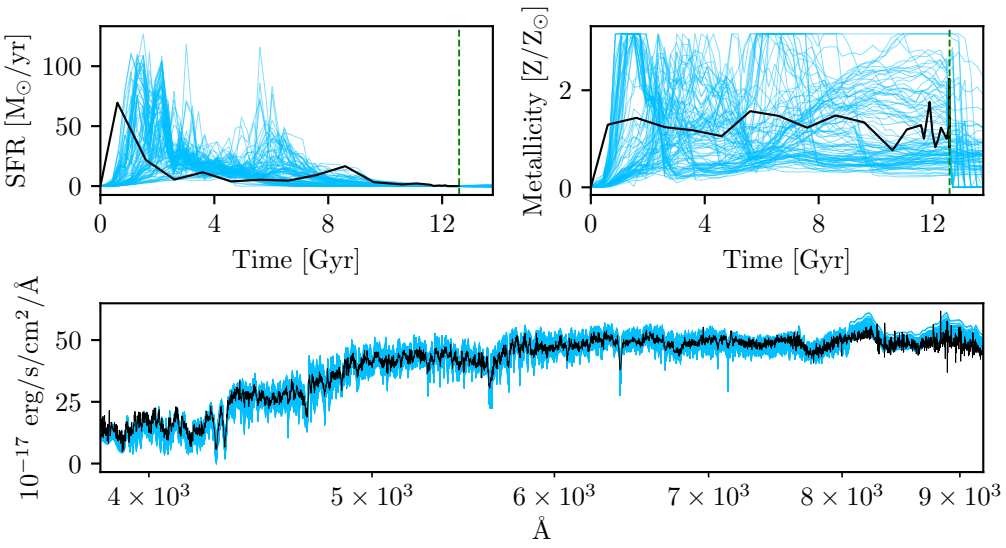

Figure 2: Top: 100 SFHs and MHs samples (blue) for MaNGA index 1-384872, and FIREFLY estimates (black). The time axis shows age of the universe and the dashed green line shows the time at redshift 0.0873. Bottom: observed SDSS spectrum (black) and posterior spectra (blue).

mean colour-excess E(B-V), which is used to calculate $A_V = R_V E(B-V)$ [31], with $R_V = 4.05$ for Calzetti extinction. The SDSS field `ivar` (inverse variance) is used as $1/\sigma_{\mathbf{y}}^2$ for $\Sigma_{\mathbf{y}} = \sigma_{\mathbf{y}}^2$ when sampling.

We apply our model to the spectrum of SDSS J082926.47+224436.2 which was an arbitrarily chosen massive galaxy, with redshift $= 0.0873$ and $A_V = 0.281$. The MaNGA index for this galaxy is 1-384872, and we refer to it as such from here on. We show the posterior samples and spectra in Figure 2 and residuals in Figure 3. A more detailed view of the spectra can be found in Appendix C. We find a cumulative stellar mass of $1.371 \pm 0.174 \times 10^{11} M_\odot$, comparable to the remnant stellar mass reported by FIREFLY: $1.571 \times 10^{11} M_\odot$. Both models find a large burst of early star formation with another burst at later times, though these events differ slightly in their timing. Because of the age-dust-metallicity degeneracy [16], both dust and high metallicities can result in dimmer spectra. As our model tends to push for higher metallicities, we investigate the effect of sampling with higher values of A_V. With $A_V = 0.339$, we find that the metallicities are more well behaved, but the mass is lowered to $1.195 \pm 0.053 \times 10^{11} M_\odot$. Increasing further to $A_V = 0.506$ further reduces the metallicity but yields a similar mass of $1.378 \pm 0.087 \times 10^{11} M_\odot$. Details can be found in Appendix D. The lopsided residuals in Figure 3 are primarily caused by exaggerated absorption features in our posterior samples, which may be a result of the high metallicity values inferred by our model, or from model assumptions. We discuss this further in Appendix E. Future work will test the agreement of our model with a larger set of galaxies from the catalogue.

4 Conclusion

We have demonstrated the potential of SBMs to infer (SFH, MH) pairs from real data. Our model improves upon other models by using simulations to train a broad and detailed prior rather than models with overly strict constraints about the form of the SFHs and MHs. By training our prior model on physics-based simulations, we generate highly realistic SFHs and MHs.

Galaxy formation is a complex process of many interacting mechanisms that make inference challenging. SED fitting codes can be used to infer properties from observed galaxy spectra, but are sensitive to the approximations made by the model. We have made advances in this regard by using a flexible and informed prior over the SFH and MH.

Future work includes generating coverage tests with mock posterior samples and exploring alternative methods for approximating the likelihood, as CLA performs best with linear models. We also plan to

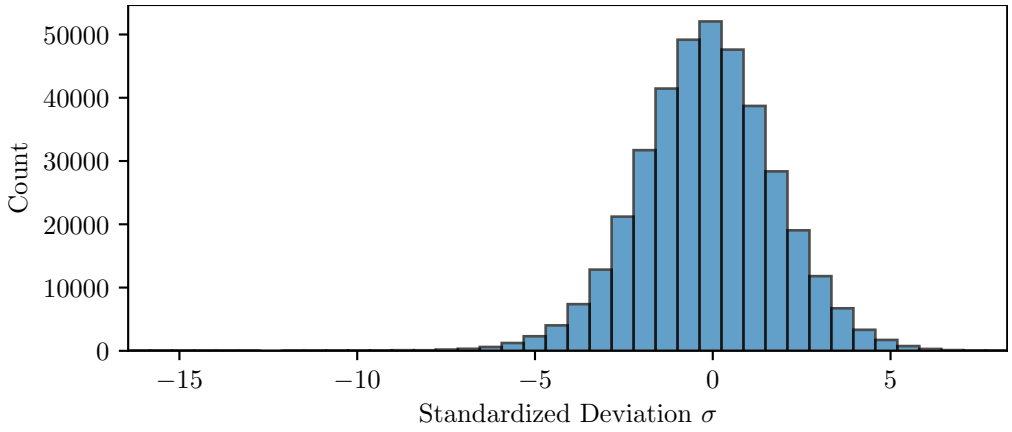

Figure 3: Residuals for the posterior samples in Figure 2.

test replacements for gradient clipping, since the likelihood is not sufficiently constraining for lower mass samples. Additionally, we will examine the effects of varying the initial mass function, dust models, and isochrone models, as well as priors informed by different hydrodynamical simulations.

