

ULISSE: Determination of the star formation rate and stellar mass based on the one-shot galaxy imaging technique

Olena Torbaniuk^{1,2,*}, Lars Doorenbos^{3,5}, Maurizio Paolillo^{4,5,6}, Stefano Cavuoti^{5,6},
Massimo Brescia^{4,5,6}, and Giuseppe Longo⁴

¹ Department of Physics and Astronomy ‘Augusto Righi’, University of Bologna, Via Gobetti 93/2, 40129 Bologna, Italy

² INAF – Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Via Gobetti 101, 40129 Bologna, Italy

³ AIMI, ARTORG Center, University of Bern, Murtenstrasse 50, 3008 Bern, Switzerland

⁴ Department of Physics ‘Ettore Pancini’, University Federico II, Strada Vicinale Cupa Cintia, 21, 80126 Napoli, Italy

⁵ INAF – Astronomical Observatory of Capodimonte, Salita Moiariello 16, 80131 Napoli, Italy

⁶ INFN – Sezione di Napoli, via Cinthia 9, 80126 Napoli, Italy

Received 22 October 2024 / Accepted 26 July 2025

ABSTRACT

Context. Modern sky surveys produce vast amounts of observational data, which makes the application of classical methods for estimating galaxy properties challenging and time-consuming. This challenge can be significantly alleviated by employing automatic machine- and deep-learning techniques.

Aims. We propose an implementation of the ULISSE algorithm to determine the physical parameters of galaxies, in particular, the star formation rates (SFR) and stellar masses (M_*), based on composite-colour images alone.

Methods. ULISSE is able to rapidly and efficiently identify candidates from a single image based on photometric and morphological similarities to a given reference object with known properties. This approach leverages features extracted from the ImageNet dataset to perform similarity searches among all objects in the sample. This eliminates the need for extensive neural-network training.

Results. Our experiments were performed on the Sloan Digital Sky Survey. They demonstrate that we are able to predict the joint SFR and M_* of the target galaxies within 1 dex in 60% to 80% of cases, depending on the investigated subsample (quiescent and star-forming galaxies, early- and late-type, etc.), and within 0.5 dex when we consider these parameters separately. This is approximately twice the fraction obtained from a random guess extracted from the parent population. Additionally, we found that ULISSE is more effective for galaxies with an active star formation than for elliptical galaxies with quenched star formation. Additionally, ULISSE performs more efficiently for galaxies with bright nuclei such as active galactic nuclei.

Conclusions. Our results suggest that ULISSE is a promising tool for a preliminary estimation of SFR and M_* for galaxies based only on single images in current and future wide-field surveys (e.g. Euclid and LSST), which target millions of sources nightly.

Key words. methods: statistical – techniques: image processing – catalogs – galaxies: elliptical and lenticular, cD – galaxies: spiral – galaxies: star formation

1. Introduction

Over the past two decades, our understanding of the Universe has been significantly enhanced by exploring vast and deep areas of the sky through multi-wavelength digital imaging surveys such as the Sloan Digital Sky Survey (SDSS; York et al. 2000), the Kilo Degree Survey (KiDS; de Jong et al. 2015), the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS; Magnier et al. 2020), the Dark Energy Survey (DES; Dark Energy Survey Collaboration 2016), and the Hyper Suprime-Cam Subaru Strategic Program (HSC SSP; Aihara et al. 2019). In the future, new multi-band wide-field surveys and projects carried out on such telescopes as the Vera C. Rubin Observatory Large Synoptic Survey Telescope (Rubin-LSST; Ivezić et al. 2019), Euclid (Euclid Collaboration; Scaramella et al. 2022; Euclid Collaboration; Mellier et al. 2025), the Nancy Grace Roman Space Telescope (formerly WFIRST; Green et al. 2012), and the James Webb Space Telescope (JWST; Álvarez-Márquez et al. 2019) are set to further increase the volume of observational data by some orders of magnitude. These forthcoming surveys

will produce photometric data for millions of sources each night. Spectroscopic follow-ups for even a small fraction of the sources in these surveys is impracticable, and therefore, algorithms are needed that can leverage photometric information to detect, classify, or measure the physical properties (redshift, masses, star formation rate, etc.) of the sources in these surveys, or that can at least identify candidates for further investigations. In response to this challenge, there has been a concerted effort in recent years to develop and refine fast self-adaptive learning methods for data prediction, classification, and visualisation. This has led to the adoption of astroinformatics solutions, in particular, machine- and deep-learning paradigms (Baron 2019; Longo et al. 2019; Fluke & Jacobs 2020; Lecun et al. 1998; D’Isanto & Polsterer 2018; Schaefer et al. 2018), which try to replace or strengthen the classical methods that are less efficient for large-scale samples.

Machine-learning algorithms can be divided into two main categories: supervised, and unsupervised. Supervised methods, which rely on labelled data for training, are more commonly used because they are easy to interpret and can be tailored to specific problems (e.g. Kim et al. 2011; Brescia et al. 2013; D’Isanto & Polsterer 2018; Kinson et al. 2021; Wenzl et al. 2021). Unsupervised methods, on the other hand, analyse data without

* Corresponding author: olena.torbaniuk@gmail.com

prior labels and use labels only for the post-analysis. They are therefore less common, but have been successfully applied in astrophysics (e.g. [Baron & Poznanski 2017](#); [Castro-Ginard et al. 2018](#); [Razim et al. 2021](#)). In addition, supervised machine learning requires a training dataset from the real observations or simulations, which often lack the comprehensive coverage of the parameter space and therefore do not provide a full picture of the physics behind it (especially for rare and poorly studied objects; see [Masters et al. 2015](#)). The latter can be solved using unsupervised methods in the way of the clustering and pre-clustering approaches ([Bishop 2006](#)).

As an intermediate approach that combines the best from supervised and unsupervised machine-learning methods, one-shot learning has been developed ([Wang et al. 2020](#)), which uses only a single-labelled sample per the studied class of objects. It allows us to eliminate not only the necessity of the extensive labelled datasets, but also addresses the challenge of rare and under-sampled objects.

Using this approach, we present aUtomatic Lightweight Intelligent System for Sky Exploration (ULISSE), a one-shot method designed to select objects closely related to a given input by directly applying it to multi-band images. For this, it transforms the image of a given source (which can be called a prototype or a target object) into a set of representative features, which are then used to search for other objects in the feature space that share physical properties with the target object. A key strength of this approach lies in its flexibility and minimal input requirements: ULISSE operates purely in image space and does not require photometry, redshifts, or spectroscopic information for the input objects. Instead, it transfers physical properties from a well-characterised reference sample via image-based visual similarity. The power of this approach is its ability to provide a reliable list of objects with similar properties (i.e. neighbours), even in the case of rare and peculiar target objects, thereby eliminating the need for a large and well-labelled training set that is essential for supervised methods. This relative matching strategy enables a robust performance even when detailed measurements are unavailable for most sources. ULISSE is therefore particularly well suited for large-scale surveys and imaging datasets.

In our previous work [Doorenbos et al. 2022](#) (hereinafter Paper I), we presented the effectiveness of this method for the detection of active galactic nucleus (AGN) candidates. As one of the results, we observed a correlation between the ULISSE output and the galaxy morphology, which indicates that our method is sensitive enough to distinguish the physical parameters of the studied galaxies. Building on this, we aim to extend the application of our method to probe different galaxy properties. This can be particularly valuable for large-scale surveys such as the LSST and Euclid by offering an efficient alternative solution for processing and analysing the vast amounts of data generated by these surveys in the nearest future.

We propose applying our method to estimate galaxy properties such as the star formation rate (SFR) and stellar mass (M_*), which play an important role in the broad range of studies of the galaxy formation and evolution ([Conselice 2014](#); [Madau & Dickinson 2014](#); [Förster Schreiber & Wuyts 2020](#)), and their gas content ([Carilli & Walter 2013](#); [Morganti & Oosterloo 2018](#); [Maiolino & Mannucci 2019](#)) and co-evolution with supermassive black holes in their centres ([Fabian 2012](#); [Kormendy & Ho 2013](#); [Heckman & Best 2014](#); [Hickox & Alexander 2018](#); [Torbaniuk et al. 2024](#)). Traditionally, the SFR and M_* were extracted based on spectroscopic data or broadband information for a fit of the spectral energy distribution (SED; [Calzetti 2013](#); [Kennicutt &](#)

[Evans 2012](#)). However, these methods can be time-consuming and are often limited by the availability of the spectra and/or broadband observations. Several studies have already proposed an alternative approach to this problem using different machine- and deep-learning techniques, including supervised ([Bonjean et al. 2019](#); [Delli Veneri et al. 2019](#); [Domínguez Sánchez et al. 2023](#)) and semi-supervised ([Humphrey et al. 2023](#)) approaches and neural networks that were pre-trained on the broadband UV, optical, and IR photometry and/or images ([Surana et al. 2020](#); [Chu et al. 2024](#); [Zeraatgari et al. 2024](#); [Zhong et al. 2024](#); [Euclid Collaboration: Kovačić et al. 2025](#)).

[Domínguez Sánchez et al. \(2023\)](#); [Zhong et al. \(2024\)](#) recently presented an approach for determining the galaxy properties based on images similar to those we present here. In contrast to our methods, however, these works were based on self-supervised pre-trained deep-learning techniques.

The paper is structured as follows: in Section 2 we summarise our method, and in Section 3 we describe our galaxy sample and present the sample of target objects we used to test the performance of our method in retrieving the galaxy properties. Section 4 presents the experiments and discusses the results. Finally, in Section 5 we conclude and describe a possible application of our results for other samples. Additional supporting material is provided in the appendices, including the distribution of retrieved neighbours in the SFR– M_* plane for different target classes (Appendices A and B) and the statistical properties of neighbour distributions (Appendix C).

2. Method

The ULISSE algorithm and its performance have already been presented in Paper I in the context of AGN selection. We provide only a brief summary here to introduce the general concept of the method. Detailed information on the pre-training and feature extraction steps, along with computation times and tests involving single-band images and recursive application, can be found in Paper I.

ULISSE uses features that were extracted from a convolutional neural network (CNN; [Schmidhuber 2015](#)) that was trained on a large-scale dataset. Training a CNN on a large and diverse dataset allows it to learn a broad spectrum of features that can be useful beyond the original task. This concept, widely known as transfer learning, has been successfully applied in numerous areas, including astronomy ([Awang Iskandar et al. 2020](#); [Martinazzo et al. 2021](#); [Cavuoti et al. 2024](#)), malware classification ([Prima & Bouhorma 2020](#)), Earth sciences ([Zou & Zhong 2018](#)), and medicine ([Ding et al. 2019](#); [Esteva et al. 2017](#); [Kim et al. 2021](#); [Menegola et al. 2017](#)). The typical large-scale dataset used for training in this context is ImageNet ([Deng et al. 2009](#)).

These features are subsequently used to identify relevant astronomical objects via a nearest-neighbour search. In order to obtain the features, the fully connected layers of the pre-trained network are discarded, and the feature-extracting component of the network is used directly. This approach removes the need for any neural network training and makes direct application to any new dataset possible. Then, in order to obtain image-level properties, we extracted the feature maps from the final convolutional layer of the pre-trained neural network. To reduce the dimensionality of these maps, we averaged over the spatial dimensions. As the deepest layers in the network exhibit the highest level of abstraction ([Goodfellow et al. 2016](#)), we assumed that objects in the dataset whose images have similar deep features to a target object possess similar morphological properties.

We used an EfficientNet-b0 as the CNN (Tan & Le 2019), resulting in a 1280-dimensional feature descriptor for each image. As these features are derived from natural images, they are not immediately interpretable in an astronomical context, although specific patterns linked to particular features can still be identified, as shown in Paper I.

ULISSE identifies objects with similar properties by performing a similarity search in the pre-trained feature space. Using the given target object (or prototype, as in Paper I), it finds the closest objects in this feature space, providing a list of candidate lookalikes. Formally, for the image of a target object \mathbf{x}_q , the nearest neighbours \mathbf{x}_i from a dataset $\{\mathbf{x}_i\}_{i=1}^N$ are retrieved by the lowest Euclidean distances $d(\mathbf{x}_q, \mathbf{x}_i) = \|\mathbf{f}_q - \mathbf{f}_i\|^2$ in the pre-trained feature space, where \mathbf{f}_i denote the features of image i . Although ULISSE does not require prior information about the dataset used to search candidates with similar properties, having a validation set allows us to evaluate the method performance by analysing the behaviour and properties of the returned n closest objects (i.e. neighbours).

3. Dataset

The estimation of the galaxy properties such as stellar mass \mathcal{M}_* and SFR is traditionally based on the analysis of certain features in galaxy spectra (e.g. both continuum and emission lines in a wide range of wavelength, see Kennicutt & Evans 2012; Calzetti 2013) or by fitting the broadband SED. Both methods have their limitations, primarily because they require either spectroscopic information or a comprehensive multi-wavelength photometric dataset (spanning from X-rays to the far infrared) for proper SED reconstruction, which is often not available for large samples. Various machine and deep learning methods offer promising alternatives for addressing the aforementioned problems. In the initial stage, however, these algorithms still need to be tested on samples with reliable estimates of SFR and \mathcal{M}_* , derived from well-established traditional methods.

3.1. Dataset details

To test the ULISSE efficiency in the determination of the galaxy characteristics, we used the same *galSpec* galaxy catalogue¹ as in Paper I, which has been produced by the MPA-JHU group as a subsample from the main galaxy catalogue of the eighth data release of the Sloan Digital Sky Survey (SDSS DR8; Brinchmann et al. 2004). The estimates of total stellar mass (\mathcal{M}_*) in the sample were obtained through Bayesian fitting of the SDSS *ugriz* photometry to a grid of models (see details in Kauffmann et al. 2003b; Tremonti et al. 2004). In the case of SFR determination, two different approaches were used, depending on the object classification according to the BPT-diagram criterion, allowing us to make an assumption on the nature of emission lines (e.g. nebular from star-forming regions or nuclear from AGN) using the ratios of certain emission lines in the galaxy spectrum (Baldwin et al. 1981). The values of SFR for star-forming galaxies were determined using the $H\alpha$ emission line luminosity, while for all other spectral classes, where the $H\alpha$ emission line is weak (e.g. unclassified objects) or contaminated by emission from non-stellar component (e.g. AGN or composite), the empirical relation between SFR and the Balmer decrement, $D4000$, was used (see details in Kauffmann et al. 2003a).

The detailed description of the sample reduction can be found in Paper I (the references therein), where the same sample

was used for testing ULISSE performance in the framework of AGN identification. Based on the experiments conducted in our previous ULISSE paper, the most optimal thumbnail size was chosen as 22×22 arcsec (or 56×56 pixels, the SDSS pixel scale is 0.396 arcsec per pixel). Because the ULISSE efficiency may vary depending on the angular size of galaxies and its ability to resolve the entire galaxy with its morphological features (e.g. disc, bulge, spiral arms, etc), we also decided to remove galaxies with too large (or too small) angular diameter to the size of the thumbnails. Thus, we limited our sample to objects with Petrosian radius (*petroR50* in *r*-band) larger than 1.5 arcsec and smaller than 8 arcsec. For the same reason, we limited our sample to sources with redshift $z > 0.01$ and up to $z \leq 0.15$, above which the morphological classification of galaxies becomes challenging due to the resolution and depth of the SDSS data. As a result, our sample contains 449 762 objects.

In Paper I, we found that ULISSE performance seems to differ for galaxies with various appearances, so to examine this point in the current work we selected objects with the available morphological classification in the second release of Galaxy Zoo (GZ2; Willett et al. 2013), which resulted in 201 626 galaxies. It should be noted that the morphology classification in the GZ2 is quite extensive and reflects not only the presence of such features as spiral arms, bulges, and bars but also considers the strength of these features. Since the goal of our work is not to find galaxies that are identical in morphological appearance, but rather to test whether ULISSE efficiency in SFR and mass prediction varies for galaxies with significantly different appearances, we decided to simplify the original GZ2 classification and instead use only five general morphological classes. The first class, ellipticals (E), contains all galaxies with smooth visual appearance, but without the GZ2 division based on the roundness of the galaxy. As a result, this class may include not only genuine early-type ‘red’ galaxies, but also star-forming ‘bluer’ disk galaxies with poorly resolved features that appear smooth due to limited resolution or projection effects. The next three classes were defined for galaxies with disk or spiral features: spiral galaxies with a bulge (S) and/or bar component (SB) and edge-on spiral galaxies (Se). Again, we did not consider the GZ2 division according to the prominence and shape of the galaxy bulge or bar or the number and relative winding of spiral arms. In addition to the four main morphological categories, we also included galaxies exhibiting visually irregular or atypical features such as rings, lenses and arcs, dust lanes, irregular or interacting morphologies. These features deviate from the idealised, symmetric morphologies of the classes discussed above but do not necessarily imply the presence of active mergers or major interactions. Nonetheless, the presence of such structural peculiarities can still correlate with deviations in physical properties relative to the more classical systems. The galaxies in this class retain their primary morphological classification (i.e. E, S, SB, Se), and we denoted them with a (d) suffix (e.g., E(d), S(d), etc.) to indicate the presence of such visual disturbances. For simplicity, we grouped these systems under their respective disturbed subclasses and refer to them collectively as ‘E(d),S(d),SB(d),Se(d)’ throughout the paper. As demonstrated in Section 4.3, this grouping has no impact on the overall result and ULISSE performance is consistent across both regular and structurally disturbed morphologies, with no systematic bias in the predicted stellar masses or SFR².

² The relatively high number of galaxies with structural irregularities (i.e. E(d), S(d), etc) in our sample (see also the target selection in Section 3.2) also arises from our initial selection strategy, in which we aimed at sampling 3–5 galaxies per stellar mass and SFR bin across

¹ https://www.sdss.org/dr12/spectro/galaxy_mpa_jhu/

Table 1. Summary of the different subsamples.

