A&A, 707, A152 (2026)
https://doi.org/10.1051/0004-6361/202556236
© The Authors 2026

Astronomy & Astrophysics

Galactic bars and active galactic nucleus fuelling in the second half of cosmic history

A. La Marca[1,2,⋆], M. T. Nardone[2,3], L. Wang[1,2], B. Margalef-Bentabol[1], S. Kruk[3], and S. C. Trager[2]

[1] SRON Netherlands Institute for Space Research, Landleven 12, 9747 AD Groningen, The Netherlands
[2] Kapteyn Astronomical Institute, University of Groningen, Postbus 800, 9700 AV Groningen, The Netherlands
[3] European Space Agency (ESA), European Space Astronomy Centre (ESAC), Camino Bajo del Castillo s/n, 28692 Villaneuva de la Cañada, Madrid, Spain

Received 3 July 2025 / Accepted 8 January 2026

ABSTRACT

We investigate the role of galactic bars in fuelling and triggering active galactic nuclei (AGN) in disc galaxies up to $z \sim 0.8$. We utilised a deep learning model, fine-tuned on Galaxy Zoo volunteer classifications, to identify (strongly and weakly) barred and unbarred disc galaxies in Hyper Suprime-Cam Subaru Strategic Program *i*-band images. We selected AGN using three independent diagnostics: mid-infrared colours, X-ray detections, and spectral energy distribution (SED) fitting. The SED analysis, performed using CIGALE, quantifies the relative AGN contribution to the total galaxy luminosity ($f_{\rm AGN}$) and the AGN luminosity ($L_{\rm disc}$). We assessed the impact of bars by comparing AGN incidence and properties in barred galaxies against carefully constructed redshift-, stellar mass-, and colour-matched unbarred control samples. Our binary AGN classification experiment demonstrates that barred disc galaxies host a higher fraction of AGN compared to their unbarred counterparts, though the significance depends on the AGN selection method, with a more modest excess for SED AGN, and control sample size. This suggests a contributing role for bars in the global AGN budget. The contribution of bars to AGN fuelling appears confined to systems where the AGN has a lower relative contribution to the host galaxy's emission ($f_{\rm AGN} < 0.75$). Crucially, we find a significant dearth of barred disc galaxies hosting AGN with $f_{\rm AGN} > 0.75$, independent of bar strength. Consistent with this, the fraction of barred galaxies among AGN hosts decreases with increasing $L_{\rm disc}$. Combined with previous results, we suggest that bars may contribute to fuelling the population of low-to-moderate-luminosity AGN, but major mergers are the principal mechanism for triggering the most powerful and dominant accretion events.

Key words. galaxies: active – galaxies: evolution – galaxies: spiral – galaxies: structure

1. Introduction

Supermassive black holes (SMBHs) reside at the centres of the most massive galaxies and, when actively accreting matter, manifest as active galactic nuclei (AGN). Understanding the mechanisms that fuel SMBHs and trigger AGN is key for comprehending galaxy-SMBH co-evolution (see Heckman & Best 2014, for a review). While major galaxy mergers were initially considered the primary drivers of AGN activity (Di Matteo et al. 2005; Hopkins et al. 2006, 2008), the role of secular processes, particularly in powering low-to-moderate-luminosity AGN, has gained significant attention (Romeo & Fathi 2015, 2016; Martin et al. 2018; Smethurst et al. 2023). Among these secular mechanisms, large-scale galactic bars are prominent candidates for funnelling gas towards the central SMBH, potentially igniting AGN in disc-dominated galaxies (Sakamoto et al. 1999; Athanassoula 2003; Lin et al. 2013). This scenario is supported by observations showing correlations between bars and central molecular gas concentrations (e.g. Yu et al. 2022). However, this framework is complicated by theoretical works showing that while such bars can efficiently channel gas to the inner hundred parsecs, their influence diminishes at smaller scales (e.g. Shlosman et al. 1989; Hopkins & Quataert 2010).

Disc-dominated galaxies, given their likely merger-free histories (Somerville & Davé 2015), are ideal laboratories for studying secular AGN triggering. Bars are common in local disc galaxies, with observed fractions at optical wavelengths of 30–60% (Barazza et al. 2008; Aguerri et al. 2009; Masters et al. 2011; Méndez-Abreu et al. 2017; Géron et al. 2021; Euclid Collaboration: Huertas-Company et al. 2026), and even higher (up to ~70%) in the infrared (Eskridge et al. 2000; Menéndez-Delmestre et al. 2007). At higher redshifts, the observed bar fraction decreases to 15–20% at $z \approx 2$ (Le Conte et al. 2024; Guo et al. 2025). However, the precise connection between bars and AGN activity remains debated. Some studies find that barred galaxies are more likely to host AGN (Galloway et al. 2015; Alonso et al. 2018; Silva-Lima et al. 2022; Garland et al. 2023, 2024; Marels et al. 2025), while others report no significant correlation or find that differences are reduced after controlling for host galaxy properties such as stellar mass (Lee et al. 2012; Cheung et al. 2013, 2015; Cisternas et al. 2015; Zee et al. 2023).

Discrepancies in the literature may arise from several factors: diverse AGN selection techniques probing different AGN populations (Alexander & Hickox 2012); the strong link between AGN and host galaxy properties (e.g. Kauffmann et al. 2003); and the correlation of bar presence with both stellar mass and galaxy colour, with barred galaxies often being redder than their unbarred counterparts at fixed mass (e.g. Masters et al. 2011; Skibba et al. 2012; Erwin 2018; Kruk et al. 2018). Additionally, the differing timescales of AGN activity (~10–100 Myr; Marconi et al. 2004) and bar persistence (several gigayears;

⋆ Corresponding author: antonio.la.marca.astro@gmail.com

Sellwood 2014; de Sá-Freitas et al. 2023, 2025) can complicate observational connections. Furthermore, AGN feedback or the intense SMBH growth phase itself might disrupt or destroy bar structures (Zee et al. 2023).

In recent work (La Marca et al. 2024, hereafter LM24), we found compelling evidence that major mergers are the primary, if not the only, triggers for the most powerful and dominant AGN. Our study introduced a novel approach, examining the AGN contribution to the total galaxy luminosity (AGN fraction, $f_{\rm AGN}$) and exploring the merger-AGN connection from a continuous perspective. We found that major mergers are the main trigger of AGN-dominated galaxies ($f_{\rm AGN} > 0.80$), while they appear less significant for fuelling less dominant AGN. This led us to hypothesise that secular processes, particularly galactic bars, could be more prominent mechanisms for these AGN.

In this paper, we investigate the galactic bar-AGN relationship in disc-dominated galaxies up to $z \sim 0.8$, using the same parent sample from LM24 but focusing on non-merging systems. To this end, we employed a deep learning (DL) model (`Zoobot`: Walmsley et al. 2023) to identify barred and unbarred disc galaxies in Hyper-Suprime-Cam Subaru Strategic Program (HSC-SSP; Aihara et al. 2018) images. Using rich multi-wavelength data, we selected AGN via mid-infrared (MIR) colours, X-ray detections, and spectral energy distribution (SED) fitting, quantifying $f_{\rm AGN}$ and AGN luminosity. We compared AGN incidence and properties in barred galaxies against carefully matched unbarred control samples to determine the role of bars in AGN fuelling across the second half of cosmic history.

