

Unsupervised Domain Adaptation for Constraining Star Formation Histories

Sankalp Gilda,^{*1} Antoine de Mathelin,^{*2} Sabine Bellstedt,³ Guillaume Richard⁴

¹ ML Collective

² Michelin - ENS Paris-Saclay, Centre Borelli

³ International Centre for Radio Astronomy Research, The University of Western Australia

⁴ EDF R&D - ENS Paris-Saclay, Centre Borelli

sankalp.gilda@gmail.com, antoine.de-mathelin-de-papigny@michelin.com, sabine.bellstedt@uwa.edu.au, guillaume-gu.richard@edf.fr

Abstract

The prevalent paradigm of machine learning today is to use past observations to predict future ones. What if, however, we are interested in knowing the past given the present? This situation is indeed one that astronomers must contend with often. To understand the formation of our universe, we must derive the time evolution of the visible mass content of galaxies. However, to observe a complete star life, one would need to wait for one billion years! To overcome this difficulty, astrophysicists leverage supercomputers and evolve simulated models of galaxies till the current age of the universe, thus establishing a mapping between observed radiation and star formation histories (SFHs). Such ground-truth SFHs are lacking for actual galaxy observations, where they are usually inferred – with often poor confidence – from spectral energy distributions (SEDs) using Bayesian fitting methods. In this investigation, we discuss the ability of unsupervised domain adaptation to derive accurate SFHs for galaxies with simulated data as a necessary first step in developing a technique that can ultimately be applied to observational data.

1 Introduction

In recent times, many transfer problems have arisen, and various methods exist to solve them. We can broadly classify these into three categories. First, supervised domain adaptation, where only a few labeled target data are available (Dai et al. 2007; de Mathelin et al. 2020; Motiiian et al. 2017a,b). Second, semi-supervised domain adaptation, where in addition to these labeled data, a large amount of unlabeled data is available (Kumar, Saha, and Daume 2010; Saito et al. 2019; Tzeng et al. 2015). Finally, unsupervised domain adaptation, where only unlabeled data is available in the target (Ganin et al. 2016; Huang et al. 2007; Richard et al. 2020; Saito et al. 2018; Sugiyama et al. 2007). Researchers have shown that adding target labels can increase the performance of the model (Motiiian et al. 2017a). This setting is also the most encountered in practice: (Cortes, Mohri, and Medina 2019), as it is often possible to label at least a few target samples.

However, there are cases in which getting supervised data is not possible. Such is the case of the present astrophysics problem of SFH prediction. Here, we aim to learn the history

of our Universe through the evolution of its galaxies’ masses throughout their individual histories. We try to obtain this history from the radiation that reaches us on earth – the Spectral Energy Distribution (SED), a function of the brightness of a galaxy with the wavelength of observation.

As galaxies evolve over billions of years, it is impossible to completely capture any single galaxy’s SED as a function of time. Astrophysicists then generate models of light through physically-motivated mechanisms, which allow building artificial star formation histories with the corresponding radiation. One of the objectives of these artificial data sets is to find the correspondence between the radiation and the SFH. We try to find a function that can robustly perform this mapping. Classically, Bayesian models have been developed : PROSPECT (Robotham et al. 2020), MAGPHYS (da Cunha, Charlot, and Elbaz 2008), CIGALE (Noll et al. 2009), BAGPIPES (Carnall et al. 2018), PROSPECTOR (Johnson and Leja 2017), and numerous others, more recently MIRKWOOD, (Gilda, Lower, and Narayanan 2021) showed that deep learning approaches have the potential to infer SFHs from SEDs.

A significant challenge with applying all the above models arises due to the ‘domain shift’ between simulated and real galaxies. Even a deep neural network (DNN) trained only on the synthetic data sets will falter when predicting the SFHs of real galaxies. Therefore, in this work, we investigate the use of unsupervised domain adaptation as an alternative method of extracting SFHs for individual galaxies without being limited by the simplifications of physical models and parametrizations of conventional SED-fitting techniques. We turn toward cosmological, hydrodynamic simulations, including EAGLE (Schaye et al. 2015a), ILLUSTRISTNG (Nelson et al. 2018; Pillepich et al. 2018) and SIMBA (Davé et al. 2019). These simulate a volume of the Universe from shortly after the Big Bang to the present day, recording the evolution of dark and baryonic matter over time. Galaxies formed within simulations can be studied as a proxy for real galaxies and can be compared directly to observations. They can thus serve to augment our understanding of the evolution of the Universe.

2 Data

Our training and test data sets consist of SEDs from three state-of-the-art cosmological galaxy formation simulations, where the true physical properties are known – including

^{*}Equal Contribution




Figure 1: Kernel density estimate plots of the log of flux densities (FD) for the first three features. The feature names at the top of each plot are the names of the filters, each centered at a different wavelength, in which photometry was simulated. We notice that for the three features shown here, all three simulations share the same support, which justifies our decision of using KLIEP.

their true star formation histories. SIMBA (Davé et al. 2019), EAGLE (Schaye et al. 2015b; Schaller et al. 2015; McAlpine et al. 2016), and ILLUSTRISTNG (Vogelsberger et al. 2014), with 1,688, 4,697, and 9,633 samples respectively, together constitute a diverse sample of galaxies with realistic growth histories. Specifically, we select galaxies at a redshift of 0 – this corresponds to simulated galaxies at the “present” epoch. Each SED is comprised of 20 measurements of the “brightness” of a galaxy at different wavelength, in the form of flux density (with units of Jansky) and are the co-variates or features that we train/learn on. The outputs/labels are star formation history (SFH) time series vectors with 29 scalar elements, in units of solar mass per year ($M_{\odot} \text{ yr}^{-1}$).

The SFH for a galaxy informs us about the net stellar mass generated as a function of time – sum of masses of all stars born, less the sum of stellar mass lost in stellar winds as stars age and eventually die. In other words, SFHs are plots of the star formation rates (SFRs) of galaxies against the *lookback time*, which for galaxies at $z = 0$ extends from 0 to the very age of our Universe (~ 13.8 Gyrs). By stacking (adding) together SFHs of several thousand galaxies from a simulation, and dividing by the volume of the simulated box (i.e. the size of the simulated universe), we can derive the cosmic star formation rate density (CSFRD), which is a well-studied global property of our Universe. Astrophysicists aim to tune various physical properties in any given simulation until the CSFRD plot derived from it matches closely to the most-widely accepted one derived from observations (Madau and Dickinson 2014).¹

We sequentially train on any two of these simulated data sets, and predict on the third, thus giving us three sets of source- and target-domain data and results.

