



LJMU Research Online

Sarmiento, R, Huertas-Company, M, Knapen, JH, Sanchez, SF, Dominguez Sanchez, H, Drory, N and Falcon-Barroso, J

Capturing the Physics of MaNGA Galaxies with Self-supervised Machine Learning

<http://researchonline.ljmu.ac.uk/id/eprint/17793/>

Article

Citation (please note it is advisable to refer to the publisher's version if you intend to cite from this work)

Sarmiento, R, Huertas-Company, M, Knapen, JH, Sanchez, SF, Dominguez Sanchez, H, Drory, N and Falcon-Barroso, J (2021) Capturing the Physics of MaNGA Galaxies with Self-supervised Machine Learning. ASTROPHYSICAL JOURNAL. 921 (2). ISSN 0004-637X

LJMU has developed [LJMU Research Online](#) for users to access the research output of the University more effectively. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LJMU Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain.

The version presented here may differ from the published version or from the version of the record. Please see the repository URL above for details on accessing the published version and note that access may require a subscription.

For more information please contact researchonline@ljmu.ac.uk

<http://researchonline.ljmu.ac.uk/>

Capturing the physics of MaNGA galaxies with self-supervised Machine Learning

REGINA SARMIENTO,^{1,2} MARC HUERTAS-COMPANY,^{1,2,3} JOHAN H. KNAPEN,^{1,2} SEBASTIÁN F. SÁNCHEZ,⁴ HELENA DOMÍNGUEZ SÁNCHEZ,^{5,6} NIV DRORY,⁷
AND JESUS FALCÓN-BARROSO^{1,2}

¹*Instituto de Astrofísica de Canarias (IAC) La Laguna, 38205, Spain*

²*Departamento de Astrofísica, Universidad de La Laguna (ULL) E-38200, La Laguna, Spain*

³*LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Universités, UPMC Univ. Paris 06, F-75014 Paris, France*

⁴*Instituto de Astronomía, Universidad Nacional Autónoma de México*

⁵*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA*

⁶*Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Magrans, E-08193 Barcelona, Spain*

⁷*McDonald Observatory, The University of Texas at Austin, 1 University Station, Austin, TX 78712, USA*

(Accepted August 14, 2021)

Submitted to ApJ

ABSTRACT

As available data sets grow in size and complexity, advanced visualization tools enabling their exploration and analysis become more important. In modern astronomy, integral field spectroscopic galaxy surveys are a clear example of increasing dimensionality and complexity of datasets, which challenge the traditional methods used to extract the physical information they contain. We present the use of a novel self-supervised Machine Learning method to visualize the multi-dimensional information on stellar population and kinematics in the MaNGA survey in a two dimensional plane. Our framework is insensitive to non-physical properties such as the size of integral field unit (IFU) and is therefore able to order galaxies according to their resolved physical properties. Using the extracted representations, we study how galaxies distribute based on their resolved and global physical properties. We show that even when using exclusively information about the internal structure, galaxies naturally cluster into two well-known categories from a purely data driven perspective: rotating main-sequence disks and massive slow rotators, hence confirming distinct assembly channels. Low-mass rotation-dominated quenched galaxies appear only as a third cluster if information about the integrated physical properties is preserved, suggesting a mixture of assembly processes for these galaxies without any particular signature in their internal kinematics that distinguishes them from the two main groups. The framework for data exploration is publicly released with this publication, ready to be used with the MaNGA or other integral field data sets.

Keywords: Galaxies, Neural networks

1. INTRODUCTION

Over the past decades, the sizes of available data sets in astronomy have grown by orders of magnitude (e.g. Sloan Digital Sky Survey (SDSS; Gunn et al. 2006), COSMOS (Taniguchi et al. 2007)) and the trend is expected to accelerate in this new decade (e.g. Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), Ivezić et al. 2019; Euclid, Laureijs et al. 2011; Nancy Grace Roman Space Telescope, Spergel et al. 2015). As data sets grow not only in size but also in complexity, efficient exploratory methods are

needed to boost discovery. Dimensionality reduction techniques are vital in this respect (Baron & Ménard 2020, Portillo et al. 2020). However, astronomy data present some specific problems such as noise and complex selection effects which challenge the use of standard techniques out of the box. This is a field where recent advances in Machine Learning can undoubtedly be of major utility.

We explore in this work the application of self-supervised Contrastive Learning as a tool to explore multi-dimensional astronomy data. We employ it to analyse maps derived from the SDSS-IV (Blanton et al. 2017) Mapping Nearby Galaxies at APO (MaNGA; Bundy et al. 2015) survey of close to 10^4 nearby galaxies with dedicated integral field spectroscopic fiber bundles. This is a particularly complex survey in which

multiple sampling sizes are used and it presents a complex selection function.

Contrastive Learning has rapidly emerged in the Machine Learning community as a powerful solution to find meaningful representations of data. As opposed to more classical dimensionality reduction approaches such as principal component analysis (PCA; Pearson 1901) or Auto Encoders (Rumelhart et al. 1986) which aim at projecting data into a low dimensional space to minimize the reconstruction error, contrastive learning maximizes the agreement between different transformations of the same data point with no reconstruction involved in the process. This is particularly well adapted to datasets where observational effects can make objects with similar physical properties appear very different (e.g. due to PSF variations, sampling, S/N etc). The representations learned can not only be used for exploration but also for downstream tasks such as anomaly detection or classification with the advantage of reducing the sizes of the training sets. A major breakthrough of last year was indeed that a supervised classification based on contrastive learning representations achieved better accuracy than a purely convolutional neural network-based supervised classification on the ImageNet database (Chen et al. 2020b). Even more recently, Abul Hayat et al. (2020) proved a similar behavior in astronomical data. They analyzed multi-band optical images from the SDSS and found that photometric redshift and galaxy morphology classifications can be estimated with the same accuracy as a completely supervised approach but using half the number of labels during training.

In this work, we use contrastive learning to explore the multi-dimensional kinematic and stellar population maps of galaxies derived from the MaNGA dataset. This is the first time contrastive learning has been applied to integral field spectroscopy data. We use prior knowledge on observational biases in order to remove these artificial features and extract physical information from the data. We aim to explore to what extent physically relevant parameters and relations can be reproduced or uncovered automatically. Testing such algorithms on MaNGA data allows us to compare our findings with results obtained within the survey project using more traditional analysis tools. This is an important validation step which must be taken before we can consider using self-supervised analysis on larger, more complicated, or newer data sets.

This paper is organised as follows. In Section 2 we describe the data we used, and in Section 3 the machine learning software applied. Section 4 describes our results in terms of visualisation, regression, feature decomposition and clustering. Implications of our work are discussed in Section 5, and summarised in Section 6.

2. DATA

We use data products from the MaNGA survey (Bundy et al. 2015, Drory et al. 2015, Aguado et al. 2019). MaNGA is one of the three core surveys in the fourth-generation SDSS (SDSS-IV; Blanton et al. 2017, Smee et al. 2013). It uses 17 fiber-bundle integral field units (IFUs) that vary in diameter from 12 arcsec (19 fibers) to 32 arcsec (127 fibers of 2 arcsec diameter) to observe a total of 10^4 nearby galaxies (Law et al. 2015, Yan et al. 2016a, Wake et al. 2017), constituting a representative sample of all types of nearby galaxies in a redshift range $0.01 < z < 0.15$. Data-cubes are obtained with IFUs of 19, 37, 61, 91 and 127 fibers, matched to the galaxy size, with spectra covering the range 360–1000nm in wavelength at a spectral resolution $R \sim 2000$ (Law et al. 2016, Yan et al. 2016b).

The data set used in this work consists of post-processed maps derived from MaNGA MPL-10 data release using the Pipe3D pipeline (v3.0.1b) (Sánchez et al. 2015, Sánchez et al. 2016). Pipe3D performs a spectral fit of a combination of three synthetic components: stellar populations, dust and gas. The pipeline produces maps by deriving the properties of averaged stellar population. To test the performance of the algorithm, we use integrated magnitudes derived from the post-processed MaNGA data cubes, which are included in the DR14Pipe3D primary catalog.

The input data for which we aim at finding low-dimensional representations consists of five binned (2×2) and stacked hexagonal maps for each galaxy, namely the V-band reconstructed images, luminosity-weighted age and metallicity maps, and radial velocity and velocity dispersion maps. Previous to the binnig, the pixels with a signal to noise ratio lower than 3 in the median intensity maps are set to zero. As the MaNGA IFUs have different numbers of fibers, the sizes of the resulting maps vary as well. Since convolutional neural networks (CNNs) typically require input maps of a fixed shape, we zero-padded the hexagonal maps to the same shape and size 32×32 . In order to deal with the different dynamic ranges, we scaled all maps linearly to the range $[0,1]$ except for the V-band image, to which we applied logarithmic and square root functions previous to the linear normalization. Two different approaches will be analyzed regarding the normalization. On the one hand, we consider a normalization that is equivalent to a change of units that is uniformly applied to all the galaxies (we will refer to it as *relative norm*), as we aim to analyze the relative physical magnitudes among the galaxies. In this case, all maps have values between 0–1, but not necessarily span the complete range. On the other hand, a normalization that forces each galaxy map to span the $[0,1]$ range is used to evaluate the effect of the spatial structures of the galaxies (the *individual norm*).

Our final dataset consists of 9507 MaNGA galaxies, of which 1099, 2235, 2235, 1122, and 2816 were observed with a 19, 37, 61, 91, and 127-fiber bundle, respectively.

3. DERIVING MEANINGFUL REPRESENTATIONS OF MANGA GALAXIES

3.1. Contrastive model

We use the Simple framework for Contrastive Learning of visual Representations (SimCLR; Chen et al. 2020a). SimCLR is a contrastive learning framework to extract meaningful representations of the data. More precisely, we use the model in Tensorflow (Abadi et al. 2015) and follow the findings on optimization by Chen et al. (2020a) with minor changes in the architecture.