		Entire		Random		
		<i>N</i>	Fraction	<i>N</i>	Fraction	
Criteria	GZ2 Morphology	Ellipticals (smooth), E	73 075	36.2%	36 173	36.2%
		Spirals, S	67 855	33.7%	33 686	33.7%
		Spirals with bar, SB	15 389	7.6%	7594	7.6%
		Edge-on spirals, Se	19 980	9.9%	9865	9.8%
		E(d),S(d),SB(d),Se(d)	25 327	12.6%	12 682	12.7%
	BPT-diagram	AGN	56 536	28.0%	28 039	28.0%
		Non-AGN	145 090	72.0%	71 961	72.0%
	SFR– \mathcal{M}_* diagram	Star-forming galaxies, SFG	109 664	54.4%	54 301	54.3%
		Quiescent galaxies, QG	91 962	45.6%	45 699	45.7%
	Redshift range	$0.01 < z < 0.05$	51 275	25.4%	25 265	25.2%
		$0.05 < z < 0.1$	102 226	50.7%	50 658	50.7%
		$0.1 < z < 0.15$	48 125	23.9%	24 077	24.1%
	Dust content	Low ($H\alpha/H\beta$) ≤ 2.86	41 779	20.7%	20 700	20.7%
		Moderate $2.86 < (H\alpha/H\beta) \leq 5.0$	110 423	54.8%	54 908	54.9%
		High ($H\alpha/H\beta$) > 5.0	34 067	16.9%	16 786	16.8%
Unknown		15 337	7.6%	7 606	7.6%	
Total number of objects		201 626	–	100 000	–	

Notes. The fractions represent the percentage of objects in each subsample defined based on the GZ2 morphology classification, the presence of an AGN according to the BPT-diagram selection criteria, the galaxy position in the SFR– \mathcal{M}_* diagram, the redshift range, and the dust content defined based on the Balmer decrement.

As we discussed in Paper I ULISSE performance in AGN identification is linked not only to its ability to retrieve the properties of their host galaxies (e.g. morphology and colour), but it seems also to be able to recognize the presence of a central nuclear source (i.e. AGN). Thus, to test for a possible difference in the ULISSE efficiency in the retrieving SFR and \mathcal{M}_* based on whether or not a galaxy hosts AGN, we decided also to add simplified AGN or non-AGN classification for our sample. So, each source in our sample was marked as AGN if it has been identified as AGN, low S/N AGN, or composite source (i.e. contribution from both an AGN and star-forming processes) according to the BPT selection criteria (see details in Brinchmann et al. 2004), while the other objects like SFG, low S/N SFG and unclassified were accordingly marked as non-AGN.

To assess the effect of dust-related reddening on the performance of our image-based method particularly given its reliance on optical g -, r -, i -bands, we also classified galaxies into dust content classes based on the Balmer decrement ($H\alpha/H\beta$), which is a widely used tracer of nebular dust extinction (Brinchmann et al. 2004; Koyama et al. 2015; Qin et al. 2019). We adopted an intrinsic, dust-free ratio of $(H\alpha/H\beta) = 2.86$, consistent with Case B recombination under typical H II region conditions (Fitzpatrick 1999; Lin & Yan 2024). Galaxies with $(H\alpha/H\beta) \leq 2.86$ (corresponding to extinction of approximately $A_V \lesssim 0.5$ mag) were

each morphological subclass, including both regular and visually disturbed systems. This approach was designed to ensure broad coverage across parameter space, rather than to reproduce the actual distribution of galaxy types in the Universe. As a result, the fraction of irregular systems in our sample is artificially elevated compared to their actual occurrence rates in the general galaxy population. We also emphasize that the galaxies marked as E(d), S(d), SB(d) and Se(d), include a wide range of visual features (e.g., rings, dust lanes, tidal structures, etc) and do not solely represent systems with ongoing mergers or strong interactions.

classified as low-dust content. Those with $2.86 < (H\alpha/H\beta) \leq 5.0$ ($A_V = 0.5$ – 1.5 mag) were considered as moderate-dust, and those with $(H\alpha/H\beta) > 5.0$ ($A_V \gtrsim 1.5$ mag) as high-dust content. Galaxies with low S/N ratio for one or both lines and providing unphysical values of $(H\alpha/H\beta)$ which may arise from poor spectral quality or line fitting issues, were excluded from the dust-based subsample analysis (marked as unknown dust content).

The numbers of objects (and their fractions in the sample) of different morphology or AGN classes are presented in Table 1. Since the goal of our work is mainly to show the validity of our method for the galaxy properties extraction and not to obtain results for the whole sample of galaxies, we created a smaller subsample just for more efficient computation. So, similarly to the approach in Paper I, we randomly shuffled the coordinates and take the first 100 000 objects. This subsample is labelled as *Random* in Table 1, and it can be seen its proportions are practically equal to those of the whole sample.

3.2. Target objects

To test ULISSE ability to retrieve galaxies with similar star formation rates and stellar masses, we need to select the set of target objects within a wide range of SFR and \mathcal{M}_* . Target objects in this study represent galaxies with unknown physical parameters (SFR and \mathcal{M}_*), which we want to measure based on a set of neighbours selected by ULISSE. For this purpose, we plotted our *Random* sample in the SFR– \mathcal{M}_* diagram. The distribution of galaxies in the SFR– \mathcal{M}_* diagram is known to show the clustering of galaxies into two main populations: the sequence of star-forming galaxies with steady processes of new stars formation (the so-called main-sequence of star-forming galaxies; MS of SFG), and quiescent galaxies with passively evolving stellar populations. We marked each galaxy in our

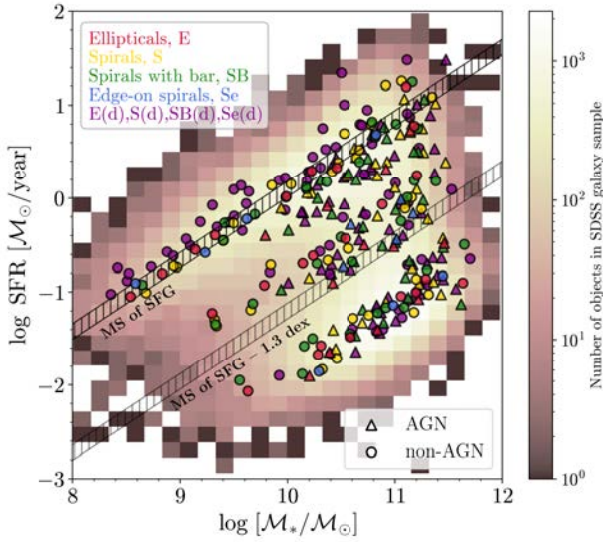


Fig. 1. Distribution of SFR vs. stellar mass for 100 000 galaxies in our sample. The individual target objects selected for our study are shown in different colours and shapes depending on the GZ2 morphology class (from red to purple) and on the AGN or non-AGN selection according to the BPT-diagram criteria (triangle and circle). The shaded black band presents the MS of SFG defined by Eq. (1). The shaded grey band shows a cut at 1.3 dex below the MS of SFG applied to separate the galaxies into star-forming and quiescent galaxy populations. Both areas correspond to the redshift interval $z = 0.01$ – 0.15 .

sample as star-forming or quiescent based on their position in the SFR– M_* diagram relative to the evolving MS of SFG defined by Aird et al. (2017) as

$$\begin{aligned} \log \text{SFR}_{\text{MS}}(z) [\text{M}_\odot \text{year}^{-1}] \\ = -7.6 + 0.76 \log [M_*/\text{M}_\odot] + 2.95 \log(1+z). \quad (1) \end{aligned}$$

The threshold between the two classes was set at 1.3 dex below the MS of SFG: galaxies that fall below this cut were classified as quiescent, while those above the line as star-forming. Note that the relation in Equation (1) is redshift-dependent, so for the classification, we used the redshift of each object.

Then, using the same distribution of galaxies in the SFR– M_* diagram, we manually selected objects uniformly distributed within the diagram. Selecting targets, we tried to cover all possible morphological classes mentioned in Table 1, which resulted in 290 objects. The distribution of the selected target objects in the SFR– M_* diagram is presented in Fig. 1. Note, to avoid introducing additional uncertainties to the final results, we did not select target objects on the weakly populated areas of the SFR– M_* diagram (i.e. mainly on the edges) where ULISSE will suffer to probe objects with similar properties due to intrinsically low number of them in our sample. The number of target objects selected among different classes is presented in Fig. 2. It should be noted that the SFR– M_* diagram has a logarithmic scale (i.e. it is the $\log \text{SFR}$ – $\log M_*$ diagram in reality), but for a matter of simplicity we will refer to both logarithmic quantities ($\log \text{SFR}$ and $\log M_*$) as SFR and M_* throughout the paper.

4. Experiments and results

In this section, we initially discuss the obtained result using as examples six targets covering all the classes discussed in Section 3.2. By covering a broad range of galaxy properties, we

	SFG	QG	AGN	non-AGN	E	S	SB	Se	E(d),S(d),SB(d),Se(d)	$0.01 < z < 0.05$	$0.05 < z < 0.1$	$0.1 < z < 0.15$	low dust cont.	moderate dust cont.	high dust cont.	unknown dust cont.
SFG	305	—	87	118	22	47	41	9	86	87	80	38	10	139	49	7
QG	—	85	43	42	12	17	19	3	32	21	39	25	26	38	14	7
AGN	87	43	130	—	12	32	27	5	54	38	66	26	14	78	33	5
non-AGN	118	42	—	160	24	32	33	7	64	70	53	37	22	99	30	9
E	22	12	12	24	36	—	—	—	—	20	9	7	8	23	4	1
S	47	17	32	32	—	64	—	—	—	20	22	22	8	43	11	2
SB	41	19	27	33	—	—	60	—	—	20	31	9	10	33	14	3
Se	9	3	5	7	—	—	—	12	—	7	3	2	1	6	4	1
E(d),S(d),SB(d),Se(d)	86	32	54	64	—	—	—	—	118	41	54	23	9	72	30	7
$0.01 < z < 0.05$	87	21	38	70	20	20	20	7	41	108	—	—	13	80	13	2
$0.05 < z < 0.1$	80	39	66	53	9	22	31	3	54	—	119	—	13	68	31	7
$0.1 < z < 0.15$	38	25	26	37	7	22	9	2	23	—	—	63	10	29	19	5
low dust cont.	10	26	14	22	8	8	10	1	9	13	13	10	36	—	—	—
moderate dust cont.	139	38	78	99	23	43	33	6	72	80	68	29	—	177	—	—
high dust cont.	49	14	33	30	4	11	14	4	30	13	31	19	—	—	63	—
unknown dust cont.	7	7	5	9	1	2	3	1	7	2	7	5	—	—	—	14

Fig. 2. Number of target objects selected among various classes mentioned in Table 1, including the classification based on the GZ2 galaxy morphology, the presence or absence of AGN, the location in the SFR– M_* diagram, the redshift range and the dust content.

try to establish whether the performance of our method is sensitive to the particular population of galaxies or depends on their properties, as well as identify possible limitations for some specific class. Then, using the entire set of target objects selected in the previous section, we performed a statistical analysis to define the general efficiency of our method and compared it with a ‘random guess’ approach.

4.1. Statistical evaluation of efficiency

As mentioned in the previous section, for each of 290 target objects we have spectroscopically estimated values of SFR and M_* (see Fig. 1) and labels regarding their GZ2 morphology class, AGN or non-AGN nature according to the BPT selection criteria and the position in the SFR– M_* diagram (star-forming or quiescent galaxy). Following a similar approach as in Paper I we performed our experiment using the g -, r -, i -colour-composite SDSS thumbnails. To visualise the obtained results, we chose six target objects with different properties and plotted them together with their retrieved nearest neighbours in the SFR– M_* diagrams in Figs. 3–8. Each figure also provides the SDSS thumbnail analysed by ULISSE as well as the summary of the targets characteristics, including its SDSS name and object ID (objid), celestial coordinates (right ascension RA and declination Dec), spectroscopic redshift z , and the object’s affiliation to the specific class according to its GZ2 morphology, AGN or non-AGN, and its position in the SFR– M_* diagram (see definition in the Section 3). In our work, we decided to set the number of objects N_{neig} retrieved by ULISSE to the 100 nearest neighbours; however, in practice, the choice of N_{neig} depends on the purpose of the user or the goals of the study, and can be changed to any number.

To quantify the ULISSE efficiency in retrieving galaxies with similar SFR and M_* , we define a set of ‘distances’ between the studied target and retrieved neighbours in SFR– M_* parameter space. These distances are divided into two general groups:

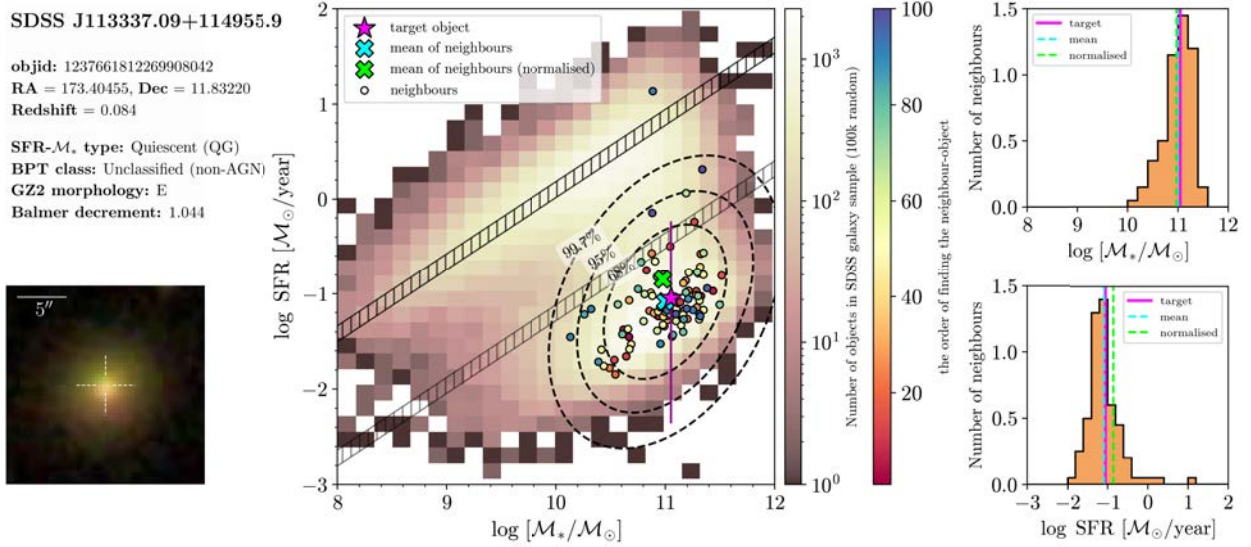


Fig. 3. *Left:* properties of target object 1 with its classification according to the criteria presented in Section 3.1 (*upper panel*) and the SDSS thumbnail used for the analysis by ULISSE (*lower panel*). *Centre:* distribution of SFR versus stellar mass for 100 000 galaxies in our sample. The position of the target in the SFR- M_* diagram (with 1σ errors) is presented as a magenta star together with 100 nearest neighbours (circles). The colour of the circles represents the order in which each neighbour was retrieved by ULISSE. The cyan and green crosses show the position of the mean and normalised mean position for 100 neighbours, respectively. Ellipses enclosing 68%, 95%, and 99.7% of the retrieved neighbours are shown by black dashed lines. The estimated distances for this target are presented in Table 2. The grey- and black-shaded areas are the same as in Fig. 1. *Right:* distributions of M_* (*upper panel*) and SFR (*lower panel*) for 100 neighbours retrieved by ULISSE, normalised to the number of neighbours. Vertical lines show the corresponding values of M_* or SFR for the target object (magenta), the arithmetic mean (cyan), and the normalised mean (green) of the neighbour distributions.

