

**Leveraging Machine Learning and Hydrodynamic Simulations of Galaxy
Formation to Illuminate the Effects of Incident Radiation Fields on Atomic Gas
Cooling and Heating Functions**

by

David B. Robinson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in the University of Michigan
2025

Doctoral Committee:

Professor Camille Avestruz, Chair

Professor Fred Adams

Professor Lia Corrales

Professor Nickolay Gnedin, University of Chicago

Professor Josh Spitz

David B. Robinson

dbrobins@umich.edu

ORCID iD: 0000-0002-3751-6145

© David B. Robinson 2025

DEDICATION

In memory of my Poppop, Bill Schmitt, who always reminded me to see the fun in studying science.

ACKNOWLEDGEMENTS

The research presented in this dissertation would not have been possible without the support of many people. Acknowledgments specific to each project are included in the corresponding chapters.

Thank you to my Ph.D. advisor Camille Avestruz, and to Nick Gnedin, for making all of the research in this dissertation possible and giving me so many opportunities to develop as a researcher. I am so grateful that the trial project on galaxy cooling and heating I started on 6 years ago led to follow-up projects and now forms the backbone of this dissertation!

Thank you to the funding sources that have helped support my research in graduate school, including the University of Michigan Physics department, the Leinweber Graduate Fellowship through the Leinweber Center for Theoretical Physics at the University of Michigan, the Michigan Space Grant Consortium Graduate Fellowship, and the U.S. Department of Energy (DoE) Office of Science Graduate Student Research Fellowship (SCGSR). The SCGSR enabled me to spend a year at Fermilab working on a project that became Chapter 3 of this dissertation.

Thank you to all the members of my committee for your helpful advice throughout the writing process. Thank you to the postdocs and other graduate students in the Avestruz Lab for Computational Cosmology and Astrophysics for many useful and interesting conversations throughout graduate school, in particular Ismael Mendoza, Raziq Noorali, and Deric Jones. I have also had the honor of working with two amazing REU students at the University of Michigan: Alexandra Wells and Ella Werre. Thank you for your patience with me through some of my first mentoring experiences, and all of the diligent work and initiative you put into your projects.

Teaching has been an integral part of my graduate school experience. Thank you to the workshop leaders and consultants at the University of Michigan Center for Research on Learning & Teaching who have helped me better understand and reflect on my teaching practices. Thank you to the professors who have been my mentors about teaching, including Ramón Torres-Isea, Wayne Lau, Josh Spitz, and especially Henriette Elvang.

As a graduate student, it's easy to get stuck in the weeds of the technical details of a project and forget about the bigger picture. Discussing science with general audiences is a

great way to avoid this trap, and has been an incredibly enriching part of my time in graduate school. The University of Michigan Museum of Natural History Science Communication Fellows program and Nerd Nite Ann Arbor were a lot of fun to participate in, and left me with some fascinating audience questions. I especially want to thank Researchers Expanding Lay-Audience Teaching and Engagement (RELATE), a graduate-student led science communication workshop at U-M. I first took their workshop online in 2020. The skills I learned have been very useful for presentations in the years since, even for the more technical ones! Since 2020, I have been on the coordinator team as the workshops transitioned from virtual back to in-person. Helping teach these workshops has forced me to think about teaching in a very different way than physics labs and discussion sections, and has (almost) always been fun! Thank you to the many graduate student coordinators who have been part of making RELATE such a positive experience.

Last but certainly not least, thank you to my family, who have been a constant source of support and encouragement along this journey. Thank you to my parents for always giving me (and the cats!) a safe place to call home, especially during the pandemic lockdown. You gave me the best start to my education I could ask for by making sure I grew up as a reader. Thank you for helping me to explore new interests throughout my life, and for doing your best not to fall asleep when I talk about what I'm working on!

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	viii
LIST OF TABLES	xiv
LIST OF ACRONYMS	xvi
LIST OF SYMBOLS	xvii
ABSTRACT	xix

CHAPTER

1 Introduction	1
1.1 Galaxy formation in a cosmological context	1
1.1.1 Structure formation and dark matter halos	2
1.1.2 Baryons within dark matter halos	5
1.2 Gas cooling and heating processes	8
1.2.1 Energy level transitions	9
1.2.2 Ionization and recombination	11
1.2.3 Free electron cooling	12
1.2.4 Non-atomic processes	13
1.3 Photoionization codes	14
1.4 Hydrodynamic simulations	14
1.4.1 Gravity and dark matter	15
1.4.2 Incorporating the gas: hydrodynamics	16
1.4.3 Subgrid physics models	19
1.4.4 Epoch of Reionization simulations and radiative transfer	21
1.4.5 Zoom-in and isolated galaxy simulations	24
1.5 Machine learning	24
1.5.1 Machine learning for interpolation	25
1.5.2 Gradient-boosted tree algorithms	26
1.5.3 Feature importance analysis with SHapley Additive exPlanation (SHAP) values	28

1.6	Outline of dissertation	29
2	Can Cooling and Heating Functions Be Modeled with Homogeneous Radiation Fields?	31
2.1	Motivation	31
2.2	Methodology	33
2.2.1	Cosmic Reionization on Computers (CROC) simulations	33
2.2.2	Cooling and heating functions	33
2.2.3	Cooling and heating functions in the simulations	35
2.2.4	Median cooling and heating rates and functions of the Interstellar Medium (ISM)	37
2.2.5	Numerical artifacts	39
2.3	Results	41
2.3.1	Actual rates vs. median functions	41
2.3.2	Redshift trends	43
2.3.3	Mass dependence	44
2.3.4	Density and metallicity ranges	47
2.4	Summary and discussion	49
2.5	Acknowledgements	53
3	Exploring the Dependence of Gas Cooling and Heating Functions on the Incident Radiation Field with Machine Learning	54
3.1	Motivation	54
3.2	Data and methods	58
3.2.1	Model data input	58
3.2.2	eXtreme Gradient Boosting (XGBoost)	60
3.2.3	Training data preparation	61
3.2.4	Principal Component Analysis (PCA) of rates	62
3.2.5	How to define an error?	66
3.2.6	Model training	66
3.2.7	Feature importances with SHAP values	67
3.2.8	Interpolation in metallicity	70
3.3	Results	71
3.3.1	Comparison on training data	71
3.3.2	Comparison on off-grid data	74
3.4	Conclusions and discussion	76
3.5	Acknowledgements	78
3.6	Appendix	79
3.6.1	Hyperparameter tuning	79
3.6.2	Comparison on training data for all metallicities	81
4	On the Minimum Number of Radiation Field Parameters to Specify Gas Cooling and Heating Functions	83
4.1	Motivation	83
4.2	Methodology	86

4.2.1	The training grid	86
4.2.2	Radiation field sampling	87
4.2.3	Machine learning models	91
4.2.4	Feature importance with SHAP values	93
4.3	Results	97
4.4	Conclusions and discussion	99
4.5	Acknowledgements	102
4.6	Appendix	102
4.6.1	Logarithmic bin spacing	102
4.6.2	Hyperparameter tuning	103
4.6.3	Model error comparisons	104
4.6.4	The performance of 2 bin models	105
5	The Effects of Different Cooling and Heating Function Models on a Simulated Analog of NGC300	111
5.1	Motivation	111
5.2	Methodology	113
5.2.1	Simulations	113
5.2.2	Cooling and heating function models	114
5.2.3	CII luminosity	117
5.3	Results	118
5.3.1	Phase diagrams and residuals	118
5.3.2	Effects on CII luminosity	123
5.4	Summary and discussion	123
5.5	Acknowledgements	126
5.6	Appendix: Convergence of the gas phase diagram	127
6	Conclusions and Future Work	129
6.1	Summary	130
6.2	Future work	132
6.2.1	Metallicity dependence of cooling and heating functions	132
6.2.2	Non-atomic cooling and heating processes	134
6.2.3	Other applications of our machine learning framework	134
6.2.4	Improved interpolation tables	135
6.2.5	Extended simulation comparisons	136
	BIBLIOGRAPHY	138

LIST OF FIGURES

FIGURE

1.1	Timeline of the universe, from the Epoch of Recombination and last scattering of the Cosmic Microwave Background (CMB) on the left to the present day universe on the right. Reproduced from Robertson et al. (2010).	2
1.2	Flowchart of interactions during galaxy evolution. The dark gray box with a dashed border represents the entire galaxy, the lighter gray boxes with solid borders represent individual (baryonic) components, and the arrows represent interactions between these components. Reproduced from Mo et al. (2010). . .	6
1.3	Atomic gas cooling function $\Lambda_{e,H}$ (here defined as the cooling rate divided by $n_e n_H$) vs. temperature for a gas with solar element abundances and no incident radiation field, calculated with the photoionization code Cloudy (Ferland et al., 1998, see section 1.3). The shaded grey band and dotted grey line show the calculation of the total cooling rate using two different versions of Cloudy. The solid and dashed red and black lines show the contributions from various elements and their associated ionization states. Reproduced from Gnat & Ferland (2012).	10
1.4	Snapshots from various hydrodynamic cosmological simulations: Illustris (Vogelsberger et al., 2014), IllustrisTNG (Springel et al., 2018), Magneticum (Bocquet et al., 2016), EAGLE (Schaye et al., 2015), Romulus25 (Tremmel et al., 2017), SIMBA (Davé et al., 2019), Massiveblack-II (Khandai et al., 2015), and Horizon-AGN (Dubois et al., 2014). Reproduced from Vogelsberger et al. (2020).	18
1.5	Slices through one realization of the CROC (Gnedin, 2014) simulations of the Epoch of Reionization (EoR) showing the neutral hydrogen fraction at $z = 8.04$ (top panel) and $z = 7.01$ (bottom panel). Reproduced from Gnedin & Kaurov (2014).	22
1.6	Slices through the disk of the Semenov et al. (2021) idealized galaxy simulation showing the gas number density (leftmost panel), temperature (middle left panel), turbulent velocity dispersion (middle right panel), and Ultraviolet (UV) radiation field strength (rightmost panel). Reproduced from Semenov et al. (2021).	25
1.7	A schematic picture of XGBoost regression models. Reproduced from Wang et al. (2019).	28
1.8	A schematic depiction of the SHAP values (see equation (1.22)) for a machine learning model with 4 features. The ϕ_i shown in this figure should be averaged over all possible orderings of the 4 features to obtain the SHAP values. Reproduced from Lundberg & Lee (2017).	29

2.1	Histograms of the CVI photoionization rate P_{CVI} vs the Lyman-Werner band photodissociation rate P_{LW} (left) or the HI photoionization rate P_{HI} (right) for cells in the 50 randomly selected central halos in the fiducial mass range at $z \sim 5$. The dashed green horizontal line shows the cutoff between the low and high radiation field parts of the distribution ($P_{\text{CVI}} = 2 \times 10^{-20} \text{ s}^{-1}$). The two populations of cells are distinguished by color. The low P_{CVI} values and correlation between P_{CVI} and P_{LW} for the low P_{CVI} population have no obvious physical reason and are likely to be numerical artifacts.	40
2.2	Comparison of actual rates (dashed blue curves) and the median ISM (dotted green) and instantaneous (solid orange) functions for cooling (top panel) and heating (bottom panel) for the 50 randomly selected central halos at $z \sim 5$ in the fiducial mass range $M > 10^{10} h^{-1} M_{\odot}$ (the same as shown in Fig. 2.1). The shaded bands show the 25th-75th percentile spread.	42
2.3	Median ISM cooling (top panel) and heating (bottom panel) functions vs. temperature for up to 50 central halos in the fiducial mass range at $z \sim 5, 8, 9, 10$ (solid lines) and for sub halos at $z \sim 5$ (dashed lines). The solid lines show the median, while the shaded bands show the 25th-75th percentile spread for central halos at $z \sim 5$. For direct comparison, we overplot the heating functions in the upper panel with faint lines; the intersection between the cooling and the heating function corresponds to the equilibrium temperature.	45
2.4	Actual rates (blue) and median ISM functions (orange) for cooling (top panel) and heating (bottom panel) for central halos in three mass bins at $z \sim 5$. The 25th-75th percentile shaded region is shown only for the middle mass bin, but the widths are similar for all bins.	46
2.5	Median ISM cooling (top panel) and heating (bottom panel) functions vs. temperature for selected cells from the random subsample of 50 central halos at $z \sim 5$ in our fiducial mass range (the same sample as previous figures) for cells with $0.01 < n_b < 0.1$ (blue), $0.1 < n_b < 1$ (orange), $1 < n_b < 10$ (green), $1 < n_b < 3$ (red), and $3 < n_b < 10 \text{ cm}^{-3}$ (purple). The bands correspond to the 25th to 75th percentile spread for each density range.	48
2.6	The median baryon number density n_b multiplied by the median ISM heating function vs. temperature for selected cells from the random subsample of 50 central halos at $z \sim 5$ in our fiducial mass range (the same sample as above) for cells with $0.01 < n_b < 0.1$ (blue), $0.1 < n_b < 1$ (orange), $1 < n_b < 10$ (green), $1 < n_b < 3$ (red), and $3 < n_b < 10 \text{ cm}^{-3}$ (purple). The shaded regions correspond to the 25th and 75th percentile n_b values respectively multiplied by the heating function evaluated for 25th and 75th percentile ISM properties.	50
2.7	Actual (solid lines) cooling and heating rates and median ISM (dashed lines) cooling (top panel) and heating (bottom panel) functions vs. temperature for selected cells from a random subsample of 50 central halos at $z \sim 5$ in our fiducial mass range (the same sample as above) for cells with $0.03 < Z/Z_{\odot} < 0.1$ (blue, as used for the above plots) and $0.1 < Z/Z_{\odot} < 0.3$ (orange).	51
3.1	Feature distributions of six scaled rates $\log(Q_j/Q_{\text{LW}})$. All have been linearly rescaled to the range $[0, 1]$, and their distributions normalized.	62

3.2	Spearman correlation matrix of the same six scaled rate features $\log(Q_j/Q_{\text{LW}})$ as Figure 3.1.	64
3.3	The relative importance c_j defined by equation (3.6) for scaled rates $\log(Q_j/Q_{\text{LW}})$, in decreasing order. For simplicity, only the 35 scaled rates with the highest relative importance are shown. We use the 21 rates with the highest significance (those to the left of the vertical dashed line) in our analysis, as all other rates have similarly small significance.	65
3.4	Mean absolute SHAP value for 500 randomly selected points in 20% of the training grid withheld from model training (see section 3.2.3) for models trained with 21 scaled rate features selected from the PCA analysis described in section 3.2.4. Cooling function models are shown in blue in the left column, with heating function models in red in the right column. The rows are for models at the 5 metallicity values in the training data. Only the 7 features with the largest mean absolute SHAP value are shown. Note that the full descriptions of the features are $\log(T/\text{K})$ for T , $\log(n_{\text{H}}/\text{cm}^{-3})$ for n_{H} , $\log(Q_{\text{LW}}/\text{cm}^3 \text{s}^{-1})$ for LW, and $\log(Q_j/Q_{\text{LW}})$ otherwise.	69
3.5	Cumulative distribution function of errors $\Delta \log \mathcal{F}$ (see equation (3.9)) on the training data (described in section 3.2.1) for the cooling function Λ (left column, in blue) and heating function Γ (right column, in red) at fixed metallicity $Z = 0$ (top row) and $Z = Z_{\odot}$ (bottom row), for the interpolation table in Gnedin & Hollon (2012) (solid lines) and XGBoost models trained using the same scaled photoionization rates $j = \text{HI, HeI, CVI}$ (dashed lines), 21 scaled rates from the PCA analysis described in section 3.2.4 (dotted lines), and the top 3 scaled rates from the SHAP value analysis, described in section 3.2.7 (dashed-dotted lines). For the “XGBoost (HI, HeI, CVI)” model, we also include the cumulative error distribution function on 20% of the training grid withheld from model training for otherwise identical models (thin dashed lines).	73
3.6	Cumulative distribution function of errors $\Delta \log \mathcal{F}$ on the evaluation data described in section 3.2.8. Left panel: Comparison of results from the interpolation table in Gnedin & Hollon (2012) and the constrained quadratic fit described in section 3.2.8 between XGBoost models trained with the same scaled rates (HI, HeI, CVI). Right panel: Comparison of the same constrained quadratic fit between XGBoost models trained with three different sets of scaled rates. . . .	75
3.7	Results from a grid search in <code>max_depth</code> and <code>n_estimators</code> for cooling function model trained with the 21 scaled rates from the PCA of section 3.2.4 at $Z/Z_{\odot} = 1$. The three panels show heatmaps of the Mean Squared Error (MSE) in $\log \mathcal{F}$ on the test set (a randomly select 20% of the training data table described in section 3.2.1 (left), the training time in seconds (middle), and the frequency of errors $\Delta \log \mathcal{F}$ larger than 0.3 on the test set (right). The heatmaps are all on a logarithmic scale.	81
3.8	Same as Figure 3.5, but now for all 5 metallicity values in the training data table (described in Table 3.1).	82

4.1	Examples of the radiation field J_ν (normalized by its overall amplitude J_0) defined by equation (4.2) for different choices of α , f_Q , and τ_0 . The dashed vertical lines indicate our choice of bin edges, at 0.5, 1, 4, 7, 10, 13, and 16 Ry.	89
4.2	The absolute value of the Pearson (left panel) and Spearman (right panel) correlation coefficients between binned radiation field intensities $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ for the values in Table 4.1, with linear bin spacing. Given the high correlation between the highest three energy bins, we exclude the two highest bins (16 – 19 Ry and 19 – 22 Ry) from our initial model training.	90
4.3	Mean absolute SHAP values on 500 randomly selected points in the test set (20% of the training grid withheld from model training) for XGBoost models trained using 6 bins (see ‘6 bins’ column of Table 4.2). Models for the cooling function (top panel) and heating function (bottom panel) at $Z/Z_\odot = 1$ are shown. Only the 7 features with the largest mean absolute SHAP values for each model are shown. At $Z/Z_\odot = 1$, the gas temperature and overall amplitude of the radiation field are most important.	94
4.4	The mean absolute SHAP values for the quantity $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$, divided by that for $\log(\langle J_{0.5-1 \text{ Ry}} \rangle \text{ cm}^{-3} / n_b / J_{\text{MW}})$. For each feature, the values are shown for the cooling function (above, blue) and the heating function (below, red). The 5 panels are for models trained at different metallicities, from $Z/Z_\odot = 0$ at the top to $Z/Z_\odot = 3$ at the bottom. Either the 1 – 4 or 13 – 16 Ry bins (top and bottom bars in each panel) are the most important for both cooling and heating at each metallicity.	96
4.5	Error distributions for the interpolation table of Gnedin & Hollon (2012) (thin solid lines) and our XGBoost models trained with 6 bins (thick solid lines, described in section 4.2.3), 3 bins (thick dashed lines, described in section 4.2.4), 4 rates (thick dash-dotted lines, described in section 4.2.4), and 3 rates (thick dotted lines, described in section 4.2.4) for models trained and evaluated on the entire training grid. See Table 4.2 for the features included in each model. Error distributions for the cooling function are shown in the left column, with the heating function in the right column. All models are trained to minimize the MSE on the training set. Our XGBoost models all have similar performance on the training data, and significantly outperform an interpolation table.	98
4.6	The absolute value of the Pearson (left panel) and Spearman (right panel) correlation coefficients between binned radiation field intensities $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ for the values in Table 4.1 with logarithmic bin spacing.	103

4.7	Error distributions for our XGBoost models trained with 6 bins (solid lines, described in section 4.2.3), 3 bins (dashed lines, described in section 4.2.4), 4 rates (dash-dotted lines, described in section 4.2.4), and 3 rates (dotted lines, described in section 4.2.4) for models trained on 80% of the training grid and evaluated on the withheld 20% test subset (see section 4.2.3). See Table 4.2 for the features included in each model. The error distributions for the interpolation table of Gnedin & Hollon (2012) on the entire training grid are overplotted as thin solid lines. Distributions for the cooling function are shown in the left column, with the heating function in the right column. All models are trained to minimize the MSE on the training set. All our XGBoost models have comparable performance on the test subset.	106
4.8	The predicted heating function for the training grid point described in Table 4.6 (thick orange curve), and the true heating functions for every training grid point with the same density, as well as binned radiation field features (see Table 4.2) within 10% of the values for the point in Table 4.6 (thin light blue curves). . . .	108
4.9	The radiation field specific intensity J_ν (equation (4.2)) for the points from Figure 4.8 with the lowest (left panel) and highest (right panel) heating function values at $\log(T/K) = 1$	109
5.1	Upper panel: Median (curves), 25th-75th percentile (darker bands), and 10th-90th percentile (lighter bands) temperatures as a function of gas density bin for GH12 (blue, solid curve) and XGB (orange, dashed curve) runs. Bottom panel: ratio of median temperature for the XGB run to the GH12 run as a function of density bin.	120
5.2	Residual of gas mass between runs using the interpolation table of GH12 and XGBoost cooling and heating functions after 5 Myr. There is a ‘critical curve’ where the residual is equal to 0.	121
5.3	Predicted cooling (blue curves) and heating functions (red curves) for the GH12 interpolation table (Gnedin & Hollon, 2012, solid curves) and XGBoost machine learning models from Chapter 3 (dashed curves) at $n_b = 1 \text{ cm}^{-3}$. The cooling and heating curves for a given model intersect at the equilibrium temperature. .	122
5.4	The CII emission ratio r_j , as defined in equation (5.3) as a function of baryon number density n_b (upper panel) excited by interactions with electrons (solid blue), atomic hydrogen or helium (dashed orange), CMB photons (dash-dotted green), and molecular hydrogen with para (dashed red) and ortho (solid purple) spins. For comparison, the bottom panel shows the overall baryon number density profile for the GH12 simulation run (solid brown).	124
5.5	Residual between residuals $\delta_{1 \text{ Myr}, 2 \text{ Myr}}$ (top, see equation (5.12)) and $\delta_{4 \text{ Myr}, 5 \text{ Myr}}$ (bottom).	128

6.1	The total (atomic and molecular, with abundances of $f_{\text{H}_2} = 10^{-5}$ for molecular hydrogen and $f_{\text{HD}} = 10^{-8}$ for hydrogen-deuterium) gas cooling function with $n_{\text{H}} = 1 \text{ cm}^{-3}$ for metal number fractions (relative to hydrogen) of 10^{-6} (black dotted), 10^{-5} (dark blue short-dashed), 10^{-4} (light blue long-dashed), and 10^{-3} (red solid) interpolated from the tables of Sutherland & Dopita (1993). Reproduced from Maio et al. (2007).	133
6.2	Galaxy scaling relations for the radius containing 50% of the total stellar luminosity $R_{*,1/2}$ (top row), dark matter fraction f_{DM} within $R < R_{*,1/2}$ (second row), dark matter mass within $R < R_{*,1/2}$ (third row), and total mass M_{tot} (bottom row) with stellar mass M_* . The circular points show predictions from IllustrisTNG simulations in the CAMELS suite (Villaescusa-Navarro et al., 2021). The shaded bands in the top three rows are observational constraints from La Barbera et al. (2010) (grey), Cappellari et al. (2013) (red), and Bundy et al. (2015) (orange). The pink band in the bottom row is a theoretical relation from Moster et al. (2013). All shaded bands are bounded by the 16th and 84th percentile, with the middle line showing the median. Reproduced from Busillo et al. (2025).	136

LIST OF TABLES

TABLE

1.1	Cosmological parameter values from Planck Collaboration et al. (2020).	3
2.1	Cell fraction with P_{HI} , P_{HeI} , or $P_{\text{CVI}} = 0$ after density and metallicity cuts.	39
2.2	Redshift evolution of low P_{CVI} cell fraction.	41
3.1	The parameters describing the training data table and the values they take. Here, $J_{\text{MW}} = 10^6 \text{ photons cm}^{-2} \text{ s}^{-1} \text{ ster}^{-1} \text{ eV}^{-1}$ (Gnedin & Hollon, 2012). The $Z = 0$ case actually uses $Z = 10^{-4} Z_{\odot}$	59
3.2	Minimum and maximum values of $\log(Q_j/Q_{\text{LW}})$ for some rates of interest attained in the training data table described in Table 3.1.	60
3.3	The three most important scaled rates, by mean absolute SHAP value, for 500 randomly sampled test points on cooling function models with the 21 scaled rate features from section 3.2.4, retrained on 80% of the training data table using the optimized hyperparameters found via the procedure in section 3.2.6.	68
3.4	Same as Table 3.4, but now for the heating function.	69
3.5	MSEs for the cooling function models shown in Figure 3.5 and Figure 3.8. For the interpolation table in Gnedin & Hollon (2012), the MSE is on the entire training grid described in section 3.2.1. The values in parentheses are the fraction of ‘catastrophic errors’ on the training set where the prediction is negative (so $\Delta \log \mathcal{F} = \infty$). We remove these points from the MSE calculation. For the three XGBoost models, the MSE is shown for a model trained and evaluated on the entire training grid (‘training’ columns), and for a model trained on 80% of the training grid and evaluated on the 20% that is withheld (‘test’ columns).	72
3.6	Same as Table 3.5, but now for heating function models.	74
3.7	MSEs on the off-grid data set described in section 3.2.8 for the models in Figure 3.6. For the GH12 column, the values in parentheses are the fraction of ‘catastrophic errors’ on the evaluation set where the prediction is negative (so $\Delta \log \mathcal{F} = \infty$). We remove these points from the MSE calculation.	76
4.1	Gas and radiation field parameters used to train XGBoost models of Cloudy-computed cooling and heating functions. We normalize the radiation field amplitude J_0 by $J_{\text{MW}} = 10^6 \text{ photons cm}^{-2} \text{ s}^{-1} \text{ ster}^{-1} \text{ eV}^{-1}$	88

4.2	The features sets used for each XGBoost model considered in this chapter. Note that all features are the logarithm of the quantities shown here. Above the line are features describing gas properties, which are the same for each of our models. Below the line are features related to the incident radiation field. The first row describes the radiation field features containing the overall amplitude of the radiation field J_0 (see equation (4.2)). The remaining rows show all other radiation field features, which are scaled to the rate or bin containing the overall amplitude.	92
4.3	MSEs (see equation (4.8)) for the cooling function XGBoost models in the left column of Figure 4.5 and Figure 4.7 (see Table 4.2). All MSEs have similar orders of magnitude.	105
4.4	MSEs (see equation (4.8)) for the heating function XGBoost models in the right column of Figure 4.5 and Figure 4.7 (see Table 4.2). All MSEs have similar orders of magnitude.	107
4.5	MSEs (see equation (4.8)) for 3 bin and 2 bin heating function models (see Table 4.2). The 2 bin model is systematically ~ 2 orders of magnitude less accurate than the 3 bin model.	107
4.6	Parameters for a point in the training grid (see Table 4.1) with $\Delta \log \Gamma > 0.8$ for a 2 bin heating function model at $Z/Z_\odot = 0$	107
5.1	Important parameters for our simulation runs	114
5.2	MSEs of the cooling and heating function approximations used in our simulations on the training grid of Cloudy calculations at gas metallicity $Z = 0.3Z_\odot$. The MSEs for the XGBoost models from Chapter 3 are 2-3 orders of magnitude smaller than those for the Gnedin & Hollon (2012) interpolation table.	116

LIST OF ACRONYMS

AGN	Active Galactic Nuclei
AMR	Adaptive Mesh Refinement
ART	Adaptive Refinement Tree
BAO	Baryon Acoustic Oscillations
CGM	Circumgalactic Medium
CIE	Collisional Ionization Equilibrium
CMB	Cosmic Microwave Background
CROC	Cosmic Reionization on Computers
EoR	Epoch of Reionization
FFT	Fast Fourier Transform
GPU	Graphics Processing Unit
IMF	Initial Mass Function
ISM	Interstellar Medium
ISRF	Interstellar Radiation Field
LIM	Line-Intensity Mapping
MSE	Mean Squared Error
PCA	Principal Component Analysis
SHAP	SHapley Additive exPlanation
SMBH	Supermassive Black Hole
SPH	Smoothed Particle Hydrodynamics
UV	Ultraviolet
XGBoost	eXtreme Gradient Boosting

LIST OF SYMBOLS

z	Redshift
a	Cosmic scale factor
H	Hubble rate
H_0	Hubble constant
h	Reduced Hubble constant
δ	Overdensity
λ_J	Jeans length
Λ	The gas cooling function
Γ	The gas heating function
\mathcal{F}	The gas cooling function or heating function
J_ν	Specific intensity of a radiation field as a function of frequency
U	Thermal energy density
T	Temperature
n	Number density
ρ	Mass density
n_b	Baryon number density
n_H	Hydrogen number density
Z	Metallicity
Z_\odot	Solar metallicity
P_j	Photoionization rate (for band j)
Q_j	Specific photoionization rate (for band j)

M_{\odot}	Solar mass
M_{vir}	Virial mass (of a dark matter halo)
R_{vir}	Virial radius (of a dark matter halo)
τ	Optical depth
σ	Cross section
μ	Molecular weight

ABSTRACT

Galaxies and stars form out of clouds of gas that collapse under the influence of gravity. This collapse stops when the gas heats up enough for its outward thermal pressure to counteract the gravitational compression. The cooling and heating functions of gas determine how efficiently gas can cool down and heat up by respectively emitting or absorbing photons. Gas thermodynamics drives the formation and evolution of galaxies. In this dissertation, we explore how locally varying radiation fields impact gas cooling and heating functions. We begin with a discussion of galaxy formation and evolution, and the atomic line emission and ionization processes that cool and heat interstellar gas. We also describe the main computational tools used in this dissertation: photoionization codes, hydrodynamic simulations, and machine learning.

We then present results from galaxies simulated by the Cosmic Reionization on Computers (CROC) project, a hydrodynamic simulation of the Epoch of Reionization (EoR). The EoR is the era when ionizing radiation from early galaxies and stars caused the gas in the intergalactic medium to transition from neutral to ionized. We find that the cooling and heating functions of gas in these simulated galaxies cannot be well-described with a spatially constant radiation field, demonstrating the necessity of accurately accounting for spatial variations in the radiation field intensity.

Next, we apply a machine learning approach to the problem of accurately approximating cooling and heating functions in a locally varying radiation field. We train models to predict cooling and heating functions at fixed metallicity using gas properties and four or more photoionization rates describing the radiation field. These models are able to compute cooling and heating functions more accurately than existing interpolation table methods. At arbitrary metallicity, we are able to reduce the frequency of the largest errors compared to an interpolation table approach. Interpolating between fixed metallicity models is the main bottleneck to further improvement. In a corollary study, we train machine learning models where the only radiation field parameters are its average intensity in various energy bins. We find that only three bins (one just below and one just above the hydrogen ionization energy, and one bin corresponding to high energy photons) are sufficient to accurately capture the cooling and heating behavior at fixed metallicity.

Finally, to investigate the effects of these various modeling choices in a physical context, we implement one set of our machine learning cooling and heating models into a hydrodynamic simulation of an isolated galaxy. To assess the effects of improved accuracy in cooling and heating function approximations, we compare the gas temperature-density phase diagrams of this simulation to the same simulation with cooling and heating rates calculated using an interpolation table.

This dissertation lays groundwork for more accurate modeling of baryonic processes in next-generation hydrodynamic simulations. This will improve our understanding of both systematic errors from baryonic uncertainties in the constraints on cosmological parameters from large-scale structure surveys and our understanding of gas physics in galaxy formation.

CHAPTER 1

Introduction

This dissertation explores how the thermodynamics of gas in and around galaxies is affected by the local radiation field environment. In this introduction, we provide a brief overview of the theoretical framework of galaxy evolution (section 1.1) and the various thermodynamic processes that heat and cool interstellar gas (section 1.2). Then, we discuss the main computational tools used in the studies presented in subsequent chapters: photoionization codes (section 1.3), hydrodynamic galaxy evolution simulations (section 1.4), and machine learning (section 1.5). Finally, we outline the rest of this dissertation in section 1.6.

1.1 Galaxy formation in a cosmological context

Galaxies are luminous, gravitationally-bound structures containing stars, gas, dust, and the mysterious invisible substance we call dark matter (Mo et al., 2010). It is now thought that almost all galaxies have a Supermassive Black Hole (SMBH) at their center (e.g. Binney, 1977; Soltan, 1982; Merritt, 2000). The universe today is full of such galaxies, including our home galaxy, the Milky Way. The field of galaxy formation and evolution studies how these structures first formed, and how they changed through time to evolve to their present state.

Galaxy formation is inextricable from *cosmology*, the study of the state and evolution of the universe on the largest scales. Almost all the galaxies we can see are receding away from the Milky Way, with more distant galaxies receding faster, indicating that our universe is expanding (Hubble, 1929). The standard model of cosmology (called Λ CDM) describes both the expansion of the universe and the evolution of small density perturbations in the very early universe into the large-scale structure we see today. Figure 1.1 is a schematic depiction of the history of the universe, including the formation of galaxies. In section 1.1.1, we discuss the formation of large-scale structure, in particular dark matter halos, in Λ CDM cosmology. We describe the formation and evolution of galaxies within these dark matter halos in section 1.1.2. Much of this overview is based on the galaxy formation textbook Mo




Figure 1.1: Timeline of the universe, from the Epoch of Recombination and last scattering of the CMB on the left to the present day universe on the right. Reproduced from Robertson et al. (2010).

et al. (2010).

1.1.1 Structure formation and dark matter halos

On cosmological scales, the universe appears to be both *homogeneous* (spatially uniform) and *isotropic* (the same in every direction). We describe spacetime in the expanding universe with *comoving* coordinates (t, x, y, z) and a scale factor $a(t)$, which depends only on the time coordinate t . If two points in space are at the comoving positions \vec{x}_1 and \vec{x}_2 , then at time t , they are separated by a physical distance $\Delta r = a(t) |\vec{x}_2 - \vec{x}_1|$. The expansion of the universe means that light emitted from any point is *redshifted* as seen from any other point. This cosmological redshift z is given by:

$$1 + z = \frac{a(t_0)}{a(t_e)}, \quad (1.1)$$

where t_e is the time the photon is emitted, and t_0 is the time it is observed (i.e. now). It is conventional to set the current scale factor $a(t_0) = 1$, so that $1 + z = a^{-1}$ (Mo et al., 2010).

Applying the equations of general relativity to this homogeneous and isotropic spacetime yields the Friedmann equation:

$$\left(\frac{\dot{a}}{a}\right) = H^2(t) = \frac{8\pi G}{3}\rho = H_0^2 [\Omega_m a^{-3} + \Omega_\Lambda + \Omega_r a^{-4} + (1 - \Omega_m - \Omega_\Lambda - \Omega_r) a^{-2}], \quad (1.2)$$

where G is Newton's gravitational constant, ρ is the energy density of the universe (including vacuum energy), and \dot{a} indicates a derivative with respect to time. We have also introduced the Hubble parameter $H(t) = \frac{\dot{a}}{a}$, the expansion rate of the universe. The current value of the Hubble parameter is the Hubble constant $H_0 = H(t_0)$. We define present-day density parameters $\Omega_i = \frac{\rho_i}{\rho_{\text{crit}}}$, where the critical density is defined as $\rho_{\text{crit}} = \frac{3H_0^2}{8\pi G}$. In Λ CDM, the

Parameter	Value
H_0	$67.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$
Ω_m	0.315
Ω_Λ	0.685
Ω_b	0.049

Table 1.1: Cosmological parameter values from Planck Collaboration et al. (2020).

only components that contribute to the energy density of the universe are: matter (both ordinary ‘baryonic’ matter and invisible dark matter, the ‘CDM’ in Λ CDM), with a density that dilutes with the expansion of the universe as a^{-3} ; the cosmological constant Λ , which has negative pressure and a constant energy density as the universe expands; and radiation (photons and neutrinos), with an energy density that dilutes as a^{-4} (the number of photons per unit volume dilutes as a^{-3} , and the energy of each photon is redshifted by an additional factor of a^{-1}). If the universe is flat, the total energy density is exactly equal to the critical density (i.e. $\Omega_m + \Omega_\Lambda + \Omega_r = 1$), so that the a^{-2} term in equation (1.2) vanishes (Mo et al., 2010).

The Hubble constant and density parameters Ω_i can be constrained from observations of the Cosmic Microwave Background (CMB; see Figure 1.1 and the discussion below). Table 1.1 shows the values inferred from Planck Collaboration et al. (2020). This table also includes the density parameter for normal ‘baryonic’ matter, Ω_b . Note that the total matter density $\Omega_m \sim 6\Omega_b$. The radiation density parameter Ω_r is not included as it is negligible (at least in this epoch). When plugged into the Friedmann equation (equation (1.2)), the values of Ω_m and Ω_Λ indicate that the expansion of the universe must currently be accelerating. This agrees with observations of the redshift and distance of Type Ia supernovae in the late universe. These measurements can also be used to constrain H_0 . However, these local universe measurements yield a much higher value of the Hubble constant, $H_0 = 73.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$. This is the so-called ‘Hubble tension’ (e.g. Brout et al., 2022).

A completely homogeneous and isotropic universe would be quite uninteresting. Fortunately, our universe is *not* homogeneous or isotropic on sufficiently small scales, as can be easily observed by looking at the night sky. It is thought that quantum mechanical processes created density fluctuations in the very early universe during an epoch called inflation, when the universe expanded at an accelerated rate. This could explain why the present-day universe is so nearly flat, i.e. why $\Omega_m + \Omega_\Lambda + \Omega_r \approx 1$. We define the *overdensity* δ through the equation:

$$\delta = \frac{\rho - \bar{\rho}}{\bar{\rho}} = \frac{\rho}{\bar{\rho}} - 1, \quad (1.3)$$

where $\bar{\rho}$ is the average density described by the Friedmann equation (equation (1.2)). When $\delta \ll 1$, the growth of the density perturbations δ is linear, and can be described in terms of their Fourier transform $\delta_{\vec{k}}$. For adiabatic initial conditions, the linear perturbations evolve according to:

$$\ddot{\delta}_{\vec{k}} + 2H\dot{\delta}_{\vec{k}} = \left[4\pi G\bar{\rho} - \frac{k^2 c_s^2}{a^2} \right] \delta_{\vec{k}}, \quad (1.4)$$

where c_s is the adiabatic sound speed:

$$c_s^2 = \left(\frac{\partial P}{\partial \rho} \right)_S, \quad (1.5)$$

the derivative of pressure with respect to density at constant entropy (Mo et al., 2010).

Solutions to equation (1.4) will propagate as sound waves if $4\pi G\bar{\rho} - \frac{k^2 c_s^2}{a^2} < 0$. This condition can be rearranged in terms of the *proper length* of the perturbation $\lambda = \frac{2\pi a}{k}$ as $\lambda < \lambda_J$, where the *Jeans length* λ_J is defined as:

$$\lambda_J \equiv c_s \sqrt{\frac{\pi}{G\bar{\rho}}}. \quad (1.6)$$

Equation (1.6) is just the distance a sound wave traveling at c_s can cross in the gravitational free-fall time $t_{\text{ff}} \propto \sqrt{\frac{1}{G\bar{\rho}}}$. If $\lambda < \lambda_J$, the thermal pressure propagates quickly enough to support the perturbation against gravitational collapse. For larger perturbations $\lambda > \lambda_J$, the gravitational collapse is too fast for the thermal pressure to counteract it, and the perturbation will grow (at least once the expansion of the universe is dominated by matter, i.e. $\Omega_m a^{-3} > \Omega_r a^{-4}$, which occurs at $z \sim 3600$, Mo et al., 2010).

The early universe was sufficiently hot that all atoms were ionized. Because photons Thomson scatter strongly off of free electrons, the photons and baryons were strongly coupled at this epoch. Sound waves called Baryon Acoustic Oscillations (BAO) could propagate in this baryon-photon fluid until the universe cooled enough for electrons and protons to combine to form neutral hydrogen at $z \sim 1100$. At this redshift, baryons and photons decouple, and the photons can free-stream as the CMB. Between inflation and $z \sim 1100$, sound waves in the baryon-photon fluid travel a characteristic distance known as the BAO scale (Mo et al., 2010).

The sound speed for baryons alone scales with their temperature as $c_s \propto T^{1/2}$. This speed is much smaller than the relativistic sound speed in the coupled baryon-photon fluid. After the baryons and photons decouple at $z \sim 1100$, density perturbations in the baryons can begin to gravitationally collapse, increasing the overdensity δ (Mo et al., 2010).

Dark matter is crucial to forming large-scale structures because the sound speed in dark

matter is always proportional to $T^{1/2}$. This means that dark matter perturbations can begin to grow gravitationally as early as $z \sim 3600$. Before this, the energy density of the universe is dominated by radiation, and the expansion of the universe is too fast for perturbations to grow significantly. After the decoupling of the baryon-photon fluid at $z \sim 1100$, baryons begin to fall into the gravitational potential wells of the dark matter perturbations (Mo et al., 2010).

When perturbations grow to $\delta \gtrsim 1$, linear perturbation theory (equation (1.4)) breaks down. Analytical calculations for geometrically simple cases show that perturbations will grow in both size and mass until they reach a critical density. At this point, non-linear gravitational collapse will begin to shrink the size of the halo. N-body simulations, where the dark matter is modeled as particles of fixed mass which interact only through gravity (see section 1.4.1), are a useful approach in this regime. These simulations show that non-linear collapse leads to virialized¹ *halos* with a near-uniform density profile, and that these halos are clustered and connected by thin filaments in a *cosmic web*. Smaller halos can merge to form larger halos, creating substructure within higher-mass halos. Λ CDM cosmology predicts the relative abundance of dark matter halos of a given mass, the *halo mass function* (Mo et al., 2010).

The only interaction included in these N-body simulations is gravity. However, gravity will pull baryons into dark matter halos, and baryons do not only interact through gravity. Indeed, these additional baryonic interactions are critical to galaxy formation inside dark matter halos (Mo et al., 2010), which is the subject of the next section.

1.1.2 Baryons within dark matter halos

Galaxies contain baryons in many forms: an evolving population of stars; the SMBHs at the centers of galaxies; and the Interstellar Medium (ISM) and Circumgalactic Medium (CGM), which consist of gas in various phases of temperature and density as well as solid dust grains. Baryons interact both with each other and with photons (Mo et al., 2010). Figure 1.2 is a flowchart illustrating the interactions between these baryonic components of galaxies.

The most visually prominent feature of galaxies is their stellar population. Stars eject energy (and momentum) as radiation and stellar winds during their lifetime, and through supernovae at the end of their lifetime. This energy helps shape the structure of galaxies. Stars also fuse hydrogen into helium and heavier elements² in their cores, ejecting both metal-enriched gas and solid dust grains in their stellar winds and supernovae. The output

¹For a system of collisionless particles (like cold dark matter) in equilibrium, the total energy $E = -K$, where K is the kinetic energy of the system.

²In astrophysics, these heavier elements are collectively referred to as ‘metals’

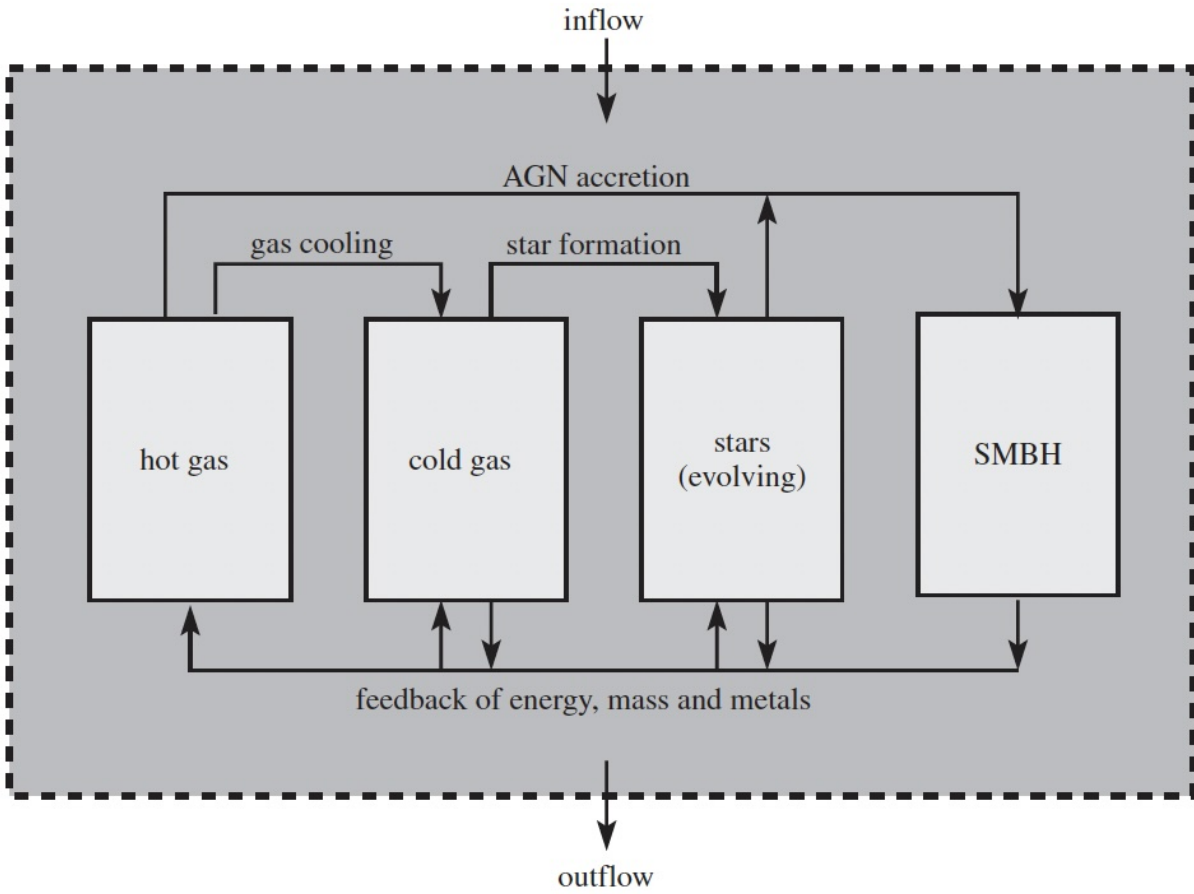


Figure 1.2: Flowchart of interactions during galaxy evolution. The dark gray box with a dashed border represents the entire galaxy, the lighter gray boxes with solid borders represent individual (baryonic) components, and the arrows represent interactions between these components. Reproduced from Mo et al. (2010).

of energy, momentum, and metals from stars are all aspects of a phenomenon called *stellar feedback* (see Figure 1.2, Mo et al., 2010).

Stars form out of the gravitational collapse of cold, dense clouds of molecular gas. The Jeans length (see equation (1.6)) depends on temperature as $\lambda_J \propto T^{1/2}$ for baryonic gas. That is, colder gas has a smaller Jeans length and so a perturbation of any given size λ will satisfy the $\lambda > \lambda_J$ condition for collapse in a sufficiently cold gas (e.g. McKee & Ostriker, 2007; Girichidis et al., 2020). Empirically, the surface density of star formation Σ_{SFR} in a galaxy has been observed to follow a *Kennicutt-Schmidt law*:

$$\Sigma_{\text{SFR}} \propto \Sigma_{\text{gas}}^N, \quad (1.7)$$

where Σ_{gas} is the surface density of (atomic or molecular) gas, and $N = 1.4 \pm 0.15$ (Kennicutt, 1998). Stars of different masses form at different rates, described by the Initial Mass Function (IMF). The IMF is an important property of galaxies because stars evolve differently depending on both their elemental composition and initial mass. So, the IMF, combined with the properties of star-forming gas, determines the evolution of the stellar population of a galaxy (e.g. Salpeter, 1955; Kroupa, 2001; Chabrier, 2003).

The ISM contains gas in multiple phases with roughly equal pressure $P \propto nT$ (Field et al., 1969), in addition to star-forming clouds dominated by molecular gas. These phases include cold neutral gas with $T \sim 100$ K and $n_{\text{H}} \sim 40 \text{ cm}^{-3}$; warm neutral gas with $T \sim 8000$ K and $n_{\text{H}} \sim 0.25 \text{ cm}^{-3}$; and diffuse, hot ionized gas with $T \gtrsim 10^{5.7}$ K and $n_{\text{H}} \sim 0.003 \text{ cm}^{-3}$. As described above, feedback from stars (as well as Active Galactic Nuclei (AGN), see below) deposits energy and momentum into the ISM. This can alter the thermal state of gas, changing the spatial structure of different ISM phases (e.g. McKee & Ostriker, 1977). So, the cooling and heating of interstellar gas is critically important to the structure and evolution of the ISM. We describe the thermal processes that contribute to this cooling and heating in section 1.2. The ISM also contains solid dust grains. Despite making up only a small fraction of its mass, dust grains are critical to the thermal evolution of the ISM (Draine, 1978).