Acknowledgments and Disclosure of Funding

The work is in part supported by computational resources provided by Calcul Québec and the Digital Research Alliance of Canada. C.S. acknowledges the support of a NSERC Postdoctoral Fellowship and a CITA National Fellowship. This work is partially supported by Schmidt Futures, a philanthropic initiative founded by Eric and Wendy Schmidt as part of the Virtual Institute for Astrophysics (VIA). Y.H. and L.P.-L. acknowledge support from the Canada Research Chairs Program, the Natural Sciences and Engineering Research Council of Canada (NSERC) through grants RGPIN-2020-05073 and 05102, and the Fonds de recherche du Québec through grant CFQCU-2024-348060. R.L. acknowledges support from the NSERC CGSD scholarship.

Funding for the Sloan Digital Sky Survey V has been provided by the Alfred P. Sloan Foundation, the Heising-Simons Foundation, the National Science Foundation, and the Participating Institutions. SDSS acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. SDSS telescopes are located at Apache Point Observatory, funded by the Astrophysical Research Consortium and operated by New Mexico State University, and at Las Campanas Observatory, operated by the Carnegie Institution for Science. The SDSS web site is `www.sdss.org`.

SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration, including the Carnegie Institution for Science, Chilean National Time Allocation Committee (CNTAC) ratified researchers, Caltech, the Gotham Participation Group, Harvard University, Heidelberg University, The Flatiron Institute, The Johns Hopkins University, L'Ecole polytechnique fédérale de Lausanne (EPFL), Leibniz-Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Extraterrestrische Physik (MPE), Nanjing University, National Astronomical Observatories of China (NAOC), New Mexico State University, The Ohio State University, Pennsylvania State University, Smithsonian Astrophysical Observatory, Space Telescope Science Institute (STScI), the Stellar Astrophysics Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Illinois at Urbana-Champaign, University of Toronto, University of Utah, University of Virginia, Yale University, and Yunnan University.

References

[1] Arianna Di Cintio, Chris B. Brook, Andrea V. Macciò, Greg S. Stinson, Alexander Knebe, Aaron A. Dutton, and James Wadsley. The dependence of dark matter profiles on the stellar-to-halo mass ratio: a prediction for cusps versus cores. *Monthly Notices of the Royal Astronomical*

Society, 437(1):415–423, 11 2013. ISSN 0035-8711. doi: 10.1093/mnras/stt1891. URL https://doi.org/10.1093/mnras/stt1891.

[2] Peter Behroozi, Risa H Wechsler, Andrew P Hearin, and Charlie Conroy. Universemachine: The correlation between galaxy growth and dark matter halo assembly from z=0-10. *Monthly Notices of the Royal Astronomical Society*, 488(3):3143–3194, 05 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz1182. URL https://doi.org/10.1093/mnras/stz1182.

[3] Edward N. Taylor, Andrew M. Hopkins, Ivan K. Baldry, Michael J. I. Brown, Simon P. Driver, Lee S. Kelvin, David T. Hill, Aaron S. G. Robotham, Joss Bland-Hawthorn, D. H. Jones, R. G. Sharp, Daniel Thomas, Jochen Liske, Jon Loveday, Peder Norberg, J. A. Peacock, Steven P. Bamford, Sarah Brough, Matthew Colless, Ewan Cameron, Christopher J. Conselice, Scott M. Croom, C. S. Frenk, Madusha Gunawardhana, Konrad Kuijken, R. C. Nichol, H. R. Parkinson, S. Phillipps, K. A. Pimbblet, C. C. Popescu, Matthew Prescott, W. J. Sutherland, R. J. Tuffs, Eelco van Kampen, and D. Wijesinghe. Galaxy and mass assembly (gama): stellar mass estimates. *Monthly Notices of the Royal Astronomical Society*, 418(3):1587–1620, 12 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.19536.x. URL https://doi.org/10.1111/j.1365-2966.2011.19536.x.

[4] Joel Leja, Adam C. Carnall, Benjamin D. Johnson, Charlie Conroy, and Joshua S. Speagle. How to measure galaxy star formation histories. ii. nonparametric models. *The Astrophysical Journal*, 876(1):3, April 2019. ISSN 1538-4357. doi: 10.3847/1538-4357/ab133c. URL http://dx.doi.org/10.3847/1538-4357/ab133c.

[5] Nicola Borghi, Michele Moresco, and Andrea Cimatti. Toward a better understanding of cosmic chronometers: A new measurement of h(z) at z $\sim$ 0.7. *The Astrophysical Journal Letters*, 928(1):L4, March 2022. ISSN 2041-8213. doi: 10.3847/2041-8213/ac3fb2. URL http://dx.doi.org/10.3847/2041-8213/ac3fb2.

[6] Adam C. Carnall, Joel Leja, Benjamin D. Johnson, Ross J. McLure, James S. Dunlop, and Charlie Conroy. How to measure galaxy star formation histories. i. parametric models. *The Astrophysical Journal*, 873(1):44, March 2019. ISSN 1538-4357. doi: 10.3847/1538-4357/ab04a2. URL http://dx.doi.org/10.3847/1538-4357/ab04a2.

[7] Sidney Lower, Desika Narayanan, Joel Leja, Benjamin D. Johnson, Charlie Conroy, and Romeel Davé. How well can we measure the stellar mass of a galaxy: The impact of the assumed star formation history model in sed fitting. *The Astrophysical Journal*, 904(1):33, November 2020. ISSN 1538-4357. doi: 10.3847/1538-4357/abbfa7. URL http://dx.doi.org/10.3847/1538-4357/abbfa7.

[8] P. Ocvirk, C. Pichon, A. Lançon, and E. Thiébaut. Steckmap: Stellar content and kinematics from high resolution galactic spectra via maximum a posteriori. *Monthly Notices of the Royal Astronomical Society*, 365(1):74–84, January 2006. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2005.09323.x. URL http://dx.doi.org/10.1111/j.1365-2966.2005.09323.x.