3 Methodology

We consider the problem of prediction of star formation history (SFH) where the learner has access to a data set $X \in \mathbb{R}^{n \times p}$ encoding the radiations of n galaxies with p the number of filters/wavelength-bins, and a data set $Y \in \mathbb{R}^{n \times T}$ giving the corresponding SFH for each galaxy with T is the

¹The observed CSFRD (especially at $z > 2$) is still a topic of active research currently, as the amount of star formation obscured by dust in the early Universe is still unknown.

time length of the history. Each SED (row of X) is comprised of $p = 20$ measurements of the “brightness” of a galaxy at different wavelength, in the form of flux density (with units of Jansky) (see Figure 1). Below are our pre-processing steps:

1. First, we create three sets of experiments. For each, we use two galaxies in the training and validation sets (with a 9:1 split) and the third galaxy in the test set.
2. Second, we normalize each SFH time series (each row of Y) by its sum and store the resultant normalized SFH (SFH_{norm}) and the sum (SFH_{sum}) separately. This step is needed because of the large dynamic range of the various star formation histories (see Appendix B): SFH curves have a large variety of scales (some increasing to more than 100 whereas others never increase over 0.1). By scaling each SFH time series, we make the learning of the curve trend easier. The learning of the SFHs is now decoupled in the learning of SFH_{norm} and SFH_{sum} .
3. Third, we further ease the learning of the SFH_{norm} by reducing the curves to their first 3 Kernel-PCA components (Soentpiet et al. 1999). The choice of this decomposition method is motivated by extensive experimentation that showed that Kernel-PCA beat both linear PCA (Jolliffe and Cadima 2016) and discrete wavelet transform (Shensa 1992) in their ability to recreate the original time series successfully. For each of the three experiments, we provide the Kernel-PCA with a wide range of hyperparameters and pick the ones that can recreate the original SFH time series back, judged according to the DILATE loss metric (Le Guen and Thome 2019) (see Appendix A). The DILATE similarity metric between two time series is defined as an equally weighted average of the dynamic time warping (DTW) and the temporal distortion index (TDI) similarity scores. The number of principal components (3) was chosen by selecting the smallest set that explains at least 80% variance in the validation set. We refer to these kernel-PCA components as SFH_{kPCA} .
4. Fourth, we normalize the input features (the columns from X) via log-scaling, and follow this up by standard scaling normalization. In Figure 1 we visualize the log-scaled flux densities, in units of Janskies, for all three simulations, in 3 out of 20 filters.
5. Finally, we derive KLIEP weights for the training samples

in all three experiments.

After performing these steps, we apply a domain adaptation method to correct the shift between the source and target input distributions. We choose the method KLIEP (Sugiyama et al. 2007), an instance-based method that reweights the sources in order to minimize the KL-divergence between the two domains. Instance-based approaches have been widely used to handle regression domain adaptation issues (Cortes and Mohri 2014; Huang et al. 2007; Mansour, Mohri, and Rostamizadeh 2009; Sugiyama et al. 2007), and are particularly robust to negative transfer (de Mathelin et al. 2020). We also visually observe on the marginal distributions for the 20 input filters that all domains have the same support in the feature space, which is the framework considered by KLIEP. Finally, KLIEP has the critical advantage of proposing an unsupervised selection procedure to select a relevant bandwidth.

For each of SFH_{kPCA} and SFH_{sum} , we train a 4 layer feed-forward DNN with drop-out and 256 nodes in each layer. We use ReLu as the activation function, and Adam (Kingma and Ba 2015) as the optimizer, with learning rate of $1e^{-3}$. We train for 200 epochs, with early stopping to prevent over-fitting.

4 Results

There are two main ways of assessing the SFH outputs derived via our machine learning implementation:

1. Comparing the derived SFH to the true SFH for individual galaxies in the test set.
2. Comparing the *predicted* ΣSFH to the *true* ΣSFH for each simulation. *True* ΣSFH is the sum of SFH for all galaxies in a simulation and is a critical metric enabling us to verify the correctness of input physics in a simulation. We know from observations (Madau and Dickinson 2014) that star formation in the observable Universe peaked about 2 Gyrs after its formation, and all hydrodynamical simulations must produce galaxies that satisfy this observation. By ensuring that the ΣSFH curve inferred from our neural networks matches those from the underlying simulations, we ground our predictions in science while simultaneously enabling appropriate tuning of model architectures, loss functions, and hyperparameters in case of mismatches. Such comparison also enables us to assess any systematic effects in modeling when training using one simulation and comparing predictions on another.

In Figure 2, we plot the true and derived ΣSFH curves for the three distinct test datasets (SIMBA, ILLUSTRISTNG, and EAGLE), where the training data consists of samples from the other two simulations. In these cases, the resulting sample SFHs (and thus ΣSFH) are most deviant from their ground truth vectors. This is unsurprising, as the different simulations have intrinsically different sample SFHs; this is a known difference between different simulations (as shown in Figure A1 of Bellstedt et al. 2020). Figures 4, 5, and 6 show examples of true and derived SFHs for individual galaxies within the simulations using KLIEP. Based on the five metrics tested within this work (MAE, RMSE, BE, DTW, and TDI),

each row shows an example of the best-performing galaxy output on the left and the worst-performing galaxy on the right. The first thing to note here is that in most cases, the best- and worst-performing galaxies are different when we use different metrics to pick them. This discrepancy highlights the fact that each metric compares time-series differently, and hence it is difficult to pick one metric as the loss function to minimize.

Another observation made from these three figures is how well our technique can reproduce the stochastic nature of the simulated galaxies’ SFHs. As star formation can be an incredibly stochastic process (as is clear from examples such as the top-left panel in Figure 4 and the top-left panel in Figure 5), SFHs can regularly fluctuate between high and low values. In general, we find that such stochastic SFHs are poorly recovered. For the sake of galaxy property analysis, the accurate recovery of overall SFH trends is more important than the recovery of individual star formation rate (SFR) epochs. In the bottom-right panel of Figure 4, for example (galaxy index 333) it can be seen that there is a star formation event early on at $\sim 10 - 12$ Gyr, and then a secondary star formation event from ~ 2 Gyr to the present day. Recovering these two main epochs is more crucial than correctly recovering the individual SFR peaks within each epoch. One way of achieving this is to temporally smooth the SFHs of individual galaxies prior to training and testing. We can justify such smoothing of simulated features given that we would never expect the derived SFHs for observed galaxies (the ultimate aim of this work) to reproduce such short-scale stochastic features. As an example of this, see Figure C1 of Robotham et al. (2020), where “good” fits to the SFHs from the semi-analytic model SHARK (Lagos et al. 2018) from the SED-fitting code PROSPECT do not recover the stochastic SFHs.

5 Future Work

Results presented here are only one part of a continuing effort to apply sophisticated data-driven methods to SED-fitting, a crucial first step for extracting galaxy properties. In future work, we will explore advanced multi-source domain adaptation approaches (Zhao et al. 2018; Richard et al. 2021) which should more robustly account for the presence of multiple source-domain datasets as well as adaptively reweight the training samples depending on the prediction task, thus providing higher quality results.