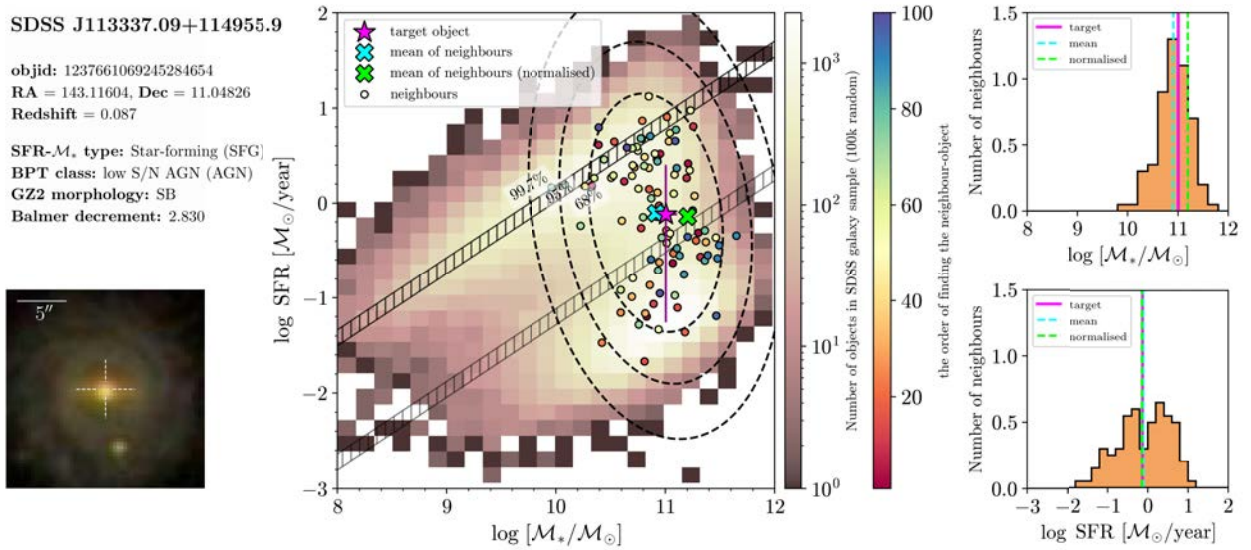


Fig. 4. Same as in Fig. 3, but for target object 2. The estimated distances for this target are presented in Table 2.

the first one represents a simple physical distance (in dex) in SFR- M_* parameter space, while the second one is the weighted version of the first one considering the accuracy of SFR and M_* estimates of the target object. For instance, we propose the so-called total distance, d_{total} , which represents the sum of statistical distances (in SFR and M_* parameter space) between the target (with M_*^{targ} and SFR^{targ} in the SFR- M_* diagram) and each separate neighbours (M_*^{neig} and SFR^{neig}) normalised by the total number of neighbours retrieved by ULISSE (i.e. $N_{\text{neig}} = 100$), and defined as

$$d_{\text{total}} = \frac{1}{N_{\text{neig}}} \sum_i^{100} \sqrt{(M_*^{\text{targ}} - M_{*,i}^{\text{neig}})^2 + (\text{SFR}^{\text{targ}} - \text{SFR}_i^{\text{neig}})^2}$$

Secondly, we introduce the mean distance, d_{mean} , the distance between the position of the target object and the average position among all neighbours retrieved by ULISSE for a specific target. The latter is presented as the mean values of stellar mass ($\langle M_*^{\text{neig}} \rangle$) and star formation rate ($\langle \text{SFR}^{\text{neig}} \rangle$) calculated for 100 neighbours, and then d_{mean} can be expressed as

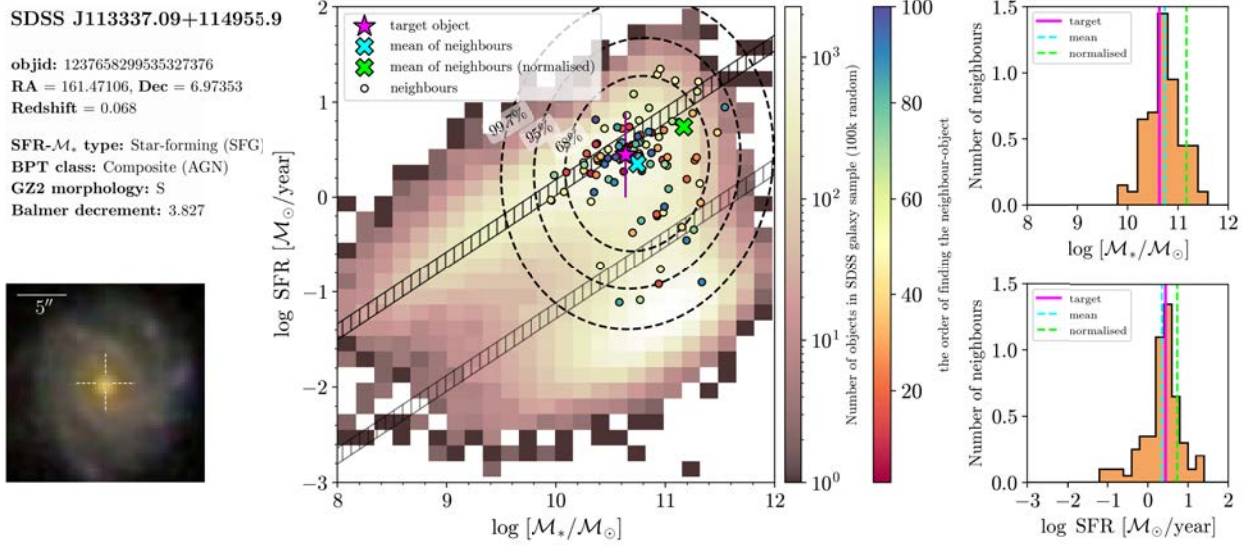


Fig. 5. Same as in Fig. 3, but for target object 3. The estimated distances for this target are presented in Table 2.

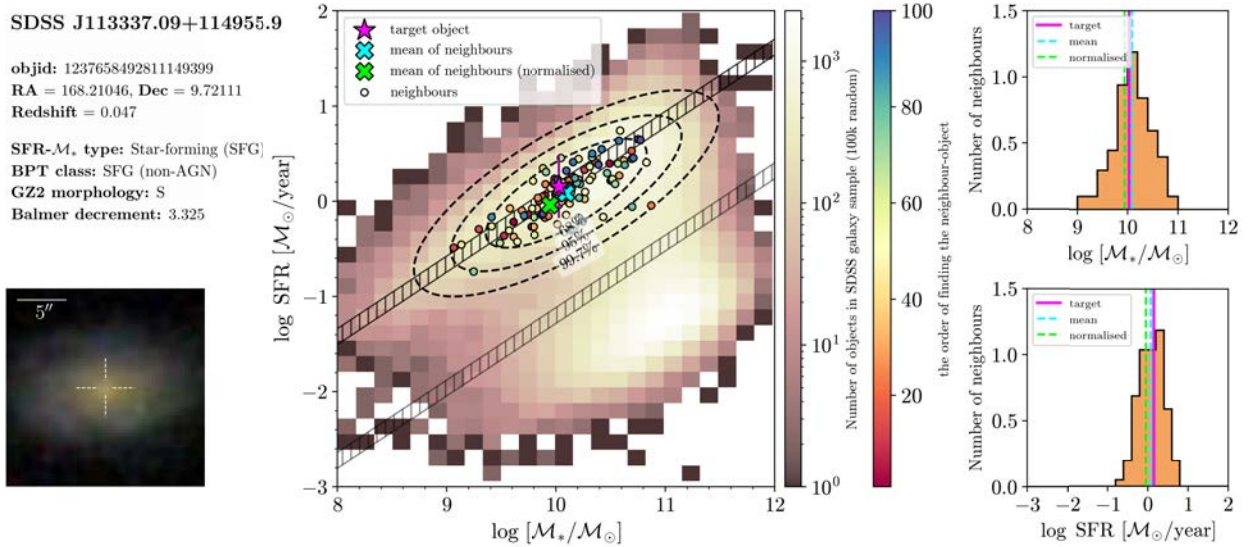


Fig. 6. Same as in Fig. 3, but for target object 4. The estimated distances for this target are presented in Table 2.

$$d_{\text{mean}} = \sqrt{(\mathcal{M}_*^{\text{targ}} - \langle \mathcal{M}_*^{\text{neig}} \rangle)^2 + (\text{SFR}^{\text{targ}} - \langle \text{SFR}^{\text{neig}} \rangle)^2}.$$

The third distance, d_{norm} , is defined similarly to d_{mean} , but before averaging the distances among all retrieved neighbours the position of each neighbour is normalised by the number of galaxies (N_{gal}) in the sample falling within SFR- M_* bin where this neighbour is located. In this way, each neighbour obtains a normalisation parameter reflecting the rarity of this object in the sample and therefore allows us to reduce the bias on the average position of the neighbours (and its distance from the target) due to denser sampling of different regions in the SFR- M_* diagram, intrinsic to any flux- or volume-limited survey. The expression for d_{norm} can be then written as

$$d_{\text{norm}} = \sqrt{(\mathcal{M}_*^{\text{targ}} - \mathcal{M}_*^{\text{norm}})^2 + (\text{SFR}^{\text{targ}} - \text{SFR}^{\text{norm}})^2},$$

where $\mathcal{M}_*^{\text{norm}} = \sum_i \mathcal{M}_{*,i}^{\text{neig}} \cdot w_i^{\text{neig}} / \sum_i w_i^{\text{neig}}$ and $\text{SFR}^{\text{norm}} = \sum_i \text{SFR}_i^{\text{neig}} \cdot w_i^{\text{neig}} / \sum_i w_i^{\text{neig}}$, and $w_i^{\text{neig}} = 1/N_{\text{gal},i}$, $N_{\text{gal},i}$ is the number of galaxies with each SFR- M_* bin for each retrieved neighbour. Note, that the size of each SFR- M_* bin in Fig. 2 is near 0.17 dex.

The second group of distances mimics the first one but takes into account the accuracy of spectroscopic SFR and M_* estimates of each particular target. For this, each distance defined above is normalised by the SFR and M_* uncertainties to obtain the weighted distances that reflect the average significance range (in σ units) of the target into which the retrieved neighbours fall. In this way, the weighted version of the total distance $d_{\text{total}}^{\text{weighted}}$ can be rewritten as

$$d_{\text{total}}^{\text{weighted}} = \frac{1}{N_{\text{neig}}} \sum_i \sqrt{\left(\frac{\mathcal{M}_*^{\text{targ}} - \mathcal{M}_{*,i}^{\text{neig}}}{\sigma_{\mathcal{M}_*}^{\text{targ}}} \right)^2 + \left(\frac{\text{SFR}^{\text{targ}} - \text{SFR}_i^{\text{neig}}}{\sigma_{\text{SFR}}^{\text{targ}}} \right)^2},$$

and the normalised mean distance $d_{\text{mean}}^{\text{weighted}}$ is as

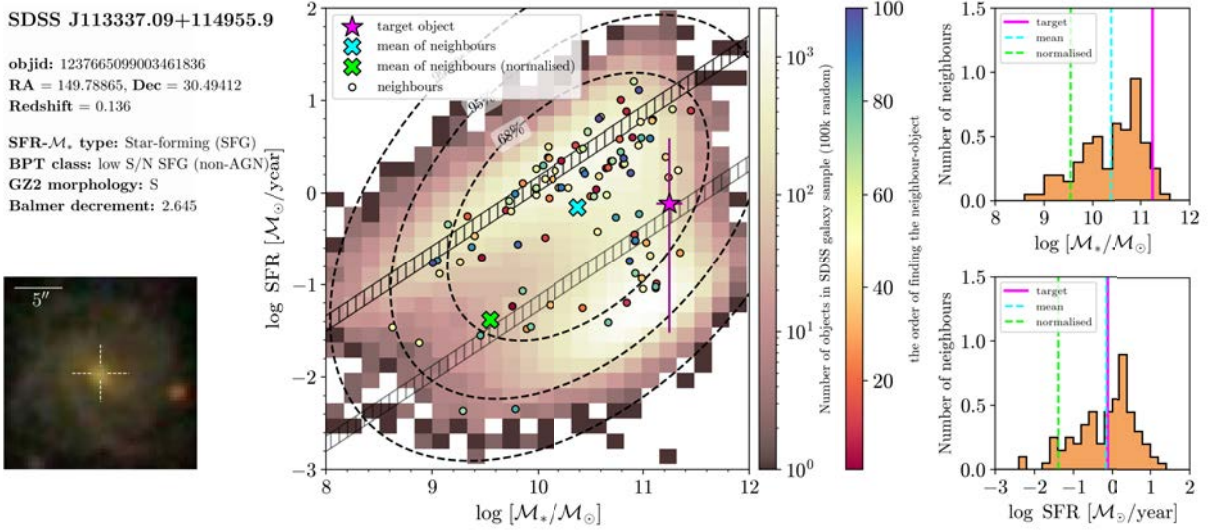


Fig. 7. Same as in Fig. 3, but for the target object 5. The estimated distances for this target are presented in Table 2.

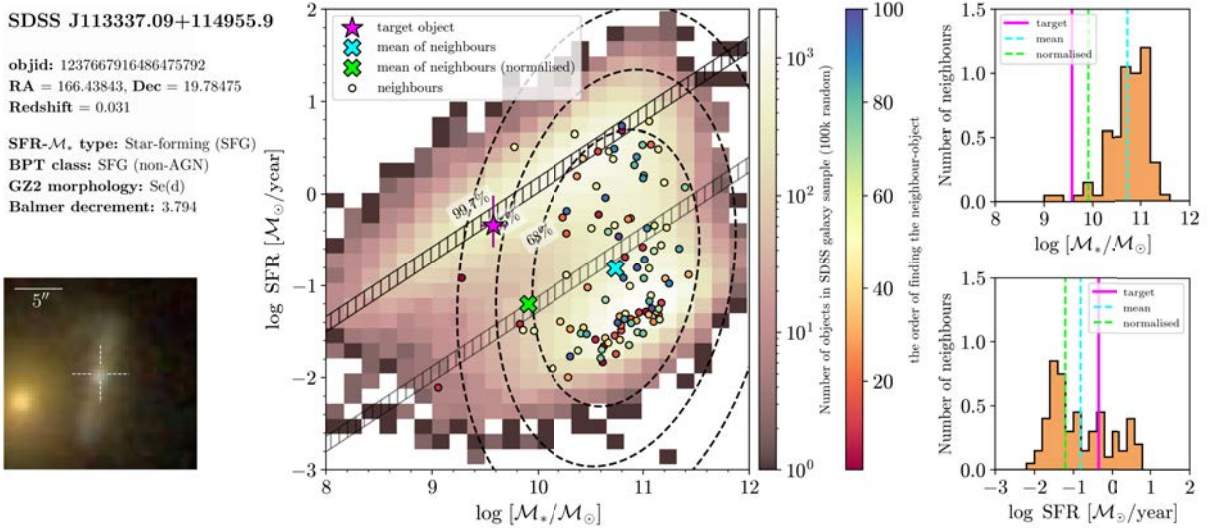


Fig. 8. Same as in Fig. 3, but for the target object 6. The estimated distances for this target are presented in Table 2.

$$d_{\text{mean}}^{\text{weighted}} = \sqrt{\left(\frac{\mathcal{M}_*^{\text{targ}} - \langle \mathcal{M}_*^{\text{neig}} \rangle}{\sigma_{\mathcal{M}_*}^{\text{targ}}}\right)^2 + \left(\frac{\text{SFR}^{\text{targ}} - \langle \text{SFR}^{\text{neig}} \rangle}{\sigma_{\text{SFR}}^{\text{targ}}}\right)^2}.$$

And finally, the normalised version of the weighted distance $d_{\text{norm}}^{\text{weighted}}$ is as

$$d_{\text{norm}}^{\text{weighted}} = \sqrt{\left(\frac{\mathcal{M}_*^{\text{targ}} - \mathcal{M}_*^{\text{norm}}}{\sigma_{\mathcal{M}_*}^{\text{targ}}}\right)^2 + \left(\frac{\text{SFR}^{\text{targ}} - \text{SFR}^{\text{norm}}}{\sigma_{\text{SFR}}^{\text{targ}}}\right)^2}.$$

In addition to the distance-based metrics we introduce a set of metrics which can help us to characterise the spatial distribution of the retrieved neighbours with respect to the target's position in the SFR- \mathcal{M}_* diagram. These metrics may serve as a valuable extension to distance-based measures, offering deeper insight into the efficiency and reliability of the results provided by ULISSE, particularly in cases where distances alone may fail to capture complex spatial patterns in the retrieved neighbours. First of all, we add the standard deviations in stellar mass and

SFR (i.e. $\Delta \mathcal{M}_*$ and ΔSFR) to quantify the spread of the neighbours along each axis, along with the two-dimensional scatter, computed as $\Delta_{2D} = \sqrt{\Delta \mathcal{M}_*^2 + \Delta \text{SFR}^2}$ that represent the individual and joint spread of neighbours serving as a compact diagnostic of how tightly or diffusely they populate the SFR- \mathcal{M}_* plane.

Additionally, we define the shape asymmetry as the ratio $A_s = \Delta \text{SFR} / \Delta \mathcal{M}_*$, quantifying the relative dispersion of neighbours along \mathcal{M}_* and SFR axes. In contrast to its common use as a metric for characterising the elongation of a distribution, we apply A_s as a diagnostic to assess retrieval imbalance: values significantly deviating from unity indicate disproportionate spread in one parameter relative to the other. In this context, shape asymmetry can be used as an indicator of reduced retrieval efficiency in either \mathcal{M}_* or SFR, revealing which axis contributes more substantially to the broadening of the neighbour distribution around the target object.

In contrast, the aspect ratio (\mathcal{R}), defined as the ratio of the principal eigenvalues of the covariance matrix of the neighbour distribution, captures the overall geometric elongation and

Table 2. Summary statistics of the distances and structural parameters for target objects 1–6.

	SDSS name	d_{total}	d_{mean}	d_{norm}	$d_{\text{total}}^{\text{weighted}}$	$d_{\text{mean}}^{\text{weighted}}$	$d_{\text{norm}}^{\text{weighted}}$	ΔM_*	ΔSFR	A_S	\mathcal{R}	C
1	J113337.09+114955.9	0.41 [0.95]	0.07 [0.65]	0.21 [0.78]	2.61 [5.41]	0.60 [4.46]	0.87 [8.02]	0.29	0.42	1.5	2.9	0.43
2	J093227.84+110253.7	0.66 [0.97]	0.10 [0.50]	0.20 [1.35]	2.91 [4.99]	0.99 [3.78]	2.07 [3.29]	0.32	0.65	2.1	4.4	0.76
3	J104553.05+065824.7	0.49 [1.17]	0.14 [0.88]	0.61 [0.88]	3.23 [6.07]	1.15 [2.18]	5.64 [6.42]	0.34	0.48	1.4	2.0	0.60
4	J111250.51+094316.0	0.42 [1.13]	0.10 [0.70]	0.21 [1.24]	3.94 [9.52]	0.83 [5.99]	1.20 [8.45]	0.38	0.30	0.8	7.3	0.28
5	J095909.27+302938.8	1.18 [1.25]	0.87 [0.93]	2.11 [1.64]	7.91 [7.81]	7.76 [7.61]	15.27 [13.22]	0.64	0.75	1.2	2.7	1.61
6	J110545.22+194705.0	1.50 [1.39]	1.25 [1.01]	0.92 [0.82]	11.67 [10.61]	11.05 [9.25]	4.46 [6.82]	0.41	0.78	1.9	3.8	1.19

Notes. The table presents the total, mean, and normalised mean distances (and their weighted versions) obtained by our method and by random guess (values in square brackets). It also includes the standard deviations in stellar mass (ΔM_*) and star formation rate (ΔSFR), along with the shape asymmetry (A_S), aspect ratio (\mathcal{R}), and compactness (C) of the neighbour distributions for these target objects.

directional alignment in the $\text{SFR}-M_*$ plane. While A_S quantifies the imbalance of scatter between the two axes, \mathcal{R} reflects the coherence of that scatter, i.e. whether neighbours are distributed isotropically or exhibit a preferred orientation. High values of \mathcal{R} indicate alignment along a dominant direction within the parameter space (e.g. the MS of SFG), suggesting that the retrieved neighbours, despite their possible spread, tend to follow an underlying intrinsic scaling relation of the galaxy population.

