

A machine learning approach to measuring the quenched fraction of low-mass satellites beyond the Local Group

Devontae C. Baxter¹,[★] M. C. Cooper¹ and Sean P. Fillingham²

¹Center for Cosmology, Department of Physics & Astronomy, University of California, Irvine, 4129 Reines Hall, Irvine, CA 92697, USA

²Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA

Accepted 2021 February 18. Received 2021 February 15; in original form 2020 December 11

ABSTRACT

Observations suggest that satellite quenching plays a major role in the build-up of passive, low-mass galaxies at late cosmic times. Studies of low-mass satellites, however, are limited by the ability to robustly characterize the local environment and star formation activity of faint systems. In an effort to overcome the limitations of existing data sets, we utilize deep photometry in Stripe 82 of the Sloan Digital Sky Survey, in conjunction with a neural network classification scheme, to study the suppression of star formation in low-mass satellite galaxies in the local Universe. Using a statistically driven approach, we are able to push beyond the limits of existing spectroscopic data sets, measuring the satellite quenched fraction down to satellite stellar masses of $\sim 10^7 M_\odot$ in group environments ($M_{\text{halo}} = 10^{13-14} h^{-1} M_\odot$). At high satellite stellar masses ($\gtrsim 10^{10} M_\odot$), our analysis successfully reproduces existing measurements of the quenched fraction based on spectroscopic samples. Pushing to lower masses, we find that the fraction of passive satellites increases, potentially signalling a change in the dominant quenching mechanism at $M_\star \sim 10^9 M_\odot$. Similar to the results of previous studies of the Local Group, this increase in the quenched fraction at low satellite masses may correspond to an increase in the efficacy of ram-pressure stripping as a quenching mechanism in groups.

Key words: galaxies: dwarf – galaxies: evolution – galaxies: formation – galaxies: star formation.

1 INTRODUCTION

The recent generation of large-scale galaxy surveys has revealed that the population of non-star-forming (i.e. ‘quiescent’ or ‘quenched’) galaxies increased by more than a factor of 2 in the past 7–10 Gyr, such that quenched systems, as opposed to their star-forming counterparts, comprise the majority of the stellar mass budget at $z \sim 0$ (Bell et al. 2004; Bundy et al. 2006; Faber et al. 2007). While the growth of the global quenched population is relatively well measured at late cosmic time, our current understanding of the physical processes responsible for the suppression (or ‘quenching’) of star formation remains woefully incomplete as evidenced by many models of galaxy formation overestimating the observed satellite quenched fraction (Kimm et al. 2009; Weinmann et al. 2012; Hirschmann et al. 2014). Nevertheless, a wide assortment of physical processes has been put forth to explain how galaxies transition from star forming to quiescent. In general, these processes are split into two distinct categories, namely, internal and environmental quenching. The former, which acts independent of local environment (i.e. on both central and satellite systems), refers to any quenching process that suppresses star formation from within a galaxy. Examples of internal quenching mechanisms include feedback from star formation (Oppenheimer & Davé 2006; Ceverino & Klypin 2009) and active galactic nuclei (AGNs; Di Matteo, Springel & Hernquist 2005; Hopkins et al. 2005; Croton et al. 2006). On the other

hand, environmental quenching, which typically applies to low-mass satellites ($\lesssim 10^{10} M_\odot$), refers to a range of quenching mechanisms that suppress star formation due to environmental factors – e.g. ram-pressure stripping (Gunn & Gott 1972; Abadi, Moore & Bower 1999), tidal stripping (Merritt 1983; Moore et al. 1999; Gnedin 2003), strangulation or starvation (Larson, Tinsley & Caldwell 1980; Kawata & Mulchaey 2008), and harassment (Moore et al. 1996; Moore, Lake & Katz 1998). In general, environmental quenching mechanisms suppress star formation either by preventing satellites from accreting gas (e.g. strangulation) or by removing pre-existing gas reservoirs through galaxy–galaxy interactions (i.e. harassment), gravitational tidal forces (e.g. tidal stripping), or interaction with the circumgalactic medium of the host (e.g. ram-pressure stripping).

At $z \sim 0$, galaxy surveys find that satellites, not centrals, comprise the largest fraction of passive systems over a wide range of stellar masses ($\lesssim 10^{10.7} M_\odot$; Wetzel et al. 2013). Furthermore, observations of low-mass ($\lesssim 10^9 M_\odot$) galaxies in the local Universe have demonstrated that nearly all field galaxies are star forming, signalling that environmental quenching is primarily responsible for suppressing star formation in the low-mass regime (Haines, Gargiulo & Merluzzi 2008; Geha et al. 2012). Altogether, these observations demonstrate the importance and ubiquity of satellite quenching at late times and especially low satellite masses. Yet, hydrodynamic and semi-analytical models, which successfully predict the fraction of quiescent centrals, continue to significantly overpredict the relative number of passive satellites, especially at low masses (Kimm et al. 2009; Hirschmann et al. 2014; Wang et al. 2014, but see also Henriques et al. 2017). This discrepancy between theoretical predictions and observations is driven by a failure to properly

[★] E-mail: dbaxter@uci.edu

† LSSTC DSFP Fellow.

model the physical processes responsible for satellite quenching. This lack of agreement between observations and theoretical models further emphasizes that understanding the details of satellite quenching is tantamount to advancing our understanding of galaxy formation.

Our current understanding of satellite quenching at low masses is largely derived from studies of dwarf galaxies ($M_* \sim 10^6\text{--}10^8 M_\odot$) in the very local Universe, including our own Local Group. First and foremost, a range of observations demonstrate that the vast majority of low-mass satellites are gas-poor and passive, in contrast to their gas-rich, star-forming counterparts in the field (e.g. Grcevich & Putman 2009; Spekkens et al. 2014; Weisz et al. 2014a, b). Furthermore, studies of the accretion history of these systems using N -body simulations demonstrate that quenching is highly efficient, such that the typical time-scale over which quenching occurs is ~ 2 Gyr at $M_* \lesssim 10^8 M_\odot$ [likely driven by an increase in the efficacy of ram-pressure stripping (Fillingham et al. 2015, 2016, 2018; Weisz et al. 2015; Wetzel, Tollerud & Weisz 2015)]. On the other hand, studies of the more massive satellites ($M_* \gtrsim 10^8 M_\odot$) in the Local Group and nearby groups/clusters find that these systems have significantly longer quenching time-scales ($\gtrsim 5$ Gyr), consistent with starvation acting as the dominant quenching mechanism (De Lucia et al. 2012; Wetzel et al. 2013; Wheeler et al. 2014). Taken together, this implies that a transition in the dominant quenching mechanism occurs at $M_* \sim 10^8 M_\odot$ (at least within Milky Way-like host haloes, $M_{\text{halo}} \sim 10^{12} M_\odot$; Fillingham et al. 2016; Rodriguez Wimberly et al. 2019).