In addition to heating the ISM, radiation from stars and AGNs can also *ionize* the gas. Recall from section 1.1.1 that, after photons and baryons decouple at $z \sim 1100$, almost all the gas in the universe is *neutral*. Thus, after the first stars and galaxies form, the radiation they emit *reionizes* the universe. This period, between about $z \sim 10$ and $z \sim 5$ is called the Epoch of Reionization (EoR, see Figure 1.1). The relative contributions of abundant faint galaxies and rarer bright galaxies to the reionization process is still an open question (e.g. Gnedin & Madau, 2022; Robertson, 2022).

The SMBHs at the centers of most galaxies are another baryonic component. These objects grow in mass by accreting surrounding gas. How the high-redshift progenitors of SMBHs form is still an unsolved problem. Gas accreting onto an SMBH forms an *accretion disk*. This system of an SMBH and accretion disk, also called an AGN, ejects radiation, energy, and momentum into the ISM (see Figure 1.2). These processes are referred to as *AGN feedback*, and can include relativistic jets emitted by the SMBH (see the review of Benson, 2010).

Baryons can also be exchanged between the galaxy itself and the surrounding CGM within the same dark matter halo. Stellar and AGN feedback (as described above) can eject material out of the galaxy. Similarly, gas from the CGM (or along filaments of the cosmic web) can accrete into a galaxy (e.g. Tumlinson et al., 2017; Faucher-Giguère & Oh, 2023).

Finally, galaxies (and their host dark matter halos) can *merge* together into a single system. This is a key part of the galaxy formation process (e.g. Press & Schechter, 1974; White & Rees, 1978). If the two galaxies have very different masses (a *minor merger*), *dynamical friction* removes energy from the smaller *satellite* galaxy and transfers it to the larger galaxy (Chandrasekhar, 1943; White, 1976). In this process, gas and stars can be removed from the satellite. The efficiency of dynamical friction increases for larger mass satellites. If the satellite is sufficiently light, its structure can persist for many orbits around the host *central* galaxy. More massive satellites will be absorbed into the central galaxy (e.g. Lacey & Cole, 1993; Navarro et al., 1995; Colpi et al., 1999; Boylan-Kolchin et al., 2008). On the other hand, if the two galaxies have similar masses (a *major merger*), then the structure of both galaxies will be disrupted. A new structure, which is usually elliptical, will be produced (e.g. Toomre & Toomre, 1972; Farouki & Shapiro, 1982; Negroponte & White, 1983; De Lucia et al., 2006). Collisions between the ISM of merging galaxies can produce shocks, heating previously cold gas and potentially spurring a burst of star formation (e.g. Joseph & Wright, 1985; Hernquist, 1989; Barnes & Hernquist, 1991).

1.2 Gas cooling and heating processes

As discussed in section 1.1.2, gas cooling and heating are vital processes in the evolution of the ISM of galaxies. In this section, we will explore the processes through which interstellar gas cools down and heats up. This section is based in large part on the Pogge (2008) lecture notes from a graduate-level course on the ISM.

The ISM is too diffuse to be in local thermodynamic equilibrium, but other kinds of equilibria where different collisional and radiative processes balance each other out can occur. Any *photoionization* process where a photon is absorbed and a free electron is released heats

the gas. If a gas particle emits a photon, this can cool the gas, but only if the gas is *optically thin* to photons of that energy so that the photon is able to escape (Pogge, 2008). In this dissertation, we will primarily be concerned with *atomic* cooling and heating processes: those that involve collisions between atoms, ions, and free electrons, as well as interactions between these particles and photons. These processes include atomic energy level transitions (section 1.2.1), ionization and recombination (section 1.2.2), and free electron processes (section 1.2.3). We can define a *cooling function* Λ and *heating function* Γ by factoring out the quadratic density dependence of the rate of radiative energy loss or gain from two-body collisional processes (e.g. Gnedin & Hollon, 2012). Different processes and ions dominate the total cooling and heating functions in different regimes. An example of the cooling function for a gas with solar element abundances and no incident radiation field from Gnat & Ferland (2012) is shown in Figure 1.3, where the quantity $\Lambda_{e,H}$ is defined by writing the rate of energy loss per unit volume as $n_e n_H \Lambda_{e,H}$, where n_e is the electron density. In this dissertation, we will adopt a slightly different cooling function definition using the overall gas number density (see equation (2.1)). In section 1.2.4, we will briefly discuss *non-atomic* cooling and heating processes.

1.2.1 Energy level transitions

Electrons bound to neutral or partially ionized atoms can be excited to higher energy levels by collisions or absorbing photons with energy equal to the gap between the initial and final states. Similarly, these electrons can be de-excited by collisions or *emitting* a photon with the same energy. The balance between these excitation and de-excitation processes determines the equilibrium population of each energy level (e.g. Bahcall & Wolf, 1968).

De-excitation photons cool the gas when the gas is sufficiently optically thin at the relevant energy for an emitted photon to escape the gas without being re-absorbed. This means that we can neglect absorption and stimulated emission of photons, and need only be concerned with spontaneous emission of photons. Hence, there is a balance between collisional excitation and *both* collisional and radiative de-excitation (e.g. Osterbrock, 1965).

Consider a simple example of a two level system with levels ‘1’ and ‘2’, colliding with free electrons. The equilibrium condition is:

$$n_e n_1 q_{12} = n_e n_2 q_{21} + n_2 A_{21}, \tag{1.8}$$

where A_{21} is the Einstein coefficient for spontaneous emission, and q_{12}, q_{21} are collisional coefficients for excitation and de-excitation. The excitation and de-excitation coefficients are related by $q_{12} = \frac{g_2}{g_1} e^{-h\nu_{21}/kT} q_{21}$, where the g_i are the degeneracies of each energy level.




Figure 1.3: Atomic gas cooling function $\Lambda_{e,H}$ (here defined as the cooling rate divided by $n_e n_H$) vs. temperature for a gas with solar element abundances and no incident radiation field, calculated with the photoionization code Cloudy (Ferland et al., 1998, see section 1.3). The shaded grey band and dotted grey line show the calculation of the total cooling rate using two different versions of Cloudy. The solid and dashed red and black lines show the contributions from various elements and their associated ionization states. Reproduced from Gnat & Ferland (2012).

The collisional de-excitation rates can be written as $q_{21} \propto T^{-1/2} \gamma_{12}(T)$, where the $\gamma_{ij}(T)$ have weak temperature dependence and are tabulated for various energy levels (Pogge, 2008).

The net cooling is the total energy released as photons: $E = n_2 A_{21} h\nu_{21}$. We can rearrange equation (1.8) to find:

$$\frac{n_2}{n_1} = \frac{q_{12}}{q_{21}} \left[1 + \frac{A_{21}}{q_{21} n_e} \right]^{-1}. \quad (1.9)$$

There is a critical density for this transition $n_{\text{crit}} = \frac{A_{21}}{q_{21}}$ (Osterbrock, 1965). For small densities $n_e \ll n_{\text{crit}}$, equation (1.9) implies that $n_2 = n_1 n_e \frac{q_{12}}{A_{21}}$, and radiative de-excitation is more prevalent than collisional de-excitation. In this regime, the cooling rate $E \propto n_1 n_e$, and depends on the density *squared*³. For large densities $n_e \gg n_{\text{crit}}$, equation (1.9) simplifies to $n_2 = n_1 \frac{q_{12}}{q_{21}} \propto n_1 e^{-h\nu_{21}/kT}$. This is just collisional, i.e. thermal, equilibrium. In this limit, the cooling $E \propto n_1$ depends only linearly on the density (Pogge, 2008). So, collisionally excited line cooling is most effective just below the critical density (Osterbrock, 1965).

Line cooling can only occur when the gas temperature is high enough for collisional excitation of the relevant energy levels. In colder $100 \text{ K} \lesssim T \lesssim 15000 \text{ K}$ gas, forbidden fine-structure lines from metals with excitation energies similar to the gas temperature (especially the $157.7 \mu\text{m}$ line of ionized carbon CII) are important coolants. At higher temperatures $20000 \text{ K} \lesssim T \lesssim 40000 \text{ K}$, neutral hydrogen (HI) and neutral helium (HeI) can be collisionally excited above the ground state, and their line emission is a dominant coolant. At $T \gtrsim 60000 \text{ K}$, non-forbidden transitions of metals (especially carbon and oxygen) can contribute to collisionally-excited line cooling as well (see Figure 1.3, Pogge, 2008).

1.2.2 Ionization and recombination

Similarly to section 1.2.1, there is a balance between collision and radiative ionization and recombination processes. Neutral and partially ionized atoms can be ionized either by collisions, or by absorbing an ultraviolet (UV) or higher-energy photon with energy above the ionization threshold. The 13.6 eV ionization potential of hydrogen, the most abundant element in interstellar gas, corresponds to a temperature of $T \sim 10^5 \text{ K}$. For this reason, collisional ionization is only relevant in very hot gas. In colder gas, photoionization must be the dominant source of ionization (Pogge, 2008).

Photoionization heats the gas by an amount $E = h\nu - \chi$, where χ is the ionization potential of the electron that was released. This is just the kinetic energy of the freed electron (Dalgarno & McCray, 1972). Note that higher energy photons heat the gas more

³Note that $n_e \approx n_{\text{H}}$ because hydrogen is the most abundant element in the gas. Electrons from helium and metal ions introduce only a small correction factor.

than photons that are close to the ionization threshold. So, the efficiency of photoionization heating is very sensitive to the energy spectrum of the incident radiation.

Recombination is a collisional process involving a free electron and a positively charged ion. The cross section for recombination depends on electron velocity as $\sigma \propto v^{-1/2}$. The kinetic energy of any recombined electrons is removed from the gas, so recombination is generically a cooling process. However, because electrons with lower kinetic energy $\frac{1}{2}mv^2$ are more likely to undergo recombination than electrons with higher kinetic energy, this process can actually increase the gas temperature (Pogge, 2008).

Each recombination also emits a photon with energy equal to the difference between the kinetic energy of the recombining electron and the energy level the electron recombines into. If the electron recombines into the ground state, the emitted photon will quickly photoionize another atom and so cannot escape to cool the gas. However, if the electron recombines to an excited state, it can then “cascade” back to the ground state through multiple transitions. The photons in this cascade can escape, leading to cooling. This cooling is important at $15000 \text{ K} \lesssim T \lesssim 20000 \text{ K}$, where the lower limit roughly corresponds to the 13.6 eV threshold for hydrogen ionization and the upper limit is the temperature above which collisionally-excited line cooling of neutral hydrogen and helium are dominant (Pogge, 2008).

1.2.3 Free electron cooling

Free electrons can also emit cooling radiation through interactions that do not involve recombination. The primary cooling mechanism of this type is free-free emission, also known as *bremstrahlung*. In this process, a free electron is accelerated due to an electromagnetic interaction with a charged ion and emits electromagnetic waves. The cooling rate \dot{E} due to this process follows $\dot{E} \propto n^2 T^{1/2}$, and is the primary coolant in very hot gas with $T \gtrsim 10^7 \text{ K}$ (Draine, 2011). This can be seen in the increase in the overall cooling rate at high temperatures in Figure 1.3 (primarily due to interactions with H and He ions).

Interactions between free electrons and interstellar magnetic fields lead to *synchrotron* emission. However, this emission is primarily at radio frequencies $\nu \lesssim 1 \text{ GHz}$, and so has low enough energy that it is not a significant coolant (Draine, 2011).

Free electrons can Compton scatter off of CMB photons, a process that cools the gas when the gas temperature is higher than the CMB temperature $T_{\text{CMB}} = T_0(1+z)^2$, where $T_0 = 2.7 \text{ K}$. Note that since this process only involves *one* gas particle, the cooling rate from Compton scattering scales linearly with the gas density (the collisional processes discussed above and in section 1.2.1 and section 1.2.2 all scale with the density squared). Compton scattering occurs more frequently at higher redshifts when the photon density of the CMB

is higher. For $z \lesssim 6$, the characteristic timescale for Compton cooling becomes longer than the age of the universe (Benson, 2010).

1.2.4 Non-atomic processes

In addition to the atomic cooling and heating processes which are the main focus of this dissertation, there are also cooling and heating processes due to other components of the ISM. A brief overview of some of these processes is presented below:

- **Molecular cooling:** Rotational and vibrational line emission from interstellar molecules can be collisionally excited, similarly to the analysis of atomic line emission in section 1.2.1. Due to the low excitation energies involved, molecular cooling is particularly important in very cold regions of the ISM, such as molecular clouds (e.g. Dalgarno & McCray, 1972; Scoville & Solomon, 1974; McKee et al., 1982; Neufeld & Kaufman, 1993).
- **Dust photoheating:** Similarly to the photoionization heating discussed in section 1.2.2, photons absorbed by solid dust grains in the ISM can eject an electron, heating the ISM. This process occurs more frequently for smaller dust grains. This turns out to be the most important heating process in the diffuse ISM (e.g. Jura, 1976; de Jong, 1977; Draine, 1978).
- **Dust emission:** Dust grains can also emit cooling radiation. This includes both thermal emission in the infrared and line emission from small grains that have recently been excited by absorbing a photon (Draine, 2004).
- **Cosmic ray heating:** High-energy ultrarelativistic particles can ionize atoms and heat the ISM. As described in section 1.2.2, the heating from ionization is dependent on the difference between the incident cosmic ray energy and the ionization threshold (e.g. Field et al., 1969; Goldsmith et al., 1969; Habing & Goldsmith, 1971). Estimates suggest that the heating rate from cosmic rays in the diffuse ISM is significantly smaller than that from dust photoheating (Pogge, 2008; Draine, 2011).

Given this census of cooling and heating processes in the ISM, we now turn to the question of how we can model the actual thermal behavior of ISM gas with particular gas properties and incident radiation field.

1.3 Photoionization codes

The net cooling or heating from the processes described in section 1.2 can be computed with a *photoionization code* that finds the emission spectrum of interstellar gas. This is accomplished by solving for the equilibrium populations of ionization states of each element through balancing recombination processes with both collisional ionization and photoionization (see section 1.2.2). The equilibrium temperature can be found by balancing all the relevant cooling and heating processes (Ercolano, 2005). In this dissertation, we use calculations from the photoionization code Cloudy (specifically `CLOUDY90`, Ferland et al., 1998).

Running photoionization codes requires inputting the properties of the gas itself (its initial temperature and density, as well as the abundances of the various elements), and the energy spectrum of any incident ionization sources (photons and cosmic rays). Early photoionization codes considered spherically symmetric gas clouds around an ionizing source, such as an HII region of ionized hydrogen around a hot, young star or a planetary nebula. Advances in computing have enabled 3D photoionization codes that consider multiple ionizing sources with nearly arbitrary geometry relative to the gas (see the review of Ercolano, 2005).

A key element of any photoionization code is the database of atomic properties used. Important properties include the temperature-dependence of collisional excitation rates $\gamma_{ij}(T)$ discussed in section 1.2.1, photoionization cross sections and recombination rates to various energy levels, and Einstein coefficients for radiative transitions. As described in section 1.2.4, dust grains can contribute significantly to the heating and infrared emission of the ISM. Photoionization codes differ in how realistically they treat dust (Ercolano, 2005).

Photoionization codes computationally model the thermal and ionization equilibrium of an interstellar gas cloud on the scale of interactions between atoms, molecules, dust grains, and ionizing photons in a given geometry. However, this is *many* orders of magnitude smaller than any individual galaxy. In the next section, we turn to the problem of simulating the formation and evolution of galaxies.

1.4 Hydrodynamic simulations

Galaxy formation and evolution involve the interplay of many different physical processes (see Figure 1.2 and section 1.1). In this section, we provide a brief overview of computational simulations of galaxy evolution. Galaxy formation simulations are often called *hydrodynamic simulations* because the baryonic gas is a critical component (Vogelsberger et al., 2020). We begin with a discussion of simulations which include only dark matter in section 1.4.1, then describe how the baryonic gas can be incorporated in section 1.4.2. In section 1.4.3, we

consider how to incorporate processes on scales too small to be spatially resolved. Finally, we introduce two specific types of hydrodynamic simulations that are utilized in this dissertation: radiative transfer simulations of the EoR (section 1.4.4) and idealized simulations of individual galaxies (section 1.4.5). In this dissertation we use the Adaptive Mesh Refinement (AMR) hydrodynamic simulation code Adaptive Refinement Tree (ART) (Kravtsov, 1999; Kravtsov et al., 2002; Rudd et al., 2008).

1.4.1 Gravity and dark matter

In Λ CDM cosmology, dark matter interacts only through gravity. So, the only ingredients we need to simulate the formation of dark matter halos (as described in section 1.1.1) are dark matter and gravity. N-body simulations model dark matter with particles of fixed mass that interact through the force of gravity. The mass of these dark matter particles is known as the *mass resolution* of the simulation (e.g. the review of N-body simulation techniques Angulo & Hahn, 2022). The ART simulations used in this dissertation use an N-body code also called ART (Kravtsov, 1999; Kravtsov et al., 2002; Rudd et al., 2008). We describe the ART N-body solver (Kravtsov et al., 1997) below.

We can only simulate a finite volume, but we do not want the gravitational field to have a sharp cutoff at the edge of this volume. We also do not want dark matter particles to disappear if their momentum takes them outside of the volume. This problem is solved by using *periodic boundary conditions*. Each side of the simulated box is equivalent to the opposite side. For example, a particle that exits the right side of the simulation immediately re-enters the left side. Put another way, the simulated universe consists of infinite repeated copies of the simulated volume (Kravtsov et al., 1997).

One difficulty that arises with simulating gravitational interactions is that the force of gravity $F \propto \frac{1}{r^2}$, where r is the distance between particles, diverges at small distances. A common solution to this is adding a *softening length* ϵ . The dependence of force on separation is re-defined as $F \propto \frac{1}{r^2 + \epsilon^2}$. For $r \gg \epsilon$, the softening has essentially no effect. However, for separations $r \lesssim \epsilon$, the softening makes the force anomalously small. So, the softening length sets the effective *spatial resolution* of an N-body simulation (Angulo & Hahn, 2022).

Calculating the gravitational force on every dark matter particle requires considering the gravitational interaction between every pair of particles. This scales as N^2 with the number of particles N , making it difficult to run higher-resolution simulations with more dark matter particles. Rather than finding the force on each particle, ART proceeds by solving the Poisson equation for the gravitational potential ϕ :

$$\nabla^2 \phi = 4\pi G a^2 (\rho - \bar{\rho}), \quad (1.10)$$

where G is Newton’s gravitational constant. The force on a particle of mass m due to this potential is $F = -am\nabla\phi$ (Kravtsov et al., 1997).

There are several approaches to solving equation (1.10) in a box with period boundary conditions used in modern N-body simulations. The ART simulations utilized in this dissertation use a technique called *particle mesh* (Kravtsov et al., 1997). Particle mesh codes discretize the gravitational potential ϕ on a 3D grid inside the simulation box. This discretization truncates the Fourier transform of ϕ at a wavenumber k_{\max} corresponding to the side length of the mesh grid. So, the Fourier series can be computed numerically with the Fast Fourier Transform (FFT) algorithm. In Fourier space, the Poisson equation and the equation for the force due to the potential ϕ are no longer differential equations and can be directly solved. Applying the inverse FFT gives the gravitational potential ϕ and force in real space (Angulo & Hahn, 2022).

These methods allow for simulations of the growth of dark matter halos and the cosmic web. However, to simulate galaxies within those halos, the baryonic gas must also be included (see section 1.1.2). This is the topic of the next section.

1.4.2 Incorporating the gas: hydrodynamics

Baryonic gas interacts through gravity with both dark matter and itself. However, as described in section 1.1, the gas is also collisional and is subject to pressure. To include this gas in *hydrodynamic simulations* of galaxy evolution, we need dark matter particles modeled as described in section 1.4.1, and the gas. The gas is generally modeled as a fluid with density ρ , velocity field \vec{v} , and energy per unit mass $e = u + \frac{1}{2}v^2$, where u is the thermal energy per unit mass and $v = |\vec{v}|$ is the speed. In addition to Poisson’s equation (still equation (1.10), but now with the density ρ including both dark matter particles and the gas density), there are additional differential equations describing the evolution the gas density, velocity, and energy:

$$\frac{\partial\rho}{\partial t} + \nabla \cdot (\rho\vec{v}) = 0, \quad (1.11)$$

$$\frac{\partial(\rho\vec{v})}{\partial t} + \nabla \cdot (\rho\vec{v} \otimes \vec{v} + P\mathbb{I}) = 0, \quad (1.12)$$

$$\frac{\partial(\rho e)}{\partial t} + \nabla \cdot (\rho e + P)\vec{v} = 0, \quad (1.13)$$

where P is the (isotropic) fluid pressure and \mathbb{I} is the identity tensor (see the technical review of galaxy formation simulations Vogelsberger et al., 2020). Equation (1.11) for the density ρ is known as the continuity equation, and equation (1.12) for the velocity \vec{v} is Euler’s equation (Mo et al., 2010).

There are two common approaches to solving equations (1.11-1.13): (1) Lagrangian Smoothed Particle Hydrodynamics (SPH) simulations, and (2) Eulerian AMR simulations (Vogelsberger et al., 2020). Figure 1.4 shows snapshots from various cosmological volume hydrodynamic simulations, including both SPH and AMR approaches. ART is an AMR code (Kravtsov, 1999; Kravtsov et al., 2002; Rudd et al., 2008).

The main difference between AMR and SPH simulations is how the gas fluid is divided into coarse-grained resolution elements. Each such element is treated as having constant properties over the volume of the resolution element. In an SPH simulation, the gas is modeled as a set of *particles* that sample the motion of the gas. That is, the particles advect with the fluid flow. This is why SPH simulations are also referred to as Lagrangian. These particles have constant mass, momentum, and energy. A region of the simulation with denser gas will have more SPH particles as compared to a region with less dense gas. This approach is very useful for any scenario where the bulk flow of gas is of interest, such as outflow from a galaxy or gas accretion along cosmic web filaments (Vogelsberger et al., 2020).

In AMR simulations (like the ART simulations used in this dissertation), the gas is discretized on a mesh grid similarly to the particle mesh N-body method described in section 1.4.1. On this grid, the differential equations (1.11-1.13) become *finite difference* equations. To conserve mass, momentum, and energy, we need to solve for the flux of each across every interface between grid cells (Vogelsberger et al., 2020). Even within the ISM of a single galaxy, the range of gas densities covers several orders of magnitude. Dense, star-forming molecular clouds have densities as high as $n_{\text{H}} \gtrsim 1000 \text{ cm}^{-3}$ while hot, diffuse ionized gas has low densities $n_{\text{H}} \sim 0.004 \text{ cm}^{-3}$ (Draine, 2011). A cosmological simulation will also include the CGM of galaxies and intergalactic gas, which both reach even lower densities. For this reason, the gas mesh grid cells in an AMR simulation are generally *adaptive*, with smaller cells in regions with denser gas. This is usually accomplished by starting with a fixed-size mesh grid, and refining the mesh once a cell becomes denser than some threshold value (Vogelsberger et al., 2020).

Note that the gas pressure P is included in the equations for the evolution of the velocity (equation (1.12)) and energy (equation (1.13)). Both the pressure P and the energy density e depend on the temperature and density of the gas. So, to solve equations (1.11-1.13) in either an AMR or SPH simulation, we need to evaluate the cooling and heating rates of the gas to determine its temperature, density, and pressure. That is, cooling and heating function calculations are critical components of a hydrodynamic simulation (Vogelsberger et al., 2020).

AMR and SPH methods describe the behavior of baryonic gas. However, we saw in section 1.1.2 that there are other important baryonic components of galaxies which are not




Figure 1.4: Snapshots from various hydrodynamic cosmological simulations: Illustris (Vogelsberger et al., 2014), IllustrisTNG (Springel et al., 2018), Magneticum (Bocquet et al., 2016), EAGLE (Schaye et al., 2015), Romulus25 (Tremmel et al., 2017), SIMBA (Davé et al., 2019), Massiveblack-II (Khandai et al., 2015), and Horizon-AGN (Dubois et al., 2014). Reproduced from Vogelsberger et al. (2020).

accounted for in equations (1.10-1.13). These include stars and SMBHs. We discuss how to incorporate these components in the next section.

1.4.3 Subgrid physics models

Most of the baryonic processes in galaxy evolution (see section 1.1.2) occur on spatial scales that are too small to resolve in galaxy evolution simulations. That is, the spatial scales involved are smaller than the most refined cell in an AMR simulation or the mass involved is smaller than the gas particle mass in an SPH simulation. This is true for both cosmological volumes and isolated galaxy simulations (described in section 1.4.5 below).

To model such processes in a hydrodynamic simulation, we need *subgrid models*. These are approximations using physics-informed parameters available in the simulation. They can be empirically informed based on either observational data or high-resolution small-scale modeling of the processes in question (Vogelsberger et al., 2020). Note that a variety of subgrid modeling choices are implemented in various simulations based on the ART hydrodynamics code (e.g. Rudd et al., 2008; Gnedin, 2014; Semenov et al., 2021). We list examples of some common subgrid models used in state-of-the-art hydrodynamical simulations below:

- **Star formation:** The molecular clouds where stars form are usually $\lesssim 10$ pc in size (Mo et al., 2010). Both these clouds and individual stars are too small to be resolved in most galaxy formation simulations. Instead, one approach is to form *star particles* describing a population of stars characterized by some IMF following an empirical Kennicutt-Schmidt law (equation 1.7) in sufficiently cold and dense gas (Vogelsberger et al., 2020).
- **Stellar feedback:** The evolution of a star particle is dictated by the IMF, along with the element abundances of the gas in which the star particle formed. The most massive stars will die as supernovae. Stellar winds and supernovae eject momentum, energy, and gas enriched with metals into the surrounding gas, potentially leading to an outflow of gas from the galaxy (see Figure 1.2 and section 1.1.2). While the overall amount of metals, energy, and momentum deposited can be calculated with a stellar evolution model, the individual stars and supernovae cannot be resolved. So, it is not obvious how the metals, energy, and momentum should be distributed into the surrounding gas. Energy can be deposited into the gas either as kinetic energy or radiation. Studies that deposit the energy as radiation often pause radiative cooling (see section 1.2) in the affected gas for $\sim 10^7$ years to prevent the gas from cooling too quickly. Kinetic feedback does not require this ad hoc procedure, but can sometimes

be insufficient to drive galactic outflows without artificially decoupling the affected gas from the gravity and hydrodynamic pressure of nearby gas. A more empirical and approximate approach to stellar feedback is incorporating a *mass loading factor* η so that the mass flowing out of a galaxy is proportional to the star-formation rate: $\dot{M}_{\text{out}} = \eta \dot{M}_*$. The free parameter η can be tuned to match observations (Vogelsberger et al., 2020).

- **SMBH accretion:** SMBHs in simulations are usually treated by placing a *seed* black hole in all galaxies above a given mass threshold, and allowing the mass of the seed to grow by accreting gas at a rate given by:

$$\dot{M}_{\text{SMBH}} = \frac{4\pi G M_{\text{SMBH}}^2 \rho_{\text{gas}}}{(c_s^2 + v_{\text{rel}}^2)^{3/2}}, \quad (1.14)$$

where M_{SMBH} is the black hole mass, ρ_{gas} is the gas density, c_s is the adiabatic sound speed (see equation (1.5)), and v_{rel} is the relative velocity between the black hole and the gas (Vogelsberger et al., 2020).

- **AGN feedback:** SMBHs also eject gas, momentum, and energy into the surrounding gas via AGN feedback (see section 1.1.2). Two different types of AGN feedback are commonly incorporated in hydrodynamic simulations as subgrid models: (1) adding energy and momentum to gas near the black hole at a rate proportional to the accretion rate as calculated in equation (1.14), and (2) relativistic jets that extend far from the SMBH when the accretion rate is below some threshold (Vogelsberger et al., 2020).
- **ISM turbulence:** If there is turbulence in the ISM, the turbulent cascade will convert flows on resolved scales to smaller flows on unresolved scales. This will continue until the turbulent energy is dissipated as heat. So, there are both resolved *and* unresolved components to the turbulent energy density. A subgrid model which conserves the overall energy is needed for the unresolved component (e.g. Schmidt et al., 2014; Semenov et al., 2016).
- **Multiphase ISM:** The ISM of galaxies contains gas in multiple temperature-density phases (see Figure 1.2 and section 1.1.2). High-resolution simulations of individual galaxies (see section 1.4.5 below) can have sufficiently high spatial resolution to separate different ISM phases, but cosmological volume simulations typically will not. In general, cold dense gas occupies a significant fraction of the ISM mass but a small fraction of the ISM volume. This is because it has a high density compared to more diffuse phases. Other complications to simulating a multiphase ISM include the many

radiation-matter interactions described in section 1.2, as well as turbulence. One approach is to assume that the gas fluid in the simulation represents a hotter, more diffuse phase and add a cold phase with an *effective equation of state* $T \propto \rho^\gamma$, where γ is the usual gas adiabatic index in thermodynamics (Vogelsberger et al., 2020).

Subgrid models generally include free parameters that must be either chosen based on heuristic physical considerations or calibrated to match observational constraints. Calibration is a difficult problem because the number of simulations that can be run is limited by the available computational resources (Vogelsberger et al., 2020). It is also challenging to compare results from simulations which use different subgrid models. Several collaborations have emerged to use common subgrid models and parameters across various simulation groups, including Aquila (Scannapieco et al., 2012), AGORA (Kim et al., 2014), and CAMELS (Villaescusa-Navarro et al., 2021).

After this overview of how hydrodynamic galaxy evolution simulations work, we turn to two particular types of such simulations that we use in this dissertation.

1.4.4 Epoch of Reionization simulations and radiative transfer

To simulate the EoR, we need to follow the ionization of intergalactic gas. This requires a large cosmological volume containing many galaxies. In addition to the subgrid models listed in section 1.4.3, simulations of the EoR also need *radiative transfer* of ionizing radiation. That is, given the flux of ionizing radiation emitted by star particles and SMBHs, the simulations must follow how this radiation propagates through the gas. This is difficult because the relevant absorption and scattering cross sections are functions of photon frequency. The high speed of light also means that short timesteps are needed to limit the distance a photon can propagate (Vogelsberger et al., 2020). Simulations sometimes use a *reduced speed of light approximation*, where a lower value \bar{c} for the speed of light is used, so that the non-relativistic limit $v/\bar{c} \ll 1$ is still satisfied for the gas speed v (Gnedin & Abel, 2001). Because of the computational expense of radiative transfer, simulations of the EoR are usually only run to the end of the EoR at $z \sim 5$ (Vogelsberger et al., 2020).

In Chapter 2, we analyze data from the Cosmic Reionization on Computers (CROC) project simulations of the EoR. These simulations are based on the ART hydrodynamic AMR code and end at $z = 5$. Figure 1.5 shows the neutral hydrogen fraction in slices through a CROC simulation at redshift $z = 8.04$ and $z = 7.01$. At $z = 8.04$, individual bubbles of ionized hydrogen are visible. By $z = 7.01$, these ionized bubbles have begun to merge together, creating a Swiss cheese-like topology (Gnedin & Kaurov, 2014).

The CROC project uses an Optically Thin Variable Eddington Tensor method (Gnedin &




Figure 1.5: Slices through one realization of the CROC (Gnedin, 2014) simulations of the EoR showing the neutral hydrogen fraction at $z = 8.04$ (top panel) and $z = 7.01$ (bottom panel). Reproduced from Gnedin & Kaurov (2014).

Abel, 2001) to implement radiative transfer (Gnedin, 2014; Gnedin & Kaurov, 2014). This approach uses the zeroth and first moments of the radiative transfer equation, which give differential equations for the zeroth and first moments of radiation field. These moments are, respectively, the energy density and flux of the radiation field (Gnedin & Abel, 2001).

On the scale of an AMR cell, the radiative transfer equation can be written as:

$$\frac{a}{c} \frac{\partial f_\nu}{\partial t} + n^i \frac{\partial f_\nu}{\partial x^i} = -\hat{\kappa}_\nu f_\nu + \psi_\nu, \quad (1.15)$$

where $f_\nu \equiv J_\nu/\bar{J}_\nu$ is the ratio of the radiation field specific intensity J_ν to its spatial and angular average \bar{J}_ν ; \vec{n} is a unit vector pointing in the direction the photon propagates; $\psi_\nu \equiv \frac{a}{c} S_\nu/\bar{J}_\nu$, where S_ν is the source term; and $\hat{\kappa}_\nu \equiv \frac{a}{c} (k_\nu - \bar{k}_\nu) + (\bar{S}_\nu/\bar{J}_\nu)$, where k_ν is the absorption coefficient. Note that equation (1.15) uses the Einstein summation convention where the repeated index i indicates that this term is summed over all spatial directions $i = \{x, y, z\}$ (Gnedin & Abel, 2001).

Taking the zeroth moment of equation (1.15) gives:

$$\frac{a}{c} \frac{\partial E_\nu}{\partial t} + \frac{\partial F_\nu^i}{\partial x^i} = -\hat{\kappa}_\nu E_\nu + \psi_\nu, \quad (1.16)$$

and taking the first moment of equation (1.15) gives:

$$\frac{a}{c} \frac{\partial F_\nu^j}{\partial t} + \frac{\partial}{\partial x^i} E_\nu h_\nu^{ij} = -\hat{\kappa}_\nu F_\nu^j, \quad (1.17)$$

where the relative energy density E_ν is the zeroth moment of f_ν is:

$$E_\nu(t, \vec{x}) \equiv \frac{1}{4\pi} \int f_\nu(t, \vec{x}, \vec{n}) d\Omega, \quad (1.18)$$

the flux F_ν^i is the first moment of f_ν is:

$$F_\nu^i(t, \vec{x}) \equiv \frac{1}{4\pi} \int n^i f_\nu(t, \vec{x}, \vec{n}) d\Omega, \quad (1.19)$$

and the Eddington tensor h_ν^{ij} is derived from the second moment of f_ν through (Gnedin & Abel, 2001):

$$E_\nu(t, \vec{x}) h_\nu^{ij}(t, \vec{x}) \equiv \frac{1}{4\pi} \int n^i n^j f_\nu(t, \vec{x}, \vec{n}) d\Omega. \quad (1.20)$$

Equation (1.16) and equation (1.17) describe the evolution of the relative energy density E_ν and flux F_ν^i in terms of the Eddington tensor h_ν^{ij} , which must be specified to solve these equations. In the Optically Thin Variable Eddington Tensor approach, h_ν^{ij} is calculated using

equation (1.20) under the assumption of optically thin gas so that no photons are absorbed and h_ν^{ij} is independent of the frequency ν (Gnedin & Abel, 2001). Similar radiative transfer methods are implemented in cutting-edge EoR simulations such as Cosmic Dawn (Ocvirk et al., 2016) and THESAN (Kannan et al., 2022).

1.4.5 Zoom-in and isolated galaxy simulations

Simulations of the EoR need a large volume to track how ionizing radiation from galaxies ionizes intergalactic gas. However, this large volume limits the resolution in any of the simulated galaxies. In this section, we discuss how to achieve higher resolution in individual galaxies. This is crucial to studying galaxy formation and evolution in more detail than is possible in cosmological volumes (e.g. Guedes et al., 2011; Wang et al., 2015; Sawala et al., 2016; Wetzel et al., 2016; Grand et al., 2017).

To achieve higher spatial resolution in a galaxy, we need higher resolution initial conditions in the region where the galaxy will form. Everywhere else, lower resolution than for a typical cosmological volume simulation is acceptable. For that reason, these types of simulations are sometimes referred to as *zoom-in* simulations, since they are essentially re-simulations of a small volume at higher resolution. Note that this usually requires first running a cosmological volume simulation to determine where the regions of interest in the initial conditions are. This initial simulation can be of relatively low resolution (Vogelsberger et al., 2020).

The computational expense of a zoom simulation scales with the mass of the galaxy being simulated and with the desired resolution inside the galaxy. If the galaxy is large and accretes a significant fraction of its mass through a major merger, the initial condition volume we need to sample at high(er) resolution could be quite large. For this reason, the smallest *dwarf galaxies* are the easiest to model with zoom-in simulations. Dwarf galaxy simulations are very useful for calibrating subgrid models as described in section 1.4.3 (Vogelsberger et al., 2020). In Chapter 5, we run and compare two versions of an ART dwarf galaxy simulation from Semenov et al. (2021) utilizing two different methods for computing the gas cooling and heating rates. Gas properties in a slice through the disk of the fiducial galaxy simulation from Semenov et al. (2021) are shown in Figure 1.6.

1.5 Machine learning

The last main computational technique we use in this dissertation is *machine learning*. In the past few decades, the size of astrophysical datasets available for analysis has increased dramatically. This has spurred a wave of research using machine learning techniques to




Figure 1.6: Slices through the disk of the Semenov et al. (2021) idealized galaxy simulation showing the gas number density (leftmost panel), temperature (middle left panel), turbulent velocity dispersion (middle right panel), and UV radiation field strength (rightmost panel). Reproduced from Semenov et al. (2021).

analyze large data sets and extract patterns. The number of published astrophysics papers per year referencing machine learning in the title, keywords, and/or abstract has risen from ~ 10 in the year 2000 to over 400 by the year 2020 (Rodríguez et al., 2022).

In Chapter 3 and Chapter 4, we use machine learning to *interpolate* the dependence of the base-10 logarithm of the atomic gas cooling and heating functions $\log \Lambda$ and $\log \Gamma$ (see section 1.2) for a gas with given element abundances using the gas temperature T , gas density n_{H} , and some set of frequency-averages of the incident radiation field intensity J_{ν} . We can calculate ‘true’ cooling and heating functions using the photoionization code Cloudy (Ferland et al., 1998, see section 1.3). These calculations are the training data for our machine learning models. The cooling and heating functions can depend on a large number of factors, such as the abundances of important metal-line coolants and the intensity of the radiation that can ionize relevant atoms and ions.

We discuss the general framework of applying machine learning to interpolation problems in section 1.5.1 and the specific machine learning algorithm we use in 1.5.2. Finally, we introduce the tool we use to *interpret* the predictions of our machine learning models in section 1.5.3.

1.5.1 Machine learning for interpolation

As described in section 1.4.2, hydrodynamic galaxy formation simulations require computing the cooling and heating functions of the gas resolution elements. These can be computed with the photoionization codes discussed in section 1.3. However, these codes are computationally expensive. A common practice is to construct an *interpolation table* of cooling and heating

functions calculated from a photoionization code such as Cloudy on a grid of gas and radiation field properties. When cooling and heating function values are needed at intermediate points not included in the table, the values can be interpolated using the adjacent grid points (e.g. Wiersma et al., 2009; Gnat & Ferland, 2012; Gnedin & Hollon, 2012; Ploeckinger & Schaye, 2020).

However, because cooling and heating functions are non-linear functions of the gas and radiation field properties, these interpolation table approximations will inevitably result in large errors (compared to calculations from the photoionization code used to construct the table) at some points in parameter space (e.g. Wiersma et al., 2009; Gnedin & Hollon, 2012). Cooling and heating calculations that include all the relevant processes (as described in section 1.2) are known to be important for capturing the thermal properties of the ISM (e.g. Kim et al., 2023). The need for accurate cooling and heating function interpolation in galaxy formation simulations makes cooling and heating function interpolation a good use case for applying machine learning interpolation techniques, with the aim of improving on the accuracy of existing interpolation table approaches.

Cooling and heating functions are *continuous* functions of gas properties and the incident radiation field. So, predicting them with machine learning is a general case of a *regression* or interpolation problem: fitting a continuous function given training data of the true values of the function on a set of input parameters called *features*. Because the training data is labeled with the true values of the function, regression models are *supervised learning* algorithms (Sarker, 2021).

The training data for fitting a multivariate function like atomic gas cooling and heating rates is essentially *tabular*: each row is one instance of the input features. Each column represents a different input feature, with one additional column for the true value of the output. The rows of this table can be swapped freely, but swapping the columns changes the meaning of the table. Studies have shown that the best-performing class of machine learning algorithm for regression problems with this type of tabular training data are *gradient-boosted tree* algorithms (Shwartz-Ziv & Armon, 2022; Grinsztajn et al., 2022), which we introduce in the next section.

1.5.2 Gradient-boosted tree algorithms

Gradient-boosted tree algorithms (Friedman, 2001) are based on regression trees. Each node of a tree is split into two nodes the next level down based on the value of one feature x_i compared to a threshold y . The nodes at the bottom level of the tree are called *leaves*. Each leaf is assigned a value. Any data point (either a labeled point in the training or test data or

an unlabeled point) will be assigned to exactly one leaf of the tree after considering the value of the relevant feature at each node. The value on this leaf is the prediction of the regression tree. The threshold values y are chosen (i.e. *trained*) by minimizing a *loss function* that is the sum of the Mean Squared Error (MSE) and a *regularization* term that penalizes more complicated models. The MSE is defined as (Chen & Guestrin, 2016):

$$\text{MSE} = \frac{1}{N} \sum_i (y_i - y_{\text{pred}}(x_i))^2, \quad (1.21)$$

where x_i are the training data features, y_i are the associated labels, N is the number of training data points, and $y_{\text{pred}}(x_i)$ is the prediction of the tree on the training point x_i . This loss function is minimized using its gradient (Chen & Guestrin, 2016).

A single regression tree would not be a very accurate regression model. Instead, gradient-boosted tree models use an *ensemble* of trees. The first tree is trained to predict the training data labels. Each subsequent tree is trained to predict the *residual* between the training labels and the prediction made by all the previous trees. The prediction of the ensemble is simply the sum of the predictions made by each tree (Chen & Guestrin, 2016). This is illustrated schematically in Figure 1.7. In Chapter 3 and Chapter 4 we train models using the gradient-boosted tree algorithm eXtreme Gradient Boosting (XGBoost) implemented with the package `xgboost` (Chen & Guestrin, 2016).

The computational complexity of (training and evaluating) gradient-boosted tree models scales very well with the number of input features, since each node of any tree only considers one feature value. Furthermore, XGBoost contains a *hyperparameter*⁴ which sets the fraction of the input features that are considered in when constructing each tree (Chen & Guestrin, 2016).

While it is technically possible to visualize any regression tree as a graph to understand how the predictions depend on the inputs, this quickly becomes infeasible for trained XGBoost models with more than a handful of trees, or a maximum tree depth⁵ larger than a few⁶. XGBoost models also interface easily with methods to assess the relative importance of the model features, which are discussed in the following section.

⁴A parameter that is manually chosen prior to model training, rather than being optimized in the training process (Chen & Guestrin, 2016).

⁵The number of nodes traversed until reaching a leaf (Chen & Guestrin, 2016).

⁶Both the number of trees and maximum tree depth are hyperparameters of XGBoost models (Chen & Guestrin, 2016).




Figure 1.7: A schematic picture of XGBoost regression models. Reproduced from Wang et al. (2019).

1.5.3 Feature importance analysis with SHAP values

Machine learning models often function as so-called ‘black boxes’. In other words, it is unclear how changing a given model input will affect the prediction of the model. An antidote to this is provided by *interpretability* methods, which provide a quantitative measure of the importance of each input feature. In Chapter 3 and Chapter 4 of this dissertation, we make use of one particular interpretability tool: SHapley Additive exPlanation (SHAP) values, which are an approximation of Shapley regression values (Lundberg & Lee, 2017). Calculation of SHAP values for XGBoost models is implemented in the Python package `shap` (Lundberg et al., 2018; Lundberg et al., 2020), making them particularly useful for this work.

Shapley regression values can be computed for each feature for any given model prediction. They describe the difference between the actual prediction of the model and what the model would predict if it ignored the feature in question. The Shapley regression value ϕ_i for a feature i is given by (Lundberg & Lee, 2017):

$$\phi_i = \sum_{S \subseteq F \setminus i} \frac{|S|!(|F| - |S| - 1)!}{|F|!} [f_{S \cup \{i\}}(x_{S \cup \{i\}}) - f_S(x_S)], \quad (1.22)$$

where F is the set of all features and $|F|$ is the total number of features. The sum is over




Figure 1.8: A schematic depiction of the SHAP values (see equation (1.22)) for a machine learning model with 4 features. The ϕ_i shown in this figure should be averaged over all possible orderings of the 4 features to obtain the SHAP values. Reproduced from Lundberg & Lee (2017).

any subset of features S which does *not* include the feature i . The model f_S is trained using only the features in the subset S , and $|S|$ is the number of features in S . Equation (1.22) is just a weighted sum of the difference between model predictions with and without the feature i in question (Lundberg & Lee, 2017).

Computing equation (1.22) requires training models using every possible subset of the available features. In realistic scenarios, this is not possible. SHAP values approximate equation (1.22) by averaging over predictions on the training data rather than training additional models (Lundberg & Lee, 2017).

SHAP values have the advantage that the sign and numerical value are both meaningful: a SHAP value is the deviation of the model prediction from a naive expectation ignoring the value of the feature in question. Furthermore, the SHAP values for each feature at a given data point always add to the difference between the model prediction and the expected value by ignoring *all* features. This is just the median prediction on the training set (Lundberg & Lee, 2017). Figure 1.8 illustrates SHAP values for a model with 4 features.

However, because SHAP values are only defined for individual model predictions, we have to average their *absolute value* over some sample of points to assess the overall importance of each feature. Since SHAP values are relatively computationally expensive to compute, we only do this on a subset of model predictions (see Chapter 3 and Chapter 4 for more details).

It is also worth noting that SHAP values are based solely on the *predictions* of the model and ignore how accurate these predictions are (Lundberg & Lee, 2017). This means that if the underlying machine learning model is biased or inaccurate, the feature importance calculations will inherit that bias.

1.6 Outline of dissertation

The body of this dissertation consists of four related studies about how the cooling and heating rates of gas (often subject to a spatially-varying incident radiation field) are modeled

in hydrodynamic galaxy formation simulations.

In Chapter 2, we analyze the cooling and heating rates of gas in the ISM of simulated EoR galaxies from the CROC project. We compare the calculations within the simulations (which includes spatial variations in the radiation field) to approximations assuming a uniform radiation field throughout the ISM. Chapter 2 is reproduced with minimal alterations from Robinson et al. (2022), published in *The Astrophysical Journal*. This paper was co-authored with Dr. Camille Avestruz and Dr. Nick Gnedin.

We develop an XGBoost framework to approximate atomic gas cooling and heating rates subject to a spatially-varying incident radiation field in Chapter 3. We train models for 5 fixed element abundance patterns, and compare models trained using different sets of photoionization rates to describe the incident radiation field. This chapter is reproduced from Robinson et al. (2024a), published in the *Monthly Notices of the Royal Astronomical Society*, and also co-authored with Dr. Camille Avestruz and Dr. Nick Gnedin.

In Chapter 4, we extend the machine learning framework of Chapter 3 to characterize the radiation field with other frequency-averages of the intensity. We use SHAP values to find *minimal sets* of radiation field parameters that result in reasonably accurate models. Chapter 4 is reproduced from a pre-print (Robinson et al., 2024b) which has been accepted for publication in the *Open Journal of Astrophysics*, and once again co-authored with Dr. Camille Avestruz and Dr. Nick Gnedin.

Chapter 5 presents the results of implementing one of the machine learning models developed in Chapter 3 into the ART isolated galaxy simulation of Semenov et al. (2021). We compare galaxy properties between the fiducial simulation and a run with our machine learning model. We focus on the temperature-density phase diagram of the gas and emission rates for the fine-structure CII line from various excitation processes. This chapter is reproduced from a preprint (Robinson et al., 2024c), co-authored with Dr. Camille Avestruz, Dr. Nick Gnedin, and Dr. Vadim Semenov, which is currently in the revision process at the *Open Journal of Astrophysics*.

Finally, we present a summary of the main conclusions from the preceding chapters and several potential directions for future work in Chapter 6.

CHAPTER 2

Can Cooling and Heating Functions Be Modeled with Homogeneous Radiation Fields?

2.1 Motivation

The physics of galaxy formation depends on the interplay between gravitational compression and dissipative cooling of the proto-galactic material. On smaller scales, this same behavior is replicated in star formation (e.g. Binney, 1977; Rees & Ostriker, 1977; Silk, 1977). For example, the cooling rate of gas in a galactic disk can affect the accretion rate of new material onto the disk, which in turn affects the star formation rate (Kannan et al., 2014).

Cooling and heating functions describe how the energy density of a gas fluid element changes with time due to radiative processes (e.g. Tucker & Gould, 1966; Cox & Tucker, 1969; Sutherland & Dopita, 1993). Hence, these functions play an important role in galaxy formation. The form of these functions depend on the ionization states and energy levels of atoms and ions, and have been extensively calculated and tabulated for solar and cosmic abundances with the assumption that the gas is in Collisional Ionization Equilibrium (CIE). CIE assumes that there is no incident ionizing radiation and the abundances of various ions are set by a balance between collisional ionization and recombination (Cox & Tucker, 1969; Sutherland & Dopita, 1993; Tucker & Gould, 1966). Under the assumption of CIE, the cooling and heating functions depend *only* on the gas temperature, density, and metallicity. Without the CIE assumption, the cooling and heating functions can depend on any physical quantity which affects the ionization states and energy levels of any chemical element in the gas. In particular, one such quantity is the incident radiation field, specified by its specific intensity J_ν . Radiation fields can ionize and change the energy levels of ions, thus affecting which atomic transitions can emit or absorb photons. Previous works indicate that the presence of ionizing radiation fields can significantly modify the cooling and heating

functions of gas (e.g. Gnedin & Hollon, 2012; Wiersma et al., 2009; Faerman et al., 2022; Romero et al., 2021).