This algorithm optimises representations that keep the semantic, or physical, information of the data by maximizing the agreement between the representations of a data object and its transformed pair (x, x^T) . This is particularly beneficial when similar objects appear to be different due to observational biases. By selecting transformations that simulate these observational biases the model is encouraged to neglect these differences while capturing the common remaining features, which correspond solely to the galaxy. Likewise, agreement is minimized when representations originate in different galaxy maps. This way, the program is self-taught to recognize meaningful patterns present in the data.

More specifically, the model consists of a CNN block followed by a set of fully connected layers after which the contrastive loss l is computed. The CNN block acts as a base encoder that outputs the representations, while the fully connected layers project the representations to the comparison space. We will refer to these last fully connected layers as head projection function. While both components are trained together, the head projection is removed to obtain the final galaxy representations. Chen et al. (2020a) found that this configuration yields better representations in the previous layers.

The CNN base encoder consists of four convolutional layers with kernel sizes 5, 3, 3, 3 and 128, 256, 512, 1024 filters each. Max-pooling layers and Exponential Linear Unit activation functions (Clevert et al. 2016) are used between the convolutional layers. The output of this set of layers encodes the array-like representations of the galaxies of 1024 features each.

The projection head is composed of three fully connected layers (512, 128 and 64 neurons each). This non-linear function takes the galaxy representations to a latent space where the contrastive loss is computed.

The contrastive loss is calculated for each L^2 -normalized representation z_i in the contrastive space with the following equation:

$$l_{i,j} = -\log \frac{\exp(\langle z_i, z_j \rangle / h)}{\sum_{k=1, k \neq i}^{2N} \exp(\langle z_i, z_k \rangle / h)}, \quad (1)$$

where \langle, \rangle is the dot product, h is a temperature factor that regulates the distribution of the output representations by de-

creasing or increasing its concentration (Wu et al. 2018, Hinton et al. 2015), in our case, h is set to 0.5. z_j denotes the representation of the augmented view of x_i , while k sweeps the other $2N - 1$ representations in the batch.

This loss, first introduced in Sohn (2016), does not assume a prior distribution of the representation space and therefore can be used to learn representations of both continuous or discrete data sets. The model benefits from large number of negative examples and long training times (Chen et al. 2020a). Our model was trained using a batch size of 1024 during 500 epochs, which represents 40 minutes of computing time on one GeForce RTX 2080 Ti GPU with NVIDIA V-450.102.04.

3.2. Augmentations

Data augmentation is a common practice in Machine Learning that consists of applying transformations to the input data during the training step. The goal is to create different views of the objects while keeping their semantic information. In a supervised approach, data augmentation can be used to enlarge the training sample, since the different views of the data will have the same labels. In the contrastive learning framework, augmentations are essential to create positive pairs of objects, these are the transformed images that originate from the same image.

When generating meaningful representations of the data in our machine learning approach, we seek that galaxies that are similar in intrinsic physical parameters will appear close to each other in the representation space. A standard but non-optimal feature decomposition of the data may reproduce features that are intrinsic to the data, but not representative of the galaxies. The possible origin of such features is diverse, and may be related to the survey design, to instrumental effects, to the pipeline used to process the data, or top observational constraints such as orientation or apparent size of the target. Therefore, the augmentation functions we used during the training of the CNN were designed for the MaNGA data set and with the purpose of mitigating non-physical dependencies as detailed below and illustrated in Figure 1.

- **Model variation:** To prevent the model from learning modelling systematics that could be introduced by the pipeline, a transformation that replaces the input set of maps by a set of differently modelled maps is included. With this transformation, the kinematic maps are replaced by those obtained from the MaNGA Data Analysis Pipeline (DAP) which uses a Voronoi binning to ensure a S/N=10 in the g-band (Cappellari & Copin 2003) and MILES template library for the stellar continuum fitting (Sánchez-Blázquez et al. 2006, Falcón-Barroso et al. 2011). The luminosity weighted age and metallicity maps are replaced by the FIREFLY catalog (Neumann et al. 2021, Goddard et al. 2017)

analogs. The V-band reconstructed image is replaced by the scale matched SDSS r-band photometric image¹, which represents a change of view of the galaxy morphological features rather than a modelling variation.

- Random flip and rotation (0° , 90° , 180° , 270°): Since there is no preference in the orientation of the galaxies, the model should be invariant to the galaxy position. We quantify this dependence with the galaxies position angle and rotation angle. While the first is used as listed in the DR14Pipe3D catalog, the second was estimated from each velocity map as the perpendicular angle to the direction defined by the mean position of the negative pixels and the mean position of the positive pixels.
- Noise perturbation: The data is affected by errors inherent to the instrument and the post-processing of the observations. This transformation aims to compensate for measurement uncertainties by shifting the values of the map pixels. The transformation consists of multiplying each value by a random normal factor $\mathcal{N}(1, \sigma)$, with σ adjusted to the expected error for each pixel of each map. σ maps are calculated as the averaged signal-to-noise maps for each channel².
- Resize: Since the model uses a fixed size for the input, IFUs with a smaller number of fibers will be more intensively zero-padded than IFUs with more fibers. Even though the IFU size used for each galaxy is linked to its apparent size, the MaNGA survey has an uneven sampling of galaxy angular sizes to IFU size. The IFU size is chosen for each galaxy by optimizing a $1.5R_e$ angular coverage of the target object in the Primary sample, and $2.5R_e$ in the Secondary sample (Bundy et al. 2015), which means that the ratio $D_{\text{IFU}}[\text{arcsec}]/R_e[\text{arcsec}]$ ranges from 1.5 to 2.5 when combining Primary and Secondary samples. To prevent the algorithm from considering the IFU size and resolution as relevant features, the resize transformation randomly enlarges or reduces the image size to up to 25% of the side size of the original frame. Cubic interpolation is used.
- Gaussian blur: A 2D Gaussian kernel of 3×3 pixels² and σ [pixels] from $\mathcal{U}(0.1, 2)$ is convolved with each layer of the data input. This encourages the model to




Figure 1. Augmentations used during training and their effect on the different input channels of an example galaxy. The leftmost column shows the original five maps of the example galaxy in adimensional units, from top to bottom: V-band reconstructed image, luminosity weighted age and metallicity, radial velocity and velocity dispersion. The columns to the right show the resulting maps after each transformation has been applied. Namely, model variation, rotation, resize, noise perturbation and blur.

become invariant to different resolutions and sharp artificial features in the data maps.

While Noise perturbation, Resize and Gaussian blur transformations have a 50% probability of being applied to the raw data input copy, Random flip and rotation is always applied for a better performance. The Model variation augmentation is applied with a lower probability since the Voronoi tessellation of the stellar maps could affect the performance of the kernels of the CNN. Therefore, this augmentation randomly replaces 25% of the copies.

4. REMOVING INSTRUMENTAL BIASES: COMPARISON WITH PCA

We now assess how the learned representations correlate with known physical properties of the data. For this purpose, the input data normalized to keep relative values between galaxies is used. We first visually explore how galaxies are distributed in the representation space. We then illustrate a clustering application to characterize groups in specific regions of the representation space.

4.1. Two-dimensional visualization of representations

In order to easily visualize the representation space we perform a dimensionality reduction of the representations of dimension 1024 to a 2D space using Uniform Manifold Approximation and Projection (UMAP; McInnes et al. 2020) and color-code using different physical and instrumental properties. This is shown in Fig. 2. We emphasize that

¹ SDSS r-band photometric image as in: <http://skyserver.sdss.org/ImgCutoutDR7/ImgCutout.asmx>

² As error maps for the V-band reconstructed image were not available, a $\sigma = 0.1$ has been assumed for this channel.

the UMAP representation is only used here for visualization purposes. The actual dimensionality of the representation space is higher than these two dimensions. Our goal is to understand how galaxies are organized in the representation space and more precisely whether non-physical dependencies have been properly removed. We therefore include, on the one hand, the parameters related to instrumental effects (leftmost and left-center panels in Fig. 2), which are (top to bottom): number of fibers in the IFU, galaxy angular size (R_e [arcsec]), number of zero pixels, position and rotation angle of the galaxy (PA and RA, respectively), angular coverage ($D_{\text{IFU}}[\text{arcsec}]/R_e[\text{arcsec}]$) and redshift z to account for physical resolution. On the other hand, we include a set of integrated physical properties derived from the input maps (right-center and rightmost panels), namely (top to bottom) effective radius (R_e [kpc]), specific angular momentum within $1.5R_e$ (λ_{R_e}), radial velocity dispersion at center of frame (σ_{cen}), luminosity weighted age at R_e and luminosity weighted metallicity at R_e normalized by the solar metallicity ($[Z/H]_{R_e}$). We also include the slopes of the gradients within 0.5 to $2R_e$ of the last two parameters (∇age_{R_e} and $\nabla[Z/H]_{R_e}$).

To better quantify the efficiency of the removal of instrumental dependencies, we contrast representations obtained by a PCA with those obtained with the SimCLR algorithm. For the first set, the five MaNGA maps are decomposed to 1024 principal components, to match the SimCLR representations dimension. In all panels, the distribution of the 2D projection of the PCA feature decomposition, on the left (labelled with "a" in the top right corner), with the representations learned by contrastive learning, on the right (labelled with "b" in the top right corner). These projections are simply obtained by applying the UMAP to the features obtained with PCA, on the one hand, and with SimCLR, on the other.

The principal components directly projected via UMAP are primarily organized based on instrumental and observational parameters (Fig. 2). The main features that drive the organization of the PCA space are indeed the number of zero pixels, the number of fibers on the IFU and the angular size of the galaxies (panels a1, a2 and a3). Although some slight dependencies on these parameters are recognizable in the SimCLR representation space, they are remarkably reduced (panels b1, b2 and b3). More precisely, we find that the principal components are locally organized according to the orientation parameters (position angle, PA, and rotation angle, RA) in panels a4 and a5, while the representations learned through the contrastive algorithm remove these orientation dependencies efficiently (panels b4 and b5). We notice, however, that the angular coverage dependencies were not completely removed (panel b6): an over-density of the MaNGA secondary sample is visible in the representation space. Panel b7 also shows a dependence of the representations with redshift, this could be due to the fact that the MaNGA survey has an over-

population of massive galaxies toward higher z (Wake et al. 2017). When not considering a volume-limited sample, using z to quantify non-physical dependencies could lead to the conclusion that representations depend on spatial resolution, whereas the features learnt actually reflect physical properties. Therefore, even if redshift indicates a physical resolution dependency, it is not a reliable indicator for this particular data set.