A major step towards increasing our understanding of satellite quenching involves determining whether the aforementioned results extend beyond the Local Group. Is a similar increase in the quenched fraction observed at low masses outside of the Local Group (and/or in more massive host haloes), indicating a corresponding increase in the efficiency of environmental (or satellite) quenching at this mass range? Unfortunately, the current generation of spectroscopic surveys lacks the necessary combination of depth, area, and/or completeness to reliably probe this mass regime. For example, at the magnitude limit of the main spectroscopic survey, the Sloan Digital Sky Survey (SDSS; York et al. 2000) can only probe galaxies with stellar masses less than $10^8 M_\odot$ at $z < 0.01$. While more recent surveys push fainter, including the Galaxy and Mass Assembly (GAMA; Driver et al. 2009, 2011) survey, the corresponding area of sky mapped is significantly smaller, again limiting the number of nearby hosts around which we can study their satellites. In contrast to spectroscopic data sets, wide and deep imaging programs are able to probe both star-forming and passive galaxies down to stellar masses of $\sim 10^7 M_\odot$ at $z \lesssim 0.1$, covering significant areas on the sky.

Herein, we present a method for measuring the satellite quenched fraction down to $M_* \sim 10^7 M_\odot$ by applying machine learning and statistical background subtraction techniques to wide and deep photometric data sets, pushing beyond the limits of current spectroscopic samples. In Section 2, we describe the spectroscopic and photometric data sets utilized in our analysis. In Section 3, we discuss the training, testing, and performance of our neural network classifier (NNC) as well as our use of the trained model to classify galaxies in our photometric sample as star forming or quenched. In Section 4, we describe our statistical background subtraction technique and use it to measure the satellite quenched fraction around nearby groups. Lastly, in Section 5, we discuss and summarize our results. When necessary, we adopt a flat Lambda cold dark matter cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.3$. All magnitudes are on the AB system (Oke & Gunn 1983).

2 DATA

2.1 Photometric sample

Our analysis utilizes the co-added images and photometry from the SDSS, focusing on the deeper Stripe 82 data set (S82; Annis et al. 2014; Bundy et al. 2015). S82 is centred on the celestial equator and is comprised of an area of $\sim 300 \text{ deg}^2$ that spans between $-50^\circ < \alpha < 60^\circ$ and $-1.25^\circ < \delta < +1.25^\circ$. The co-added images in S82 reach a depth ~ 2 mag deeper in *ugriz* relative to the SDSS single-pass data. Overall, the wide area and impressive depth ($r \sim 22.4$, 95 per cent complete for galaxies) of S82 make it well suited for studying the properties of low-mass galaxies in the local Universe. For the purpose of our analysis, we limit the S82 sample to only include galaxies (defined using the SDSS TYPE parameter) with $13 < r < 21.5$. This apparent r -band magnitude cut is applied to ensure that galaxies in our sample are below the SDSS saturation limit and above the 95 per cent completeness limit for galaxies in the *gri* passbands.

We exclude the shallower and less complete u and z bands throughout our analysis. Furthermore, as discussed in Bundy et al. (2015), the TYPE-based galaxy classification is contaminated with a non-negligible fraction of stars (~ 10 per cent), which is attributed to point spread function (PSF) characterization issues in the co-added images. Based on visual inspections, we find that the fraction of stars misclassified as galaxies is higher and more pronounced at brighter magnitudes. To eliminate stars from the photometric sample, we remove sources at $15 < r < 18$ that are classified as stars in the corresponding single-pass SDSS images. At the very brightest magnitudes ($r < 15$), where number counts are lower, we remove stars based on a visual inspection of the single-pass SDSS images. Combined, these two procedures remove ~ 6 per cent of sources at $r < 18$ from our sample, such that our final catalogue includes 1, 293, and 392 galaxies with non-extinction-corrected *gri* photometry in S82. Accounting for Galactic extinction does not change our qualitative results, in part due to the low extinction in the S82 field (Schlegel, Finkbeiner & Davis 1998).

2.2 Spectroscopic training set

To train our classification scheme, which aims to identify galaxies as star forming or quenched, we use spectroscopic data products from the Max Planck Institute for Astrophysics and Johns Hopkins University Data Release 7 (DR7) catalogue (MPA-JHU; Kauffmann et al. 2003; Brinchmann et al. 2004) along with photometry from SDSS DR7 (Aihara et al. 2011). The MPA-JHU catalogue is a value-added data set derived from the spectroscopic SDSS DR7, containing stellar mass and star formation rate (SFR) estimates for nearly a million galaxies up to $z \sim 0.3$. When available, the SFRs are derived using the extinction-corrected $H\alpha$ emission line luminosities. For galaxies that lack emission lines, the SFRs are estimated using a relationship between SFR and the 4000 Å break index (D_{4000} ; Bruzual 1983; Hamilton 1985; Brinchmann et al. 2004). Likewise, the stellar masses are computed using model fits to the broad-band *ugriz* photometry (Kauffmann et al. 2003).

We match galaxies in the MPA-JHU and SDSS DR7 catalogues using their unique MJD, plate ID, and fiber ID to construct a cross-matched catalogue that includes both photometric and spectroscopic galaxy properties. These properties include *gri* model magnitudes, specific star formation rates (sSFR; SFR divided by stellar mass), redshifts, and stellar masses. Furthermore, we limit our cross-matched catalogue to only include galaxies in which CLEAN = 1 and RELI-




Figure 1. A heatmap displaying the correlation between observed colours, apparent magnitudes, and sSFRs for galaxies in our spectroscopic training set. In general, supervised NNCs rely heavily on an existing correlation between input features and target variables (i.e. quenched or star-forming label). For our sample, we detect a relatively strong correlation between observed colour and sSFR.

ABLE $\neq 0$. The former is a photometric flag that removes sources suffering from saturation, deblending, and/or interpolation issues. The latter is a spectroscopic flag that omits galaxies with unreliable line profiles and physical parameters. Overall, these cuts remove roughly 3 per cent of galaxies from the original MPA-JHU catalogue. Finally, we limit our sample to only include galaxies at $z < 0.1$ and $M_\star > 10^{6.5} M_\odot$, with measured sSFRs. Overall, our final sample includes $\sim 240\,000$ galaxies, with a median redshift of 0.07, median stellar mass of $2.7 \times 10^{10} M_\odot$, and median r -band magnitude of 17.