[9] Dylan Nelson, Annalisa Pillepich, Volker Springel, Rüdiger Pakmor, Rainer Weinberger, Shy Genel, Paul Torrey, Mark Vogelsberger, Federico Marinacci, and Lars Hernquist. First results from the tng50 simulation: galactic outflows driven by supernovae and black hole feedback. *Monthly Notices of the Royal Astronomical Society*, 490(3):3234–3261, August 2019. ISSN 1365-2966. doi: 10.1093/mnras/stz2306. URL http://dx.doi.org/10.1093/mnras/stz2306.

[10] Annalisa Pillepich, Dylan Nelson, Volker Springel, Rüdiger Pakmor, Paul Torrey, Rainer Weinberger, Mark Vogelsberger, Federico Marinacci, Shy Genel, Arjen van der Wel, and Lars Hernquist. First results from the tng50 simulation: the evolution of stellar and gaseous discs across cosmic time. *Monthly Notices of the Royal Astronomical Society*, 490(3):3196–3233, September 2019. ISSN 1365-2966. doi: 10.1093/mnras/stz2338. URL http://dx.doi.org/10.1093/mnras/stz2338.

[11] Juna A. Kollmeier, Hans-Walter Rix, Conny Aerts, James Aird, Pablo Vera Alfaro, Andrés Almeida, Scott F. Anderson, Óscar Jiménez Arranz, Stefan M. Arseneau, Roberto Assef, Shir

Aviram, Catarina Aydar, Carles Badenes, Avrajit Bandyopadhyay, Kat Barger, Robert H. Barkhouser, Franz E. Bauer, Chad Bender, Felipe Besser, Binod Bhattarai, Pavaman Bilgi, Jonathan Bird, Dmitry Bizyaev, Guillermo A. Blanc, Michael R. Blanton, John Bochanski, Jo Bovy, Christopher Brandon, William Nielsen Brandt, Joel R. Brownstein, Johannes Buchner, Joseph N. Burchett, Joleen Carlberg, Andrew R. Casey, Lesly Castaneda-Carlos, Priyanka Chakraborty, Julio Chanamé, Vedant Chandra, Brian Cherinka, Igor Chilingarian, Johan Comparat, Maren Cosens, Kevin Covey, Jeffrey D. Crane, Nicole R. Crumpler, Katia Cunha, Tim Cunningham, Xinyu Dai, Jeremy Darling, James W. Davidson Jr., Megan C. Davis, Nathan De Lee, Niall Deacon, José Eduardo Méndez Delgado, Sebastian Demasi, Mariia Demianenko, Mark Derwent, Elena D'Onghia, Francesco Di Mille, Bruno Dias, John Donor, Niv Drory, Tom Dwelly, Oleg Egorov, Evgeniya Egorova, Kareem El-Badry, Mike Engelman, Mike Eracleous, Xiaohui Fan, Emily Farr, Logan Fries, Peter Frinchaboy, Cynthia S. Froning, Boris T. Gänsicke, Pablo García, Joseph Gelfand, Nicola Pietro Gentile Fusillo, Simon Glover, Katie Grabowski, Eva K. Grebel, Paul J Green, Catherine Grier, Pramod Gupta, Aidan C. Gray, Maximilian Häberle, Patrick B. Hall, Randolph P. Hammond, Keith Hawkins, Albert C. Harding, Viola Hegedűs, Tom Herbst, J. J. Hermes, Paola Rodríguez Hidalgo, Thomas Hilder, David W Hogg, Jon A. Holtzman, Danny Horta, Yang Huang, Hsiang-Chih Hwang, Hector Javier Ibarra-Medel, Julie Imig, Keith Inight, Arghajit Jana, Alexander P. Ji, Paula Jofre, Matt Johns, Jennifer Johnson, James W. Johnson, Evelyn J. Johnston, Amy M Jones, Ivan Katkov, Anton M. Koekemoer, Marina Kounkel, Kathryn Kreckel, Dhanesh Krishnarao, Mirko Krumpe, Nimisha Kumari, Thomas Kupfer, Ivan Lacerna, Chervin Laporte, Sebastien Lepine, Jing Li, Xin Liu, Sarah Loebman, Knox Long, Alexandre Roman-Lopes, Yuxi Lu, Steven Raymond Majewski, Dan Maoz, Kevin A. McKinnon, Ilija Medan, Andrea Merloni, Dante Minniti, Sean Morrison, Natalie Myers, Szabolcs Mészáros, Kirpal Nandra, Prasanta K. Nayak, Melissa K Ness, David L. Nidever, Thomas O'Brien, Micah Oeur, Audrey Oravetz, Daniel Oravetz, Jonah Otto, Gautham Adamane Pallathadka, Povilas Palunas, Kaike Pan, Daniel Pappalardo, Rakesh Pandey, Castalia Alenka Negrete Peñaloza, Marc H. Pinsonneault, Richard W. Pogge, Manuchehr Taghizadeh Popp, Adrian M. Price-Whelan, Nadiia Pulatova, Dan Qiu, Solange Ramirez, Amy Rankine, Claudio Ricci, Jessie C. Runnoe, Sebastian Sanchez, Mara Salvato, Natascha Sattler, Andrew K. Saydjari, Conor Sayres, Kevin C. Schlaufman, Donald P. Schneider, Matthias R. Schreiber, Axel Schwope, Javier Serna, Yue Shen, Cristóbal Sifón, Amrita Singh, Amaya Sinha, Stephen Smee, Ying-Yi Song, Diogo Souto, Keivan G. Stassun, Matthias Steinmetz, Alexander Stone-Martinez, Guy Stringfellow, Amelia Stutz, José, Sá, nchez Gallego, Jonathan C. Tan, Jamie Tayar, Riley Thai, Ani Thakar, Yuan-Sen Ting, Andrew Tkachenko, Gagik Tovmasian, Benny Trakhtenbrot, José G. Fernández-Trincado, Nicholas Troup, Jonathan Trump, Sarah Tuttle, Roeland P. van der Marel, Sandro Villanova, Jaime Villaseñor, Stefanie Wachter, Zachary Way, Anne-Marie Weijmans, David Weinberg, Adam Wheeler, John Wilson, Alessa I. Wiggins, Tony Wong, Qiaoya Wu, Dominika Wylezalek, Xiang-Xiang Xue, Qian Yang, Nadia Zakamska, Eleonora Zari, Gail Zasowski, Grisha Zeltyn, Catherine Zucker, Carlos G. Román Zúñiga, and Rodolfo de J. Zermeño. Sloan digital sky survey-v: Pioneering panoptic spectroscopy, 2025. URL `https://arxiv.org/abs/2507.06989`.