This paper is organised as follows. Section 2 presents a brief summary of the multi-wavelength observations and the characteristics of the galaxy sample used in this work. Section 3 describes in detail our method for detecting galactic bars, selecting AGN, and constructing control samples. Section 4 presents our results; Sect. 5 discusses the impact of major mergers on our analysis; and Sect. 6 summarises our conclusions. Throughout the paper we assume a flat Λ-Cold Dark Matter (λCDM) universe with $\Omega_M = 0.2865$, $\Omega_\Lambda = 0.7135$, and $H_0 = 69.32\,{\rm km\,s^{-1}\,Mpc^{-1}}$ (Hinshaw et al. 2013). Unless otherwise stated, all magnitudes are in the AB system.

2. Data

In this section, we first introduce the multi-wavelength dataset used to construct the galaxy sample. Second, we briefly illustrate how we previously measured the galaxies' physical properties. Then, we describe the Galaxy Zoo (GZ) projects and the HSC-SSP imaging data used for identifying galactic bars.

2.1. Multi-wavelength catalogue

This investigation utilises the galaxy sample constructed in LM24. This parent sample was drawn from a sky area of approximately $65\,{\rm deg}^2$ within the Kilo-Degree Survey (KiDS; de Jong et al. 2013) North-West 2 (N-W2) equatorial field (Kuijken et al. 2019). This field benefits from rich multi-wavelength coverage and was almost entirely observed by the Galaxy And Mass Assembly spectroscopic survey (GAMA; Driver et al. 2011), whose extensive spectroscopic redshifts facilitated the calibration of photometric redshifts for the broader galaxy sample. The catalogue compiled in LM24 incorporated data from numerous surveys. These included: X-ray data (0.2–2.3 keV) from the extended ROentgen Survey with an Imaging Telescope Array (eROSITA; Predehl et al. 2021) as part of the eROSITA Final Equatorial Depth Survey (eFEDS; Brunner et al. 2022); optical (*ugri*) and near-infrared (NIR; $ZYJHK_S$) photometry from the combined VISTA Kilo-degree INfrared Galaxy survey (KiDS-VIKING, hereafter KV; Kuijken et al. 2019; Edge et al. 2013) and the HSC-SSP survey (DR2, *grizy*; Aihara et al. 2018); mid-infrared (MIR) data (centred at 3.4, 4.6, 12, and 22 µm) from the NASA Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010); and far-IR and sub-millimetre data (centred at 100, 160, 250, 350, 500 µm) from the Herschel Astrophysical Terahertz Large Area Survey (H-ATLAS; Valiante et al. 2016).

To construct this parent galaxy sample, LM24 began by selecting all detected sources in the KiDS-N-W2 field from the final nine-band photometric KV catalogue, which had photometric redshifts (Kuijken et al. 2019). These sources were then cross-matched with detections from the other aforementioned surveys. For detailed information on the cross-matching procedure, we refer the reader to LM24. Subsequent data cleaning involved selecting objects within the redshift range $0.1 \leq z \leq 1.0$ and removing clearly identified stars, objects with unreliable photometry, and problematic detections (see Kuijken et al. 2019, for more details). After this cleaning, the sample comprised just over one million galaxies. The analysis in LM24, and consequently this work, focuses on the redshift range $0.1 \leq z \leq 0.76$.

2.2. Galaxy properties

The galaxy catalogue inherited from LM24 contains physical properties derived from SED decomposition. We employed the SED fitting and modelling tool Code Investigating GALaxy Emission (CIGALE; Burgarella et al. 2005; Noll et al. 2009; Boquien et al. 2019) to estimate key properties, including stellar mass ($M_\star$), and AGN fraction ($f_{\rm AGN}$). The $f_{\rm AGN}$ parameter is defined as the AGN contribution to the total galaxy luminosity within a specific rest-frame wavelength range (in this case the MIR range 3–30 µm, denoted as $L_{\rm AGN}/L_{\rm tot}$). In addition to these SED-derived properties, the LM24 catalogue also includes AGN selected using three distinct diagnostics (detailed in Sect. 3.2) and classifications for major mergers. The following paragraphs summarise the main aspects of the SED fitting procedure employed in LM24.

We utilised `X-CIGALE` version 2022.1 (Yang et al. 2020, 2022) for the derivation of physical properties. The star formation history was modelled using a delayed-τ component plus an optional recent (1–150 Myr) exponential starburst. The delayed component models the bulk of the stellar population (large τ values for late-type galaxies, small τ for early-types), while the exponential starburst provides some flexibility to the module, allowing it to represent the latest episode of star formation (e.g. Małek et al. 2018). A Bruzual & Charlot (2003) single stellar population model was adopted, assuming a Chabrier (2003) initial mass function and solar metallicity. For dust attenuation, we selected the Calzetti et al. (2000) law, and for the dust emission component, we chose the Draine et al. (2014) models. The SKIRTOR model (Stalevski et al. 2012, 2016) was used as the AGN template. This model assumes a clumpy two-phase dusty torus with a flared disc geometry, consisting of high-density clumps embedded in a lower-density medium. Crucially, `X-CIGALE` incorporates an X-ray module, enabling the simultaneous fitting of X-ray emission from both AGN and host galaxy components (such as hot gas and X-ray binaries). For a comprehensive list of all parameters used in the `X-CIGALE` configuration, we refer the reader to Table 2 of LM24.

To ensure robust constraints on the $f_{\rm AGN}$ parameter, particularly in the MIR regime where AGN emission can dominate,

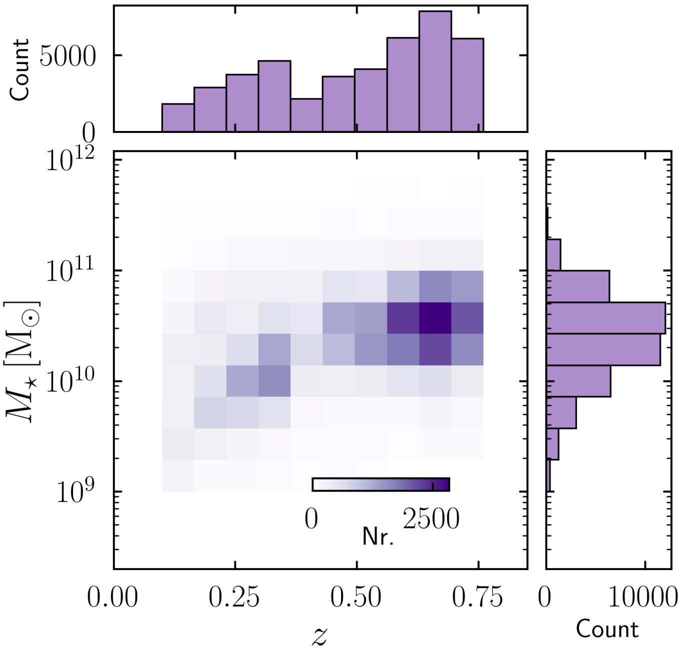

Fig. 1. 2D histogram (linear scaling) representation of the stellar mass–redshift distribution for the LM24 parent sample used in this work. The marginal histograms display the individual z and $M_\star$ distributions.