Finally, we define the compactness (C) as the area of the 1σ confidence ellipse enclosing the neighbour distribution, normalised by π , the area of a reference circle with a radius of 1 dex in logarithmic space. This dimensionless quantity provides a standardised measure of how concentrated or dispersed the neighbours are relative to a baseline scale (e.g. circle with 1 dex radius) commonly considered representative of broad distributions in the $\text{SFR}-M_*$ diagram. Values of C smaller than unity indicate a distribution more compact than the assumed baseline scale, reflecting tighter clustering around the target, whereas larger values denote greater spread and less localised neighbour retrieval. This reference radius can be adjusted to other values (e.g. 0.5 dex) depending on the desired sensitivity to clustering scale, the compactness metric naturally rescales with the chosen normalisation, allowing flexibility for more stringent or relaxed definitions of concentration.

4.2. Retrieval efficiency of the physical parameter for individual target objects.

We calculated the six distances defined above for each of 290 target objects introduced in Section 3.2. In the following, we start presenting the results for six template targets, each one representing a different galaxy class and various positions in the $\text{SFR}-M_*$ diagram. To simplify the discussion of the results, we refer to each target with the order in which they are presented in Figs. 3–8 (i.e. target 1, 2, etc).

In Fig. 3, we show the ULISSE results for the elliptical galaxy SDSS J113337.09+114955.9 (i.e. target object 1), which represents a galaxy with smooth morphology (E) according to the GZ2 catalogue and shows no evidence of AGN activity (i.e. this object is unclassified based on the BPT selection criterion meaning that its spectrum shows weak or no emission lines). The neighbours retrieved by ULISSE are well clustered around the target object position in the $\text{SFR}-M_*$ diagram. The distances presented in Table 2 are relatively small indicating a close correspondence between the position of the target and averaged values of SFR and M_* of the neighbours, which thus represent good tracers of the target properties. As seen in Table 2 the scatter in M_* and SFR is modest ($\Delta M_* = 0.29$ and $\Delta \text{SFR} = 0.42$), suggesting a relatively homogeneous sample. The compactness is

significantly below unity ($C = 0.43$), indicating that the neighbours are tightly clustered in the parameter space. The shape asymmetry ($A_S = 1.5$) indicates a moderately larger spread in SFR than in stellar mass, though without a strong dominance of one parameter over the other. The aspect ratio $\mathcal{R} = 2.9$ confirms the absence of strong elongation or preferred orientation. Overall, the combination of these parameters indicates a structurally coherent set of retrieved neighbours, reinforcing the robustness of ULISSE estimates for target 1.

The target object 2 presented in Fig. 4 is the star-forming galaxy SDSS J093227.84+110253.7 with spiral morphology revealing the presence of a bar component (i.e. SB morphology) and a low S/N AGN. The retrieved neighbours appear to be located in the narrow range of stellar masses around M_*^{target} , while they are significantly scattered in SFR. This scatter can be caused by the mixed appearance of the target 2 showing redder nuclear and/or bar component of the galaxy (i.e. lower level of star formation) simultaneously with blue (and more star-forming) spiral disk structure of the galaxy, implying that ULISSE finds similarities with both passive and star-forming sources. However, analysing Table 2, we can see that both ‘physical’ and weighted distances are small for this target. Moreover, the estimated weighted distances are lying within 2σ uncertainty of spectroscopic SFR and M_* indicating that on average the retrieved SFR and M_* still trace well the properties of the target. Furthermore, the large uncertainty in the spectroscopic SFR demonstrates that for such targets the estimate of the stellar formation rate is intrinsically difficult even with spectroscopic SDSS data. The values of $\Delta M_* = 0.32$ and $\Delta \text{SFR} = 0.65$ reflect a relatively tight clustering in stellar mass, while the increased scatter in SFR highlights the intrinsic complexity of the target’s SF activity. The compactness of 0.76 indicates a moderately concentrated neighbour distribution, though less compact than the previous case. The shape asymmetry $A_S = 2.1$ and the relatively high aspect ratio $\mathcal{R} = 4.4$ together indicate a neighbour distribution that is elongated predominantly along the SFR axis, reflecting the dual nature of SF in the target’s less active nuclear and/or bar region and more star-forming spiral disk component.

A similar result was obtained for star-forming galaxy SDSS J104553.05+065824.7 (the target object 3) which presents a spiral morphology (S) and likely hosts an AGN (see Figure 5). As in the previous case, some scatter of the retrieved neighbours toward lower SFR values occurs, most likely due to the presence of the strong red nuclear component. The mean distances for retrieved neighbours are also small (see Table 2) pointing to a good agreement between the properties of the target 3 and the average retrieved parameters. It should be noted that the normalised mean position for neighbours is slightly overestimated with respect to the position of the target (and out of its

5σ significance range), although it remains along the MS of SFG. Such a shift toward higher stellar masses can be possibly brought by the presence of the strong nuclear component (i.e. AGN) in the discussed galaxy or by the fact that ULISSE found a similarity with some rare sources and the normalisation overemphasises their contribution. In the latter case, the remedy would be to use fewer neighbours, as discussed later. The neighbour distribution shows moderate spreads in both \mathcal{M}_* and SFR ($\Delta\mathcal{M}_* = 0.34$ and $\Delta\text{SFR} = 0.48$) indicating a reasonably consistent clustering around the target galaxy. The compactness parameter ($C = 0.60$) reflects a fairly concentrated spatial distribution in the SFR– \mathcal{M}_* plane. The shape asymmetry ($A_S = 1.4$) points to a near-equilibrium dispersion between SFR and \mathcal{M}_* .

This assumption can be tested on the target object 4, star-forming galaxy SDSS J11250.51+094316.0 belonging to the same morphological class S, but without the presence of a bright nucleus or an AGN according to the BPT diagram (see Fig. 6). As a result, the positions of the retrieved neighbours are less scattered than for the two previous objects (2 and 3) and are concentrated along the MS. As can be seen in Table 2 the mean distances are small, indicating an excellent agreement between the SFR and \mathcal{M}_* of the target and those derived from the neighbours. We also point out that, for all previous targets (1–4) the weighted distances are within 3σ of the target (except $d_{\text{mean}}^{\text{norm}}$ for 3 as discussed above), indicating good performance of our method when also considering the uncertainties in the spectroscopic estimates of SFR and \mathcal{M}_* . The neighbour distribution shows moderate spreads in \mathcal{M}_* and SFR ($\Delta\mathcal{M}_* = 0.38$ and $\Delta\text{SFR} = 0.30$) both indicating good retrieval efficiency. The relatively high aspect ratio ($\mathcal{R} = 7.3$) indicates a pronounced elongation of the neighbour distribution, consistent with visual alignment along the MS of SFG (see Fig. 6). Despite this elongation, the compactness parameter ($C = 0.28$) demonstrates that the neighbours remain highly concentrated overall, reflecting a tightly clustered distribution around the target.

Nevertheless, for some selected targets the performance of our method was not as effective as shown above. For instance, for the star-forming galaxy SDSS J095909.27+302938.8 with a spiral (S) morphology and a relatively bright nucleus (the target 5 in Figure 7) the neighbours are spread over a broad range of stellar masses and SFR with a slight concentrating around the MS. The distances presented in Table 2 are also large, indicating a significant mismatch between the average position of the neighbours and the target object. Such poor performance may be the result of the more complex morphology of target 5 and the fact that it fills the thumbnail, but could also indicate that the spectroscopic values of the mass and SFR (which has a large uncertainty) are biased toward the nuclear properties (possibly due to the SDSS fibre size). In fact, most neighbours cluster around the main sequence, suggesting a higher SFR more consistent with the observed spiral morphology). The relatively high dispersions $\Delta\mathcal{M}_* = 0.64$ and $\Delta\text{SFR} = 0.75$ show that the neighbours are widely scattered across the parameter space, reflecting a less homogeneous grouping. This is supported by the large $C = 1.61$, which quantifies the broad spatial distribution of neighbours, likely influenced by the complex morphology of the target and potential biases from limited spectroscopic aperture.

Another example of ULISSE poor performance is presented in Fig. 8. Target object 6 is an edge-on spiral galaxy SDSS J110545.22+194705.0, which shows the presence of disturbed morphology in the GZ2 catalogue (Se(d) class in Fig. 8). The location of the target is close to the MS at relatively low masses, while the neighbours are mostly retrieved as quiescent galaxies with higher stellar masses in the SFR– \mathcal{M}_* diagram. The

reason for this discrepancy is most likely the presence of the elliptical galaxy near our target within our thumbnail. Moreover, this elliptical companion appears brighter than the target object, which may lead ULISSE to focus on its properties rather than those of the target. This confirmed also by the substantial scatter in $\Delta\text{SFR} = 0.78$, while the stellar mass dispersion remains moderate ($\Delta\mathcal{M}_* = 0.41$). Together with a relatively broad neighbour distribution ($C = 1.19$) these parameters show that a retrieval efficiency is influenced by environmental complexity.

The issues highlighted above, which may significantly degrade the performance of our method, can be partly addressed by using larger thumbnails or masking nearby and/or overlapping sources.

In order to compare the efficiency of our method with a purely random guess, for each target object we randomly selected 100 sources from our sample. The distribution of the randomly selected sources in the SFR– \mathcal{M}_* diagram and their total, mean and normalised mean distances for targets 1–6 are presented in Fig. 9 and Table 2. As expected, the random sources are distributed in the most populated areas of the SFR– \mathcal{M}_* diagram, and therefore, their means lie in the ‘valley’ between the populations of the star-forming and quiescent galaxies with $10^9 - 10^{11} M_\odot$. While Table 2 shows that the distances remain within 1 dex from the target, this trend is mainly due to targets originally located near the ‘centre’ of the SFR– \mathcal{M}_* diagram. In fact, the distances derived for the random samples are generally larger than those estimated for actual neighbours and are comparable only in the case of the objects 5 and 6, for which our method has shown a lower performance (see discussion above).

4.3. General efficiency of the method

To evaluate the general ULISSE efficiency we studied the cumulative distribution of the distances defined in Sect. 4.1 for all targets in our sample (see Figs. A.1–A.4). Using these distributions the ULISSE efficiency can be defined as the fraction of target objects (of a certain class) having the corresponding physical distance smaller than a certain value, i.e. the position in the SFR– \mathcal{M}_* diagram estimated based on the retrieved neighbours, is within some threshold value from the position of the target estimated based on spectroscopic data. A similar approach can be applied to the weighted distances in order to obtain the significance threshold of the offset based on the uncertainties on \mathcal{M}_* and SFR of the target. We chose the values 1 dex and 5σ as cuts to define the ULISSE efficiency in retrieving galaxy properties; however, the threshold can be adapted to the specific science case of interest (favouring, e.g. completeness versus accuracy) using the distributions in Figs. A.1–A.4. Table 3 presents the fractions of targets falling below the defined thresholds.

Analysing the results reported in Table 3 (and figures in Appendix A) we found that our algorithm, based on the average properties of the neighbours found by ULISSE, is able to provide an estimate of the target galaxy properties within 1 dex of the expected value for 60–80% of the studied galaxies. At the same time, a disjoint analysis of star-forming and quiescent galaxies reveals a slightly increased efficiency for quiescent galaxies (78% on average for the three distances) compared to SFG (72%), which can be due to the fact that quiescent galaxies represent a homogeneous population of massive and old galaxies in the SFR– \mathcal{M}_* diagram, while SFG exhibit the variety of morphological features (e.g. spirals, bars, etc.) and a wider range of physical parameters. At the same time, the spectroscopically estimated SFR for SFG with a more complex appearance (i.e.

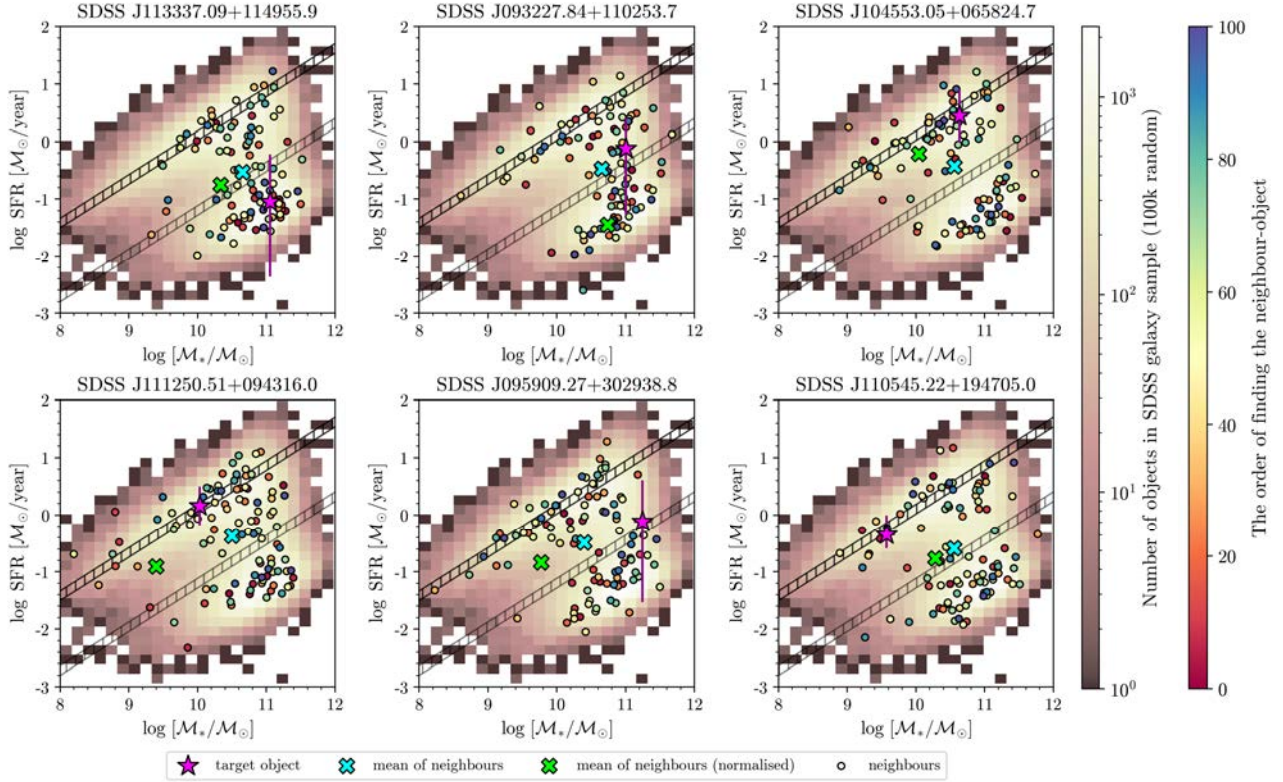


Fig. 9. Similar SFR vs. M_* distribution for the target objects presented in Figs. 3–8, but 100 neighbours have been retrieved randomly from the sample.

Table 3. ULISSE efficiency in retrieving SFR and M_* for different galaxy classes.

Target object sample	N	Within 1 dex			Within 5σ		
		d_{total}	d_{mean}	d_{norm}	$d_{\text{total}}^{\text{weighted}}$	$d_{\text{mean}}^{\text{weighted}}$	$d_{\text{norm}}^{\text{weighted}}$
All	290	69.0% [14.8%]	80.0% [60.3%]	72.4% [34.8%]	54.5% [15.2%]	65.2% [44.1%]	65.2% [29.0%]
SFG	205	66.8% [16.1%]	76.6% [54.6%]	72.7% [39.5%]	46.8% [12.2%]	60.5% [41.0%]	58.5% [29.3%]
QG	85	74.1% [11.8%]	88.2% [71.8%]	71.8% [23.5%]	69.4% [12.9%]	76.5% [51.8%]	80.0% [24.7%]
AGN	130	73.8% [22.3%]	86.2% [76.2%]	73.8% [33.8%]	63.8% [9.2%]	74.6% [57.7%]	73.1% [26.2%]
non-AGN	160	65.6% [10.6%]	75.0% [47.5%]	73.1% [35.6%]	43.8% [13.1%]	57.5% [33.1%]	58.5% [28.1%]
E	36	63.9% [16.7%]	66.7% [58.3%]	75.0% [41.7%]	41.7% [5.6%]	61.1% [41.7%]	69.4% [22.2%]
S	64	70.3% [10.9%]	84.4% [59.4%]	73.4% [18.8%]	53.1% [7.8%]	70.3% [46.9%]	60.9% [17.2%]
SB	60	68.3% [16.7%]	83.3% [61.7%]	80.0% [45.0%]	58.3% [11.7%]	66.7% [50.0%]	68.3% [28.3%]
Se	12	66.7% [25.0%]	75.0% [66.7%]	83.3% [41.7%]	50.0% [16.7%]	50.0% [50.0%]	66.7% [25.0%]
E(d), S(d), SB(d), Se(d)	118	67.8% [16.9%]	79.7% [59.3%]	70.3% [36.4%]	55.1% [16.1%]	66.9% [44.1%]	62.7% [32.2%]
$0.01 < z < 0.05$	108	65.7% [12.0%]	76.9% [50.0%]	82.4% [43.5%]	35.2% [6.5%]	50.9% [38.9%]	57.4% [36.1%]
$0.05 < z < 0.1$	119	77.3% [23.5%]	88.2% [71.4%]	75.6% [39.5%]	68.1% [20.2%]	84.0% [64.7%]	71.4% [31.9%]
$0.1 < z < 0.15$	63	52.4% [4.8%]	68.3% [52.4%]	57.1% [12.7%]	46.0% [4.8%]	61.9% [15.9%]	61.9% [3.2%]
$(H\alpha/H\beta) \leq 2.86$	36	63.9% [11.1%]	77.8% [58.3%]	75.0% [19.4%]	58.3% [0.0%]	61.1% [44.4%]	58.3% [22.2%]
$2.86 < (H\alpha/H\beta) \leq 5.0$	177	68.4% [14.1%]	80.8% [56.5%]	72.3% [38.4%]	49.2% [11.9%]	67.8% [40.7%]	64.4% [28.2%]
$(H\alpha/H\beta) > 5.0$	63	68.3% [15.9%]	73.0% [61.9%]	71.4% [39.7%]	52.4% [7.9%]	61.9% [49.2%]	60.3% [30.2%]
Unknown dust content	14	78.6% [35.7%]	92.9% [85.7%]	71.4% [28.6%]	92.9% [42.9%]	92.9% [71.4%]	85.7% [21.4%]

Notes. Efficiency is defined as the fraction of target objects whose physical distance (total, mean, or normalised mean) is smaller than 1 dex or whose weighted distances are smaller than 5σ (see details in the text).

the presence of multiple stellar components) can be easily miscalculated due to the limitations in the SDSS fiber size and its positioning on the particular galaxy component. However, the higher efficiency for quiescent galaxies must be interpreted with caution because our primary sample is limited. The quiescent sample primarily includes high-mass, red-sequence galaxies with very low SFRs, occupying a relatively narrow region in

the SFR– M_* diagram (typically 1.5–2.5 dex in both axes). In such a constrained parameter space, even simple similarity-based retrievals are likely to succeed. Therefore, part of the observed performance may reflect the limited diversity of our quiescent sample, rather than indicating universally high efficiency of the method across all early-type systems, and as a result require more detailed studies on a broader and more diverse sample of

Table 4. Average structural and statistical properties of neighbour distributions for different classes of target galaxies.