The rightmost column of Figure 2 shows that the SimCLR representation space smoothly transitions as a function of physical properties: age, metallicity and central velocity dispersion (panels b11, b12, b10). Furthermore, parameters that are less trivial to obtain from the maps, such as λ_{R_e} (b9) and ∇age_{R_e} (b13), are also learned by the model and show smooth transitions in the SimCLR representations. In general, the distribution of the principal components also shows a clear dependence on these parameters, but in this space galaxies that share the same properties tend to appear in separate areas of the diagram due to instrumental biases.

These results confirm that the SimCLR representations efficiently remove the complex instrumental and observational biases present in the MaNGA survey data and primarily organize galaxies based on physical parameters.

4.2. Unsupervised Clustering of kinematics and stellar populations of galaxies

Based on the encouraging results of the previous subsection, we further investigate how the sample of galaxies is clustered in the representation space by dividing the sample into groups using the K -means clustering algorithm (Pedregosa et al. 2011, MacQueen 1967) on the 1024 dimensions of the representation space. K -means is a standard iterative algorithm to cluster unlabelled data by minimizing the sum of squared euclidean distances of the data-points within each cluster to their average position, or centroid. The number of clusters provided by the user will determine the number of centroids that will be initialized. We assume the same distance metric in the representation space as the one used to calculate the contrastive loss in the head projection space and therefore, normalize the representations with L^2 norm such that the euclidean and cosine (as $[1 - \langle \mathbf{a}, \mathbf{b} \rangle]$, where \mathbf{a}, \mathbf{b} are two array-like representations) distances between representations are equivalent.

The resulting clusters represent a first-order division of representation space and will allow a further understanding of the underlying astrophysical properties that drive the space distribution by studying the galaxy groups. Since our representation space is continuous in nature, this exercise will not lead to a sharply defined, discrete, classification. Nevertheless, we define a criterion to obtain a repeatable and physically driven classification, allowing a further study of the galaxy groups thus identified. To define the number of

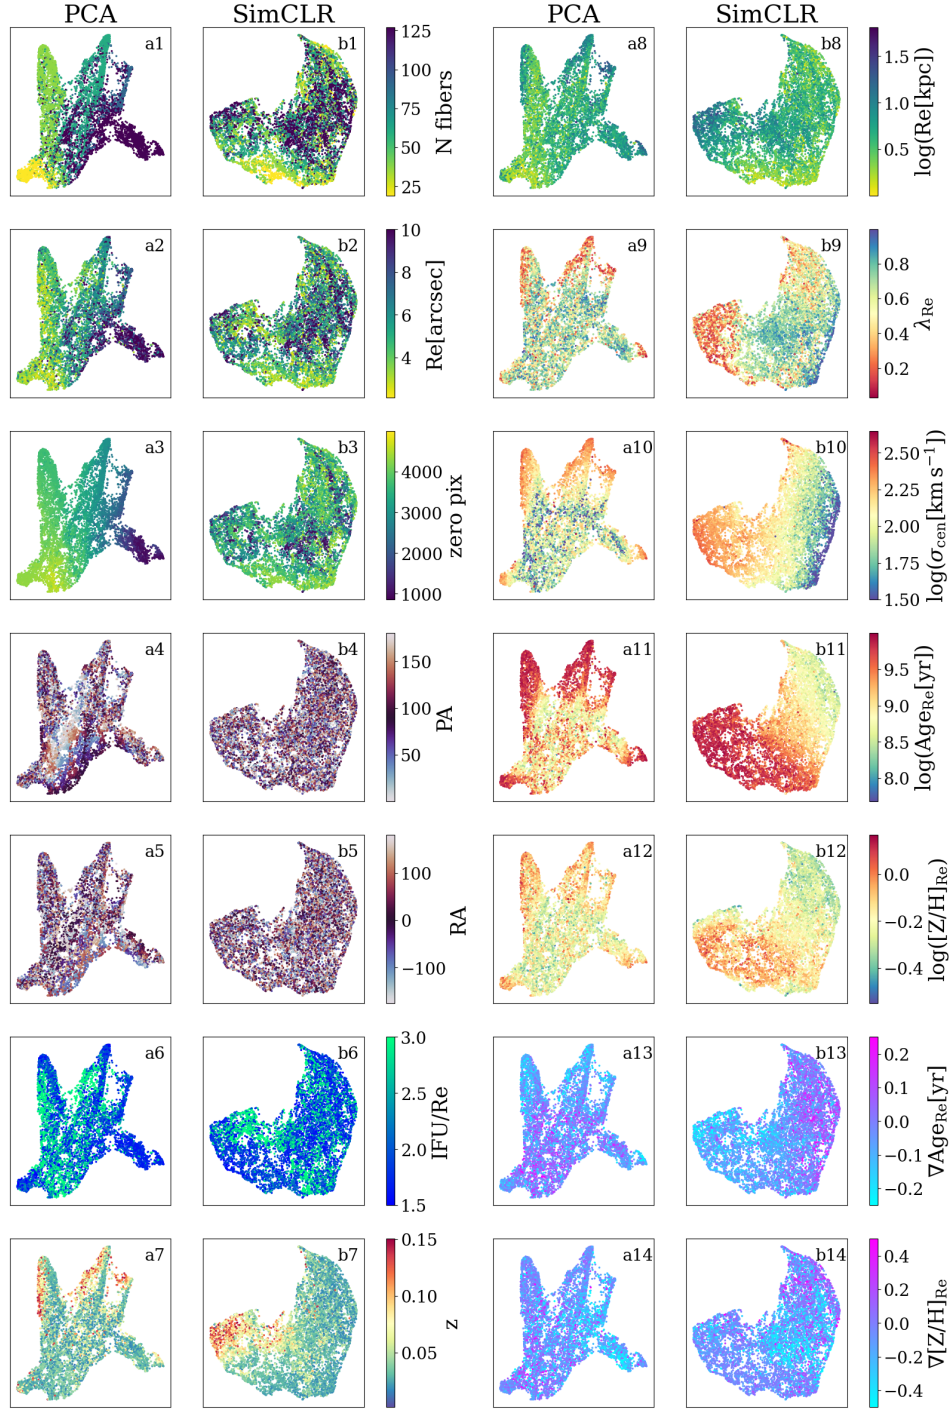


Figure 2. 2D UMAP visualization of PCA decomposition (columns 1 and 3) and SimCLR representations (columns 2 and 4) color-coded according to different labels. The two left columns show dependencies with non-physical labels, from top to bottom: number of fibers in the IFU, galaxy angular size, number of zero pixels, position and rotation angle of the galaxy, angular coverage and redshift. While the two columns on the right are color-coded according to integrated parameters derived from the input maps, from top to bottom: effective radius, specific angular momentum, central velocity dispersion, luminosity weighted age and metallicity and their gradients. Principal components are dominantly distributed in the 2D projection according to the number of fibers in their IFUs. The SimCLR representations show poorer correlations with undesired dependencies, while presenting smooth transitions with physical parameters.

clusters we require the division to be robust, i.e. it should be independent from the training run and initialization parameters. Additionally, the clusters must be independent of features that are not physically relevant.

To ensure a robust division, we train the SimCLR framework ten independent times such that the input maps are perturbed differently in each run. By construction, the features obtained in each run are not exactly the same since the training has some inherent stochastic behavior. Nevertheless, our assumption is that the representation spaces must encode the same information irrespectively of the initialization, such that neighboring galaxies in one run should remain neighbors in the other resulting spaces. We set the first run as the reference run and maximize matching clusters of the remaining runs with the linear assignment algorithm implemented in SciPy (Virtanen et al. 2020, Crouse 2016). As the clusters of each run will not necessarily be identified under the same labels by the clustering algorithm, the linear assignment algorithm provides a new set of cluster labels such that the number of galaxies that fall in the same groups as the reference run is maximized. An agreement score is defined as the fraction of objects that are consistently placed in the same group in at least 90% of the runs. Although this score might appear too restrictive when a large number of clusters is considered, it shows how reproducible the resulting groups are. We consider divisions to have good agreement when the agreement score is over 85 %.

Finally, to account for undesired dependencies, e.g. on instrumental parameters, we compute the linear contribution of the non-physical parameters to the classifications found. We perform a linear regression fit of the parameters to a binary classification for each cluster. One class contains galaxies that belong to the given cluster while the other class is a random selection of non-member galaxies. The fit is presented with the same number of examples of both classes while a spare 20% of the sample is used to obtain the validation accuracy. The resulting coefficient associated to each parameter is then normalized to indicate its contribution percentage to the classification. We reject clusters when the accuracy of this fit is higher than 65%, in which case we assume that the cluster is the result of instrumental and/or non-physical effects.

For comparison purposes, we also perform a K -means clustering on a representation space derived through PCA on the five MaNGA maps. We note that the metrics chosen are compatible with the SimCLR framework, while this might not be the case for PCA. Therefore, a more favorable additional test is performed for the PCA decomposition. In this case, only the first 10 PCA parameters are selected and each parameter is scaled to have a 0 mean and standard deviation of 1.

The agreement score and the most accurate fit to the parameters that are not physically relevant are shown in Ta-

Table 1. Cluster divisions performance.