2.3 Host sample

Our spectroscopically confirmed host sample is selected from the group catalogue of Yang et al. (2007). We select groups within the S82 footprint at $z < 0.1$ and $10^{13} < M_{\text{halo}} h^{-1} M_\odot < 10^{14}$, excluding groups that are located within 0.5 deg of the edges of the S82 field. Our final sample consists of 110 hosts, with a median redshift of 0.077 and a median halo mass of $1.6 \times 10^{13} M_\odot$. The central galaxies in these groups have a median stellar mass of $1.3 \times 10^{11} M_\odot$.

3 NEURAL NETWORK CLASSIFIER

3.1 Feature selection and pre-processing

The first step in constructing our training set for supervised machine learning involves selecting the appropriate features that will enable our machine learning model to accurately classify galaxies as either star forming or quenched. Moreover, we can only include photometric features since we ultimately seek to apply our NNC to galaxies without spectra. To that end, we construct a heatmap to visualize the degree of correlation, as measured by the Pearson correlation coefficient, between the sSFR of the MPA-JHU galaxies and their




Figure 2. sSFR versus stellar mass for galaxies in our spectroscopic training set. The contours highlight the star-forming and quenched galaxy population within our sample. We divide the galaxy sample at $\text{sSFR} = 10^{-11} \text{ yr}^{-1}$, such that galaxies above this threshold are labelled as star forming and galaxies below this threshold are labelled as quenched.

photometric properties. As shown in Fig. 1, we find a relatively strong negative correlation between the optical colours of the galaxies and their sSFRs, which implies that optically blue (red) galaxies tend to have higher (lower) sSFRs. With this correlation in mind, we construct our training set using only the $g-r$, $r-i$, and $g-i$ observed colours as features. The inclusion of magnitude information (i.e. apparent gri magnitudes) has a negligible effect on the resulting classifications, and as such was not utilized in the final configuration.

The second important step in constructing our training set for supervised machine learning involves systematically labelling galaxies as either star forming or quenched. We achieve this by taking advantage of the strong bimodality in $\text{sSFR}-M_\star$ space, which for our MPA-JHU sample is illustrated in Fig. 2. In particular, we adopt a cut of $\text{sSFR} = 10^{-11} \text{ yr}^{-1}$ as our quenching threshold, such that galaxies above (below) this threshold are labelled as star forming (quenched). This results in a balanced training set where 49 per cent (51 per cent) of galaxies are classified as quenched (star forming). This is important because imbalanced training sets can result in uninformative models that naively overpredict the majority class and underpredict the minority class. Furthermore, we standardize the features of our training set to have a mean of zero and standard deviation of one according to $X_{\text{st}} = (X - \mu)/\sigma$, where X , μ , and σ are the input feature, mean, and standard deviation of the sample, respectively. This pre-processing procedure is implemented to optimize the performance and stability of the NNC, which assumes that the inputs are standardized.

Lastly, to construct our validation set, we remove 6600 out of the $\sim 240\,000$ galaxies in our training set. The validation set is composed of a subset of those galaxies cross-matched between the MPA-JHU and the S82 photometric catalogues using a search radius of 1 arcsec. We omit these galaxies from the training and testing process, so that they can ultimately be used to evaluate the performance of the resultant NNC.




Figure 3. Classification accuracy in the validation set as a function of redshift and stellar mass (*left*) along with observed $g-r$ and $g-i$ colour (*right*). The accuracy of the NNC is largely independent of host halo redshift and weakly dependent on stellar mass, with high-mass galaxies more likely to be incorrectly classified. Despite the training set being largely composed of high-mass galaxies (~ 90 per cent of the training set has $M_* > 10^{9.5} M_\odot$), we find that the overall classification accuracy as well as our primary satellite quenched results remain qualitatively unchanged when high-mass systems are omitted from the training set. As expected, the NNC is less reliable at classifying galaxies at intermediate colours (i.e. in the ‘green valley’ of the colour bimodality) precisely due to the binary nature of the classification scheme.

3.2 Supervised neural network classifier

The supervised NNC is a machine learning model that is trained using labelled observations in order to learn a mapping function between input features and output targets. The utility of these models is that once they are trained they can be readily used to classify unlabelled observations. Moreover, NNCs are constructed using a variety of hyperparameters that influence the overall performance of the machine learning model. The optimal hyperparameters for our NNC are obtained using a K-fold cross-validation grid search. The names and values of these hyperparameters are as follows: (i) the number of hidden layers is 2; (ii) the number of nodes in the first and second hidden layer is 8 and 4, respectively; (iii) the batch size is 64; (iv) the number of epochs is 10; (v) the dropout is 20 per cent. As is standard for binary classification, we use the rectified linear unit (ReLU) activation function for the input and hidden layers, while the sigmoid activation function is used for the output layer. Our model is compiled using a binary cross-entropy loss function and stochastic gradient descent with a learning rate of 0.01. Lastly, we use a stratified K-fold cross validation procedure with $k = 5$ to determine the average accuracy and logarithmic loss of our model.

3.3 Performance of neural network classifier

The K-fold cross validation yields an average classification accuracy of 0.94 and logarithmic loss of 0.17. Here, the accuracy measures the fraction of galaxies that are correctly classified during the training/testing process, while the logarithmic loss measures the uncertainty of the predictions made by the NNC. Therefore, the high average classification accuracy and low logarithmic loss suggest that our NNC returns both accurate and precise classifications. Another diagnostic for determining the reliability of the NNC involves applying the trained model to labelled data that was not utilized during

the training or testing process. In our case, we use our validation set that is composed of a subset of the galaxies cross-matched between our spectroscopic training set and the S82 photometric sample. Upon applying the NNC to our validation set, we find that 93 per cent of galaxies in the validation set are correctly classified as quenched, and 95 per cent of star-forming galaxies in the validation set are correctly classified as star forming. Moreover, we find that the true quenched fraction for the validation set is reproduced by the NNC with an average percent error of ~ 2 per cent, largely independent of stellar mass.

In addition to classifying the galaxies in the validation set, the NNC also provides a classification probability (CP) between 0 and 1 for each prediction such that the CP equals 0 (1) when the model is 100 per cent certain that a given galaxy is quenched (star forming). With this information, we define the classification confidence to be equal to the classification probability when $CP > 0.5$ and equal to $1 - CP$ when $CP < 0.5$. The mean and median classification confidence are 0.927 and 0.98, respectively. Overall, these results provide further confidence in the reliability and accuracy of the predictions made by the N.