[12] Astropy Collaboration, T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan, W. E. Kerzendorf, A. Conley, N. Crighton, K. Barbary, D. Muna, H. Ferguson, F. Grollier, M. M. Parikh, P. H. Nair, H. M. Unther, C. Deil, J. Woillez, S. Conseil, R. Kramer, J. E. H. Turner, L. Singer, R. Fox, B. A. Weaver, V. Zabalza, Z. I. Edwards, K. Azalee Bostroem, D. J. Burke, A. R. Casey, S. M. Crawford, N. Dencheva, J. Ely, T. Jenness, K. Labrie, P. L. Lim, F. Pierfederici, A. Pontzen, A. Ptak, B. Refsdal, M. Servillat, and O. Streicher. Astropy: A community Python package for astronomy. *A&A*, 558:A33, October 2013. doi: 10.1051/0004-6361/201322068.

[13] Astropy Collaboration, A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva, A. Ginsburg, J. T. Vand erPlas, L. D. Bradley, D. Pérez-Suárez, M. de Val-Borro, T. L. Aldcroft, K. L. Cruz, T. P. Robitaille, E. J. Tollerud, C. Ardelean, T. Babej, Y. P. Bach, M. Bachetti, A. V. Bakanov, S. P. Bamford, G. Barentsen, P. Barmby, A. Baumbach, K. L. Berry, F. Biscani, M. Boquien, K. A. Bostroem, L. G. Bouma, G. B. Brammer, E. M. Bray, H. Breytenbach, H. Buddelmeijer, D. J. Burke, G. Calderone, J. L. Cano Rodríguez, M. Cara, J. V. M. Cardoso, S. Cheedella, Y. Copin, L. Corrales, D. Crichton, D. D'Avella, C. Deil, É. Depagne, J. P. Dietrich, A. Donath, M. Droettboom, N. Earl, T. Erben, S. Fabbro, L. A. Ferreira, T. Finethy, R. T. Fox, L. H.

Garrison, S. L. J. Gibbons, D. A. Goldstein, R. Gommers, J. P. Greco, P. Greenfield, A. M. Groener, F. Grollier, A. Hagen, P. Hirst, D. Homeier, A. J. Horton, G. Hosseinzadeh, L. Hu, J. S. Hunkeler, Ž. Ivezić, A. Jain, T. Jenness, G. Kanarek, S. Kendrew, N. S. Kern, W. E. Kerzendorf, A. Khvalko, J. King, D. Kirkby, A. M. Kulkarni, A. Kumar, A. Lee, D. Lenz, S. P. Littlefair, Z. Ma, D. M. Macleod, M. Mastropietro, C. McCully, S. Montagnac, B. M. Morris, M. Mueller, S. J. Mumford, D. Muna, N. A. Murphy, S. Nelson, G. H. Nguyen, J. P. Ninan, M. Nöthe, S. Ogaz, S. Oh, J. K. Parejko, N. Parley, S. Pascual, R. Patil, A. A. Patil, A. L. Plunkett, J. X. Prochaska, T. Rastogi, V. Reddy Janga, J. Sabater, P. Sakurikar, M. Seifert, L. E. Sherbert, H. Sherwood-Taylor, A. Y. Shih, J. Sick, M. T. Silbiger, S. Singanamalla, L. P. Singer, P. H. Sladen, K. A. Sooley, S. Sornarajah, O. Streicher, P. Teuben, S. W. Thomas, G. R. Tremblay, J. E. H. Turner, V. Terrón, M. H. van Kerkwijk, A. de la Vega, L. L. Watkins, B. A. Weaver, J. B. Whitmore, J. Woillez, V. Zabalza, and Astropy Contributors. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. *The Astronomical Journal*, 156(3): 123, September 2018. doi: 10.3847/1538-3881/aabc4f.

[14] Astropy Collaboration, Adrian M. Price-Whelan, Pey Lian Lim, Nicholas Earl, Nathaniel Starkman, Larry Bradley, David L. Shupe, Aarya A. Patil, Lia Corrales, C. E. Brasseur, Maximilian N"othe, Axel Donath, Erik Tollerud, Brett M. Morris, Adam Ginsburg, Eero Vaher, Benjamin A. Weaver, James Tocknell, William Jamieson, Marten H. van Kerkwijk, Thomas P. Robitaille, Bruce Merry, Matteo Bachetti, H. Moritz G"unther, Thomas L. Aldcroft, Jaime A. Alvarado-Montes, Anne M. Archibald, Attila B'odi, Shreyas Bapat, Geert Barentsen, Juanjo Baz'an, Manish Biswas, M'ed'eric Boquien, D. J. Burke, Daria Cara, Mihai Cara, Kyle E. Conroy, Simon Conseil, Matthew W. Craig, Robert M. Cross, Kelle L. Cruz, Francesco D'Eugenio, Nadia Dencheva, Hadrien A. R. Devillepoix, J"org P. Dietrich, Arthur Davis Eigenbrot, Thomas Erben, Leonardo Ferreira, Daniel Foreman-Mackey, Ryan Fox, Nabil Freij, Suyog Garg, Robel Geda, Lauren Glattly, Yash Gondhalekar, Karl D. Gordon, David Grant, Perry Greenfield, Austen M. Groener, Steve Guest, Sebastian Gurovich, Rasmus Handberg, Akeem Hart, Zac Hatfield-Dodds, Derek Homeier, Griffin Hosseinzadeh, Tim Jenness, Craig K. Jones, Prajwel Joseph, J. Bryce Kalmbach, Emir Karamehmetoglu, Mikolaj Kaluszy'nski, Michael S. P. Kelley, Nicholas Kern, Wolfgang E. Kerzendorf, Eric W. Koch, Shankar Kulumani, Antony Lee, Chun Ly, Zhiyuan Ma, Conor MacBride, Jakob M. Maljaars, Demitri Muna, N. A. Murphy, Henrik Norman, Richard O'Steen, Kyle A. Oman, Camilla Pacifici, Sergio Pascual, J. Pascual-Granado, Rohit R. Patil, Gabriel I. Perren, Timothy E. Pickering, Tanuj Rastogi, Benjamin R. Roulston, Daniel F. Ryan, Eli S. Rykoff, Jose Sabater, Parikshit Sakurikar, Jes'us Salgado, Aniket Sanghi, Nicholas Saunders, Volodymyr Savchenko, Ludwig Schwardt, Michael Seifert-Eckert, Albert Y. Shih, Anany Shrey Jain, Gyanendra Shukla, Jonathan Sick, Chris Simpson, Sudheesh Singanamalla, Leo P. Singer, Jaladh Singhal, Manodeep Sinha, Brigitta M. SipHocz, Lee R. Spitler, David Stansby, Ole Streicher, Jani Sumak, John D. Swinbank, Dan S. Taranu, Nikita Tewary, Grant R. Tremblay, Miguel de Val-Borro, Samuel J. Van Kooten, Zlatan Vasovi'c, Shresth Verma, Jos'e Vin'icius de Miranda Cardoso, Peter K. G. Williams, Tom J. Wilson, Benjamin Winkel, W. M. Wood-Vasey, Rui Xue, Peter Yoachim, Chen Zhang, Andrea Zonca, and Astropy Project Contributors. The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package. *The Astrophysical Journal*, 935(2):167, August 2022. doi: 10.3847/1538-4357/ac7c74.