Since radiation fields impact cooling and heating functions, which, in turn, impact galaxy formation, incident radiation fields can affect galaxy formation in a non-trivial way. Some examples of this include suppressing of cooling, leading to fewer dwarf galaxies forming (Efsthathiou, 1992) or preventing cooling of gas at $T \sim 10^{4-4.5}$ K through suppression of cooling via photoionization of neutral hydrogen (Ceverino et al., 2014).

Proper modeling of cooling and heating functions therefore requires an understanding of the dependence of cooling and heating functions on the incident radiation field J_ν . For example, Wiersma et al. (2009) demonstrates that photoionization by extragalactic radiation can modify cooling functions by about an order of magnitude in realistic conditions. There are various ways to approach modeling this dependence, with varying accuracy. The most conceptually straightforward approach is to compute the cooling and heating functions using the full radiation field via a radiative transfer code such as Cloudy (Ferland et al., 1998). However, this is too computationally expensive to be viable in simulations. A more practical approach is to tabulate results from Cloudy with gas temperature, density, and metallicity assuming a spatially constant but temporally varying extragalactic background (e.g. Haardt & Madau, 2001; Faucher-Giguère et al., 2009; Faucher-Giguère, 2020). Both Illustris (Vogelsberger et al., 2014, 2013) and EAGLE (Schaye et al., 2015), two state of the art cosmological hydrodynamic simulations, use such an approach. While this approach does incorporate some effects of photoionizing radiation on cooling and heating, it does not account for the radiation field from local sources within galaxies, which is by far the dominant contribution for ISM gas (Schaye et al., 2015).

Gnedin & Hollon (2012) describe a straightforward, albeit approximate, way to incorporate local J_ν contributions by approximating the full radiation spectrum dependence with several key photoionization rates. Hence, one can tabulate the values for cooling and heating functions in a reasonably sized grid in temperature, density, metallicity, and photoionization rates. We will describe this approach in further detail in section 2.2.2. Other works approach including the effect of the local radiation field in different ways. For example, Ploeckinger & Schaye (2020) includes an interstellar radiation field with the same shape as that of the Milky Way, scaled by the local gas properties. Thus, their approximation scheme does *not* include any radiation field parameters (such as photoionization rates).

In this chapter, we examine the cooling and heating functions of simulated galaxies from the CROC project (Gnedin, 2014). In these simulations, the dependence of the cooling and heating functions on the local radiation field is computed using the approximation scheme of Gnedin & Hollon (2012). We compare three approaches to finding the median cooling

and heating function across halos of similar mass at a given redshift to see whether the forms of these functions vary, and if so, how. In section 2.2, we discuss the methodology: the simulations we use, how we calculate the cooling and heating functions, and the three different cooling and heating function medians we consider. In section 2.3, we demonstrate our results comparing the cooling and heating function medians. Finally, we close with a summary and discussion of our results and potential future work in section 2.4.

2.2 Methodology

2.2.1 CROC simulations

For this work, we use simulations from the CROC project, a program of simulations utilizing the ART code (Kravtsov, 1999; Kravtsov et al., 2002; Rudd et al., 2008). The simulations include many physical processes expected to be necessary to model cosmic reionization self-consistently. These processes include gravity, gas dynamics, star formation, stellar feedback, the formation of molecular hydrogen, ionizing radiation from stars and other sources, radiative transfer, and metallicity and radiation field-dependent cooling and heating computed using the approximation scheme of Gnedin & Hollon (2012), which excludes cooling and heating effects due to cosmic rays, molecules, and dust. The ionizing radiation due to stars is the only ionizing radiation source fully calculated self-consistently. The simulation incorporates other sources in the radiation background as seen by all regions of the simulation, instead of calculating the contribution from those sources locally. For more details on the CROC simulations, see Gnedin (2014).

For this study, we use one $20h^{-1}$ comoving Mpc box size simulation realization (denoted box A), which has a spatial resolution of 100 pc. CROC resolves the radiation field in both space and time, allowing us to study both the space and time-dependence of cooling and heating functions.

2.2.2 Cooling and heating functions

We can divide the rate of change of the gas energy density due to radiative processes into two pieces: processes which increase the energy density of the gas (heating processes), and processes which decrease the energy density of the gas (cooling processes). Guided by this distinction, we can write:

$$\left. \frac{dU}{dt} \right|_{\text{rad}} = n_b^2 [\Gamma(T, \dots) - \Lambda(T, \dots)], \quad (2.1)$$

where U is the energy density of the gas, n_b is the number density of baryons (here, baryons refer to hydrogen, helium, and metal nuclei), and Γ, Λ are the respective heating and cooling functions of the gas; T is the gas temperature and ‘...’ in the cooling and heating functions indicate that these functions generally depend on additional variables besides T . The factor of n_b^2 accounts for the baryon number density dependence of collisional processes involving two gas particles. In CIE, where collisional ionization is balanced by electron recombination, the prefactor of n_b^2 ensures that both Γ and Λ are independent of n_b in the absence of three-body processes (Gnedin & Hollon, 2012). Note that the left-hand side of equation (2.1) only includes changes in energy density due to *radiative* processes, excluding other processes (e.g. gas heating due to adiabatic compression).

While Γ and Λ are independent of density n_b and depend only on gas temperature and metallicity for gas in CIE, an incident radiation field J_ν can modify the distribution of energy levels and ionization states of gas particles, changing Γ and Λ . For an arbitrary J_ν , calculating Γ and Λ would require a radiative transfer code such as Cloudy (Ferland et al., 1998). In order to calculate the radiation field-dependent cooling and heating functions, we use the approximation scheme of Gnedin & Hollon (2012). This is the same approximation that CROC uses to follow gas cooling and heating within the simulation, making this procedure self-consistent (Gnedin, 2014).

This approximation assumes that J_ν follows a general form including contributions from stars and AGN. The implementation approximates the J_ν dependence with 4 key photoionization rates: HI (neutral hydrogen), HeI (neutral helium), CVI (quintuply-ionized carbon), and the Lyman and Werner bands photodissociation rate for molecular hydrogen (parameterized by $Q_{\text{HI}}, Q_{\text{HeI}}, Q_{\text{CVI}}$, and Q_{LW} , where $Q_i = P_i/n_b$, and P_i is the photoionization rate for the relevant band). The cooling and heating functions also depend on the local gas temperature T , baryon number density n_b , and metallicity Z . The cooling and heating functions are computed using the radiative transfer code Cloudy for a table of 4000 values of these parameters, and linear interpolation in each parameter is used for intermediate values (Gnedin & Hollon, 2012).

This approximation scheme fares well compared to Cloudy calculations within the parameter space of the table described above. When evaluated on a uniformly log-spaced grid of parameter values (different from the table described above) and compared to the cooling and heating functions calculated by Cloudy, the approximation scheme leads to fractional errors > 2 in fewer than 1 in 10^3 cases. Fractional errors of about 6 occur in around 1 in 10^6 cases (Gnedin & Hollon, 2012).

2.2.3 Cooling and heating functions in the simulations

We now consider how to describe the cooling and heating functions in a simulation. To put this question in a broader context, imagine that someone wants to simulate a set of individual galaxies similar to ones modeled in our cosmological simulations. What cooling and heating functions should they adopt?

There are several options of increasing level of complexity:

- A** The simplest option would be to adopt a single set of cooling and heating functions $\mathcal{F}(T)$ (i.e. function of temperature only) for the whole simulation. This would be the case if one used, for example, the CIE only cooling function (Sutherland & Dopita, 1993). It is rarely a good approximation, since cooling and heating functions also depend on other gas properties.
- B** The next step is to take the cooling and heating functions $\mathcal{F}(T, n, Z)$ that depend on the gas density and metallicity, but ignore the dependence on the spatial variations in the radiation field (such an approximation may still depend on the uniform radiation field, like cosmic background, e.g. Kravtsov et al., 2002; Wiersma et al., 2009). *This is the most commonly used approximation.* It is the approach used in Illustris-TNG (Vogelsberger et al., 2014, 2013) and EAGLE (Schaye et al., 2015), which both include the effects of a time-varying but spatially homogeneous extragalactic radiation field.
- C** Finally, a simulation can adopt the full dependence of the cooling and heating functions $\mathcal{F}(T, n, Z, J_\nu)$ on gas density, metallicity, and the radiation field at every spatial location. This is the approach used here, and has also been adopted in other simulations such as DRAGONS (Mutch et al., 2016) and Cosmic Dawn (Ocvirk et al., 2016).

The main question we consider in this chapter is how well approximations for cases A and B describe the cooling and heating functions from the case C that is used in our simulations. In order to simplify this question, we consider limited bins in gas density and metallicity within which we can treat densities and metallicities as being approximately constant at the median values n_0 and Z_0 , respectively. Within each bin, cases A and B are identical.

In the simulation, resolution elements (in our case grid cells) within a single bin in density and metallicity have a range of temperatures and radiation fields, and hence also a range of cooling and heating rates \mathcal{F}_i that the cells i have at a given timestep. There are several possible ways one can turn a set of such rates into a single set of cooling and heating functions $\mathcal{F}(T)$.

The most direct way one can define $\mathcal{F}(T)$ is for each value at T to be an average or a median of all cells i that have temperatures T_i sufficiently close to T (say, within a bin of

half-width ΔT , i.e.

$$\begin{aligned}\bar{\mathcal{F}}_{\text{R}}(T) &= \langle \mathcal{F}(T_i, n_i, Z_i, J_{\nu,i}) \rangle_{|T-T_i| < \Delta T} \\ &\approx \langle \mathcal{F}(T_i, n_0, Z_0, J_{\nu,i}) \rangle_{|T-T_i| < \Delta T}\end{aligned}\tag{2.2}$$

The second line of equation (2.2) emphasizes that, hereafter, we only use cells within single bins of density and metallicity, $|n_0 - n_i| < \Delta n$ and $|Z_0 - Z_i| < \Delta Z$. We omit dependence on n_i and Z_i for brevity.

Since the cooling and heating rates span several orders of magnitude, we choose the median rather than the average. We refer to the resulting function from the median rates as the *actual rates*, as it represents the actual cooling and heating rates seen by the cells in the simulation. Note that the actual cooling and heating rates \mathcal{F}_{R} do not necessarily correspond to cooling and heating functions for some “typical” values of density, metallicity, and the radiation field.

The latter can be defined as the cooling and heating functions for the median values of density, metallicity, and the radiation field,

$$\bar{\mathcal{F}}_{\text{M}}(T) = \mathcal{F}(T, n_0, Z_0, \langle J_{\nu,i} \rangle).\tag{2.3}$$

Note that such a median ignores all the information about gas cell temperatures T_i , only cell density, metallicity, and the radiation field contribute to that definition. We will refer to this median as the *median ISM* cooling and heating functions.

One can also define another median, which we call “instantaneous”. Imagine a cell i with given values of T_i , n_i , Z_i , and $J_{\nu,i}$. If it heats/cools instantaneously, without any change in n_i , Z_i , or $J_{\nu,i}$ (which actually change on dynamical and/or star formation timescales), then the temperature of cell i changes with the cooling/heating functions $\mathcal{F}(T, n_i, Z_i, J_{\nu,i})$. Hence, one can define the median instantaneous cooling and heating functions,

$$\bar{\mathcal{F}}_{\text{I}}(T) = \langle \mathcal{F}(T, n_0, Z_0, J_{\nu,i}) \rangle,\tag{2.4}$$

where the median is taken over all cells in the narrow bins of density and metallicity but with any values of temperature or the radiation field.

The distinction between the three definitions of the cooling and heating functions above is subtle but important. The median ISM cooling and heating functions $\bar{\mathcal{F}}_{\text{M}}(T)$ are indeed cooling and heating functions in the canonical sense - they can be computed with, say, Cloudy, for fixed values of n_0 , Z_0 , and $J_{\nu,0} = \langle J_{\nu,i} \rangle$, while the actual cooling and heating rates and the median instantaneous cooling and heating functions may not be. The physical

interpretations of the actual cooling and heating rates and median ISM functions are more clear than for the instantaneous case.

In order to illustrate that a collection of cooling and heating rates is not equal to the cooling and heating function, let us consider a single fluid element with the given density n_i , metallicity Z_i , and the radiation field J_i . Its cooling function is $\Lambda(T, n_i, Z_i, J_i(\nu))$. As that fluid element cools thermodynamically (let us assume its heating function is initially small), it may also evolve dynamically. In general, the density of that element will change (for example, if the cooling is isobaric or adiabatic), unless the cooling is isochoric and no new gas is introduced into the element. The metallicity and radiation field may also change. The cooling rate along the flowline of that fluid element is $\dot{E}_C(t) = \Lambda(T_i(t), n_i(t), Z_i(t), J_i(\nu, t))$ and the collection of these cooling rates in some time interval, $\{\dot{E}_C\}_{t_0 < t < t_1}$ may not correspond to any cooling function $\Lambda(T_i(t), n_*, Z_*, J_{\nu*})$ for some fixed parameters n_* , Z_* , and $J_{\nu*}$. We show in subsequent sections that this is in fact the case, and highlight the comparison for a range of halo masses and gas metallicities.

2.2.4 Median cooling and heating rates and functions of the ISM

We ultimately want to examine how the actual cooling and heating rates and median ISM and instantaneous cooling and heating functions of the interstellar medium of galaxies in our cosmological simulations vary with cosmologically relevant quantities, such as redshift z and host halo mass. Before we can quantify this behavior, we first need to examine differences in median cooling and heating rates and functions to capture the average behavior of galaxies in a given mass range and at a given redshift.

In practice, we consider dark matter halos within the cosmological simulation, rather than individual galaxies. For each halo identified by the ROCKSTAR (Behroozi et al., 2013) halo catalog at the relevant redshift, we select all gas cells within one virial radius R_{vir} of the center of each halo. For the mass of the halo, we use the virial mass M_{vir} from the ROCKSTAR halo catalog. To account for potential effects due to variations in spatial clustering, we also distinguish between central halos and subhalos using the ‘parent ID’ from the ROCKSTAR halo catalog. The particular simulation from the CROC project we use for this work has redshift snapshots ranging between $z \sim 5$ and $z \sim 10$. Except for when considering the mass dependence of the cooling and heating functions, we choose a fiducial mass range of $M_{\text{vir}} > 10^{10} h^{-1} M_{\odot}$ (for the rest of this chapter, we write M for M_{vir} , as virial masses are the only masses used here).

For each combination of mass range, redshift, and choice of central or subhalos, we randomly select 50 halos when they are available, and otherwise select all the halos. To find

and utilize the temperature, number density, metallicity, and photoionization rates of cells within each halo, we used the toolkit `yt` (Turk et al., 2011). We combine the cell data from each of the selected halos, then compute the median cooling and heating functions using the procedures described in section 2.2.3 for specific bins in density and metallicity (see below). At the density scales explored in this chapter, the hydrodynamic cells have approximately constant mass. Hence, in physical terms, the median over hydrodynamic cells is effectively mass-weighted.

We also apply cell selection criteria when calculating the cooling and heating functions. We examine the cooling and heating functions for cells in the density range of $1 < n_b < 10 \text{ cm}^{-3}$, the metallicity range $0.03 < Z/Z_\odot < 0.1$, where $Z_\odot \approx 0.02$ is the metallicity of gas in the solar neighborhood (Gnedin & Hollon, 2012), and with radiation field values $P_{\text{HI}}, P_{\text{HeI}} > 0$, and $P_{\text{C VI}} > 2 \times 10^{-20} \text{ s}^{-1}$. The density range corresponds to the typical galactic ISM. We limit the metallicity to a relatively narrow range since cooling and heating functions can depend on metallicity. Generally, both cooling and heating functions increase with metallicity across all temperatures, since adding more metals to the gas increases the available pathways for radiative cooling and heating. Increasing metallicity can also introduce new features (i.e. peaks) to the cooling and heating functions due to elements besides hydrogen and helium. This range ($0.03 < Z/Z_\odot < 0.1$) ensures that we have sufficient cells for our analysis between redshifts $z \sim 10$ and $z \sim 5$. The photoionization rate cuts are discussed in more detail in the next section. Unless otherwise noted, all gas cells used in the analysis below are within these parameter ranges.

For each approach described in section 2.2.3, we compute the median cooling and heating rates and functions over all cells remaining after the cuts described above from the selected halos at a given redshift in our fiducial mass bin. As a measure of spread, we compute the 25th and 75th percentiles. For the spread of actual rates and the instantaneous functions, we compute percentiles within all cells in each temperature bin. For the median ISM cooling and heating functions, we compute the spread by evaluating the cooling and heating functions (in each temperature bin) at the 25th and 75th percentile ISM (that is, the 25th and 75th percentiles of $n_b, Z, P_{\text{LW}}, P_{\text{HI}}, P_{\text{HeI}}$, and $P_{\text{C VI}}$). This spread is dominated by spread between cells rather than spread between the medians of different halos. To choose temperature bins, we begin by finding the minimum and maximum cell temperatures for the halos under consideration. We construct logarithmic temperature bins between these values, with 20 bins per decade. For ease of comparison, we evaluate the median ISM and instantaneous cooling and heating functions for the same halos at the logarithmic center of the bins used for the actual cooling and heating rates (the geometric mean of the bin edges).

Redshift z	Fraction of discarded cells (%)
5	1.5
8	2.1
9	3.2
10	0.7

Table 2.1: Cell fraction with P_{HI} , P_{HeI} , or $P_{\text{CVI}} = 0$ after density and metallicity cuts.

2.2.5 Numerical artifacts

We exclude cells with any of P_{HI} , P_{HeI} , P_{CVI} equal to 0 because such cells can yield anomalously large cooling and heating function values with the approximation described in section 2.2.2. The table used to interpolate between photoionization rate values in the approximation does not extend to P_{HI} , P_{HeI} , P_{CVI} this low, so we interpret these cooling and heating function values as numerical artifacts. Some simulated cells where P_{HI} , P_{HeI} , or $P_{\text{CVI}} = 0$ are obvious numerical artifacts (that is, none of these three photoionization should be 0 for that cell). It is possible that some cells with zero photoionization rates are *not* numerical artifacts of the simulation, but there is no straightforward way to separate such cells, so we choose to exclude all such cells in this analysis. This restriction removes a few percent of the cells with density and metallicity in the ranges described in section 2.2.4. The fraction of cells removed for the (up to) 50 randomly selected central halos in our fiducial mass bin at various redshifts is shown in Table 2.1.

Furthermore, we also exclude cells with $P_{\text{CVI}} < 2 \times 10^{-20} \text{ s}^{-1}$. We find that the distribution of the CVI photoionization rates, P_{CVI} , of gas cells in all galaxies in the fiducial mass range is distinctly bimodal across the entire range of simulated redshifts, with no difference in behavior between galaxies hosted by central or subhalos. A representative plot is shown for our 50 randomly selected central halos at $z \sim 5$ with masses $M > 10^{10} h^{-1} M_{\odot}$ in Figure 2.1. We term the two distinct regions of the distribution ‘low P_{CVI} ’ for gas populating the distribution in blue and ‘high P_{CVI} ’ for gas populating the distribution in orange. A cutoff of $2 \times 10^{-20} \text{ s}^{-1}$ (as shown in the mid-plot horizontal axis in Figure 2.1) separates the two modes at all the redshifts we examined. The range of P_{CVI} for the low P_{CVI} distribution does not vary strongly with redshift. However, the location of the narrow range of high P_{CVI} values systematically increases with time from $\sim 7 \times 10^{-20} \text{ s}^{-1}$ at $z \sim 10$ to $\sim 3 \times 10^{-17} \text{ s}^{-1}$ at $z \sim 5$ as the cosmic X-ray background gradually builds up. The cells in each mode are not separated in the phase space of gas properties. That is, cells in both parts of the P_{CVI} distribution are similarly distributed in temperature, number density, and metallicity; P_{CVI} is the only feature which cleanly separates the two distributions. The correlation between P_{CVI} and P_{LW} for the low P_{CVI} population seen in Figure 2.1 is surprising. CVI is ionized

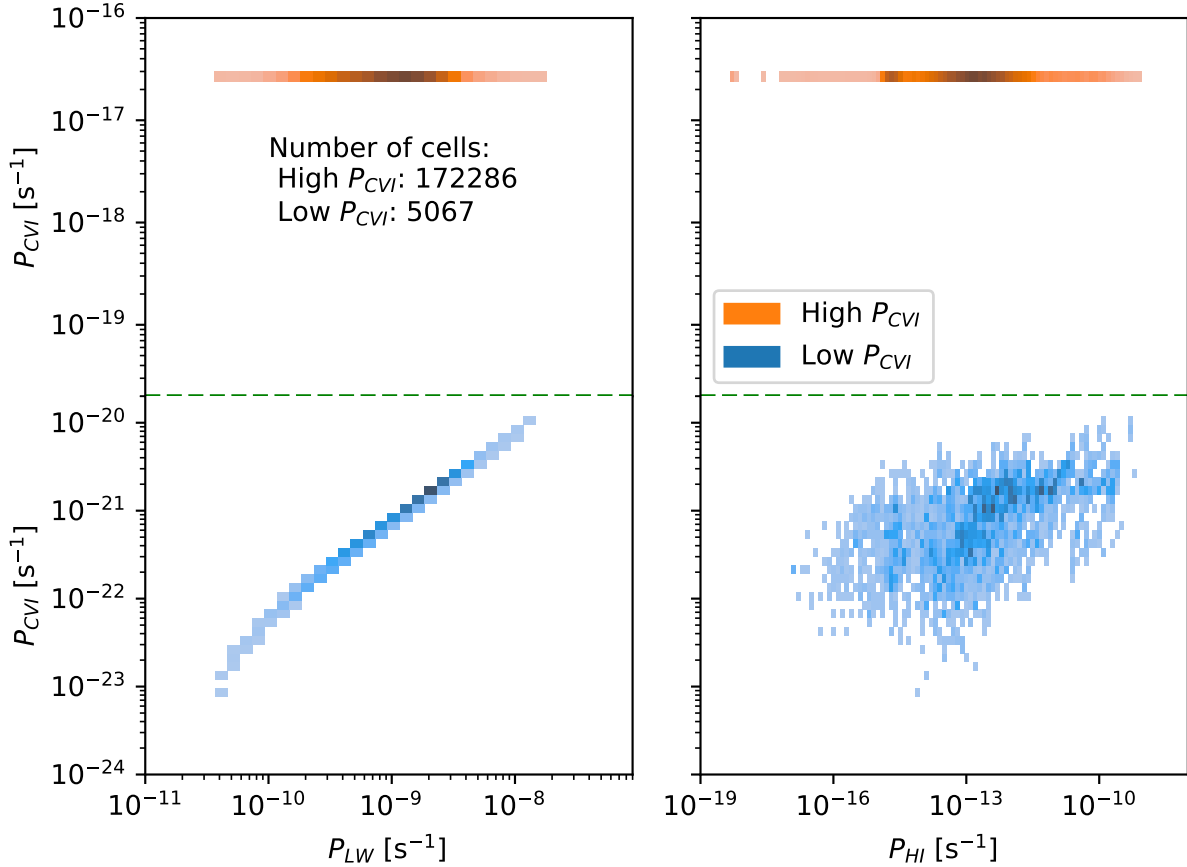


Figure 2.1: Histograms of the CVI photoionization rate P_{CVI} vs the Lyman-Werner band photodissociation rate P_{LW} (left) or the HI photoionization rate P_{HI} (right) for cells in the 50 randomly selected central halos in the fiducial mass range at $z \sim 5$. The dashed green horizontal line shows the cutoff between the low and high radiation field parts of the distribution ($P_{\text{CVI}} = 2 \times 10^{-20} \text{ s}^{-1}$). The two populations of cells are distinguished by color. The low P_{CVI} values and correlation between P_{CVI} and P_{LW} for the low P_{CVI} population have no obvious physical reason and are likely to be numerical artifacts.

Redshift z	Fraction of discarded low P_{CVI} cells (%)
5	2.9
8	2.8
9	7.9
10	8.8

Table 2.2: Redshift evolution of low P_{CVI} cell fraction.

by X-rays, sourced only by the spatially constant quasar background in the simulation. On the other hand, the UV radiation in the Lyman-Werner (LW) bands comes from both local stellar sources and a spatially constant UV background (Gnedin, 2014). Since these rates originate from disparate sources, there is no obvious reason why they should be correlated. We examined some of these cells visually, and while some of them are located in very high density regions and could be optically thick to X-rays, not all such cells are obviously located in these regions. Hence, we also choose to exclude the cells with low P_{CVI} , since we cannot clearly show that they are not a previously unidentified numerical artifact of the radiative transfer solver.

Note, no more than a few percent of gas cells (for halos in our fiducial mass range at a given redshift) are cut from our cooling and heating function calculation with these radiation field selection criteria (after the density and metallicity selection criteria have already been applied).

We find that the fraction of cells with low P_{CVI} values (which we remove) varies with redshift. In Table 2.2, we show the fraction of total gas cells in (up to) 50 randomly selected central halos after the density, metallicity, and photoionization rate cuts described above with low P_{CVI} across various redshifts. The fraction of low- P_{CVI} cells generally increases with increasing redshift.

2.3 Results

2.3.1 Actual rates vs. median functions

We first compare the behavior of the median ISM and instantaneous cooling and heating functions against the actual cooling and heating rates for the 50 randomly selected central halos at $z \sim 5$ in our fiducial mass range, shown in Figure 2.2. From Figure 2.2, we see that the instantaneous and median ISM cooling and heating functions differ from the actual cooling and heating rates by more than the 25th-75th percentile spread below $T \sim 10^3$ K. They even have different qualitative shapes at these lower temperatures. Above $T \sim 10^4$




Figure 2.2: Comparison of actual rates (dashed blue curves) and the median ISM (dotted green) and instantaneous (solid orange) functions for cooling (top panel) and heating (bottom panel) for the 50 randomly selected central halos at $z \sim 5$ in the fiducial mass range $M > 10^{10} h^{-1} M_{\odot}$ (the same as shown in Fig. 2.1). The shaded bands show the 25th-75th percentile spread.

K, there is quantitative agreement for the cooling functions, but only qualitative agreement between the shapes of the heating functions. On the other hand, the median ISM and instantaneous cooling and heating functions agree very well across all temperature bins, the main difference being that the median ISM cooling and heating functions have slightly smaller spread. Since the median ISM cooling and heating functions have a clearer theoretical interpretation than the instantaneous case, and these two agree very closely here, we will now focus on the comparison between actual cooling and heating rates against the median ISM cooling and heating functions.

The median ISM cooling and heating functions and actual cooling and heating rates are *not* the same. This suggests that the actual cooling and heating rates experienced by actual gas cells as their temperature changes are not given by the cooling and heating functions evaluated at median ISM properties. Hence, we find that the actual cooling and heating rates are *not* expressible in the form $\mathcal{F}(T, n_*, Z_*, J_*(\nu))$ for fixed $n_*, Z_*, J_*(\nu)$ (as explained in section 2.2.3). For example, the actual heating rates do not monotonically decrease with temperature, but all $\Gamma(T, n_*, Z_*, J_*(\nu))$ do. Since the actual cooling and heating rates explicitly describe the thermal evolution of gas cells, we conclude that the thermodynamics of the simulated gas cells cannot be well-described by a single set of cooling and heating functions of the form $\mathcal{F}(T, n_*, Z_*, J_*(\nu))$, which is a necessary approximation for simulations which use cooling and heating functions with a homogeneous radiation field computed with tools such as Cloudy. This illustrates how the assumption of a spatially fixed radiation field (the case B approximation discussed in section 2.2.3) can break down in the presence of local radiation fields. Since we have restricted $1 < n_b < 10 \text{ cm}^{-3}$ and $0.03 < Z/Z_\odot < 0.1$ for these cells, the spread in median ISM cooling and heating functions is primarily due to variations in photoionization rates. We discuss this in further detail below.

2.3.2 Redshift trends

The actual cooling and heating rates do not vary strongly with redshift. For simplicity, we do not overplot these. On the other hand, there is redshift evolution for the median ISM functions. Since the number of cells in each temperature bin varies and some temperature bins may contain *no cells at all*, the actual cooling and heating rates are much more ‘jagged’ and less smooth than the median ISM cooling and heating functions (see Figure 2.2, Figure 2.4, and Figure 2.7). This effect tends to obscure the effects of the redshift evolution seen for the median ISM, discussed below. Figure 2.3 shows the median ISM cooling (top panel) and heating (lower panel) functions for gas in central halos in our fiducial mass range at $z \sim 5, 8, 9$, and 10 in the solid curves. Note that the temperature range for each curve

is given by the minimum and maximum temperatures of the selected (i.e. within the density, metallicity, and photoionization rate ranges specified in section 2.2.4) cells in the up to 50 halos used at the given redshift. We overplot the heating function in faint lines in the top panel to illustrate the evolution of the intersection point (i.e. equilibrium temperature) between the cooling and heating functions. The equilibrium temperature increases with redshift: from $\sim 10^4$ K at $z \sim 5$ to 2×10^4 K at $z \sim 10$. For comparison of populations, the dashed lines correspond to the same curves for gas in subhalos at $z \sim 5$, where we have a sufficient number of reasonably resolved subhalos to illustrate a comparison; the central and subhalo populations have indistinguishable median ISM functions within the spread.

From Figure 2.3, we see that the median ISM cooling functions increase with redshift for $T \lesssim 10^4$ K and the median ISM heating functions increase with redshift for all temperatures. At $z \sim 9, 10$, the median cooling and heating functions are about three orders of magnitude larger than at lower redshifts, reflecting the weakness of the extragalactic X-ray background in the earliest galaxies before the onset of reionization, reflected in the lower value of P_{CVI} (see section 2.2.5). As the X-ray background increases, it will heat up and ionize gas that was previously cooler and less ionized. Cooler gas heats more efficiently than hotter gas with the same other properties. Consistent with the expected decrease in cooling efficiency of hotter gas, Figure 2.2 shows that the median ISM heating function is a monotonically decreasing function of temperature. This overall suggests that the median ISM heating function should decrease as the X-ray background builds up with decreasing redshift, as we see in Figure 2.3. At all the redshifts shown, the cooling function has two peaks (above $\sim 10^4$ K) corresponding to hydrogen and helium. The location of the first peaks shifts to higher temperatures at $z \sim 9, 10$ compared with $z \sim 5, 8$.

2.3.3 Mass dependence

Our fiducial halo mass range is relatively wide: $M > 10^{10} h^{-1} M_{\odot}$. To justify this choice, we examine the actual cooling and heating rates and median ISM cooling and heating functions in narrow mass bins across the range of halo masses at $z \sim 5$. We choose mass bins of $(1.0 - 1.1) \times 10^{10} h^{-1} M_{\odot}$, $(5.0 - 5.5) \times 10^{10} h^{-1} M_{\odot}$, and $(1.0 - 1.1) \times 10^{11} h^{-1} M_{\odot}$. We show the cooling and heating functions in these mass bins at $z \sim 5$ in Figure 2.4.

We observe no halo mass dependence for either the actual cooling and heating rates or median ISM cooling and heating functions at $z \sim 5$. While there are slight offsets in the curves for the three mass bins at low T ($\lesssim 10^5$ K for heating and $\lesssim 10$ K for cooling), these offsets are much smaller than the 25th-75th percentile bands. Despite the wide choice of fiducial mass bin, neither the actual cooling and heating rates nor the median ISM cooling




Figure 2.3: Median ISM cooling (top panel) and heating (bottom panel) functions vs. temperature for up to 50 central halos in the fiducial mass range at $z \sim 5, 8, 9, 10$ (solid lines) and for sub halos at $z \sim 5$ (dashed lines). The solid lines show the median, while the shaded bands show the 25th-75th percentile spread for central halos at $z \sim 5$. For direct comparison, we overplot the heating functions in the upper panel with faint lines; the intersection between the cooling and the heating function corresponds to the equilibrium temperature.




Figure 2.4: Actual rates (blue) and median ISM functions (orange) for cooling (top panel) and heating (bottom panel) for central halos in three mass bins at $z \sim 5$. The 25th-75th percentile shaded region is shown only for the middle mass bin, but the widths are similar for all bins.

and heating functions depend on the halo mass within our fiducial mass range.

2.3.4 Density and metallicity ranges

To consider our choices of $1 < n_b < 10 \text{ cm}^{-3}$ and $0.03 < Z/Z_\odot < 0.1$ for density and metallicity ranges, we explore the median ISM cooling and heating functions and actual cooling and heating rates for other density and metallicity ranges (using the same halos as above).

First, we consider the density range. Similar to the lack of redshift dependence, the actual cooling and heating rates also do not vary strongly with the choice of density bin, so we do not overplot these curves. Instead, we explore how much of the median ISM spread is due to the width of our chosen density range. In Figure 2.5, we plot the median ISM cooling and heating functions for density ranges of $1 < n_b < 3$, and $3 < n_b < 10 \text{ cm}^{-3}$, in addition to $1 < n_b < 10 \text{ cm}^{-3}$. For comparison, we also include the median ISM cooling and heating functions for density ranges $0.01 < n_b < 0.1$ and $0.1 < n_b < 1 \text{ cm}^{-3}$, which are of the same logarithmic width as our $1 < n_b < 10 \text{ cm}^{-3}$. These low density ranges explore the density range of the CGM rather than the ISM.

As shown in Figure 2.5, the 25th to 75th percentile spreads in the median ISM cooling and heating functions for $1 < n_b < 3$ and $3 < n_b < 10 \text{ cm}^{-3}$ have comparable width to that for $1 < n_b < 10 \text{ cm}^{-3}$, and all three ranges overlap very strongly. This suggests that the spread within our adopted density range of $1 < n_b < 10 \text{ cm}^{-3}$ is *not* primarily due to the range of cell densities, but rather to the spread in photoionization rates. However, the median ISM cooling and heating function spreads for $0.01 < n_b < 0.1$, $0.1 < n_b < 1$, and $1 < n_b < 10 \text{ cm}^{-3}$ overlap only weakly, if at all. For the lowest density range shown with the blue dotted line, $0.01 < n_b < 0.1 \text{ cm}^{-3}$, the peaks in the cooling function that correspond to hydrogen and helium shift to higher temperatures (but lower cooling function values). The shift to higher temperatures is due to the fact that fewer atomic collisions occur in low gas density regions, e.g. the CGM. This gas therefore requires higher temperatures to excite a sufficient number of hydrogen atoms for radiative cooling, which is ultimately less efficient than radiative cooling in a higher density medium.

In general, the heating function includes contributions from direct photoionization heating. Since this process involves one photon and one atom, the rate of change of energy density is proportional to n_b rather than n_b^2 (Sutherland & Dopita, 1993). Due to the definition of the heating function $\Gamma(T, \dots)$ in equation (2.1), we can expect $\Gamma(T, \dots) \propto n_b^{-1}$ if photoionization is the only contribution to the heating function. Thus, for a density range of $1 < n_b < 10 \text{ cm}^{-3}$, a one dex spread in the cooling and heating functions (at fixed T, Z ,




Figure 2.5: Median ISM cooling (top panel) and heating (bottom panel) functions vs. temperature for selected cells from the random subsample of 50 central halos at $z \sim 5$ in our fiducial mass range (the same sample as previous figures) for cells with $0.01 < n_b < 0.1$ (blue), $0.1 < n_b < 1$ (orange), $1 < n_b < 10$ (green), $1 < n_b < 3$ (red), and $3 < n_b < 10$ cm^{-3} (purple). The bands correspond to the 25th to 75th percentile spread for each density range.

and radiation field, i.e. the median ISM case) can be explained by the width of the density range. It is therefore also useful to compare the product $n_b\Gamma$ for the density ranges discussed above, since $n_b\Gamma$ would be constant across all density ranges if $\Gamma \propto n_b^{-1}$ from photoheating were the only n_b dependence. We show this product for the median ISM in Figure 2.6. Here, we see that $n_b\Gamma$ is indistinguishable for temperatures below $T \sim 4 \times 10^5$ K to within the 25th to 75th percentile spread for all density ranges except $0.01 < n_b < 0.1 \text{ cm}^{-3}$. The 25th to 75th percentile spread in $n_b\Gamma$ for all density ranges is considerable (larger than 1 dex at the lowest temperatures). The scatter in $n_b\Gamma$ is primarily due to the scatter in ISM properties other than n_b , i.e. the photoionization rates, since we also select a relatively narrow range for the metallicity, Z . By definition, the radiation field for the median ISM cooling and heating functions is constant. The photoionization heating rate is therefore also constant. Thus, if photoionization is the dominant contribution to the heating function (say, over a given temperature range), then the median ISM heating function will also be constant in that regime.

To consider any potential dependence on our adopted metallicity range of $0.03 < Z/Z_\odot < 0.1$, we examine a different, non-overlapping metallicity range of similar logarithmic width, $0.1 < Z/Z_\odot < 0.3$. We compare both the actual cooling and heating rates and median ISM cooling and heating functions in Figure 2.7. The qualitative shapes are similar for both metallicity ranges for both actual rates and median ISM functions, but the normalization differs for $T \lesssim 10^4$ K for cooling functions and all temperatures for heating. Our choice of fiducial metallicity range has some effect on the numerical values of the cooling and heating functions at a given redshift, but does not affect our conclusions about the differences between the actual cooling and heating rates and median ISM cooling and heating functions.

2.4 Summary and discussion

In this chapter, we compare the actual cooling and heating rates with the median ISM and instantaneous cooling and heating functions of halos in a simulation from the CROC project to assess the validity of universal (i.e. independent of redshift and halo mass) cooling and heating function formulations with a spatially fixed radiation field. Such formulations are often used in cosmological hydrodynamic simulations, as described in the case B approximation discussed in section 2.2.3. Here, cooling and heating functions are computed with tools such as Cloudy using a fixed extragalactic background radiation field. The main conclusions from this work are:

- The actual cooling and heating rates cannot be described by the canonical cooling and heating functions of temperature for median values of gas density, metallicity, and the




Figure 2.6: The median baryon number density n_b multiplied by the median ISM heating function vs. temperature for selected cells from the random subsample of 50 central halos at $z \sim 5$ in our fiducial mass range (the same sample as above) for cells with $0.01 < n_b < 0.1$ (blue), $0.1 < n_b < 1$ (orange), $1 < n_b < 10$ (green), $1 < n_b < 3$ (red), and $3 < n_b < 10$ cm⁻³ (purple). The shaded regions correspond to the 25th and 75th percentile n_b values respectively multiplied by the heating function evaluated for 25th and 75th percentile ISM properties.




Figure 2.7: Actual (solid lines) cooling and heating rates and median ISM (dashed lines) cooling (top panel) and heating (bottom panel) functions vs. temperature for selected cells from a random subsample of 50 central halos at $z \sim 5$ in our fiducial mass range (the same sample as above) for cells with $0.03 < Z/Z_{\odot} < 0.1$ (blue, as used for the above plots) and $0.1 < Z/Z_{\odot} < 0.3$ (orange).

radiation field (see Figure 2.2). In fact, the heating rates depend non-monotonically on the gas temperature and thus cannot be described by a heating function at *any* fixed values of gas density, metallicity, and the radiation field. Specifically, the commonly used approximation of the gas cooling and heating functions with dependence only on T, n_b, Z , and a spatially fixed radiation field is inadequate to describe the actual cooling and heating rates of the ISM illuminated with a spatially varying radiation field (such as with the simulations used here). While one can still parameterize these rates as functions of temperature, such parameterizations do not have the form of cooling and heating functions which could be computed with typical, fixed values of the gas density, metallicity, and the radiation field as might be done with codes such as Cloudy.

- The median ISM cooling and heating functions (a single cooling and heating function evaluated at the median n_b, Z , and photoionization rates) and instantaneous cooling and heating functions (the median of the cooling and heating functions evaluated for the n_b, Z , and photoionization rates of each cell) are nearly identical (see Figure 2.2).
- While there is no significant redshift evolution for the actual cooling and heating rates, the median ISM cooling functions increase with redshift for $T \lesssim 10^4$ K and the median ISM heating function increases with redshift for all temperatures (see Figure 2.3), reflecting the weakness of the extragalactic X-ray background from quasars prior to the onset of recombination. Note, 50% of the hydrogen becomes ionized between $z \sim 7 - 8$ in the CROC project simulations considered here (Gnedin & Kaurov, 2014).
- Galaxies hosted by central halos and subhalos exhibit no systematic difference in either their actual cooling and heating rates or their median ISM cooling and heating functions.
- There is no significant mass trend in either the actual cooling and heating rates or median ISM cooling and heating functions across a decade of mass within our fiducial mass bin (see Figure 2.4).
- The scatter in cooling and heating rates is dominated by the scatter within individual halos rather than the scatter between different halos.

In our analysis we discovered that a few percent of all cells in CROC simulations have anomalous photoionization rates ($P_{\text{HI}} = P_{\text{HeI}} = P_{\text{CVI}} = 0$ or P_{CVI} is suspiciously low), resulting in anomalously high cooling and heating function values in such cells. In future work, we will consider eliminating these numerical artifacts by extending the interpolation

tables in the approximation of Gnedin & Hollon (2012) to smaller values of P_{HI} , P_{HeI} , and P_{CVI} . In addition to extending the interpolation tables, it may also be possible to improve the approximation (and extend its range of validity further) by using a different combination of photoionization rates in place of the four described in equation (9) of Gnedin & Hollon (2012). Machine learning provides a promising tool to assess which combinations of photoionization rates might most impact cooling and heating functions. A natural follow-up project would be to use these rates to construct new interpolation tables for use in cosmological simulations with radiation fields. We explore interpretable machine learning approaches to interpolating cooling and heating functions in Chapter 3 and Chapter 4.

2.5 Acknowledgements

This chapter has been co-authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. This work used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. This research is also part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. This research was also supported in part through computational resources and services provided by Advanced Research Computing (ARC), a division of Information and Technology Services (ITS) at the University of Michigan, Ann Arbor. Co-author CA acknowledges support from the Leinweber Center for Theoretical Physics at the University of Michigan. I acknowledge support from the Physics Department Fellowship at the University of Michigan.

CHAPTER 3

Exploring the Dependence of Gas Cooling and Heating Functions on the Incident Radiation Field with Machine Learning

3.1 Motivation

Galaxy formation involves many interacting processes, of which the primary one is the gravitational collapse of baryonic gas into the potential wells of dark matter halos after the decoupling of baryons and photons. Since baryonic gas can provide thermal pressure support, the rate at which the gas can dissipate energy (via cooling) is a critical factor in determining the density and temperature at which the gravitational collapse stops (Rees & Ostriker, 1977). Cooling and heating functions are a traditional way of describing how the internal energy of the gas changes due to radiative processes (e.g. Cox & Tucker, 1969; Sutherland & Dopita, 1993; Lykins et al., 2013; Wang et al., 2014), and determine the thermal evolution of gas (e.g. Dalgarno & McCray, 1972; Gnat & Sternberg, 2007). Coupled with other relevant processes, cooling and heating functions help determine the overall evolution of the gas (e.g. Martínez-Serrano et al., 2008; Richings et al., 2014; Galligan et al., 2019; Romero et al., 2021). Hence, cooling and heating functions are important to theoretical modeling of galaxy formation (see the review of Benson, 2010). For example, the comparison between the gas cooling time and the gravitational freefall time introduces characteristic scales where the two are equal (e.g. Rees & Ostriker, 1977; Silk, 1977; White & Frenk, 1991; Kauffmann et al., 1993). Cooling also regulates how efficiently additional gas can accrete onto a (proto)galaxy (e.g. Binney, 1977; Bertschinger, 1985; Cole et al., 1994; Croton et al., 2006; Brooks et al., 2009), and the radiation emitted by the accreting gas (Fardal et al., 2001). Gas mass can be lost from a (forming) galaxy due to heating by ionizing radiation from outside the galaxy (Okamoto et al., 2008).

Given an atomic and molecular composition, and using values for relevant atomic prop-

erties, the equilibrium populations of various ionization states and energy levels in a gas can be computed (e.g. Spitzer, 1962; Arnaud & Rothenflug, 1985; Ferland, 1993, 2009). These populations can then be used to determine the cooling and heating rates of the gas (e.g. Cox & Tucker, 1969; Dalgarno & McCray, 1972; Sutherland & Dopita, 1993; Ferland et al., 1998). In the absence of external sources of ionizing radiation, these populations are set by CIE and become well-known and relatively simple functions of gas density, metallicity, and temperature only (Cox & Tucker, 1969; Sutherland & Dopita, 1993).

An incident radiation field can change the distribution of ionization states and energy level populations, and hence can alter cooling and heating functions from the CIE limit (e.g. Wiersma et al., 2009; Gnedin & Hollon, 2012; Galligan et al., 2019, and Chapter 2). Computing cooling and heating functions for an arbitrary radiation field is a highly complex and computational expensive effort, requiring major development efforts and sophisticated software, in particular photoionization codes (see the review of Kallman, 2001) such as Cloudy (Ferland et al., 1997, 1998, 2013, 2017; Chatzikos et al., 2023). Photoionization codes require a comprehensive database of atomic data related to ionization and recombination, such as XSTAR (Bautista & Kallman, 2001). Photoionization codes have been developed for a variety of geometries, such as a spherical cloud with an illuminating source at the center (Kallman & McCray, 1982), a slab illuminated from either side (Dumont et al., 2000), a cone illuminated from the apex (Kinkhabwala et al., 2003), pseudo-3D geometries with non-extended sources (Morisset et al., 2005), and fully general 3D geometries (e.g. Ercolano et al., 2003; Wood et al., 2004; Baes et al., 2005).

In this chapter, we aim to accurately predict cooling and heating functions for an atomic gas in the presence of an incident radiation field, which may include a local component in addition to the extragalactic background. Local contributions can dominate the radiation field inside galaxies and in the CGM (Draine, 1978).

Since gas cooling (and heating) is a vital ingredient in galaxy formation, cooling and heating functions need to be included in numerical simulations of galaxy formation. However, the full calculation of level and ionization state populations is much too complex (and, hence, slow) for modeling cooling and heating functions on the fly in galaxy formation simulations. Hence, for simulation use, the cooling and heating functions must be approximated with a less computationally intensive method. One traditional approach to this is pre-computing cooling and heating functions with Cloudy on a grid of relevant parameters, and interpolating between table values in the simulation (e.g. Kravtsov, 2003; Smith et al., 2008; Hopkins et al., 2011; Vogelsberger et al., 2014). Cooling and heating tables for gas in CIE generally include dimensions of temperature, density, and metallicity (e.g. Cox & Tucker, 1969; Sutherland & Dopita, 1993; Smith et al., 2008). The same table dimensions can be used for

gas with a constant ionizing background at fixed redshift (e.g. Hopkins et al., 2011). Incorporating a spatially constant UV background adds one additional dimension, the redshift z (or radiation field strength), to the tables (e.g. Kravtsov, 2003; Robertson & Kravtsov, 2008; Smith et al., 2017; Gutcke et al., 2021; Schaye et al., 2023). In particular, various simulations incorporate the redshift-dependent cooling tables of Wiersma et al. (2009), including simulations described in Thomas et al. (2009); Schaye et al. (2015); McCarthy et al. (2017). More recent simulations, such as those of Gutcke et al. (2021) and Schaye et al. (2023) incorporate the cooling tables of Ploeckinger & Schaye (2020). These tables include an interstellar radiation field with a fixed frequency dependence and normalization depending on the gas column density, in addition to a spatially constant UV background. Tables with more dimensions can also be constructed for gas illuminated by an incident radiation field which includes local contributions from a synthesized stellar spectrum and a quasar-like power law (Gnedin & Hollon, 2012).

Interpolation tables are not the only approach to approximating cooling and heating functions. One alternate approach is simplified chemical networks which follow only the most relevant atoms and molecules, which *can* be evaluated on-the-fly (e.g. Anninos et al., 1997; Grassi et al., 2011; Richings et al., 2014; Salz et al., 2015; Bovino et al., 2016). Other works use Cloudy to calculate cooling and heating functions for gas properties sampled from a numerical galaxy formation simulation of interest (e.g. Robertson & Kravtsov, 2008; Wiersma et al., 2010; Romero et al., 2021, and Chapter 2). These ‘true’ cooling and heating functions can then be used to train a neural network to predict cooling and heating functions (at least for gas properties which are reasonably consistent with those seen in the training data) (e.g. Grassi et al., 2011; Galligan et al., 2019).

Cooling and heating functions depend on several gas parameters, as well as the incident radiation field, which is in general a function of frequency. The effect of the incident radiation field on the gas can be described by the photoionization rates for all possible ionization states of elements in the gas. Since an interpolation table in very large dimension is computationally impractical, we need to find a smaller number of photoionization rates that are still able to capture the dependence of cooling and heating functions on the incident radiation field. This is an example of the general class of dimensionality reduction problems. Such problems are common in machine learning, and several machine learning techniques exist to approach these problems (see reviews by Zebari et al., 2020; Jia et al., 2022).