SimCLR			
N clusters	% Agree	Acc _{max} Non-phys	Acc _{mean} Phys
2	89.419 ± 0.006	64 ± 2	93.1 ± 0.5
3	93.63 ± 0.03	60.5 ± 0.5	89.3 ± 0.5
4	62.5 ± 0.1	63 ± 2	86.4 ± 0.3
5	84.43 ± 0.08	68 ± 1	85.5 ± 0.6
6	85.04 ± 0.09	68 ± 1	84.1 ± 0.5
7	60 ± 1	70 ± 1	83.1 ± 0.5
8	58 ± 5	72 ± 1	82.4 ± 0.3
9	56 ± 2	72.9 ± 0.7	82.6 ± 0.5
10	58 ± 2	71 ± 1	81.9 ± 0.5
11	65 ± 3	71 ± 2	81.5 ± 0.9
12	47 ± 2	71 ± 2	80.9 ± 0.7
13	50 ± 3	72 ± 3	80.1 ± 0.7
14	51 ± 5	73 ± 3	79.1 ± 0.8
15	64 ± 8	75 ± 4	78.7 ± 0.9
PCA 1024			
N clusters		Acc _{max} Non-phys	Acc _{mean} Phys
2		92.13 ± 0.09	72.28 ± 0.07
3		94.9 ± 0.8	71.3 ± 0.2
4		92.7 ± 0.7	64.7 ± 0.2
5		91.9 ± 0.3	69.5 ± 0.5
PCA 10			
N clusters		Acc _{max} Non-phys	Acc _{mean} Phys
2		95 ± 3	72 ± 2
3		93 ± 2	67.4 ± 0.6
4		92 ± 1	68 ± 1
5		94.1 ± 0.5	64.2 ± 0.5

NOTE—For each division into N clusters an agreement score is calculated from ten different runs of the contrastive algorithm. Third column indicates the most accurate fit to non-physical parameters among the N clusters, while the last column shows the average accuracy of the fit to integrated astrophysical parameters. Uncertainty for agreement score is estimated as the standard deviation due to K -means initialization with ten different seeds, while accuracy uncertainties were estimated as the standard deviation corresponding to five SimCLR independent runs with five different seeds each in the clustering step.

ble 1 for 14 divisions of the SimCLR representations and four divisions of the 1024 and 10 most relevant PCA features. For comparison, the relevance of physical and non-physical parameters in each cluster—computed through linear regression—are also listed in the last two columns.

In all cases, the SimCLR representations yield group divisions that are more dependent on physical features than on those that are not physically relevant, thus confirming the

qualitative results of the previous section. In contrast, the divisions derived from the PCA are highly dependent on non-physical features ($> 90\%$ correlation). In particular, the number of fibers in the IFU and the number of zeros in the maps show the strongest correlation with the group divisions in the PCA features.

The maximum number of clusters that fulfills our agreement and non-physical accuracy requirements is three. This number also corresponds with the natural divisions of the UMAP maps of Fig. 2 which clearly show a large cluster in the bottom left corner and a strong age trend in the right blob.

In the following we only consider these three clusters and analyze their physical properties. Table 2 shows the correlation of the clusters with several integrated physical parameters. The strongest correlations are with velocity dispersion, age, physical size and λ_{R_c} . For a better visualization, we show in Fig. 3 four well studied planes for galaxies: $\log M_* - \log \text{SFR}$, $\epsilon - \lambda_R$, $\log M_* - [Z/H]$ and $\log \sigma - \log \text{age}$.

We clearly observe that the three different clusters identified in a purely unsupervised way populate different regions in the different planes. While star-forming discs (with $\langle \log \text{age}[\text{yr}] \rangle = 9.9$ and $\langle \epsilon \rangle = 0.36$) with low metallicity ($\langle [Z/H] \rangle = -0.25$) group in cluster 2, quenched fast-rotating galaxies with low to intermediate masses ($\log M_*/M_\odot \lesssim 10.9$) can be found in cluster 1. Lastly, cluster 3 includes the massive ($\log M_*/M_\odot \gtrsim 10.4$) slow-rotating galaxies. We emphasize that the star formation rate is not included in the maps used for the representation. Nevertheless, most of main sequence galaxies are associated to one unique cluster which naturally suggests that star forming galaxies share similar kinematic and stellar population properties over several decades of stellar mass. We will discuss this further in Section 6. For comparison, the bottom row of Fig. 3 shows the same four physical planes obtained through clustering on the PCA space. Interestingly, we see that there is very little dependence on physical properties which again highlights the effectiveness in extracting physical information of the SimCLR representations when compared to those of the PCA.

Additionally, we analyze the morphological features of the galaxies in each cluster using the MaNGA Deep Learning DR-17 morphological catalogue (MDLM-VAC; private communication, Domínguez Sánchez et al. in prep). This catalogue is an extension of the Deep learning DR-15 catalogue presented in Fischer et al. (2019) and it was obtained with deep learning models trained on SDSS-DR7 images (see Domínguez Sánchez et al. (2018) for more details on the methodology and models performance). Figure 4 shows how the galaxies distribute in each cluster according to the morphological parameter T-Type.

We find that galaxies that have been classified as early types (T-Type < 0) belong predominantly to the cluster con-

taining high-mass slow-rotating galaxies, while a smaller fraction belongs to the intermediate-mass quenched galaxy cluster. Galaxies with T-Types > 0 distribute in the clusters that include less massive galaxies, with the latest T-Types belonging to the star-forming cluster.

5. INTERNAL STRUCTURE

In the previous analysis the relative values of the observed maps were kept, since the normalization applied works as a uniform change of units across all the galaxy maps. We now analyze how the representation space organizes when only the spatial features of the input maps are considered. For this purpose, a linear normalization is applied for each map individually, such that the values of each map span the $[0, 1]$ range (individual norm). With this normalization, information about the absolute values of the different quantities is removed and the model is forced to focus exclusively on the internal structure of the observed galaxies. The training process is repeated on the newly normalized input data using the same augmentations and model settings. We present the analysis of the representation space performed as in Section 4.

5.1. Two-dimensional visualization of representations

As in Section 4.1, the SimCLR representations obtained after training the model are projected to a 2D space with UMAP (McInnes et al. 2020) to visualize how the representation space correlates with known physical parameters. For comparison, the projections of both sets of normalized maps are presented together in Figure 5.

While the relative norm representations show smoother transitions for central velocity dispersion, projected angular momentum, age and metallicity (panels R2, R3, R4 and R5), the correlation between these parameters and the individual norm representation space is not negligible. Furthermore, panels I3, I4 and I5 show that young and metal poor galaxies with low velocity dispersion tend to populate the left side of the maps, while old metal rich galaxies with higher velocity dispersion primarily locate on the right-hand side.

5.2. Unsupervised Clustering of the representation space

As the representation space distribution changes when applying different norms to the input data, we analyse the dominant drivers of the new space by performing an unsupervised clustering.

In this representation space, we find that only a division into two clusters fulfills the agreement requirement defined in Section 4.2 with an agreement score of 90.89 ± 0.03 , while more partitions yield scores $< 78\%$. The non-physical parameter fits have comparable accuracy to the ones obtained with the relative norm. Specifically for the two cluster division, this value is 67.6 ± 1.3 , which is only slightly larger than the 65% threshold imposed in the previous section. To avoid

Table 2. Cluster dependence with integrated parameters.

Cluster	$\log(\sigma)$	LW Age $_{R_c}$	$\log(R_c[\text{kpc}])$	λ_{R_c}	$\log(\text{SFR}_{H\alpha})$	LW ZH $_{R_c}$	$\nabla\text{LW ZH}_{R_c}$	$\nabla\text{LW Age}_{R_c}$
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1	$13.4 \pm 3.7\%$	$35.2 \pm 1.6\%$	$14.6 \pm 1.4\%$	$17.9 \pm 1.9\%$	$6.8 \pm 0.6\%$	$4.5 \pm 0.8\%$	$6.4 \pm 0.3\%$	$1.1 \pm 0.3\%$
2	$3.7 \pm 2.2\%$	$61.5 \pm 3.1\%$	$8.9 \pm 0.4\%$	$4.5 \pm 1.1\%$	$10.9 \pm 1.2\%$	$7.4 \pm 1.4\%$	$1.9 \pm 0.8\%$	$0.9 \pm 0.5\%$
3	$43.6 \pm 1.6\%$	$14.2 \pm 1.6\%$	$11.2 \pm 1.2\%$	$16.7 \pm 1.5\%$	$1.3 \pm 0.5\%$	$4.9 \pm 0.7\%$	$3.9 \pm 0.4\%$	$4.0 \pm 0.3\%$

NOTE—Dependence with integrated parameters of the three groups found with the K -means clustering in the SimCLR representation space.