[15] P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, E. Battaner, R. Battye, K. Benabed, A. Benoît, A. Benoit-Lévy, J.-P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J.-F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, R.-R. Chary, H. C. Chiang, J. Chluba, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, C. Combet, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, F.-X. Désert, E. Di Valentino, C. Dickinson, J. M. Diego, K. Dolag, H. Dole, S. Donzelli, O. Doré, M. Douspis, A. Ducout, J. Dunkley, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, M. Farhang, J. Fergusson, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frejsel, S. Galeotta, S. Galli, K. Ganga, C. Gauthier, M. Gerbino, T. Ghosh, M. Giard, Y. Giraud-Héraud, E. Giusarma, E. Gjerløw, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, J. Hamann, F. K. Hansen, D. Hanson, D. L. Harrison, G. Helou, S. Henrot-Versillé, C. Hernández-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson,

W. A. Holmes, A. Hornstrup, W. Hovest, Z. Huang, K. M. Huffenberger, G. Hurier, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, E. Keihänen, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J.-M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, J. P. Leahy, R. Leonardi, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Linden-Vørnle, M. López-Caniego, P. M. Lubin, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marchini, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Masi, S. Matarrese, P. McGehee, P. R. Meinhold, A. Melchiorri, J.-B. Melin, L. Mendes, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M.-A. Miville-Deschênes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Nørgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, T. J. Pearson, O. Perdereau, L. Perotto, F. Perrotta, V. Pettorino, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, L. Popa, G. W. Pratt, G. Prézeau, S. Prunet, J.-L. Puget, J. P. Rachen, W. T. Reach, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, I. Ristorcelli, G. Rocha, C. Rosset, M. Rossetti, G. Roudier, B. Rouillé d'Orfeuil, M. Rowan-Robinson, J. A. Rubiño-Martín, B. Rusholme, N. Said, V. Salvatelli, L. Salvati, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, M. D. Seiffert, P. Serra, E. P. S. Shellard, L. D. Spencer, M. Spinelli, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, D. Sutton, A.-S. Suur-Uski, J.-F. Sygnet, J. A. Tauber, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, T. Trombetti, M. Tucci, J. Tuovinen, M. Türler, G. Umana, L. Valenziano, J. Valiviita, F. Van Tent, P. Vielva, F. Villa, L. A. Wade, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Wilkinson, D. Yvon, A. Zacchei, and A. Zonca. Planck2015 results: Xiii. cosmological parameters. *Astronomy & Astrophysics*, 594:A13, September 2016. ISSN 1432-0746. doi: 10.1051/0004-6361/201525830. URL http://dx.doi.org/10.1051/0004-6361/201525830.

[16] Charlie Conroy. Modeling the panchromatic spectral energy distributions of galaxies. *Annual Review of Astronomy and Astrophysics*, 51(1):393–455, August 2013. ISSN 1545-4282. doi: 10.1146/annurev-astro-082812-141017. URL http://dx.doi.org/10.1146/annurev-astro-082812-141017.

[17] Ben Johnson, Dan Foreman-Mackey, Jonathan Sick, Joel Leja, Mike Walmsley, Erik Tollerud, Henry Leung, Spencer Scott, and Minjung Park. dfm/python-fsps: v0.4.7, 2024.

[18] Charlie Conroy, James E Gunn, and Martin White. The propagation of uncertainties in stellar population synthesis modeling. i. the relevance of uncertain aspects of stellar evolution and the initial mass function to the derived physical properties of galaxies. *Astrophys. J.*, 699(1): 486–506, July 2009.

[19] Charlie Conroy and James E Gunn. The propagation of uncertainties in stellar population synthesis modeling. iii. model calibration, comparison, and evaluation. *Astrophys. J.*, 712(2): 833–857, April 2010.

[20] Daniela Calzetti, Lee Armus, Ralph C. Bohlin, Anne L. Kinney, Jan Koornneef, and Thaisa Storchi-Bergmann. The dust content and opacity of actively star-forming galaxies. *The Astrophysical Journal*, 533(2):682–695, April 2000. ISSN 1538-4357. doi: 10.1086/308692. URL http://dx.doi.org/10.1086/308692.

[21] A C Carnall, R J McLure, J S Dunlop, and R Davé. Inferring the star formation histories of massive quiescent galaxies with bagpipes: evidence for multiple quenching mechanisms. *Monthly Notices of the Royal Astronomical Society*, 480(4):4379–4401, August 2018. ISSN 1365-2966. doi: 10.1093/mnras/sty2169. URL http://dx.doi.org/10.1093/mnras/sty2169.