Several additional factors which we do not include in this chapter can also impact gas cooling and heating, such as cosmic rays, dust, and molecules. Cosmic rays can ionize gas particles, heating the gas. Dust grains contribute to the heating function via photoheating, and to cooling via collisions between gas particles and dust grains (Ferland, 1993). Dust

also impacts the cooling function indirectly, via the accretion of gas-phase metals onto dust grains. In general, this accretion will modify the relative abundances of metals (Dwek, 1998). Note that UV radiation fields can destroy dust grains through sublimation (Guhathakurta & Draine, 1989). Hence, low incident ionizing radiation is needed to have significant dust mass. Molecules contribute to heating via photodissociation, and cooling via vibrational and rotational line emission (Richings et al., 2014). In a hydrodynamic simulation, additional numerical approximations for the contributions of molecules, dust, and cosmic rays to the total cooling and heating functions could be added to the atomic cooling and heating approximations described in this chapter.

In this work, we use the machine learning algorithm XGBoost (Chen & Guestrin, 2016)) to model the dependence of cooling and heating functions on gas properties and various radiation field parameters. We train XGBoost models on tabulated cooling and heating functions calculated with Cloudy (Ferland et al., 1998). XGBoost is known to perform well on this type of tabular training data (Shwartz-Ziv & Armon, 2022; Grinsztajn et al., 2022). We also explore the SHAP (Lundberg & Lee, 2017)) values (Lundberg et al., 2018; Lundberg et al., 2020) of our model inputs, which we use to determine what radiation field parameters are most predictive of cooling and heating functions. The XGBoost algorithm has been used in the astrophysics literature for a variety of tasks in classification (e.g. Tamayo et al., 2016; Ivanov et al., 2021; Lucey et al., 2023; Luo et al., 2023), such as identifying pulsar candidates in gamma ray data (e.g. Mirabal et al., 2016; Wang et al., 2019) and separating stars, galaxies, and quasars in photometric data (e.g. Jin et al., 2019; Chang et al., 2021; Fu et al., 2021; Golob et al., 2021; Li et al., 2021; Nakoneczny et al., 2021; Hughes et al., 2022); and regression (as here, see also Calderon & Berlind, 2019; Hayden et al., 2022; Machado Poletti Valle et al., 2021; Dang et al., 2022; Andrae et al., 2023), including for predicting quasar redshifts from photometry (so called “photo-z”s, e.g. Jin et al., 2019; Nakoneczny et al., 2021; Kunsági-Máté et al., 2022). Some works include the use of SHAP values to analyze feature importances (e.g. Machado Poletti Valle et al., 2021; Heyl et al., 2023).

In this chapter, we first describe our methodology, including the Cloudy calculations we use to train and evaluate XGBoost models, the input features we use and their distributions and correlations, how we determine which combinations of radiation field features to use, and how we compare and evaluate XGBoost models. Next, we present results for trained models compared to each other and the interpolation table in Gnedin & Hollon (2012), on both the training data grid and an independent off-grid sample. Finally, we discuss our conclusions and potential future directions.

3.2 Data and methods

3.2.1 Model data input

In order to train machine learning models to predict cooling and heating functions, we use computations from the photoionization code Cloudy (Ferland et al., 1998). The training data used here is very similar to that used in Gnedin & Hollon (2012), with some additional information included.

For a given parcel of gas, we define cooling and heating functions through equation (3.1):

$$\left. \frac{dU}{dt} \right|_{\text{rad}} = n_b^2 [\Gamma(T, \dots) - \Lambda(T, \dots)], \quad (3.1)$$

where U is the thermal energy density of the gas, $n_b = n_H + 4n_{\text{He}} + \dots$ is the number of density of baryons (protons and neutrons), T is the gas temperature, and Γ, Λ are, respectively, the heating and cooling functions. The prefactor n_b^2 is included because of the importance of collisional two-body processes in radiative cooling and heating. In the limit of CIE, where there is no external ionizing radiation field on the gas, and accounting for two-body processes only, Λ and Γ are independent of n_b . However, even in CIE, multi-electron processes such as Auger ionization and dielectronic recombination can contribute to the cooling and heating functions when the gas contains metals (Ferland et al., 1998), introducing n_b dependence to Λ and Γ . Here, we consider the more general case of an external radiation field, and account for other relevant physical processes beyond two-body collisions. Hence, Λ and Γ also depend on n_b , metallicity, the incident radiation field, and may depend on other physical characteristics - this complex dependence is captured by the dots in equation (3.1).

The input data we use consists of Cloudy computations of Λ and Γ evaluated on a grid of values for the gas properties from Gnedin & Hollon (2012). The relevant gas properties used are temperature T , hydrogen number density n_H (since Cloudy uses the hydrogen number density n_H as its density parameter instead of the baryon density n_b , Ferland et al., 1998), metallicity Z , and 4 parameters describing a model for the incident radiation field:

$$J_\nu = J_0 \left[\frac{1}{1 + f_Q} s_\nu + \frac{f_Q}{1 + f_Q} x^{-\alpha} \right] e^{-\tau_\nu}, \quad (3.2)$$

where $x = h\nu/(1\text{Ry})$ is the photon energy in Rydbergs. The parameter J_0 in equation (3.2) accounts for the overall amplitude of the radiation field. The expression in square brackets contains two terms: a stellar spectrum s_ν and an AGN-like power law with slope α . The parameter f_Q quantifies the ratio of the stellar and AGN-like contributions to J_ν at 1 Ry.

Parameter	Values
$\log(T/\text{K})$	1, 1.1, 1.2, \dots , 9,
$\log(n_H/\text{cm}^{-3})$	-6, -5, -4, \dots , 6,
Z/Z_\odot	0, 0.1, 0.3, 1, 3,
$\log(J_0 \text{ cm}^{-3}/n_b/J_{\text{MW}})$	-5, -4.5, -4, \dots , 7,
$\log f_Q$	-3, -2.5, -2, \dots , 1,
$\log \tau_0$	-1, -0.5, 0, \dots , 3,
α	0, 0.5, 1, \dots , 3,

Table 3.1: The parameters describing the training data table and the values they take. Here, $J_{\text{MW}} = 10^6 \text{ photons cm}^{-2} \text{ s}^{-1} \text{ ster}^{-1} \text{ eV}^{-1}$ (Gnedin & Hollon, 2012). The $Z = 0$ case actually uses $Z = 10^{-4}Z_\odot$.

The stellar spectrum s_ν is given by:

$$s_\nu = \frac{1}{5.5} \begin{cases} 5.5 & x < 1, \\ x^{-1.8} & 1 < x < 2.5, \\ 0.4x^{-1.8} & 2.5 < x < 4, \\ 2 \times 10^{-3} \frac{x^3}{e^{x/1.4} - 1} & x > 4, \end{cases} \quad (3.3)$$

and is a fit to stellar spectra from the Starburst99 spectral synthesis library (Leitherer et al., 1999). The radiation field J_ν in equation (3.2) is also subject to attenuation due to neutral hydrogen and helium with optical depth τ_ν given by:

$$\tau_\nu = \frac{\tau_0}{\sigma_{\text{HI},0}} [0.76\sigma_{\text{HI}}(\nu) + 0.06\sigma_{\text{HeI}}(\nu)], \quad (3.4)$$

where τ_0 describes the overall optical depth, $\sigma_{\text{HI},0}$ is the photoionization cross section of neutral hydrogen at its ionization threshold, and $\sigma_{\text{HI}}(\nu), \sigma_{\text{HeI}}(\nu)$ are the photoionization cross sections for neutral hydrogen and helium, respectively, as functions of frequency. A loose justification for this parameterization choice is that a random place in the universe may be irradiated by both stellar and AGN radiation, and the strongest nearby source may happen to lie behind a sufficiently dense absorbing cloud. More complex scenarios with multiple nearby sources of approximately equal strengths can obviously be imagined, but they are unlikely to be common.

In total, the input data we use is fully described by 7 parameters: $T, n_H, Z, J_0, f_Q, \tau_0$, and α . We define our training data on a grid of these parameters, described in Table 3.1.

The 4 radiation field parameters (J_0, f_Q, τ_0 , and α) are only well-defined for a radiation field with the functional form given by equation (3.2), and cannot necessarily be determined

	HI	HeI	CVI	CaXX	CI	NaXI
Min	-6.91	-5.55	-9.02	-13.20	0.68	-11.12
Max	0.48	0.72	-1.06	-2.11	1.29	-1.59

Table 3.2: Minimum and maximum values of $\log(Q_j/Q_{\text{LW}})$ for some rates of interest attained in the training data table described in Table 3.1.

for a given radiation field from a galaxy formation simulation. Since photoionization of atoms in the gas play an important role in radiative cooling and heating, we choose to represent the radiation field via the specific photoionization rates Q_j defined as:

$$Q_j = \frac{c}{n_H} \int_0^\infty \sigma_j(\nu) n_\nu d\nu, \quad (3.5)$$

where n_ν is the number density of photons with frequency ν (Gnedin & Hollon, 2012). We explore 58 different such Q_j , including rates for all the ionization states of magnesium (MgI – MgXII) and iron (FeI – FeXXVI). We also include the photodissociation rate of molecular hydrogen Q_{LW} in order to sample the radiation field at UV energies below the threshold for hydrogen ionization. Since the subsequent ionization states of chemical elements usually have increasing ionization thresholds, ionization rates of all ionization states of a given element sample a wide range of photon energies and hence may serve as good proxies of a radiation field spectrum. These Q_j sample the radiation field across different wavelength ranges, and are determined both by the minimum photon frequency required for the particular ionization and the shape of the ionization cross section $\sigma_j(\nu)$ for higher frequencies. The minimum and maximum values attained given the radiation field parameter ranges from Table 3.1 are given for rates of particular interest for our analysis in Table 3.2.

3.2.2 XGBoost

For this work, we use the gradient-boosted tree algorithm XGBoost. XGBoost uses an ensemble of trees, where the trees are trained sequentially to predict the residual between the true value and the prediction of the previous trees (Chen & Guestrin, 2016). XGBoost is known to perform very well, including outperforming deep learning models, on tabular training data (Shwartz-Ziv & Armon, 2022; Grinsztajn et al., 2022). Given the tabular setup of our training data as described in section 3.2.1, XGBoost is an appropriate choice of machine learning model. XGBoost is also well-suited to the high-dimensional input data described in section 3.2.1 because the structure of regression trees means that the splitting at any node of any tree is determined by the value of only one feature. Furthermore, XGBoost

includes a hyperparameter which limits the fraction of the input features used for each tree (Chen & Guestrin, 2016). An additional reason for the choice of XGBoost over other machine learning models is the availability of tools to compute feature importances, as discussed in section 3.2.7 below. We train XGBoost regression models on a Graphics Processing Unit (GPU) using the tree method `gpu_hist`. Explanations of the XGBoost hyperparameters we varied from their default values can be found in section 3.6.1.

3.2.3 Training data preparation

The true cooling and heating functions Λ and Γ vary over several orders of magnitude on the training data, so we train our XGBoost models to predict $\log \Lambda$ and $\log \Gamma$ as functions of the gas temperature T , density n_H , metallicity Z , and photoionization rates Q_j .

Since we only have very sparse sampling in Z/Z_\odot , we do not use metallicity as an input feature for our XGBoost models. Instead, we train separate models at each of the 5 metallicity values in Table 3.1, and interpolate between model predictions (with the input features fixed) at intermediate metallicities. This interpolation is discussed further in section 3.2.8. We explored incorporating the metallicity Z as an additional XGBoost feature, but found that we could predict cooling and heating function more accurately by combining fixed-metallicity XGBoost models with interpolation in metallicity. This is somewhat unsurprising, since XGBoost models perform a piecewise constant interpolation in each input feature (Chen & Guestrin, 2016). With only 5 data points in metallicity, we can outperform this by choosing a manual fitting function with some knowledge of the expected behavior.

We utilize gas temperature $\log(T/\text{K})$ and density $\log(n_H/\text{cm}^{-3})$ as input features for all of our XGBoost models. To describe the radiation field J_ν , we use photoionization rates Q_j to capture how the incident radiation field impacts various elements in the gas. As suggested by the structure of the interpolation table in Gnedin & Hollon (2012), we expect to need at least 4 distinct Q_j to adequately describe the radiation field dependence. However, all rates Q_j scale linearly with the radiation field amplitude J_0 . Hence, to keep our features as uncorrelated as possible, we use the photodissociation rate of molecular hydrogen Q_{LW} as a reference rate. That is, we use Q_{LW} as an input to all of our XGBoost models, and scale all other photoionization rates as Q_j/Q_{LW} to remove the overall amplitude. The radiation field features we use are $\log(Q_{\text{LW}}/\text{cm}^3 \text{s}^{-1})$ and various sets of $\log(Q_j/Q_{\text{LW}})$ for $j \neq \text{LW}$.

As seen in Table 3.1 and Table 3.2, these features have large, and very different, ranges. For example $1 \leq \log(T/\text{K}) \leq 9$ and $-6 \leq \log(n_H/\text{cm}^{-3}) \leq 6$. To avoid this, we linearly rescale all features to be between 0 and 1 for the training data, using `MinMaxScaler()` from the `preprocessing` module of `scikit-learn` (Pedregosa et al., 2011). The resulting




Figure 3.1: Feature distributions of six scaled rates $\log(Q_j/Q_{\text{LW}})$. All have been linearly rescaled to the range $[0, 1]$, and their distributions normalized.

distributions for 6 $\log(Q_j/Q_{\text{LW}})$ features are shown in Figure 3.1.

3.2.4 PCA of rates

The 58 photoionization rates Q_j we consider have similar ionization thresholds and frequency dependence for their absorption cross-section, $\sigma_j(\nu)$. This means that many of the wavelength ranges sampled overlap with each other. As such, we should expect many of the Q_j to be highly correlated with each other. As discussed in section 3.2.3, we always use $\log(Q_{\text{LW}}/\text{cm}^3 \text{s}^{-1})$ as a feature to encode the overall radiation field amplitude J_0 , and consider the scaled rates $\log(Q_j/Q_{\text{LW}})$ for $j \neq \text{LW}$ as candidate features. There are still 57 of these scaled rates, with highly overlapping wavelength ranges and similar absorption cross sections. An example of this can be seen visually in the right-hand column of Figure 3.1, where the feature distributions of $\log(Q_j/Q_{\text{LW}})$ for $j = \text{CVI}$ and NaXI appear very similar by eye.

In order to measure how strongly correlated the photoionization rate features are, we eval-

uate the correlation matrix of $\log(Q_j/Q_{\text{LW}})$, $j \neq \text{LW}$ for the data described in section 3.2.1. We use the Spearman correlation coefficient, so that the correlations would be the same regardless of whether or not we take the logarithm of the scaled rates (since the logarithm is a monotonically increasing function). The correlation matrix (for the same scaled rate features $\log(Q_j/Q_{\text{LW}})$ as Figure 3.1) is shown in Figure 3.2. Note the high Spearman correlations of 0.95 or above for several pairs of features, and that the Spearman correlation is positive for all pairs.

To limit both the number of photoionization rate features we need to include and the correlations between them, we utilize tools from PCA. In particular, we determine the eigenvalues λ_i and normalized eigenvectors \vec{v}_i of the 57×57 Spearman correlation matrix described above. We find that the principal components of this correlation matrix depend non-trivially on almost all of these rates. That is, the correlation matrix is close to degenerate, and the principal components are poorly constrained. For this reason, we depart from PCA and instead rank each $\log(Q_j/Q_{\text{LW}})$ by a “relative importance” c_j defined as the sum of the magnitude of the j -th component of each eigenvector $\vec{v}_i(j)$, weighted by the corresponding eigenvalue:

$$c_j = \sum_i \lambda_i |\vec{v}_i(j)|. \quad (3.6)$$

As defined by equation (3.6), the relative importance for a given scaled rate is essentially a weighted average of the contribution of the scaled rate to each principal component of the correlation matrix, where the weights are the eigenvalue of each principal component. However, since we do not normalize the sum in equation (3.6), only comparisons between different c_j are physically meaningful, and not the magnitude of any individual c_j . In particular, a scaled rate with a larger relative importance c_j , on average, contributes more to more important principal components of the correlation matrix. Limiting the scaled rates we use as features for our XGBoost models to those with sufficiently large relative importance actually limits the number of scaled rates utilized in our analysis. In contrast, even selecting only the principal component with the largest eigenvalue would technically require the use of all 57 scaled rates, due to the near-degeneracy of the correlation matrix.

We show the relative importance, c_j , in Figure 3.3. We choose to select all rates with relative importance c_j larger than that of $j = \text{FeII}$ (i.e. left of the dashed line in Figure 3.3), using a natural break in the distribution of c_j . Note that all the photoionization rates for iron, and most magnesium rates (with the exception of MgI – III and MgV) have very low relative importance c_j and are excluded from the remainder of our analysis.

This leaves 21 scaled rates, which we include as features in a set of XGBoost models at each metallicity. We perform a feature importance analysis, discussed in section 3.2.7, on




Figure 3.2: Spearman correlation matrix of the same six scaled rate features $\log(Q_j/Q_{\text{LW}})$ as Figure 3.1.




Figure 3.3: The relative importance c_j defined by equation (3.6) for scaled rates $\log(Q_j/Q_{LW})$, in decreasing order. For simplicity, only the 35 scaled rates with the highest relative importance are shown. We use the 21 rates with the highest significance (those to the left of the vertical dashed line) in our analysis, as all other rates have similarly small significance.

these models to guide further downselection of the set of scaled rate features. We use these models as a baseline comparison for models incorporating fewer scaled rate features.

3.2.5 How to define an error?

To train and quantitatively evaluate the performance of our XGBoost models, we must define the error in a predicted cooling or heating function. We generally define this as:

$$\Delta \log \mathcal{F} = |(\log \mathcal{F})_{\text{pred}} - (\log \mathcal{F})_{\text{true}}|, \quad (3.7)$$

where $\mathcal{F} = \Lambda$ or Γ .

In order to train our XGBoost models, we must choose a ‘loss function’ to minimize on the training set (which is described below in section 3.2.6). This loss function compares the XGBoost predictions to the true values in the data table described in section 3.2.1, averaged over a set of predictions. For this purpose, we use the MSE:

$$\text{MSE} = \langle (\Delta \log \mathcal{F})^2 \rangle, \quad (3.8)$$

where the average is over all points in the specific data set in question.

However, we must also keep in mind that for any set of model predictions we will have a *distribution* of errors $\Delta \log \mathcal{F}$, which cannot be completely characterized by the MSE in general. Instead, we can consider the *cumulative error distribution function* $P(> \Delta \log \mathcal{F})$, which is defined by:

$$P(> \Delta \log \mathcal{F}) = \frac{\text{Number of points with error above } \Delta \log \mathcal{F}}{\text{Total number of points}}. \quad (3.9)$$

In particular, the cumulative error distribution function defined in equation (3.9) may include very large ‘catastrophic errors’. While it may have a minimal effect on the overall MSE, minimizing both the frequency and magnitude of these catastrophic errors is crucial for robust prediction of cooling and heating functions.

3.2.6 Model training

For each set of features, we must perform model validation (hyperparameter optimization), training, and testing for both the cooling function and heating function at $Z/Z_{\odot} = 0, 0.1, 0.3, 1, 3$. That is, we validate, train, and test 10 different XGBoost models for each set of features.

After calculating the photoionization rates, evaluating the cooling or heating function at the desired metallicity, and scaling the features as described above, we split the data into validation, training, and test sets. Note that we scale the features *before* splitting, so that the minimum and maximum values used to scale each feature could be in any of the testing, training, or validation data. To split the data, we use `train_test_split()` from the `model_selection` module of `scikit-learn`. We split 20% of the data as the testing set. Of the remaining 80% of the data, we use 10% for hyperparameter validation (i.e. 8% of the total data). Our hyperparameter tuning procedure is described in section 3.6.1. The remaining 72% of the data is used for model training using hyperparameters obtained from the validation step. These fractions are chosen so that the set used for model training contains the majority of the data, the testing set is large enough to be reasonably qualitatively representative of the feature distributions in the entire data set, and the separation of the validation set does not significantly reduce the amount of data available for model training. We do not expect small numerical changes to these fractions to significantly affect our results, as long as they satisfy these qualitative constraints. We use the same train-test-validation split for all models.

Using a suitable set of hyperparameters found as described in section 3.6.1, we then retrain XGBoost models on the 80% training set (but without splitting off any of this data for hyperparameter validation). We also retrain these XGBoost models on the entire grid for interpolating in Z , where we have independent testing data at intermediate values of Z .

3.2.7 Feature importances with SHAP values

To analyze which scaled photoionization rates have the most impact on the cooling and heating functions, we calculate SHAP values (Lundberg & Lee, 2017; Lundberg et al., 2018; Lundberg et al., 2020), using models with the 21 scaled rate features described in section 3.2.4 and retrained with optimal hyperparameters using 80% of the training data table, as described in section 3.2.6. We evaluate SHAP values for all features on 500 randomly sampled points from the 20% test set using the Python package `shap` (Lundberg & Lee, 2017; Lundberg et al., 2020).

The formula for computing the exact SHAP value for a particular model prediction and input feature comes from game theory, and describes the difference between the true model prediction and what would be expected without including the feature in question. These feature importances have several desirable properties, including exactly matching the model prediction in question, and always giving a value of 0 for a feature that does not affect the model prediction. However, the exact computation involves evaluating machine learning

$Z/Z_{\odot} = 10^{-4}$	$Z/Z_{\odot} = 0.1$	$Z/Z_{\odot} = 0.3$	$Z/Z_{\odot} = 1$	$Z/Z_{\odot} = 3$
MgII (3)	CI (4)	CI (4)	HI (4)	MgII (4)
CI (4)	HI (5)	HI (5)	MgII (5)	HI (5)
HeI (5)	MgII (6)	MgII (6)	CI (6)	CVI (6)

Table 3.3: The three most important scaled rates, by mean absolute SHAP value, for 500 randomly sampled test points on cooling function models with the 21 scaled rate features from section 3.2.4, retrained on 80% of the training data table using the optimized hyperparameters found via the procedure in section 3.2.6.

models analogous to the one we wish to explain *using every possible subset of features*. Since this is computationally unfeasible for complex models with many features, SHAP values must often be approximated in practice (Lundberg & Lee, 2017). The specific approximation for tree-based models implemented in the `shap` package which we use here is described in Lundberg et al. (2020).

To compare the importance of the various features, we evaluate the mean absolute SHAP value (since SHAP values can be positive or negative) across the 500 random test points. Ranking these (in decreasing order) gives a description of the most important features (for a given XGBoost model predicting cooling or heating, at a given metallicity). The 7 most important features by this ranking are shown for all models trained using the 21 rates from the PCA described in section 3.2.4 in Figure 3.4. For all these models, temperature $\log(T/K)$ and radiation field amplitude (described by $\log(Q_{\text{LW}}/\text{cm}^3 \text{s}^{-1})$) are the two most important features. For the cooling function at metallicities $Z/Z_{\odot} \geq 0.1$, density $\log(n_H/\text{cm}^{-3})$ is the third-most important feature by SHAP value, but it is less important at $Z/Z_{\odot} = 0$ and for the heating function at all metallicities. The three most important scaled photoionization rates, along with their SHAP value importance rank, are shown in Table 3.3 for the cooling function and Table 3.4 for the heating function. Several scaled rates (MgII, CI, HI, HeI) have high SHAP values at all metallicities for both cooling and heating, while CaXX has high SHAP values for heating at all metallicities except for $Z/Z_{\odot} = 0$. At the highest metallicities ($Z/Z_{\odot} = 1, 3$), the rate CVI reaches high SHAP values. Since the models in Figure 3.4 are trained at fixed metallicity values $Z/Z_{\odot} = \{0, 0.1, 0.3, 1, 3\}$, these trends hold only in the metallicity range $0 \leq Z/Z_{\odot} \leq 3$, and should not be extrapolated beyond this range. Within this metallicity range, we also do not determine the functional form of the dependence of SHAP values on metallicity, and instead describe only *trends* of increasing or decreasing SHAP values with metallicity.

To limit the data size of the XGBoost models, and to compare with the 4 photoionization rates used in the interpolation table in Gnedin & Hollon (2012) ($Q_{\text{LW}}, Q_{\text{HI}}, Q_{\text{HeI}}, Q_{\text{CVI}}$), we

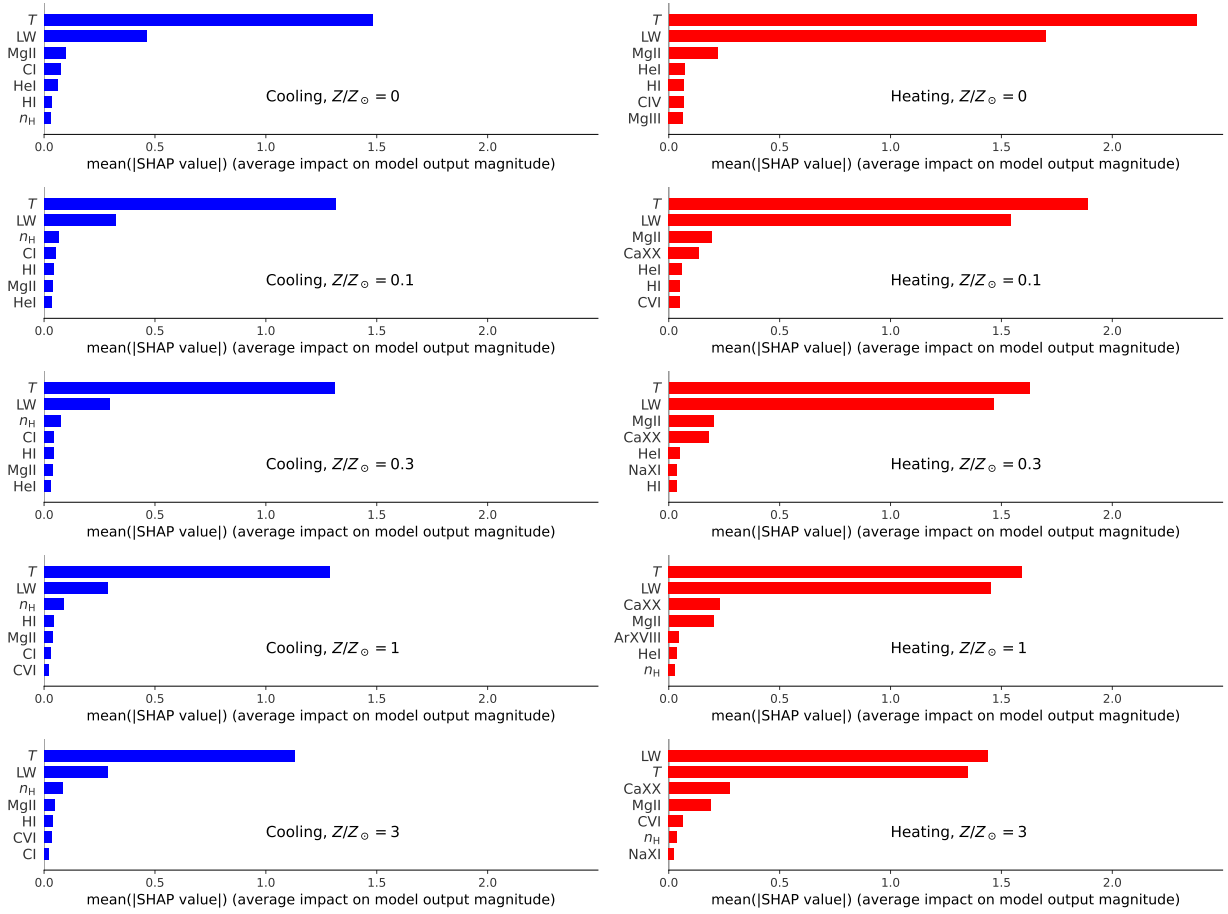


Figure 3.4: Mean absolute SHAP value for 500 randomly selected points in 20% of the training grid withheld from model training (see section 3.2.3) for models trained with 21 scaled rate features selected from the PCA analysis described in section 3.2.4. Cooling function models are shown in blue in the left column, with heating function models in red in the right column. The rows are for models at the 5 metallicity values in the training data. Only the 7 features with the largest mean absolute SHAP value are shown. Note that the full descriptions of the features are $\log(T/\text{K})$ for T , $\log(n_{\text{H}}/\text{cm}^{-3})$ for n_{H} , $\log(Q_{\text{LW}}/\text{cm}^3 \text{ s}^{-1})$ for LW , and $\log(Q_j/Q_{\text{LW}})$ otherwise.

$Z/Z_{\odot} = 10^{-4}$	$Z/Z_{\odot} = 0.1$	$Z/Z_{\odot} = 0.3$	$Z/Z_{\odot} = 1$	$Z/Z_{\odot} = 3$
MgII (3)	MgII (3)	MgII (3)	CaXX (3)	CaXX (3)
HeI (4)	CaXX (4)	CaXX (4)	MgII (4)	MgII (4)
HI (5)	HeI (5)	HeI (5)	ArXVIII (5)	CVI (5)

Table 3.4: Same as Table 3.4, but now for the heating function.

choose to use the 3 scaled rates with the highest SHAP values to train new models (as above, the feature $\log(Q_{\text{LW}}/\text{cm}^3 \text{ s}^{-1})$ is always used). For cooling and heating at each metallicity, we train models using the 3 most important scaled rate features from our SHAP value analysis in Figure 3.4, Table 3.3, and Table 3.4 using the procedure described in section 3.2.6. We also select several candidate sets of 3 scaled rates that appear frequently in Table 3.3, Table 3.4, or both.

3.2.8 Interpolation in metallicity

To approximate Λ and Γ at values of Z besides the sample points, we must interpolate between the predictions of the 5 fixed- Z models evaluated at the relevant feature values. Following Gnedin & Hollon (2012), we perform a quadratic fit to Λ and Γ at the five metallicity values $Z/Z_{\odot} = \{0, 0.1, 0.3, 1, 3\}$. However, we found that their quadratic fit sometimes yields unphysical *negative* predictions for Λ or Γ .

For this reason, we trained our XGBoost models to predict Λ and Γ at the five metallicities for which we have cooling and heating functions calculated with Cloudy, and perform a quadratic fit between XGBoost predictions that is constrained so as to never predict negative values on the metallicity domain of interest, i.e. $0 \leq Z/Z_{\odot} \leq 3$. Specifically, we fit a quadratic in metallicity:

$$f(Z|A, B, C) = A + B \left(\frac{Z}{Z_{\odot}} \right) + C \left(\frac{Z}{Z_{\odot}} \right)^2. \quad (3.10)$$

We find the best fit values of the parameters (A, B, C) by numerically minimizing the loss function:

$$L(A, B, C) = \chi^2(A, B, C) + g(A, B, C). \quad (3.11)$$

The χ^2 term is the error between \mathcal{F} (Γ or Λ) predicted by XGBoost, and the quadratic fit, calculated as:

$$\chi^2(A, B, C) = \sum_i [\mathcal{F}_{\text{pred}}(Z_i) - f(Z_i|A, B, C)]^2, \quad (3.12)$$

where the sum is over $Z_i/Z_{\odot} = \{0, 0.1, 0.3, 1, 3\}$. The constraint function $g(A, B, C)$ is defined by:

$$g(A, B, C) = \begin{cases} 0 & \begin{aligned} & f(Z|A, B, C) > \min_i \mathcal{F}_{\text{pred}}(Z_i) \text{ and} \\ & f(Z|A, B, C) < \max_i \mathcal{F}_{\text{pred}}(Z_i) \forall i \\ & \text{and } \forall 0 \leq Z/Z_{\odot} \leq 3, \end{aligned} \\ 10^{10} & \text{otherwise,} \end{cases} \quad (3.13)$$

where again $Z_i/Z_\odot = \{0, 0.1, 0.3, 1, 3\}$. This constraint ensures that the quadratic fit is bounded by the minimum and maximum XGBoost predictions at fixed metallicity. Our XGBoost models actually predict $\log \mathcal{F}$, so the values of \mathcal{F} predicted by XGBoost are always positive. Hence, the constraint $g(A, B, C)$ also ensures that the quadratic fit always yields positive values for the metallicity range where we make predictions ($0 \leq Z/Z_\odot \leq 3$).

To evaluate the performance of our XGBoost models combined with a quadratic fit in metallicity, we use a sample of off-grid data, which is randomly sampled in $\log(n_H/\text{cm}^{-3})$, $\log(J_0 \text{ cm}^{-3}/n_b/J_{\text{MW}})$, $\log f_Q$, $\log \tau_0$, and α for the same ranges as in Table 3.1, and in the range $[-3, \log 3]$ for $\log(Z/Z_\odot)$. We use 8666 such random samples, evaluated at the same 81 values of $\log(T/\text{K})$ as in Table 3.1, for a total of 701,946 evaluation points.

3.3 Results

3.3.1 Comparison on training data

We first compare the performance of our XGBoost models at fixed metallicity to the interpolation table in Gnedin & Hollon (2012) on the training data described in section 3.2.1. To make this comparison, we use XGBoost models trained on the entirety of this data set, as Gnedin & Hollon (2012) uses that same set to construct their interpolation table. To understand how accurately the XGBoost models can generalize outside of the training data, we also consider the performance of models trained on only 80% of the training grid (as described in section 3.2.6) on the 20% test set withheld from model training.

To quantify the performance of these models, we consider the cumulative distribution function of errors $P(\Delta > \log \mathcal{F})$ defined in equation (3.9). This is just the frequency of errors *at least as large as* $\Delta \log \mathcal{F}$. We also consider the mean squared error (MSE, defined in equation (3.8)).

At each fixed metallicity, we compare the interpolation table in Gnedin & Hollon (2012) with 3 different XGBoost models: a model trained using the same inputs as Gnedin & Hollon (2012), i.e. using the scaled photoionization rates $\log(Q_j/Q_{\text{LW}})$ for $j = \{\text{HI}, \text{HeI}, \text{CVI}\}$; a model trained using the 21 scaled rates identified from the PCA analysis described in section 3.2.4; and a model trained with the top 3 scaled rates from the SHAP value analysis of section 3.2.7.

The approach of Gnedin & Hollon (2012) involves a quadratic fit to Cloudy calculations at the 5 metallicity values in the training data table described in section 3.2.1, $Z/Z_\odot = \{0, 0.1, 0.3, 1, 3\}$. Here, we focus on $Z/Z_\odot = 0$ and 1 for comparison with our XGBoost models.

Cooling	GH12		XGBoost (HI, HeI, CVI)	
	Training	Test	Training	Test
$Z/Z_{\odot} = 0$	3.27×10^{-3}	(0)	1.49×10^{-5}	3.76×10^{-5}
$Z/Z_{\odot} = 0.1$	1.65×10^{-3}	(1.63×10^{-4})	1.36×10^{-5}	3.58×10^{-5}
$Z/Z_{\odot} = 0.3$	1.69×10^{-3}	(2.98×10^{-4})	1.28×10^{-5}	4.04×10^{-5}
$Z/Z_{\odot} = 1$	1.69×10^{-3}	(0)	2.29×10^{-5}	5.94×10^{-5}
$Z/Z_{\odot} = 3$	1.91×10^{-3}	(2.49×10^{-5})	3.95×10^{-5}	8.06×10^{-5}
	XGBoost (21 scaled rates)		XGBoost (Top 3 SHAP scaled rates)	
	Training	Test	Training	Test
$Z/Z_{\odot} = 0$	2.11×10^{-5}	4.19×10^{-5}	1.37×10^{-5}	3.60×10^{-5}
$Z/Z_{\odot} = 0.1$	1.97×10^{-5}	4.33×10^{-5}	2.86×10^{-5}	5.54×10^{-5}
$Z/Z_{\odot} = 0.3$	1.93×10^{-5}	4.56×10^{-5}	4.39×10^{-5}	8.39×10^{-5}
$Z/Z_{\odot} = 1$	2.51×10^{-5}	6.02×10^{-5}	9.80×10^{-5}	1.74×10^{-4}
$Z/Z_{\odot} = 3$	6.83×10^{-5}	1.11×10^{-4}	3.40×10^{-5}	7.84×10^{-5}

Table 3.5: MSEs for the cooling function models shown in Figure 3.5 and Figure 3.8. For the interpolation table in Gnedin & Hollon (2012), the MSE is on the entire training grid described in section 3.2.1. The values in parentheses are the fraction of ‘catastrophic errors’ on the training set where the prediction is negative (so $\Delta \log \mathcal{F} = \infty$). We remove these points from the MSE calculation. For the three XGBoost models, the MSE is shown for a model trained and evaluated on the entire training grid (‘training’ columns), and for a model trained on 80% of the training grid and evaluated on the 20% that is withheld (‘test’ columns).

Figure 3.5 shows the cumulative distribution function of training errors for the interpolation table in Gnedin & Hollon (2012) and the three types of XGBoost model described above for the cooling and heating function at $Z/Z_{\odot} = 0$ and 1 (for the same plots at all metallicities in the training data, see section 3.6.2). The MSEs for all these models are shown for the cooling function in Table 3.5 and heating function in Table 3.6.

From Figure 3.5, Figure 3.8, Table 3.5, and Table 3.6, we see that all 3 types of XGBoost model greatly outperform the interpolation table in Gnedin & Hollon (2012) in all cases, with lower MSEs and a smaller cumulative error distribution function for all values shown on the plots ($10^{-2} < \Delta \log \mathcal{F} < 2$). The MSE on the 20% of the training table withheld from model training is always larger than that for analogous models trained on the entire table by a factor of order unity. However, these MSE values are still always smaller than for the interpolation table in Gnedin & Hollon (2012), and the cumulative error distribution function is similarly lower at the same error values.

The comparisons between the 3 different types of XGBoost model are not so clear. In particular, we can see that the cumulative error distribution functions for the XGBoost models (on the entire training table) sometimes intersect in Figure 3.5 and Figure 3.8 (see

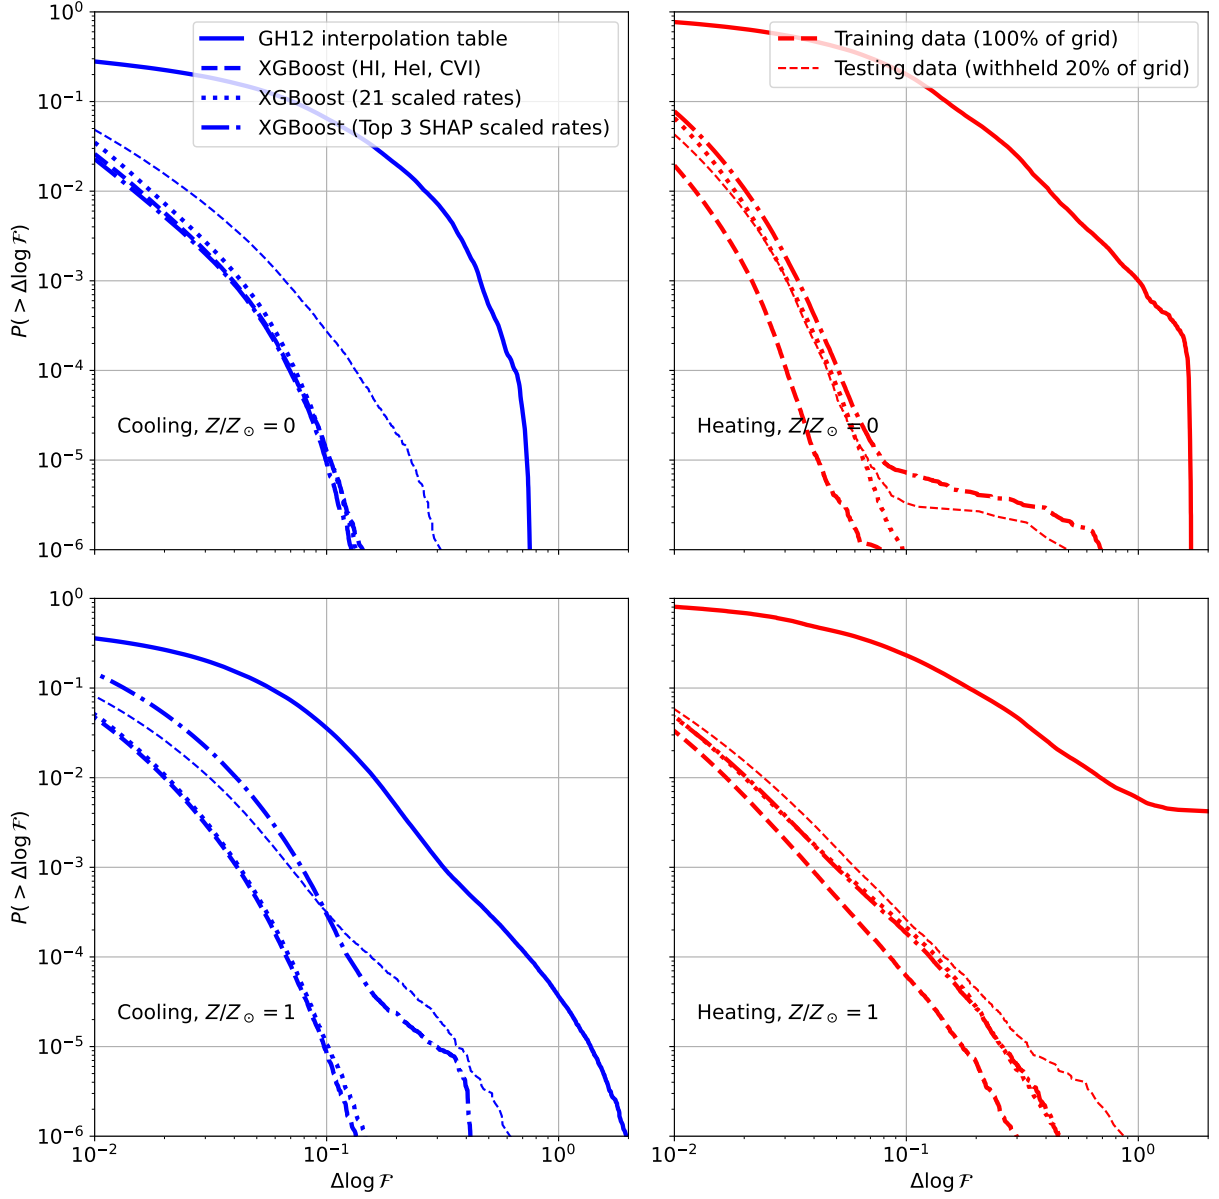


Figure 3.5: Cumulative distribution function of errors $\Delta \log \mathcal{F}$ (see equation (3.9)) on the training data (described in section 3.2.1) for the cooling function Λ (left column, in blue) and heating function Γ (right column, in red) at fixed metallicity $Z = 0$ (top row) and $Z = Z_{\odot}$ (bottom row), for the interpolation table in Gnedin & Hollon (2012) (solid lines) and XGBoost models trained using the same scaled photoionization rates $j = \text{HI, HeI, CVI}$ (dashed lines), 21 scaled rates from the PCA analysis described in section 3.2.4 (dotted lines), and the top 3 scaled rates from the SHAP value analysis, described in section 3.2.7 (dashed-dotted lines). For the “XGBoost (HI, HeI, CVI)” model, we also include the cumulative error distribution function on 20% of the training grid withheld from model training for otherwise identical models (thin dashed lines).

Heating	GH12		XGBoost (HI, HeI, CVI)	
	Training	Test	Training	Test
$Z/Z_{\odot} = 0$	1.27×10^{-2}	(0)	1.10×10^{-5}	2.08×10^{-5}
$Z/Z_{\odot} = 0.1$	3.26×10^{-2}	(0)	4.51×10^{-5}	6.74×10^{-5}
$Z/Z_{\odot} = 0.3$	2.19×10^{-2}	(0)	3.22×10^{-5}	5.79×10^{-5}
$Z/Z_{\odot} = 1$	9.51×10^{-2}	(1.84×10^{-3})	1.96×10^{-5}	4.38×10^{-5}
$Z/Z_{\odot} = 3$	2.05×10^{-2}	(1.05×10^{-5})	3.28×10^{-5}	5.82×10^{-5}
	XGBoost (21 scaled rates)		XGBoost (Top 3 SHAP scaled rates)	
	Training	Test	Training	Test
$Z/Z_{\odot} = 0$	3.03×10^{-5}	4.05×10^{-5}	3.41×10^{-5}	5.03×10^{-5}
$Z/Z_{\odot} = 0.1$	2.89×10^{-5}	6.13×10^{-5}	4.45×10^{-5}	8.24×10^{-5}
$Z/Z_{\odot} = 0.3$	3.95×10^{-5}	7.44×10^{-5}	3.09×10^{-5}	6.13×10^{-5}
$Z/Z_{\odot} = 1$	3.28×10^{-5}	6.12×10^{-5}	3.09×10^{-5}	5.75×10^{-5}
$Z/Z_{\odot} = 3$	1.68×10^{-5}	4.32×10^{-5}	2.16×10^{-5}	4.27×10^{-5}

Table 3.6: Same as Table 3.5, but now for heating function models.

the cooling function at $Z/Z_{\odot} = 0$ in the upper left of both figures), meaning that we cannot unambiguously identify a single ‘best-performing’ XGBoost model.

3.3.2 Comparison on off-grid data

Next, we compare the performance of our XGBoost models combined with the interpolation in metallicity described in section 3.2.8 to the interpolation table in Gnedin & Hollon (2012) on the sample of off-grid data (also described in section 3.2.8). Note that we are now evaluating the XGBoost models on points not seen in model training, and introducing interpolation between individual XGBoost predictions. Both of these factors decrease model performance. Since the rate of change of the energy density of a gas cloud is proportional to the heating function minus the cooling function, as seen equation (3.1), the physical evolution will usually be dominated by whichever of the cooling or heating function is larger. Hence, we will also consider the error in the maximum of the cooling and heating function at each point in the evaluation data.

For the interpolation in metallicity, we must select a common set of scaled rate features to use in the XGBoost models at each Z value. Here, we consider two different sets of 3 scaled rates, $\{\text{CaXX}, \text{HI}, \text{CVI}\}$ and $\{\text{Cl}, \text{HI}, \text{NaXI}\}$, in addition to the set $\{\text{HI}, \text{HeI}, \text{CVI}\}$ that makes XGBoost take the same inputs as the interpolation table in Gnedin & Hollon (2012). The choice of these sets of scaled rates is motivated by considering rates which have high SHAP value importance across the range of metallicity values, and by the heuristic of needing a photoionization rate with X-ray threshold to capture the quasar part of the

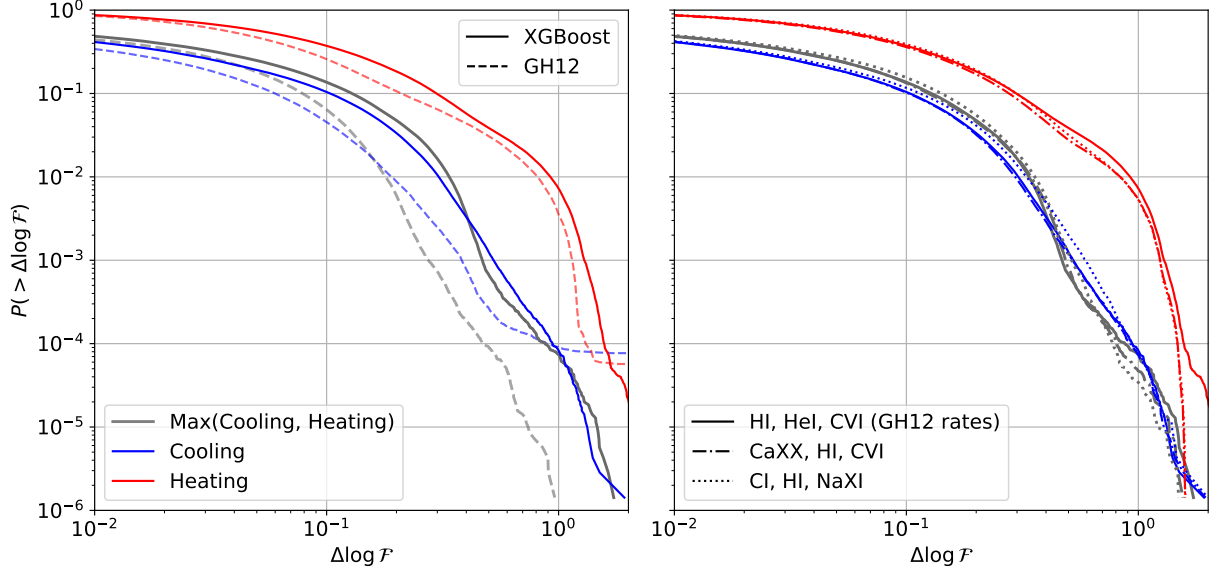


Figure 3.6: Cumulative distribution function of errors $\Delta \log \mathcal{F}$ on the evaluation data described in section 3.2.8. Left panel: Comparison of results from the interpolation table in Gnedin & Hollon (2012) and the constrained quadratic fit described in section 3.2.8 between XGBoost models trained with the same scaled rates (HI, HeI, CVI). Right panel: Comparison of the same constrained quadratic fit between XGBoost models trained with three different sets of scaled rates.

ionizing radiation field (parameterized by the power-law slope α).