We are also actively analyzing and attempting to correct systematic biases induced by our various design choices. How do our SFH inferences for galaxies of low mass compare to those of high mass? Is our modeling sufficient to accurately infer SFH from both old and young galaxies, unlike traditional parametric approaches which suffer at earlier epochs in the Universe’s history? Answers to such questions will enable us to discover and correct potential biases in our predictions.

Finally, we also plan on vastly increasing the size of our simulated data to account for the immense diversity in the physics of galaxy formation and evolution modeling—different initial mass functions (IMFs) (for example Salpeter 1955; Kroupa 2001; Chabrier 2003), nebular emission models (for example Bruzual and Charlot 2003), mass-metallicity relationships (such as those presented by Tremonti et al. 2004;




Figure 2: Global SFH predictions for the three experiments. The curves correspond to the sums over all SFH or predicted SFH.

Test Simulation		MAE (↓)	RMSE (↓)	BE (↓)	DTW (↓)	TDI (↓)
ILLUSTRISTNG	Baseline	0.23 ^{+0.04} _{-0.03}	0.31 ^{+0.04} _{-0.03}	0.05 ^{+0.07} _{-0.10}	1.06 ^{+0.14} _{-0.07}	3.21 ^{+2.56} _{-1.54}
	UDA	0.20 ^{+0.01} _{-0.01}	0.27 ^{+0.01} _{-0.01}	0.03 ^{+0.03} _{-0.02}	0.94 ^{+0.03} _{-0.02}	2.56 ^{+0.31} _{-0.72}
	PCA	0.11	0.16	0.00	0.74	0.23
EAGLE	Baseline	0.32 ^{+0.02} _{-0.01}	0.46 ^{+0.02} _{-0.02}	-0.10 ^{+0.02} _{-0.02}	1.84 ^{+0.11} _{-0.12}	2.69 ^{+0.27} _{-0.42}
	UDA	0.29 ^{+0.02} _{-0.02}	0.42 ^{+0.02} _{-0.02}	-0.08 ^{+0.02} _{-0.02}	1.70 ^{+0.11} _{-0.13}	2.37 ^{+0.55} _{-0.33}
	PCA	0.12	0.18	0.00	0.77	0.15
SIMBA	Baseline	0.61 ^{+0.03} _{-0.02}	0.93 ^{+0.04} _{-0.02}	0.09 ^{+0.04} _{-0.07}	3.23 ^{+0.14} _{-0.12}	2.22 ^{+0.54} _{-0.34}
	UDA	0.55 ^{+0.02} _{-0.02}	0.90 ^{+0.02} _{-0.02}	0.07 ^{+0.06} _{-0.05}	3.15 ^{+0.14} _{-0.18}	2.10 ^{+0.77} _{-0.42}
	PCA	0.31	0.49	0.03	2.23	0.19

Table 1: Forecasting results, with EAGLE and SIMBA as source domains and ILLUSTRISTNG as the target domain. Predictions for all 9,633 samples in ILLUSTRISTNG have been averaged, and 16th, 50th, and 84th quantile values are drawn from the 100 predictions per galaxy, corresponding to the 100 neural networks used. Leading method between UDA and no-UDA (baseline) are shown in bold; PCA is by definition always the best, and is considered the ‘oracle’ prediction.

Test Simulation		MAE (↓)	RMSE (↓)	BE (↓)	DTW (↓)	TDI (↓)
ILLUSTRISTNG	Baseline	0.84 ^{+0.39} _{-0.27}	1.01 ^{+0.52} _{-0.31}	0.50 ^{+0.66} _{-0.92}	3.57 ^{+2.06} _{-1.62}	0.41 ^{+0.42} _{-0.27}
	UDA	0.61 ^{+0.16} _{-0.19}	0.74 ^{+0.25} _{-0.21}	0.27 ^{+0.28} _{-0.20}	2.43 ^{+0.70} _{-1.33}	0.20 ^{+0.11} _{-0.16}
	PCA	0.17	0.25	0.02	1.11	0.02
EAGLE	Baseline	0.98 ^{+0.12} _{-0.13}	1.31 ^{+0.11} _{-0.15}	-0.48 ^{+0.11} _{-0.11}	4.75 ^{+0.54} _{-0.66}	0.66 ^{+0.45} _{-0.22}
	UDA	0.87 ^{+0.13} _{-0.16}	1.25 ^{+0.16} _{-0.17}	-0.39 ^{+0.12} _{-0.14}	3.93 ^{+0.73} _{-0.77}	0.60 ^{+0.41} _{-0.27}
	PCA	0.10	0.15	0.01	0.79	0.01
SIMBA	Baseline	0.31 ^{+0.08} _{-0.06}	0.39 ^{+0.08} _{-0.07}	0.15 ^{+0.06} _{-0.11}	1.16 ^{+0.34} _{-0.41}	0.31 ^{+0.17} _{-0.10}
	UDA	0.27 ^{+0.09} _{-0.08}	0.35 ^{+0.09} _{-0.11}	0.12 ^{+0.09} _{-0.09}	1.07 ^{+0.58} _{-0.29}	0.25 ^{+0.16} _{-0.11}
	PCA	0.08	0.10	0.04	0.40	0.08

Table 2: Forecasting results for Σ SFH (total star formation history) with EAGLE and SIMBA as source domains and ILLUSTRISTNG as the target domain. Predictions for all 9,633 samples in ILLUSTRISTNG have been added, and 16th, 50th, and 84th quantile values are drawn from the 100 predictions per galaxy, corresponding to the 100 neural networks used. Leading method between UDA and no-UDA (baseline) are shown in bold; PCA is by definition always the best, and is considered the ‘oracle’ prediction. Metrics are scaled (MAE \times 1000, RMSE \times 1000, BE \times 1000, DTW \times 1000 for readability).

Jimmy et al. 2015; Lara-Lopez et al. 2013), among others. While this task has a relatively longer horizon—generating and saving new simulations is a compute-intensive task—it is crucial to undertake before data-driven methods such as the one proposed in this paper become trusted by astronomers.