[22] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

[23] B. Remy, F. Lanusse, N. Jeffrey, J. Liu, J.-L. Starck, K. Osato, and T. Schrabback. Probabilistic mass-mapping with neural score estimation. *Astronomy & Astrophysics*, 672:A51, March 2023. ISSN 1432-0746. doi: 10.1051/0004-6361/202243054. URL http://dx.doi.org/10.1051/0004-6361/202243054.

[24] Alexandre Adam, Adam Coogan, Nikolay Malkin, Ronan Legin, Laurence Perreault-Levasseur, Yashar Hezaveh, and Yoshua Bengio. Posterior samples of source galaxies in strong gravitational lenses with score-based priors, 2022. URL https://arxiv.org/abs/2211.03812.

[25] Alexandre Adam, Adam Coogan, Nikolay Malkin, Ronan Legin, Laurence Perreault-Levasseur, Yashar Hezaveh, and Yoshua Bengio. Posterior samples of source galaxies in strong gravitational lenses with score-based priors. *arXiv e-prints*, art. arXiv:2211.03812, November 2022. doi: 10.48550/arXiv.2211.03812.

[26] Ronan Legin, Alexandre Adam, Yashar Hezaveh, and Laurence Perreault-Levasseur. Beyond Gaussian Noise: A Generalized Approach to Likelihood Analysis with Non-Gaussian Noise. *The Astrophysical Journal Letters*, 949(2):L41, June 2023. doi: 10.3847/2041-8213/acd645.

[27] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.

[28] Justus Neumann, Daniel Thomas, Claudia Maraston, Lewis Hill, Lorenza Nanni, Oliver Wenman, Jianhui Lian, Johan Comparat, Violeta Gonzalez-Perez, Kyle B Westfall, Renbin Yan, Yanping Chen, Guy S Stringfellow, Matthew A Bershady, Joel R Brownstein, Niv Drory, and Donald P Schneider. The manga <scp>firefly</scp> value-added-catalogue: resolved stellar populations of 10,010 nearby galaxies. *Monthly Notices of the Royal Astronomical Society*, May 2022. ISSN 1365-2966. doi: 10.1093/mnras/stac1260. URL http://dx.doi.org/10.1093/mnras/stac1260.

[29] David M. Wilkinson, Claudia Maraston, Daniel Goddard, Daniel Thomas, and Taniya Parikh. firefly (fitting iteratively for likelihood analysis): a full spectral fitting code. *Monthly Notices of the Royal Astronomical Society*, 472(4):4297–4326, September 2017. ISSN 1365-2966. doi: 10.1093/mnras/stx2215. URL http://dx.doi.org/10.1093/mnras/stx2215.

[30] David A. Wake, Kevin Bundy, Aleksandar M. Diamond-Stanic, Renbin Yan, Michael R. Blanton, Matthew A. Bershady, José R. Sánchez-Gallego, Niv Drory, Amy Jones, Guinevere Kauffmann, David R. Law, Cheng Li, Nicholas MacDonald, Karen Masters, Daniel Thomas, Jeremy Tinker, Anne-Marie Weijmans, and Joel R. Brownstein. The sdss-iv manga sample: Design, optimization, and usage considerations. *The Astronomical Journal*, 154(3):86, aug 2017. doi: 10.3847/1538-3881/aa7ecc. URL https://doi.org/10.3847/1538-3881/aa7ecc.

[31] Samir Salim and Desika Narayanan. The dust attenuation law in galaxies. *Annual Review of Astronomy and Astrophysics*, 58(1):529–575, August 2020. ISSN 1545-4282. doi: 10.1146/annurev-astro-032620-021933. URL http://dx.doi.org/10.1146/annurev-astro-032620-021933.

[32] Aaron Dotter. Mesa isochrones and stellar tracks (mist) 0: Methods for the construction of stellar isochrones. *The Astrophysical Journal Supplement Series*, 222(1):8, January 2016. ISSN 1538-4365. doi: 10.3847/0067-0049/222/1/8. URL http://dx.doi.org/10.3847/0067-0049/222/1/8.

[33] Jieun Choi, Aaron Dotter, Charlie Conroy, Matteo Cantiello, Bill Paxton, and Benjamin D. Johnson. Mesa isochrones and stellar tracks (mist). i. solar-scaled models. *The Astrophysical Journal*, 823(2):102, May 2016. ISSN 1538-4357. doi: 10.3847/0004-637x/823/2/102. URL http://dx.doi.org/10.3847/0004-637X/823/2/102.

[34] P. Sanchez-Blazquez, R. F. Peletier, J. Jimenez-Vicente, N. Cardiel, A. J. Cenarro, J. Falcon-Barroso, J. Gorgas, S. Selam, and A. Vazdekis. Medium-resolution isaac newton telescope library of empirical spectra. *Monthly Notices of the Royal Astronomical Society*, 371(2):703–718, September 2006. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2006.10699.x. URL http://dx.doi.org/10.1111/j.1365-2966.2006.10699.x.

[35] David W. Hogg. Distance measures in cosmology, 2000. URL https://arxiv.org/abs/astro-ph/9905116.

Table 1: FSPS SSP parameters

Parameter	Value
z_continuous	1
sfh	0
dust_type	0
dust1_index	0
dust1	0
dust2	0
frac_nodust	1
frac_obrun	1
add_neb_emission	False
add_dust_emission	False
add_agb_dust_model	False
add_neb_continuum	False

[36] Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=OnD9zGAGT0k.

A Forward model details

A.1 SSP generation

Our FSPS installation uses the MIST [32, 33] isochrone library, and the MILES [34] stellar library. The parameters used to create the SSPs are listed in Table 1, with all other parameters set to the defaults. All dust fields are set to 0 to allow handling them through A_V instead.

We generate the grid of SSPs for 512 stellar ages and 64 metallicities. The ages are evenly spaced from 0 to the age of the universe, with the 0-age stellar population being replaced with the minimum stellar age of FSPS. The metallicities are evenly binned between the min and max values specified by MIST: $\log_{10} Z/Z_\odot \in [-2.5, 0.5]$.

A.2 Redshift and dust attenuation

We use `astropy` to convert the CSP from solar luminosities $L_\odot$ to $erg/s/cm^2/\mathring{A}$ [35] based on the value of redshift. The wavelengths are then shifted by a factor of (1 + redshift) and interpolated back to the target range.