In the left panel of Figure 3.6, we plot the cumulative error distribution function on the off-grid evaluation data described in section 3.2.8 for the interpolation table in Gnedin & Hollon (2012) and our constrained quadratic interpolation between XGBoost models using the same inputs. In the right panel, we compare the cumulative error distribution function for our constrained quadratic interpolation between XGBoost models with the three different sets of scaled rates described above. In Table 3.7, we show the MSEs corresponding to each of the cumulative error distribution functions in Figure 3.6.

From the left panel of Figure 3.6, we see that our XGBoost models plus constrained quadratic interpolation can reduce the frequency of the largest cooling and heating function errors compared to the interpolation table in Gnedin & Hollon (2012), using the same inputs. From Table 3.7, we see that these XGBoost models (using the scaled rates {HI, HeI, CVI}) yields comparable MSE values to the interpolation table in Gnedin & Hollon (2012), while (by construction) completely removing the occurrence of ‘catastrophic errors’ for cooling or heating. However, this improvement *does not* affect the prediction of the maximum of the cooling or heating functions. Indeed, both the cumulative error distribution function at each

	GH12	XGBoost (HI, HeI, CVI)
Max(C, H)	2.29×10^{-3} (0)	6.89×10^{-3}
Cooling (C)	2.02×10^{-3} (7.55×10^{-5})	5.26×10^{-3}
Heating (H)	3.00×10^{-2} (5.56×10^{-5})	4.49×10^{-2}
	XGBoost (CaXX, HI, CVI)	XGBoost (CI, HI, NaXI)
Max(C, H)	6.84×10^{-3}	7.94×10^{-3}
Cooling (C)	5.07×10^{-3}	6.32×10^{-3}
Heating (H)	3.71×10^{-2}	4.06×10^{-2}

Table 3.7: MSEs on the off-grid data set described in section 3.2.8 for the models in Figure 3.6. For the GH12 column, the values in parentheses are the fraction of ‘catastrophic errors’ on the evaluation set where the prediction is negative (so $\Delta \log \mathcal{F} = \infty$). We remove these points from the MSE calculation.

error value shown ($10^{-2} < \Delta \log \mathcal{F} < 2$) and the overall MSE are higher for XGBoost than Gnedin & Hollon (2012).

The performance of all three XGBoost calculations in the right panel of Figure 3.6 are similar. Neither rate set in this figure produces significant improvement in the performance for the maximum of the cooling or heating function compared to {HI, HeI, CVI}, in either MSE or the cumulative error distribution function. However, the scaled rate set {CaXX, HI, CVI} results in lower MSE, and slightly lower cumulative error distribution in the high-error tail for cooling and heating individually. Hence, we are able to produce some improvement in overall performance with judicious choices of scaled rate features.

Comparison of Figure 3.5 and Figure 3.6 suggests that, for the training data grid used in this project and Gnedin & Hollon (2012), there is a bottleneck in model performance on off-grid data due to the interpolation in Z . This is perhaps unsurprising, given that we have only 5 values of Z to interpolate over, and that we should expect the true dependence of Γ and Λ on Z to be more complicated than a quadratic.

3.4 Conclusions and discussion

In this chapter, we train XGBoost models to predict cooling and heating functions at the fixed metallicities $Z/Z_{\odot} = \{0, 0.1, 0.3, 1, 3\}$, using the same pre-computed Cloudy tables as Gnedin & Hollon (2012). We perform a constrained quadratic fit (with the input features held constant) between these XGBoost models to predict cooling and heating functions at arbitrary metallicity. Using model features selected from a PCA and SHAP value feature importances, we compare both fixed metallicity performance and performance at arbitrary metallicity with results from the interpolation tables in Gnedin & Hollon (2012) in order

to assess the potential of XGBoost as an alternative approach for evaluating cooling and heating functions in the presence of generalized radiation fields. We summarize our findings as follows:

- Using the same photoionization rates as Gnedin & Hollon (2012) ($j = \text{LW, HI, HeI, CVI}$) our XGBoost models are able to significantly reduce the frequencies of errors $0.01 < \Delta \log \mathcal{F} < 2$, as well as the mean squared error, at all 5 fixed metallicities. This can be seen in the comparison between solid and dashed thick lines in Figure 3.5 and Figure 3.8.
- Using an eigenvalue and eigenvector analysis of the correlation matrix of scaled photoionization rates, we identified 21 particularly significant scaled photoionization rates (see Figure 3.3).
- For models trained using these 21 rates, we identified the most important features using SHAP values (see Figure 3.4). Many rates have consistently high importance for both cooling and heating across all five metallicity values, especially MgII, CI, HI, and HeI. The rate CaXX has high importance for the heating function, but not cooling. The rate CVI has low importance at low metallicities, but emerges as one of the more important rates for the highest metallicity $Z/Z_{\odot} = 3$ for both the cooling and heating functions. We expect trends in mean SHAP values with metallicity to be robust for the metallicity range we sample, $0 \leq Z/Z_{\odot} \leq 3$. Higher mean SHAP values for a given rate signifies that the wavelengths sampled by that photoionization rate affect cooling or heating more strongly at that metallicity.
- The shape of the cumulative error distribution functions for various fixed Z models (see the dashed, dotted, and dot-dashed lines in Figure 3.5) cannot be adequately described by a single parameter. While the MSE (see the corresponding values in Table 3.5 and Table 3.6) describes the overall normalization of these curves, they also can differ in overall shape, particularly in how steeply the cumulative error distribution function falls off at high errors. This generally prevents determination of a single clearly ‘best-performing’ model.
- Using the same scaled rates as Gnedin & Hollon (2012), we are able to reduce the frequency of the largest errors $\log \mathcal{F} \geq 1$ for cooling or heating individually. We were able to achieve further reduction with other rate sets suggested by a combination of SHAP values and physical heuristics, such as $\{\text{CaXX, HI, CVI}\}$ (see Figure 3.6).

The reduction in the frequency of cooling or heating function errors $\Delta \log \mathcal{F} \geq 1$ that we achieve at arbitrary metallicity (see Figure 3.6 and Table 3.7) is much more modest than that seen for our individual fixed metallicity XGBoost models (see Figure 3.5, Table 3.5, and Table 3.6). This is likely due to the small number of samples we have to perform a fit in metallicity (the five Z/Z_{\odot} values in Table 3.1).

Our metallicity fit, seen in equation (3.10), assumes that the dependence of cooling and heating functions on metallicity (with all other parameters fixed) is quadratic. If the only processes involved were two-body collisions, then this assumption would be exactly correct *for the cooling and heating functions computed by Cloudy*. However, the Cloudy training data includes additional processes involving multiple electrons (Ferland et al., 1998). Furthermore, even small errors in our XGBoost predictions at fixed metallicity imply that the dependence of these predictions on metallicity will not exactly match the metallicity dependence of the Cloudy calculations they approximate. Also, we introduced an additional unphysical constraint on our quadratic fits in order to prevent negative predicted cooling and heating functions. These all suggest that a different, more complex metallicity fitting function is needed to accurately capture the metallicity dependence.

Hence, our findings suggest future work running Cloudy at intermediate metallicity values between the five used here, in order to implement a more flexible and accurate fit in metallicity and better understand the metallicity dependence of cooling and heating functions. This additional Cloudy data could also enable further exploration of how the importances of various radiation field parameters changes with metallicity.

3.5 Acknowledgements

This chapter has been co-authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. This research was also supported in part through computational resources and services provided by Advanced Research Computing (ARC), a division of Information and Technology Services (ITS) at the University of Michigan, Ann Arbor, in particular the Great Lakes cluster and the U-M Research Computing Package.

I acknowledge funding provided by the Leinweber Graduate Fellowship at the University of Michigan; the National Aeronautics and Space Administration (NASA), under award number 80NSSC20M0124, Michigan Space Grant Consortium (MSGC); and the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists, Office of Science Graduate Student Research (SCGSR) program. The SCGSR program is administered by the Oak Ridge Institute for Science and Education for the DOE under

contract number DE-SC0014664. Co-author CA acknowledges support from the Leinweber Center for Theoretical Physics at the University of Michigan and Department of Energy grant DE-SC009193. The authors would like to thank the anonymous reviewer for their helpful and constructive comments.

The XGBoost pipeline is maintained on GitHub at [this link](#). The Cloudy tables used for model training and evaluation, and our trained XGBoost models, can be made available upon request.

3.6 Appendix

3.6.1 Hyperparameter tuning

There are many hyperparameters that can be tuned to adjust how precisely an XGBoost model can be fit to its training data¹. For this chapter, we consider non-default values for only 7 of the available hyperparameters:

- **max_depth**: the maximum depth of each tree. Default: 6, must be a positive integer.
- **min_child_weight**: the minimum weight required to split a node, where smaller means a more complex model. Default: 1, must be non-negative.
- **subsample**: the fraction of training data rows used in each tree. Default: 1, must be between 0 and 1 (inclusive).
- **colsample_bytree**: the fraction of features used in each tree. Default: 1, must be between 0 and 1 (inclusive).
- **gamma**: the cost for adding each new node in the objective function. Default: 0, must be non-negative.
- **eta**: the learning rate. Default: 0.3, must be positive.
- **n_estimators**: the number of trees in the ensemble (results of each tree are added together). Since the trees are trained sequentially, this is equivalent to the number of training rounds. Default: 100, must be a positive integer.

To find the ‘best’ hyperparameters using the validation set (described in section 3.2.6), we use 5-fold cross validation, as implemented in `scikit-learn` and the Bayesian optimization

¹Detailed descriptions of these hyperparameters can be found in the XGBoost documentation: <https://xgboost.readthedocs.io/en/stable/index.html>

scheme `BayesSearchCV` from `scikit-optimize` (Head et al., 2021). We set up the search space as follows.

- `min_child_weight`: Real, log-uniform prior between 0.1 and 2.
- `subsample`: Real, uniform prior between 0.6 and 1.
- `colsample_bytree`: Real, uniform prior between 0.6 and 1.
- `gamma`: Real, uniform prior between 0 and 1.
- `eta`: Real, log-uniform prior between 0.03 and 0.3.

We use the negative MSE as the scoring method for the Bayesian hyperparameter search.

In principle, the error on the training set (defined in section 3.2.6) can be made arbitrary close to 0 with an overfitted model. However, we do not necessarily need to find the lowest possible error, but rather are looking for a sufficiently well-performing set of hyperparameters. For these reasons, rather than incorporating them into the Bayesian search, we perform a manual grid search in `max_depth` and `n_estimators` (the two hyperparameters which we found to have the strongest effect on model performance), where we explore the Cartesian product of:

$$\begin{aligned} \text{max_depth} &= \{8, 12, 16, 20\}, \\ \text{n_estimators} &= \{1, 10, 100, 300, 1000\}. \end{aligned}$$

At each of these 40 combinations of `{max_depth, n_estimators}`, we perform the Bayesian optimization of the other 5 hyperparameters as described above. We compare both the mean squared error and the cumulative error distribution evaluated at $\Delta \log \mathcal{F} = 0.3$, $P(\Delta \log \mathcal{F} > 0.3)$ (the frequency of errors $\Delta \log \mathcal{F} > 0.3$) on the test set, to constrain the frequency of large errors. The latter quantity is more meaningful for discriminating between similar well-performing models, because the mean squared error is dominated by the more frequent smaller errors and is not very sensitive to changes in the magnitude of the largest, most infrequent errors. As an additional comparison, we evaluate the time needed to train each of the 40 models. In general, pairs `{max_depth, n_estimators}` which produce lower MSE and/or frequency of large errors on the test set have longer training times.

An example of such a grid search, for a cooling function model trained using the 21 scaled rates selected by the PCA described in section 3.2.4 at $Z/Z_{\odot} = 1$ is shown in Fig. 3.7. From this grid search, we selected `max_depth = 20, n_estimators = 100` as our choice, because it has the shortest training time (middle panel) among choices of `max_depth` and `n_estimators`




Figure 3.7: Results from a grid search in `max_depth` and `n_estimators` for cooling function model trained with the 21 scaled rates from the PCA of section 3.2.4 at $Z/Z_{\odot} = 1$. The three panels show heatmaps of the MSE in $\log \mathcal{F}$ on the test set (a randomly select 20% of the training data table described in section 3.2.1 (left), the training time in seconds (middle), and the frequency of errors $\Delta \log \mathcal{F}$ larger than 0.3 on the test set (right). The heatmaps are all on a logarithmic scale.

with similar mean squared error and frequency of large errors (left- and right- most panels in Fig. 3.7 to the best cases in the grid search.

We find that the results of such grid searches, and in particular the choice of `max_depth` = 20, `n_estimators` = 100, are consistent across models trained using the 21 scaled rates identified in the PCA analysis described in section 3.2.4 for both cooling and heating functions at $Z/Z_{\odot} = 0$ and 1, and using both the 21 scaled rates from section 3.2.4 and only the top 3 scaled rates from the feature importance analysis discussed in section 3.2.7.

After confirming that the choice of `max_depth` = 20, `n_estimators` = 100 is robust across metallicity, cooling vs. heating functions, and for models with very different numbers of scaled rate features, we adopted `max_depth`= 20, `n_estimators`= 100 for all other models, without performing further grid searches in `max_depth` and `n_estimators` (but still performing a Bayesian optimization of the other 5 hyperparameters we consider).

3.6.2 Comparison on training data for all metallicities

In Figure 3.8, we extend the comparison in Figure 3.5 to all metallicities in the training data table (described in Table 3.1): $Z/Z_{\odot} = \{0, 0.1, 0.3, 1, 3\}$.

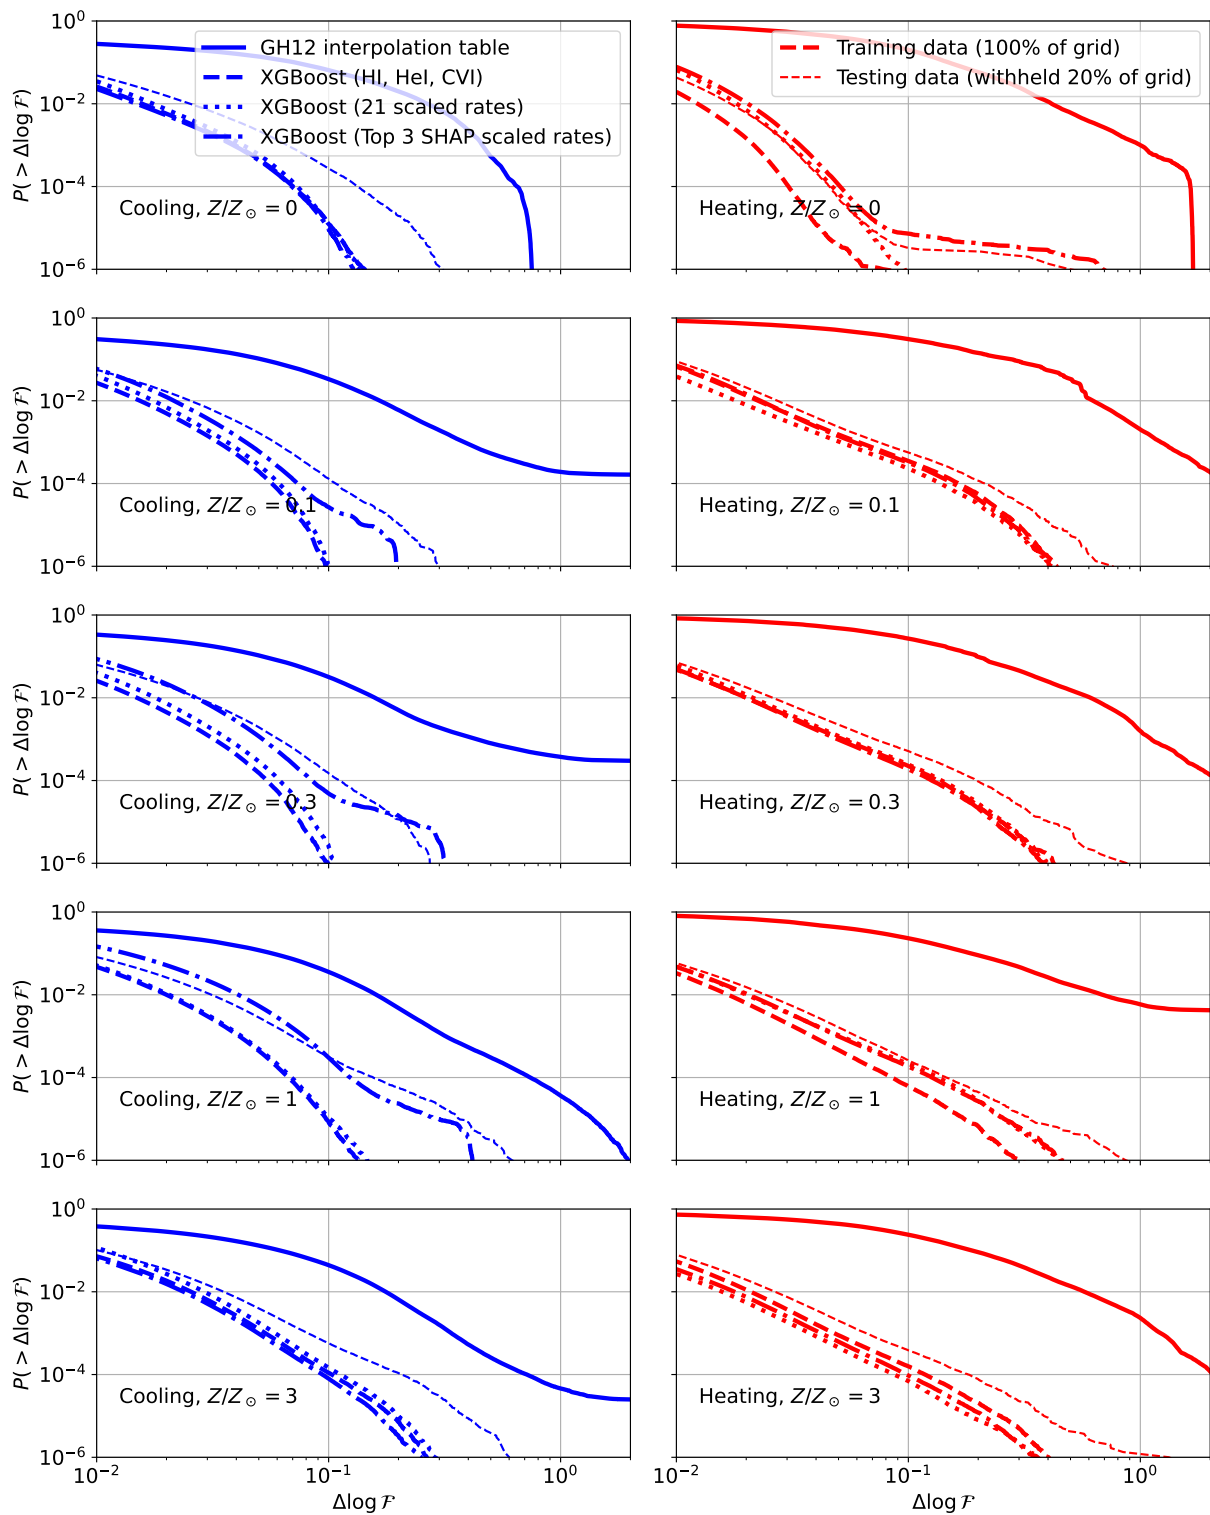


Figure 3.8: Same as Figure 3.5, but now for all 5 metallicity values in the training data table (described in Table 3.1).

CHAPTER 4

On the Minimum Number of Radiation Field Parameters to Specify Gas Cooling and Heating Functions

4.1 Motivation

The various feedback processes through which stars and active galactic nuclei deposit energy, momentum, and gas enriched with metals and dust into the CGM and back into the ISM from which they formed, are crucial parts of galaxy formation and evolution. One particular mode of feedback is direct photoionization and photoheating of gas by the local Interstellar Radiation Field (ISRF), which modifies the thermodynamics of the gas (see, e.g., the reviews of Benson, 2010; Naab & Ostriker, 2017).

Cooling and heating functions describe how the energy density of a cloud of gas changes in time due to radiative processes (e.g. Cox & Tucker, 1969; Sutherland & Dopita, 1993), and so will be affected by photoionization and photoheating feedback. Cooling and heating functions are critical to understanding galaxy formation and evolution (e.g., the reviews of Benson, 2010; Naab & Ostriker, 2017). The energy density of a gas cloud can also change due to mechanical expansion or compression. In an idealized scenario without star formation, the evolution of the size of the cloud is determined by the balance between (net) radiative cooling and gravitational compression (e.g. Lee, 2024). This same balance also occurs in the process of star formation. Including the effects of photoionization feedback in the cooling function has been shown to change the stellar mass formed in a simulated Milky Way-like galaxy (Kannan et al., 2014).

The incident radiation field in galaxies, which affects the cooling and heating functions, can have contributions from a variety of sources. External galaxies and quasars combine to produce a mildly undulating (but time-varying) UV background of ionizing radiation. Theoretical models for (the spatial average of) this background include Faucher-Giguère

et al. (2009) and Haardt & Madau (2012). The interstellar radiation field within the Milky Way depends on location through the spatial positions of sources (including stars) and absorbers (such as gas and dust, Popescu et al., 2017).

In a cosmological or galaxy evolution simulation with radiative transfer, we can calculate the radiation field in various energy bands. Some recent works have focused on modeling the radiation field in simulations across a wide range of wavelengths (e.g. Romero et al., 2023), or in specific wavelength ranges such as polycyclic aromatic hydrocarbon emission in the infrared (e.g., Narayanan et al., 2023); the ultraviolet Lyman and Werner bands for the photodissociation of molecular hydrogen (e.g., Incatasciato et al., 2023); and ultraviolet bands containing the photoionization thresholds for neutral hydrogen (HI), neutral helium (HeI), and singly ionized helium (HeII), among others (e.g., Kannan et al., 2020; Baumschlager et al., 2024). Baumschlager et al. (2024) uses narrow energy bins of width 0.1 eV at the edges of the HI, HeI, and HeII-ionizing bins, in order to model the radiation field spectrum in each of those bins as a power law (as opposed to a constant).

Given an arbitrary incident radiation field, we can evaluate a network of differential equations for the creation and destruction of each relevant ion by chemical reactions, collisions, and absorption and emission of radiation. The net emission and absorption of photons at each relevant energy determine the cooling and heating functions. This is implemented in both photoionization codes like Cloudy (Ferland et al., 1998), and chemical networks such as GRACKLE (Smith et al., 2017) and KROME (Grassi et al., 2014). These networks can have varying degrees of complexity (i.e. numbers of ions and reactions incorporated). Including more ions makes the calculations more exact at the cost of increased computational complexity. Including fewer ions decreases the accuracy, but can make the networks fast enough to utilize in a hydrodynamic simulation (e.g. Grassi et al., 2014; Smith et al., 2017). Another approach to simplifying the computational complexity of large chemical networks is to emulate the network with a machine learning algorithm (e.g. Branca & Pallottini, 2023, 2024).

In general, approximating cooling and heating functions subject to an incident radiation field requires making some assumptions to simplify the radiation field to a finite number of parameters. The simplest possible radiation field, $J_\nu = 0$, has no parameters. In this limit, known as CIE, no ionizing radiation is incident on the gas, and the cooling and heating functions depend only on properties of the gas (density, temperature, and metallicity, e.g., Cox & Tucker, 1969; Sutherland & Dopita, 1993). Since the extragalactic UV background radiation depends only on redshift, the cooling and heating functions of a gas irradiated by this background depend on redshift, in addition to gas properties (e.g., Kravtsov, 2003; Wiersma et al., 2009). Gnedin & Hollon (2012) assume a 4-parameter model for the local

radiation field: an arbitrary mixture of a synthesized stellar spectrum and a quasar-like power law, optionally attenuated by an additional optical depth due to the local absorption by H I and He I (this spectrum is discussed in more detail in section 4.2.1 and 4.2.2), and tabulate the resulting cooling and heating functions using photoionization rates calculated from that spectrum. Another possibility is to use the average intensity of the radiation field in various energy bins as the parameters, as is done in Branca & Pallottini (2024). Rather than interpolating on a grid of parameter values, Galligan et al. (2019) uses gas and radiation field properties from a zoom-in galaxy simulation with radiative transfer to run Cloudy models to calculate the cooling and heating functions, and trains a machine learning algorithm to interpolate the results.

In this chapter, we utilize a machine learning framework to approximate cooling and heating functions that we developed in Chapter 3. We parameterize the radiation field with its average intensity in various energy bins instead of the photoionization rates used in Chapter 3. Using radiation field bins is useful because many radiative transfer simulations calculate the radiation field in such bins (see above). These bins also summarize the radiation field in a different way than photoionization rates. In particular, they do not include the frequency dependence of photoionization cross sections. So, comparing the performance of machine learning models trained using binned radiation field intensities with those trained using photoionization rates allows us to gain more insight into what aspects of the incident radiation field are most important to cooling and heating functions.

Following the approach of Chapter 3, we use a training table from Gnedin & Hollon (2012), which consists of cooling and heating functions calculated with Cloudy across a range of temperatures, densities, metallicities, and 4 radiation field parameters (see Table 4.1). We apply the machine learning algorithm XGBoost (Chen & Guestrin, 2016) to predict the cooling and heating functions *at fixed metallicity*. XGBoost is an example of a gradient-boosted tree algorithm, where an ensemble of trees are trained *sequentially*. Each new tree is trained to predict the difference between the true value and the prediction from the previously-trained trees. The final prediction is then the sum of the predictions from each tree. At each tree node, the branches split based on the value of one feature (Chen & Guestrin, 2016). Gradient-boosted tree algorithms have been shown to outperform deep-learning methods on problems where the training data is in tabular form (Shwartz-Ziv & Armon, 2022; Grinsztajn et al., 2022), such as the training data used here.

XGBoost also interfaces easily with SHAP values (Lundberg & Lee, 2017; Lundberg et al., 2018; Lundberg et al., 2020), which we use to study the relative importance of various radiation field bins. For any given model prediction, the SHAP value for each feature is the difference between the actual model prediction and an expected prediction ignoring the

feature in question. This definition ensures that the SHAP values for every feature add up to the exact point prediction in question, and gives 0 for any feature that does not affect the model prediction. But, training models on all subsets of features is impractical for models with a significant number of features (Lundberg & Lee, 2017), leading to the development of tools to *approximate* SHAP values for various kinds of models (Lundberg et al., 2018; Lundberg et al., 2020).

In this chapter we present our methodology in section 4.2, including our training data, how we choose radiation field bins, our machine learning pipeline, and using SHAP values to compare the importance of various model features. Then, in section 4.3, we evaluate the performance of the various machine learning models trained in this chapter. We conclude with a summary and discussion in section 4.4.

4.2 Methodology

4.2.1 The training grid

Similarly to Chapter 4, we train our machine learning models to predict atomic gas cooling and heating functions computed with the photoionization code Cloudy (Ferland et al., 1998), using the same data table as Gnedin & Hollon (2012). We define the gas heating function Γ and cooling function Λ through the equation

$$\left. \frac{dU}{dt} \right|_{\text{rad}} = n_b^2 [\Gamma(T, n_b, Z, J_\nu) - \Lambda(T, n_b, Z, J_\nu)], \quad (4.1)$$

where U is the thermal energy density of the gas, Z is its metallicity, $n_b = n_{\text{H}} + 4n_{\text{He}} + \dots$ is its baryon number density, and J_ν is the specific intensity of the incident radiation field. The factor of n_b^2 accounts for the expected density dependence in the limit of collisional ionization equilibrium, and with only two-body collisions. However, in the more general case considered here, the cooling and heating functions will depend on n_b (and also Z and J_ν).

We assume a radiation field model with specific intensity given by

$$J_\nu = J_0 \left[\frac{1}{1 + f_Q} s_\nu + \frac{f_Q}{1 + f_Q} x^{-\alpha} \right] e^{-\tau_\nu}, \quad (4.2)$$

where the overall amplitude J_0 is a free parameter, and $x = h\nu/(1 \text{ Ry})$ is the photon energy in Rydbergs (1 Ry = 13.6 eV, the ionization threshold for neutral hydrogen HI). The first term in the square brackets is a stellar spectrum fit from the spectral synthesis library Starburst99

(Leitherer et al., 1999):

$$s_\nu = \frac{1}{5.5} \begin{cases} 5.5 & x < 1, \\ x^{-1.8} & 1 < x < 2.5, \\ 0.4x^{-1.8} & 2.5 < x < 4, \\ 2 \times 10^{-3} \frac{x^3}{e^{x/1.4} - 1} & x > 4. \end{cases} \quad (4.3)$$

The second term is a quasar-like power law, with index α (another free parameter). The free parameter f_Q is the strength of the quasar-like component relative to the stellar spectrum at $x = 1$. The entire spectrum is also attenuated by neutral hydrogen and helium (HeI) absorption with optical depth

$$\tau_\nu = \frac{\tau_0}{\sigma_{\text{HI},0}} [0.76\sigma_{\text{HI}}(\nu) + 0.06\sigma_{\text{HeI}}(\nu)], \quad (4.4)$$

where $\sigma_j(\nu)$ is the photoionization cross section as a function of frequency for ion j , computed using the fits from Verner et al. (1996), and $\sigma_{j,0}$ is that same cross section evaluated at the threshold frequency. Finally, τ_0 is a dimensionless free parameter which scales the overall optical depth. That is, it describes the gas column density attenuating the radiation field.

The training data consists of Cloudy models evaluated on a grid of values for the gas temperature T , hydrogen number density n_{H} , and metallicity Z , as well as the radiation field parameters J_0 , f_Q , τ_0 , and α . We describe the values used in Table 4.1. For $Z/Z_\odot \lesssim 0.01$, the cooling and heating functions do not depend on metallicity (metal-line cooling is unimportant, Boehringer & Hensler, 1989). Hence, we choose $Z/Z_\odot = 0.1$ as the lowest non-zero metallicity value.¹ We also require $Z/Z_\odot \leq 3$ as higher metallicity values are unlikely to occur in the ISM or CGM. This results in $81 \times 13 \times 5 \times 25 \times 9 \times 9 \times 7 \approx 7.5 \times 10^7$ points.

4.2.2 Radiation field sampling

Note that if a radiation field J_ν does not have precisely the form given by equation (4.2), the field may not have well-defined values of the parameters J_0 , f_Q , τ_0 , and α . To represent the radiation field, we instead compress the spectrum J_ν into a few parameters r_j by integrating over the frequency domain with some weight functions $w_j(\nu)$,

$$r_j = \int_0^\infty J_\nu w_j(\nu) d\nu. \quad (4.5)$$

¹The $Z/Z_\odot = 0$ case actually has $Z/Z_\odot = 10^{-4}$, but that value is small enough to be indistinguishable from the primordial gas for our purposes.

Parameter	Values
$\log(T/\text{K})$	1, 1.1, 1.2, \dots , 9
$\log(n_{\text{H}}/\text{cm}^{-3})$	-6, -5, -4, \dots , 6
Z/Z_{\odot}	0, 0.1, 0.3, 1, 3
$\log(J_0 \text{ cm}^{-3}/n_b/J_{\text{MW}})$	-5, -4.5, -4, \dots , 7
$\log(f_Q)$	-3, -2.5, -2, \dots , 1
$\log(\tau_0)$	-1, -0.5, 0, \dots , 3
α	0, 0.5, 1, \dots , 3

Table 4.1: Gas and radiation field parameters used to train XGBoost models of Cloudy-computed cooling and heating functions. We normalize the radiation field amplitude J_0 by $J_{\text{MW}} = 10^6 \text{ photons cm}^{-2} \text{ s}^{-1} \text{ ster}^{-1} \text{ eV}^{-1}$.

In Chapter 3, we used weights $w_j(\nu) = 4\pi\sigma_j(\nu)/(h\nu)$, where $\sigma_j(\nu)$ is the frequency-dependent cross section for photoionization of the ion j . With this choice, equation (4.5) yields the photoionization rates

$$P_j = c \int_0^{\infty} \sigma_j(\nu) n_{\nu} d\nu, \quad (4.6)$$

where $n_{\nu} = 4\pi J_{\nu}/(ch\nu)$ is the number density of photons at frequency ν .

However, we could in principle choose *any* sufficiently large set of weight function $w_j(\nu)$, and we are not aware of a rigorous way to find an optimal choice for such a set. In the absence of a way to define optimal weighting functions, we choose to use the simple case of constant weights $w(\nu) = \text{const}$ within a fixed interval $\nu_a < \nu < \nu_b$. More specifically, we compute the average intensity in a frequency bin:

$$\langle J_{\nu_a-\nu_b} \rangle = \frac{1}{\nu_b - \nu_a} \int_{\nu_a}^{\nu_b} J_{\nu} d\nu, \quad (4.7)$$

for several bins as our parameters for representing the radiation field.

This definition still leaves a choice of bin edges ν_a and ν_b . To choose these bin edges, we consider the correlations between different $\langle J_{\nu_a-\nu_b} \rangle$ values. Plugging equation (4.2) into equation (4.7) indicates that $\langle J_{\nu_a-\nu_b} \rangle$ scales linearly with the overall amplitude J_0 for all bins $\nu_a - \nu_b$. To account for this, we choose the bin $0.5 - 1 \text{ Ry}$, just below the ionization threshold for HI, and divide all other $\langle J_{\nu_a-\nu_b} \rangle$ by $\langle J_{0.5-1 \text{ Ry}} \rangle$ to remove the dependence on J_0 . This is directly analogous to our use of the photodissociation rate of molecular hydrogen P_{LW} to scale all other photoionization rates P_j in Chapter 3. The photons that dissociate molecular hydrogen have energies of about $0.82 - 1 \text{ Ry}$ (Draine & Bertoldi, 1996), lying within this $0.5 - 1 \text{ Ry}$ bin.

We are now faced with the question of how to choose bin edges $\nu_a - \nu_b$ for the values




Figure 4.1: Examples of the radiation field J_ν (normalized by its overall amplitude J_0) defined by equation (4.2) for different choices of α , f_Q , and τ_0 . The dashed vertical lines indicate our choice of bin edges, at 0.5, 1, 4, 7, 10, 13, and 16 Ry.

$\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$. Two options for selecting bin edges are linear spacing and logarithmic spacing. Linear spacing produces bins of constant width, while logarithmic spacing creates bins that increase in width exponentially with increasing frequency. Our choices here are guided by the correlations between bins and heuristic physical considerations. For example, the ionization energies for helium are approximately 1.8 Ry for HeI and 4 Ry for doubly ionized helium HeII (Verner et al., 1996). We select a 1 – 4 Ry bin to include the ionization thresholds for HI and HeI, but exclude HeII ionizing radiation.²

To shed more light on the correlations between various $\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$ values, we consider how the radiation field J_ν given by equation (4.2) changes with different choices of f_Q , τ_0 , and α . Three examples are shown in Figure 4.1, along with our choice of bin edges discussed below. Comparing the three panels of Figure 4.1 shows that J_ν/J_0 is strongly dependent on α at energies above about 16 Ry, but has only weak dependence on f_Q and τ_0 . So, we should expect all $\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$ to be highly correlated for $h\nu_a \gtrsim 16 \text{ Ry}$. The correlations between bins with logarithmically spaced edges are explored further in section 4.6.1, but we find linear bins to work better.

The Pearson and Spearman correlation matrices of $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ for bins of constant width $h\nu_b - h\nu_a = 3 \text{ Ry}$, up to $h\nu_b = 22 \text{ Ry}$ are shown in Figure 4.2. The Spearman correlation only considers the *ranks* of points. Thus, these correlations are unaffected by taking the logarithm of the scaled intensities (the logarithm is a monotonically increasing function). In contrast, the Pearson correlation is sensitive to the values of individual points, not just their ranks. The bins with energies above 16 Ry are very highly correlated, with both Pearson and Spearman correlations of at least 0.99. We explore energies up to 2048 Ry

²See appendix 4.6.1 for a discussion of using two energy bins in the 1 – 4 Ry range.




Figure 4.2: The absolute value of the Pearson (left panel) and Spearman (right panel) correlation coefficients between binned radiation field intensities $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ for the values in Table 4.1, with linear bin spacing. Given the high correlation between the highest three energy bins, we exclude the two highest bins (16 – 19 Ry and 19 – 22 Ry) from our initial model training.

(with logarithmic bin spacing) in section 4.6.1, but, as expected from Figure 4.1, all bins above 16 Ry are highly correlated with each other. The correlations are more modest for energies below 13 Ry.

For any $f_Q > 0$, the radiation field spectrum given by equation (4.2) becomes dominated by the quasar-like power law $x^{-\alpha}$ at sufficiently high photon energies x . Increasing f_Q (moving left to right in Figure 4.1) decreases the energy at which this power law becomes dominant. Since the optical depth τ_ν described in equation (4.4) includes absorption by HI and HeI, the primary effect of increasing the optical depth scaling τ_0 (moving left to right in Figure 4.1) is to reduce J_ν in the 1 – 4 Ry bin containing the ionization energies for HI and HeI. However, sufficiently large values of τ_0 also reduce the specific intensity in bins *above* the photon energy 4 Ry. This can be seen in the 4 – 7 and 7 – 10 Ry bins in the rightmost panel of Figure 4.1. So, increasing τ_0 decreases the flux in every bin above 1 Ry, but decreases the flux in the 1 – 4 Ry bin most strongly. Recall here that the Pearson correlation coefficient is sensitive to the values of each $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$, while the Spearman correlation coefficient depends only on their ranks. This suggests that the optical depth τ_0 should cause the Pearson correlation coefficient for $\log(\langle J_{1-4 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ and $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ with $h\nu_a > 4 \text{ Ry}$ to be smaller than the corresponding Spearman correlations (ignoring the

effects of varying f_Q and α). From Figure 4.2, we see that this is true for all but the 4 – 7 Ry bin.

For the remainder of this chapter, we focus on the less-correlated values of $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1\text{Ry}} \rangle)$ from Figure 4.2: 1 – 4, 4 – 7, 7 – 10, 10 – 13 and 13 – 16 Ry (as well as $\langle J_{0.5-1\text{Ry}} \rangle$, which incorporates the overall amplitude J_0).

4.2.3 Machine learning models

Gradient-boosted tree algorithms like XGBoost are known to outperform other machine learning methods for tabular training data like the the grid described in Table 4.1 (Shwartz-Ziv & Armon, 2022; Grinsztajn et al., 2022). XGBoost provides a useful framework for exploring a variety of radiation field parameters. Because each individual tree in an XGBoost model only utilizes a fixed fraction of the available features, the data size and computational cost scale weakly with the number of features. Furthermore, XGBoost interfaces easily with the feature importance tools described in Section 4.2.4.

We train the machine learning algorithm XGBoost (Chen & Guestrin, 2016) to predict $\log(\Gamma)$ and $\log(\Lambda)$ at fixed values of the metallicity Z . This results in 10 distinct models (cooling and heating at $Z/Z_\odot = \{0, 0.1, 0.3, 1, 3\}$). As discussed in Chapter 3, we train separate models at each metallicity because the 5 metallicity values in the training grid (see Table 4.1) are not enough samples for XGBoost to outperform a manual quadratic interpolation in metallicity. Indeed, Chapter 3 found that this metallicity interpolation is the main bottleneck for accurately predicting cooling and heating functions at arbitrary metallicities. The models are trained to optimize the MSE, defined as:

$$\text{MSE} = \langle [\log \mathcal{F}_{\text{true}} - (\log \mathcal{F})_{\text{pred}}]^2 \rangle, \quad (4.8)$$

where $\mathcal{F} = \Gamma$ or Λ .

The 8 dimensionless, logarithmically scaled inputs we use are described in the ‘6 bins’ column of Table 4.2. Even after taking the logarithm, these features vary over quite different ranges (e.g., see Table 4.1). To ameliorate this, we linearly rescale each feature to the range $[0, 1]$ for the values in Table 4.1.

Before training the ‘final’ XGBoost models we use for our analysis below, we perform an initial hyperparameter validation step. We use a random subsample of 80% of the input data (described in section 4.2.1) for model training. We subselect 10% of this training subset (i.e. 8% of the total data) for hyperparameter validation and train models with the optimal hyperparameter on the remaining 72%. The remaining 20% is used as a test set for final model evaluation. These fractions were chosen to ensure that the training, hyperparameter

	6 bins	3 bins
Temperature	T/K	T/K
Density	$n_{\text{H}}/\text{cm}^{-3}$	$n_{\text{H}}/\text{cm}^{-3}$
Amplitude	$\langle J_{0.5-1 \text{ Ry}} \rangle \text{ cm}^{-3}/n_{\text{b}}/J_{\text{MW}}$	$\langle J_{0.5-1 \text{ Ry}} \rangle \text{ cm}^{-3}/n_{\text{b}}/J_{\text{MW}}$
Other	$\langle J_{1-4 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$ $\langle J_{4-7 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$ $\langle J_{7-10 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$ $\langle J_{10-13 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$ $\langle J_{13-16 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$	$\langle J_{1-4 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$ $\langle J_{13-16 \text{ Ry}} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle$
	2 bins	
Temperature	T/K	
Density	$n_{\text{H}}/\text{cm}^{-3}$	
Amplitude	$\langle J_{1-4 \text{ Ry}} \rangle \text{ cm}^{-3}/n_{\text{b}}/J_{\text{MW}}$	
Other	$\langle J_{13-16 \text{ Ry}} \rangle / \langle J_{1-4 \text{ Ry}} \rangle$	
	4 rates	3 rates
Temperature	T/K	T/K
Density	$n_{\text{H}}/\text{cm}^{-3}$	$n_{\text{H}}/\text{cm}^{-3}$
Amplitude	$P_{\text{LW}} \text{ cm}^{-3}/\text{s}^{-1}$	$P_{\text{LW}} \text{ cm}^{-3}/\text{s}^{-1}$
Other	$P_{\text{HI}}/P_{\text{LW}}$ $P_{\text{HeI}}/P_{\text{LW}}$ $P_{\text{C VI}}/P_{\text{LW}}$	$P_{\text{HI}}/P_{\text{LW}}$ $P_{\text{C VI}}/P_{\text{LW}}$

Table 4.2: The features sets used for each XGBoost model considered in this chapter. Note that all features are the logarithm of the quantities shown here. Above the line are features describing gas properties, which are the same for each of our models. Below the line are features related to the incident radiation field. The first row describes the radiation field features containing the overall amplitude of the radiation field J_0 (see equation (4.2)). The remaining rows show all other radiation field features, which are scaled to the rate or bin containing the overall amplitude.

validation, and test sets are all large enough to accurately sample the distributions of each feature and that the training set contains a majority of the available data. We use the same data splitting for all the models discussed below. See section 4.6.2 for a brief description of the hyperparameter optimization procedure, which we more fully laid out in section 3.6.1.

Finally, we retrain models using the optimized hyperparameters (found as described in section 4.6.2) on both the same 80% training subset described above (so we can analyze model performance on a 20% test set withheld from model training) and on the entire input data. The model evaluation on the 20% test set illustrates how well our models generalize to data they have not yet seen, providing an unbiased measure of how well our models perform. Training models on the entire training grid (and testing on the same data) provides a quantification of an overfitted result and a potential best-case scenario of how well our models might do with a larger training sample.

4.2.4 Feature importance with SHAP values

To evaluate the impact each feature has on model predictions, we use SHAP values (Lundberg & Lee, 2017; Lundberg et al., 2018; Lundberg et al., 2020). SHAP values can be calculated for each feature for any given point prediction of a machine learning model. SHAP values describe how the model prediction differs from what would be expected from ignoring the value of the feature in question (Lundberg & Lee, 2017).

We use the Python package `shap` (Lundberg et al., 2018; Lundberg et al., 2020), which incorporates functions to approximate SHAP values for tree-based models like XGBoost, to calculate SHAP values for predictions on the 20% test set from the 10 models trained on only 80% of the input data. We randomly select 500 points from the 20% test set for each model. Since SHAP values can be either positive or negative (i.e. a model prediction can be either increased or decreased, given the value of a specific feature), we consider the *mean absolute SHAP value* over the 500 test points for each feature. As an example of this, Figure 4.3 shows the 7 features with the largest mean absolute SHAP values for cooling and heating function models at $Z/Z_\odot = 1$ (trained using 6 bins, as described in the ‘6 bins’ column of Table 4.2).

As seen in Figure 4.3, the temperature feature $\log(T/\text{K})$ and the 0.5 – 1 Ry bin that includes the overall radiation field amplitude, $\log(\langle J_{0.5-1\text{Ry}} \rangle \text{cm}^{-3}/n_b/J_{\text{MW}})$, are the two most important features for all 10 models. The density feature, $\log(n_{\text{H}}/\text{cm}^{-3})$, tends to have relatively low mean absolute SHAP values, especially for the heating function.

In order to examine the SHAP values for the scaled radiation field bin features more closely, we divide the mean absolute SHAP values for each $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1\text{Ry}} \rangle)$ feature

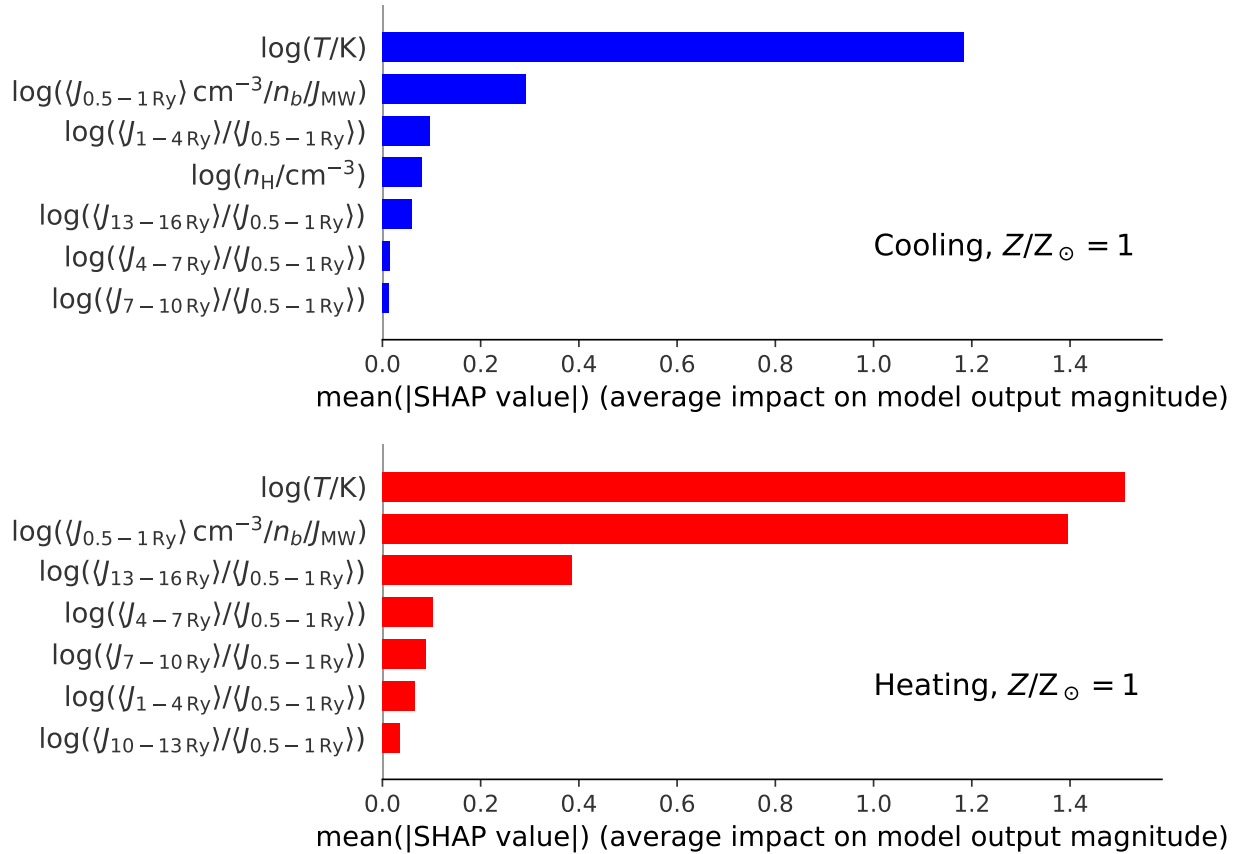


Figure 4.3: Mean absolute SHAP values on 500 randomly selected points in the test set (20% of the training grid withheld from model training) for XGBoost models trained using 6 bins (see ‘6 bins’ column of Table 4.2). Models for the cooling function (top panel) and heating function (bottom panel) at $Z/Z_{\odot} = 1$ are shown. Only the 7 features with the largest mean absolute SHAP values for each model are shown. At $Z/Z_{\odot} = 1$, the gas temperature and overall amplitude of the radiation field are most important.

by the corresponding mean absolute SHAP values for $\log(\langle J_{0.5-1\text{Ry}} \rangle \text{cm}^{-3}/n_b/J_{\text{MW}})$. These scaled SHAP values are shown in Figure 4.4 for all 10 models described in section 4.2.3 above.