the sample selection in LM24 was restricted to galaxies with significant MIR detections. Specifically, we selected objects with a signal-to-noise ratio (S/N) > 1 in both the WISE W1 (3.4 µm) and W2 (4.6 µm) bands, and additionally required $S/N > 3$ in either W1 or W2. Furthermore, to avoid saturation issues, only sources fainter than the saturation limits of $W1 > 8$ mag and $W2 > 7$ mag (Vega system) were included. This MIR pre-selection ensures that CIGALE has sufficient information to reliably model the AGN component. Furthermore, to maintain a sample of reliable SED fits, only galaxies with a reduced chi-squared ($\chi_{\rm red}2$) < 5 from the fitting process were retained.

In LM24, we constructed stellar mass–limited samples within three redshift bins:

- $0.1 \leq z < 0.31$, with $M_\star \geq 10^9\ M_\odot$;
- $0.31 \leq z < 0.52$, with $M_\star \geq 10^9\ M_\odot$;
- $0.52 \leq z < 0.76$, with $M_\star \geq 2.5 \cdot 10^9\ M_\odot$.

These mass limits were adopted from Wright et al. (2019) to ensure mass completeness. After applying this mass selection, the final mass-complete sample comprises 9 880, 10 364, and 22 593 galaxies in redshift bins 1, 2, and 3, respectively. Figure 1 shows the redshift and stellar mass distributions of the final mass-complete sample. The main panel displays stellar mass versus redshift, while the marginal histograms show the overall redshift distribution and stellar mass distribution for the sample used in this study.

It is important to note that for the primary analysis presented in this paper (Sect. 4), we explicitly exclude galaxies classified as major mergers in LM24. This exclusion ensures that our investigation focuses on the influence of bars in relatively isolated, disc-dominated systems, reducing potential contamination from merger-induced AGN activity. The impact of including major mergers on our findings is explored separately in Sect. 5.

2.3. HSC-SSP imaging and Galaxy Zoo classifications for bar identification

To visually identify galactic bars, we leveraged citizen science classifications from multiple GZ projects applied to HSC-SSP survey images. We utilised *i*-band coadded images (which include sky subtraction) from the HSC-SSP Public Data Release 3 (PDR3; Aihara et al. 2022), selected for their generally superior seeing compared to other bands. The HSC-SSP *Wide* survey layer reaches a depth of approximately $i \sim 26$ mag (5σ for point sources) with an average *i*-band seeing of 0.61″ (Aihara et al. 2022). For each galaxy in our parent sample (Sect. 2.2), we downloaded an image cutout with a semi-width and semi-height equal to four times the *i*-band Kron radius, as measured by the HSC-SSP PDR3 pipeline (`i kronflux radius`).

To ensure reliable morphological classification, we performed stringent image quality control. Each selected galaxy cutout was required not to have any of the following critical flags raised in the HSC-SSP PDR3 catalogue: `i pixelflags edge` (source is outside usable exposure area), `i pixelflags bad` (bad pixel in the source footprint), `i pixelflags crcenter` (cosmic ray in the source centre), `i pixelflags saturatedcenter` (saturated pixel in the source centre), `i pixelflags interpolatedcenter` (interpolated pixel in the source centre), `i kronflux flag` (general failure flag for Kron fit), `i kronradius flag bad radius` (bad Kron radius), or `i cmodel flag` (flag set if the final cmodel fit, or any previous fit, failed). We also removed images if their associated mask images contained any of the following flags: `MP BAD` (pixel is physically bad), `MP SAT` (pixel flux exceeded full-well), `MP NO DATA` (pixel has no input data), or `MP UNMASKEDNAN` (a NaN occurred in this pixel in instrument signature removal).

2.3.1. Galaxy Zoo classifications for bar identification

Bar identifications were sourced from four GZ science projects: GZ Hubble (GZH; Willett et al. 2017), GZ2 (Hart et al. 2016), GZ DECaLS (Dark Energy Camera Legacy Survey; Walmsley et al. 2022), and GZ GAMA-KiDS (Holwerda et al. 2024). In these projects, citizen scientists visually classified galaxies by answering a series of questions[1] organised in a decision tree structure (an example relevant to bar identification is shown in Fig. 2; Willett et al. 2013). Most GZ projects provide a ‘weighted’ and an ‘unweighted’ vote fraction for each possible answer; we consistently utilised the ‘unweighted’ vote fractions, as the impact of weighting is generally minimal (Willett et al. 2017). Given potential overlaps between the GZ survey footprints, we established the following priority for sourcing galaxy classifications: GZ DECaLS, GZ GAMA-KiDS, GZH, and finally GZ2. The GZ DECaLS project was prioritised as it employed a modified decision tree specifically aimed at improving the identification of weak bars (Walmsley et al. 2022), which is pertinent to this study.

2.3.2. Defining barred and unbarred galaxies

To construct the training dataset for our bar detection model, we focused on specific answers from the GZ decision trees. To select barred galaxies, we were primarily interested in responses to the question, ‘Is there a sign of a bar feature through the centre of the galaxy?’ (Fig. 2). Volunteers could only answer this question if they first identified the galaxy as disc-like or if it shows any type

[1] More information is available at https://data.galaxyzoo.org/

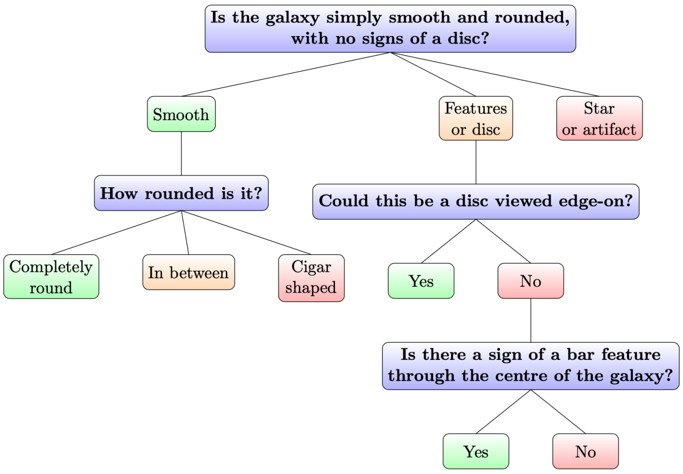

Fig. 2. Simplified representation of the relevant questions from the Galaxy Zoo decision tree (adapted from Willett et al. 2013) used to classify galaxies for the training set.

of feature (i.e. is not smooth), and is not edge-on. We considered only those galaxies for which at least 20 unique users answered the bar or no-bar question. If a galaxy received fewer than 20 votes in its highest-priority GZ project (as defined in Sect. 2.3.1), we consulted the next GZ project in our priority list, ensuring a single, robust bar classification per galaxy. The bar likelihood ($p_{\rm bar}$) for each galaxy represents the fraction of positive ('yes, there is a bar') answers to this question[2].