References

- Bellstedt, S.; Robotham, A. S. G.; Driver, S. P.; Thorne, J. E.; Davies, L. J. M.; Lagos, C. d. P.; Stevens, A. R. H.; Taylor, E. N.; Baldry, I. K.; Moffett, A. J.; Hopkins, A. M.; and Phillipps, S. 2020. Galaxy And Mass Assembly (GAMA): A forensic SED reconstruction of the cosmic star formation history and metallicity evolution by galaxy type. *Monthly Notices of the Royal Astronomical Society*.
- Bruzual, G.; and Charlot, S. 2003. Stellar population synthesis at the resolution of 2003. *Monthly Notices of the Royal Astronomical Society*, 344: 1000–1028.
- Carnall, A. C.; McLure, R. J.; Dunlop, J. S.; and Davé, R. 2018. Inferring the star formation histories of massive quiescent galaxies with BAGPIPES: evidence for multiple quenching mechanisms. *Monthly Notices of the Royal Astronomical Society*, 480(4): 4379–4401.
- Chabrier, G. 2003. Galactic Stellar and Substellar Initial Mass Function. *Publications of the Astronomical Society of the Pacific*, 115(809): 763–795.
- Cortes, C.; and Mohri, M. 2014. Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519.
- Cortes, C.; Mohri, M.; and Medina, A. M. n. 2019. Adaptation Based on Generalized Discrepancy. *J. Mach. Learn. Res.*, 20(1): 1–30.
- da Cunha, E.; Charlot, S.; and Elbaz, D. 2008. A simple model to interpret the ultraviolet, optical and infrared emission from galaxies. *Monthly Notices of the Royal Astronomical Society*, 388: 1595–1617.
- Dai, W.; Yang, Q.; Xue, G.-R.; and Yu, Y. 2007. Boosting for Transfer Learning. In *Proceedings of the 24th International Conference on Machine Learning*, volume 227, 193–200.
- Davé, R.; Anglés-Alcázar, D.; Narayanan, D.; Li, Q.; Rafieferantsoa, M. H.; and Appleby, S. 2019. SIMBA: Cosmological simulations with black hole growth and feedback. *Monthly Notices of the Royal Astronomical Society*, 486(2): 2827–2849.
- de Mathelin, A.; Richard, G.; Mougeot, M.; and Vayatis, N. 2020. Adversarial weighting for domain adaptation in regression. *arXiv preprint arXiv:2006.08251*.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial Training of Neural Networks. *J. Mach. Learn. Res.*, 17(1): 2096–2030.
- Gilda, S.; Lower, S.; and Narayanan, D. 2021. mirkwood: Fast and Accurate SED Modeling Using Machine Learning. *arXiv e-prints*, arXiv:2101.04687.
- Huang, J.; Gretton, A.; Borgwardt, K.; Schölkopf, B.; and Smola, A. J. 2007. Correcting Sample Selection Bias by Unlabeled Data. In Schölkopf, B.; Platt, J. C.; and Hoffman, T., eds., *Advances in Neural Information Processing Systems 19*, 601–608. MIT Press.
- Jimmy; Tran, K.-V.; Saintonge, A.; Accurso, G.; Brough, S.; and Oliva-Altamirano, P. 2015. The Gas Phase Mass Metallicity Relation for Dwarf Galaxies: Dependence on Star Formation Rate and HI Gas Mass. *The Astrophysical Journal*, 812(2): 98.
- Johnson, B.; and Leja, J. 2017. Bd-J/Prospector: Initial Release.
- Jolliffe, I. T.; and Cadima, J. 2016. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society of London Series A*, 374(2065): 20150202.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kroupa, P. 2001. On the variation of the initial mass function. *Monthly Notices of the Royal Astronomical Society*, 322(2): 231–246.
- Kumar, A.; Saha, A.; and Daume, H. 2010. Co-regularization based semi-supervised domain adaptation. *Advances in neural information processing systems*, 23: 478–486.
- Lagos, C. d. P.; Tobar, R. J.; Robotham, A. S. G.; Obreschkow, D.; Mitchell, P. D.; Power, C.; and Elahi, P. J. 2018. Shark: introducing an open source, free, and flexible semi-analytic model of galaxy formation. *Monthly Notices of the Royal Astronomical Society*, 481: 3573–3603.
- Lara-Lopez, M. A.; Hopkins, A. M.; Lopez-Sanchez, A. R.; Brough, S.; Colless, M.; Bland-Hawthorn, J.; Driver, S.; Foster, C.; Liske, J.; Loveday, J.; Robotham, A. S. G.; Sharp, R. G.; Steele, O.; and Taylor, E. N. 2013. Galaxy and mass assembly (GAMA): the connection between metals, specific SFR and hi gas in galaxies: the Z-SSFR relation. *Monthly Notices of the Royal Astronomical Society*, 433: L35–L39.
- Le Guen, V.; and Thome, N. 2019. Shape and Time Distortion Loss for Training Deep Time Series Forecasting Models. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Madau, P.; and Dickinson, M. 2014. Cosmic Star-Formation History. *Annual Review of Astronomy and Astrophysics*, 52: 415–486.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009. Domain Adaptation: Learning Bounds and Algorithms. In *COLT*.
- McAlpine, S.; Helly, J. C.; Schaller, M.; Trayford, J. W.; Qu, Y.; Furlong, M.; Bower, R. G.; Crain, R. A.; Schaye, J.; Theuns, T.; Dalla Vecchia, C.; Frenk, C. S.; McCarthy, I. G.; Jenkins, A.; Rosas-Guevara, Y.; White, S. D. M.; Baes, M.; Camps, P.; and Lemson, G. 2016. The EAGLE simulations of galaxy formation: Public release of halo and galaxy catalogues. *Astronomy and Computing*, 15: 72–89.