Figure 4.4 shows that the 13–16 Ry bin has high importance *relative to the other radiation field bins* for the heating function at metallicities $Z/Z_{\odot} > 0$. Other bins have similarly low importance for the heating function. The 1–4 Ry bin has high relative importance for the cooling function at all metallicities. The 13–16 Ry bin has similar relative importance to the 1–4 Ry bin at the highest metallicities ($Z/Z_{\odot} = 1, 3$). Overall, the 1–4 Ry and 13–16 Ry are the most important scaled bins. For $Z/Z_{\odot} > 0.1$, the cooling function for $10^4 \lesssim T/\text{K} \lesssim 10^7$ is dominated by high ionization states of heavy elements like carbon, oxygen, neon, and iron (Wiersma et al., 2009). The relevant ionization thresholds are $\gtrsim 13$ Ry, and so are sampled by the 13–16 Ry bin. For example, FeIX becomes an important coolant in highly photoionized gas (e.g. Cohen et al., 2000; Cantalupo, 2010) and has an ionization potential of ~ 12 Ry (Kramida et al., 2024).

Motivated by these observations, we also train an additional 10 models (cooling and heating at each metallicity) using only the two scaled bins 1–4 Ry and 13–16 Ry. These models are trained using exactly the same pipeline as described in section 4.2.3. We label these models as ‘3 bins’, since they incorporate 3 radiation field bins (0.5–1, 1–4, and 13–16 Ry), in contrast to the ‘6 bins’ models described in section 4.2.3 (see Table 4.2 for a comparison between the features used in these models).

Figure 4.4 shows that 1–4 and 13–16 Ry are the most important scaled bins at all metallicities. Photoionization of HI and HeI by photons with energies greater than 1 Ry are vital contributions to the heating function (Ferland et al., 1998). So, we decide to train additional heating function models using the 1–4 Ry bin to include the overall radiation field amplitude J_0 , and scale the intensity in the 13–16 Ry bin by that value. We label these models as ‘2 bins’, since only the 1–4 and 13–16 Ry bins are utilized. The features used are described in corresponding column of Table 4.2. Again, these models are trained using the pipeline of section 4.2.3.

For comparison with these models, we also consider two sets of models trained with photoionization rates P_j (see equation (4.6)) as described in Chapter 3. First, we use models trained with P_j for $j = \text{LW, HI, HeI, and CVI}$ (the same rates used in the interpolation table of Gnedin & Hollon (2012)), which we label ‘4 rates’. As demonstrated in Chapter 3, these models perform comparably well to models trained with other sets of 4 photoionization rates, as well as larger sets of photoionization rates. We also train new models with ‘3 rates’ (P_j for $j = \text{LW, HI, and CVI}$) to compare with our ‘3 bins’ models described above. Note that the choices of rates and bins for these models are entirely analogous. The radiation responsible

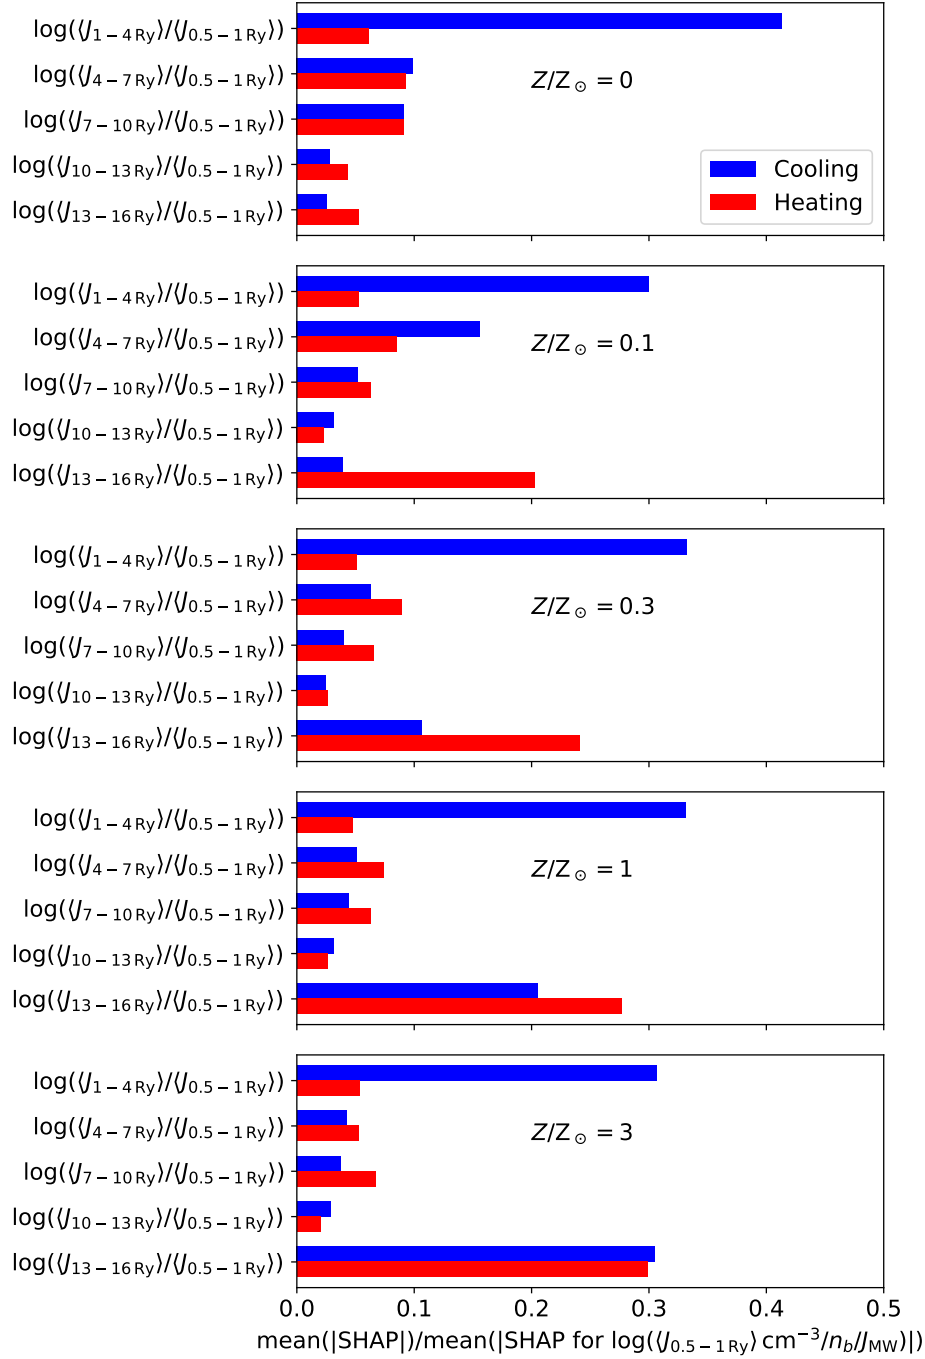


Figure 4.4: The mean absolute SHAP values for the quantity $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1\text{Ry}} \rangle)$, divided by that for $\log(\langle J_{0.5-1\text{Ry}} \rangle \text{ cm}^{-3} / n_b / J_{\text{MW}})$. For each feature, the values are shown for the cooling function (above, blue) and the heating function (below, red). The 5 panels are for models trained at different metallicities, from $Z/Z_\odot = 0$ at the top to $Z/Z_\odot = 3$ at the bottom. Either the 1 – 4 or 13 – 16 Ry bins (top and bottom bars in each panel) are the most important for both cooling and heating at each metallicity.

for photodissociation of molecular hydrogen lies in the 0.5 – 1 Ry bin, as described above. The 1 – 4 Ry contains HI-ionizing radiation, and the 13 – 16 Ry bin and CVI photoionization rate both sample J_ν at X-ray energies. The features for these ‘4 rates’ and ‘3 rates’ models are described in Table 4.2.

4.3 Results

As seen in Chapter 3, XGBoost models trained on the data described in section 4.2.1 occasionally make very large ‘catastrophic errors’. To track these, the primary metric we use to evaluate the performance of our models is the *distribution* of all errors in a given sample, rather than the MSE. Specifically, we define the error distribution function

$$P(> \Delta \log \mathcal{F}) = \frac{\text{Points with error above } \Delta \log \mathcal{F}}{\text{Total number of points}}, \quad (4.9)$$

where, once again, $\mathcal{F} = \Gamma$ or Λ .

We examine error distributions for all the cooling and heating function models described in Table 4.2, at each metallicity value. We show the error distribution on the entire input data for models trained on the entire input data, which assesses how well the models are able to fit this data, in Figure 4.5. The error distributions on the 20% withheld test set for models trained on 80% of the input data, which assesses how predictive the models are, are discussed in section 4.6.3. We compare the performance of our XGBoost models on the entire training grid with that of the interpolation table from Gnedin & Hollon (2012) (see Figure 4.5). The corresponding MSE values are shown in section 4.6.3. We also compare the MSEs for 3 bin and 2 bin heating function models in section 4.6.4. Since both the training and test set MSEs are about two orders of magnitude higher for the 2 bin models than the 3 bin models, we do not include the 2 bin heating function models in Figure 4.5.

For both cooling and heating functions, and at each metallicity value, the performance of the models trained with 3 or 6 radiation field bin features and 3 or 4 photoionization rates is comparable, and much better than the interpolation table of Gnedin & Hollon (2012). This can be seen through the similarity of the thick solid, dashed, dash-dotted, and dotted training data error distribution curves in Figure 4.5. Note that the interpolation table of Gnedin & Hollon (2012) uses a quadratic Taylor expansion in metallicity. The errors for the interpolation table are dominated by errors in the metallicity interpolation (see Chapter 3). Because of this, our XGBoost models are able to outperform the interpolation table at fixed metallicity. However, the Gnedin & Hollon (2012) interpolation scheme is more general, and works at all metallicities (within the range $0 \leq Z/Z_\odot \leq 3$ of the training data in Table 4.1).

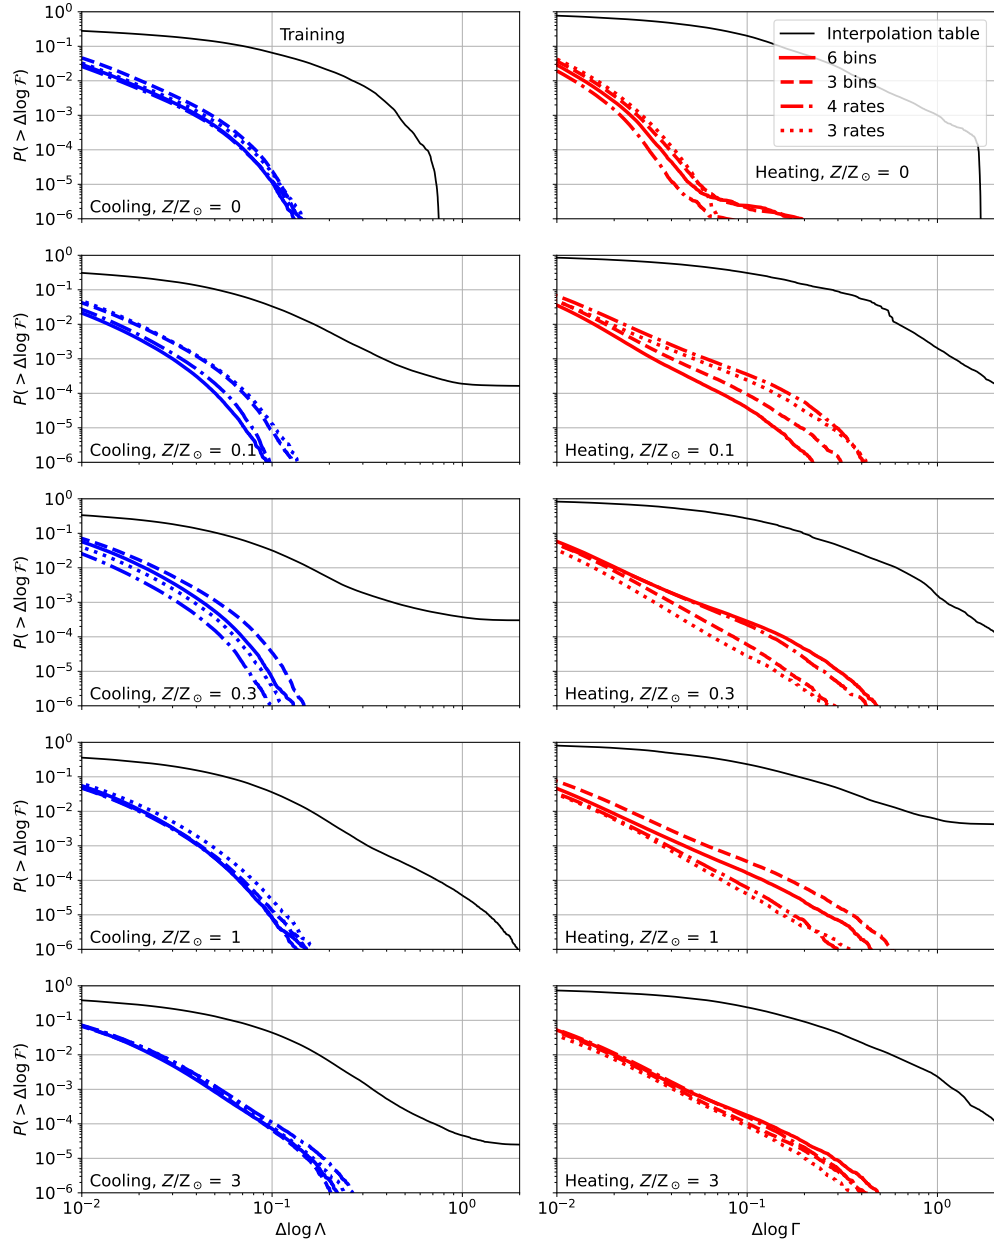


Figure 4.5: Error distributions for the interpolation table of Gnedin & Hollon (2012) (thin solid lines) and our XGBoost models trained with 6 bins (thick solid lines, described in section 4.2.3), 3 bins (thick dashed lines, described in section 4.2.4), 4 rates (thick dash-dotted lines, described in section 4.2.4), and 3 rates (thick dotted lines, described in section 4.2.4) for models trained and evaluated on the entire training grid. See Table 4.2 for the features included in each model. Error distributions for the cooling function are shown in the left column, with the heating function in the right column. All models are trained to minimize the MSE on the training set. Our XGBoost models all have similar performance on the training data, and significantly outperform an interpolation table.

The MSEs for 3 and 6 bin and 3 and 4 rate models differ only by factors of order unity (on either the entire training grid, or the test subset). The MSEs on the test set are always larger than the corresponding training set MSEs by a factor of order unity (again, see section 4.6.3). The heating function models using 2 radiation field bin features have much larger MSEs than the 3 bin models for both the training and test sets (see section 4.6.4, where this is explored further).

For the heating function at $Z/Z_\odot = 0.3$, and both cooling and heating at $Z/Z_\odot = 3$, we observe that the training data MSEs for the 3 bin models are rather counterintuitively *smaller* than the training data MSEs for the corresponding 6 bin models (see section 4.6.3). Note that, since the dependence of the cooling functions and heating function on radiation field bin features $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1\text{Ry}} \rangle)$ is non-linear, training on a finite set of data results only in the *approximately* optimal dependence of the cooling and heating functions on these binned radiation field features. In other words, while the XGBoost training procedure minimizes the MSE *on a given set of training data*, it is entirely possible that the model performance could be improved with additional training data (that better samples the dependence of the cooling or heating function at the given metallicity on the features being used). For the models in question, our training data table may not be large enough to take full advantage of the the extra information provided by the additional 3 scaled radiation field bin features in the 6 bin models. Also, for a given number of trees and maximum tree depth (see section 4.6.2), the predictions are likely to be more precise for a model with *fewer* features. This is because each feature will be used to split the input data more often than in a model with more features.

4.4 Conclusions and discussion

In this chapter, we used machine learning to analyze the cooling and heating functions of gas in the presence of a generalized radiation field. The radiation field contains a synthesized stellar spectrum, a quasar-like power law, and absorption by neutral hydrogen and helium. We describe the radiation field with averaged intensities in discretized photon energy bins. We use the 0.5 – 1 Ry bin to encode the overall amplitude of the radiation field and scale all other binned intensities by that value. Then, we examine the correlations between such binned intensities. We train machine learning models to predict the true (i.e. Cloudy-computed) cooling and heating functions of the gas at 5 fixed metallicities with the 6 least correlated radiation field intensity bins, and the subsequently identified 3 most important radiation field bins, to predict cooling and heating functions. We also compare these models to models from Chapter 3 trained using 3 and 4 photoionization rates. The main conclusions

from this analysis are:

- For our radiation field model (equation (4.2)), very little additional information is included above a photon energy of 16 Ry (see Figure 4.2). This is because, at such high energies, the radiation field (up to a factor of the overall amplitude J_0) is strongly dependent on the quasar-like power law index α , and only weakly dependent on the other parameters (f_Q, τ_0 , see Figure 4.1).
- For XGBoost models trained with 6 radiation field intensity bins, the 0.5 – 1 Ry bin (containing the overall radiation field amplitude) and temperature are always the two most important features by mean absolute SHAP values (e.g., see Figure 4.3).
- Among the scaled radiation field bin intensities, the 13 – 16 Ry bin has the largest mean absolute SHAP value for heating function models at $Z/Z_\odot > 0$. The 13 – 16 and/or 1 – 4 Ry bins have significantly larger mean absolute SHAP values than the other bins for cooling function models at all metallicities (see Figure 4.4).
- The MSEs and error distributions (on both the training data, and a 20% subset withheld from training) are very similar for models trained with 3 and 6 radiation field bins, and also for 3 and 4 photoionization rates. These errors are $\gtrsim 10$ times smaller than for the interpolation table of Gnedin & Hollon (2012) applied to the same data. However, the MSEs are significantly worse for the heating function models trained with only 2 bins (see Figure 4.5 and section 4.6.3).
- We conclude that just 3 photon energy bins (0.5 – 1, 1 – 4, and 13 – 16 Ry) or 3 photoionization rates (P_j for $j = \text{LW, HI, and CVI}$) are sufficient to capture the dependence of gas cooling and heating functions on an incident radiation field with a spectrum given by equation (4.2).

We find that XGBoost models trained using both photoionization rates and energy bins to describe the incident radiation field are able to interpolate Cloudy calculations of gas cooling and heating functions more accurately than the interpolation table of Gnedin & Hollon (2012) at fixed metallicity. The ‘4 rates’ XGBoost model is slower to evaluate than the interpolation table while utilizing the same input parameters. However, adding new radiation field parameters to an interpolation table means increasing the dimension and data size of the table. Because each tree in an XGBoost model only uses a given subset of the available features, the data size of XGBoost models (with the same hyperparameters) scales only weakly with the number of features. This allows us to construct models (such as the ‘6 bins’ models presented here) which incorporate more radiation field parameters

than the Gnedin & Hollon (2012) interpolation table without significantly increasing the evaluation time from a model with fewer features.

Given that our radiation field model (see equation (4.2)) has 4 free parameters, it is surprising that 3 radiation field features are enough to capture the dependence of cooling and heating functions on that radiation field. In particular, the high energy power-law tail of the radiation spectrum is set by the power-law slope α and the amplitude J_0 , and these two parameters cannot be captured by the single value in the 13 – 16 Ry band.

Our conclusion that 3 energy bins are enough to accurately capture the cooling and heating functions with XGBoost relies on our specific choices of energy bins. We expect the most important energy bins to be robust to changing the atomic data set or using a different photoionization code. The overall amplitude of the radiation field and H and He ionization radiation in the 1-4 Ry bin will always be important, and cooling at $Z/Z_\odot > 0.1$ is dominated by metal-line cooling (e.g. Wiersma et al., 2009), ensuring that radiation in the quasar-dominated part of the spectrum capable of accessing the highest ionization states of metals such as carbon, oxygen, and iron will have a significant effect. However, the machine learning pipeline used in this work (section 4.2.3) could be used to train new cooling and heating function models with *any* set of radiation field energy bins, and in particular a set used in a radiative transfer simulation.

The radiation field sampling in equation (4.5) can be done with any frequency bounds ν_a and ν_b , and any weight functions $w_j(\nu)$. In this work, we have only consider two cases, constant weights $w_j(\nu) = \text{const}$ and photoionization cross sections $w_j(\nu) \propto \sigma_j(\nu)$. As shown in Figure 4.5 and section 4.6.3, 3 such radiation field samples are sufficient to describe the cooling and heating functions reasonably accurately in both cases. It is of course possible that with different, more optimal weighting functions $w_j(\nu)$, one could find a set of only 2 radiation field samples sufficient to specify the gas cooling and heating functions or a set of 3 samples with significantly smaller errors than the samples presented here. However, finding such weighting functions is a non-trivial exercise.

We have also assumed the radiation field spectrum is well-described by equation (4.2). While this is fairly general, it of course does not include every possible radiation field. For example, it only include absorption by neutral hydrogen and helium, and not other elements or ions. Our machine learning framework is capable of handling a more complex radiation field model, and future work could involve running more Cloudy models to expand or add more dimensions to the training data in Table 4.1. The Cloudy calculations in the training data used here include only *atomic* cooling and heating processes. Additional processes such as cooling and heating from dust grains and cosmic ray heating could be incorporated by adding additional dimensions (e.g. describing the dust abundance and cosmic ray ionization

rate) into the training data in Table 4.1. These processes might introduce additional important radiation field energy bins, perhaps including photons capable of ejecting an electron from dust grains and heating the gas. Non-solar element abundances would also change the cooling and heating functions from the values in the training data. Abundances for a few particularly important elements could also be included as additional dimensions of the training data in Table 4.1. While incorporating non-solar abundances would require training new XGBoost models, we expect our qualitative conclusions about the most significant energy bins and photoionization rates to be robust as long as there are non-trivial amounts of any important metal coolants whose higher ionization states are excited by photons in the quasar-dominated region of the radiation field spectrum described by the 13 – 16 Ry energy bin.

4.5 Acknowledgements

I, and co-author CA, acknowledge support from the Leinweber Foundation. Co-author CA acknowledges support from DOE grant DE-SC009193. This chapter has been co-authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. This research was also supported in part through computational resources and services provided by Advanced Research Computing (ARC), a division of Information and Technology Services (ITS) at the University of Michigan, Ann Arbor, in particular the Great Lakes cluster and the U-M Research Computing Package.

This work utilizes many Python packages, including `xgboost` (Chen & Guestrin, 2016), `shap` Lundberg et al. (2018); Lundberg et al. (2020), `numpy` (Harris et al., 2020), `pandas` (pandas development team, 2020), `matplotlib` (Hunter, 2007), `seaborn` (Waskom, 2021), `scikit-learn` (Pedregosa et al., 2011), and `scikit-optimize` (Head et al., 2021). The code pipeline to train and evaluate these models can be found at https://github.com/davidbrobins/ml_chf.

4.6 Appendix

4.6.1 Logarithmic bin spacing

In addition to the bin edges with linear spacing shown in Figure 4.2, we also consider the correlations of $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ for bins with logarithmic spacing, from 1 – 2 Ry up to 1024 – 2048 Ry. The Pearson and Spearman correlation matrices for the first 7 of these bins




Figure 4.6: The absolute value of the Pearson (left panel) and Spearman (right panel) correlation coefficients between binned radiation field intensities $\log(\langle J_{\nu_a-\nu_b} \rangle / \langle J_{0.5-1 \text{ Ry}} \rangle)$ for the values in Table 4.1 with logarithmic bin spacing.

are shown in Figure 4.6 (the remaining five bins are all highly correlated with the 64 – 128 Ry bin).

From Figure 4.6, we see that the 1 – 2 and 2 – 4 Ry bins have large correlations with each other for both Pearson (0.99) and Spearman (0.94). Thus, it is reasonable to combine these into a single 1 – 4 Ry, as we do in section 4.2.2.

The correlation matrices for both linear (Figure 4.2) and logarithmic (Figure 4.6) bins show that all bins above 16 Ry are highly correlated with each other. As discussed in section 4.2.2, this is unsurprising because the radiation field given by equation (4.2) is dominated by the quasar-like power law at these energies (see Figure 4.1). This power law depends on only one parameter, the index α . The logarithmic bin spacing used in Figure 4.6 results in only 4 bins below 16 Ry, while the linear spacing we use in section 4.2.2 gives 5 bins below 16 Ry (exactly the bins used for our 6 bin models).

4.6.2 Hyperparameter tuning

In the XGBoost hyperparameter validation step described in section 4.2.3, we optimize 5 XGBoost parameters:³

³These, as well as other XGBoost hyperparameters, are described in more detail in the code documentation: <https://xgboost.readthedocs.io/en/stable/index.html>

- `min_child_weight`: minimum ‘weight’ required for each tree leaf. Larger values result in simpler models. Default value: 1 (cannot be negative).
- `subsample`: the fraction of the training data points used to train each tree. Default value: 1 (cannot be negative or larger than 1).
- `colsample_bytree`: the fraction of features used to train each tree. Default value: 1 (cannot be negative or larger than 1).
- `gamma`: the minimum improvement in model performance required to split a node. Larger values result in simpler models. Default value: 0 (cannot be negative).
- `eta`: the learning rate. Default value: 0.3 (cannot be negative or larger than 1).

Similarly to Chapter 3, we fix the maximum depth (number of nodes to reach the leaves of the tree) to `max_depth = 20` and the number of trees to `n_estimators = 100`.

To optimize `min_child_weight`, `subsample`, `colsample_bytree`, `gamma`, and `eta`, we use 5-fold cross-validation. That is, we split the validation subset (as described in section 4.2.3, this is 8% of the input data) into 5 disjoint subsets. For each of these 5 ‘folds’, we train a model on the remaining 4 folds and evaluate it on the fold in question. The performance of each model is the average MSE across each of the 5 folds. To minimize this value, we perform a Bayesian search, with the following priors:

- `min_child_weight`: log-uniform prior from 0.1 to 2.
- `subsample`: uniform prior from 0.6 to 1.
- `colsample_bytree`: uniform prior from 0.6 to 1.
- `gamma`: uniform prior from 0 to 1.
- `eta`: log-uniform prior from 0.03 and 0.3.

This entire procedure is implemented with the function `BayesSearchCV` in the `scikit-optimize` package (Head et al., 2021).

4.6.3 Model error comparisons

In this section, we present further details on the performance on our XGBoost models. The error distributions (see equation (4.9)) on the 20% withheld test set for models trained on 80% of the training grid are shown in Figure 4.7. We also overplot the error distributions for the interpolation table of Gnedin & Hollon (2012) on the entire data grid (described

Cooling	6 bins		3 bins	
	Training	Test	Training	Test
$Z/Z_{\odot} = 0$	1.65×10^{-5}	4.05×10^{-5}	2.52×10^{-5}	5.18×10^{-5}
$Z/Z_{\odot} = 0.1$	1.10×10^{-5}	4.62×10^{-5}	2.15×10^{-5}	4.94×10^{-5}
$Z/Z_{\odot} = 0.3$	2.72×10^{-5}	5.83×10^{-5}	3.63×10^{-5}	6.69×10^{-5}
$Z/Z_{\odot} = 1$	2.58×10^{-5}	6.63×10^{-5}	2.39×10^{-5}	6.66×10^{-5}
$Z/Z_{\odot} = 3$	3.48×10^{-5}	9.10×10^{-5}	3.46×10^{-5}	8.64×10^{-5}

	4 rates		3 rates	
	Training	Test	Training	Test
$Z/Z_{\odot} = 0$	1.49×10^{-5}	3.76×10^{-5}	1.89×10^{-5}	3.85×10^{-5}
$Z/Z_{\odot} = 0.1$	1.36×10^{-5}	3.58×10^{-5}	2.41×10^{-5}	4.77×10^{-5}
$Z/Z_{\odot} = 0.3$	1.28×10^{-5}	4.04×10^{-5}	1.96×10^{-5}	4.56×10^{-5}
$Z/Z_{\odot} = 1$	2.29×10^{-5}	5.94×10^{-5}	3.25×10^{-5}	6.82×10^{-5}
$Z/Z_{\odot} = 3$	3.95×10^{-5}	8.06×10^{-5}	3.84×10^{-5}	8.02×10^{-5}

Table 4.3: MSEs (see equation (4.8)) for the cooling function XGBoost models in the left column of Figure 4.5 and Figure 4.7 (see Table 4.2). All MSEs have similar orders of magnitude.

in section 4.2.1). Similarly to Figure 4.5, the error distributions for all 4 XGBoost models shown in Figure 4.7 are similar, and much better than the interpolation table of Gnedin & Hollon (2012).

The MSEs for the XGBoost models in Figure 4.5 and Figure 4.7 are shown for cooling function models in Table 4.3 and for heating function models in Table 4.4. As discussed in section 4.3, all of the MSEs in each row of Table 4.3 or Table 4.4 are of the same order of magnitude.

4.6.4 The performance of 2 bin models

We compare the MSEs of our 2 bin and 3 bin heating function models in Table 4.5. Models trained with only 2 bins have MSEs that are about 2 orders of magnitude higher than the corresponding 3 bin models, both for models evaluated on the entire training grid and on a 20% withheld test set (see Table 4.2 for the specific features used in each model).

To understand why the 2 bin models have such large MSEs, we choose a single point from the training data where the heating function error at $Z/Z_{\odot} = 0$ satisfies $\Delta \log \Gamma > 0.8$. The values of each input to the training table for the selected point are shown in Table 4.6.

Then we select all other points in the training grid with the same density, and values for the two binned radiation field features, $\log(\langle J_{1-4\text{Ry}} \rangle \text{cm}^{-3}/n_b/J_{\text{MW}})$ and $\log(\langle J_{13-16\text{Ry}} \rangle/\langle J_{1-4\text{Ry}} \rangle)$ within 10% of the values of the point described in Table 4.6 *after*




Figure 4.7: Error distributions for our XGBoost models trained with 6 bins (solid lines, described in section 4.2.3), 3 bins (dashed lines, described in section 4.2.4), 4 rates (dash-dotted lines, described in section 4.2.4), and 3 rates (dotted lines, described in section 4.2.4) for models trained on 80% of the training grid and evaluated on the withheld 20% test subset (see section 4.2.3). See Table 4.2 for the features included in each model. The error distributions for the interpolation table of Gnedin & Hollon (2012) on the entire training grid are overplotted as thin solid lines. Distributions for the cooling function are shown in the left column, with the heating function in the right column. All models are trained to minimize the MSE on the training set. All our XGBoost models have comparable performance on the test subset.

Heating	6 bins		3 bins	
	Training	Test	Training	Test
$Z/Z_\odot = 0$	1.53×10^{-5}	2.59×10^{-5}	1.81×10^{-5}	2.93×10^{-5}
$Z/Z_\odot = 0.1$	1.94×10^{-5}	6.16×10^{-5}	2.64×10^{-5}	5.88×10^{-5}
$Z/Z_\odot = 0.3$	3.83×10^{-5}	7.80×10^{-5}	2.58×10^{-5}	6.57×10^{-5}
$Z/Z_\odot = 1$	2.97×10^{-5}	7.33×10^{-5}	4.99×10^{-5}	8.69×10^{-5}
$Z/Z_\odot = 3$	3.23×10^{-5}	7.68×10^{-5}	2.75×10^{-5}	6.79×10^{-5}

	4 rates		3 rates	
	Training	Test	Training	Test
$Z/Z_\odot = 0$	1.10×10^{-5}	2.08×10^{-5}	1.97×10^{-5}	3.15×10^{-5}
$Z/Z_\odot = 0.1$	4.51×10^{-5}	6.74×10^{-5}	3.25×10^{-5}	5.46×10^{-5}
$Z/Z_\odot = 0.3$	3.22×10^{-5}	5.79×10^{-5}	1.79×10^{-5}	4.65×10^{-5}
$Z/Z_\odot = 1$	1.96×10^{-5}	4.38×10^{-5}	1.77×10^{-5}	4.17×10^{-5}
$Z/Z_\odot = 3$	3.28×10^{-5}	5.82×10^{-5}	2.17×10^{-5}	4.61×10^{-5}

Table 4.4: MSEs (see equation (4.8)) for the heating function XGBoost models in the right column of Figure 4.5 and Figure 4.7 (see Table 4.2). All MSEs have similar orders of magnitude.

Heating	3 bins		2 bins	
	Training	Test	Training	Test
$Z/Z_\odot = 0$	1.81×10^{-5}	2.93×10^{-5}	2.01×10^{-3}	2.79×10^{-3}
$Z/Z_\odot = 0.1$	2.64×10^{-5}	5.88×10^{-5}	4.95×10^{-3}	6.82×10^{-3}
$Z/Z_\odot = 0.3$	2.58×10^{-5}	6.57×10^{-5}	5.90×10^{-3}	9.35×10^{-3}
$Z/Z_\odot = 1$	4.99×10^{-5}	8.69×10^{-5}	7.03×10^{-3}	1.09×10^{-2}
$Z/Z_\odot = 3$	2.75×10^{-5}	6.79×10^{-5}	7.68×10^{-3}	1.18×10^{-2}

Table 4.5: MSEs (see equation (4.8)) for 3 bin and 2 bin heating function models (see Table 4.2). The 2 bin model is systematically ~ 2 orders of magnitude less accurate than the 3 bin model.

Parameter	Values
$\log(T/\text{K})$	1
$\log(n_{\text{H}}/\text{cm}^{-3})$	-6
Z/Z_\odot	0
$\log(J_0 \text{ cm}^{-3}/n_b/J_{\text{MW}})$	-5
$\log(f_Q)$	1
$\log(\tau_0)$	-2
α	0

Table 4.6: Parameters for a point in the training grid (see Table 4.1) with $\Delta \log \Gamma > 0.8$ for a 2 bin heating function model at $Z/Z_\odot = 0$.




Figure 4.8: The predicted heating function for the training grid point described in Table 4.6 (thick orange curve), and the true heating functions for every training grid point with the same density, as well as binned radiation field features (see Table 4.2) within 10% of the values for the point in Table 4.6 (thin light blue curves).

scaling to the interval 0 – 1 (0.512886 and 0.306554, respectively). In Figure 4.8, we plot the true heating function (i.e. as a function of temperature) from Cloudy for each such point, along with the predicted heating function for the point in Table 4.5.

While the true heating functions in Figure 4.8 are fairly similar for $T \gtrsim 10^{6.5}$ K, they span almost 3 orders of magnitude for $T \lesssim 10^4$ K. So, any single prediction would by necessity have $\Delta \log \Gamma \sim 1$ at $T \lesssim 10^4$ K for some of the true heating functions shown in Figure 4.8. This indicates that points with similar values of the radiation field features for our 2 bin models (see Table 4.2) can have very different heating functions, suggesting that these two radiation field features are insufficient to capture the behavior of the heating function.

To illustrate this, we plot the radiation field specific intensity J_ν for the points with the highest and lowest heating function values at $\log(T/K) = 1$ from Figure 4.8. These spectra




Figure 4.9: The radiation field specific intensity J_ν (equation (4.2)) for the points from Figure 4.8 with the lowest (left panel) and highest (right panel) heating function values at $\log(T/K) = 1$.

are shown in Figure 4.9.

The two spectra in Figure 4.9 have the same values of optical depth scaling $\tau_0 = 100$ and quasar-like power law index $\alpha = 0$, but very different values of the fractional contribution of quasars, f_Q . The value of J_ν plateaus at energies above about 10 Ry in both spectra. In the right panel, the high value of f_Q means that J_ν is as large above 10 Ry as it is below 1 Ry. However, in the left panel, the value of J_ν above 10 Ry is almost 3 orders of magnitude smaller than that below 1 Ry. This difference would be captured by the radiation field features of our 3 and 6 bin models (see Table 4.2), which utilize the 0.5 – 1 Ry bin to scale all other bins. However, it is not captured by scaling with the 1 – 4 Ry bin used in our 2 bin heating function models.

The intensity in the 0.5 – 1 Ry bin (used to include J_0 and for scaling all other bins in our 3 and 6 bin models) has no dependence on the optical depth parameter τ_0 , because only photons with energies greater than 1 Ry can ionize HI or HeI (see equation (4.4)). On the other hand, the intensity in the 1 – 4 Ry is *strongly* sensitive to the value of τ_0 (see Figure 4.1). The 13 – 16 Ry bin, should only have weak dependence on τ_0 . Scaling by the 1 – 4 Ry bin (as is done in our 2 bin models) introduces additional dependence on τ_0 . This creates a degeneracy in the radiation field J_ν implied by our 2 bin model features (see Table 4.2). This is the same degeneracy seen in the spread of the heating function at low

temperatures in Figure 4.8.

CHAPTER 5

The Effects of Different Cooling and Heating Function Models on a Simulated Analog of NGC300

5.1 Motivation

Gas in the ISM and CGM of galaxies exists in multiple interacting phases with different temperatures and densities. These phases include a diffuse hot ionized medium (Spitzer, 1956), clouds of an atomic cold neutral medium (Ewen & Purcell, 1951), a warm ionized medium surrounding the cold clouds (Hjellming et al., 1969), and very cold molecular clouds (Cheung et al., 1968). Gas cools and heats through the respective emission and absorption of radiation, providing a critical mechanism for gas to transition between these phases (e.g. McKee & Ostriker, 1977). Star formation (primarily) occurs in cold molecular clouds, and the star formation rate can depend on both the thermal state *and turbulent properties* of the molecular gas (see Dobbs (2023) and Hennebelle & Grudić (2024) for recent reviews).

Gas cooling and heating rates thereby play a vital component of galaxy evolution models, including hydrodynamic simulations. However, there are many different prescriptions in use for calculating the cooling and heating rates as functions of the gas properties and the galactic environment (see Kim et al. (2023) for an overview).

The effects of different cooling and heating function models on galaxy formation can be tested using galaxy formation simulations. Simulation comparisons from the Aquila and AGORA collaborations typically focus on the effects of differing stellar feedback models (Scannapieco et al., 2012; Kim et al., 2014). But, differences between varied gas thermodynamic prescriptions in simulations is relatively understudied. Previous work has explored the effect of varying the assumed initial temperature and density profiles of halo gas in semi-analytic models, which changes how the gas subsequently cools (e.g. Monaco et al., 2014; Hou et al., 2018, 2019). However, each of these works assumes a non-evolving cooling

function which does not depend on the ionization state of particles in the gas.

Some hydrodynamic simulations now incorporate *non-equilibrium* chemistry, where the simulations calculate element abundances on-the-fly using a chemical network. We can also calculate gas cooling and heating rates from these non-equilibrium abundances. The assumption of gas photoionization equilibrium yields an alternate set of cooling and heating rates. Richings & Schaye (2016) compare isolated galaxy simulations run with gas cooling calculated from non-equilibrium against those run with equilibrium chemical abundances. Capelo et al. (2018) perform a similar comparison, but with the non-equilibrium metal abundances only used to calculate cooling in cold ($T < 10^4$ K) gas. Both works find that non-equilibrium cooling has little effect on the overall star formation rate of the galaxy, but can change the overall amount of molecular gas.

In this chapter, we consider the isolated galaxy simulations of an NGC300 analog presented in Semenov et al. (2021). These simulations were originally used to study the spatial decorrelation between dense molecular gas and sites of recent star formation, and include subgrid models for gas turbulence and turbulence-regulated variable star formation efficiency. We use these simulations to compare two different prescriptions for approximating cooling and heating functions: the interpolation table approach from Gnedin & Hollon (2012) and the machine learning approximation from Chapter 3, which yields more accurate cooling and heating function approximations at fixed metallicity.

Since the only differences between the two simulation runs are in the gas cooling and heating functions, we focus on the comparison between the resulting temperature-density phase diagrams of the simulated gas. These phase diagrams describe the gas fraction in various phases of the ISM and CGM. We also compare CII emission between each simulation run. The luminosity in various ionic emission lines depend on the thermal state (i.e. temperature and density) of the gas. The $157.7\ \mu\text{m}$ fine-structure line emitted by ionized carbon (CII) is a particularly important case. This line is often used to trace molecular, star-forming gas in galaxies (e.g. Zhao et al., 2024; Casavecchia et al., 2025). The CII fine-structure line is also a candidate for Line-Intensity Mapping (LIM) surveys mapping the 3-D structure of the universe. LIM surveys could probe CII emission across a wide range of redshifts $3 \lesssim z \lesssim 9$, stretching back into the EoR (Kovetz et al., 2017).

We explain our methodology, including the simulation code and how we analyze simulated snapshots, in section 5.2. We compare gas phase diagrams and CII emission rates from our two simulation runs in section 5.3, and present conclusions and further discussion in section 5.4.

5.2 Methodology

5.2.1 Simulations

In this work, we use the isolated galaxy simulation of Semenov et al. (2021). This simulation is run using ART, a Eulerian adaptive mesh refinement hydrodynamics code (Kravtsov, 1999; Kravtsov et al., 2002; Rudd et al., 2008).

The specific simulation we use includes a sub-grid prescription for turbulence and a star formation efficiency that depends on the velocity dispersion (including both turbulent and thermal components) through the local virial parameter (Padoan et al., 2012; Semenov et al., 2016, 2021). Different versions of the isolated galaxy simulation incorporate a constant value for star formation efficiency below a maximum virial parameter (or above a minimum density), different contributions to feedback (type II supernovae and/or pre-supernova stellar winds), and iterations with and without radiative transfer (Semenov et al., 2021).

The metallicity-dependent atomic gas cooling and heating functions in the simulation are evaluated using the interpolation table of Gnedin & Hollon (2012), which interpolates between exact cooling and heating rates evaluated with the photoionization code Cloudy (Ferland et al., 1998). This approximation depends on the temperature T , baryon number density n_b , and metallicity Z of the gas, and on rates calculated from the local radiation field: P_{LW} (the photodissociation rate of molecular hydrogen), and the photoionization rates P_{HI} , P_{HeI} , and P_{LW} , in units of $[\text{s}^{-1}]$ (Gnedin & Hollon, 2012).

The thermal energy density U of the gas evolves due to radiative processes as:

$$\left. \frac{dU}{dt} \right|_{\text{rad}} = n_b^2 [\Gamma - \Lambda], \quad (5.1)$$

where Γ is the heating function and Λ is the cooling function. Note that the cooling and heating functions also include additional contributions from non-atomic components of the gas such as molecules that are not included in the interpolation table of Gnedin & Hollon (2012), as they depend on additional properties of the gas. Examples include heating from the photodissociation of molecular hydrogen, and cooling from vibrational and rotational transitions in molecular hydrogen, which both depend on the molecular gas fraction. These contributions are treated separately in the simulation (Gnedin & Kravtsov, 2011), and are not varied in this chapter.

The initial conditions of the fiducial simulation are chosen to approximate observed structural properties of the NGC300 galaxy from Westmeier et al. (2011). The dark matter particles are initialized with a Navarro-Frenk-White profile with mass $M_{200c} \approx 8.3 \times 10^{10} M_{\odot}$ (defined as the mass contained in a sphere with average density equal to 200 times the critical

Parameter	Value
Radiative transfer	Off
Q_{LW}	$2 \times 10^{-11} \text{ s}^{-1}$
Q_{HI}	$2 \times 10^{-17} \text{ s}^{-1}$
Q_{HeI}	$3 \times 10^{-16} \text{ s}^{-1}$
$Q_{\text{C VI}}$	$9 \times 10^{-18} \text{ s}^{-1}$
Z	$0.3Z_{\odot}$
Timestep	1 Myr

Table 5.1: Important parameters for our simulation runs

density) and concentration $c_{200c} \approx 15.4$. The stellar and gas components are initialized as (independent) exponential scale disks. The stellar disk has mass $M = 10^9 M_{\odot}$, scale radius 1.39 kpc, and scale height 0.28 kpc. The gas disk has mass $2.29 \times 10^9 M_{\odot}$ and scale radius 3.44 kpc (the scale height is determined by the ISM pressure gradient). The minimum gas cell size reached in the simulation run is $\Delta = 10$ pc. The fiducial simulation run includes radiative transfer, all feedback processes, and a star formation efficiency depending on the local virial parameter (Semenov et al., 2021).

As described in section 5.2.2, we turn off radiative transfer and use constant photoionization rates. We also fix the gas metallicity. Following Semenov et al. (2021), we initialize our simulations with a snapshot from a fiducial simulation run at $t \approx 600$ Myr. The fiducial simulation turns on various physical processes (such as radiative transfer, gas cooling, and star formation) in stages, allowing the galaxy to ‘settle’ after each stage. By $t \approx 600$ Myr, each relevant process has been incorporated (Semenov et al., 2021). After the initial snapshot, we save snapshots every $\Delta t = 1$ Myr.

Key parameters for our simulations are shown in Table 5.1. We run both simulations for the same duration in simulated cosmic time and find that the thermodynamic properties of the gas in the two simulations have sufficiently converged relative to each other after 5 Myr. This convergence is discussed in more detail in section 5.6.

5.2.2 Cooling and heating function models

The Semenov et al. (2021) simulations calculate metallicity-dependent atomic gas cooling and heating functions using the interpolation table of Gnedin & Hollon (2012), which incorporates a local radiation field including contributions from a synthesized stellar spectrum, quasar-like power law, and absorption by neutral hydrogen and helium. Gnedin & Hollon (2012) constructed these tables by interpolating between exact calculations of cooling and heating function from the Cloudy photoionization code (Ferland et al., 1998). In Chapter 3,

we trained machine learning models on the same Cloudy calculations, with the same gas properties and radiation field parameters as inputs. So, we can directly replace the Gnedin & Hollon (2012) interpolation table with machine learning models from Chapter 3 in the simulation code. The interpolation table of Gnedin & Hollon (2012) implements a quadratic interpolation in metallicity that can lead to occasional unphysical *negative* predicted cooling or heating functions. The machine learning models from Chapter 3 are constructed to always predict positive cooling and heating functions.

Additionally, the interpolation table of Gnedin & Hollon (2012) is known to make ‘catastrophic errors’ at some points in parameter space, due to interpolating a non-linear function on a grid of parameters with fixed spacing. To assess the impact of these catastrophic errors on the simulation, we replace the interpolation table with analogous machine learning models from Chapter 3, constructed using the gradient-boosted tree algorithm XGBoost (Chen & Guestrin, 2016). We label the simulation run using the cooling and heating function interpolation table from Gnedin & Hollon (2012) as ‘GH12’ and the run using the analogous XGBoost models from Chapter 3 as ‘XGB’.

More specifically, we use *fixed metallicity* cooling and heating function models from Chapter 3 with input parameters that parallel the dimensions of the Gnedin & Hollon (2012) interpolation table: temperature T , *hydrogen* number density n_{H} , and rates Q_{LW} , Q_{HI} , Q_{HeI} , and $Q_{\text{C VI}}$, where $Q_j = P_j/n_{\text{H}}$. Separate models are trained to predict the cooling function and the heating function. Both the interpolation table of Gnedin & Hollon (2012) and the machine learning models of Chapter 3 use quadratic interpolation in metallicity to make predictions at arbitrary metallicities, and in Chapter 3 we found that this quadratic interpolation (with the 5 metallicity values present in the training data) is the main limitation to improving accuracy at intermediate metallicities. In order to avoid the complications of the metallicity interpolation, we fix the gas metallicity at $Z = 0.3Z_{\odot}$ (this is one of the metallicities with exact Cloudy calculations we used to train the machine learning models in Chapter 3).

Since the inputs (gas temperature, gas density, and 4 photoionization rates) and outputs (cooling function and heating function) at fixed metallicity are the same, we can perform a one-to-one replacement of the Gnedin & Hollon (2012) interpolation table for the machine learning models of Chapter 3 in the simulation code. For this initial study, because the machine learning models of Chapter 3 are more computationally expensive than the Gnedin & Hollon (2012) interpolation table, we turn off radiative transfer in the ART code. So, the photoionization rates are spatially constant throughout the simulation box. For these rates, we use ISM averages computed from the fiducial simulation run in Semenov et al. (2021) with radiative transfer: $Q_{\text{LW}} = 2 \times 10^{-11} \text{ s}^{-1}$, $Q_{\text{HI}} = 2 \times 10^{-17} \text{ s}^{-1}$, $Q_{\text{HeI}} = 3 \times 10^{-16} \text{ s}^{-1}$, and

	GH12	XGB
Cooling	1.69×10^{-3}	1.28×10^{-5}
Heating	2.19×10^{-2}	3.22×10^{-5}

Table 5.2: MSEs of the cooling and heating function approximations used in our simulations on the training grid of Cloudy calculations at gas metallicity $Z = 0.3Z_{\odot}$. The MSEs for the XGBoost models from Chapter 3 are 2-3 orders of magnitude smaller than those for the Gnedin & Hollon (2012) interpolation table.

$Q_{\text{CVI}} = 9 \times 10^{-18} \text{ s}^{-1}$ (see Table 5.1).