Target object sample	$\langle \Delta M_* \rangle$	$\langle \Delta \text{SFR} \rangle$	$\langle \Delta_{2D} \rangle$	$\langle A_S \rangle$	$\langle \mathcal{R} \rangle$	$\langle C \rangle$
All	0.38	0.61	0.47	1.7	4.6	0.79
SFG	0.40	0.60	0.48	1.6	4.8	0.80
QG	0.33	0.62	0.45	1.9	4.4	0.75
AGN	0.36	0.66	0.48	1.9	4.3	0.85
non-AGN	0.39	0.57	0.47	1.5	4.9	0.74
E	0.37	0.55	0.45	1.5	3.9	0.72
S	0.38	0.60	0.47	1.7	4.8	0.77
SB	0.36	0.62	0.46	1.8	4.5	0.76
Se	0.40	0.66	0.51	1.7	4.5	0.91
E(d), S(d), SB(d), Se(d)	0.39	0.62	0.48	1.7	4.9	0.82
$0.01 < z < 0.05$	0.41	0.57	0.47	1.5	5.2	0.75
$0.05 < z < 0.1$	0.37	0.63	0.48	1.8	4.2	0.82
$0.1 < z < 0.15$	0.35	0.63	0.46	1.9	4.4	0.79
$(\text{H}\alpha/\text{H}\beta) \leq 2.86$	0.36	0.57	0.44	1.7	4.3	0.72
$2.86 < (\text{H}\alpha/\text{H}\beta) \leq 5.0$	0.38	0.60	0.47	1.7	5.0	0.77
$(\text{H}\alpha/\text{H}\beta) > 5.0$	0.38	0.65	0.49	1.8	3.8	0.88
Unknown dust content	0.34	0.65	0.46	2.0	4.9	0.78

Notes. The table includes the average values of the standard deviations in stellar mass (ΔM_*) and SFR (ΔSFR), the two-dimensional scatter (Δ_{2D}), and the shape asymmetry (A_S), aspect ratio (\mathcal{R}), and compactness (C) for the distributions of the retrieved neighbours. The definitions of all parameters are provided in Section 4.1.

quiescent galaxies in the future. This is also evident from the average values of the scatter parameters presented in Table 4. For instance, the average spread in M_* is slightly higher for SFGs ($\Delta M_* = 0.40$) than for QGs ($\Delta M_* = 0.33$), while the SFR scatter is comparably large for both populations ($\Delta \text{SFR} \approx 0.6$).

At the same time, the presence of an AGN in the galaxy core seems to increase only slightly the efficiency of our method (77.9% for AGN versus 71.2% for non-AGN target objects). The average scatter in stellar mass (ΔM_*) is comparable between AGN and non-AGN target objects. In contrast, ΔSFR appears to be larger for AGN hosts, possibly due to contamination from AGN emission affecting the galaxy’s integrated colours and thereby biasing the inferred star formation properties.

Comparing the fractions of galaxies with different morphology classes, we see that elliptical galaxies (E) seem to have a relatively lower efficiency (nearly 68.5% on average for three distances) compared to all other classes (76.0%, 77.2%, 75.0%, and 72.6% on average for objects with S, SB, Se, and a disturbed morphology, respectively). This result shows the apparent contradiction with the results presented above for the case of quiescent and star-forming galaxies defined by their position below the MS of SFG (in Section 3.2), where quiescent galaxies appear to have a higher efficiency than SFGs (see Table 3). As we pointed out in Section 3.2, however, the elliptical (E) class in the GZ catalogue is based purely on their smooth visual appearance and does not strictly correspond to the classical elliptical morphology. As a result, this group includes not only true early-type galaxies but also a fraction of blue, star-forming systems whose features are poorly resolved due to limited image resolution (which can be also seen from the distribution of E galaxies in the SFR- M_* diagram in Fig. 2). At the same time, both ΔM_* and ΔSFR appear to be similar across target objects with different morphological classes.

The further investigation of ULISSE performance against dust-related reddening shows that the retrieving efficiency remains relatively constant for low-, moderate-, and high-dust

content subsamples of galaxies (72.2%, 73.8%, and 70.9% on average, respectively). The same trend is seen for the average scatter of retrieved stellar mass (ΔM_* in Table 4) for different dust content samples, while ΔSFR scatter is showing a modest increase with dust content of galaxies (from 57% for low dust content to 65% for high dust). This can be likely due to the fact that ULISSE selection of similar features also accounts for typical dust-induced reddening features present in galaxies within the sample. Consequently, dusty galaxies tend to be matched with analogues of comparable dust content and morphological appearance, reducing biases caused by dust obscuration and age-metallicity degeneracies. And as a result, it allows ULISSE to maintain stable performance even when relying solely on optical g -, r -, i -band images.

In addition, the efficiency is dependent on the redshift, with the highest efficiency (near 80.4% on average) for the central $0.05 < z < 0.1$ redshift interval with respect to 75.0% and 59.3% for $0.01 < z < 0.05$ and $0.1 < z < 0.15$ intervals, respectively. This is a by-product of the dependence of the angular size of galaxies from redshift, considering that we use a fixed thumbnail size. This leads to situations where at low redshifts we may be missing part of the information as we are cutting the edges of the galaxy, while at higher z the resolution of our images may not be high enough to reveal the same number of morphological features. Across different redshift ranges, ΔSFR remains relatively constant, while ΔM_* shows a tendency to increase with redshift. This trend may reflect the greater sensitivity of M_* estimates to image resolution and structural detail, while SFR estimates, which are more closely tied to integrated colours, appear less affected across the redshift range considered.

The relations between ΔM_* and M_* , as well as between ΔSFR and SFR, are presented in the supplementary material (see Fig. C.1).

A summary of the neighbour distributions retrieved by ULISSE across different subsamples is presented in Table 4. Across all classes (defined by galaxy morphology, AGN contribution, dust content, and redshift) the distributions tend to be compact, with compactness values $C < 1$, indicating tightly clustered neighbour distribution in the SFR- M_* space. The shape asymmetry A_S , which characterises whether scatter is more dominant in SFR or stellar mass, varies between subsamples: it is generally higher for quiescent galaxies compared to star-forming ones, reflecting larger uncertainties in SFR estimates for passive systems. A similar trend is seen for AGN-hosting galaxies and spiral morphologies relative to non-AGN and smooth (elliptical-like) systems, respectively, suggesting greater physical diversity or complexity in their star formation histories, structural features, or spectral energy distributions. The shape ratio \mathcal{R} shows moderate values across all subsamples (typically ~ 4), although its interpretation as elongation along the MS is only valid for targets located near the MS of star-forming galaxies. In off-sequence regions, \mathcal{R} still provides a useful tracer of the neighbour distribution’s geometry, but its physical interpretation may differ. A full distribution of these parameters for all 290 target objects is presented in Fig. C.2 and can serve as an additional diagnostic to assess the performance of ULISSE on an object-by-object basis.

Analysing the distributions of distances for the random sample (see dashed lines in Figs. A.1–A.4) we found that the fraction of targets with distances below 1 dex are in general half of those obtained by our method, meaning that ULISSE is at least twice as efficient than a random guess. However, we point out that using the mean distance (d_{mean}), the efficiency for the random sample is significantly larger compared to the other distances. This is mostly due to the fact that this distance indicator is less

Table 5. ULISSE efficiency in separately retrieving SFR and \mathcal{M}_* for different classes of galaxies.

Target object sample	$\mathcal{M}_*^{\text{targ}} - \langle \mathcal{M}_*^{\text{neig}} \rangle$	$\mathcal{M}_*^{\text{targ}} - \mathcal{M}_*^{\text{norm}}$	$\text{SFR}^{\text{targ}} - \langle \text{SFR}^{\text{neig}} \rangle$	$\text{SFR}^{\text{targ}} - \text{SFR}^{\text{norm}}$
All	94.1% [81.4%]	93.8% [61.0%]	81.0% [70.7%]	73.4% [57.2%]
SFG	92.2% [79.5%]	91.2% [62.9%]	80.0% [70.2%]	76.1% [54.1%]
QG	98.8% [85.9%]	100.0% [56.5%]	83.5% [71.8%]	67.1% [63.5%]

Notes. The fraction of target objects with the mean, $\langle \mathcal{M}_*^{\text{neig}} \rangle$ or $\langle \text{SFR}^{\text{neig}} \rangle$, or normalised mean, $\mathcal{M}_*^{\text{norm}}$ or SFR^{norm} , values for the retrieved neighbours within 1 dex / $\sqrt{2}$ from the $\mathcal{M}_*^{\text{targ}}$ or SFR^{targ} estimations of the target.

sensitive to outliers, and since most sources in the SFR- \mathcal{M}_* diagram are clustered around the MS or in the quiescent galaxies region, (about 1.5×1.5 dex in size), even a random guess will often return close neighbours. Thus, the mean distance is not the best estimator to quantify the efficiency of our approach and we recommend using the normalised mean distance instead, which considers the varying density of objects in our sample across the SFR- \mathcal{M}_* diagram, and also performs well using lower distance thresholds (e.g. 0.5 dex instead of 1 dex, Figs. A.1–A.4).

The distribution of the weighted distances presents a similar trend, but the fractions of neighbours retrieved within 5σ significance range are generally lower compared to the fractions obtained using a threshold of 1 dex. As can be seen in Table 3, independently of the galaxy subsample, our method always returns more than 50% of neighbours within 5σ range for SFR and \mathcal{M}_* of the targets. This is mostly due to the relatively small formal errors on the spectroscopic mass estimates (see Figs. 3–8), so that in practical applications one may relax this constrain depending on the quality of the reference spectroscopic sample.

Our tests so far have considered the position of the neighbours of each target in the SFR- \mathcal{M}_* diagram. However, as highlighted above and in Section 4.2, the uncertainties in both the expected mass of the target (based on SDSS spectroscopy) and the retrieved one (based on the neighbour dispersion) are usually much smaller than the one on the SFR. Thus, to further characterise our efficiency, in Fig. B.1 (see Appendix B) and Table 5 we present the retrieval fractions within 1 dex / $\sqrt{2}$ ³ when the SFR and \mathcal{M}_* are considered separately. Taking into account that the results presented in Table 3 are relatively similar for different galaxy subsamples, we decided to present results only for all, star-forming and quiescent galaxies.

As expected, we found that our approach is more efficient in retrieving \mathcal{M}_* than SFR, i.e. the fraction of neighbours returned by our method for the entire sample of target objects (i.e. all galaxies) is near 94% within 1 dex / $\sqrt{2}$ range (with $\sim 71\%$ for the random approach), and it even reaches 100% applying the method for the subsample of quiescent galaxies (with $\sim 71\%$ for the random approach). On the contrary, for the SFR the efficiency of our method is close to 77% on average. The efficiency of SFR determination for quiescent galaxies is compatible with the random method which is a consequence of the definition of the class, clustering at very low SFR within less than 1 dex / $\sqrt{2}$ from the average value. In fact, we retrieved fractions corresponding to 1 dex in the SFR- \mathcal{M}_* diagram within 0.5 dex in each individual quantity making the method even more robust.

³ Assuming a uniform distribution. For a Rayleigh distribution, however, that is, for a normal distribution in both quantities, we obtained a slightly lower factor of 0.66.

4.4. Applications and limitations of the method ULISSE

A central advantage of ULISSE lies in its ability to estimate key galaxy properties such as \mathcal{M}_* and SFR using only single-band images, without requiring any additional photometric, spectroscopic, or redshift information. This makes the method particularly well suited for application to current and upcoming wide-field imaging surveys (e.g. LSST and Euclid), where some data may be incomplete, expensive to acquire, or entirely unavailable in early stages. Unlike most machine learning pipelines that rely on redshift-dependent training sets and regress directly on physical properties, ULISSE operates by retrieving visually and structurally similar galaxies from a pre-characterised reference sample. This image-similarity-based approach captures morphological, structural, and colour features simultaneously, linking a galaxy’s appearance directly to physical properties in a data-driven way.

It is important to emphasize that ULISSE does not perform absolute regression or classification, but rather a relative matching of each input galaxy to a reference sample whose properties are derived through photometric or spectroscopic measurements. This template-based comparison strategy allows ULISSE to effectively transfer the accuracy and reliability of these measurements to larger imaging-only datasets. However, the precision of the stellar mass and SFR estimates is inherently tied to the quality and representativeness of the reference sample, because the method relies on the properties of the matched target objects. In addition to estimating physical properties, ULISSE can also be used in a reverse mode as a similarity-based retrieval tool. For instance, given a small sample of rare or atypical galaxies (e.g. post-mergers, ring galaxies, or extreme starbursts) ULISSE can identify other objects with similar appearance within large surveys. This capability allows for systematic identification of analogues to specific galaxy populations and supports the construction of statistically robust or morphologically uniform samples of rare systems, thereby providing a foundation for targeted follow-up observations and population analyses.

It worth to mention that the presence of underlying physical correlations between galaxy properties can enhance the performance of ULISSE. A clear example is the star-forming main sequence, a well-established relation between \mathcal{M}_* and SFR for star-forming galaxies. Along this sequence, structural and colour properties tend to co-evolve, meaning that galaxies with similar visual appearance also share similar physical characteristics. This is seen in the left panel of Fig. C.2, where neighbour distributions that are compact ($C \approx 1$) and symmetric ($A_S \approx 1$) also exhibit high shape ratios \mathcal{R} , indicative of elongated distributions aligned with the MS of SFG. Conversely, in regions of parameter space that are less well-populated in the reference sample (such as the green valley, rare morphologies, or low-mass systems) the lack of a consistent underlying relation

and limited statistical support can reduce the reliability of the retrieved parameters.

To explore the limitations of ULISSE in handling targets located in sparsely populated regions of the SFR– \mathcal{M}_* diagram, we examined the relationship between the distance ratio and various characteristics of the neighbour distributions. The distance ratio is defined as the ratio between the ‘unnormalised’ mean distance of neighbours and the normalised mean distance ($d_{\text{mean}}/d_{\text{norm}}$), where the latter incorporates a weighting factor that reflects the rarity of the target object location within the SFR– \mathcal{M}_* plane (see definition in Section 4.1). In this context, values of $d_{\text{mean}}/d_{\text{norm}}$ significantly greater than unity indicate targets located in sparsely populated regions of parameter space, where the unnormalised mean distance d_{mean} is biased by the limited number of similar objects. As a result, the normalised mean distance d_{norm} provides a more reliable measure of the neighbour retrieval efficiency in these cases.

Intuitively, it might be expected that targets with high distance ratios (i.e. those in low-density or edge regions of the parameter space) would display neighbour distributions that are more scattered or asymmetric, reflecting lower efficiency of ULISSE. However, our analysis (see right panel of Fig. C.2) reveals no significant correlation between the distance ratio and two-dimensional scatter Δ_{2D} (or shape asymmetry A_S) of the neighbour distributions. This result implies that ULISSE retrieval efficiency is not determined solely by the density of the reference sample. Instead, other factors (e.g. image quality, resolution or other observational limitations) are likely to influence the fidelity of neighbour retrieval. This complexity highlights the intrinsic difficulty of correlating image-based similarity metrics with underlying physical galaxy properties, and indicates that simplistic dispersion measures may be insufficient to reliably flag targets prone to less accurate parameter estimation.

As shown in Table 4, ULISSE achieves typical uncertainties of approximately 0.3–0.4 dex in \mathcal{M}_* and 0.5–0.6 dex in SFR (see Fig. C.1). Conventional techniques such as SED fitting (Bruzual & Charlot 2003; Salim et al. 2007; Conroy 2013) or spectroscopic line diagnostics (e.g. Brinchmann et al. 2004) typically achieve uncertainties on the order of 0.1–0.3 dex for \mathcal{M}_* and less than 0.3–0.5 dex for SFR, but these methods require high-quality multi-wavelength photometry or spectroscopy and are subject to systematic uncertainties associated with assumptions about star formation histories, dust attenuation, and stellar population models. Moreover, spectroscopic SFR indicators can be unreliable in quiescent galaxies or AGN-dominated systems due to weak or contaminated emission lines. While ULISSE precision does not match that of physically motivated methods, its efficiency remains significant given that it relies solely on imaging data. Moreover, it provides robust estimates across diverse galaxy populations, including quiescent, star-forming, and AGN-hosting systems, and is relatively insensitive to dust effects or spectral contamination, making it a computationally efficient, minimal-input alternative for physical parameter estimation.

Another distinctive strength of ULISSE is its independence from the presence of redshift information, which is an uncommon feature among modern machine learning (ML) approaches for estimating galaxy properties. Most traditional and ML-based techniques rely on luminosity-sensitive observables, requiring either spectroscopic redshifts or distance proxies as inputs (Bonjean et al. 2019; Domínguez Sánchez et al. 2023; Zeraatgari et al. 2024). In contrast, ULISSE operates purely in image space, leveraging structural and colour information without requiring availability of redshift. This makes it particularly well suited

for application in early stages of wide-field imaging surveys and datasets where distance information may be incomplete or unavailable. Nonetheless, while redshift is not an explicit input to our method, its performance is not entirely redshift-independent, as indicated by subtle differences in ULISSE results across redshift ranges shown in Table 3. This is expected because features such as angular size, surface brightness, and apparent morphology are affected by redshift-driven effects like cosmological dimming and resolution loss.