- Motiiian, S.; Jones, Q.; Iranmanesh, S. M.; and Doretto, G. 2017a. Few-Shot Adversarial Domain Adaptation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, 6673–6683. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510860964.
- Motiiian, S.; Piccirilli, M.; Adjero, D. A.; and Doretto, G. 2017b. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, 5715–5725.
- Nelson, D.; Pillepich, A.; Springel, V.; Weinberger, R.; Hernquist, L.; Pakmor, R.; Genel, S.; Torrey, P.; Vogelsberger, M.; Kauffmann, G.; Marinacci, F.; and Naiman, J. 2018. First results from the IllustrisTNG simulations: the galaxy colour bimodality. *Monthly Notices of the Royal Astronomical Society*, 475(1): 624–647.
- Noll, S.; Burgarella, D.; Giovannoli, E.; Buat, V.; Marcillac, D.; and Muñoz-Mateos, J. C. 2009. Analysis of galaxy spectral energy distributions from far-UV to far-IR with CIGALE: studying a SINGS test sample. *Astronomy & Astrophysics*, 507(3): 1793–1813.
- Pillepich, A.; Nelson, D.; Hernquist, L.; Springel, V.; Pakmor, R.; Torrey, P.; Weinberger, R.; Genel, S.; Naiman, J. P.; Marinacci, F.; and Vogelsberger, M. 2018. First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies. *Monthly Notices of the Royal Astronomical Society*, 475(1): 648–675.
- Richard, G.; de Mathelin, A.; Hébrail, G.; Mougeot, M.; and Vayatis, N. 2020. Unsupervised Multi-source Domain Adaptation for Regression. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2020, Ghent, Belgium, September 14-18, 2020, Proceedings, Part I*, volume 12457 of *Lecture Notes in Computer Science*, 395–411. Springer.
- Richard, G.; Mathelin, A.; Hébrail, G.; Mougeot, M.; and Vayatis, N. 2021. *Unsupervised Multi-source Domain Adaptation for Regression*, 395–411. abc. ISBN 978-3-030-67657-5.
- Robotham, A. S. G.; Bellstedt, S.; Lagos, C. d. P.; Thorne, J. E.; Davies, L. J.; Driver, S. P.; and Bravo, M. 2020. ProSpect: generating spectral energy distributions with complex star formation and metallicity histories. *Monthly Notices of the Royal Astronomical Society*, 495(1): 905–931.
- Saito, K.; Kim, D.; Sclaroff, S.; Darrell, T.; and Saenko, K. 2019. Semi-Supervised Domain Adaptation via Minimax Entropy. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 8049–8057.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3723–3732.
- Salpeter, E. E. 1955. The Luminosity Function and Stellar Evolution. *The Astrophysical Journal*, 121: 161.
- Schaller, M.; Dalla Vecchia, C.; Schaye, J.; Bower, R. G.; Theuns, T.; Crain, R. A.; Furlong, M.; and McCarthy, I. G. 2015. The EAGLE simulations of galaxy formation: the importance of the hydrodynamics scheme. *Monthly Notices of the Royal Astronomical Society*, 454(3): 2277–2291.
- Schaye, J.; Crain, R. A.; Bower, R. G.; Furlong, M.; Schaller, M.; Theuns, T.; Dalla Vecchia, C.; Frenk, C. S.; McCarthy, I. G.; Helly, J. C.; Jenkins, A.; Rosas-Guevara, Y. M.; White, S. D. M.; Baes, M.; Booth, C. M.; Camps, P.; Navarro, J. F.; Qu, Y.; Rahmati, A.; Sawala, T.; Thomas, P. A.; and Trayford, J. 2015a. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *Monthly Notices of the Royal Astronomical Society*, 446(1): 521–554.
- Schaye, J.; Crain, R. A.; Bower, R. G.; Furlong, M.; Schaller, M.; Theuns, T.; Dalla Vecchia, C.; Frenk, C. S.; McCarthy, I. G.; Helly, J. C.; Jenkins, A.; Rosas-Guevara, Y. M.; White, S. D. M.; Baes, M.; Booth, C. M.; Camps, P.; Navarro, J. F.; Qu, Y.; Rahmati, A.; Sawala, T.; Thomas, P. A.; and Trayford, J. 2015b. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *Monthly Notices of the Royal Astronomical Society*, 446(1): 521–554.
- Shensa, M. J. 1992. The discrete wavelet transform: wedding the a trous and Mallat algorithms. *IEEE Transactions on Signal Processing*, 40(10): 2464–2482.
- Soentpiet, R.; et al. 1999. *Advances in kernel methods: support vector learning*. MIT press.
- Sugiyama, M.; Nakajima, S.; Kashima, H.; Büna, P. v.; and Kawanabe, M. 2007. Direct Importance Estimation with Model Selection and Its Application to Covariate Shift Adaptation. In *Proceedings of the 20th International Conference on Neural Information Processing Systems, NIPS'07*, 1433–1440. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781605603520.
- Tremonti, C. A.; Heckman, T. M.; Kauffmann, G.; Brinchmann, J.; Charlot, S.; White, S. D. M.; Seibert, M.; Peng, E. W.; Schlegel, D. J.; Uomoto, A.; Fukugita, M.; and Brinkmann, J. 2004. The Origin of the Mass-Metallicity Relation: Insights from 53,000 Star-forming Galaxies in the Sloan Digital Sky Survey. *The Astrophysical Journal*, 613(2): 898–913.
- Tzeng, E.; Hoffman, J.; Darrell, T.; and Saenko, K. 2015. Simultaneous Deep Transfer Across Domains and Tasks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 4068–4076.
- Vogelsberger, M.; Genel, S.; Springel, V.; Torrey, P.; Sijacki, D.; Xu, D.; Snyder, G.; Nelson, D.; and Hernquist, L. 2014. Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe. *Monthly Notices of the Royal Astronomical Society*, 444(2): 1518–1547.
- Zhao, H.; Zhang, S.; Wu, G.; Moura, J. M. F.; Costeira, J. P.; and Gordon, G. J. 2018. Adversarial Multiple Source Domain Adaptation. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 8559–8570. Curran Associates, Inc.

A Selection of Time Series Reduction Method

B Best and Worst Predictions According to Different Metrics

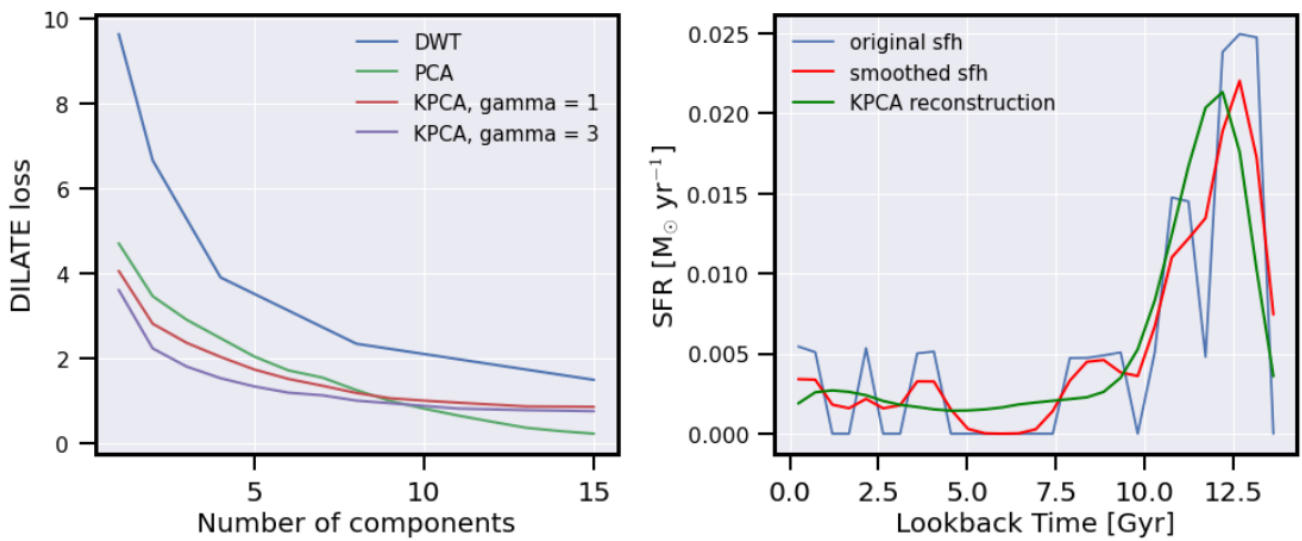


Figure 3: Comparison of SFH reduction approaches for SIMBA. On the left, the evolution of DILATE loss between the true sfh and the reconstructed signal after one of the three transformations: DWT, PCA, kernelPCA. On the right, reconstructed signal for one of the SIMBA SFH.