The exact Cloudy calculations used to train both approximations we use are done on a grid of data points with gas temperature $10 \leq T/\text{K} \leq 10^9$ and *hydrogen* number density $10^{-6} \leq n_{\text{H}}/\text{cm}^{-3} \leq 10^6$ (Gnedin & Hollon, 2012). We approximate the conversion between the number densities of hydrogen and baryons as:

$$n_b = \frac{1.4n_{\text{H}}}{1 - 0.02(Z/Z_{\odot})}. \quad (5.2)$$

For $Z = 0.3Z_{\odot}$, $n_b \approx 1.41n_{\text{H}}$. Outside of these ranges, both models are extrapolating from the Cloudy calculations. The table of Gnedin & Hollon (2012) performs a linear extrapolation from the nearest tabulated value. The machine learning approximations of Chapter 3 are piecewise constant, so will simply output the nearest tabulated value.

To compare the approximations on the entire grid of Cloudy computations (which is used to train both models), we compute the MSE $\Delta = \langle (\log \mathcal{F}_{\text{true}} - \log \mathcal{F}_{\text{pred}})^2 \rangle$, where \mathcal{F} is either the cooling function or the heating function. Note that a fraction 2.98×10^{-4} of the grid points result in negative predicted cooling functions using the Gnedin & Hollon (2012) interpolation table. These points are removed from the MSE calculation. The MSEs for the approximations at $Z = 0.3Z_{\odot}$ are shown in Table 5.2, which shows that both models have small errors, but the machine learning models of Chapter 3 produce MSEs that are 2-3 orders of magnitude smaller at this metallicity. More details about the performance comparison can be found in Chapter 3. However, this comparison is only on the data used to construct both models, so does not realistically predict their performance on new data. In the rest of the chapter, we investigate the physical effect of these differences to assess how significant how they are in a simulated galaxy.

Note that, while we can directly replace calls of the interpolation table of Gnedin & Hollon (2012) with evaluations of machine learning models from Chapter 3 in the ART code, this has a major computational cost. Evaluating each timestep of the Semenov et al. (2021) isolated galaxy simulation (see section 5.2.1 for details) is nearly 20 times slower

with the machine learning models. While it is possible that this slowdown could be reduced with careful optimization of the simulation run with machine learning models, a one-to-one replacement for simulations run over a longer timescale or a cosmological volume is simply not practical. For these applications, an interpolation table that can be evaluated quickly *and an understanding of the effects of the errors made by this interpolation* are crucial. But, it is much easier to incorporate additional photoionization rates as inputs in the machine learning framework of Chapter 3 (which includes models using as many as 22 photoionization rate features) than to add an additional dimension to an interpolation table.

5.2.3 CII luminosity

From the simulation data, we can compute quantities that depend on the density and temperature of the gas. As a representative example, here we consider CII emission. When collisionally excited, CII ions can emit a $157.7 \mu\text{m}$ photon. Setting this photon energy equal to kT yields a temperature of 91.2 K (Draine, 2011).

To compare the two simulation runs, we compute the ratio of CII emission due to various excitation channels, j :

$$r_j = \frac{\int f_{\text{XGB}}(n_b, T) R_j(T) dT}{\int f_{\text{GH12}}(n_b, T) R_j(T) dT}. \quad (5.3)$$

where $f(n_b, T)$ is the density-temperature distribution function (the phase diagram, normalized by the total gas mass) and $R_j(T)$ is a temperature-dependent rate factor that encodes the likelihood of a collision inducing the emission of a $157.7 \mu\text{m}$ photon, with units $[\text{erg cm}^3 \text{s}^{-1}]$. The excitation channels j include free electrons, elemental hydrogen and helium, molecular hydrogen H_2 , and CMB photons (Draine, 2011). Note that the only difference between the numerator and denominator of equation (5.3) is the simulation run used to determine the distribution function, while $R_j(T)$ remains the same.

Since $R_j(T)$ appears inside a temperature integral in both in the numerator and denominator of equation (5.3), we only need the temperature dependence of $R_j(T)$ and can neglect multiplicative constants. For $j = \text{H}_2$ (i.e. for collisions with molecular hydrogen), we separate the two spins of molecular hydrogen, and have (Draine, 2011):

$$R_{\text{H}_2 \text{ para}}(T) \propto \left(\frac{T}{100 \text{ K}} \right)^{0.124 - 0.018 \ln(T/(100 \text{ K}))}, \quad (5.4)$$

and

$$R_{\text{H}_2 \text{ ortho}}(T) \propto \left(\frac{T}{100 \text{ K}} \right)^{0.095 + 0.023 \ln(T/(100 \text{ K}))}. \quad (5.5)$$

For collisions with atomic hydrogen, we use (Barinovs et al., 2005):

$$R_{\text{H}}(T) \propto e^{-91.2\text{K}/T} \left(16 + 0.344 \sqrt{\frac{T}{1\text{K}}} - \frac{47.7\text{K}}{T} \right). \quad (5.6)$$

The rate for helium $R_{\text{He}}(T) = 0.38R_{\text{H}}(T)$ (Draine, 2011), and so has identical dependence on the gas temperature.

Collisions with free electrons result in the rate (Tayal, 2009; Draine, 2011):

$$R_e(T) \propto e^{-91.2\text{K}/T} \sqrt{\frac{1\text{K}}{T}}. \quad (5.7)$$

Finally, CII emission can be excited by CMB photons at a rate which depends on the CMB temperature T_{CMB} (Draine, 2011):

$$R_{\text{CMB}}(T) \propto e^{-91.2\text{K}/T_{\text{CMB}}}. \quad (5.8)$$

Since the CMB temperature is independent of the local gas temperature, $R_{\text{CMB}}(T) = \text{const.}$

5.3 Results

Here, we present results comparing our two simulation runs, using the cooling and heating function interpolation table from Gnedin & Hollon (2012), labelled ‘GH12’, and the analogous XGBoost models from Chapter 3, labelled ‘XGB’. We show results from our final snapshot, 5 Myr after our starting snapshot. This is sufficiently long for the new steady-state temperature-density distribution to settle (see section 5.6).

5.3.1 Phase diagrams and residuals

We begin by calculating residuals between the temperature-density phase diagrams of gas in each simulation. Without radiative transfer, simulation snapshots include the gas mass density ρ , thermal energy density U , and molecular weight μ . From these, we calculate the baryon number density n_b and temperature T as follows

$$n_b = \frac{\rho}{m_p}, \quad (5.9)$$

$$T = \frac{2}{3} \frac{\mu}{k n_b} U, \quad (5.10)$$

where m_p is the mass of a proton and k is Boltzmann’s constant.

We select 175 logarithmically spaced bins in temperature with $10 < T/\text{K} < 10^8$ and 200 logarithmically spaced bins in baryon number density with $10^{-7} < n_b/\text{cm}^{-3} < 10^3$ to include nearly all of the gas in the simulation box.

In each density bin, we calculate the median gas temperature, as well as the 10th, 25th, 75th, and 90th percentiles to show the spread. These temperature quantiles are shown for both simulation runs as a function of gas density in the upper panel of Figure 5.1. The ratio of the median temperatures for the two runs is shown in the bottom panel of Figure 5.1. The distributions are generally similar, but the temperature for the XGB run is systematically higher for low gas densities $-3 \lesssim \log(n_b/\text{cm}^{-3}) \lesssim -1$. The bottom panel of Figure 5.1 shows that the median temperature for the XGB run is generally *lower* than for the GH12 run at higher gas densities $-1 \lesssim \log(n_b/\text{cm}^{-3}) \lesssim 1$. However, at these densities, the 25th-75th percentile spread is larger than the difference between median temperatures, so the two simulations do not have systematically different temperatures here.

To further examine the difference in gas thermodynamics between the two simulations, we evaluate the residual in each temperature-density bin:

$$\Delta = \frac{m_{\text{GH12}} - m_{\text{XGB}}}{m_{\text{GH12}} + m_{\text{XGB}}}, \quad (5.11)$$

where m_j is the gas mass in a temperature-density bin for simulation j . The sign of the residual Δ indicates which simulation run has a higher gas mass in a given bin. These residuals are shown for our final snapshot in Figure 5.2.

Figure 5.2 shows that there is a ‘critical curve’ where $\Delta = 0$, indicating that the two simulation runs have identical gas mass in the bins along this curve. The critical curve runs from temperatures of $T \sim 10^4$ K at a density of $n_b \sim 10^{-4} \text{ cm}^{-3}$ down to $T \sim 10$ K at $n_b \sim 10^2 \text{ cm}^{-3}$, and has a non-trivial shape. Just above this curve, the GH12 run has higher gas masses ($\Delta > 0$), while just below it, the XGB run has higher gas masses ($\Delta < 0$). There is also a band at $T \sim 10^4$ K across several orders of magnitude in density where the XGB run has higher gas masses ($\Delta < 0$). This band and the residuals on either side of the critical curve are the most prominent structures in Figure 5.2. At temperatures $T \gtrsim 10^4$ K, there is noise with no clear structure in phase space.

To better understand the critical curve in Figure 5.2, we examine the temperature dependence of cooling and heating functions at fixed density $n_b = 1 \text{ cm}^{-3}$ calculated using the same approach as in the GH12 and XGB simulation runs. This is shown in Figure 5.3. Note that the gas reaches thermal equilibrium at a temperature T_{equib} when $\Gamma(T_{\text{equib}}) = \Lambda(T_{\text{equib}})$ (i.e., when the cooling and heating functions are equal) if no external work is done on it.

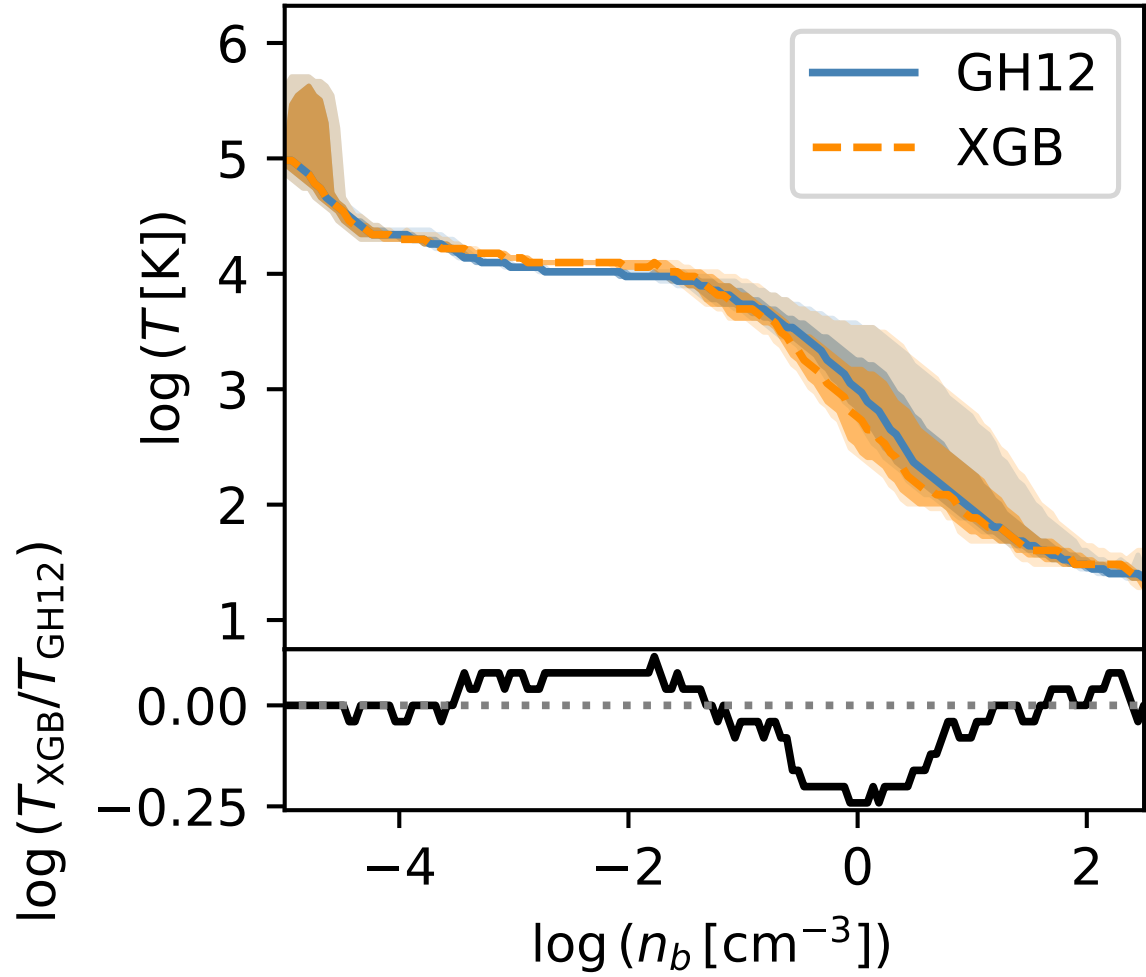


Figure 5.1: Upper panel: Median (curves), 25th-75th percentile (darker bands), and 10th-90th percentile (lighter bands) temperatures as a function of gas density bin for GH12 (blue, solid curve) and XGB (orange, dashed curve) runs. Bottom panel: ratio of median temperature for the XGB run to the GH12 run as a function of density bin.




Figure 5.2: Residual of gas mass between runs using the interpolation table of GH12 and XGBoost cooling and heating functions after 5 Myr. There is a ‘critical curve’ where the residual is equal to 0.




Figure 5.3: Predicted cooling (blue curves) and heating functions (red curves) for the GH12 interpolation table (Gnedin & Hollon, 2012, solid curves) and XGBoost machine learning models from Chapter 3 (dashed curves) at $n_b = 1 \text{ cm}^{-3}$. The cooling and heating curves for a given model intersect at the equilibrium temperature.

Thus, the curves for a given model in Figure 5.3 intersect at the equilibrium temperature for that model.

Figure 5.3 shows that XGBoost predicts a lower equilibrium temperature T_{eq} than does GH12 for $\log(n_b/\text{cm}^{-3}) = 0$. At some value between the two predicted equilibrium temperatures, the two simulation runs will have equal gas masses (the values are consistent with the ‘critical curve’ in Figure 5.2). Just below this critical temperature, $\Gamma > \Lambda$ for GH12, so the gas in the GH12 simulation heats up and gets above the critical temperature. For the XGB simulation, $\Lambda > \Gamma$, so the gas can cool down to its equilibrium temperature. So, the XGBoost simulation should have higher gas mass just below the critical curve, as we see in Figure 5.2. A similar argument predicts that the GH12 simulation should have higher gas mass just above the critical curve.

5.3.2 Effects on CII luminosity

In addition to seeing differences in the temperature-density phase diagram of simulated gas in section 5.3.1, we would also like to understand how these differences affect observable properties of the gas. To investigate this, we plot the ratio r_j defined in equation (5.3) versus n_b in the upper panel of Figure 5.4 for the various processes contributing to the CII emission described in section 5.2.3. In the lower panel, we show the overall density profile for the GH12 simulation run (the profile for the XGB run is nearly identical) to compare the size of the difference in CII emission with where most of the gas lies in phase space.

As shown in Figure 5.4, the density profile is nearly flat between $-4 \lesssim \log(n_b/\text{cm}^{-3}) \lesssim 2$, and sharply decreases for $\log(n_b/\text{cm}^{-3}) \gtrsim 2$. The ratios r_j are always of order 1 for all CII emission processes. The largest values $r_j \approx 1.4$ are reached for collisions with electrons and atomic hydrogen or helium at densities $\log(n_b/\text{cm}^{-3}) \approx 2$, where the gas mass has already begun to decrease from its value at lower densities. At the lower densities where the gas mass is at its plateau for both simulation runs, r_j differs noticeably from 1 for collisions with electrons (as low as 0.9 and as large as 1.2), ortho H_2 (as low as 0.95 and as large as 1.05), and atomic hydrogen and helium (as low as 0.9). At a density of $n_b \sim 10^2 \text{ cm}^{-3}$, all processes shown have $r_j \gtrsim 1$. So, we would expect the actual CII luminosity to be different between XGB and GH12 runs for gas at these densities (and, in particular, for the XGB gas to have a higher CII luminosity). At other densities, some processes have $r_j \gtrsim 1$ while others have $r_j \lesssim 1$, so it is not clear whether the two simulations would have systematically different CII luminosities.

5.4 Summary and discussion

In this chapter, we implement the machine learning approximation scheme for radiation-field-dependent cooling and heating functions (at a fixed metallicity) of Chapter 3 into the ART hydrodynamic simulation code, as an alternative to the interpolation table of Gnedin & Hollon (2012). We compare runs of an isolated galaxy simulation of an NGC300 analog using the two different approximations, without radiative transfer and at fixed gas metallicity. We start the simulations from a fiducial simulation snapshot after the galaxy has ‘settled’ once all relevant physical processes have been added to the simulation. We compare the gas thermodynamics in the two simulations after 5 Myr using their temperature-density phase diagrams. Our main conclusions are:

- The gas in the XGB simulation is systematically hotter for low-density gas with $-3 \lesssim \log(n_b/\text{cm}^{-3}) \lesssim -1$ (see Figure 5.1).

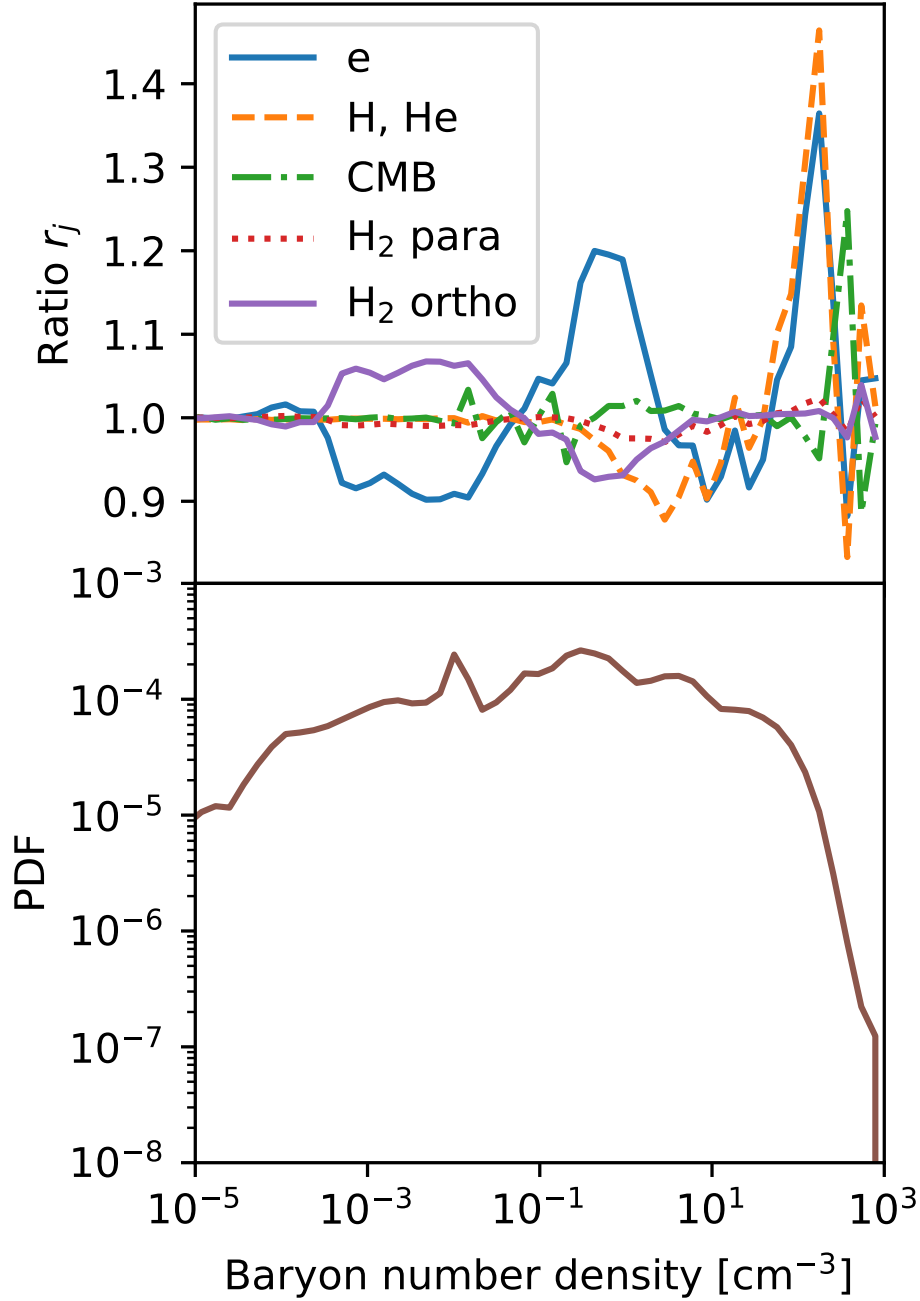


Figure 5.4: The CII emission ratio r_j , as defined in equation (5.3) as a function of baryon number density n_b (upper panel) excited by interactions with electrons (solid blue), atomic hydrogen or helium (dashed orange), CMB photons (dash-dotted green), and molecular hydrogen with para (dashed red) and ortho (solid purple) spins. For comparison, the bottom panel shows the overall baryon number density profile for the GH12 simulation run (solid brown).

- There is a ‘critical curve’ where the two simulation runs have identical gas masses. The largest differences in simulated gas mass occur at temperatures just above and below this critical curve (see Figure 5.2).
- At a given density, the critical curve lies at a temperature between the equilibrium temperatures corresponding to the predicted GH12 and XGB cooling and heating functions at that density (see Figure 5.3)
- The differences in gas thermodynamics result in small, but not necessarily negligible (10-20%) differences in the integrated CII rate for some emission processes (see Figure 5.4).

The net cooling function (the cooling function minus the heating function) determines the temperature of a gas cloud. The thermal properties of the gas are one of the factors that determine its velocity dispersion. In the simulations that we use, the local star formation efficiency is calculated using the local velocity dispersion (Semenov et al., 2016, 2021). So, the differences between the thermal properties of the gas in the two simulations can propagate to differences in the local star formation rate and stellar feedback. To explore these effects in a more realistic setting, we would need simulations including radiative transfer and varying metallicity, over a long enough timescale to look for effects on the star formation history of the galaxy. With radiative transfer, it would also be possible to directly compute CII luminosities for the two simulated galaxies. Here, we take a step towards this goal by investigating the simplified case of gas thermodynamics in a ‘settled’ isolated galaxy simulation with spatially constant metallicity and photoionization rates. This case is interesting because it allows us to explore the direct impact of changes in the gas cooling and heating functions. In a more realistic simulation, this would be difficult to disentangle from the effects of other processes that depend on gas thermodynamics (such as star formation and feedback). Even in this simplified case, we see differences in the gas thermodynamics and CII emission efficiencies for several emission processes.

In a more realistic case with varying gas metallicity across the simulated galaxy, we would expect the difference between the XGB and GH12 runs to be smaller than seen here. This is because the difference in performance between the two approximations is smaller across a sample of points with arbitrary metallicity [see Chapter 3]. However, future machine learning cooling and heating approximations (with access to more training data in metallicity) will likely be able to improve the performance at arbitrary metallicity.

As described in section 5.2.2, the direct implementation of the XGBoost models from Chapter 3 we used in the simulation code resulted in timesteps that were an order of magnitude slower than for simulation using the interpolation table of Gnedin & Hollon (2012). So,

utilizing the models from Chapter 3 in simulations of more massive galaxies or cosmological volumes is computationally infeasible. However, the machine learning approach of Chapter 3 much more easily accommodates additional radiation field dimensions. Adding an additional input to an interpolation table of Gnedin & Hollon (2012) would require increasing the dimension of the table by one, increasing the size of the table that must be stored in memory. Every node of the regression trees used in XGBoost models only considers the value of one input feature (Chen & Guestrin, 2016). The tree depth and number of trees used in the trained models of Chapter 3 do not increase with the number of inputs. This makes it feasible to train models with large numbers of photoionization rate features. So, in cases where more features describing the radiation field are needed, machine learning models could be less computationally expensive than an interpolation table. The machine learning setup of Chapter 3 also allows for the calculation of ‘feature importance’ values describing how much each input affects model predictions. In Chapter 3, we used these feature importances to construct additional machine learning models using different sets of 4 photoionization rates than the set used by the Gnedin & Hollon (2012) interpolation table. This approach can inform the identification of new sets of radiation field properties for constructing new, more accurate cooling and heating function interpolation tables in various regimes.

With these computational constraints, we only considered an isolated galaxy simulation in this work. However, the effects of different cooling and heating function models could depend on galaxy mass, galaxy environment, and redshift. To explore these effects, future work could extend the comparison of simulations with different cooling and heating function models to both a suite of multiple isolated galaxies, and a cosmological volume.

5.5 Acknowledgements

We would like to thank Raziq Noorali for his assistance calculating CII emission rates due to various processes from simulation data.

I, and co-author CA, acknowledge support from the Leinweber Foundation. Co-author CA acknowledges support from DOE grant DE-SC009193. This chapter has been co-authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. Support for co-author VS was provided by Harvard University through the Institute for Theory and Computation Fellowship. This research was also supported in part through computational resources and services provided by Advanced Research Computing (ARC), a division of Information and Technology Services (ITS) at the University of Michigan, Ann Arbor, in particular the Great Lakes cluster and the U-M Research Computing Package.

This work utilizes several Python packages, including XGBoost (Chen & Guestrin, 2016), yt (Turk et al., 2011), and numpy (Harris et al., 2020). The code used to analyze the simulation data and produce the plots can be found at https://github.com/davidbrobins/ngc300_analysis.

5.6 Appendix: Convergence of the gas phase diagram

As discussed in section 5.2.1, we start our simulation runs from fiducial snapshots after the simulated galaxy has settled. So, we need only be concerned with the convergence of the gas thermodynamics due to the different cooling and heating function models. To assess this, we compute a residual between Δ (see equation 5.11) across two snapshots a and b (note that the two simulation runs have snapshots at identical times):

$$\delta_{a,b} = \frac{\Delta_b - \Delta_a}{|\Delta_a + \Delta_b|}. \quad (5.12)$$

Since the residuals at the two timesteps Δ_a and Δ_b can be positive or negative (see Figure 5.2), the absolute value in the denominator is crucial to ensure that $-1 \leq \delta_{a,b} \leq 1$. For convergence, the phase space map of $\delta_{a,b}$ should have no structure. That is, the value of $\delta_{a,b}$ in nearby bins should be as uncorrelated as possible.

To demonstrate the convergence of gas thermodynamics by 5 Myr, we show $\delta_{1\text{Myr},2\text{Myr}}$ at top and $\delta_{4\text{Myr},5\text{Myr}}$ at bottom in Figure 5.5. Going from 1 to 2 Myr in the top panel of Figure 5.5, clear structure can be seen in vertical stripes between $-6 \lesssim \log(n_b/\text{cm}^{-3}) \lesssim -4$ and $4 \lesssim \log(T/\text{K}) \lesssim 6$, as well as a similar structure to the critical curve seen in Figure 5.2. By the timestep from 4 to 5 Myr in the bottom panel of Figure 5.5, the vertical stripes are no longer visible, and the critical curve structure is both much thinner and peaks at lower values of $|\delta|$. The remaining regions in Figure 5.5 show noise with no clear phase space structure. So, we can say that the gas thermodynamics of our simulation runs have sufficiently converged after 5 Myr, and all of the analysis in section 5.3 use the final 5 Myr timestep.




Figure 5.5: Residual between residuals $\delta_{1 \text{ Myr}, 2 \text{ Myr}}$ (top, see equation (5.12)) and $\delta_{4 \text{ Myr}, 5 \text{ Myr}}$ (bottom).

CHAPTER 6

Conclusions and Future Work

Surveys of galaxies can be used to measure the large-scale structure of the universe and constrain cosmological parameters (see section 1.1). These surveys and associated data analysis techniques have reached sufficiently high precision that one of the limiting factors is now systematic uncertainties in modeling how baryonic physics within galaxies affects the observations (e.g. Chisari et al., 2019; Abbott et al., 2022; Krolewski et al., 2025). With recent developments in observational techniques, astronomers are now able to probe the dynamics of star formation in the ISM of individual galaxies at smaller spatial scales than ever before (e.g. Lee et al., 2023; Nagy et al., 2023; Saravia et al., 2025). As discussed in section 1.1, baryonic processes in the ISM and CGM are critical to our understanding of galaxy formation and evolution.

To take full advantage of current and future galaxy observations, we will need accurate and computationally inexpensive models of baryonic processes to implement in high-resolution hydrodynamic galaxy evolution simulations. Gas cooling and heating, the focus of this dissertation, are two such baryonic processes. In this dissertation, we have developed tools to more accurately interpolate atomic gas cooling and heating functions in simulations, and to better understand the effects different cooling and heating function approximations have on the thermal properties of simulated galaxies. These tools can aid in the development of next-generation galaxy formation simulations that can be compared with high-resolution galaxy imaging to constrain the subgrid models in the simulation, or used to model the effects of uncertainties in the baryonic models on large-scale structure survey observations.

In this final chapter, we present a brief summary of the key conclusions from this dissertation in section 6.1. Then we introduce possible paths for future work based on the results and techniques of this dissertation in section 6.2.

6.1 Summary

Thermal processes in the ISM and CGM are a critical part of the process of galaxy formation and evolution (see section 1.1). Interstellar atomic gas radiatively cools predominantly through collisionally-excited line emission, recombination cascades, and free-free emission from free electrons. The gas is radiatively heated through photoionization of atoms, as well as other ionization processes involving solid dust grains and cosmic rays (see section 1.2). The *cooling function* Λ and *heating function* Γ as defined in equation (2.1) are a useful way to parameterize these processes in the context of galaxy evolution simulations.

Gas cooling and heating functions can be evaluated with a photoionization code that calculates the equilibrium gas temperature and distribution of various ionization states, and the associated radiative emission and photoionization (see section 1.3). However, photoionization codes are computationally expensive enough to run that galaxy formation simulations (see section 1.4) often use *interpolation tables* to approximate the gas cooling and heating functions without needing to run a photoionization code on the fly.

Gas cooling and heating functions depend on both the properties of the gas (i.e. its temperature, density, and element abundances) and the incident radiation field. Constructing an interpolation table for these functions requires making assumptions that reduce the number of free parameters necessary to describe the spectrum of the radiation field. A commonly-used example is assuming that the radiation field is spatially uniform (i.e. it consists of a spatially constant extragalactic background).

During the EoR, variations in the radiation field between galaxies inside bubbles that have already been re-ionized by radiation from nearby galaxies and galaxies residing in gas that has yet to be re-ionized are expected to violate this assumption of a spatially uniform (extragalactic) radiation field. In Chapter 2, we calculate the actual cooling and heating rates of the ISM of simulated EoR galaxies from the CROC project. We compare these to ‘median ISM’ cooling and heating functions calculated using the median radiation field, which is spatially constant. We find that the ‘median ISM’ cooling and heating functions differ from the actual cooling and heating rates (see Figure 2.2). This conclusion demonstrates the importance of accounting for spatial variations in the radiation field when approximating atomic gas cooling and heating functions in simulations.

However, while the cooling and heating function interpolation table used in the CROC simulations accounts for the spatially-varying radiation field, it is known to produce large errors at some points in the parameter space due to the non-linearity of cooling and heating functions (Gnedin & Hollon, 2012). Machine learning interpolation techniques (see section 1.5) are a promising avenue to improve on the accuracy of cooling and heating function

interpolation.

In Chapter 3, we develop a machine learning framework to interpolate atomic gas cooling and heating functions more accurately and flexibly than with an interpolation tables. We demonstrate that such an approach can outperform the accuracy of an interpolation table *at fixed gas metallicity* (see Figure 3.5).

Beyond these improvements in accuracy, we also use *interpretable* machine learning tools to compare the *feature importance* of various radiation field parameters and explore what radiation field parameters are best suited to predicting atomic gas cooling and heating functions. In Chapter 4, we extend the machine learning framework of Chapter 3 to incorporate different choices of radiation field parameters. We use feature importance values from trained models to find minimal machine learning approximations that only use 3 radiation field parameters (see Figure 4.4 and Figure 4.5). These 3 parameters describe the radiation field just below the hydrogen ionization threshold (0.5 – 1 Ry), the hydrogen- and neutral helium-ionizing radiation emitted predominantly by stars (1 – 4 Ry), and photons from the quasar-like power law at higher energy which are capable of accessing high-ionization states of metal ions like C VI (13 – 16 Ry). In the future, these feature importance insights can be used to help construct smaller, more computationally efficient interpolation tables for use in galaxy formation simulations without sacrificing accuracy (this is discussed more in section 6.2.4 below).

Because cooling and heating functions are used to thermally evolve gas resolution elements in hydrodynamic simulations, changing how the cooling and heating functions are approximated should change the thermal properties of simulated gas. We test this in an idealized ART simulation of an isolated galaxy simulation in Chapter 5 by comparing the gas temperature-density phase diagrams from two simulation runs, one using the interpolation table of Gnedin & Hollon (2012), and the other using a machine learning approximation from Chapter 3. Our machine learning model predicts slightly hotter gas in regions with low-density regions with $-3 \lesssim \log(n_b/\text{cm}^{-3}) \lesssim -1$ (see Figure 5.1).

We explore the effects of these different gas thermal properties on emission rates for the CII fine-structure line, and find that certain CII emitting-processes have efficiencies that differ by 10-20% between the two simulation runs (see Figure 5.4). This could potentially result in different overall CII luminosities, an observable signature that could be compared with observations of similar galaxies as a test of the modeling assumptions made in the simulations. More realistic (less idealized) future simulations can explore this further, as discussed below in section 6.2.5. We discuss this and other potential directions for future work suggested by the studies in this dissertation in the following section.

6.2 Future work

6.2.1 Metallicity dependence of cooling and heating functions

Gas cooling and heating functions can vary strongly with metallicity. An example of the total cooling function (i.e. including *molecular* cooling) from Maio et al. (2007) is shown in Figure 6.1. The cooling function varies by up to 3 orders of magnitude across this metallicity range at some temperatures, illustrating the importance of accounting for how cooling and heating functions depend on metallicity.

In Chapter 3 and Chapter 4, we trained independent machine learning values at 5 fixed values of the metallicity Z , rather than training cooling or heating function models with Z as a feature. We did this because 5 points in Z (at fixed values of the other features) is too few for XGBoost to fit the metallicity dependence well. However, we found in Chapter 3 that interpolating in Z is the main bottleneck for accurately predicting atomic gas cooling and heating function. This is unsurprising given the strong metallicity dependence seen in Figure 6.1.

We can improve the metallicity interpolation with additional training data. That is, by running additional Cloudy models at intermediate values of the metallicity between those sampled in the existing training data (see Table 3.1). However, it would be computationally expensive to run new models at several values of Z for all values of the other features (i.e. to extend the Cartesian cross product parameter grid described by Table 3.1). The number of additional training points needed could potentially be reduced by fitting Gaussian process regressions (e.g. Neal, 1997) to the previous training data in metallicity with all other parameters fixed. The Gaussian process fits would allow us to determine the metallicity values where the predicted cooling or heating function is most uncertain. Running new Cloudy models at these metallicities would optimally constrain the metallicity dependence (again, with all other parameters fixed). With this extended training data, we could train new XGBoost models, similarly to Chapter 3 and Chapter 4, but now with metallicity as an input feature. This would result in one cooling function model and one heating function model per set of radiation field features.

Different elements and ions have a range of ionization energies. This means that they are photoionized by a variety of photon frequencies and collisionally ionized at various temperatures, so they can have very different cooling and heating rates (see section 1.2). Thus, the total cooling and heating functions can be dependent on the abundances of *individual* elements in addition to the overall metallicity, at least for elements that have ions with important cooling lines or which are easily photoionized.

Freely varying the abundances of individual elements is a formidable challenge. The




Figure 6.1: The total (atomic and molecular, with abundances of $f_{\text{H}_2} = 10^{-5}$ for molecular hydrogen and $f_{\text{HD}} = 10^{-8}$ for hydrogen-deuterium) gas cooling function with $n_{\text{H}} = 1 \text{ cm}^{-3}$ for metal number fractions (relative to hydrogen) of 10^{-6} (black dotted), 10^{-5} (dark blue short-dashed), 10^{-4} (light blue long-dashed), and 10^{-3} (red solid) interpolated from the tables of Sutherland & Dopita (1993). Reproduced from Maio et al. (2007).

abundance of each element would need to be a new dimension of the training data table and a new input to the XGBoost models. Instead of allowing arbitrary variations of each element, Wiersma et al. (2009) suggests approximating the cooling (or heating) due to a specific element by taking the difference between the cooling (or heating) function with all elements present and the cooling (or heating) function with the element in question removed. This approximation effectively assumes that the cooling (or heating) due to each element is *linear* in the abundance of that element, which means that only *one* additional training point is needed for each set of fixed other features.

6.2.2 Non-atomic cooling and heating processes

The machine learning models trained in Chapter 3 and Chapter 4 only included *atomic* gas cooling and heating processes. However, as described in section 1.2.4, there are several non-atomic processes which heat and cool interstellar gas, including molecular cooling, dust photoheating, cooling via gas-grain collisions, and cosmic ray heating.

These non-atomic processes can be included in cooling and heating calculations from the photoionization code Cloudy (Ferland et al., 1998) which was used to construct the training data for the machine learning models of Chapter 3 and Chapter 4. Molecular cooling, dust cooling, dust photoheating, and cosmic ray heating can be incorporated into our machine learning framework by running additional Cloudy models that include these processes and adding them to the training data. Incorporating non-atomic cooling and heating processes would also add additional input features to the machine learning models. These could potentially include the molecular gas fraction, dust-to-metal ratio, and cosmic ray ionization rate.

6.2.3 Other applications of our machine learning framework

The interpretable machine learning framework developed to identify the most important features for predicting gas cooling and heating functions in Chapter 3 and Chapter 4 could also be applied to other contexts.

Predicting other gas properties that are calculated by the photoionization code Cloudy would be a fairly straightforward application. Emissivities for ion lines (including the fine-structure CII line discussed in section 5.2.3) could be particularly useful, as these lines are observable signatures from a variety of astrophysical sources.

The framework could also be applied to any process where we can run detailed and/or high-resolution models to construct training data, but need a less computationally expensive phenomenological approximation to model the process in a larger context. It would be

especially useful if there are many plausible phenomenological parameters to choose from, so that the feature importance techniques developed in Chapter 3 and Chapter 4 could be applied to reduce the dimensionality. The general class of subgrid models for hydrodynamic simulations, particularly those for AGN feedback, discussed in section 1.4.3 fall under this umbrella.

6.2.4 Improved interpolation tables

Interpolation tables are still used to approximate cooling and heating functions in many state-of-the-art galaxy formation simulations (e.g. Schaye et al., 2015; McCarthy et al., 2017; Schaye et al., 2023). The main challenge in constructing such an interpolation table for a general local radiation field is choosing appropriate radiation field parameters as table dimensions. The interpolation needs to be both sufficiently accurate and take up little enough space in memory to be easily deployed in simulations.

The procedure we used to construct machine learning models with minimal numbers of features in Chapter 4 provides a useful tool to find suitable radiation field parameters. With this procedure and trained machine learning models as described in section 6.2.1 and/or section 6.2.2, we can select candidate sets of radiation field parameters to use as interpolation table dimensions. For each set of candidate radiation field dimensions, we can construct cooling and heating function interpolation tables similarly to (Gnedin & Hollon, 2012). Comparing the performance of these interpolation tables would require an independent sample of Cloudy calculations off of the training grid (similar to that used in section 3.2.8).

It is possible that better performance could be obtained by choosing different sets of radiation field parameters in different regimes. For example, Figure 3.4 shows that some photoionization rates for highly ionized metals like CaXX and NaXI can have significantly higher SHAP values for heating function models than for cooling function models. This suggests that different sets of radiation field features could be ‘optimal’ for interpolating heating and cooling functions. Different radiation field parameters at high ($T \geq 10^4$ K) and low ($T < 10^4$ K) temperatures might also improve the accuracy of the interpolation.

Heating function interpolation tables could potentially be improved by changing the definition of the heating function Γ (equation (2.1)) to:

$$\left. \frac{dU}{dt} \right|_{\text{rad}} = n_b \Gamma(T, \dots) - n_b^2 \Lambda(T, \dots), \quad (6.1)$$

where equation (6.1) only factors out *one* factor of the density from the heating term. This is because the heating function is dominated by ionization of ISM particles by photons

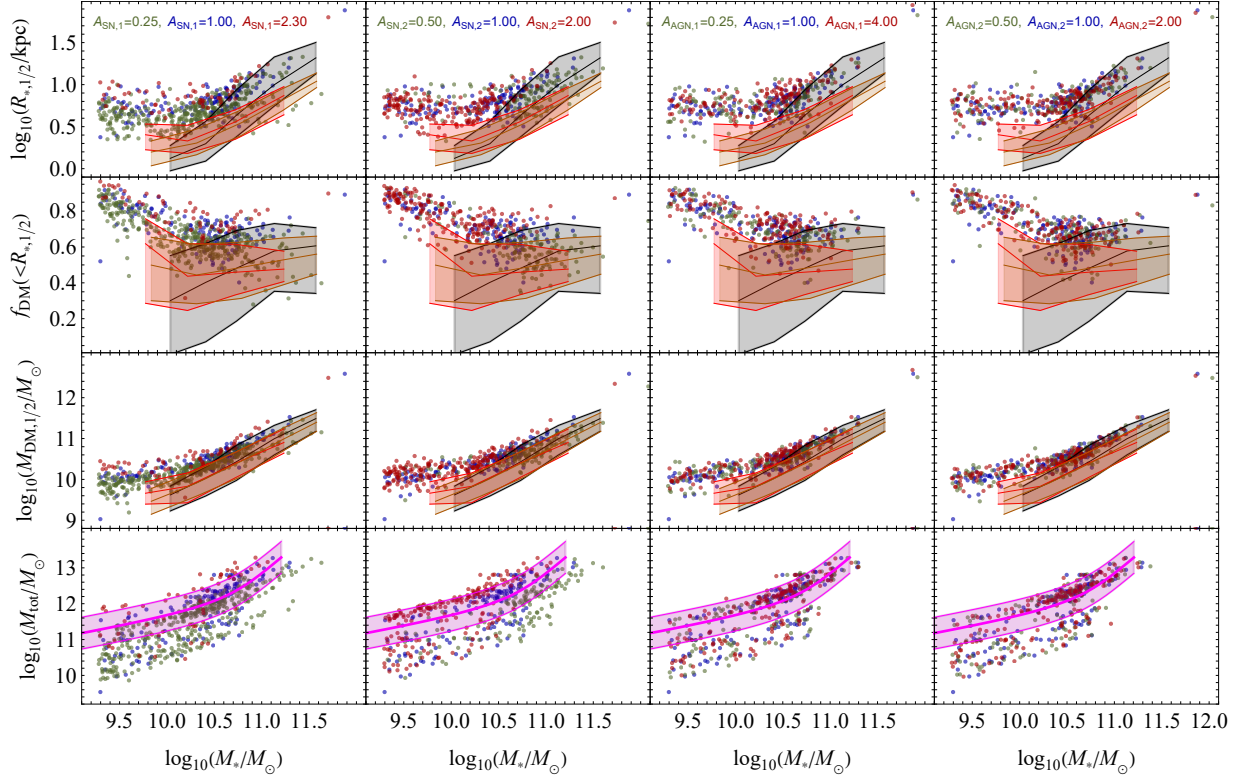


Figure 6.2: Galaxy scaling relations for the radius containing 50% of the total stellar luminosity $R_{*,1/2}$ (top row), dark matter fraction f_{DM} within $R < R_{*,1/2}$ (second row), dark matter mass within $R < R_{*,1/2}$ (third row), and total mass M_{tot} (bottom row) with stellar mass M_* . The circular points show predictions from IllustrisTNG simulations in the CAMELS suite (Villaescusa-Navarro et al., 2021). The shaded bands in the top three rows are observational constraints from La Barbera et al. (2010) (grey), Cappellari et al. (2013) (red), and Bundy et al. (2015) (orange). The pink band in the bottom row is a theoretical relation from Moster et al. (2013). All shaded bands are bounded by the 16th and 84th percentile, with the middle line showing the median. Reproduced from Busillo et al. (2025).

or cosmic rays (see section 1.2), which occur at a rate that is linear in the ISM density rather than quadratic (Pogge, 2008). Using the definition in equation (6.1) should result in a heating function that depends less strongly on the gas density compared to the heating function defined by equation (2.1).

6.2.5 Extended simulation comparisons

The simulation comparison in Chapter 5 uses very idealized galaxy simulations without radiative transfer. We make this choice here because the XGBoost cooling and heating function models from Chapter 3 are more computationally expensive to evaluate than the interpolation table of Gnedin & Hollon (2012) used in the fiducial simulation. Without

radiative transfer, we are able to compare the CII fine-structure line emission rate for various excitation channels between the two simulations, but cannot directly compute the actual CII luminosity in either simulation.

Machine learning-informed interpolation tables such as those described in section 6.2.4 above provide an ideal opportunity to extend similar comparisons to more realistic simulations. With radiative transfer, we will be able to compute observable properties, including the CII luminosity, from these simulations. This will help identify observations that could eventually constrain cooling and heating models. The simulation comparison could also be extended to include a suite of simulations with varying halo masses to search for any trends with the halo mass or other galaxy properties.

Furthermore, interpolation tables make it easy to determine the regimes of parameter space where approximations disagree. This will enable tracing differences in simulation results to particular differences in the interpolation tables.

A long-term goal for these kinds of simulation comparisons is to directly parameterize the cooling and heating function approximations, so that the free parameters of the simulation can be (jointly) *inferred* from comparing the results of simulations with galaxy observations. This approach to inferring simulation parameters is being actively pursued across a variety of hydrodynamic simulation codes by the CAMELS collaboration (e.g. Villaescusa-Navarro et al., 2021). Joint inference is critical because gas thermodynamics is tied to other baryonic processes such as star formation and stellar feedback (see Figure 1.2). Figure 6.2 shows the variation of four scaling relations from IllustrisTNG simulations in the CAMELS suite due to independently changing four subgrid model astrophysical parameters related to supernova feedback ($A_{\text{SN},1}$, $A_{\text{SN},2}$) and AGN feedback ($A_{\text{AGN},1}$, $A_{\text{AGN},2}$), compared to observational data (Busillo et al., 2025). There are degeneracies in the effects each parameter has on the scaling relations, illustrating the necessity of inferring the value of each parameter jointly rather than constraining each parameter independently.