This also raises the question of whether ULISSE, although not designed for this purpose, might implicitly capture some redshift information through its neighbour selection process. To explore this, we conducted a simple test comparing the spectroscopic redshift of each target galaxy with the average redshift of its retrieved neighbours (see Fig. C.3). The obtained average absolute redshift differences shows to be relatively small, typically in the range of 0.02–0.03. Although this level of agreement is promising, it requires careful interpretation due to potential underlying uncertainties. The narrow redshift range of our dataset ($0.01 < z < 0.15$) limits morphological evolution and restricts the parameter space, which likely contributes to the observed consistency. Interestingly, Fig. C.3 reveals a correlation between ULISSE performance and redshift consistency: targets with larger offsets between their redshift and the mean redshift of their neighbours tend to exhibit larger deviations in estimated \mathcal{M}_* and SFR (i.e. higher $\Delta\mathcal{M}_*$ and ΔSFR). This suggests that the accuracy of the method can be improved by applying ULISSE within narrower redshift intervals or by incorporating a redshift prior (when available) into the neighbour selection process. Such a refinement would be analogous to known practices in SED fitting, where including redshift priors is known to improve the robustness of derived galaxy properties (Bolzonella et al. 2000; Ilbert et al. 2006; Conroy 2013).

Additionally, the evolving nature of galaxy populations with redshift, particularly regarding their SFR and \mathcal{M}_* , further complicates the interpretation. As galaxies evolve, the distribution of them in the SFR– \mathcal{M}_* plane changes, meaning a reference sample drawn from a limited redshift range may not fully represent these variations at higher redshifts. Observational limitations, such as survey sensitivity and instrument capabilities, can influence this effect by preferentially excluding low-SFR or low-mass galaxies at higher redshifts. These factors reduce the representativeness of the reference set in some regions of parameter space, potentially impacting the reliability of neighbour matching and any implicit redshift inference.

Furthermore, redshift-dependent degeneracies in image appearance (e.g. surface brightness dimming and limited angular resolution) complicate the interpretation of similarity-based matching in this context. Since ULISSE selects analogues based on visual similarity, features correlated with redshift may bias neighbour selection, making it challenging to disentangle intrinsic physical resemblance from redshift-driven projection effects.

Consequently, while these findings suggest that ULISSE may exhibit some degree of redshift consistency through structural matching, it is not designed to serve as a redshift estimator. A more systematic assessment of ULISSE potential in this role would require a dedicated study covering a broader redshift range and incorporating appropriate calibration techniques. More detailed comparisons with established photometric redshift methods (Ilbert et al. 2006; Way et al. 2009; Hildebrandt et al. 2010; Cavuoti et al. 2017; Soo et al. 2018; Pasquet et al. 2019; Euclid Collaboration: Desprez et al. 2020; Pathi et al. 2025) would be also necessary to rigorously evaluate its applicability for redshift estimation.

5. Conclusions

We applied ULISSE to predict the stellar mass and star formation rate of resolved galaxies. Our method relies on a single composite-colour image of a target galaxy with unknown physical properties, and it is based on a pre-trained convolutional neural network. It extracts a set of representative features from the image without requiring any specific astrophysical knowledge. Applied to a sample of galaxies with known properties (e.g. derived from spectroscopy), ULISSE sorts all objects according to the distance in this feature space (i.e. from the most to the least similar) from the target object and thus allows us to retrieve a specified set of sources with similar properties.

In Paper I we have already applied our method to selection AGN candidates based on a sample of SDSS galaxies and their composite-colour images. We here tested the performance of our method in predicting the galaxy M_* and SFR, assuming that the average properties of the retrieved neighbours are representative of those of the target galaxy. Based on the results obtained running ULISSE on a sample of 290 sources with different stellar masses, SFR, morphologies (e.g. elliptical, spiral with or without bar structure based on the GZ2 survey), and the lack or presence of AGN signatures, we reached the conclusions listed below:

- The efficiency of our method in estimating the physical properties of the target galaxies, defined as a fraction of objects whose predicted position in the SFR– M_* diagram is within 1 dex from the true position (estimated from spectroscopy), ranges from ~60% and up to 88% depending on the typology of a galaxy (quiescent or star-forming, morphology, and AGN or non-AGN). In general, the method is at least twice as efficient as a random guess, and it provides a typical scatter of ~0.3–0.4 dex in M_* and ~0.5–0.6 dex in SFR;
- Our tests showed that the performance of the method mainly depends on the colour and morphology of the target galaxy. For instance, the efficiency for galaxies with low and high stellar formation (i.e. quiescent and SFG) is ~78% and ~72% (averaged for all studied distances), while for targets with a relatively featureless and smooth appearance, it is ~68% on average and increases to ~76% for galaxies with spiral and bulge and/or bar structures;
- The presence of a bright nucleus also affects the performance of our method. Its efficiency for galaxies identified as AGN according to the BPT criteria (~77% among all distances on average) is higher than for non-AGN target objects (~71%), which appears to agree with the results of Paper I;
- The analysis of the separate efficiencies in predicting just M_* or SFR reveals a higher efficiency in M_* retrieval (nearly 94%) than for the SFR (nearly 77%); this is expected because the mass mainly depends on the total luminosity, while the SFR is harder to measure, both photometrically and spectroscopically, and is mainly linked to the galaxy colour.

We conclude that while traditional methods for estimating galaxy properties, such as the SFR and M_* , require time-consuming spectroscopic observations or multi-band photometry for the SED fitting, the use of artificial intelligence algorithms represents a viable and faster alternative, but is somewhat less accurate. Considering the results presented by Paper I and in this work, ULISSE emerges as a promising approach for either selecting specific classes of objects (e.g. AGN, quiescent, or star-forming galaxies) and for predicting their properties in current and upcoming wide-field surveys, such as Euclid and the LSST, which target millions of sources every single night.

Acknowledgments. L.D. acknowledges support from research grant 200021_192285 ‘Image data validation for AI systems’ funded by the Swiss National Science Foundation (SNSF). M.P. and O.T. also acknowledge financial support from the agreement ASI-INAF no. 2017-14-H.O. M.P. also acknowledge the financial contribution from PRIN-MIUR 2022 and from the Timedomes grant within the ‘INAF 2023 Finanziamento della Ricerca Fondamentale’. L.D., S.C. and M.B. also acknowledge financial contributions from the agreement ASI-INAF 2018-23-HH.0, Euclid ESA mission – Phase D. Funding for the Sloan Digital Sky Survey V has been provided by the Alfred P. Sloan Foundation, the Heising-Simons Foundation, the National Science Foundation, and the Participating Institutions. SDSS acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. SDSS telescopes are located at Apache Point Observatory, funded by the Astrophysical Research Consortium and operated by New Mexico State University, and at Las Campanas Observatory, operated by the Carnegie Institution for Science. The SDSS web site is www.sdss.org. SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration, including Caltech, The Carnegie Institution for Science, Chilean National Time Allocation Committee (CNTAC) ratified researchers, The Flatiron Institute, the Gotham Participation Group, Harvard University, Heidelberg University, The Johns Hopkins University, L’Ecole polytechnique fédérale de Lausanne (EPFL), Leibniz-Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Extraterrestrische Physik (MPE), Nanjing University, National Astronomical Observatories of China (NAOC), New Mexico State University, The Ohio State University, Pennsylvania State University, Smithsonian Astrophysical Observatory, Space Telescope Science Institute (STScI), the Stellar Astrophysics Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Illinois at Urbana-Champaign, University of Toronto, University of Utah, University of Virginia, Yale University, and Yunnan University. This research has made use of Galaxy Zoo classifications (<https://www.zooniverse.org/projects/zookeeper/galaxy-zoo/>), so the authors thank the Galaxy Zoo team and the tens of thousands of volunteers who have contributed to the project.

References

- Aihara, H., AlSayyad, Y., Ando, M., et al. 2019, *PASJ*, 71, 114
Aird, J., Coil, A. L., & Georgakakis, A. 2017, *MNRAS*, 465, 3390
Álvarez-Márquez, J., Colina, L., Marques-Chaves, R., et al. 2019, *A&A*, 629, A9
Awang Iskandar, D. N. F., Zijlstra, A. A., McDonald, I., et al. 2020, *Galaxies*, 8, 88
Baldwin, J. A., Phillips, M. M., & Terlevich, R. 1981, *PASP*, 93, 5
Baron, D. 2019, arXiv e-prints, [arXiv:1904.07248]
Baron, D., & Poznanski, D. 2017, *MNRAS*, 465, 4530
Bishop, C. M. 2006, *Pattern Recognition and Machine Learning* (Springer)
Bolzonella, M., Miralles, J. M., & Pelló, R. 2000, *A&A*, 363, 476
Bonjean, V., Aghanim, N., Salomé, P., et al. 2019, *A&A*, 622, A137
Brescia, M., Cavuoti, S., D’Abrusco, R., Longo, G., & Mercurio, A. 2013, *ApJ*, 772, 140
Brinchmann, J., Charlot, S., White, S., et al. 2004, *MNRAS*, 351, 1151
Bruzual, G., & Charlot, S. 2003, *MNRAS*, 344, 1000
Calzetti, D. 2013, in *Secular Evolution of Galaxies*, eds. J. Falcón-Barroso, & J. H. Knapen, 419
Carilli, C. L., & Walter, F. 2013, *ARA&A*, 51, 105
Castro-Ginard, A., Jordi, C., Luri, X., et al. 2018, *A&A*, 618, A59
Cavuoti, S., Amaro, V., Brescia, M., et al. 2017, *MNRAS*, 465, 1959
Cavuoti, S., De Cicco, D., Doorenbos, L., et al. 2024, *A&A*, 687, A246
Chu, J., Tang, H., Xu, D., Lu, S., & Long, R. 2024, *MNRAS*, 528, 6354
Conroy, C. 2013, *ARA&A*, 51, 393
Conselice, C. J. 2014, *ARA&A*, 52, 291
Dark Energy Survey Collaboration (Abbott, T., et al.) 2016, *MNRAS*, 460, 1270
de Jong, J. T. A., Verdoes Kleijn, G. A., Boxhoorn, D. R., et al. 2015, *A&A*, 582, A62
Delli Veneri, M., Cavuoti, S., Brescia, M., Longo, G., & Riccio, G. 2019, *MNRAS*, 486, 1377
Deng, J., Dong, W., Socher, R., et al. 2009, in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 248
Ding, Y., Sohn, J. H., Kawczynski, M. G., et al. 2019, *Radiology*, 290, 456
D’Isanto, A., & Polsterer, K. L. 2018, *A&A*, 609, A111
Domínguez Sánchez, H., Bernardi, M., & Huertas-Company, M. 2023, *Memorie della Società Astronomica Italiana*, 94, 82
Doorenbos, L., Torbaniuk, O., Cavuoti, S., et al. 2022, *A&A*, 666, A171
Esteva, A., Kuprel, B., Novoa, R. A., et al. 2017, *Nature*, 542, 115
Euclid Collaboration (Desprez, G., et al.) 2020, *A&A*, 644, A31
Euclid Collaboration (Scaramella, R., et al.) 2022, *A&A*, 662, A112

- Euclid Collaboration (Kovačić, I., et al.) 2025, *A&A*, 695, A284
- Euclid Collaboration (Mellier, Y., et al.) 2025, *A&A*, 697, A1
- Fabian, A. C. 2012, *ARA&A*, 50, 455
- Fitzpatrick, E. L. 1999, *PASP*, 111, 63
- Fluke, C. J., & Jacobs, C. 2020, *WIREs Data Mining Knowledge Discov.*, 10, e1349
- Förster Schreiber, N. M., & Wuyts, S. 2020, *ARA&A*, 58, 661
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, *Deep Learning* (MIT Press)
- Green, J., Schechter, P., Baltay, C., et al. 2012, arXiv e-prints, [arXiv:1208.4012]
- Heckman, T. M., & Best, P. N. 2014, *ARA&A*, 52, 589
- Hickox, R. C., & Alexander, D. M. 2018, *ARA&A*, 56, 625
- Hildebrandt, H., Arnouts, S., Capak, P., et al. 2010, *A&A*, 523, A31
- Humphrey, A., Cunha, P. A. C., Paulino-Afonso, A., et al. 2023, *MNRAS*, 520, 305
- Ilbert, O., Arnouts, S., McCracken, H. J., et al. 2006, *A&A*, 457, 841
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
- Kauffmann, G., Heckman, T. M., Tremonti, C., et al. 2003a, *MNRAS*, 346, 1055
- Kauffmann, G., Heckman, T. M., White, S. D. M., et al. 2003b, *MNRAS*, 341, 33
- Kennicutt, R. C., & Evans, N. J. 2012, *ARA&A*, 50, 531
- Kim, D.-W., Protopapas, P., Byun, Y.-I., et al. 2011, *ApJ*, 735, 68
- Kim, Y. J., Bae, J. P., Chung, J.-W., et al. 2021, *Sci. Rep.*, 11, 3605
- Kinson, D. A., Oliveira, J. M., & van Loon, J. T. 2021, *MNRAS*, 507, 5106
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, 51, 511
- Koyama, Y., Kodama, T., Hayashi, M., et al. 2015, *MNRAS*, 453, 879
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proc. IEEE*, 86, 2278
- Lin, Z., & Yan, R. 2024, *A&A*, 691, A201
- Longo, G., Merényi, E., & Tiño, P. 2019, *PASP*, 131, 100101
- Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415
- Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., et al. 2020, *ApJS*, 251, 6
- Maiolino, R., & Mannucci, F. 2019, *A&A Rev.*, 27, 3
- Martinazzo, A., Espadoto, M., & Hirata, N. S. 2021, in *2020 25th International Conference on Pattern Recognition (ICPR)*, IEEE, 4176
- Masters, D., Capak, P., Stern, D., et al. 2015, *ApJ*, 813, 53
- Menegola, A., Fornaciali, M., Pires, R., et al. 2017, in *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, IEEE, 297
- Morganti, R., & Oosterloo, T. 2018, *A&A Rev.*, 26, 4
- Pasquet, J., Bertin, E., Treyer, M., Arnouts, S., & Fouchez, D. 2019, *A&A*, 621, A26
- Pathi, I. M., Soo, J. Y. H., Wee, M. J., et al. 2025, *J. Cosmology Astropart. Phys.*, 2025, 097
- Prima, B., & Bouhorma, M. 2020, *ISPRS – International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4443, 343
- Qin, J., Zheng, X. Z., Wuyts, S., Pan, Z., & Ren, J. 2019, *ApJ*, 886, 28
- Razim, O., Cavuoti, S., Brescia, M., et al. 2021, *MNRAS*, 507, 5034
- Salim, S., Rich, R. M., Charlot, S., et al. 2007, *ApJS*, 173, 267
- Schaefer, C., Geiger, M., Kuntzer, T., & Kneib, J.-P. 2018, *A&A*, 611, A2
- Schmidhuber, J. 2015, *Neural Netw.*, 61, 85
- Soo, J. Y. H., Moraes, B., Joachimi, B., et al. 2018, *MNRAS*, 475, 3613
- Surana, S., Wadadekar, Y., Bait, O., & Bhosale, H. 2020, *MNRAS*, 493, 4808
- Tan, M., & Le, Q. 2019, in *International Conference on Machine Learning*, PMLR, 6105
- Torbaniuk, O., Paolillo, M., D’Abrusco, R., et al. 2024, *MNRAS*, 527, 12091
- Tremonti, C. A., Heckman, T. M., Kauffmann, G., et al. 2004, *ApJ*, 613, 898
- Wang, Y., Yao, Q., Kwok, J. T., & Ni, L. M. 2020, *ACM Comput. Surv.*, 53, 1
- Way, M. J., Foster, L. V., Gazis, P. R., & Srivastava, A. N. 2009, *ApJ*, 706, 623
- Wenzl, L., Schindler, J.-T., Fan, X., et al. 2021, *AJ*, 162, 72
- Willett, K. W., Lintott, C. J., Bamford, S. P., et al. 2013, *MNRAS*, 435, 2835
- York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, *AJ*, 120, 1579
- Zeraatgari, F. Z., Hafezianzadeh, F., Zhang, Y. X., Mosallanezhad, A., & Zhang, J. Y. 2024, *A&A*, 688, A33
- Zhong, J., Deng, Z., Li, X., et al. 2024, *MNRAS*, 531, 2011
- Zou, M., & Zhong, Y. 2018, *Sensing Imaging*, 19, 6