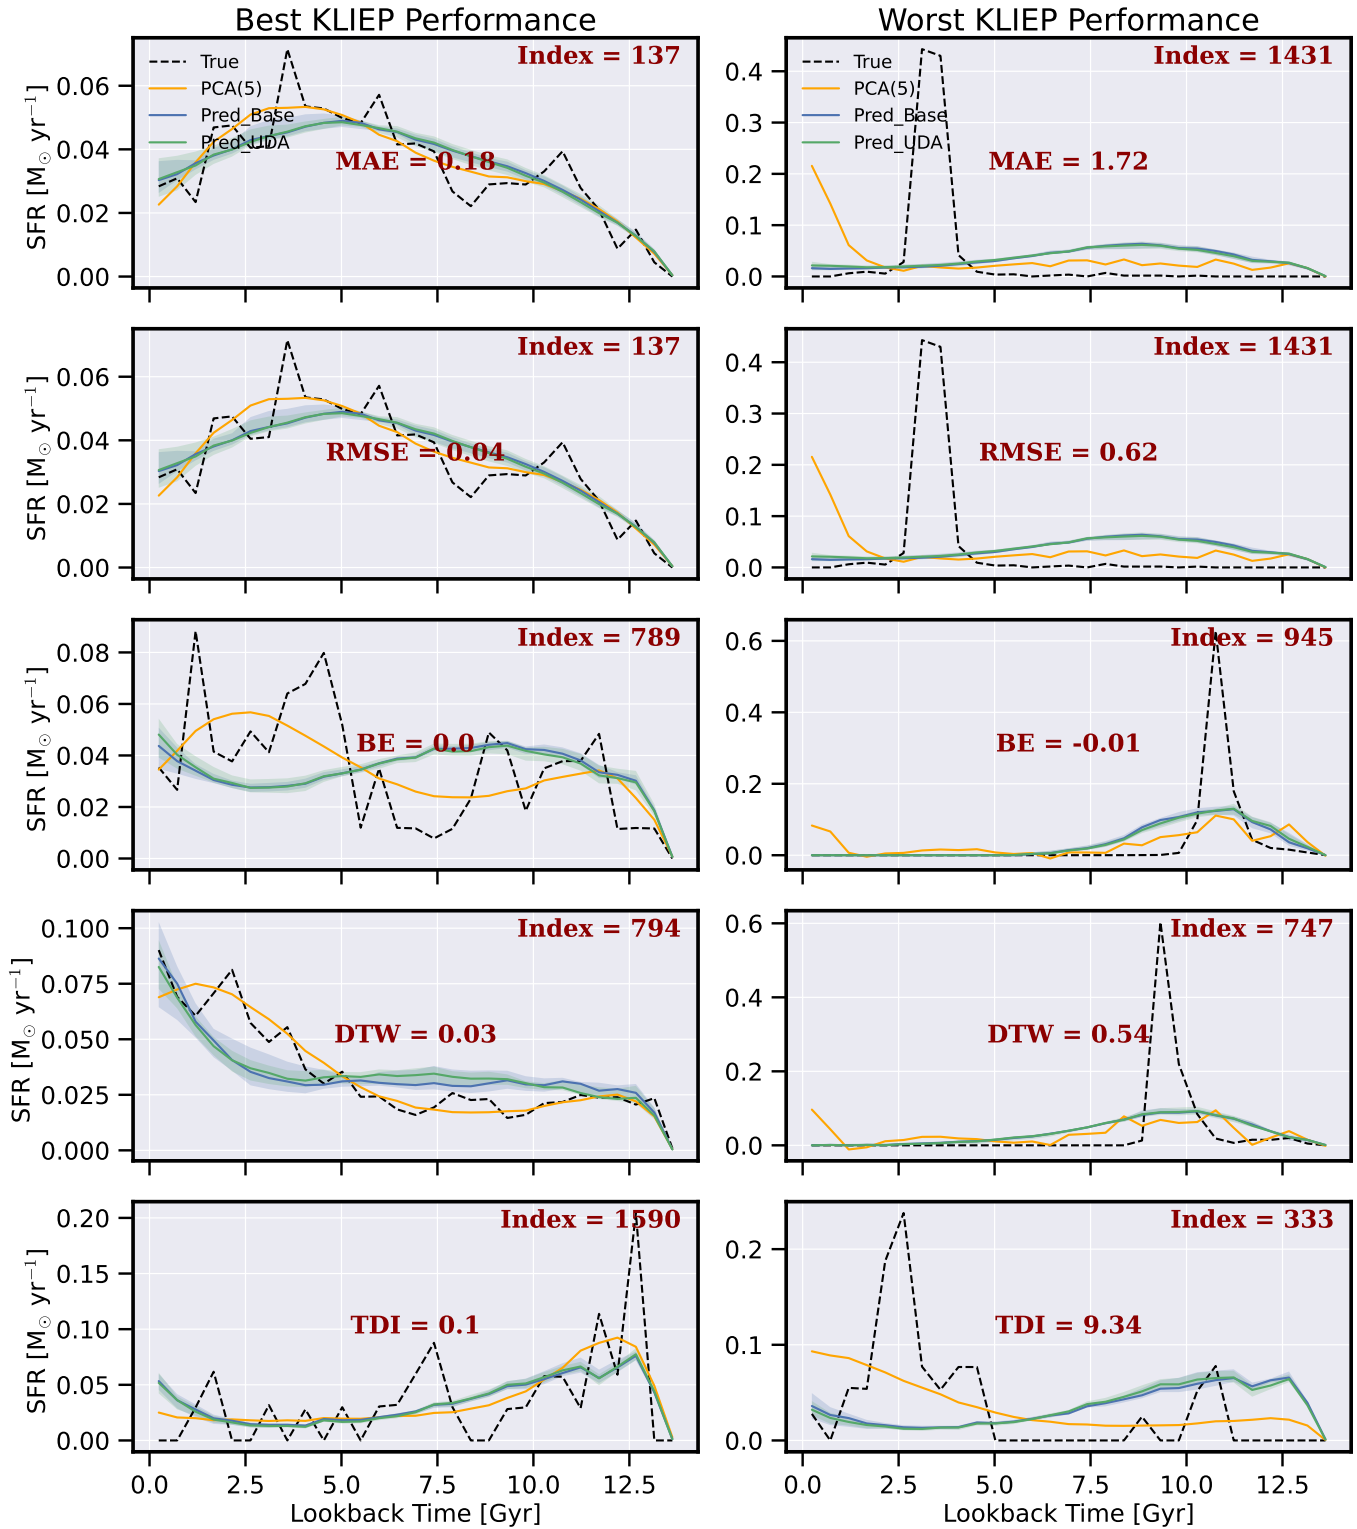


Figure 4: Worst and best galaxies for SIMBA, based on the five metrics.