BIBLIOGRAPHY

- Abbott, T. M. C., Aguena, M., Alarcon, A., et al. 2022, *Physical Review D*, 105, 023520, doi: [10.1103/PhysRevD.105.023520](https://doi.org/10.1103/PhysRevD.105.023520)
- Andrae, R., Rix, H.-W., & Chandra, V. 2023, *The Astrophysical Journal Supplement Series*, 267, 8, doi: [10.3847/1538-4365/acd53e](https://doi.org/10.3847/1538-4365/acd53e)
- Angulo, R. E., & Hahn, O. 2022, *Living Reviews in Computational Astrophysics*, 8, doi: [10.1007/s41115-021-00013-z](https://doi.org/10.1007/s41115-021-00013-z)
- Anninos, P., Zhang, Y., Abel, T., & Norman, M. L. 1997, *New Astronomy*, 2, 209, doi: [10.1016/S1384-1076\(97\)00009-2](https://doi.org/10.1016/S1384-1076(97)00009-2)
- Arnaud, M., & Rothenflug, R. 1985, *Astronomy & Astrophysics Supplement Series*, 60, 425. <https://articles.adsabs.harvard.edu/abs/1985A%26AS...60..425A>
- Baes, M., Dejonghe, H., & Davies, J. I. 2005, in *American Institute of Physics Conference Series*, Vol. 761, *The Spectral Energy Distributions of Gas-Rich Galaxies: Confronting Models with Data*, ed. C. C. Popescu & R. J. Tuffs, 27–38, doi: [10.1063/1.1913913](https://doi.org/10.1063/1.1913913)
- Bahcall, J. N., & Wolf, R. A. 1968, *The Astrophysical Journal*, 152, 701, doi: [10.1086/149589](https://doi.org/10.1086/149589)
- Barinova, Ğ., van Hemert, M. C., Krems, R., & Dalgarno, A. 2005, *The Astrophysical Journal*, 620, 537, doi: [10.1086/426860](https://doi.org/10.1086/426860)
- Barnes, J. E., & Hernquist, L. E. 1991, *The Astrophysical Journal Letters*, 370, L65, doi: [10.1086/185978](https://doi.org/10.1086/185978)
- Baumschlager, B., Shen, S., & Wadsley, J. W. 2024, *Astronomy & Astrophysics*, 691, A219, doi: [10.1051/0004-6361/202348164](https://doi.org/10.1051/0004-6361/202348164)
- Bautista, M. A., & Kallman, T. R. 2001, *The Astrophysical Journal Supplement Series*, 134, 139, doi: [10.1086/320363](https://doi.org/10.1086/320363)
- Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, *The Astrophysical Journal*, 762, 109, doi: [10.1088/0004-637x/762/2/109](https://doi.org/10.1088/0004-637x/762/2/109)
- Benson, A. J. 2010, *Physics Reports*, 495, 33, doi: [10.1016/j.physrep.2010.06.001](https://doi.org/10.1016/j.physrep.2010.06.001)
- Bertschinger, E. 1985, *The Astrophysical Journal Supplement Series*, 58, 39, doi: [10.1086/191028](https://doi.org/10.1086/191028)

- Binney, J. 1977, *The Astrophysical Journal*, 215, 483, doi: [10.1086/155378](https://doi.org/10.1086/155378)
- Bocquet, S., Saro, A., Dolag, K., & Mohr, J. J. 2016, *Monthly Notices of the Royal Astronomical Society*, 456, 2361, doi: [10.1093/mnras/stv2657](https://doi.org/10.1093/mnras/stv2657)
- Boehringer, H., & Hensler, G. 1989, *Astronomy & Astrophysics*, 215, 147. <https://ui.adsabs.harvard.edu/abs/1989A%26A...215..147B>
- Bovino, S., Grassi, T., Capelo, P. R., Schleicher, D. R. G., & Banerjee, R. 2016, *Astronomy & Astrophysics*, 590, A15, doi: [10.1051/0004-6361/201628158](https://doi.org/10.1051/0004-6361/201628158)
- Boylan-Kolchin, M., Ma, C.-P., & Quataert, E. 2008, *Monthly Notices of the Royal Astronomical Society*, 383, 93, doi: [10.1111/j.1365-2966.2007.12530.x](https://doi.org/10.1111/j.1365-2966.2007.12530.x)
- Branca, L., & Pallottini, A. 2023, *Monthly Notices of the Royal Astronomical Society*, 518, 5718, doi: [10.1093/mnras/stac3512](https://doi.org/10.1093/mnras/stac3512)
- . 2024, *Astronomy & Astrophysics*, 684, A203, doi: [10.1051/0004-6361/202449193](https://doi.org/10.1051/0004-6361/202449193)
- Brooks, A. M., Governato, F., Quinn, T., Brook, C. B., & Wadsley, J. 2009, *The Astrophysical Journal*, 694, 396, doi: [10.1088/0004-637X/694/1/396](https://doi.org/10.1088/0004-637X/694/1/396)
- Brout, D., Scolnic, D., Popovic, B., et al. 2022, *The Astrophysical Journal*, 938, 110, doi: [10.3847/1538-4357/ac8e04](https://doi.org/10.3847/1538-4357/ac8e04)
- Bundy, K., Bershadsky, M. A., Law, D. R., et al. 2015, *The Astrophysical Journal*, 798, 7, doi: [10.1088/0004-637X/798/1/7](https://doi.org/10.1088/0004-637X/798/1/7)
- Busillo, V., Tortora, C., Covone, G., et al. 2025, *Astronomy & Astrophysics*, 693, A112, doi: [10.1051/0004-6361/202451702](https://doi.org/10.1051/0004-6361/202451702)
- Calderon, V. F., & Berlind, A. A. 2019, *Monthly Notices of the Royal Astronomical Society*, 490, 2367, doi: [10.1093/mnras/stz2775](https://doi.org/10.1093/mnras/stz2775)
- Cantalupo, S. 2010, *Monthly Notices of the Royal Astronomical Society*, 403, L16, doi: [10.1111/j.1745-3933.2010.00806.x](https://doi.org/10.1111/j.1745-3933.2010.00806.x)
- Capelo, P. R., Bovino, S., Lupi, A., Schleicher, D. R. G., & Grassi, T. 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 3283, doi: [10.1093/mnras/stx3355](https://doi.org/10.1093/mnras/stx3355)
- Cappellari, M., McDermid, R. M., Alatalo, K., et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 1862, doi: [10.1093/mnras/stt644](https://doi.org/10.1093/mnras/stt644)
- Casavecchia, B., Maio, U., Péroux, C., & Ciardi, B. 2025, *Astronomy & Astrophysics*, 693, A119, doi: [10.1051/0004-6361/202452282](https://doi.org/10.1051/0004-6361/202452282)
- Ceverino, D., Klypin, A., Klimek, E. S., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 442, 1545–1559, doi: [10.1093/mnras/stu956](https://doi.org/10.1093/mnras/stu956)
- Chabrier, G. 2003, *Publications of the Astronomical Society of the Pacific*, 115, 763, doi: [10.1086/376392](https://doi.org/10.1086/376392)

- Chandrasekhar, S. 1943, *The Astrophysical Journal*, 97, 255, doi: [10.1086/144517](https://doi.org/10.1086/144517)
- Chang, Y.-Y., Hsieh, B.-C., Wang, W.-H., et al. 2021, *The Astrophysical Journal*, 920, 68, doi: [10.3847/1538-4357/ac167c](https://doi.org/10.3847/1538-4357/ac167c)
- Chatzikos, M., Bianchi, S., Camilloni, F., et al. 2023, *Revista Mexicana de Astronomía y Astrofísica*, 59, 327, doi: [10.22201/ia.01851101p.2023.59.02.12](https://doi.org/10.22201/ia.01851101p.2023.59.02.12)
- Chen, T., & Guestrin, C. 2016, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785, doi: [10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785)
- Cheung, A. C., Rank, D. M., Townes, C. H., Thornton, D. D., & Welch, W. J. 1968, *Physical Review Letters*, 21, 1701, doi: [10.1103/PhysRevLett.21.1701](https://doi.org/10.1103/PhysRevLett.21.1701)
- Chisari, N. E., Alonso, D., Krause, E., et al. 2019, *The Astrophysical Journal Supplement Series*, 242, 2, doi: [10.3847/1538-4365/ab1658](https://doi.org/10.3847/1538-4365/ab1658)
- Cohen, D. H., Cassinelli, J. P., Macfarlane, J. J., & Owocki, S. P. 2000, in *Astronomical Society of the Pacific Conference Series*, Vol. 204, *Thermal and Ionization Aspects of Flows from Hot Stars*, ed. H. Lamers & A. Sapar, 65, doi: [10.48550/arXiv.astro-ph/0008351](https://doi.org/10.48550/arXiv.astro-ph/0008351)
- Cole, S., Aragon-Salamanca, A., Frenk, C. S., Navarro, J. F., & Zepf, S. E. 1994, *Monthly Notices of the Royal Astronomical Society*, 271, 781, doi: [10.1093/mnras/271.4.781](https://doi.org/10.1093/mnras/271.4.781)
- Colpi, M., Mayer, L., & Governato, F. 1999, *The Astrophysical Journal*, 525, 720, doi: [10.1086/307952](https://doi.org/10.1086/307952)
- Cox, D. P., & Tucker, W. H. 1969, *The Astrophysical Journal*, 157, 1157, doi: [10.1086/150144](https://doi.org/10.1086/150144)
- Croton, D. J., Springel, V., White, S. D. M., et al. 2006, *Monthly Notices of the Royal Astronomical Society*, 365, 11, doi: [10.1111/j.1365-2966.2005.09675.x](https://doi.org/10.1111/j.1365-2966.2005.09675.x)
- Dalgarno, A., & McCray, R. A. 1972, *Annual Review of Astronomy and Astrophysics*, 10, 375, doi: [10.1146/annurev.aa.10.090172.002111](https://doi.org/10.1146/annurev.aa.10.090172.002111)
- Dang, Y., Chen, Z., Li, H., & Shu, H. 2022, *Applied Artificial Intelligence*, 36, doi: [10.1080/08839514.2022.2074129](https://doi.org/10.1080/08839514.2022.2074129)
- Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 486, 2827, doi: [10.1093/mnras/stz937](https://doi.org/10.1093/mnras/stz937)
- de Jong, T. 1977, *Astronomy & Astrophysics*, 55, 137. <https://ui.adsabs.harvard.edu/abs/1977A%26A....55..137D>
- De Lucia, G., Springel, V., White, S. D. M., Croton, D., & Kauffmann, G. 2006, *Monthly Notices of the Royal Astronomical Society*, 366, 499, doi: [10.1111/j.1365-2966.2005.09879.x](https://doi.org/10.1111/j.1365-2966.2005.09879.x)
- Dobbs, C. 2023, *Frontiers in Astronomy and Space Sciences*, 10, 1272771, doi: [10.3389/fspas.2023.1272771](https://doi.org/10.3389/fspas.2023.1272771)

- Draine, B. T. 1978, *The Astrophysical Journal Supplement Series*, 36, 595, doi: [10.1086/190513](https://doi.org/10.1086/190513)
- Draine, B. T. 2004, in *The Cold Universe*, ed. A. W. Blain, F. Combes, B. T. Draine, D. Pfenniger, & Y. Revaz, 213, doi: [10.48550/arXiv.astro-ph/0304488](https://doi.org/10.48550/arXiv.astro-ph/0304488)
- . 2011, *Physics of the Interstellar and Intergalactic Medium* (Princeton University Press)
- Draine, B. T., & Bertoldi, F. 1996, *The Astrophysical Journal*, 468, 269, doi: [10.1086/177689](https://doi.org/10.1086/177689)
- Dubois, Y., Pichon, C., Welker, C., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 444, 1453, doi: [10.1093/mnras/stu1227](https://doi.org/10.1093/mnras/stu1227)
- Dumont, A. M., Abrassart, A., & Collin, S. 2000, *Astronomy & Astrophysics*, 357, 823, doi: [10.48550/arXiv.astro-ph/0003220](https://doi.org/10.48550/arXiv.astro-ph/0003220)
- Dwek, E. 1998, *The Astrophysical Journal*, 501, 643, doi: [10.1086/305829](https://doi.org/10.1086/305829)
- Efstathiou, G. 1992, *Monthly Notices of the Royal Astronomical Society*, 256, 43P, doi: [10.1093/mnras/256.1.43P](https://doi.org/10.1093/mnras/256.1.43P)
- Ercolano, B. 2005, in *American Institute of Physics Conference Series*, Vol. 804, *Planetary Nebulae as Astronomical Tools*, ed. R. Szczerba, G. Stasińska, & S. K. Gorny (AIP), 35–43, doi: [10.1063/1.2146222](https://doi.org/10.1063/1.2146222)
- Ercolano, B., Barlow, M. J., Storey, P. J., & Liu, X. W. 2003, *Monthly Notices of the Royal Astronomical Society*, 340, 1136, doi: [10.1046/j.1365-8711.2003.06371.x](https://doi.org/10.1046/j.1365-8711.2003.06371.x)
- Ewen, H. I., & Purcell, E. M. 1951, *Nature*, 168, 356, doi: [10.1038/168356a0](https://doi.org/10.1038/168356a0)
- Faerman, Y., Pandya, V., Somerville, R. S., & Sternberg, A. 2022, *The Astrophysical Journal*, 928, 37, doi: [10.3847/1538-4357/ac4ca6](https://doi.org/10.3847/1538-4357/ac4ca6)
- Fardal, M. A., Katz, N., Gardner, J. P., et al. 2001, *The Astrophysical Journal*, 562, 605, doi: [10.1086/323519](https://doi.org/10.1086/323519)
- Farouki, R. T., & Shapiro, S. L. 1982, *The Astrophysical Journal*, 259, 103, doi: [10.1086/160151](https://doi.org/10.1086/160151)
- Faucher-Giguère, C.-A. 2020, *Monthly Notices of the Royal Astronomical Society*, 493, 1614–1632, doi: [10.1093/mnras/staa302](https://doi.org/10.1093/mnras/staa302)
- Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, *The Astrophysical Journal*, 703, 1416–1443, doi: [10.1088/0004-637x/703/2/1416](https://doi.org/10.1088/0004-637x/703/2/1416)
- Faucher-Giguère, C.-A., & Oh, S. P. 2023, *Annual Review of Astronomy and Astrophysics*, 61, 131, doi: [10.1146/annurev-astro-052920-125203](https://doi.org/10.1146/annurev-astro-052920-125203)
- Ferland, G. J. 1993, in *Planetary Nebulae*, ed. R. Weinberger & A. Acker, Vol. 155, 123. <https://ui.adsabs.harvard.edu/abs/1993IAUS..155..123F>

- Ferland, G. J. 2009, *Astronomy & Astrophysics*, 500, 299, doi: [10.1051/0004-6361/200912165](https://doi.org/10.1051/0004-6361/200912165)
- Ferland, G. J., Korista, K. T., & Verner, D. A. 1997, in *Astronomical Society of the Pacific Conference Series*, Vol. 125, *Astronomical Data Analysis Software and Systems VI*, ed. G. Hunt & H. Payne, 213. <https://ui.adsabs.harvard.edu/abs/1997ASPC..125..213F>
- Ferland, G. J., Korista, K. T., Verner, D. A., et al. 1998, *Publications of the Astronomical Society of the Pacific*, 110, 761, doi: [10.1086/316190](https://doi.org/10.1086/316190)
- Ferland, G. J., Porter, R. L., van Hoof, P. A. M., et al. 2013, *Revista Mexicana de Astronomia y Astrofisica*, 49, 137, doi: [10.48550/arXiv.1302.4485](https://doi.org/10.48550/arXiv.1302.4485)
- Ferland, G. J., Chatzikos, M., Guzmán, F., et al. 2017, *Revista Mexicana de Astronomia y Astrofisica*, 53, 385, doi: [10.48550/arXiv.1705.10877](https://doi.org/10.48550/arXiv.1705.10877)
- Field, G. B., Goldsmith, D. W., & Habing, H. J. 1969, *The Astrophysical Journal*, 155, L149, doi: [10.1086/180324](https://doi.org/10.1086/180324)
- Friedman, J. H. 2001, *The Annals of Statistics*, 29, 1189, doi: [10.1214/aos/1013203451](https://doi.org/10.1214/aos/1013203451)
- Fu, Y., Wu, X.-B., Yang, Q., et al. 2021, *The Astrophysical Journal Supplement Series*, 254, 6, doi: [10.3847/1538-4365/abe85e](https://doi.org/10.3847/1538-4365/abe85e)
- Galligan, T. P., Katz, H., Kimm, T., et al. 2019, *arXiv e-prints*, arXiv:1901.01264, doi: [10.48550/arXiv.1901.01264](https://doi.org/10.48550/arXiv.1901.01264)
- Girichidis, P., Offner, S. S. R., Kritsuk, A. G., et al. 2020, *Space Science Reviews*, 216, 68, doi: [10.1007/s11214-020-00693-8](https://doi.org/10.1007/s11214-020-00693-8)
- Gnat, O., & Ferland, G. J. 2012, *The Astrophysical Journal Supplement Series*, 199, 20, doi: [10.1088/0067-0049/199/1/20](https://doi.org/10.1088/0067-0049/199/1/20)
- Gnat, O., & Sternberg, A. 2007, *The Astrophysical Journal Supplement Series*, 168, 213, doi: [10.1086/509786](https://doi.org/10.1086/509786)
- Gnedin, N. Y. 2014, *The Astrophysical Journal*, 793, 29, doi: [10.1088/0004-637X/793/1/29](https://doi.org/10.1088/0004-637X/793/1/29)
- Gnedin, N. Y., & Abel, T. 2001, *New Astronomy*, 6, 437, doi: [10.1016/S1384-1076\(01\)00068-9](https://doi.org/10.1016/S1384-1076(01)00068-9)
- Gnedin, N. Y., & Hollon, N. 2012, *The Astrophysical Journal Supplement Series*, 202, 13, doi: [10.1088/0067-0049/202/2/13](https://doi.org/10.1088/0067-0049/202/2/13)
- Gnedin, N. Y., & Kaurov, A. A. 2014, *The Astrophysical Journal*, 793, 30, doi: [10.1088/0004-637X/793/1/30](https://doi.org/10.1088/0004-637X/793/1/30)
- Gnedin, N. Y., & Kravtsov, A. V. 2011, *The Astrophysical Journal*, 728, 88, doi: [10.1088/0004-637X/728/2/88](https://doi.org/10.1088/0004-637X/728/2/88)
- Gnedin, N. Y., & Madau, P. 2022, *Living Reviews in Computational Astrophysics*, 8, 3, doi: [10.1007/s41115-022-00015-5](https://doi.org/10.1007/s41115-022-00015-5)

- Goldsmith, D. W., Habing, H. J., & Field, G. B. 1969, *The Astrophysical Journal*, 158, 173, doi: [10.1086/150181](https://doi.org/10.1086/150181)
- Golob, A., Sawicki, M., Goulding, A. D., & Coupon, J. 2021, *Monthly Notices of the Royal Astronomical Society*, 503, 4136, doi: [10.1093/mnras/stab719](https://doi.org/10.1093/mnras/stab719)
- Grand, R. J. J., Gómez, F. A., Marinacci, F., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 467, 179, doi: [10.1093/mnras/stx071](https://doi.org/10.1093/mnras/stx071)
- Grassi, T., Bovino, S., Schleicher, D. R. G., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 439, 2386, doi: [10.1093/mnras/stu114](https://doi.org/10.1093/mnras/stu114)
- Grassi, T., Krstic, P., Merlin, E., et al. 2011, *Astronomy & Astrophysics*, 533, A123, doi: [10.1051/0004-6361/200913779](https://doi.org/10.1051/0004-6361/200913779)
- Grinsztajn, L., Oyallon, E., & Varoquaux, G. 2022, in *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22 No. 37* (Red Hook, NY, USA: Curran Associates Inc.), 507–520. <https://dl.acm.org/doi/10.5555/3600270.3600307>
- Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, *The Astrophysical Journal*, 742, 76, doi: [10.1088/0004-637X/742/2/76](https://doi.org/10.1088/0004-637X/742/2/76)
- Guhathakurta, P., & Draine, B. T. 1989, *The Astrophysical Journal*, 345, 230, doi: [10.1086/167899](https://doi.org/10.1086/167899)
- Gutcke, T. A., Pakmor, R., Naab, T., & Springel, V. 2021, *Monthly Notices of the Royal Astronomical Society*, 501, 5597, doi: [10.1093/mnras/staa3875](https://doi.org/10.1093/mnras/staa3875)
- Haardt, F., & Madau, P. 2001, in *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*, ed. D. M. Neumann & J. T. V. Tran, 64. <https://arxiv.org/abs/astro-ph/0106018>
- Haardt, F., & Madau, P. 2012, *The Astrophysical Journal*, 746, 125, doi: [10.1088/0004-637X/746/2/125](https://doi.org/10.1088/0004-637X/746/2/125)
- Habing, H. J., & Goldsmith, D. W. 1971, *The Astrophysical Journal*, 166, 525, doi: [10.1086/150979](https://doi.org/10.1086/150979)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Hayden, M. R., Sharma, S., Bland-Hawthorn, J., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 517, 5325, doi: [10.1093/mnras/stac2787](https://doi.org/10.1093/mnras/stac2787)
- Head, T., Kumar, M., Nahrstaedt, H., Louppe, G., & Shcherbatyi, I. 2021, *scikit-optimize/scikit-optimize*, v0.9.0, Zenodo, doi: [10.5281/zenodo.5565057](https://doi.org/10.5281/zenodo.5565057)
- Hennebelle, P., & Grudić, M. Y. 2024, *Annual Review of Astronomy and Astrophysics*, 62, 63, doi: [10.1146/annurev-astro-052622-031748](https://doi.org/10.1146/annurev-astro-052622-031748)

- Hernquist, L. 1989, *Nature*, 340, 687, doi: [10.1038/340687a0](https://doi.org/10.1038/340687a0)
- Heyl, J., Butterworth, J., & Viti, S. 2023, *Monthly Notices of the Royal Astronomical Society*, 526, 404, doi: [10.1093/mnras/stad2814](https://doi.org/10.1093/mnras/stad2814)
- Hjellming, R. M., Gordon, C. P., & Gordon, K. J. 1969, *Astronomy & Astrophysics*, 2, 202. <https://ui.adsabs.harvard.edu/abs/1969A%26A.....2..202H>
- Hopkins, P. F., Quataert, E., & Murray, N. 2011, *Monthly Notices of the Royal Astronomical Society*, 417, 950, doi: [10.1111/j.1365-2966.2011.19306.x](https://doi.org/10.1111/j.1365-2966.2011.19306.x)
- Hou, J., Lacey, C. G., & Frenk, C. S. 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 543, doi: [10.1093/mnras/stx3218](https://doi.org/10.1093/mnras/stx3218)
- . 2019, *Monthly Notices of the Royal Astronomical Society*, 486, 1691, doi: [10.1093/mnras/stz730](https://doi.org/10.1093/mnras/stz730)
- Hubble, E. 1929, *Proceedings of the National Academy of Science*, 15, 168, doi: [10.1073/pnas.15.3.168](https://doi.org/10.1073/pnas.15.3.168)
- Hughes, A. C. N., Bailer-Jones, C. A. L., & Jamal, S. 2022, *Astronomy & Astrophysics*, 668, A99, doi: [10.1051/0004-6361/202244859](https://doi.org/10.1051/0004-6361/202244859)
- Hunter, J. D. 2007, *Computing in Science & Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Incatasciato, A., Khochfar, S., & Oñorbe, J. 2023, *Monthly Notices of the Royal Astronomical Society*, 522, 330, doi: [10.1093/mnras/stad1008](https://doi.org/10.1093/mnras/stad1008)
- Ivanov, S., Tsizh, M., Ullmann, D., Panos, B., & Voloshynovskiy, S. 2021, *Astronomy and Computing*, 36, 100473, doi: [10.1016/j.ascom.2021.100473](https://doi.org/10.1016/j.ascom.2021.100473)
- Jia, W., Sun, M., Lian, J., & Hou, S. 2022, *Complex & Intelligent Systems*, 8, 2663–2693, doi: [10.1007/s40747-021-00637-x](https://doi.org/10.1007/s40747-021-00637-x)
- Jin, X., Zhang, Y., Zhang, J., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 485, 4539, doi: [10.1093/mnras/stz680](https://doi.org/10.1093/mnras/stz680)
- Joseph, R. D., & Wright, G. S. 1985, *Monthly Notices of the Royal Astronomical Society*, 214, 87, doi: [10.1093/mnras/214.2.87](https://doi.org/10.1093/mnras/214.2.87)
- Jura, M. 1976, *The Astrophysical Journal*, 204, 12, doi: [10.1086/154145](https://doi.org/10.1086/154145)
- Kallman, T. R. 2001, in *Astronomical Society of the Pacific Conference Series*, Vol. 247, *Spectroscopic Challenges of Photoionized Plasmas*, ed. G. Ferland & D. W. Savin, 175. <https://ui.adsabs.harvard.edu/abs/2001ASPC..247..175K>
- Kallman, T. R., & McCray, R. 1982, *The Astrophysical Journal Supplement Series*, 50, 263, doi: [10.1086/190828](https://doi.org/10.1086/190828)
- Kannan, R., Garaldi, E., Smith, A., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 511, 4005, doi: [10.1093/mnras/stab3710](https://doi.org/10.1093/mnras/stab3710)

- Kannan, R., Marinacci, F., Vogelsberger, M., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 499, 5732, doi: [10.1093/mnras/staa3249](https://doi.org/10.1093/mnras/staa3249)
- Kannan, R., Stinson, G. S., Macciò, A. V., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 437, 2882–2893, doi: [10.1093/mnras/stt2098](https://doi.org/10.1093/mnras/stt2098)
- Kauffmann, G., White, S. D. M., & Guiderdoni, B. 1993, *Monthly Notices of the Royal Astronomical Society*, 264, 201, doi: [10.1093/mnras/264.1.201](https://doi.org/10.1093/mnras/264.1.201)
- Kennicutt, Jr., R. C. 1998, *The Astrophysical Journal*, 498, 541, doi: [10.1086/305588](https://doi.org/10.1086/305588)
- Khandai, N., Di Matteo, T., Croft, R., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 450, 1349, doi: [10.1093/mnras/stv627](https://doi.org/10.1093/mnras/stv627)
- Kim, J.-G., Gong, M., Kim, C.-G., & Ostriker, E. C. 2023, *The Astrophysical Journal Supplement Series*, 264, 10, doi: [10.3847/1538-4365/ac9b1d](https://doi.org/10.3847/1538-4365/ac9b1d)
- Kim, J.-h., Abel, T., Agertz, O., et al. 2014, *The Astrophysical Journal Supplement Series*, 210, 14, doi: [10.1088/0067-0049/210/1/14](https://doi.org/10.1088/0067-0049/210/1/14)
- Kinkhabwala, A., Behar, E., Sako, M., et al. 2003, arXiv e-prints, astro, doi: [10.48550/arXiv.astro-ph/0304332](https://doi.org/10.48550/arXiv.astro-ph/0304332)
- Kovetz, E. D., Viero, M. P., Lidz, A., et al. 2017, arXiv e-prints, arXiv:1709.09066, doi: [10.48550/arXiv.1709.09066](https://doi.org/10.48550/arXiv.1709.09066)
- Kramida, A., Yu. Ralchenko, Reader, J., & and NIST ASD Team. 2024, NIST Atomic Spectra Database (ver. 5.12), [Online]. Available: <https://physics.nist.gov/asd> [2025, April 3]. National Institute of Standards and Technology, Gaithersburg, MD.
- Kravtsov, A. V. 1999, PhD thesis, New Mexico State University
- . 2003, *The Astrophysical Journal Letters*, 590, L1, doi: [10.1086/376674](https://doi.org/10.1086/376674)
- Kravtsov, A. V., Klypin, A., & Hoffman, Y. 2002, *The Astrophysical Journal*, 571, 563, doi: [10.1086/340046](https://doi.org/10.1086/340046)
- Kravtsov, A. V., Klypin, A. A., & Khokhlov, A. M. 1997, *The Astrophysical Journal Supplement Series*, 111, 73, doi: [10.1086/313015](https://doi.org/10.1086/313015)
- Krolewski, A., Yu, J., Ross, A. J., et al. 2025, *Journal of Cosmology and Astroparticle Physics*, 2025, 147, doi: [10.1088/1475-7516/2025/01/147](https://doi.org/10.1088/1475-7516/2025/01/147)
- Kroupa, P. 2001, *Monthly Notices of the Royal Astronomical Society*, 322, 231, doi: [10.1046/j.1365-8711.2001.04022.x](https://doi.org/10.1046/j.1365-8711.2001.04022.x)
- Kunsági-Máté, S., Beck, R., Szapudi, I., & Csabai, I. 2022, *Monthly Notices of the Royal Astronomical Society*, 516, 2662, doi: [10.1093/mnras/stac2411](https://doi.org/10.1093/mnras/stac2411)
- La Barbera, F., de Carvalho, R. R., de La Rosa, I. G., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 408, 1313, doi: [10.1111/j.1365-2966.2010.16850.x](https://doi.org/10.1111/j.1365-2966.2010.16850.x)

- Lacey, C., & Cole, S. 1993, *Monthly Notices of the Royal Astronomical Society*, 262, 627, doi: [10.1093/mnras/262.3.627](https://doi.org/10.1093/mnras/262.3.627)
- Lee, J. C., Sandstrom, K. M., Leroy, A. K., et al. 2023, *The Astrophysical Journal Letters*, 944, L17, doi: [10.3847/2041-8213/acaaae](https://doi.org/10.3847/2041-8213/acaaae)
- Lee, Y.-N. 2024, *Astronomy & Astrophysics*, 684, A48, doi: [10.1051/0004-6361/202346533](https://doi.org/10.1051/0004-6361/202346533)
- Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, *The Astrophysical Journal Supplement Series*, 123, 3, doi: [10.1086/313233](https://doi.org/10.1086/313233)
- Li, C., Zhang, Y., Cui, C., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 506, 1651, doi: [10.1093/mnras/stab1650](https://doi.org/10.1093/mnras/stab1650)
- Lucey, M., Al Kharusi, N., Hawkins, K., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 523, 4049, doi: [10.1093/mnras/stad1675](https://doi.org/10.1093/mnras/stad1675)
- Lundberg, S. M., Erion, G. G., & Lee, S.-I. 2018, arXiv e-prints, arXiv:1802.03888, doi: [10.48550/arXiv.1802.03888](https://doi.org/10.48550/arXiv.1802.03888)
- Lundberg, S. M., & Lee, S.-I. 2017, in *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17* (Red Hook, NY, USA: Curran Associates Inc.), 4768–4777. <https://dl.acm.org/doi/10.5555/3295222.3295230>
- Lundberg, S. M., Erion, G., Chen, H., et al. 2020, *Nature Machine Intelligence*, 2, 56, doi: [10.1038/s42256-019-0138-9](https://doi.org/10.1038/s42256-019-0138-9)
- Luo, J.-W., Wang, F.-F., Zhu-Ge, J.-M., et al. 2023, *The Astrophysical Journal*, 959, 44, doi: [10.3847/1538-4357/ad03ec](https://doi.org/10.3847/1538-4357/ad03ec)
- Lykins, M. L., Ferland, G. J., Porter, R. L., et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 429, 3133, doi: [10.1093/mnras/sts570](https://doi.org/10.1093/mnras/sts570)
- Machado Poletti Valle, L. F., Avestruz, C., Barnes, D. J., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 507, 1468, doi: [10.1093/mnras/stab2252](https://doi.org/10.1093/mnras/stab2252)
- Maio, U., Dolag, K., Ciardi, B., & Tornatore, L. 2007, *Monthly Notices of the Royal Astronomical Society*, 379, 963, doi: [10.1111/j.1365-2966.2007.12016.x](https://doi.org/10.1111/j.1365-2966.2007.12016.x)
- Martínez-Serrano, F. J., Serna, A., Domínguez-Tenreiro, R., & Mollá, M. 2008, *Monthly Notices of the Royal Astronomical Society*, 388, 39, doi: [10.1111/j.1365-2966.2008.13383.x](https://doi.org/10.1111/j.1365-2966.2008.13383.x)
- McCarthy, I. G., Schaye, J., Bird, S., & Le Brun, A. M. C. 2017, *Monthly Notices of the Royal Astronomical Society*, 465, 2936, doi: [10.1093/mnras/stw2792](https://doi.org/10.1093/mnras/stw2792)
- McKee, C. F., & Ostriker, E. C. 2007, *Annual Review of Astronomy and Astrophysics*, 45, 565, doi: [10.1146/annurev.astro.45.051806.110602](https://doi.org/10.1146/annurev.astro.45.051806.110602)
- McKee, C. F., & Ostriker, J. P. 1977, *The Astrophysical Journal*, 218, 148, doi: [10.1086/155667](https://doi.org/10.1086/155667)

- McKee, C. F., Storey, J. W. V., Watson, D. M., & Green, S. 1982, *The Astrophysical Journal*, 259, 647, doi: [10.1086/160200](https://doi.org/10.1086/160200)
- Merritt, D. 2000, in *Astronomical Society of the Pacific Conference Series*, Vol. 197, *Dynamics of Galaxies: from the Early Universe to the Present*, ed. F. Combes, G. A. Mamon, & V. Charmandaris, 221, doi: [10.48550/arXiv.astro-ph/9910546](https://doi.org/10.48550/arXiv.astro-ph/9910546)
- Mirabal, N., Charles, E., Ferrara, E. C., et al. 2016, *The Astrophysical Journal*, 825, 69, doi: [10.3847/0004-637X/825/1/69](https://doi.org/10.3847/0004-637X/825/1/69)
- Mo, H., van den Bosch, F. C., & White, S. 2010, *Galaxy Formation and Evolution* (Cambridge University Press)
- Monaco, P., Benson, A. J., De Lucia, G., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 441, 2058, doi: [10.1093/mnras/stu655](https://doi.org/10.1093/mnras/stu655)
- Morisset, C., Stasińska, G., & Peña, M. 2005, *Monthly Notices of the Royal Astronomical Society*, 360, 499, doi: [10.1111/j.1365-2966.2005.09049.x](https://doi.org/10.1111/j.1365-2966.2005.09049.x)
- Moster, B. P., Naab, T., & White, S. D. M. 2013, *Monthly Notices of the Royal Astronomical Society*, 428, 3121, doi: [10.1093/mnras/sts261](https://doi.org/10.1093/mnras/sts261)
- Mutch, S. J., Geil, P. M., Poole, G. B., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 250, doi: [10.1093/mnras/stw1506](https://doi.org/10.1093/mnras/stw1506)
- Naab, T., & Ostriker, J. P. 2017, *Annual Review of Astronomy and Astrophysics*, 55, 59, doi: [10.1146/annurev-astro-081913-040019](https://doi.org/10.1146/annurev-astro-081913-040019)
- Nagy, D., Dessauges-Zavadsky, M., Messa, M., et al. 2023, *Astronomy & Astrophysics*, 678, A183, doi: [10.1051/0004-6361/202346951](https://doi.org/10.1051/0004-6361/202346951)
- Nakoneczny, S. J., Bilicki, M., Pollo, A., et al. 2021, *Astronomy & Astrophysics*, 649, A81, doi: [10.1051/0004-6361/202039684](https://doi.org/10.1051/0004-6361/202039684)
- Narayanan, D., Smith, J. D. T., Hensley, B. S., et al. 2023, *The Astrophysical Journal*, 951, 100, doi: [10.3847/1538-4357/accf8d](https://doi.org/10.3847/1538-4357/accf8d)
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1995, *Monthly Notices of the Royal Astronomical Society*, 275, 56, doi: [10.1093/mnras/275.1.56](https://doi.org/10.1093/mnras/275.1.56)
- Neal, R. M. 1997, arXiv e-prints, physics/9701026, doi: [10.48550/arXiv.physics/9701026](https://doi.org/10.48550/arXiv.physics/9701026)
- Negroponte, J., & White, S. D. M. 1983, *Monthly Notices of the Royal Astronomical Society*, 205, 1009, doi: [10.1093/mnras/205.4.1009](https://doi.org/10.1093/mnras/205.4.1009)
- Neufeld, D. A., & Kaufman, M. J. 1993, *The Astrophysical Journal*, 418, 263, doi: [10.1086/173388](https://doi.org/10.1086/173388)
- Ocvirk, P., Gillet, N., Shapiro, P. R., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 463, 1462, doi: [10.1093/mnras/stw2036](https://doi.org/10.1093/mnras/stw2036)

- Okamoto, T., Gao, L., & Theuns, T. 2008, *Monthly Notices of the Royal Astronomical Society*, 390, 920, doi: [10.1111/j.1365-2966.2008.13830.x](https://doi.org/10.1111/j.1365-2966.2008.13830.x)
- Osterbrock, D. E. 1965, *The Astrophysical Journal*, 142, 1423, doi: [10.1086/148427](https://doi.org/10.1086/148427)
- Padoan, P., Haugbølle, T., & Nordlund, Å. 2012, *The Astrophysical Journal Letters*, 759, L27, doi: [10.1088/2041-8205/759/2/L27](https://doi.org/10.1088/2041-8205/759/2/L27)
- pandas development team, T. 2020, pandas-dev/pandas: Pandas, latest, Zenodo, doi: [10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134)
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *Journal of Machine Learning Research*, 12, 2825. <https://dl.acm.org/doi/10.5555/1953048.2078195>
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *Astronomy & Astrophysics*, 641, A6, doi: [10.1051/0004-6361/201833910](https://doi.org/10.1051/0004-6361/201833910)
- Ploeckinger, S., & Schaye, J. 2020, *Monthly Notices of the Royal Astronomical Society*, 497, 4857, doi: [10.1093/mnras/staa2172](https://doi.org/10.1093/mnras/staa2172)
- Pogge, R. 2008, *Physics of the Interstellar Medium*, Ohio State University. <https://www.astronomy.ohio-state.edu/pogge.1/Ast871/Notes/index.html>
- Popescu, C. C., Yang, R., Tuffs, R. J., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 470, 2539, doi: [10.1093/mnras/stx1282](https://doi.org/10.1093/mnras/stx1282)
- Press, W. H., & Schechter, P. 1974, *The Astrophysical Journal*, 187, 425, doi: [10.1086/152650](https://doi.org/10.1086/152650)
- Rees, M. J., & Ostriker, J. P. 1977, *Monthly Notices of the Royal Astronomical Society*, 179, 541, doi: [10.1093/mnras/179.4.541](https://doi.org/10.1093/mnras/179.4.541)
- Richings, A. J., & Schaye, J. 2016, *Monthly Notices of the Royal Astronomical Society*, 458, 270, doi: [10.1093/mnras/stw327](https://doi.org/10.1093/mnras/stw327)
- Richings, A. J., Schaye, J., & Oppenheimer, B. D. 2014, *Monthly Notices of the Royal Astronomical Society*, 440, 3349, doi: [10.1093/mnras/stu525](https://doi.org/10.1093/mnras/stu525)
- Robertson, B. E. 2022, *Annual Review of Astronomy and Astrophysics*, 60, 121, doi: [10.1146/annurev-astro-120221-044656](https://doi.org/10.1146/annurev-astro-120221-044656)
- Robertson, B. E., Ellis, R. S., Dunlop, J. S., McLure, R. J., & Stark, D. P. 2010, *Nature*, 468, 49, doi: [10.1038/nature09527](https://doi.org/10.1038/nature09527)
- Robertson, B. E., & Kravtsov, A. V. 2008, *The Astrophysical Journal*, 680, 1083, doi: [10.1086/587796](https://doi.org/10.1086/587796)
- Robinson, D., Avestruz, C., & Gnedin, N. Y. 2022, *The Astrophysical Journal*, 936, 50, doi: [10.3847/1538-4357/ac85e1](https://doi.org/10.3847/1538-4357/ac85e1)
- . 2024a, *Monthly Notices of the Royal Astronomical Society*, 528, 255, doi: [10.1093/mnras/stad3880](https://doi.org/10.1093/mnras/stad3880)

- . 2024b, arXiv e-prints, arXiv:2406.19446, doi: [10.48550/arXiv.2406.19446](https://doi.org/10.48550/arXiv.2406.19446)
- Robinson, D., Avestruz, C., Gnedin, N. Y., & Semenov, V. A. 2024c, arXiv e-prints, arXiv:2412.15324, doi: [10.48550/arXiv.2412.15324](https://doi.org/10.48550/arXiv.2412.15324)
- Rodríguez, J.-V., Rodríguez-Rodríguez, I., & Woo, W. L. 2022, WIREs Data Mining and Knowledge Discovery, 12, e1476, doi: <https://doi.org/10.1002/widm.1476>
- Romero, M., Ascasibar, Y., Palouš, J., Wunsch, R., & Mollá, M. 2021, Monthly Notices of the Royal Astronomical Society, 505, 5301, doi: [10.1093/mnras/stab1660](https://doi.org/10.1093/mnras/stab1660)
- Romero, M., Corcho-Caballero, P., Millán-Irigoyen, I., Mollá, M., & Ascasibar, Y. 2023, Monthly Notices of the Royal Astronomical Society, 521, 1727, doi: [10.1093/mnras/stad454](https://doi.org/10.1093/mnras/stad454)
- Rudd, D. H., Zentner, A. R., & Kravtsov, A. V. 2008, The Astrophysical Journal, 672, 19, doi: [10.1086/523836](https://doi.org/10.1086/523836)
- Salpeter, E. E. 1955, The Astrophysical Journal, 121, 161, doi: [10.1086/145971](https://doi.org/10.1086/145971)
- Salz, M., Banerjee, R., Mignone, A., et al. 2015, Astronomy & Astrophysics, 576, A21, doi: [10.1051/0004-6361/201424330](https://doi.org/10.1051/0004-6361/201424330)
- Saravia, A., Rodas-Quito, E., Barcos-Muñoz, L., et al. 2025, The Astrophysical Journal, 979, 217, doi: [10.3847/1538-4357/ad9bab](https://doi.org/10.3847/1538-4357/ad9bab)
- Sarker, I. H. 2021, SN Computer Science, 2, doi: [10.1007/s42979-021-00592-x](https://doi.org/10.1007/s42979-021-00592-x)
- Sawala, T., Frenk, C. S., Fattahi, A., et al. 2016, Monthly Notices of the Royal Astronomical Society, 457, 1931, doi: [10.1093/mnras/stw145](https://doi.org/10.1093/mnras/stw145)
- Scannapieco, C., Wadepuhl, M., Parry, O. H., et al. 2012, Monthly Notices of the Royal Astronomical Society, 423, 1726, doi: [10.1111/j.1365-2966.2012.20993.x](https://doi.org/10.1111/j.1365-2966.2012.20993.x)
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, Monthly Notices of the Royal Astronomical Society, 446, 521–554, doi: [10.1093/mnras/stu2058](https://doi.org/10.1093/mnras/stu2058)
- Schaye, J., Kugel, R., Schaller, M., et al. 2023, Monthly Notices of the Royal Astronomical Society, 526, 4978, doi: [10.1093/mnras/stad2419](https://doi.org/10.1093/mnras/stad2419)
- Schmidt, W., Almgren, A. S., Braun, H., et al. 2014, Monthly Notices of the Royal Astronomical Society, 440, 3051, doi: [10.1093/mnras/stu501](https://doi.org/10.1093/mnras/stu501)
- Scoville, N. Z., & Solomon, P. M. 1974, The Astrophysical Journal, 187, L67, doi: [10.1086/181398](https://doi.org/10.1086/181398)
- Semenov, V. A., Kravtsov, A. V., & Gnedin, N. Y. 2016, The Astrophysical Journal, 826, 200, doi: [10.3847/0004-637X/826/2/200](https://doi.org/10.3847/0004-637X/826/2/200)
- . 2021, The Astrophysical Journal, 918, 13, doi: [10.3847/1538-4357/ac0a77](https://doi.org/10.3847/1538-4357/ac0a77)

- Shwartz-Ziv, R., & Armon, A. 2022, *Information Fusion*, 81, 84, doi: [10.1016/j.inffus.2021.11.011](https://doi.org/10.1016/j.inffus.2021.11.011)
- Silk, J. 1977, *The Astrophysical Journal*, 211, 638, doi: [10.1086/154972](https://doi.org/10.1086/154972)
- Smith, B., Sigurdsson, S., & Abel, T. 2008, *Monthly Notices of the Royal Astronomical Society*, 385, 1443, doi: [10.1111/j.1365-2966.2008.12922.x](https://doi.org/10.1111/j.1365-2966.2008.12922.x)
- Smith, B. D., Bryan, G. L., Glover, S. C. O., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 466, 2217, doi: [10.1093/mnras/stw3291](https://doi.org/10.1093/mnras/stw3291)
- Soltan, A. 1982, *Monthly Notices of the Royal Astronomical Society*, 200, 115, doi: [10.1093/mnras/200.1.115](https://doi.org/10.1093/mnras/200.1.115)
- Spitzer, Lyman, J. 1956, *The Astrophysical Journal*, 124, 20, doi: [10.1086/146200](https://doi.org/10.1086/146200)
- Spitzer, L. 1962, *Physics of fully ionized gases*, 2nd edn. (New York: Interscience Publishers New York)
- Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 676, doi: [10.1093/mnras/stx3304](https://doi.org/10.1093/mnras/stx3304)
- Sutherland, R. S., & Dopita, M. A. 1993, *The Astrophysical Journal Supplement Series*, 88, 253, doi: [10.1086/191823](https://doi.org/10.1086/191823)
- Tamayo, D., Silburt, A., Valencia, D., et al. 2016, *The Astrophysical Journal Letters*, 832, L22, doi: [10.3847/2041-8205/832/2/L22](https://doi.org/10.3847/2041-8205/832/2/L22)
- Tayal, S. S. 2009, *Astronomy & Astrophysics*, 501, 381, doi: [10.1051/0004-6361/200810055e](https://doi.org/10.1051/0004-6361/200810055e)
- Thomas, R. M., Zaroubi, S., Ciardi, B., et al. 2009, *Monthly Notices of the Royal Astronomical Society*, 393, 32, doi: [10.1111/j.1365-2966.2008.14206.x](https://doi.org/10.1111/j.1365-2966.2008.14206.x)
- Toomre, A., & Toomre, J. 1972, *The Astrophysical Journal*, 178, 623, doi: [10.1086/151823](https://doi.org/10.1086/151823)
- Tremmel, M., Karcher, M., Governato, F., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 470, 1121, doi: [10.1093/mnras/stx1160](https://doi.org/10.1093/mnras/stx1160)
- Tucker, W. H., & Gould, R. J. 1966, *The Astrophysical Journal*, 144, 244, doi: [10.1086/148601](https://doi.org/10.1086/148601)
- Tumlinson, J., Peebles, M. S., & Werk, J. K. 2017, *Annual Review of Astronomy and Astrophysics*, 55, 389, doi: [10.1146/annurev-astro-091916-055240](https://doi.org/10.1146/annurev-astro-091916-055240)
- Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *The Astrophysical Journal Supplement Series*, 192, 9, doi: [10.1088/0067-0049/192/1/9](https://doi.org/10.1088/0067-0049/192/1/9)
- Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, *The Astrophysical Journal*, 465, 487, doi: [10.1086/177435](https://doi.org/10.1086/177435)

- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., et al. 2021, *The Astrophysical Journal*, 915, 71, doi: [10.3847/1538-4357/abf7ba](https://doi.org/10.3847/1538-4357/abf7ba)
- Vogelsberger, M., Genel, S., Sijacki, D., et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 436, 3031–3067, doi: [10.1093/mnras/stt1789](https://doi.org/10.1093/mnras/stt1789)
- Vogelsberger, M., Marinacci, F., Torrey, P., & Puchwein, E. 2020, *Nature Reviews Physics*, 2, 42, doi: [10.1038/s42254-019-0127-2](https://doi.org/10.1038/s42254-019-0127-2)
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014, *Nature*, 509, 177, doi: [10.1038/nature13316](https://doi.org/10.1038/nature13316)
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 444, 1518–1547, doi: [10.1093/mnras/stu1536](https://doi.org/10.1093/mnras/stu1536)
- Wang, L., Dutton, A. A., Stinson, G. S., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 454, 83, doi: [10.1093/mnras/stv1937](https://doi.org/10.1093/mnras/stv1937)
- Wang, Y., Ferland, G. J., Lykins, M. L., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 440, 3100, doi: [10.1093/mnras/stu514](https://doi.org/10.1093/mnras/stu514)
- Wang, Y., Pan, Z., Zheng, J., Qian, L., & Li, M. 2019, *Astrophysics and Space Science*, 364, 139, doi: [10.1007/s10509-019-3602-4](https://doi.org/10.1007/s10509-019-3602-4)
- Waskom, M. L. 2021, *Journal of Open Source Software*, 6, 3021, doi: [10.21105/joss.03021](https://doi.org/10.21105/joss.03021)
- Westmeier, T., Braun, R., & Koribalski, B. S. 2011, *Monthly Notices of the Royal Astronomical Society*, 410, 2217, doi: [10.1111/j.1365-2966.2010.17596.x](https://doi.org/10.1111/j.1365-2966.2010.17596.x)
- Wetzel, A. R., Hopkins, P. F., Kim, J.-h., et al. 2016, *The Astrophysical Journal Letters*, 827, L23, doi: [10.3847/2041-8205/827/2/L23](https://doi.org/10.3847/2041-8205/827/2/L23)
- White, S. D. M. 1976, *Monthly Notices of the Royal Astronomical Society*, 174, 467, doi: [10.1093/mnras/174.2.467](https://doi.org/10.1093/mnras/174.2.467)
- White, S. D. M., & Frenk, C. S. 1991, *The Astrophysical Journal*, 379, 52, doi: [10.1086/170483](https://doi.org/10.1086/170483)
- White, S. D. M., & Rees, M. J. 1978, *Monthly Notices of the Royal Astronomical Society*, 183, 341, doi: [10.1093/mnras/183.3.341](https://doi.org/10.1093/mnras/183.3.341)
- Wiersma, R. P. C., Schaye, J., Dalla Vecchia, C., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 409, 132, doi: [10.1111/j.1365-2966.2010.17299.x](https://doi.org/10.1111/j.1365-2966.2010.17299.x)
- Wiersma, R. P. C., Schaye, J., & Smith, B. D. 2009, *Monthly Notices of the Royal Astronomical Society*, 393, 99–107, doi: [10.1111/j.1365-2966.2008.14191.x](https://doi.org/10.1111/j.1365-2966.2008.14191.x)
- Wood, K., Mathis, J. S., & Ercolano, B. 2004, *Monthly Notices of the Royal Astronomical Society*, 348, 1337, doi: [10.1111/j.1365-2966.2004.07458.x](https://doi.org/10.1111/j.1365-2966.2004.07458.x)

Zebari, R., Abdulazeez, A., Zeebaree, D., Zebari, D., & Saeed, J. 2020, Journal of Applied Science and Technology Trends, 1, 56 , doi: [10.38094/jastt1224](https://doi.org/10.38094/jastt1224)

Zhao, Y., Liu, J., Zhang, Z.-Y., & Bisbas, T. G. 2024, The Astrophysical Journal, 977, 46, doi: [10.3847/1538-4357/ad8c3e](https://doi.org/10.3847/1538-4357/ad8c3e)