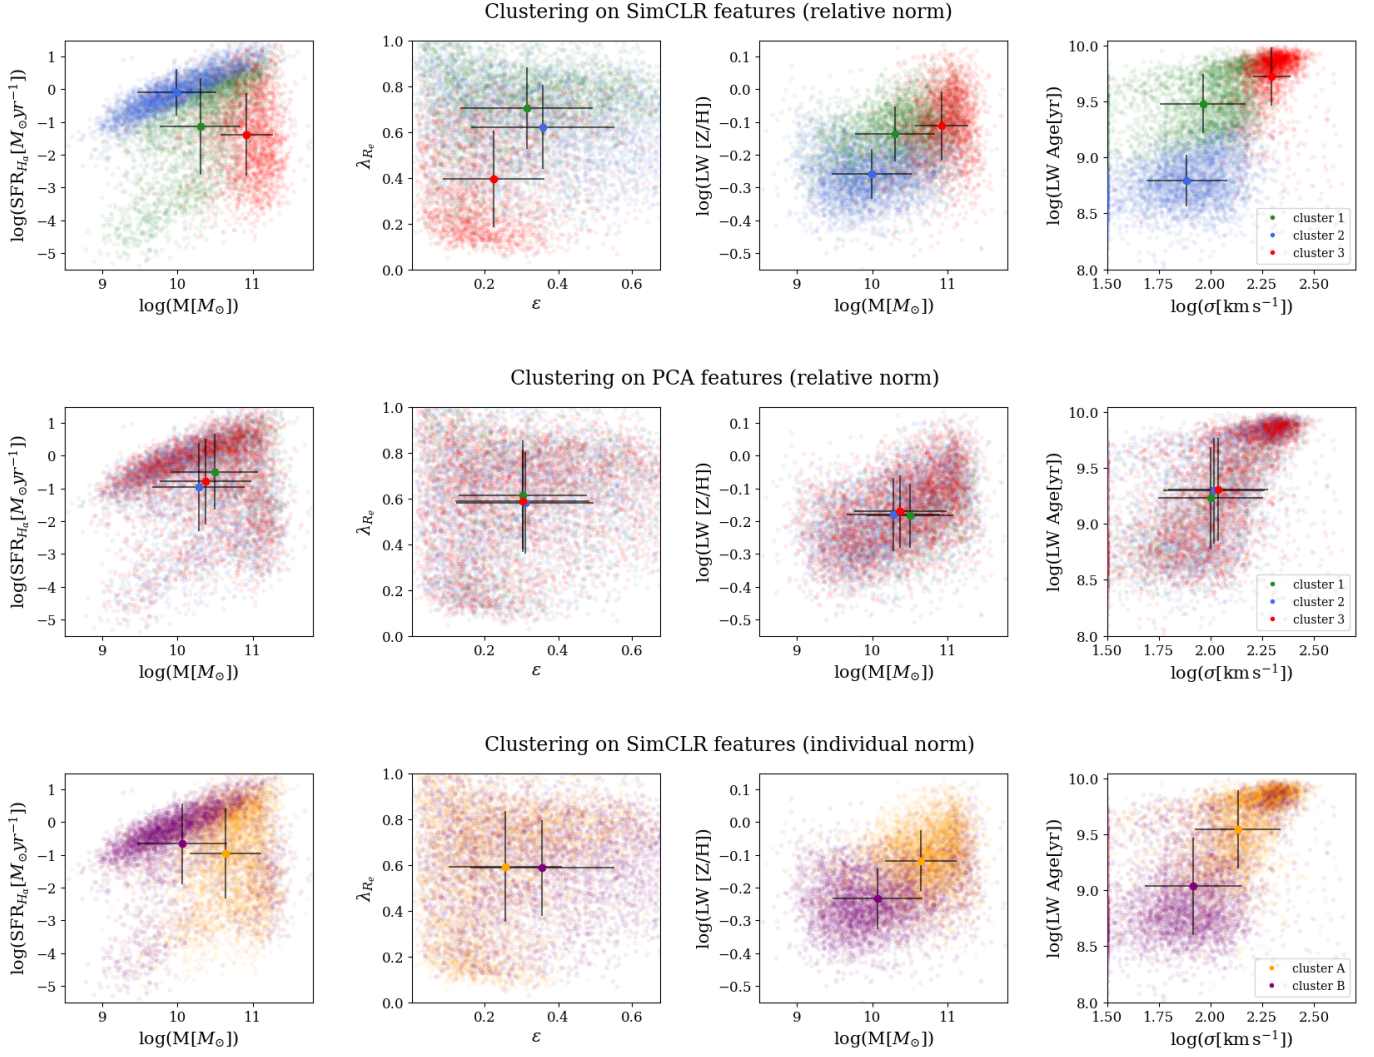


Figure 3. From left to right: stellar mass vs. SFR plane, ellipticity vs. specific angular momentum, stellar mass vs. LW metallicity, central velocity dispersion vs. LW age. Top row panels show integrated parameters color-coded according to the 3 clusters found in the SimCLR representation space when the relative norm (see Sec. 2) is applied to the raw data. In the second row panels, the cluster division was performed on the PCA-1024 feature space. The bottom row shows the division into two clusters of the SimCLR representation space when the individual norm is applied to the input data. Mean values of each cluster are displayed in solid colors, errorbars account for the standard deviation of the cluster. Groups found with SimCLR representations show less overlap in all four physical planes than groups found by clustering PCA features.




Figure 4. T-Type classification of all MaNGA galaxies (grey) and how they distribute across the unsupervised clusters (red, green and blue). The colors of the clusters have been chosen as the first row panels in Figure 3. Early type galaxies (T-Type < 0) are predominantly in the cluster that includes the massive slow-rotating galaxies. Galaxies belonging to the star forming main sequence cluster (blue) mostly distribute around the largest values of T-Type, while intermediate mass quenched galaxies show intermediate T-Type values.

confusion with the clusters of the previous section, we refer to the new groups as clusters *A* and *B*.

The distribution of the two clusters in the physical planes $\log M_* - \log \text{SFR}$, $\epsilon - \lambda_R$, $\log M_* - [Z/H]$ and $\log \sigma - \log \text{age}$ is shown on the bottom row of Figure 3. While cluster *A* is constrained in the planes towards higher mass, older and more metal rich regions (with average values $\langle \log M_*/M_\odot \rangle = 10.6$, $\langle \log \text{age}[\text{yr}] \rangle = 9.5$ and $\langle [Z/H] \rangle = -0.11$), cluster *B* spreads across all the planes with lower concentration in the overlapping regions (with average values $\langle \log M_*/M_\odot \rangle = 10$, $\langle \log \text{age}[\text{yr}] \rangle = 9$ and $\langle [Z/H] \rangle = -0.23$). The low-mass and star forming galaxies are contained primarily in the last cluster. Although there is overlap between the two groups, a physical difference is captured and is clearly visible in age, mass and metallicity. This is interesting because the information about the absolute values has been removed, however, we still see that the main clusters found correlate with integrated properties such as stellar mass and velocity dispersion.

To further visualize the typical spatial features found in each group, we select the five galaxies whose representations are closer to the centroid of each cluster. The example galaxies of each cluster are shown in Figure 6. While the example galaxies from cluster *A* show an early type morphology and negative gradients in velocity dispersion and metallicity maps, galaxies in cluster *B* tend to have inverted gradients to those in cluster *A* and their V-band reconstructed images

show structure and disk features compatible with later type morphology.

As this division roughly separates early type galaxies from late types, the morphological features encoded in the V-band reconstructed images could be driving the clustering. Therefore, in Appendix A we evaluate the influence of the kinematic and composition maps in the division. When the V-band reconstructed images are excluded from the input maps, we find that the division can be recovered with good agreement ($88.621 \pm 0.007\%$ overlap with the previous clusters) indicating that the information of the four remaining maps is enough to find the two groups.

6. SUMMARY AND DISCUSSION

6.1. Unsupervised exploration of multi-dimensional datasets

A key property of contrastive learning-based dimensional reduction, in contrast with other techniques, is that representations are encouraged to become invariant to a set of transformations. We have shown that this feature has a lot of potential when the dataset is non-homogeneous and affected by known biases, since these can be tackled with convenient transformations in the contrastive framework.

We confirm that the representations obtained in this work efficiently reduce their impact and focus on physical properties, even though MaNGA maps include important instrumental effects. This allows one to explore the physical correlations existing in the high dimensional space.

Therefore, the presented framework might constitute an important tool to explore and discover patterns in future datasets. We have illustrated in this work how the representations organize based on five kinematics and stellar population maps. However, the method has no limit on the number of channels in the input data, thus it is possible to extend the analysis to a larger number of maps including for example gas properties or even complete IFU data cubes without post-processing.

Additionally, because transformations included in the training procedure can be tuned, one could eventually combine datasets of different origins and selection effects. A natural extension of this work is thus the combination of IFU data from different surveys (i.e. MUSE, Henault et al. 2003; CALIFA, Sánchez et al. 2012; MaNGA, Bundy et al. 2015; SAMI, Croom et al. 2012) represented in a unique low-dimension space. Combining observations and predictions from recent hydrodynamic cosmological simulations is also a promising avenue to identify possible discrepancies between the physical properties of observed and simulated galaxies.

In addition to visualization, the representations learnt can be used in a variety of downstream tasks. Chen et al. (2020b) and Abul Hayat et al. (2020) found that this framework can obtain the same levels of accuracy for classification than a su-

Normalization effect in SimCLR representations

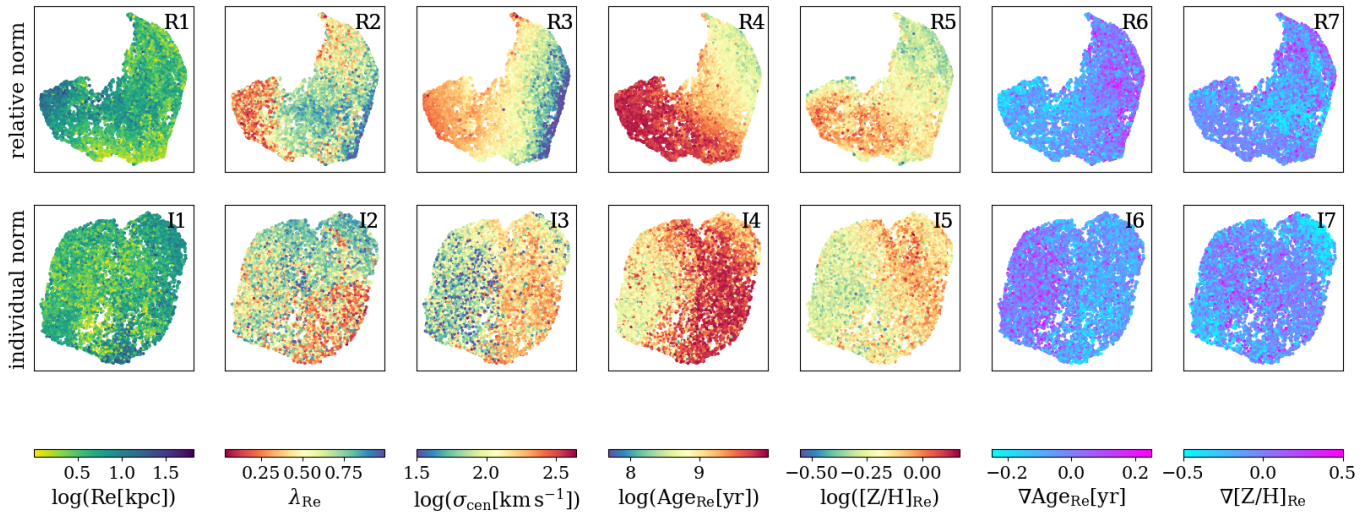


Figure 5. As Fig. 2, 2D UMAP projection of SimCLR representations using relative norm (top row) and individual norm (bottom row). Color coding corresponds to physical properties derived from the maps, from left to right: effective radius, specific angular momentum, central velocity dispersion, luminosity weighted age and metallicity and their gradients. Although the individual norm removes the relative information of these parameters from the maps, the spatial features still correlate with them.

pervised approach but reducing the number of training labels. We presented in this work one application that makes use of the representation space as a tool to explore the dataset.