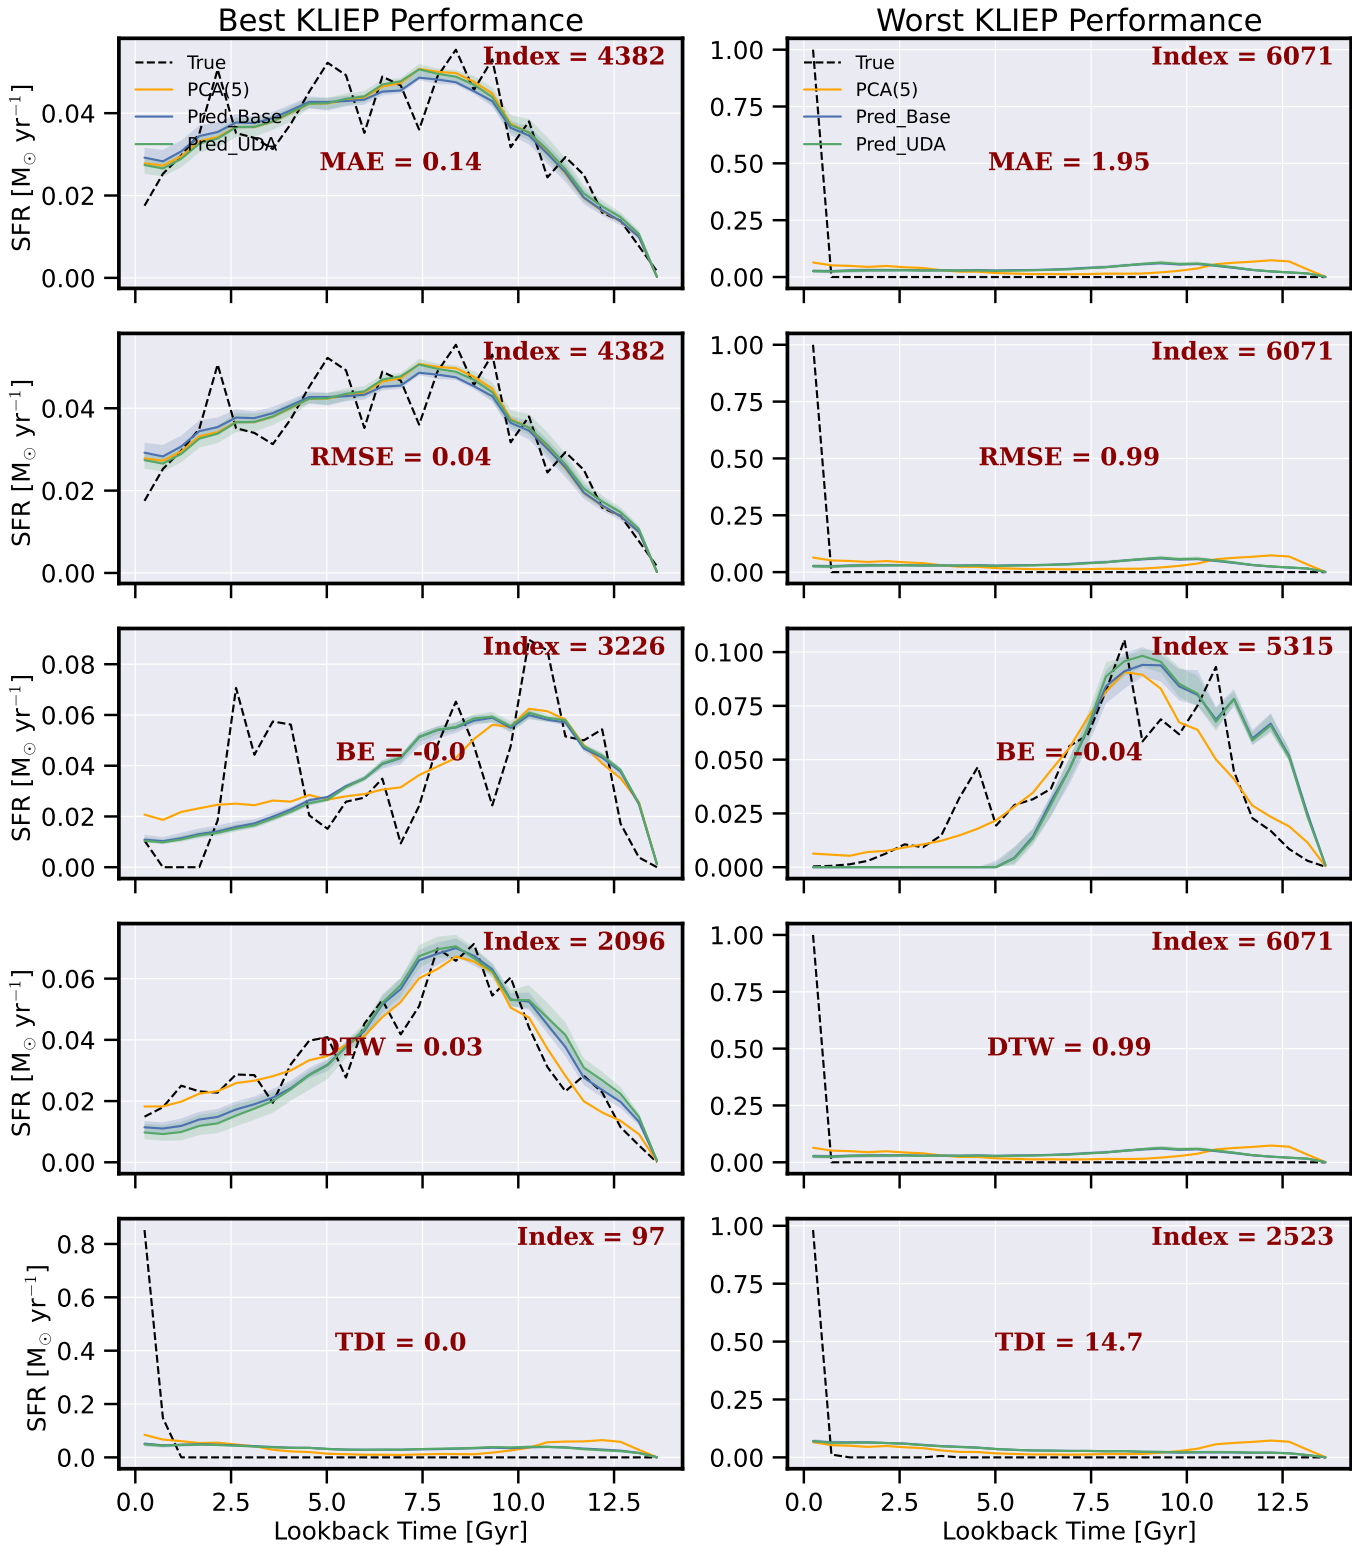


Figure 5: Worst and best galaxies for ILLUSTRISTNG, based on the five metrics.