Emission spectra are generated using `BAGPIPES` [21] assuming (qpah=3.5, umin=1, gamma=0.01), as per FSPS defaults. Dust extinction and emission are applied to the CSP using a `PyTorch` re-implementation of `BAGPIPES` codes.

B SBM details

B.1 Model training

The model was trained for ~ 10000 epochs (around 7 days). We show the parameters used in the `score_models` package for training in Table 2. All other parameters were set to their default values.

B.2 CLA approximation

We approximate $AA^\top$ by defining a $1000 \times m$ matrix B and taking

$$\begin{aligned} B_i &= S(\alpha), \alpha \sim \mathcal{N}(0, \mathbb{I}_n) \\ \text{Diag}(AA^\top) &\approx \text{Var}_\text{i}[B_{i,:}]. \end{aligned} \tag{3}$$

Table 2: SBM parameters

Parameter	Value
`sigma_min`	0.001
`sigma_max`	16
`nf`	64
`ch_mult`	$[1, 2, 2, 4]$

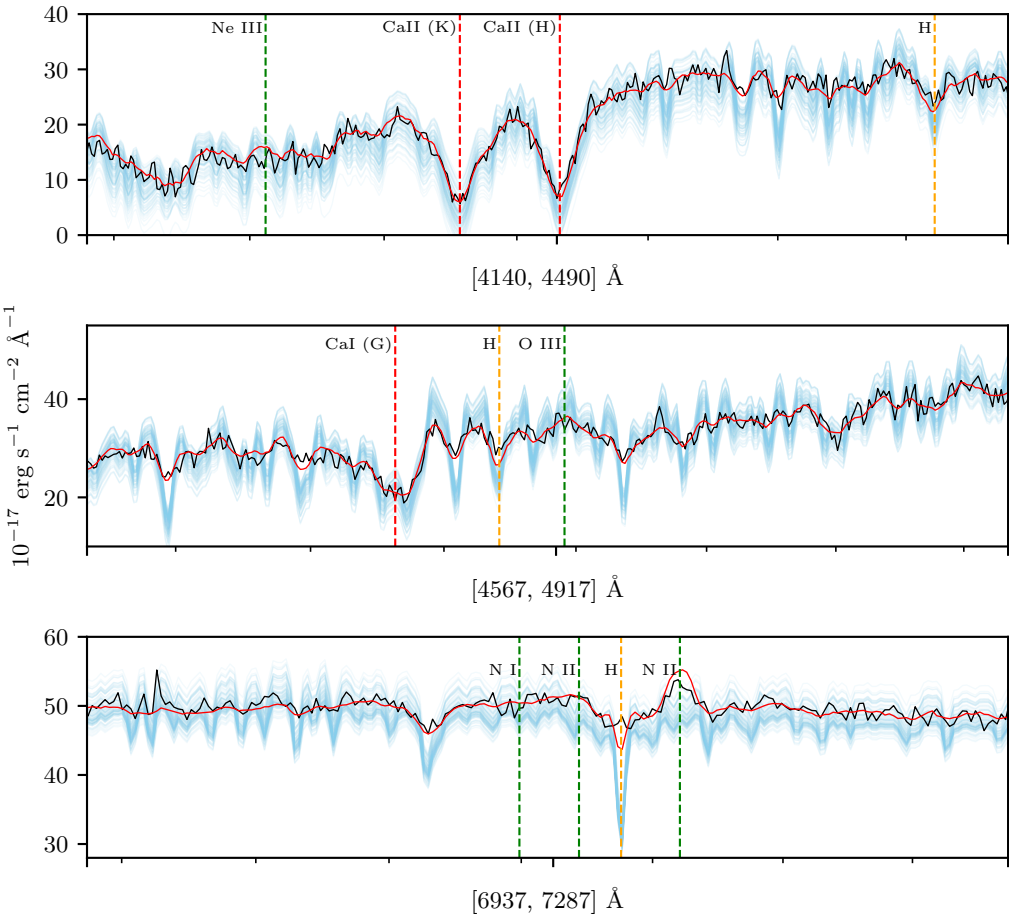

Figure 4: Close-up views of posterior spectra (blue) and true observation (black) for MaNGA index 1-384872. The full spectra are shown in Figure 2. We plot a subset of emission (green), absorption (red) and Balmer (orange) lines identified by SDSS. The best fit model from SDSS is shown in red.

This approximates $\text{Diag}(AA^\top)$ which we use in place of the full $AA^\top$ matrix.

Another approximation of the likelihood that has proven to be applicable to non-linear models is Tweedie's approximation [36]. When comparing the two approximations in our experiments, we found that CLA worked best.

C Detailed posterior spectra

We show close-up regions of the inferred posterior spectra in Figure 4, with best fit model and spectral lines from SDSS.

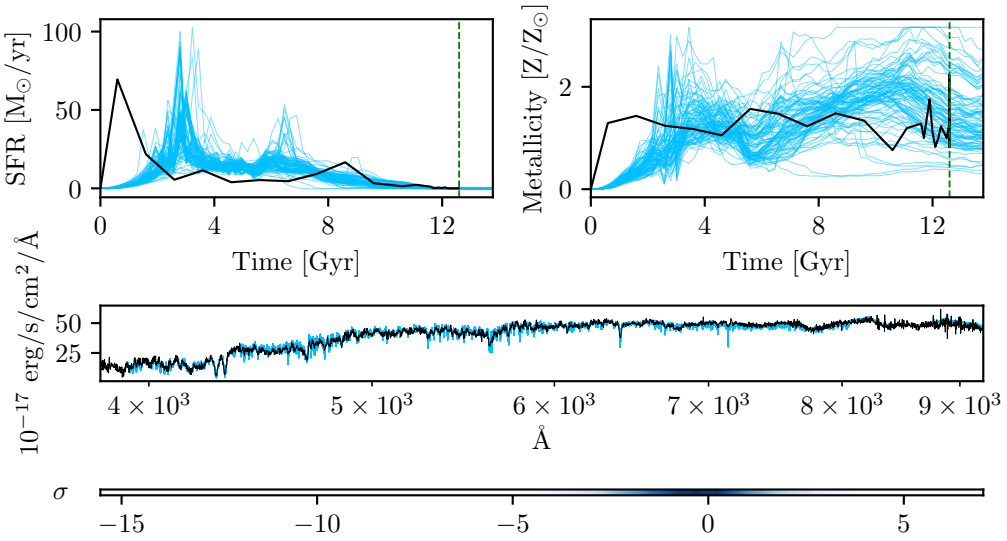

Figure 5: Top: 100 SFHs and MHs samples (blue) for MaNGA index 1-384872, with the FIREFLY estimates (black) with $A_V = 0.339$. The time axis shows age of the universe and the dashed green line shows the time at redshift 0.0873. Middle: observed spectrum (black) and posterior spectra (blue). Bottom: density of residuals.