Previous GZ studies (e.g. Masters et al. 2012; Willett et al. 2013) established thresholds for $p_{\rm bar}$ based on comparisons with expert visual classifications (e.g. Nair & Abraham 2010). A $p_{\rm bar} \geq 0.5$ is generally considered reliable for identifying strongly barred galaxies. Masters et al. (2012) suggested that a likelihood of $0.2 < p_{\rm bar} < 0.5$ often corresponds to weakly barred features, while Willett et al. (2013) proposed $p_{\rm bar} = 0.3$ as an optimal threshold to include both strong and weak bars. To create our training classes, we defined:

– Strongly barred discs: $p_{\rm bar} \geq 0.5$;
– Weakly barred discs: $0.2 < p_{\rm bar} < 0.5$;
– Unbarred discs: $p_{\rm bar} \leq 0.2$.

To effectively train a machine learning algorithm for identifying barred galaxies within a diverse galaxy population, it is essential to provide examples of various morphological types. Therefore, in addition to the unbarred disc galaxies defined by $p_{\rm bar} \leq 0.2$, we also included 'smooth' and 'edge-on' galaxies in our training dataset.

– Smooth galaxies: These were defined as galaxies for which volunteers answered 'smooth' to the initial GZ question, 'Is the galaxy simply smooth and rounded, with no signs of a disc?' (Fig. 2). We further refined this class by excluding galaxies likely to be 'cigar-shaped' (which can be confused with edge-on discs) by requiring $p_{\rm cigar} \leq 0.1$ in the follow-up 'How rounded is it?' question. This ensures the smooth class predominantly contains elliptical and S0-type galaxies.
– Edge-on disc galaxies: These were selected using a high threshold in the 'Could this be a disc viewed edge-on?' question, specifically requiring $p_{\rm edge-on} \geq 0.8$, with at least ten votes cast for this question.

[2] In the GZ DECaLS project, possible answers to the bar question included 'strong bar', 'weak bar', and 'no bar'. For consistency with other GZ projects, we defined $p_{\rm bar}$ as the sum of the vote fractions for 'strong bar' and 'weak bar' in GZ DECaLS.

Table 1. Number of galaxies in the five classes in the final GZ dataset.

Class	N
Strong bar	4105
Weak bar	4390
Unbarred disc	4525
Smooth galaxy	4591
Edge-on galaxy	4236
All samples	21 847

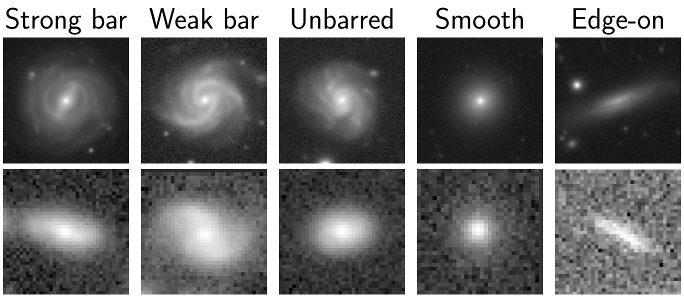

Fig. 3. Examples of GZ training galaxies for each class. Top row: Galaxies with $z < 0.5$. Bottom row: Galaxies at $z \geq 0.5$. Each cutout has a size of 8×8 Kron radius (measured in the i-band). Images are displayed using a logarithmic scaling.

Under these criteria, edge-on galaxies have a median vote count of 17, compared to 30 for the other classes. The inclusion of these distinct classes (strongly barred, weakly barred, unbarred discs, smooth, and edge-on galaxies) ensures our training set captured a comprehensive range of galaxy morphologies relevant to the bar identification task.

To create a balanced training sample, crucial for optimal machine learning model performance, we aimed for a similar number of objects in each of the five defined morphological classes. The 'strongly barred' class contained the fewest galaxies, numbering approximately 5300. Consequently, we limited all other classes to this size by randomly selecting 5300 galaxies from each. During this process, we ensured that no galaxy appeared in more than one class, removing any duplicates. Finally, we downloaded the HSC-SSP i-band image cutouts for these selected galaxies and applied the image quality cuts detailed in Sect. 2.3. The final GZ training dataset consists of 21 847 galaxies with redshifts up to $z \sim 1$, almost equally divided among the five classes: strongly barred, weakly barred, unbarred disc, smooth, and edge-on. The final classifications were sourced from the GZ projects in the following proportions: 51% GZH, 35% GZ DECaLS, 12% GZ GAMA-KiDS, and 2% GZ2. Table 1 reports the precise number of objects in each class. Figure 3 shows examples of galaxies from these different classes at various redshifts, illustrating the visual characteristics captured in our training set.

3. Methodology

In this section, we first describe the DL approach used to identify galactic bars in the HSC-SSP survey images. Secondly, we outline the three different methods employed to select AGN within our sample. Finally, we explain the construction of control samples, essential for robustly assessing the bar-AGN connection.

3.1. Detecting bars with Zoobot

To perform the multi-class morphological classification task defined by the five galaxy types established in Sect. 2.3.2 (strongly and weakly barred, unbarred discs, smooth, and edge-on galaxies), we employed convolutional neural networks (CNNs). CNNs are a class of DL models particularly effective for image analysis, featuring multiple hidden layers capable of automatically learning and extracting hierarchical features from input images (Lecun et al. 1998; Krizhevsky et al. 2012). The output of a CNN is typically a set of scores, one for each predefined class, which are then used for classification. Training a CNN involves passing a large number of labelled images through its architecture and iteratively adjusting the network's weights to minimise the difference between its predictions and the true labels.

We utilised the `Zoobot` Python package (Walmsley et al. 2022), which provides several DL models pre-trained on vast datasets of GZ volunteer responses. For this work, we fine-tuned the 'zoobot-encoder-convnext_nano' pre-trained model. This model architecture is based on the ConvNeXt family of CNNs (Liu et al. 2022), and was originally trained by Walmsley et al. (2023) on approximately 820 000 images and over 100 million volunteer responses from all major GZ campaigns, including GZ DECaLS, GZ2, and GZH. Fine-tuning allowed us to adapt this powerful, general-purpose morphology model to our specific five-class problem using our curated GZ training dataset.

We randomly split our final GZ dataset of 21 847 galaxies (Table 1) into three subsets: 70% for training (15 293 galaxies), 15% for validation (3277 galaxies), and 15% for testing (3277 galaxies). The training set was used to adjust the model weights, the validation set to monitor performance during training and optimise hyperparameters (preventing overfitting to the training data), and the test set for a final evaluation of the model's performance on unseen data. Input images were resized to 224 × 224 pixels. To increase the effective size of our training set and make the model more robust to variations, we implemented data augmentation during the training phase. This involved applying a random horizontal flip (with a 50% probability) and a random rotation by multiples of 90 degrees (also with a 50% probability) to each image in a batch before feeding it to the model.