Using the validation set, we also explore how the classification accuracy varies with galaxy properties. As shown in Fig. 3, we find that the classification accuracy remains relatively constant across our specified redshift range, such that host haloes at slightly lower redshift are not biased relative to their higher- z counterparts within the sample. We do find, however, a modest correlation between classification accuracy and stellar mass, such that the NNC achieves higher levels of accuracy when classifying lower mass galaxies ($M_* \lesssim 10^9 M_\odot$). While the spectroscopic training set is dominated by more massive galaxies (~ 90 per cent of the spectroscopic training set has $M_* > 10^{9.5} M_\odot$), the classification accuracy – and our primary results regarding the satellite quenched fraction – is qualitatively unchanged when limiting the spectroscopic training set to systems with $10^{6.5} M_\odot < M_* < 10^{9.5} M_\odot$.




Figure 4. An illustration of our background subtraction technique, in which we measure the radial number density of galaxies around our host galaxies (*left*) and in randomly selected background fields (*middle*). Using the photometric sample, we compute the mean number density profile as a function of projected radial distance, averaged over our sample of hosts and 6×10^4 background fields (*right*). The error bars for both radial profiles correspond to 1σ Poisson errors in the measured surface density of galaxies.

As suggested in Fig. 1, $g - r$ and $g - i$ colour are the most informative features with respect to predicting whether a galaxy is star forming or quenched. In Fig. 3, we explore the relationship between the classification accuracy and these two features. As expected, the classification accuracy is highest for very blue and red galaxies, with a modest decrease for galaxies residing in the green valley. This is in part due to the binary nature of our classification scheme (i.e. the lack of a transitory phase between star forming and quenched) along with the overlap between dusty star-forming galaxies and quiescent systems in rest-frame optical colour (e.g. Yan et al. 2006; Maller et al. 2009; Williams et al. 2009).

3.4 Classification of galaxies in S82

Using the hyperparameters discussed in Section 3.1, we train our NNC on the entirety of the spectroscopic training set (Section 2.2). Next, we standardize the apparent $g - r$, $g - i$, and $r - i$ colours of the galaxies in our S82 photometric sample to have a mean of zero and standard deviation of one using the same procedure outlined in Section 3.1. Classifying our photometric sample using the trained NNC, we find that 66 per cent of the galaxies in our photometric sample are classified as quenched while 34 per cent are classified as star forming. We recognize that the fraction of passive galaxies in our photometric sample is biased high due to the inevitable inclusion of high- z galaxies. Many of these high- z sources have red apparent colours, and are more likely to be classified as quenched. Ultimately, the success of our approach relies on correctly classifying the low- z sources (i.e. the satellites of our targeted group sample). With that objective, the next step in our analysis involves combining the classification results with a statistical background subtraction technique to ultimately determine the satellite quenched fraction of our low- z host sample.

4 ANALYSIS OF THE S82 SAMPLE

4.1 Statistical background subtraction

While deep imaging allows the satellite population around nearby hosts to be detected and our NNC is able to robustly classify sources

as star forming or quenched, identifying the true satellites amongst the sea of background sources remains a challenge. This is primarily due to the lack of highly complete line-of-sight velocity information for our photometric sample, which is required to cleanly determine if a particular source is truly a satellite of a given host. However, instead of identifying properties of individual satellites, we employ a statistically driven background subtraction technique that enables us to robustly measure the average properties of the satellite population. Fig. 4 illustrates our methodology, by which we compare the radial distribution of galaxies around nearby hosts to that measured in random positions on the sky. By subtracting the random background, we are able to measure the average properties (e.g. radial profile, rest-frame colour distribution, and quenched fraction) of the underlying satellite population.

This statistical approach has proven effective in previous studies of satellites at intermediate redshift (Nierenberg et al. 2011, 2012; Tal et al. 2013; Kawinwanichakij et al. 2014). In general, the background subtraction procedure utilized in these studies involves measuring the radial distribution of galaxies around spectroscopically confirmed hosts and subtracting the contribution from the background/foreground galaxies. For our analysis, we utilize the 110 centrals from the Yang et al. (2007) group catalogue that overlap with the S82 footprint as our sample of spectroscopically confirmed host galaxies. As stated in Section 2.3, our host sample is situated at $z < 0.1$ and have halo masses between 10^{13} and $10^{14} h^{-1} M_{\odot}$.

Our technique for estimating the contribution from background galaxies involves measuring the radial distribution of galaxies at random positions within the S82 footprint. In particular, we generate 10^6 random positions within S82, assigning each a corresponding redshift between $0.02 < z < 0.1$ as randomly drawn from a uniform distribution. As was done for the host sample, the random positions are also required to be less than 0.5 deg from the edges of the S82 field. We have also considered requiring the random points to be sufficiently far away from the spectroscopic hosts (e.g. $d_{\text{proj}} > 1\text{--}2$ Mpc). However, we ultimately omitted this constraint since both scenarios return qualitatively similar results.

We partition our hosts and random positions into six evenly spaced redshift bins between $0.02 < z < 0.1$. For a given redshift bin, we count the number of quenched and star-forming galaxies in annuli




Figure 5. Comparison between our estimated stellar masses and those provided by the MPA–JHU catalogue. Our stellar mass estimator, which we infer by fitting galaxies in the MPA–JHU catalogue using equation (1), provides robust mass measurements in the absence of multiband photometry. In comparison to the MPA–JHU measurements, the median stellar mass difference is -0.030 dex with a standard deviation of 0.22 dex.

centred on the hosts in bins of r -band magnitude. This procedure is repeated at the location of the background pointings, for which we count the number of quenched and star-forming galaxies in bins of r -band magnitude within annuli centred on 100 random positions. Specifically, the photometric sample is partitioned into seven r -band magnitude bins between $13 < r < 21.5$ and the galaxies are counted in five annuli between 15 and 1000 kpc. For each r -band magnitude and redshift combination, we calculate the average number of quenched and star-forming galaxies per annuli for both the background and the spectroscopically confirmed centrals. Moreover, for each individual host/random position, we calculate the 1σ Poisson error associated with our measurement and propagate this error in the calculation of the average number of galaxies per annuli. Increasing the number of random pointings used to determine the background (i.e. >100) yields no significant change in our results.