6.2. Clustering

Since contrastive learning locates objects with similar physical properties in nearby regions of the representation space, it can be used to identify groups of objects with shared physical properties in a purely data driven manner. In Sections 3 and 5 we illustrated this procedure with two different normalization of the MaNGA maps: with and without information about the absolute values of the physical properties. This allows us to study the interplay between integrated properties of galaxies and their resolved internal structure.

6.2.1. Absolute physical properties

When applying the relative norm to the input data, the three groups identified naturally divide galaxies into three previously well-defined categories using integrated properties: main sequence rotating discs, low mass quiescent rotating galaxies and massive quiescent slow-rotating galaxies. This highlights the potential of our proposed framework for future discoveries.

Although no direct star formation rate indicators were included in the input, the clustering algorithm naturally puts most of the main-sequence galaxies in the same group, i.e. the main sequence is *rediscovered* in a data-driven way. It suggests that galaxies in the main sequence share kinematic and stellar populations properties and therefore evolve with similar physical processes despite covering several decades in stellar mass.

The well-known slow-rotating galaxies are primarily included in this group, which have been put forward by several previous IFU surveys (Emsellem et al. 2007, Emsellem et al. 2011, Graham et al. 2018, Falc3n-Barroso et al. 2019). The so-called slow rotators are thought to have an assembly history more dominated by major mergers (Cappellari 2016) which therefore increases their stellar mass and velocity dispersion. Our framework naturally isolates these objects.

There is a third interesting group highlighted in Fig. 3 which is composed of low-mass quiescent galaxies. Interestingly, these galaxies are rotating fast as opposed to the slow rotators of the previous cluster. It suggests that the quenching mechanisms for this population is different. This could be predominantly a population of satellite galaxies quenched through strangulation processes (Peng et al. 2015) without significant impact on the kinematics of stars.

While the three clusters found are well-known groups from previous morpho-kinematic works, our methodology found this classification with additional stellar population information: the luminosity weighted age and metallicity maps. In Appendix B we analyze if the clusters can be recovered by the framework when only the V-band image and kinematic maps are considered. While the clusters are highly dependent of these input maps, without the age and metallicity maps the classification cannot be fully recovered ($\sim 25\%$ of the galaxies are misplaced in the new division). Furthermore, not only an expected smaller dependence with age and metallicity is found in the average cluster properties, but a greater overlap of the clusters in T-Type classification is noticeable. We note, however, that the massive early-type galaxies tend to be more

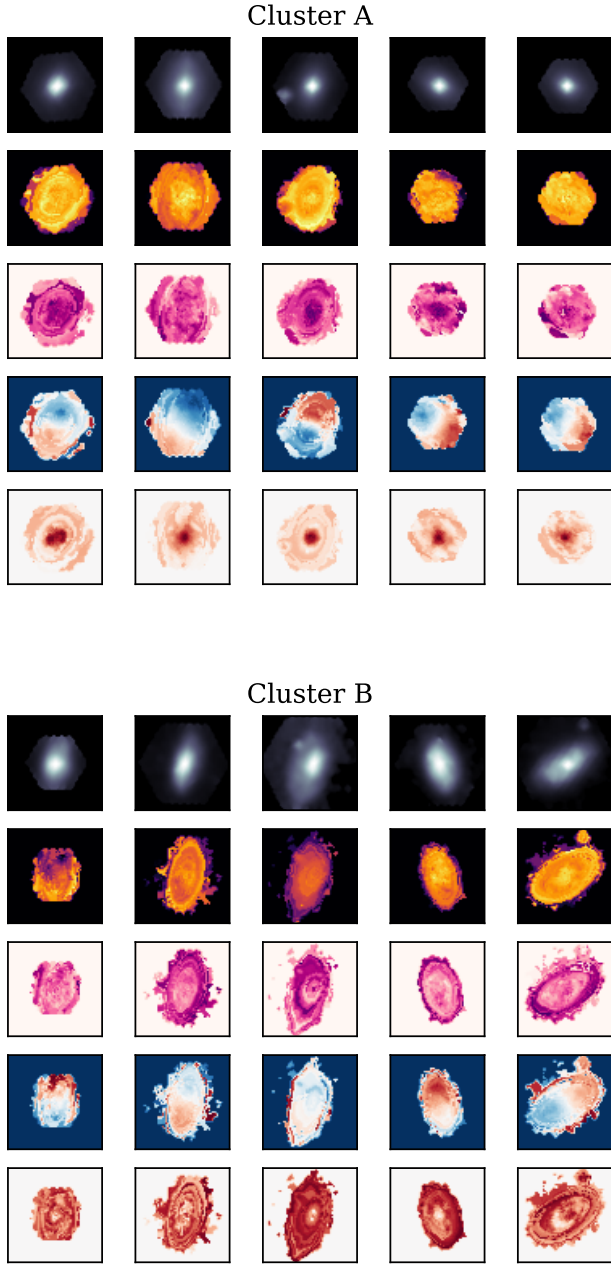


Figure 6. Example galaxies from clusters *A* and *B* (top and bottom panels, respectively). The V-band reconstructed image, age, metallicity, radial velocity and dispersion maps (rows) were normalized with the individual norm for each galaxy (columns).

consistently classified in the same group, again highlighting that the distinctive kinematic features of these galaxies isolate them from the rest. Interestingly, removing the V-band reconstructed image from the input data and including the four remaining maps has no significant effect on the T-Type classification per cluster (see Appendix A).

We conclude that the divisions that naturally arise from this framework confirm extensively studied groups of galaxies.

The clusters found reflect the dominant differences between three regions of the representation space. The clusters share main properties with well-known galaxy categories, confirming that these groups of galaxies present intrinsic differences. This again highlights the potential of our proposed framework to discover the physical mechanisms that shape galaxies.

The results might, however, still be affected by selection effects. The MaNGA sample is not a volume-limited sample and this is not taken into account in the representations. Although a set of weights to obtain a statistically representative sample of a single volume are provided in Wake et al. (2017), limiting the sample would mean a significant reduction of the number of galaxies available to train the Deep Learning model. This may conflict with the method as Deep Learning models require a large training set to avoid an overfit of the network’s parameters. Although the different groups do exist, their relative fractions or relevance might not be representative of the local Universe. Also small groups of galaxies might not be detected.

6.2.2. Internal structure

We have repeated the cluster analysis with a different normalization that removes information about the absolute values of the physical properties. This allows us to quantify the relation between integrated and resolved internal structure of nearby galaxies. Although the previous clusters cannot be completely recovered with this new normalization, it is possible to find two clusters with good agreement that are separated primarily by the spatial features present in the maps. These clusters correlate with the mean and integrated physical values derived from the maps, suggesting that the internal structure of galaxies serve as a footprint that is tightly related to the observed integrated properties. It suggests that these two galaxy types are different not only because they differ in the absolute values, but also because their internal structure is different, pointing towards distinct assembly histories, as highlighted in previous works.

The first cluster found when only spatial features are considered contains intermediate to high-mass quenched galaxies. These galaxies show in general early type morphologies and present negative slopes in the velocity dispersion and metallicity maps (Figure 6). The second cluster primarily includes the star forming main sequence of galaxies. Figure 6 shows that galaxies in this cluster tend to present flat metallicity and velocity dispersion distributions and more disk like stellar structure.

The comparison between the clusters obtained with the different normalizations reveals interesting trends. Figure 7 shows that $\sim 80\%$ and $\sim 70\%$ of the galaxies that were in different clusters when including the information about the absolute values are also separated by internal structure. It




Figure 7. Overlap of clusters found in using relative (clusters 1,2 and 3 - vertical) and individual norms (clusters A and B - horizontal). Percentages are calculated over the relative norm clusters.

confirms that the distinction between high velocity and low velocity dispersion systems reveals different assembly histories and it is not only a reflection of a stellar mass bimodality. Interestingly, the third cluster found in the relative normalization run is distributed among the two main ones with $\sim 35\%$ being closer to the star-forming rotating disks and $\sim 65\%$ being more similar to the massive slow rotating systems. It suggests that this intermediate mass quenched population is formed through a mixture of smooth and more violent assembly processes but there is no a particular signature in the internal kinematics of the galaxies that clearly distinguishes them from the two other groups. A fraction of the objects also change cluster when the normalization is modified. We will study these systems in more detail in future work.

7. CONCLUSIONS

We applied a method based on the SimCLR deep contrastive learning approach for obtaining meaningful representations of galaxy resolved maps from the MaNGA survey. The method produces a low-dimension representation of the data input while preserving meaningful physical information and reducing the impact of the instrumental effects present in the data.

We showed that the proposed representation learning scheme outperforms, at least for this particular dataset, more standard dimensionality reduction algorithms such as PCA which is more sensitive to non-physical correlations.

We have shown that the representation space can be efficiently used for visualization, clustering and exploring the properties of sub-categories of objects. We argue that contrastive learning techniques represent a promising avenue to explore and discover patterns in future high-dimensional astronomical datasets.