D Dust experiments

We test the effect of increasing the value of A_V for our inference on MaNGA index 1-384872. In Figure 5, we show the results of sampling with $A_V = 0.339$. This value was calculated by removing zero-values when taking the mean of A_V from the MaNGA data. In Figure 6 we show the results of sampling with $A_V = 0.506$, arbitrarily chosen, remarking that the overall shape of the SFHs is modified. When testing on mock observations, we find that our model pushes for higher metallicities when fixing A_V to smaller values, indicating that further work on our model is required to achieve consistency at all values of A_V.

E Exaggerated spectral features

The lopsided residuals are driven by localized features in the data that are difficult to explain under our current model. By masking out these features, shown in Figure 7, we achieve more balanced residuals. In Figure 8, we zoom in on these areas. In the left panel, the SDSS model shows a much shallower Balmer line, suggesting that our model's high metallicity estimates may be driving the overly deep features seen in the posterior samples. The deep lines may also be a result of the spectral library used. Future work will investigate the effect on sampling when using different forward model configurations than what we specify in Appendix A.1.

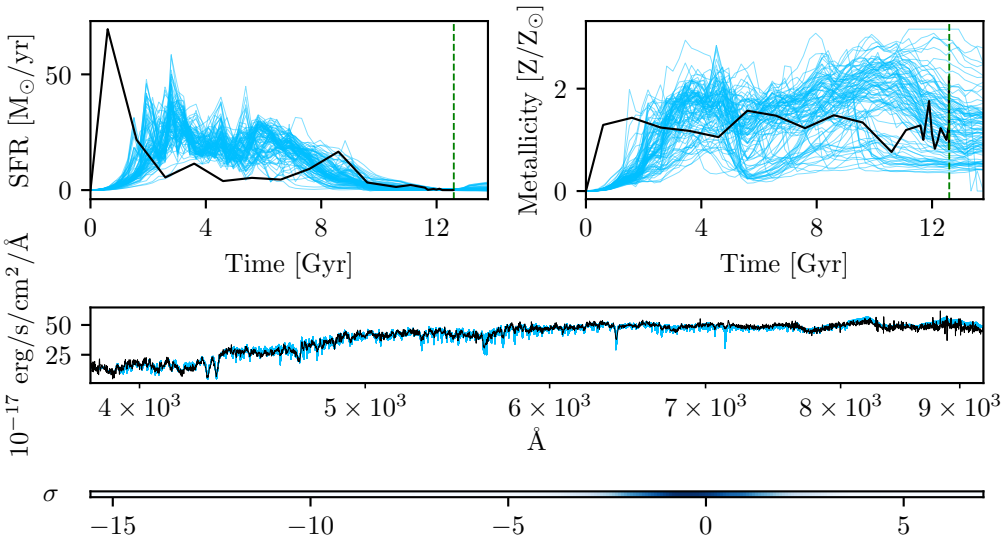

Figure 6: Top: 100 SFHs and MHs samples (blue) assuming $A_V = 0.506$ for MaNGA index 1-384872, and FIREFLY estimates (black). The time axis shows age of the universe and the dashed green line shows the time at redshift 0.0873. Middle: observed spectrum (black) and posterior spectra (blue). Bottom: density of residuals.

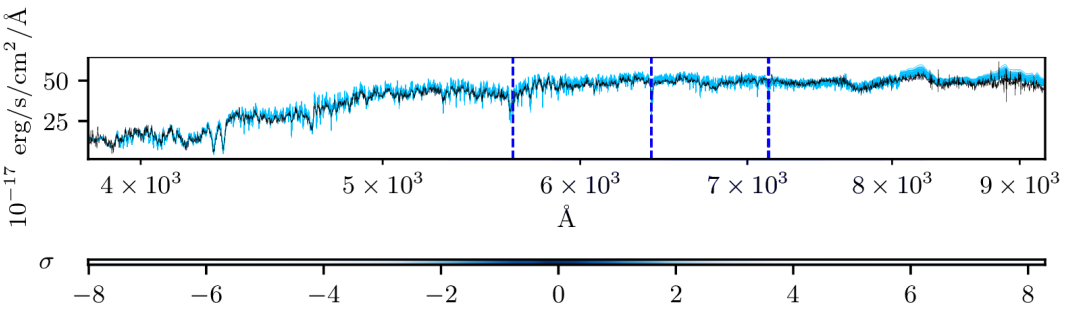

Figure 7: The three areas (dark blue) with maximum negative residuals are masked out, yielding more symmetric residuals.

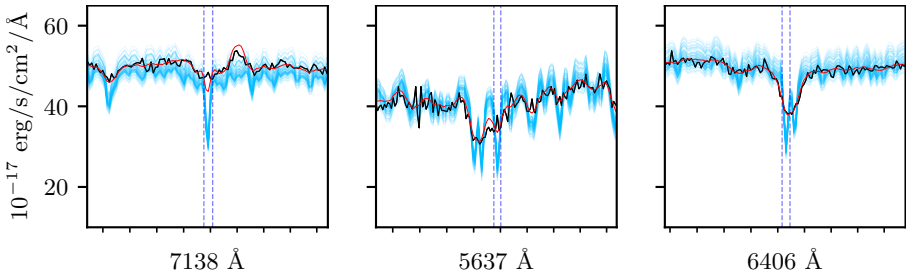

Figure 8: An up-close look of the three areas masked out in Figure 7. The SDSS best fit model is plotted in red.