We employed an early stopping criterion during training, with a patience of ten epochs. This means training was halted if the validation loss did not improve for ten consecutive epochs, preventing overfitting. Hyperparameter optimisation was performed via a grid search over different values for batch size, the number of unfrozen blocks in the pre-trained encoder, and the learning rate. The learning rate decay factor was kept fixed at 0.5. The best performance, as determined by the minimum validation loss, was achieved by unfreezing and training the last four (out of five) blocks of the ConvNeXt encoder, using a batch size of 128, and an initial learning rate of 10^{-3}.

3.1.1. Model predictions and performance

The fine-tuned `Zoobot` model outputs five 'prediction probabilities' for each input galaxy, corresponding to the five morphological classes, which sum to unity. A galaxy was classified as belonging to the class with the highest prediction probability. The overall accuracy of our model – defined as the ratio of correctly classified objects to the total number of objects – is 69.0%, calculated on a balanced test set. The balanced test set was created by randomly selecting an equal number of instances from each true class. To provide a more detailed evaluation of the classifier's performance for each class, we calculated precision, recall, and the F_1-score on the balanced test set. Precision for a given class is the ratio of correctly identified galaxies among all galaxies predicted to belong to that class. Recall (or sensitivity) measures the fraction of correctly identified galaxies among all true instances of that class. The F_1-score is the harmonic mean of precision and recall, providing a single metric that balances both. These metrics for each class, evaluated on the test set, are reported in Table 2. The model's performance is also visualised as a confusion matrix in Fig. 4. Each cell in the confusion matrix shows the raw count of galaxies, the precision (as a percentage of predictions for that class), and the recall (as a percentage of true labels for that class). The diagonal elements highlight how well the model predicts each specific class.

Table 2. Performance of the `Zoobot` model trained on the five-class task.

Class	Precision (%)	Recall (%)	F_1-score (%)
Strong bar	81.41	60.77	69.59
Weak bar	46.12	50.24	48.09
Unbarred disc	52.76	63.96	57.82
Smooth galaxy	81.90	75.76	78.71
Edge-on galaxy	91.93	94.42	93.15

Notes. Precision, recall, and F_1-score for each class.

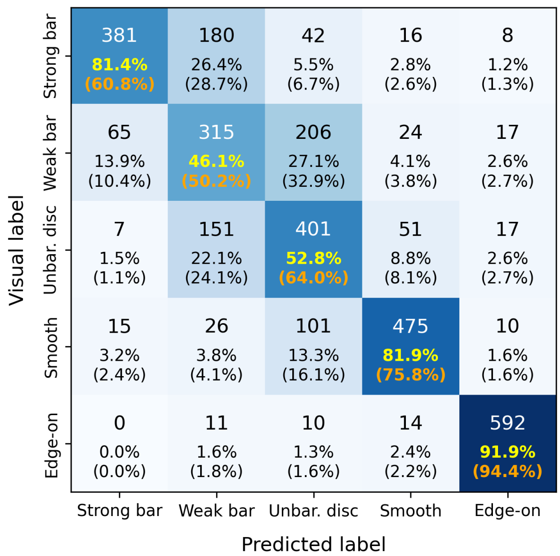

Fig. 4. Confusion matrix for the fine-tuned `Zoobot` model, evaluated on the balanced test set. The matrix is colour-coded relative to the total number of galaxies in each true morphological class (rows). Each cell displays the raw count of galaxies, the column-wise percentage (precision, yellow boldface text on the diagonal), and the row-wise percentage (recall, orange text in brackets on the diagonal).

Overall, the model demonstrates reasonable performance. As seen in Table 2 and Fig. 4, the 'Edge-on galaxy' class achieves the highest F_1-score (93.15%), with high precision (91.93%) and recall (94.42%), indicating these are relatively easy to distinguish. 'Strong bar' and 'Smooth galaxy' classes also show good F_1-scores (69.59% and 78.71%, respectively). The 'Unbarred disc' class has an F_1-score of 57.82%. The 'Weak bar' class proves to be the most challenging, with an F_1-score of 48.09%.

This is likely due to the inherent subtlety of weak bar features, which are difficult to distinguish from other central structures (such as prominent bulges or even some unbarred discs features) leading to fuzzy distinctions and consequently lower classification performance. The confusion matrix (Fig. 4) further illustrates these trends, showing for example that a notable fraction of true 'Weak bar' galaxies are misclassified as 'Unbarred disc' or vice versa.

3.1.2. Selecting barred and unbarred discs

After training and evaluating the `Zoobot` model, we applied it to the full mass-complete galaxy sample (described in Sect. 2.2) to identify barred and unbarred disc galaxies for our scientific analysis of the bar-AGN connection. The criteria for this selection were optimised to ensure high purity for the defined classes, minimising contamination that could affect our statistical conclusions. We classified galaxies in our sample based on the `Zoobot` output probabilities ($P_{\rm Strong\,bar}$, $P_{\rm Weak\,bar}$, $P_{\rm Unbar.\,disc}$, $P_{\rm Smooth}$, $P_{\rm Edge-on}$) as follows. A galaxy was classified as barred ($A_{\rm bar}$) if the sum of probabilities for being either strongly or weakly barred exceeded 0.65, $P_{\rm Bar} = P_{\rm Strong\,bar} + P_{\rm Weak\,bar} \geq 0.65$. Within the $A_{\rm bar}$ sample, a galaxy was further classified as strongly barred ($S_{\rm bar}$) if $P_{\rm Strong\,bar} \geq P_{\rm Weak\,bar}$. Conversely, a barred galaxy was classified as weakly barred ($W_{\rm bar}$) if $P_{\rm Strong\,bar} < P_{\rm Weak\,bar}$.

A galaxy was classified as an unbarred disc ($U_{\rm bar}$) if its probability of being an unbarred disc satisfied $P_{\rm Unbar.\,disc} > 0.45$ and this probability exceeded the next highest probability among the other four classes. Specifically, we required $\Delta P = P_{\rm Unbar.\,disc} - P_{\rm second\,highest} > 0.1$, where $P_{\rm second\,highest}$ is the probability of the class with the second-highest score from the `Zoobot` output. This additional criterion was imposed to create a sample of unbarred discs with high purity, minimising contamination from galaxies that might have ambiguous classifications (e.g. potentially weak bars or smooth galaxies with disc-like features). Examples of galaxies classified as strongly barred, weakly barred, and unbarred discs using these criteria are shown in Appendix A (Fig. A.1).

To assess the impact of these selection criteria on classification performance, we evaluated them on the same balanced test set used in Sect. 3.1. The results are presented as confusion matrices in Fig. 5. The top panel of Fig. 5 shows the confusion matrix when considering all barred galaxies as a single class ($A_{\rm bar}$) versus unbarred discs ($U_{\rm bar}$). Note that we randomly select $A_{\rm bar}$ galaxies to have the same number as $U_{\rm bar}$ galaxies. For the $A_{\rm bar}$ class, the model achieves a precision of 81.3% and a recall of 86.8%. For the $U_{\rm bar}$ class, the precision is 85.9%, and the recall is 80.1%. These values indicate a good ability to distinguish barred from unbarred discs with relatively high purity for both classes. The stricter criteria for $U_{\rm bar}$ (especially the ΔP cut) contribute to its high precision, minimising the fraction of contaminants.