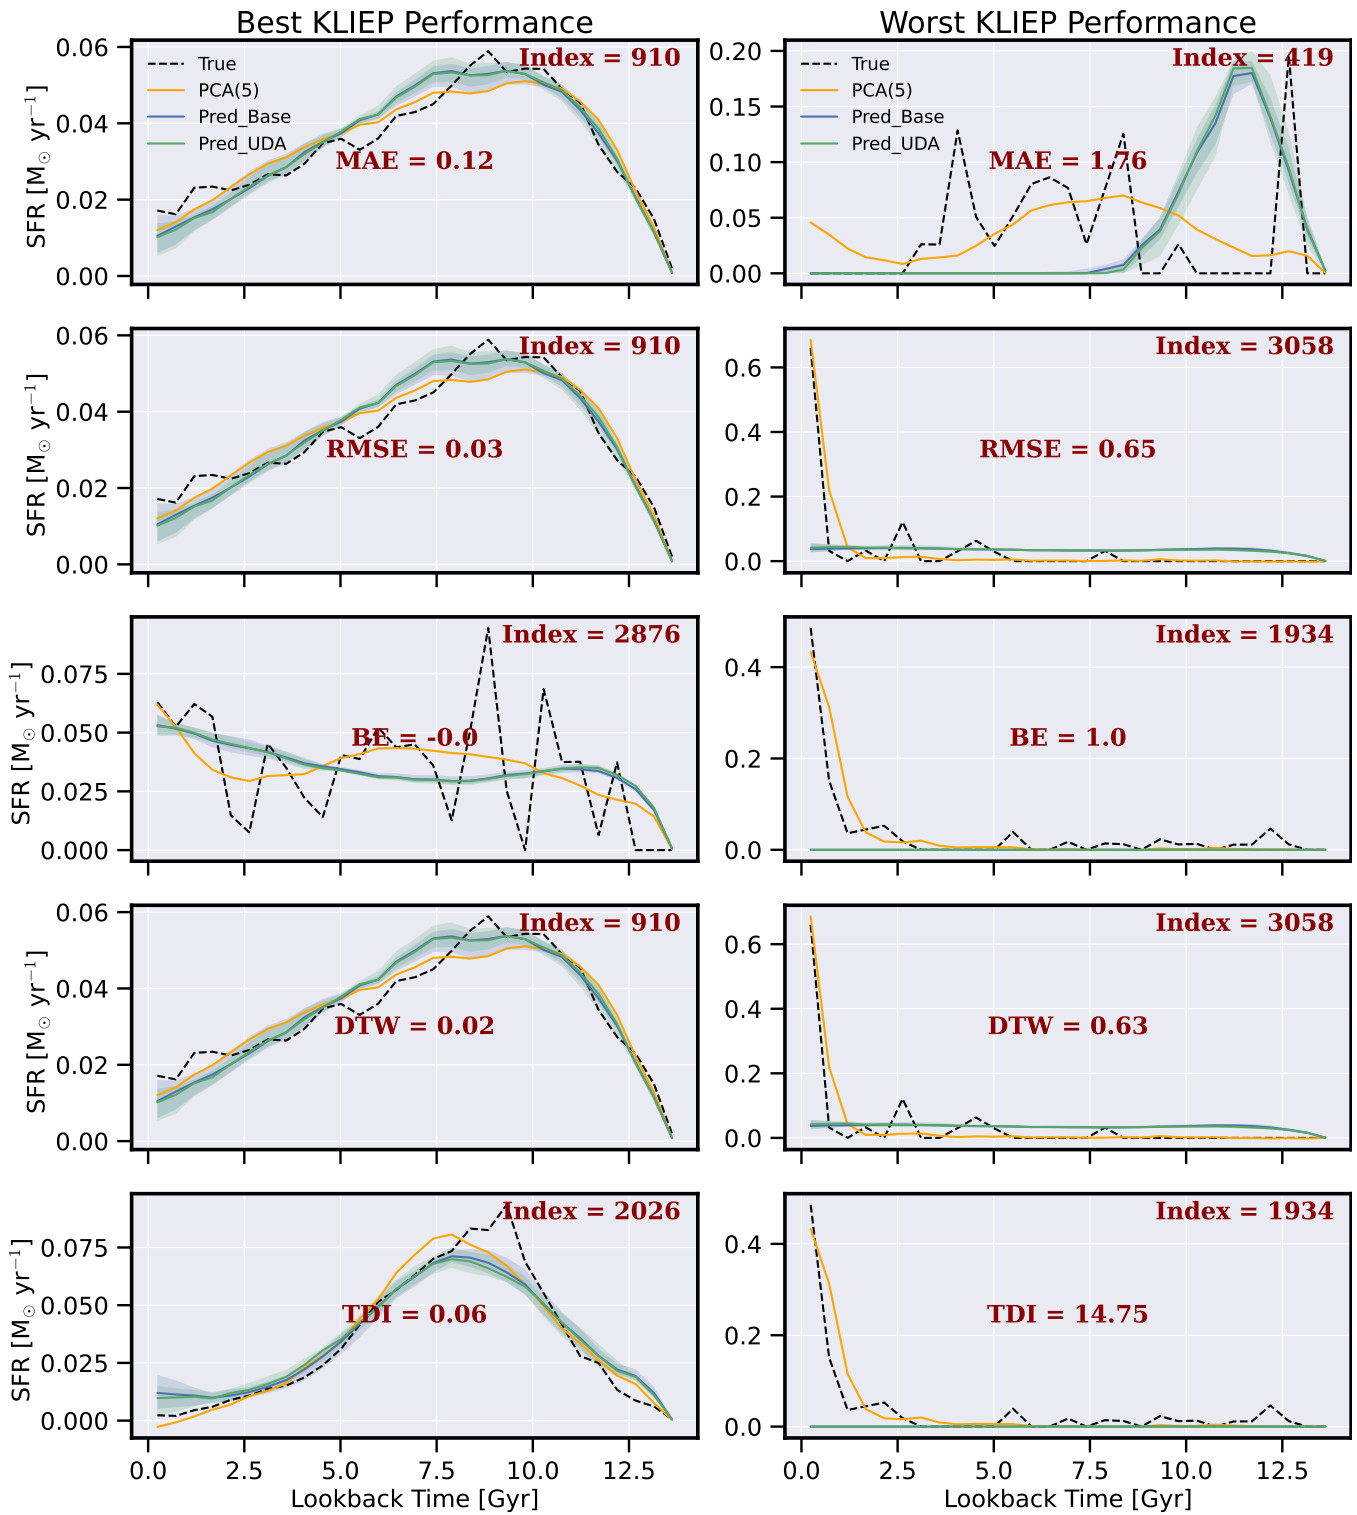


Figure 6: Worst and best galaxies for EAGLE, based on the five metrics.