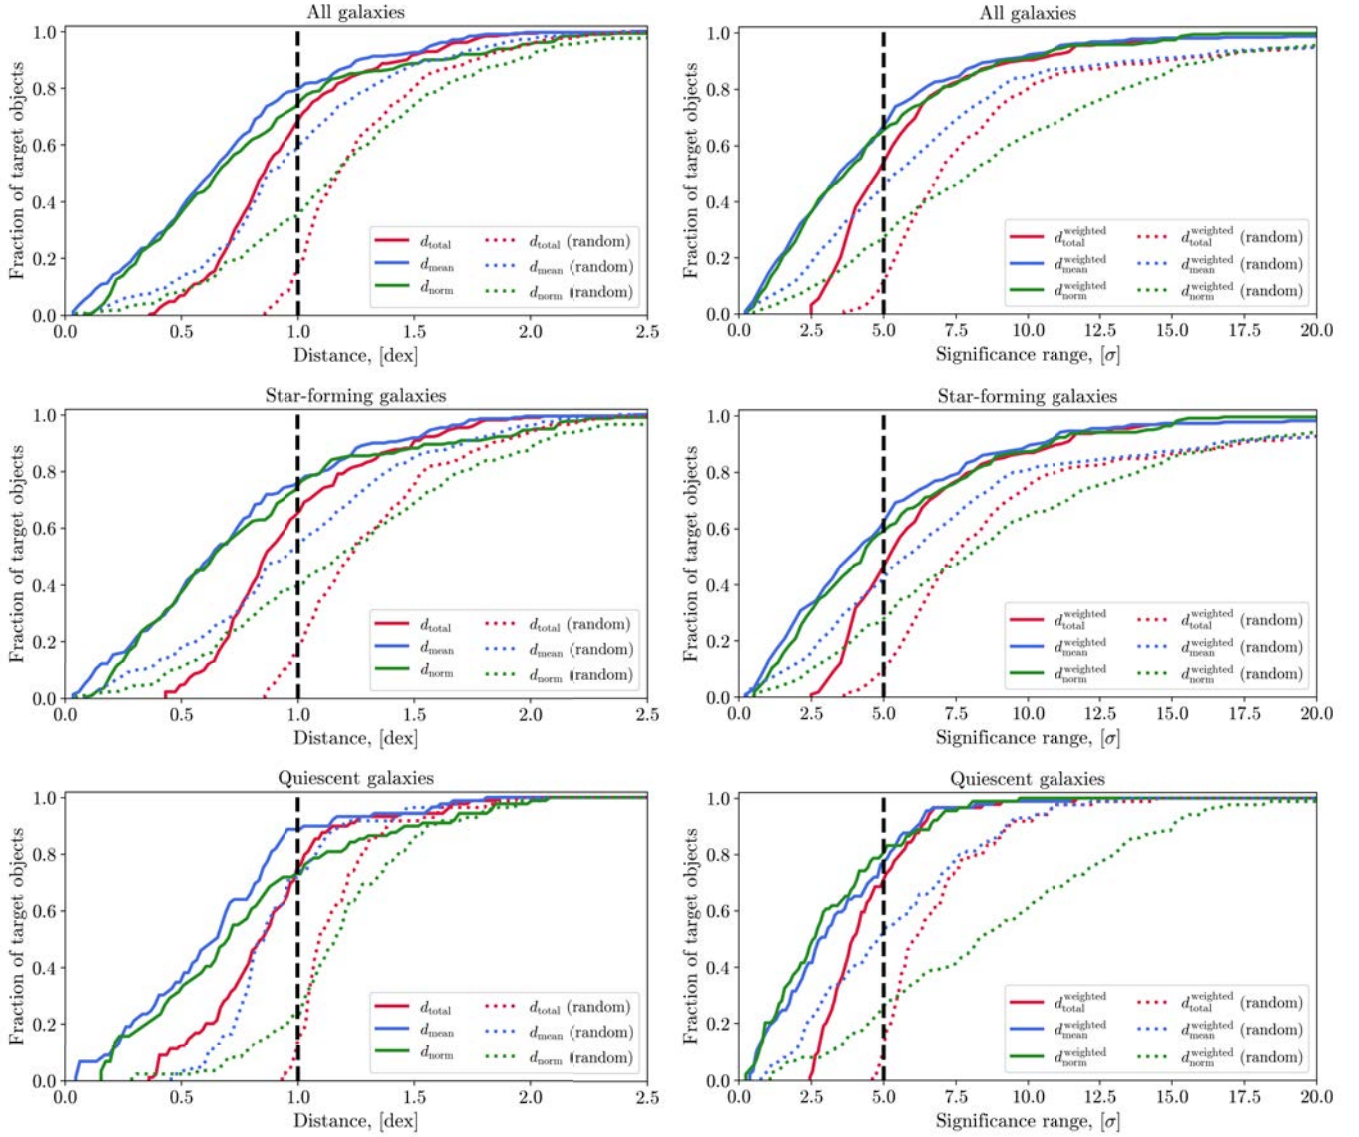
Appendix A: The distribution of the retrieved neighbours in the SFR- M_* diagram with respect to the target objects of different classes.


Fig. A.1. *Left panel:* The distribution of the total (d_{total} by red colour), mean (d_{mean} by blue colour) and normalised mean (d_{norm} by green colour) distances for all, star-forming and quiescent subsamples of target objects. The black dashed line represents the cut at 1 dex used to define ULISSE efficiency in retrieving galaxy properties (see description in Section 4.3). *Right panel:* The similar distribution as on the left panel, but for the total, mean and normalised mean distances weighted for the uncertainties of SFR and M_* parameters of the target objects. The black dashed line is the cut at 5σ .

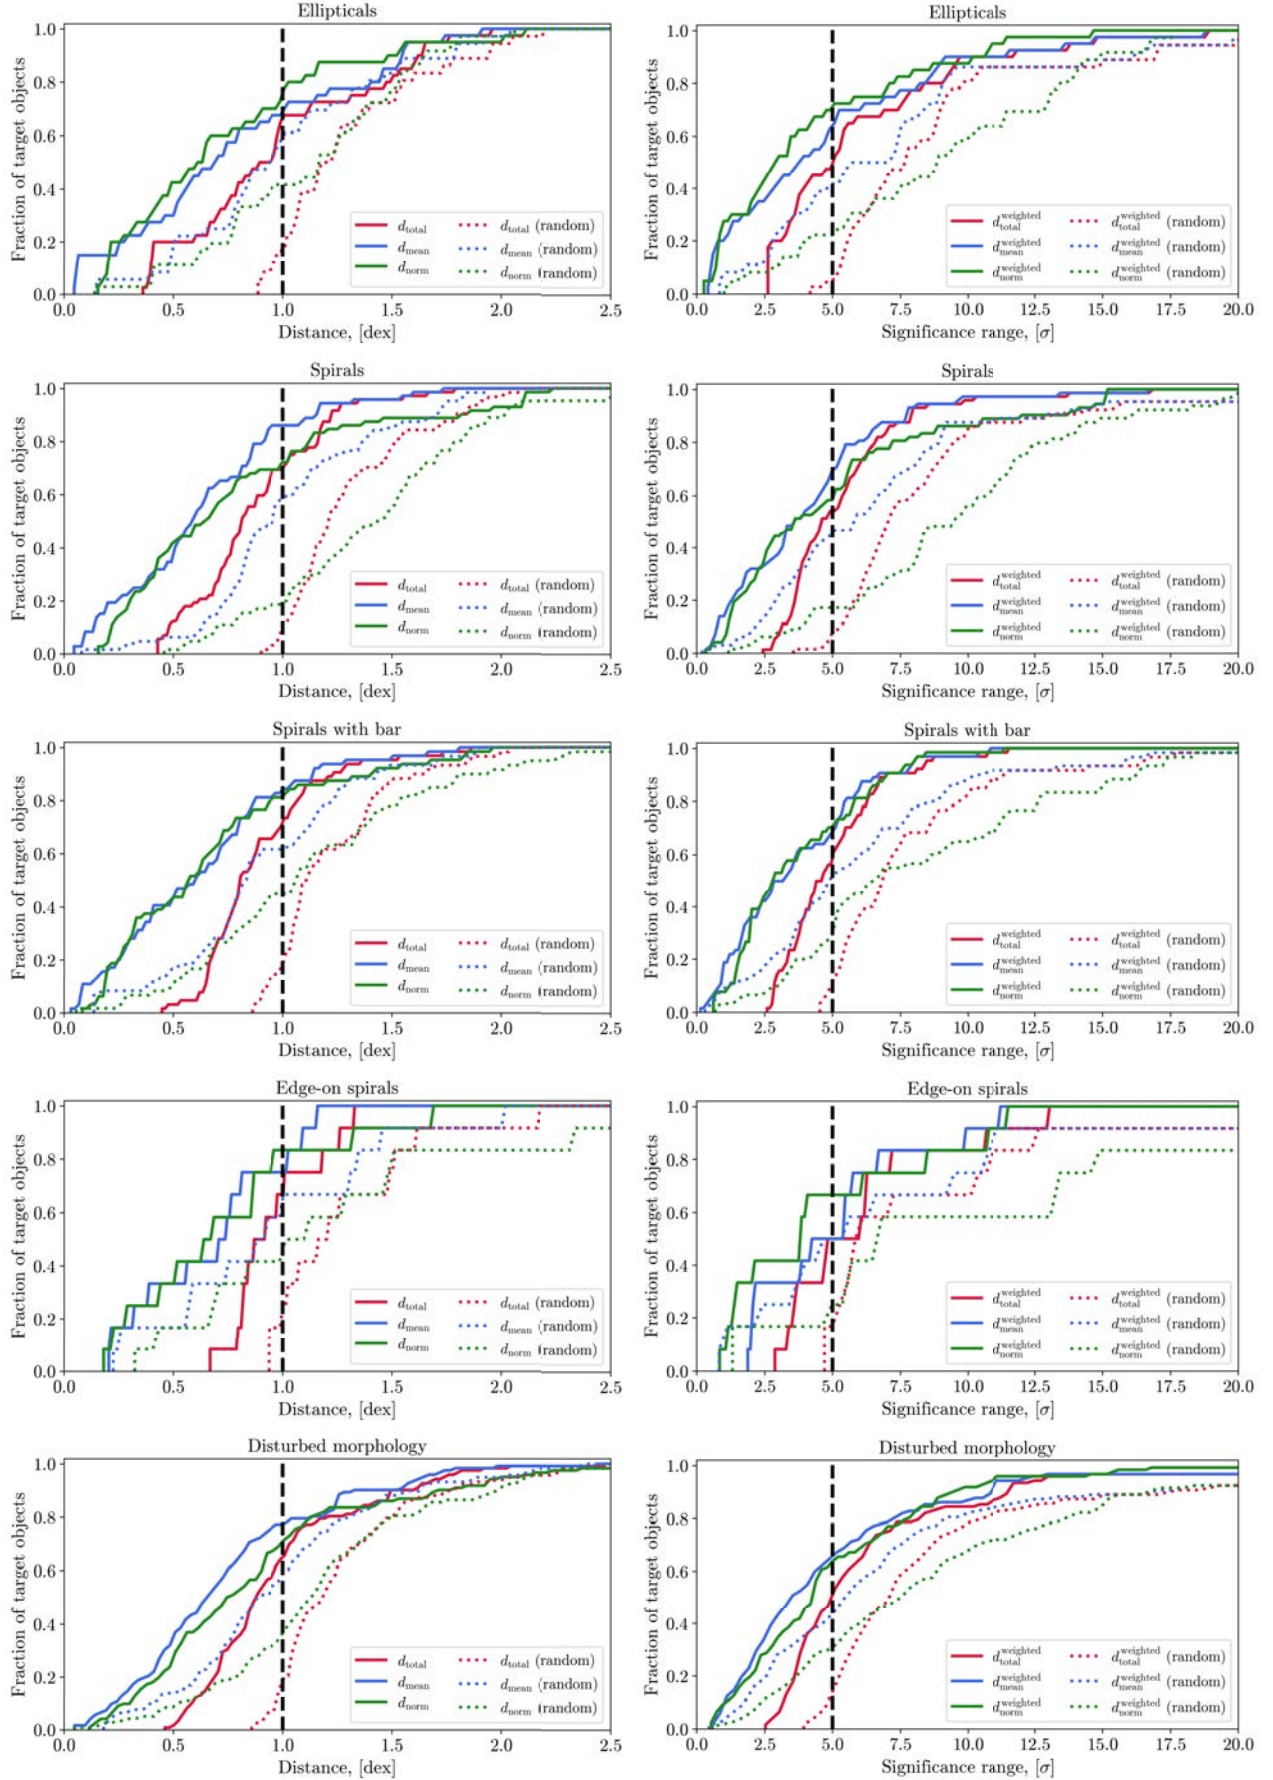


Fig. A.2. The same as in Fig. A.1 for subsamples of target objects with elliptical (E), spiral with bulge (S) and bar (SB) structure, spiral edge-on (Se) and with E(d), S(d), SB(d), Se(d) ‘disturbed’ morphology.

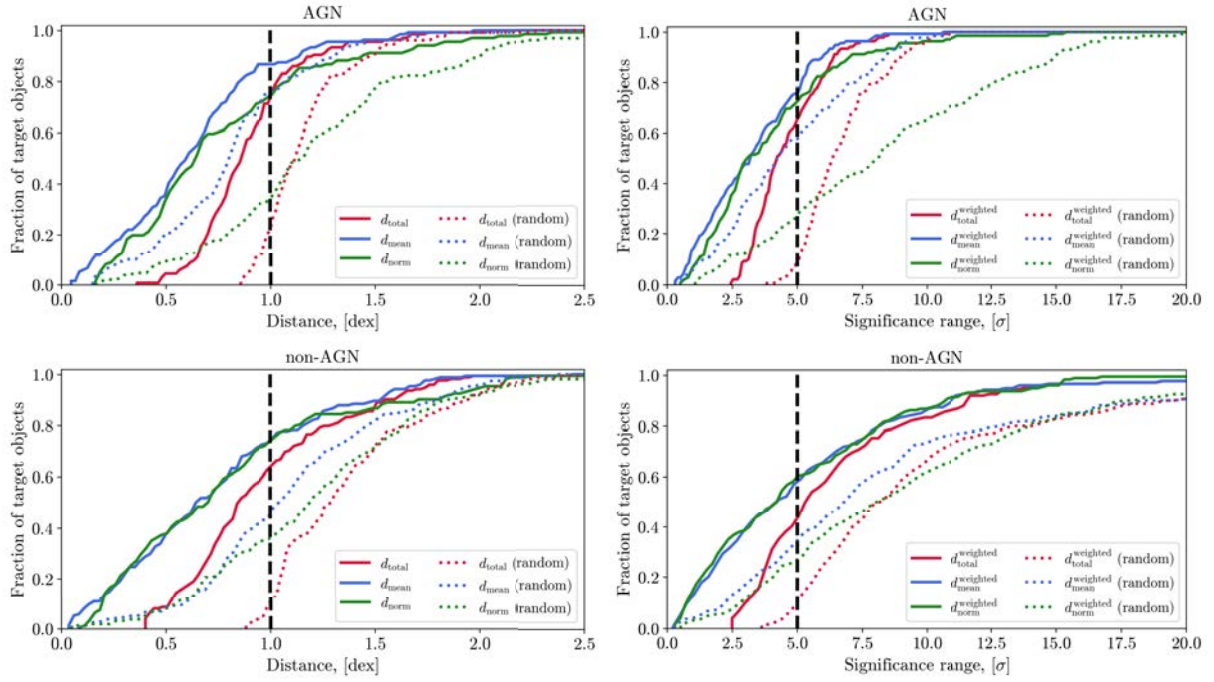


Fig. A.3. The same as in Fig. A.1 for target objects identified AGN or non-AGN according to the BPT selection criteria.

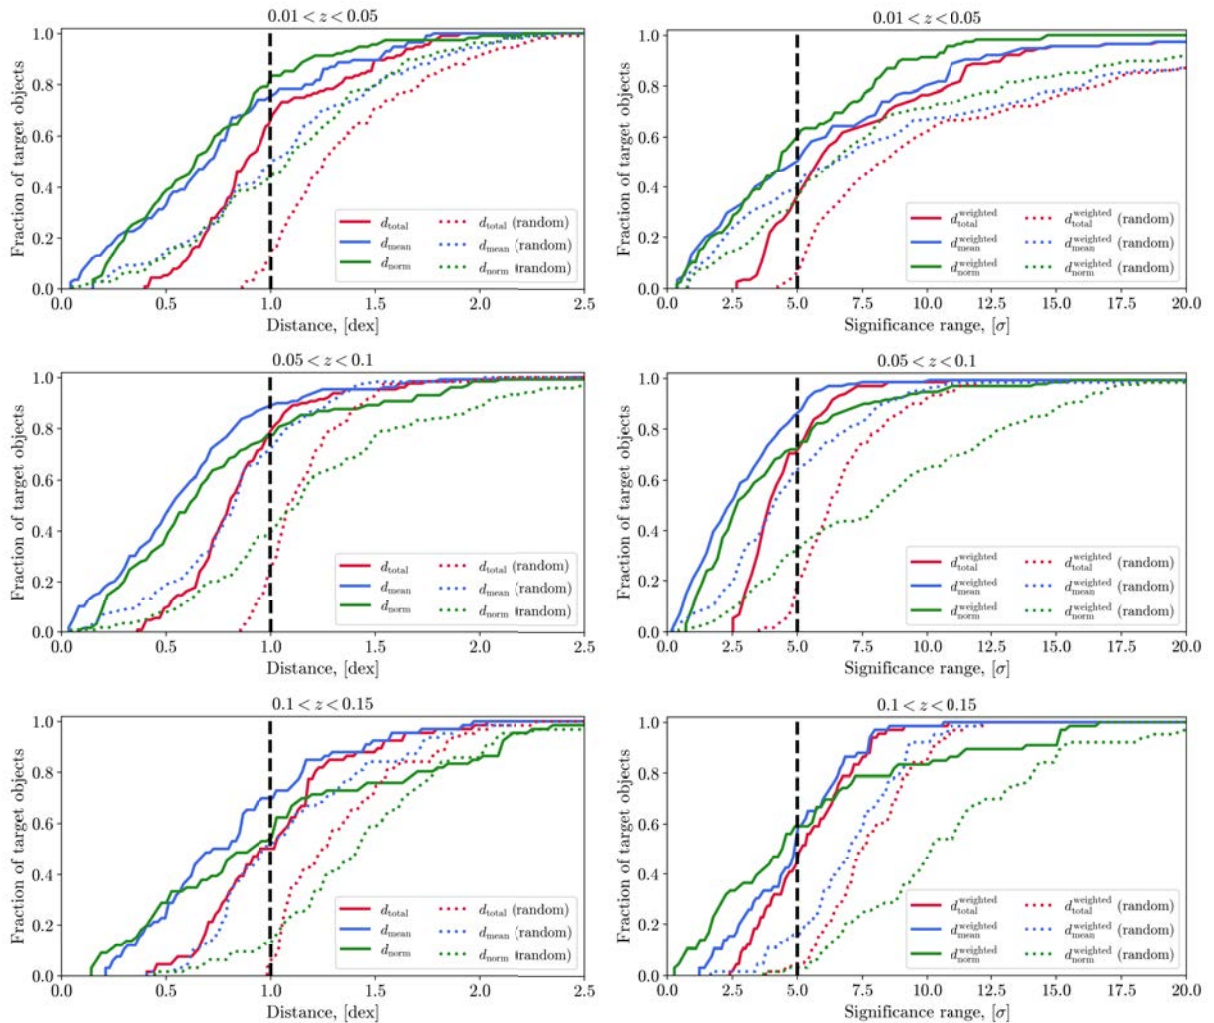


Fig. A.4. The same as in Fig. A.1 for target objects within three redshift ranges.