ACKNOWLEDGMENTS

We thank Sara Ellison, Mallory Thorp and David Patton for comments on a previous version of this manuscript. We acknowledge financial support from the European Union’s Horizon 2020 research and innovation programme under Marie Skłodowska-Curie grant agreement No 721463 to the SUNDIAL ITN network, from the State Research Agency (AEI-MCINN) of the Spanish Ministry of Science and Innovation under the grant “The structure and evolution of galaxies and their central regions” with reference PID2019-105602GB-I00/10.13039/501100011033, and from IAC project P/300724, financed by the Ministry of Science and Innovation, through the State Budget and by the Canary Islands Department of Economy, Knowledge and Employment, through the Regional Budget of the Autonomous Community. J. F-B acknowledges support through the RAVET project by the grant PID2019-107427GB-C32 from the Spanish Ministry of Science, Innovation and Universities (MCIU), and through the IAC project TRACES which is partially supported through the state budget and the regional budget of the Consejería de Economía, Industria, Comercio y Conocimiento of the Canary Islands Autonomous Community.

This project makes use of the MaNGA-Pipe3D dataproducs. We thank the IA-UNAM MaNGA team for creating this catalogue, and the Conacyt Project CB-285080 for supporting them.

Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High Performance Computing at the University of Utah. The SDSS website is www.sdss.org. SDSS-IV is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, Center for Astrophysics — Harvard & Smithsonian, the Chilean Participation Group, the French Participation Group, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU) / University of Tokyo, the Korean Participation Group, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatório Nacional / MCTI, The Ohio State University, Pennsylvania State Uni-

versity, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado

Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University.

REFERENCES

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <http://tensorflow.org/>
- Abul Hayat, M., Stein, G., Harrington, P., Lukić, Z., & Mustafa, M. 2020, arXiv e-prints, arXiv:2012.13083. <https://arxiv.org/abs/2012.13083>
- Aguado, D. S., Ahumada, R., Almeida, A., et al. 2019, ApJS, 240, 23, doi: [10.3847/1538-4365/aaf651](https://doi.org/10.3847/1538-4365/aaf651)
- Baron, D., & Ménard, B. 2020, arXiv e-prints, arXiv:2006.13948. <https://arxiv.org/abs/2006.13948>
- Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, AJ, 154, 28, doi: [10.3847/1538-3881/aa7567](https://doi.org/10.3847/1538-3881/aa7567)
- Bundy, K., Bershady, M. A., Law, D. R., et al. 2015, ApJ, 798, 7, doi: [10.1088/0004-637X/798/1/7](https://doi.org/10.1088/0004-637X/798/1/7)
- Cappellari, M. 2016, Annual Review of Astronomy and Astrophysics, 54, 597, doi: [10.1146/annurev-astro-082214-122432](https://doi.org/10.1146/annurev-astro-082214-122432)
- Cappellari, M., & Copin, Y. 2003, MNRAS, 342, 345, doi: [10.1046/j.1365-8711.2003.06541.x](https://doi.org/10.1046/j.1365-8711.2003.06541.x)
- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. 2020a, A Simple Framework for Contrastive Learning of Visual Representations. <https://arxiv.org/abs/2002.05709>
- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., & Hinton, G. 2020b, Big Self-Supervised Models are Strong Semi-Supervised Learners. <https://arxiv.org/abs/2006.10029>
- Clevert, D.-A., Unterthiner, T., & Hochreiter, S. 2016, Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). <https://arxiv.org/abs/1511.07289>
- Croom, S. M., Lawrence, J. S., Bland-Hawthorn, J., et al. 2012, MNRAS, 421, 872, doi: [10.1111/j.1365-2966.2011.20365.x](https://doi.org/10.1111/j.1365-2966.2011.20365.x)
- Crouse, D. F. 2016, IEEE Transactions on Aerospace and Electronic Systems, 52, 1679, doi: [10.1109/TAES.2016.140952](https://doi.org/10.1109/TAES.2016.140952)
- Domínguez Sánchez, H., Huertas-Company, M., Bernardi, M., Tuccillo, D., & Fischer, J. L. 2018, MNRAS, 476, 3661, doi: [10.1093/mnras/sty338](https://doi.org/10.1093/mnras/sty338)
- Drory, N., MacDonald, N., Bershady, M. A., et al. 2015, AJ, 149, 77, doi: [10.1088/0004-6256/149/2/77](https://doi.org/10.1088/0004-6256/149/2/77)
- Emsellem, E., Cappellari, M., Krajnović, D., et al. 2007, Monthly Notices of the Royal Astronomical Society, 379, 401, doi: [10.1111/j.1365-2966.2007.11752.x](https://doi.org/10.1111/j.1365-2966.2007.11752.x)
- Emsellem, E., Cappellari, M., Krajnović, D., et al. 2011, Monthly Notices of the Royal Astronomical Society, 414, 888
- Falcón-Barroso, J., Sánchez-Blázquez, P., Vazdekis, A., et al. 2011, A&A, 532, A95, doi: [10.1051/0004-6361/201116842](https://doi.org/10.1051/0004-6361/201116842)
- Falcón-Barroso, J., van de Ven, G., Lyubenova, M., et al. 2019, A&A, 632, A59, doi: [10.1051/0004-6361/201936413](https://doi.org/10.1051/0004-6361/201936413)
- Fischer, J. L., Domínguez Sánchez, H., & Bernardi, M. 2019, MNRAS, 483, 2057, doi: [10.1093/mnras/sty3135](https://doi.org/10.1093/mnras/sty3135)
- Goddard, D., Thomas, D., Maraston, C., et al. 2017, MNRAS, 466, 4731, doi: [10.1093/mnras/stw3371](https://doi.org/10.1093/mnras/stw3371)
- Graham, M. T., Cappellari, M., Li, H., et al. 2018, MNRAS, 477, 4711, doi: [10.1093/mnras/sty504](https://doi.org/10.1093/mnras/sty504)
- Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, AJ, 131, 2332, doi: [10.1086/500975](https://doi.org/10.1086/500975)
- Henault, F., Bacon, R., Bonneville, C., et al. 2003, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes, ed. M. Iye & A. F. M. Moorwood, 1096–1107, doi: [10.1117/12.462334](https://doi.org/10.1117/12.462334)
- Hinton, G., Vinyals, O., & Dean, J. 2015, Distilling the Knowledge in a Neural Network. <https://arxiv.org/abs/1503.02531>
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111, doi: [10.3847/1538-4357/ab042c](https://doi.org/10.3847/1538-4357/ab042c)
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv e-prints, arXiv:1110.3193. <https://arxiv.org/abs/1110.3193>
- Law, D. R., Yan, R., Bershady, M. A., et al. 2015, AJ, 150, 19, doi: [10.1088/0004-6256/150/1/19](https://doi.org/10.1088/0004-6256/150/1/19)
- Law, D. R., Cherinka, B., Yan, R., et al. 2016, AJ, 152, 83, doi: [10.3847/0004-6256/152/4/83](https://doi.org/10.3847/0004-6256/152/4/83)
- MacQueen, J. 1967, in Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics (Berkeley, Calif.: University of California Press), 281–297. <https://projecteuclid.org/euclid.bsm/1200512992>
- McInnes, L., Healy, J., & Melville, J. 2020, UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. <https://arxiv.org/abs/1802.03426>
- Neumann, J., Thomas, D., Maraston, C., et al. 2021, in Extragalactic Spectroscopic Surveys: Past, 21, doi: [10.5281/zenodo.4721540](https://doi.org/10.5281/zenodo.4721540)
- Pearson, K. 1901, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 2, 559, doi: [10.1080/14786440109462720](https://doi.org/10.1080/14786440109462720)
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825

- Peng, Y., Maiolino, R., & Cochrane, R. 2015, *Nature*, 521, 192, doi: [10.1038/nature14439](https://doi.org/10.1038/nature14439)
- Portillo, S. K. N., Parejko, J. K., Vergara, J. R., & Connolly, A. J. 2020, *AJ*, 160, 45, doi: [10.3847/1538-3881/ab9644](https://doi.org/10.3847/1538-3881/ab9644)
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, *Learning Internal Representations by Error Propagation* (Cambridge, MA, USA: MIT Press), 318–362
- Sánchez-Blázquez, P., Peletier, R. F., Jiménez-Vicente, J., et al. 2006, *MNRAS*, 371, 703, doi: [10.1111/j.1365-2966.2006.10699.x](https://doi.org/10.1111/j.1365-2966.2006.10699.x)
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, 146, 32, doi: [10.1088/0004-6256/146/2/32](https://doi.org/10.1088/0004-6256/146/2/32)
- Sohn, K. 2016, in *Advances in Neural Information Processing Systems*, ed. D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett, Vol. 29 (Curran Associates, Inc.), 1857–1865. <https://proceedings.neurips.cc/paper/2016/file/6b180037abbeba991d8b1232f8a8ca9-Paper.pdf>
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv e-prints, arXiv:1503.03757. <https://arxiv.org/abs/1503.03757>
- Sánchez, S. F., Kennicutt, R. C., Gil de Paz, A., et al. 2012, *A&A*, 538, A8, doi: [10.1051/0004-6361/201117353](https://doi.org/10.1051/0004-6361/201117353)
- Sánchez, S. F., Pérez, E., Sánchez-Blázquez, P., et al. 2015, Pipe3D, a pipeline to analyze Integral Field Spectroscopy data: I. New fitting philosophy of FIT3D. <https://arxiv.org/abs/1509.08552>
- . 2016, Pipe3D, a pipeline to analyse integral field spectroscopy data: II. Analysis sequence and CALIFA dataproducs. <https://arxiv.org/abs/1602.01830>
- Taniguchi, Y., Scoville, N., Murayama, T., et al. 2007, *ApJS*, 172, 9, doi: [10.1086/516596](https://doi.org/10.1086/516596)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Wake, D. A., Bundy, K., Diamond-Stanic, A. M., et al. 2017, *The Astronomical Journal*, 154, 86, doi: [10.3847/1538-3881/aa7ecc](https://doi.org/10.3847/1538-3881/aa7ecc)
- Wu, Z., Xiong, Y., Yu, S., & Lin, D. 2018, *Unsupervised Feature Learning via Non-Parametric Instance-level Discrimination*. <https://arxiv.org/abs/1805.01978>
- Yan, R., Bundy, K., Law, D. R., et al. 2016a, *AJ*, 152, 197, doi: [10.3847/0004-6256/152/6/197](https://doi.org/10.3847/0004-6256/152/6/197)
- Yan, R., Tremonti, C., Bershady, M. A., et al. 2016b, *AJ*, 151, 8, doi: [10.3847/0004-6256/151/1/8](https://doi.org/10.3847/0004-6256/151/1/8)

APPENDIX

To analyze the dependence of the representation space with the input maps we repeat the clustering analysis with different combinations of maps. In particular, we consider two variations to the proposed methodology that tackle specific questions. On the one hand, we examine the cluster division when morphology information is excluded. For this purpose, we perform K -means clustering on a representation space obtained by excluding the V-band reconstructed image. On the other hand, we analyze if the three clusters found in Section 4.2 can be recovered solely from the morpho-kinematic input data.