The lower panel of Fig. 5 further divides the barred galaxies into $S_{\rm bar}$ and $W_{\rm bar}$. For $S_{\rm bar}$, the precision is 75.9%, with a recall of 75.9% from the visually labelled strong bars in the test set. Most confusion for $S_{\rm bar}$ arises from $W_{\rm bar}$ galaxies, which is expected. For $W_{\rm bar}$, the precision is 56.2% and the recall is 50.8%. This class remains the most challenging, with significant confusion with both $S_{\rm bar}$ and $U_{\rm bar}$ galaxies. This is expected given the continuous nature of bar strengths and reflects the inherent difficulty in distinguishing subtle weak bar features. For $U_{\rm bar}$, the precision remains high at 73.0% (slightly lower than when $A_{\rm bar}$ was a single class), and the recall is 80.1%. A higher contamination of visually labelled $W_{\rm bar}$ galaxies emerges.

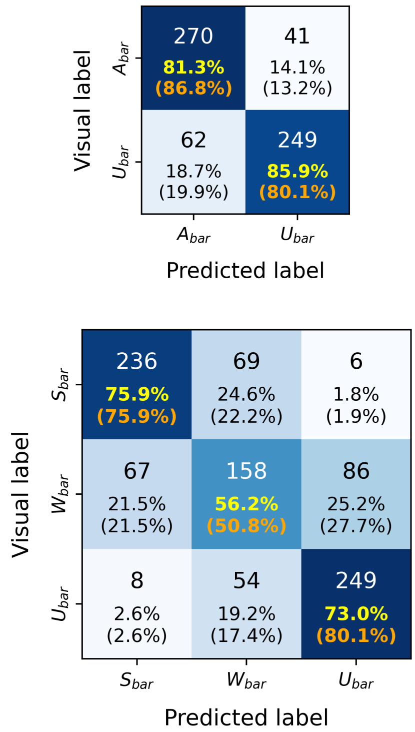

Fig. 5. Confusion matrices for the sample selected using the criteria in Sect. 3.1.2, colour-coded by the total number of galaxies in each row. The content of each cell is the same as in Fig. 4. *Top*: Confusion matrix considering all bars as a single class ($A_{\rm bar}$). We selected the same number of $A_{\rm bar}$ and $U_{\rm bar}$ examples. *Lower*: Confusion matrix dividing the bars into $S_{\rm bar}$ and $W_{\rm bar}$.

Overall, these criteria provide a robust separation, particularly between the general 'barred' ($A_{\rm bar}$) and 'unbarred' ($U_{\rm bar}$) categories, and a reasonably pure sample of 'strongly barred' ($S_{\rm bar}$) galaxies. While the 'weakly barred' ($W_{\rm bar}$) class has lower purity, its inclusion enables a more nuanced investigation of bar strength effects. In Appendix A, we further demonstrate how the adopted definitions, especially for $U_{\rm bar}$, effectively mitigate confusion with the 'smooth' galaxy class (Fig. A.2). With this adopted classification scheme, the final mass-complete galaxy sample used for the subsequent analysis (from Sect. 2.2) contains 3174 $U_{\rm bar}$ galaxies and 2 405 $A_{\rm bar}$ galaxies. Of the barred galaxies, 1 261 are classified as $S_{\rm bar}$ and 1144 as $W_{\rm bar}$.

3.2. AGN selections

To investigate the bar-AGN connection, we identified AGN within our mass-complete disc galaxy sample using three different diagnostics (MIR colours, X-ray detections, and SED fitting), leveraging the multi-wavelength data available. We selected AGN based on their MIR colours using the criterion $W1 - W2 > 0.8$ mag (Vega system), as proposed by Stern et al. (2012). This selection was applied to all galaxies in our sample with WISE $W2 \leq 15$ mag (Vega) and a $S/N \geq 5$ in both the $W1$ and $W2$ bands. This method effectively identifies AGN whose hot dust emission dominates the MIR spectrum, yielding a sample of 15 MIR AGN within our final disc galaxy sample. We identified X-ray AGN by cross-matching our galaxy sample with the eFEDS main catalogue (Salvato et al. 2022). Galaxies coincident with

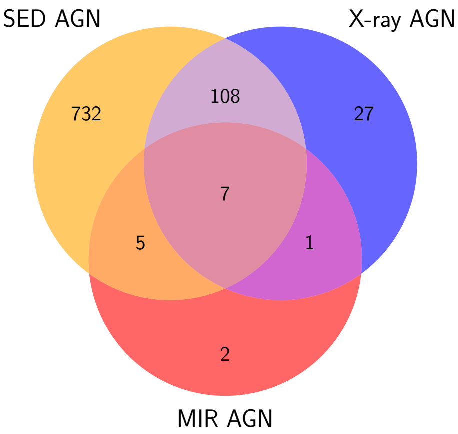

Fig. 6. Venn diagram showing the number of unique and overlapping AGN identified by the three selection methods within the combined sample of barred and unbarred disc galaxies analysed in this work.

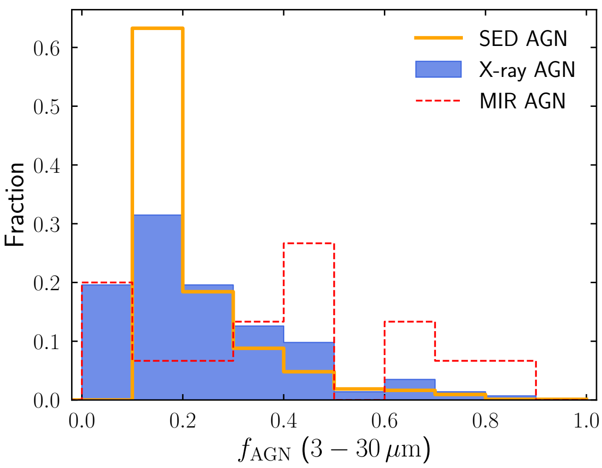

Fig. 7. Distributions of AGN fraction, $f_{\rm AGN}(3\text{–}30\,\mu\text{m})$, for the MIR-selected (dashed red line), X-ray-selected (blue-filled histogram), and SED-selected (solid orange line) AGN samples.