The galaxies are counted in the manner outlined above because it allows us to robustly estimate stellar masses for our statistical satellite population by capitalizing on the strong correlation between apparent r -band magnitude and stellar mass at fixed redshift. To determine this mapping from r and z to stellar mass, we fit the following relation to galaxies in the MPA–JHU catalogue:

$$M_*(r, z) = \gamma * r + b(z), \quad (1)$$

where γ and $b(z)$ correspond to the slope and y-intercept of the fit in a given redshift bin. In particular, we fit this relation in redshift bins (with typical width of $\Delta z = 0.005$), such that the statistically inferred galaxy counts as a function of r -band magnitude (following background subtraction) can be mapped to stellar mass based on the redshift of the host system. In Fig. 5, we compare the stellar masses estimated using our best-fitting parameters for equation (1) to the corresponding stellar masses from the MPA–JHU catalogue, which are based on fitting the multiband photometry to model spectral




Figure 6. The average number of satellites as a function of stellar mass in projected distance bins. The vertical error bars give the standard deviation in the distribution of the number of satellites after repeating the background subtraction procedure 100 times, whereas the horizontal error bars represent the standard deviation within the stellar mass bin. For our analysis, we limit our satellite population to systems at $d_{\text{proj}} < 400$ kpc.

energy distributions. Our stellar mass estimates, inferred solely from the observed r -band magnitude, are relatively accurate with a median difference of -0.030 dex and a 1σ scatter of 0.22 dex. There is a slight bias towards our method underpredicting and overpredicting the masses of high-mass and low-mass galaxies, respectively. Not surprisingly, the fits to equation (1) are best at intermediate stellar masses, where the spectroscopic training set is more abundant. Tuning our fits to better reproduce the stellar masses of low-mass systems does not yield a significant change in our results, with our measurements of the satellite quenched fraction computed in bins of stellar mass that exceed the typical measurement uncertainty.

Altogether, the statistical background procedure provides us with a measure of the average number of quenched and star-forming galaxies as a function of projected distance and stellar mass at the location of both the spectroscopically confirmed host galaxies and the random background positions. With these galaxy counts and classifications, we compute the average number of quenched and star-forming satellites as a function of stellar mass and projected distance according to

$$\bar{N}_{\text{sats}}(d_{\text{proj}}, M_*) = \sum (\bar{N}_{\text{back+sats}} - \bar{N}_{\text{back}}), \quad (2)$$

where $\bar{N}_{\text{back+sats}}$ and \bar{N}_{back} are the average number of galaxies measured in annuli centred on the spectroscopically confirmed centrals and random positions, respectively.

In Fig. 6, we show the resulting average number of satellites as a function of stellar mass and projected host-centric distance. We adopt 400 kpc as the outer extent of our groups (roughly R_{200}) based upon a comparison to similar haloes in the IllustrisTNG project (Naiman et al. 2018; Nelson et al. 2018, 2019; Marinacci et al. 2018; Pillepich et al. 2018; Springel et al. 2018). For host haloes at $z = 0$ and $10^{13} h^{-1} M_\odot < M_{200} < 10^{14} h^{-1} M_\odot$ within the TNG300

simulation, a sample of >2000 haloes with a median mass of $M_{200} \sim 1.85 \times 10^{13} h^{-1} M_{\odot}$, the median R_{200} is $430 h^{-1} \text{ kpc}$ with a 1σ scatter of $97 h^{-1} \text{ kpc}$. Given that our measurements are made in projection, we limit our selection of the satellite population to projected distances of $<400 \text{ kpc}$. While this excludes a subset of satellites at hostcentric distances of $400 \text{ kpc} < R < R_{200}$, it also reduces contamination from objects in the surrounding infall regions ($R \sim 1-2 R_{200}$). As discussed in Section 4.2, our results are qualitatively unchanged when including satellites out to projected distances of 600 or 800 kpc.

Selecting satellites within 400 kpc, we find excellent agreement between our inferred satellite stellar mass function and that measured for the IllustrisTNG hosts. As shown in Fig. 7, our integrated satellite counts are very tightly bracketed by the corresponding predicted counts in the TNG100 and TNG300 simulations, where we select satellites at projected distances of $<400 \text{ kpc}$ for hosts with $M_{200} = 10^{13-14} h^{-1} M_{\odot}$. In addition, we compare to the observed satellite mass function from Yang, Mo & van den Bosch (2008, 2009), based on a sample of spectroscopically confirmed satellites in $\sim 300\,000$ low- z groups (see also Vázquez-Mata et al. 2020). Overall, our measured satellite mass function is in remarkably good agreement, especially at low masses (or faint magnitudes). While our background-subtraction technique is unable to identify individual satellite galaxies, it is quite robust at indirectly identifying the satellite population, such that its properties may be characterized.

4.2 Measuring the satellite quenched fraction

As a benchmark for comparison, we measure the quenched fraction as a function of satellite stellar mass for the spectroscopically confirmed satellites in the Yang et al. (2007) group catalogue. We limit our sample of host haloes to those with $10^{13} h^{-1} M_{\odot} < M_{\text{halo}} < 10^{14} h^{-1} M_{\odot}$, identifying satellites as quenched according to the sSFR cut of 10^{-11} yr^{-1} described in Section 3.1. Unlike our photometric analysis, however, we include groups across the entire SDSS spectroscopic footprint – i.e. both within and beyond the S82 footprint. From this parent population, we then select two subsamples at $z < 0.06$ and at $z < 0.1$. The lower- z ($z < 0.06$) sample includes ~ 1500 groups with $\sim 14\,000$ satellite galaxies, complete down to a stellar mass of $\sim 10^{10} M_{\odot}$. The higher- z sample includes more host systems (~ 8000 groups with $\sim 40\,000$ satellites), but only probes down to $\sim 10^{10.5} M_{\odot}$. In agreement with many previous studies of satellite properties at $z \sim 0$ (e.g. Baldry et al. 2006; Wetzel et al. 2013; Woo et al. 2013; Hirschmann et al. 2014; Omand, Balogh & Poggianti 2014), we find that the satellite quenched fraction decreases with decreasing satellite stellar mass, such that nearly all satellites are quenched at $>10^{11} M_{\odot}$ with a quenched fraction of < 50 per cent at $\sim 10^{10} M_{\odot}$.

In an effort to push measurements of the satellite quenched fraction to lower masses (i.e. $<10^{10} M_{\odot}$), we use the background subtraction technique described in Section 4.1 as applied to our photometric sample in S82. Accordingly, we compute the satellite quenched fraction as a function of stellar mass as

$$f_q^{\text{sats}}(d_{\text{proj}}, M_{\star}) = \frac{\bar{N}_{\text{sats,q}}}{\bar{N}_{\text{sats,q}} + \bar{N}_{\text{sats,sf}}}, \quad (3)$$

where $\bar{N}_{\text{sats,sf}}$ and $\bar{N}_{\text{sats,q}}$ are the average number of star-forming and quenched satellites detected at $d_{\text{proj}} < 400 \text{ kpc}$, respectively. As discussed in Section 4.1, we adopt 400 kpc as the outer extent of our groups based upon a comparison to comparable haloes in the IllustrisTNG simulation suite. Our resulting satellite quenched fraction, however, remains qualitatively unchanged when integrating satellite counts out to 600 or 800 kpc.