A. AGE, METALLICITY AND KINEMATICS

We train the SimCLR model on the age, metallicity, radial velocity and dispersion maps. We later use this trained model to extract the representations of this set of maps and perform K -means on them. When the input data is normalized with the relative norm the clusters overlap a $92.04 \pm 0.03\%$ (error estimated as the standard deviation of ten different K -means initialisation). This value is slightly smaller than the agreement score for the clustering on the representation space obtained when the V-band reconstructed image is included, and therefore this channel does contribute to the 3 cluster division. We find that the clusters show more discrepancy in the later-type galaxies (T-Type > 0 , as seen the left panel of Figure 8). However, we conclude that the overlap is significant since it is indeed greater than the agreement threshold (85%) defined in Section 4.2.

When this analysis is repeated on the maps normalized with the individual norm, an overlap of $88.621 \pm 0.007\%$ is recovered. After removing the V-band image, the clusters tend to equalize their T-Type distributions in later type galaxies (Fig. 8). However, the agreement of the clusters before and after excluding the V-band image is significant and the overall trends of both clusters are still clear. This indicates that the division does not only depend of the galaxy morphology, but the remaining maps have sufficient information to separate the two groups.

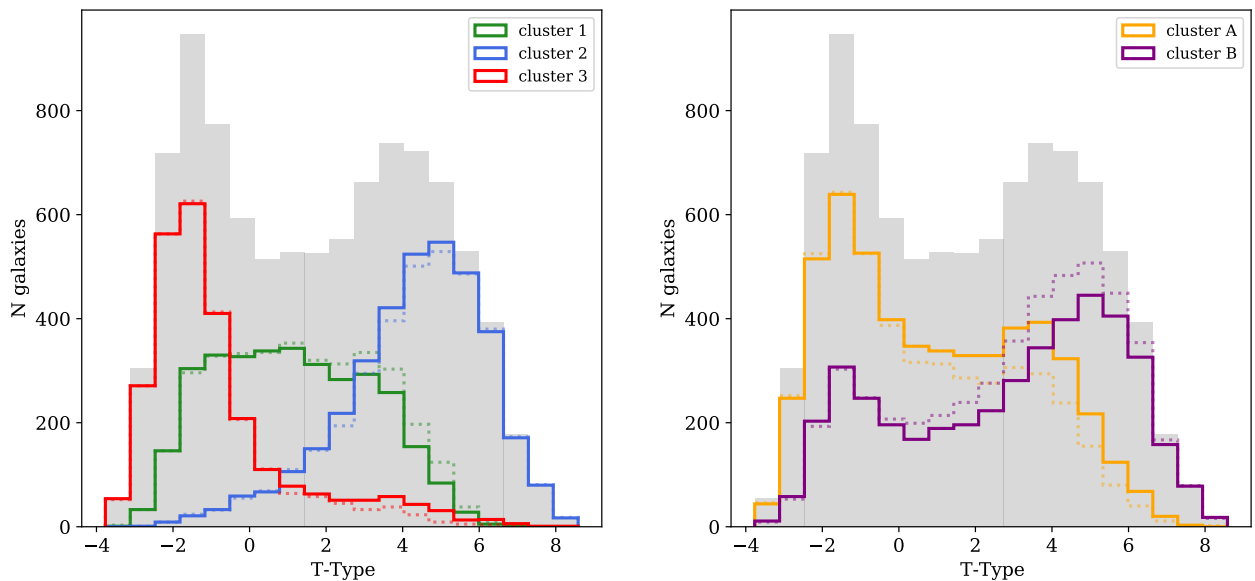


Figure 8. As Figure 4, the grey filled histograms show the T-type classification of all MaNGA galaxies and colored line histograms correspond to unsupervised clusters. Spot line histograms correspond to clusters obtained with the five maps as input: V-band reconstructed image, age, metallicity, radial velocity and velocity dispersion maps. Solid lines show the cluster divisions when the the V-band reconstructed image is excluded. Left panel: colors correspond to the 3 clusters found when the *relative* norm is applied to the input data. Right panel: colors correspond to the 2 clusters found when the *individual* norm is applied. In both panels the clusters' trends are maintained after removing the V-band image from the input. The discrepancy between clusters is clearer toward later-type galaxies.

B. MORPHO-KINEMATICS

While the three clusters found in Section 4.2 are well known groups from previous morpho-kinematic works, our methodology found these classification with additional stellar population information: the luminosity weighted age and metallicity maps. Therefore, we analyze if the clusters can be recovered by the framework when only the morpho-kinematic maps are considered.

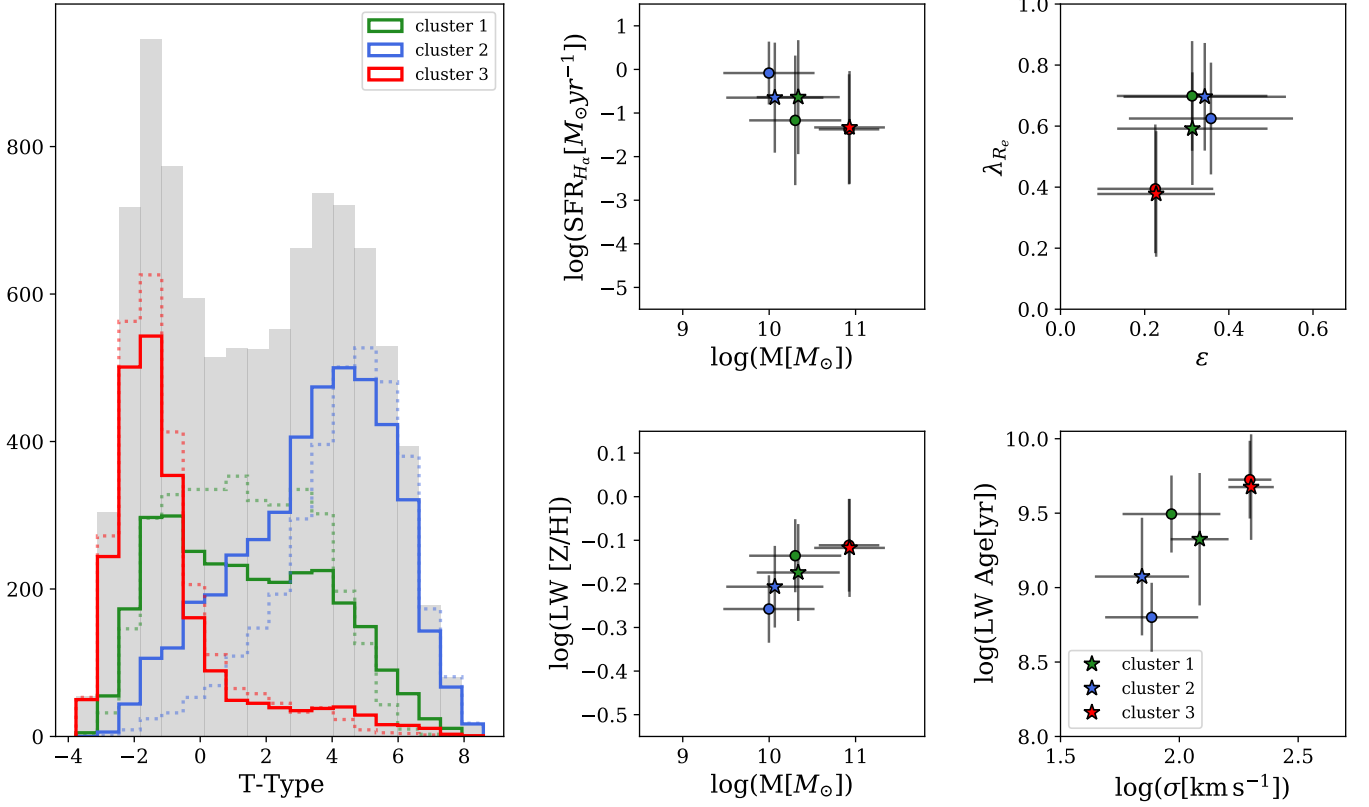


Figure 9. Left panel: T-Type classification. As Fig. 8, spotted line histograms correspond to clusters obtained when five maps are used as input. Solid line histograms . Small-sized panels on the right-hand side show the average integrated properties of the clusters. 5-map clusters are identified with circle markers, while star markers were used for the morpho-kinematic data.

The clusters that arise when the age and metallicity maps are excluded overlap in a $76.95 \pm 0.02\%$ with those clusters found in Section 4.2. While the morpho-kinematic maps greatly influence the clusters found in Section 4.2, we consider that the initial groups cannot be fully recovered because $\sim 25\%$ of the galaxies are misplaced when the age and metallicity maps are excluded.

We further inspect how the average properties of each cluster vary in Figure 9. While cluster 3 is the least affected of all three groups, clusters 1 and 2 show more overlap in T-Type classification and their mean group properties tend to equalize. However, there are parameters that show a larger separations in the physical planes (center and right panels of Fig. 9), like σ_{cen} and λ_{Re} . The fact that the difference between the clusters is enhanced only for these parameters is consistent with these parameters being derived from the maps used in this test while excluding the age and metallicity information. These results are consistent with the correlations analyzed in Table 2, which suggest that the cluster with the most massive galaxies is primarily separated due to their kinematics, while the low to intermediate mass galaxies are separated into old/metal-rich and young/metal-poor groups.