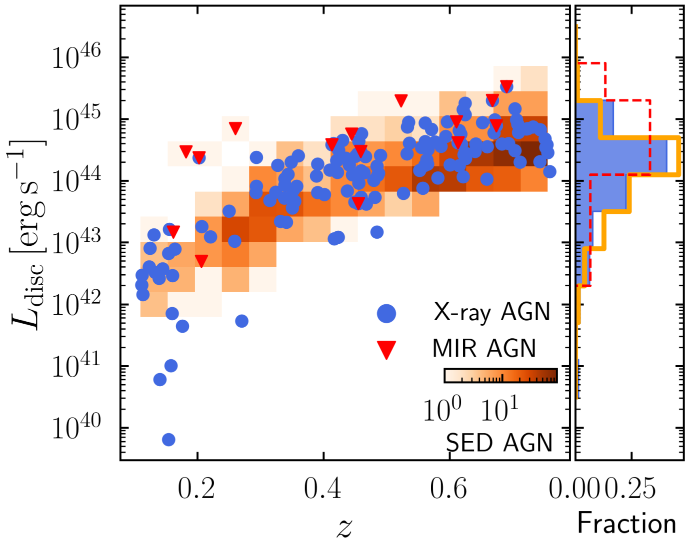

Fig. 8. Accretion disc luminosity ($L_{\rm disc}$) as a function of redshift for MIR AGN (red triangles), X-ray AGN (blue circles), and SED AGN (2D histogram, colour-coded by the number of sources). The marginal histogram on the right shows the $L_{\rm disc}$ distributions for MIR-selected (dashed red line), X-ray-selected (blue-filled histogram), and SED-selected (solid orange line) AGN.

an eFEDS X-ray source classified as an AGN in that catalogue were flagged as X-ray AGN. This selection resulted in 143 X-ray AGN. Following the methodology in LM24, we identified AGN based on the results of our `CIGALE` SED fitting. In this work, however, we adopted a more conservative selection criterion, classifying a galaxy as an SED AGN if its AGN fraction in the 3–30, µm range satisfies $f_{\rm AGN} \geq 0.1$. This threshold was chosen to ensure a more robust identification of AGN-dominated sources and is motivated by dedicated tests showing that galaxies with $f_{\rm AGN} \geq 0.1$ exhibit a significantly worse SED fit when the AGN component is omitted (see Appendix B). Using this criterion, we identified 852 SED AGN.

Figure 6 presents a Venn diagram illustrating the overlap between these three AGN selection methods for the total sample of barred and unbarred disc galaxies used in this work. While there is some overlap – with seven AGN identified by all three methods and 108 common to SED and X-ray selections – each technique also uniquely identifies a distinct subset of AGN. This highlights the complementary nature of these selection methods in capturing different facets of the AGN population.

The primary AGN properties derived from the CIGALE SED fitting that we explore in this paper are the AGN fraction, $f_{\rm AGN}$ (3–30 µm), and the AGN accretion disc luminosity, $L_{\rm disc}$, which is computed as the viewing angle-averaged accretion disc luminosity. The accretion disc luminosity considered is the intrinsic luminosity, i.e. before reprocessing by the torus (Yang et al. 2018). Figure 7 displays the distribution of $f_{\rm AGN}$ (3–30 µm) for each of the three AGN types. SED AGN, by selection, have $f_{\rm AGN} \geq 0.1$ and show a distribution that generally decreases towards higher $f_{\rm AGN}$ values, with most having $f_{\rm AGN} < 0.3$. X-ray AGN exhibit a similar $f_{\rm AGN}$ distribution with a peak around $f_{\rm AGN} \approx 0.2$, but include objects with $f_{\rm AGN} < 0.1$. Although the MIR AGN sample is small, they typically display high $f_{\rm AGN}$ values, consistent with the selection criterion targeting MIR-dominant sources (see LM24).

Figure 8 shows $L_{\rm disc}$ as a function of redshift for the three AGN types. The main panel illustrates that for all types, a higher number of AGN, particularly the more luminous ones, are found at relatively higher redshifts ($z > 0.4$) within our sample range (up to $z = 0.76$). The marginal histogram on the right of Fig. 8 displays the overall $L_{\rm disc}$ distributions for each AGN type. Both SED and X-ray AGN span a broad range from approximately 10^{42} to $\sim 10^{45}$ erg s^{-1}. SED AGN present a slightly higher fraction of fainter sources compared to X-ray AGN. The bulk of SED and X-ray AGN have $L_{\rm disc}$ between 10^{44} and 10^{45} erg s^{-1}. MIR AGN are consistently the most luminous in our sample, with approximately 85% of them having $L_{\rm disc} > 10^{44}$ erg s^{-1}.

3.3. Control pools

To robustly assess the impact of bars on AGN activity, it is crucial to construct appropriate control samples of unbarred galaxies. This is because AGN occurrence is known to correlate with host galaxy properties such as stellar mass (e.g. Kauffmann et al. 2003), and bar incidence itself can depend on stellar mass and colour (e.g. Masters et al. 2011). Therefore, for each barred galaxy ($S_{\rm bar}$ or $W_{\rm bar}$), we identified a pool of unbarred ($U_{\rm bar}$) control galaxies. These control galaxies were required to simultaneously match the barred galaxy in redshift (z), stellar mass ($M_\star$), and rest-frame ($g-r$) colour, according to the following

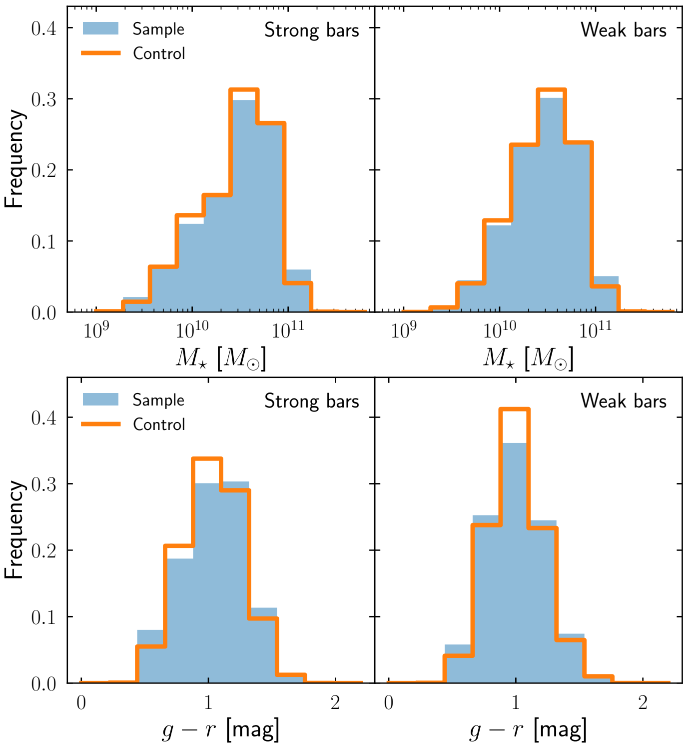

Fig. 9. $M_\star$ (*top panels*) and $(g-r)$ colour (*bottom panels*) distributions for strongly and weakly barred galaxies (blue-filled histograms), compared to their corresponding unbarred control samples (thick orange lines).

criteria:

$$|z_{\rm control} - z_{\rm sample}| \leq \Delta z \times (1 + z_{\rm sample}), \quad (1)$$

$$|\log M_{*,\rm control} - \log M_{*,\rm sample}| \leq \Delta M_*, \quad (2)$$

$$|(g-r)_{\rm control} - (g-r)_{\rm sample}| \leq \Delta(g-r), \quad (3)$$

where the initial matching tolerances were set to $\Delta z = 0.05$, $\Delta M_\star = 0.1$ dex, and $\Delta(g-r) = 0.1$. These values were chosen based on the typical photo-z precision, the median uncertainties in our CIGALE-derived $M_\star$, and the median photometric uncertainties in $(g-r)$ colour, respectively. For each barred galaxy in the sample, we required at least ten unique $U_{\rm bar}$ control counterparts satisfying these criteria. If fewer than ten controls were found, the tolerances were iteratively increased by a factor of 1.5 (up to a maximum of three iterations). If, after these iterations, fewer than ten controls were still found, the original barred galaxy was excluded from analyses requiring matched controls. When more than ten controls were available, ten were randomly selected to form the control pool for that specific barred galaxy. This larger population of unbarred galaxies compared to the barred sample minimises potential biases related to host galaxy properties.