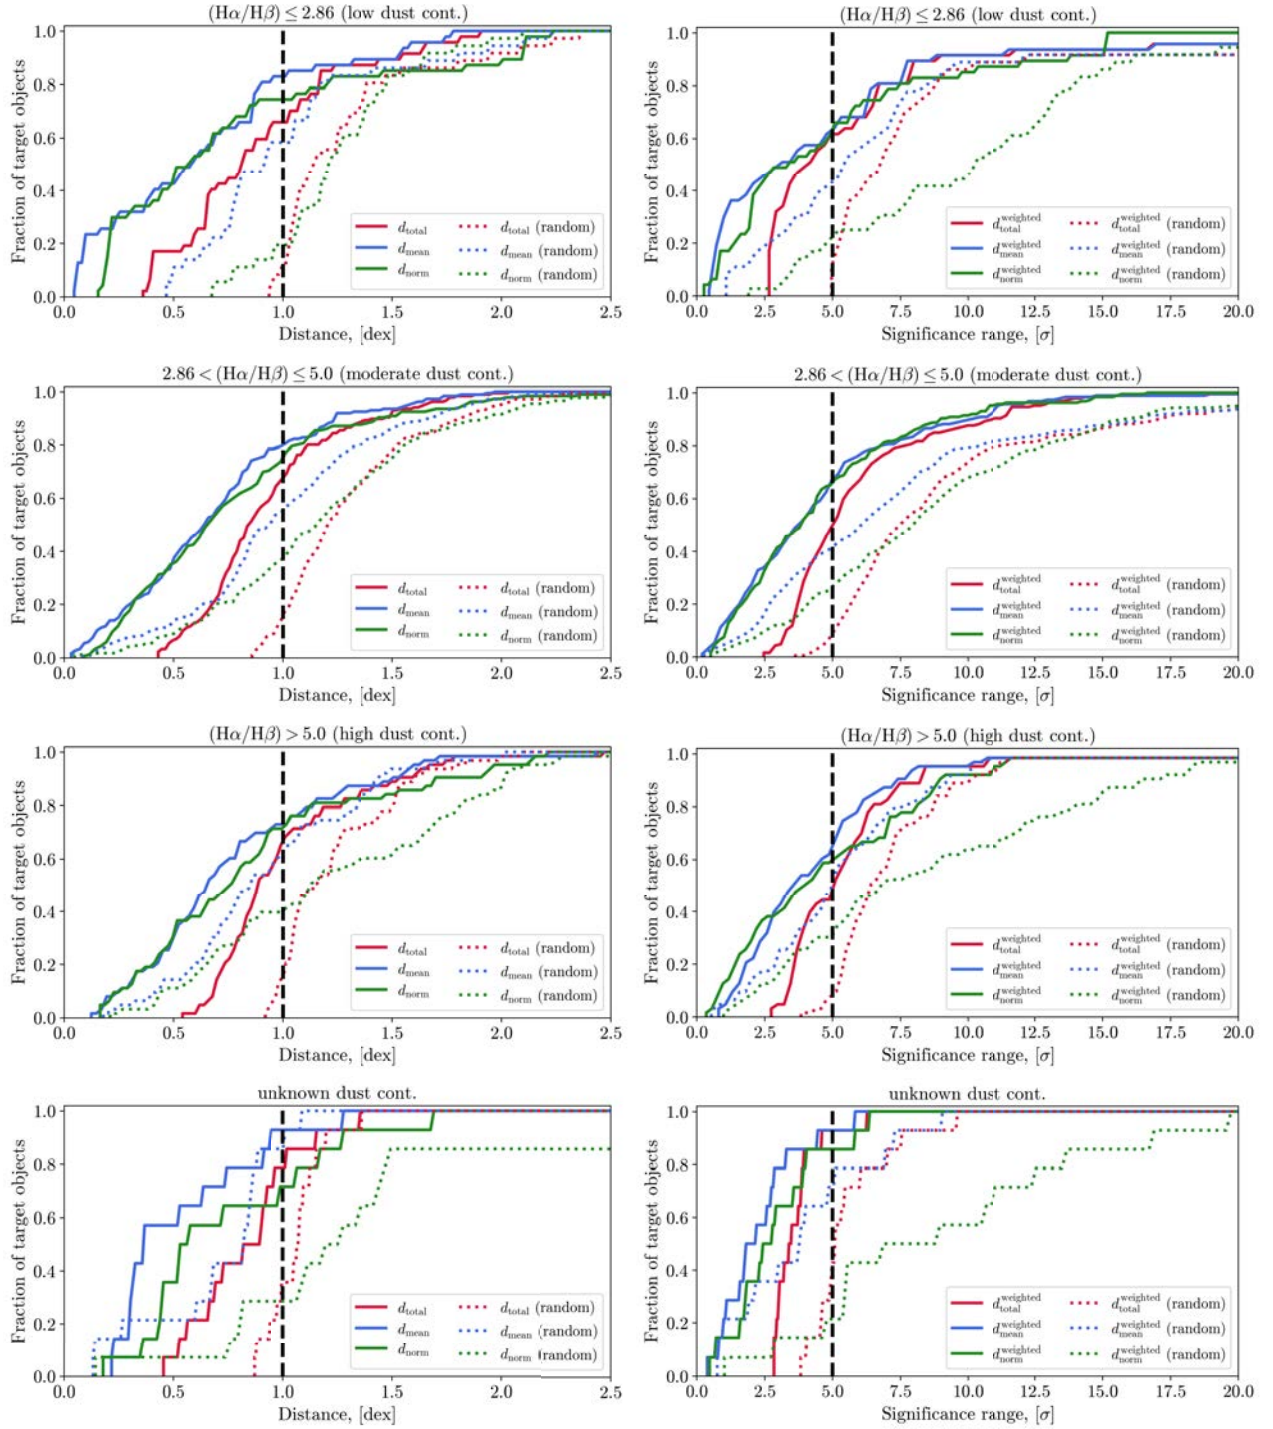


Fig. A.5. The same as in Fig. A.1 for target objects with low ($H\alpha/H\beta \leq 2.86$), moderate ($2.86 < H\alpha/H\beta \leq 5.0$), high ($H\alpha/H\beta > 5.0$) and unknown dust content.

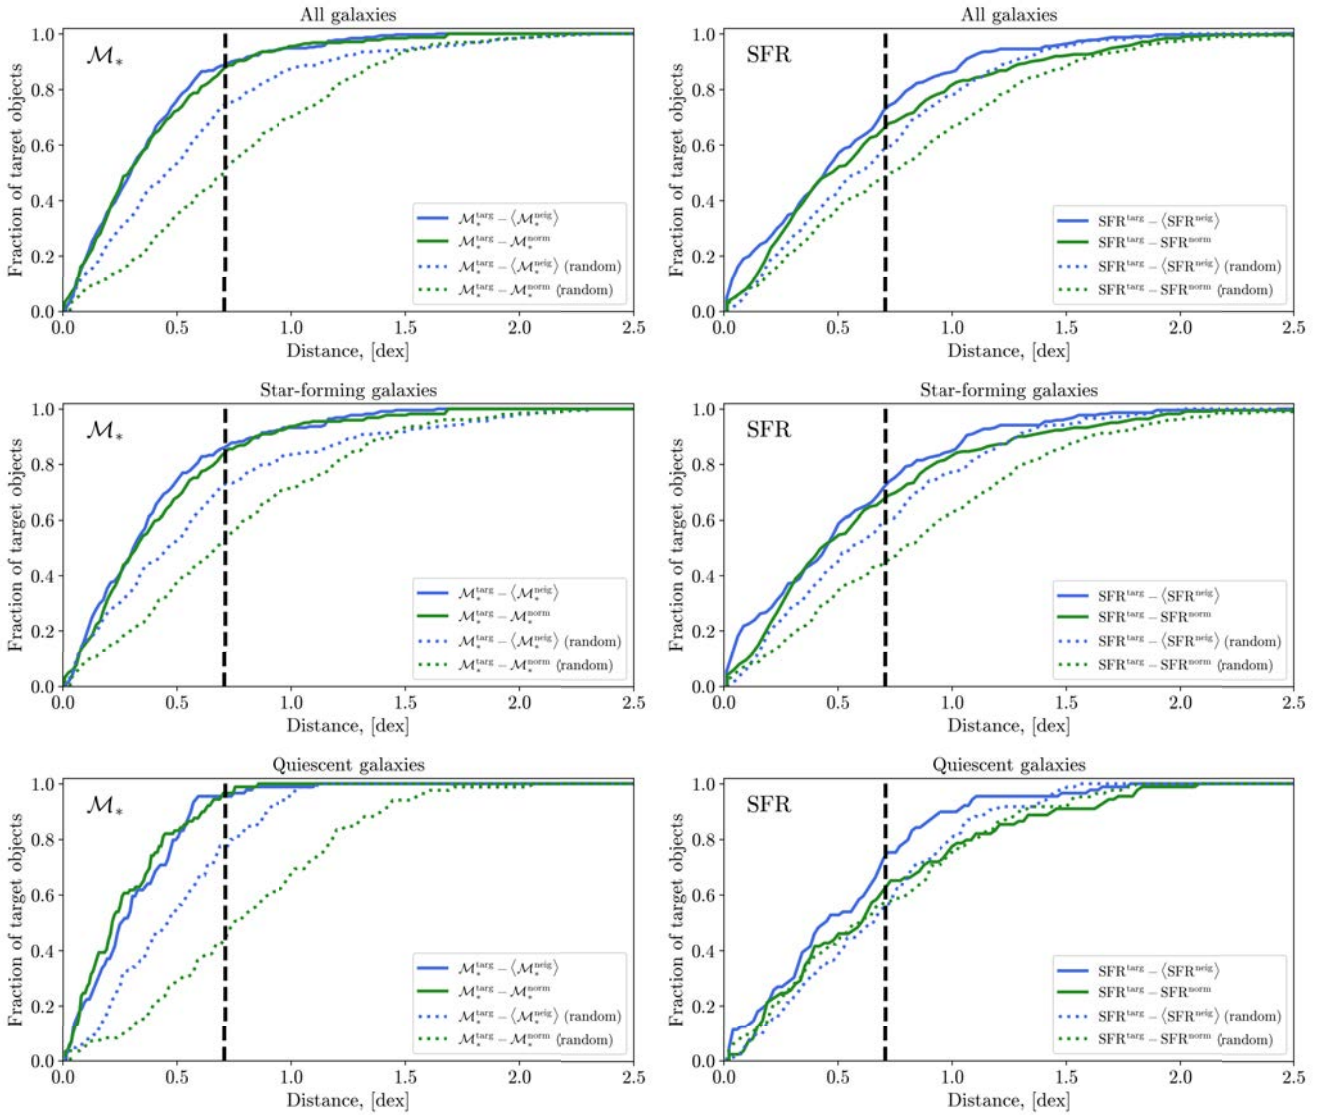
Appendix B: The distribution of the retrieved neighbours in the SFR (or \mathcal{M}_*) space with respect to the target objects.


Fig. B.1. *Left panel:* The distributions of distances between the mean and the normalised mean \mathcal{M}_* for the retrieved neighbours and the target object, i.e. $\mathcal{M}_*^{\text{larg}} - \langle \mathcal{M}_*^{\text{neig}} \rangle$ (by blue colour) and $\mathcal{M}_*^{\text{larg}} - \mathcal{M}_*^{\text{norm}}$ (by green colour), for all, star-forming and quiescent subsamples of target objects. The black dashed line represents the cut at $1 \text{ dex} / \sqrt{2}$ used to define ULISSE efficiency (see description in Section 4.3). *Right panel:* The same as on the left panel, but for $\text{SFR}^{\text{larg}} - \langle \text{SFR}^{\text{neig}} \rangle$ and $\text{SFR}^{\text{larg}} - \text{SFR}^{\text{norm}}$.

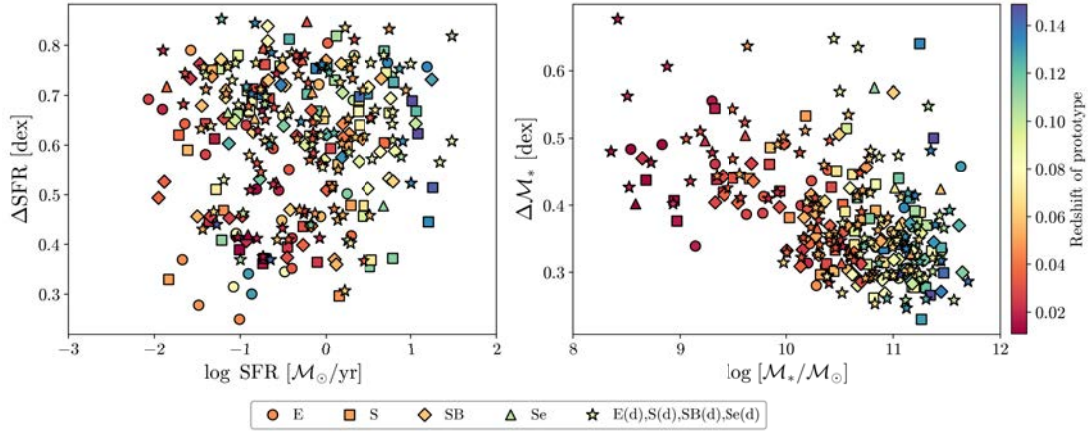
Appendix C: Statistical properties of neighbour distributions retrieved by ULISSE


Fig. C.1. Scatter of retrieved neighbours in SFR (ΔSFR , *left panel*) and M_* (ΔM_* , *right panel*) as a function of the target properties for all 290 target objects. Marker shapes indicate different morphological classes based on the GZ2 catalogue, and the colour scale reflects the redshift of the target object.

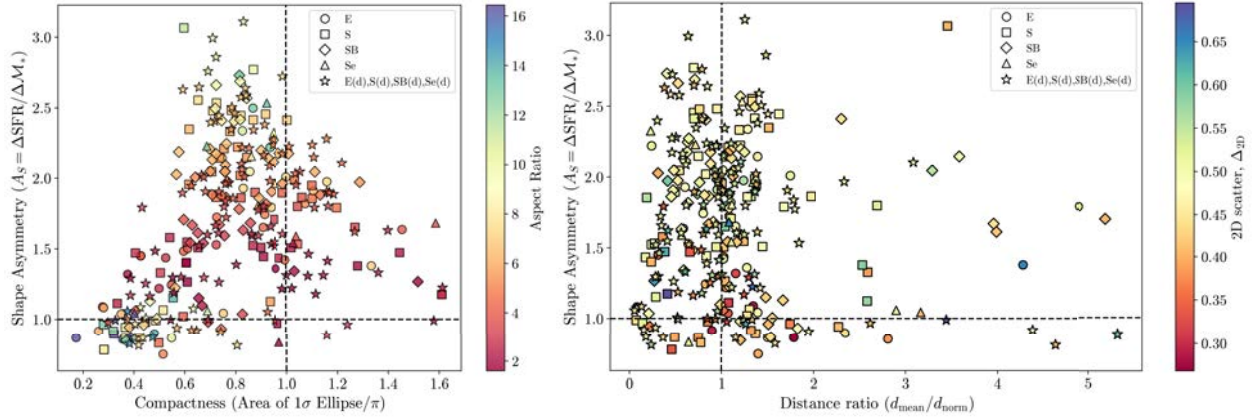


Fig. C.2. *Left panel:* Shape asymmetry A_S versus compactness C of the neighbour distributions retrieved by ULISSE for all 290 target objects. The colour scale represents the aspect ratio \mathcal{R} of each corresponding neighbour distribution. The black dashed lines corresponds to key values: compactness $C = 1$ separates compact ($C < 1$) from dispersed ($C > 1$) distributions, while shape asymmetry $A_S = 1$ corresponds to equal scatter in M_* and SFR. Deviations from unity indicate dominant spread along either M_* ($A_S < 1$) or SFR ($A_S > 1$). *Right panel:* Shape asymmetry A_S versus the ratio between mean and normalised mean distances of the neighbour distributions. Colour scale shows the two-dimensional scatter Δ_{2D} of the neighbours. The black dashed line indicates $d_{\text{mean}}/d_{\text{norm}} = 1$, i.e. cases where the normalisation to the ‘rarity’ of the target object in the sample has no effect on the mean position of the retrieved neighbours. Values significantly above 1 reflect targets located in sparsely populated regions of parameter space, where d_{mean} is biased due to the limited availability of similar objects. Marker symbols in both panels indicate the GZ2 morphological class of each target galaxy.

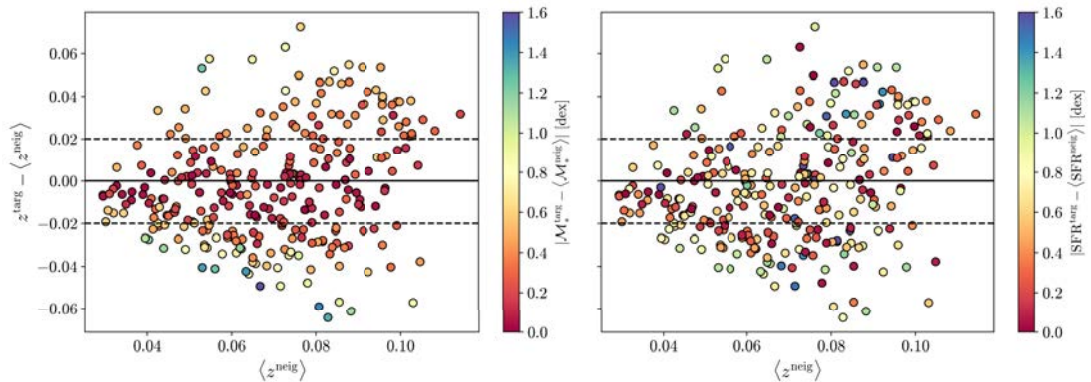


Fig. C.3. Relative redshift offset $z^{\text{targ}} - \langle z^{\text{neig}} \rangle$ as a function of the average redshift of retrieved neighbours. The colour shows the difference between the target and averaged values for neighbours in M_* (*left panel*) and SFR (*right panel*), respectively.