Figure 7. The cumulative satellite stellar mass function based on our statistical background subtraction technique in comparison to that from spectroscopic observations and simulations. The dark and light crimson bands show our satellite counts (per group) at $R < 400$ and $R < 600 \text{ kpc}$, respectively. The purple band corresponds to the satellite stellar mass function for groups with $M_{\text{halo}} = 10^{13.2-13.8} h^{-1} M_{\odot}$ from Yang et al. (2009), while the black dashed and dotted lines denote the satellite counts for host haloes with $M_{200} = 10^{13-14} h^{-1} M_{\odot}$ and satellites at projected distances of $<400 \text{ kpc}$ within the TNG100 and TNG300 simulations, respectively. We find excellent agreement between our inferred satellite counts and those based on simulations and shallower spectroscopic samples.

Fig. 8 shows the measured satellite quenched fraction as a function of satellite stellar mass using the spectroscopic group membership and our photometric analysis. For the stellar mass range at which both approaches overlap (i.e. $M_{\star} > 10^{10} M_{\odot}$), we find excellent agreement between the independent measurements. This serves as a strong validation of the background-subtraction technique and our classification model.

Using the deeper photometry in S82, we are able to push our measurements of the satellite quenched fraction down to $\sim 10^7 M_{\odot}$, probing satellite quenching in group environments across four orders of magnitude in satellite stellar mass. In contrast to measurements in the high-mass regime ($>10^{10} M_{\odot}$), we find that the satellite quenched fraction in $M_{\text{halo}} \sim 10^{13-14} h^{-1} M_{\odot}$ groups increases below satellite stellar masses of $\sim 10^9 M_{\odot}$. This transition in the quenched fraction suggests a change in the quenching efficiency (and possibly dominant quenching mechanism), such that the suppression of star formation in low-mass satellites is increasingly efficient at $M_{\star} \lesssim 10^9 M_{\odot}$.

5 SUMMARY AND DISCUSSION

We have utilized a combination of supervised machine learning and statistical background subtraction to measure the satellite quenched fraction in group environments across four orders of magnitude in satellite stellar mass ranging from $M_{\star} \sim 10^7-11 M_{\odot}$. Our analysis utilizes an NNC trained on a spectroscopic training set to label galaxies in the co-added S82 photometric catalogue as either star forming or quenched based solely on their $g-r$, $g-i$, and $r-i$ colours. The results from this procedure were subsequently used to statistically identify the quenched and star-forming satellite populations around spectroscopically confirmed hosts within S82 with halo masses of $10^{13-14} h^{-1} M_{\odot}$. The main results from this analysis are as follows:

- (i) Using our photometric approach, we successfully reproduce the measured satellite quenched fraction at $M_{\star} \gtrsim 10^{10} M_{\odot}$, as derived




Figure 8. The satellite quenched fraction as a function of stellar mass for group environments with $M_{\text{halo}} = 10^{13-14} h^{-1} M_{\odot}$. The solid red points represent the median quenched fraction for our statistically derived satellite population. The vertical error bars correspond to the 1σ Poisson error in the quenched fraction, while the horizontal error bars denote the standard deviation of the binned stellar masses. The shaded grey (light grey) band represents the quenched fraction for the spectroscopic members of the Yang et al. (2007) groups at $z < 0.1$ ($z < 0.06$). Our statistically driven approach using S82 photometry successfully reproduces the satellite quenched fraction results at high masses ($> 10^{10} M_{\odot}$), and pushes beyond previous studies to probe satellite quenching down to $10^7 M_{\odot}$. We find an increase in the quenched fraction at low masses ($\lesssim 10^9 M_{\odot}$), potentially indicating an increase in the efficiency of quenching in the low-mass regime.

from spectroscopic studies in the local Universe. We find that the satellite quenched fraction increases with increasing satellite mass at $M_{\star} \gtrsim 10^{10} M_{\odot}$.

(ii) We measure the satellite quenched fraction down to $M_{\star} \sim 10^7 M_{\odot}$, pushing measurements of satellite quenching in $\sim 10^{13-14} h^{-1} M_{\odot}$ haloes to a new regime that is not readily probed outside of the Local Group.

(iii) We find that the satellite quenched fraction increases towards lower satellite masses below $\sim 10^9 M_{\odot}$.

(iv) The increase in satellite quenching at low masses potentially indicates a change in the dominant quenching mechanism at $\sim 10^9 M_{\odot}$, where ram-pressure stripping begins to become increasingly effective (see discussion that follows).

Given that low-mass field galaxies are almost entirely star forming as a population, the increase in the satellite quenched fraction at $< 10^9 M_{\odot}$ can be interpreted as a corresponding increase in the satellite quenching efficiency within $10^{13-14} h^{-1} M_{\odot}$ haloes. This increase is similar to that observed in the Local Group, where there is an apparent transition in the dominant quenching mechanism at $\sim 10^8 M_{\odot}$ with lower mass satellites quenched more efficiently following infall. Both hydrodynamic simulations and analytical modelling of the satellite population find that ram-pressure stripping is much more efficient below $10^8 M_{\odot}$ within Milky Way-like galaxies (Mayer et al. 2007; Fillingham et al. 2016; Simpson et al. 2018; Akins et al. 2020), while more massive satellites are primarily quenched via starvation (Fillingham et al. 2015). Given that our host sample is more massive ($M_{\text{halo}} = 10^{13-14} M_{\odot}$) relative to Milky Way-like haloes, it is expected that an increase in infall velocities and the density of the circumgalactic medium would cause this transition mass to increase, such that starvation is the primary driver of satellite

quenching above $\sim 10^{9.5} M_{\odot}$ and ram-pressure stripping becoming increasingly important in the low-mass regime. A more detailed study of the potential quenching mechanisms at play requires further analysis of the time-scales on which the observed satellites are quenched following infall to the host haloes. In future work (Baxter et al., in preparation), we aim to bridge this gap by combining the measured satellite quenched fractions from this work with the accretion and orbital histories determined using high-resolution cosmological simulations, to estimate the typical quenching time-scale as a function of satellite mass.

The satellite quenched fractions that we obtain at low masses ($M_{\star} < 10^9 M_{\odot}$) are generally lower than what have been reported in studies of dwarf galaxies in more massive nearby clusters. For example, Weinmann et al. (2011) studied the satellite galaxy population in the nearby Virgo ($M_{\text{halo}} \sim 1.4-4 \times 10^{14} M_{\odot}$), Coma ($M_{\text{halo}} \sim 1.3 \times 10^{15} M_{\odot}$), and Perseus ($M_{\text{halo}} \sim 6.7 \times 10^{14} M_{\odot}$) clusters, finding red fractions between 70 and 80 per cent at stellar masses of $\sim 10^{8-10} M_{\odot}$ (see also Boselli et al. 2016). At slightly higher redshift ($z \sim 0.2$), analysis of the satellite population in Abell 209 ($M_{\text{halo}} \sim 10^{15} M_{\odot}$) by Annunziatella et al. (2016) also finds an elevated quenched fraction relative to our results in less massive haloes. Interestingly, while the study of Annunziatella et al. (2016) only probes down to $\sim 10^{8.6} M_{\odot}$ in satellite stellar mass, the results show a quenched fraction that decreases from near unity (~ 95 per cent) at $M_{\star} \sim 10^{10.5} M_{\odot}$ to ~ 75 per cent at $M_{\star} \sim 10^9 M_{\odot}$ (see also Sarrouh et al., in preparation). Naively, if there is a transition in the dominant quenching mechanism (or efficiency) in these massive clusters similar to that found in the Local Group and our group sample, we would expect the transition scale to occur at higher satellite masses (e.g. $\gtrsim 10^{9.5} M_{\odot}$) as ram-pressure stripping (and other cluster-specific processes) should be increasingly effective

in hosts with $M_{\text{halo}} \sim 10^{15} M_{\odot}$. Extrapolations of the mass functions from Annunziatella et al. (2016), however, do not support this picture.

Finally, we report satellite quenched fractions in the low-mass regime ($<10^8 M_{\odot}$) that are potentially lower than expected when compared to studies of satellite quenching in the Local Group, where ~ 90 per cent of satellites with $M_{\star} < 10^8 M_{\odot}$ are passive. As discussed above, environmental quenching mechanisms are expected to be more efficient in our more massive host haloes relative to the Local Group. Of course, our results are based on a study of ~ 100 groups, whereas studies of the Local Group satellites sample only two host haloes. While observations of the nearby M81 group yield a satellite quenched fraction comparable to that measured for the Local Group (Kaisin & Karachentsev 2013; Karachentsev, Makarov & Kaisina 2013), various studies also indicate that the Local Group satellites may be outliers relative to the cosmic mean (e.g. Boylan-Kolchin et al. 2010; Busha et al. 2011; Tollerud et al. 2011; Ibata et al. 2013; Pawlowski & Kroupa 2020). Moreover, recent results from the Satellites Around Galactic Analogs (SAGA) survey (Geha et al. 2017; Mao et al. 2021) find lower satellite quenched fractions (~ 20 per cent) around hosts with halo masses comparable to those of the Milky Way and M31. We contend that the application of the methodology presented in this work to Milky Way-like hosts is an intriguing way to better place the Local Group into a cosmological context and constrain the quenching of satellites around hosts with $M_{\text{halo}} \sim 10^{12} M_{\odot}$.

ACKNOWLEDGEMENTS

DCB thanks the LSSTC Data Science Fellowship Programme, which is funded by LSSTC, NSF Cybertraining Grant #1829740, the Brinson Foundation, and the Moore Foundation; participation in the program has greatly benefited this work. MCC thanks Benedetta Vulcani for helpful discussions that improved this work. We also thank the anonymous referee for helping to improve the clarity of this work.

This work was supported in part by NSF grants AST-1815475 and AST-1518257. Additional support was provided by NASA through grant AR-14289 from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555. This research made extensive use of ASTROPY, a community-developed core PYTHON package for astronomy (Astropy Collaboration 2013, 2018). Additionally, the PYTHON packages NUMPY (Van Der Walt, Colbert & Varoquaux 2011), IPYTHON (Pérez & Granger 2007), SCIPY (Virtanen et al. 2020), SCIKIT-LEARN (Pedregosa et al. 2011), KERAS (Chollet et al. 2015), and MATPLOTLIB (Hunter 2007) were utilized for our data analysis and presentation.

DATA AVAILABILITY

Data sharing is not applicable to this article as no new data were created or analysed in this study.

REFERENCES

Abadi M. G., Moore B., Bower R. G., 1999, *MNRAS*, 308, 947
 Aihara H. et al., 2011, *ApJS*, 193, 29
 Akins H. B., Christensen C. R., Brooks A. M., Munshi F., Applebaum E., Angelhardt A., Chamberland L., 2020, preprint (arXiv:2008.02805)
 Annis J. et al., 2014, *ApJ*, 794, 120
 Annunziatella M. et al., 2016, *A&A*, 585, A160
 Astropy Collaboration, 2013, *A&A*, 558, A33

Astropy Collaboration, 2018, *AJ*, 156, 123
 Baldry I. K., Balogh M. L., Bower R. G., Glazebrook K., Nichol R. C., Bamford S. P., Budavari T., 2006, *MNRAS*, 373, 469
 Bell E. F. et al., 2004, *ApJ*, 608, 752
 Boselli A. et al., 2016, *A&A*, 585, A2
 Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., 2010, *MNRAS*, 406, 896
 Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, 351, 1151
 Bruzual A. G., 1983, *ApJ*, 273, 105
 Bundy K. et al., 2006, *ApJ*, 651, 120
 Bundy K. et al., 2015, *ApJS*, 221, 15
 Busha M. T., Wechsler R. H., Behroozi P. S., Gerke B. F., Klypin A. A., Primack J. R., 2011, *ApJ*, 743, 117
 Ceverino D., Klypin A., 2009, *ApJ*, 695, 292
 Chollet F. et al., 2015, Keras. Available at: <https://keras.io> (last accessed March 8, 2021)
 Croton D. J. et al., 2006, *MNRAS*, 365, 11
 De Lucia G., Weinmann S., Poggianti B. M., Aragón-Salamanca A., Zaritsky D., 2012, *MNRAS*, 423, 1277
 Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
 Driver S. P. et al., 2009, *Astron. Geophys.*, 50, 5.12
 Driver S. P. et al., 2011, *MNRAS*, 413, 971
 Faber S. M. et al., 2007, *ApJ*, 665, 265
 Fillingham S. P., Cooper M. C., Wheeler C., Garrison-Kimmel S., Boylan-Kolchin M., Bullock J. S., 2015, *MNRAS*, 454, 2039
 Fillingham S. P., Cooper M. C., Pace A. B., Boylan-Kolchin M., Bullock J. S., Garrison-Kimmel S., Wheeler C., 2016, *MNRAS*, 463, 1916
 Fillingham S. P., Cooper M. C., Boylan-Kolchin M., Bullock J. S., Garrison-Kimmel S., Wheeler C., 2018, *MNRAS*, 477, 4491
 Geha M., Blanton M. R., Yan R., Tinker J. L., 2012, *ApJ*, 757, 85
 Geha M. et al., 2017, *ApJ*, 847, 4
 Gnedin O. Y., 2003, *ApJ*, 582, 141
 Grcevich J., Putman M. E., 2009, *ApJ*, 696, 385
 Gunn J. E., Gott J. R. III, 1972, *ApJ*, 176, 1
 Haines C. P., Gargiulo A., Merluzzi P., 2008, *MNRAS*, 385, 1201
 Hamilton D., 1985, *ApJ*, 297, 371
 Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R. E., Guo Q., Lemson G., Wang W., 2017, *MNRAS*, 469, 2626
 Hirschmann M., De Lucia G., Wilman D., Weinmann S., Iovino A., Cucciati O., Zibetti S., Villalobos Á., 2014, *MNRAS*, 444, 2938
 Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Martini P., Robertson B., Springel V., 2005, *ApJ*, 630, 705
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Ibata R. A. et al., 2013, *Nature*, 493, 62
 Kaisin S. S., Karachentsev I. D., 2013, *Astrophysics*, 56, 305
 Karachentsev I. D., Makarov D. I., Kaisina E. I., 2013, *AJ*, 145, 101
 Kauffmann G. et al., 2003, *MNRAS*, 341, 33
 Kawata D., Mulchaey J. S., 2008, *ApJ*, 672, L103
 Kawinwanichakij L. et al., 2014, *ApJ*, 792, 103
 Kimm T. et al., 2009, *MNRAS*, 394, 1131
 Larson R. B., Tinsley B. M., Caldwell C. N., 1980, *ApJ*, 237, 692
 Maller A. H., Berlind A. A., Blanton M. R., Hogg D. W., 2009, *ApJ*, 691, 394
 Mao Y.-Y., Geha M., Wechsler R. H., Weiner B., Tollerud E. J., Nadler E. O., Kallivayalil N., 2021, *ApJ*, 907, 85
 Marinacci F. et al., 2018, *MNRAS*, 480, 5113
 Mayer L., Kazantzidis S., Mastropietro C., Wadsley J., 2007, *Nature*, 445, 738
 Merritt D., 1983, *ApJ*, 264, 24
 Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, *Nature*, 379, 613
 Moore B., Lake G., Katz N., 1998, *ApJ*, 495, 139
 Moore B., Lake G., Quinn T., Stadel J., 1999, *MNRAS*, 304, 465
 Naiman J. P. et al., 2018, *MNRAS*, 477, 1206
 Nelson D. et al., 2018, *MNRAS*, 475, 624
 Nelson D. et al., 2019, *Comput. Astrophys. Cosmol.*, 6, 2
 Nierenberg A. M., Auger M. W., Treu T., Marshall P. J., Fassnacht C. D., 2011, *ApJ*, 731, 44

- Nierenberg A. M., Auger M. W., Treu T., Marshall P. J., Fassnacht C. D., Busha M. T., 2012, *ApJ*, 752, 99
- Oke J. B., Gunn J. E., 1983, *ApJ*, 266, 713
- Omand C. M. B., Balogh M. L., Poggianti B. M., 2014, *MNRAS*, 440, 843
- Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
- Pawlowski M. S., Kroupa P., 2020, *MNRAS*, 491, 3042
- Pedregosa F. et al., 2011, *J. Mach. Lear. Res.*, 12, 2825
- Pérez F., Granger B. E., 2007, *Comput. Sci. Eng.*, 9, 21
- Pillepich A. et al., 2018, *MNRAS*, 475, 648
- Rodríguez Wimberly M. K., Cooper M. C., Fillingham S. P., Boylan-Kolchin M., Bullock J. S., Garrison-Kimmel S., 2019, *MNRAS*, 483, 4031
- Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525
- Simpson C. M., Grand R. J. J., Gómez F. A., Marinacci F., Pakmor R., Springel V., Campbell D. J. R., Frenk C. S., 2018, *MNRAS*, 478, 548
- Spekkens K., Urbancic N., Mason B. S., Willman B., Aguirre J. E., 2014, *ApJ*, 795, L5
- Springel V. et al., 2018, *MNRAS*, 475, 676
- Tal T., van Dokkum P. G., Franx M., Leja J., Wake D. A., Whitaker K. E., 2013, *ApJ*, 769, 31
- Tollerud E. J., Boylan-Kolchin M., Barton E. J., Bullock J. S., Trinh C. Q., 2011, *ApJ*, 738, 102
- Van Der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
- Vázquez-Mata J. A. et al., 2020, *MNRAS*, 499, 631
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Wang W., Sales L. V., Henriques B. M. B., White S. D. M., 2014, *MNRAS*, 442, 1363
- Weinmann S. M., Lisker T., Guo Q., Meyer H. T., Janz J., 2011, *MNRAS*, 416, 1197
- Weinmann S. M., Pasquali A., Oppenheimer B. D., Finlator K., Mendel J. T., Crain R. A., Macciò A. V., 2012, *MNRAS*, 426, 2797
- Weisz D. R., Dolphin A. E., Skillman E. D., Holtzman J., Gilbert K. M., Dalcanton J. J., Williams B. F., 2014a, *ApJ*, 789, 147
- Weisz D. R., Dolphin A. E., Skillman E. D., Holtzman J., Gilbert K. M., Dalcanton J. J., Williams B. F., 2014b, *ApJ*, 789, 148
- Weisz D. R., Dolphin A. E., Skillman E. D., Holtzman J., Gilbert K. M., Dalcanton J. J., Williams B. F., 2015, *ApJ*, 804, 136
- Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013, *MNRAS*, 432, 336
- Wetzel A. R., Tollerud E. J., Weisz D. R., 2015, *ApJ*, 808, L27
- Wheeler C., Phillips J. I., Cooper M. C., Boylan-Kolchin M., Bullock J. S., 2014, *MNRAS*, 442, 1396
- Williams R. J., Quadri R. F., Franx M., van Dokkum P., Labbé I., 2009, *ApJ*, 691, 1879
- Woo J. et al., 2013, *MNRAS*, 428, 3306
- Yan R., Newman J. A., Faber S. M., Konidaris N., Koo D., Davis M., 2006, *ApJ*, 648, 281
- Yang X., Mo H. J., van den Bosch F. C., Pasquali A., Li C., Barden M., 2007, *ApJ*, 671, 153
- Yang X., Mo H. J., van den Bosch F. C., 2008, *ApJ*, 676, 248
- Yang X., Mo H. J., van den Bosch F. C., 2009, *ApJ*, 695, 900
- York D. G. et al., 2000, *AJ*, 120, 1579

This paper has been typeset from a $\mathrm{T}_{\mathrm{E}}\mathrm{X}/\mathrm{L}^{\mathrm{A}}\mathrm{T}_{\mathrm{E}}\mathrm{X}$ file prepared by the author.