We compared the $M_\star$ and $(g-r)$ distributions for barred galaxies and relative unbarred controls to verify the effectiveness of our control sample matching procedure. Figure 9 displays the stellar mass ($M_\star$, top panels) and observed $(g-r)$ colour (bottom panels) distributions for the $S_{\rm bar}$ and $W_{\rm bar}$ samples (blue-filled histograms) compared to their respective $U_{\rm bar}$ control samples (thick orange lines). In both cases, the distributions demonstrate excellent agreement between the samples and their controls after the matching process. This confirms that our methodology successfully removes biases related to stellar mass, redshift, and colour when comparing barred and unbarred populations.

4. Results

This section presents our findings on the bar-AGN connection. We first analyse this relationship by comparing AGN frequency in barred versus unbarred disc galaxies (binary classification). Subsequently, we explore how relative AGN power ($f_{\rm AGN}$) and absolute AGN luminosity ($L_{\rm disc}$) vary with bar presence and type.

4.1. Bar-AGN connection using a binary AGN classification

To assess if galactic bars contribute to AGN fuelling, we compared AGN frequencies (i.e. the fraction of galaxies hosting an AGN) in $S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ samples against their respective $U_{\rm bar}$ control samples, which were carefully matched in redshift, stellar mass, and $(g-r)$ colour (Sect. 3.3). The ratio of the AGN frequency in a barred sample to that in its corresponding control sample gives the 'AGN excess', where a value greater than 1 indicates an enhancement of AGN activity in the presence of a bar. Results are detailed in Table 3 and Fig. 10.

Our analysis consistently reveals an enhanced AGN fraction in barred galaxies across all bar types and AGN selection methods. MIR-selected AGN show a notable enhancement in barred galaxies – though statistically limited due to small numbers – with an AGN excess of approximately 6 for the three bar samples (≈0.3 ± 0.1% vs. ≈0.05 ± 0.02% in controls). For the X-ray AGN, this enhancement is statistically more robust: the $A_{\rm bar}$ sample hosts X-ray AGN at a frequency of 1.8 ± 0.3%, significantly higher than the 0.7 ± 0.1% in controls, yielding an AGN excess ratio of 2.7±0.5. Similar excesses are observed for $S_{\rm bar}$ (2.9±0.6) and $W_{\rm bar}$ (2.5 ± 0.6). For SED AGN, the most prevalent type, the enhancement in barred galaxies is only modest. The $S_{\rm bar}$ sample has an SED AGN frequency of 11 ± 1%, compared to 9.6 ± 0.3% in the matched control sample (excess ratio of 1.2 ± 0.1). The absolute difference in frequency is 1.4 percentage points. Propagating the binomial errors in quadrature yields an error on this difference of ≈1.04%, corresponding to a significance of approximately 1.3σ. While this is below the conventional threshold for high statistical significance, a Pearson's chi-squared test on the raw galaxy counts (see Table 3) yields a p-value of 0.04, indicating a marginal statistical rejection of the null hypothesis that AGN fraction is independent of the presence of a strong bar. We therefore consider this a tentative detection.

The $W_{\rm bar}$ sample shows a higher frequency of 15 ± 1% versus 14.1 ± 0.3% in controls (excess of 1.07 ± 0.08). However, this difference is not statistically significant. For $A_{\rm bar}$, the frequency is 13 ± 1% compared to 11.8 ± 0.2% in their $U_{\rm bar}$ controls, yielding an overall excess of 1.12 ± 0.06. The absolute difference in frequency is 1.4 percentage points, with an error of ≈1.02%, corresponding to a significance of ≈1.4σ. Similarly to the $S_{\rm bar}$ sample, a chi-squared test on the raw counts yields a p-value of 0.027. This strengthens the argument that, while the absolute effect is modest, there is a detectable enhancement of SED-selected AGN in the general population of barred galaxies within your sample.

To ensure that the imbalance in sample size – with control samples being ten times larger than the barred samples – does not introduce a bias, we performed a resampling test. We generated 1000 random subsets of the $U_{\rm bar}$ control samples, each with a size matching the corresponding barred sample ($S_{\rm bar}$, $W_{\rm bar}$, or $A_{\rm bar}$). We found that the median SED AGN fraction of these subsets remained identical, confirming that the full control sample provides an unbiased reference. However, due to the reduced statistical power (and larger error bars) of the smaller subsets, individual p-values frequently exceed the 0.05 threshold. For $S_{\rm bar}$

Table 3. MIR, X-ray, and SED AGN frequencies in $S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ disc galaxies, compared to their respective matched $U_{\rm bar}$ control samples.

Type	AGN freq. in $S_{\rm bar}$	AGN freq. in $U_{\rm bar}$ contr.	Excess	AGN freq. in $W_{\rm bar}$	AGN freq. in $U_{\rm bar}$ contr.	Excess	AGN freq. in $A_{\rm bar}$	AGN freq. in $U_{\rm bar}$ contr.	Excess
MIR	0.4 ± 0.2% (5/1137)	0.05 ± 0.02% (6/11 370)	8 ± 5	0.2 ± 0.1% (2/1049)	0.05 ± 0.02% (5/10 490)	4 ± 3	0.3 ± 0.1% (7/2186)	0.05 ± 0.2% (11/21 860)	6 ± 3
X-ray	2.0 ± 0.4% (23/1137)	0.7 ± 0.1% (79/11 370)	2.9 ± 0.7	1.6 ± 0.4% (17/1049)	0.7 ± 0.1% (69/10 490)	2.5 ± 0.7	1.8 ± 0.3% (40/2186)	0.7 ± 0.1% (148/21 860)	2.7 ± 0.5
SED	11 ± 1% (130/1137)	9.6 ± 0.3% (1097/11 370)	1.2 ± 0.1	15 ± 1% (158/1049)	14.1 ± 0.3% (1481/10 490)	1.07 ± 0.08	13.2 ± 0.7% (288/2186)	11.8 ± 0.2% (2578/21 860)	1.12 ± 0.06

Notes. The 'Excess' column reports the ratio of AGN frequency in the barred sample to that in its control sample. Errors are calculated using binomial statistics. Numbers in brackets indicate the raw count of AGN over the total number of galaxies in each specific sample.

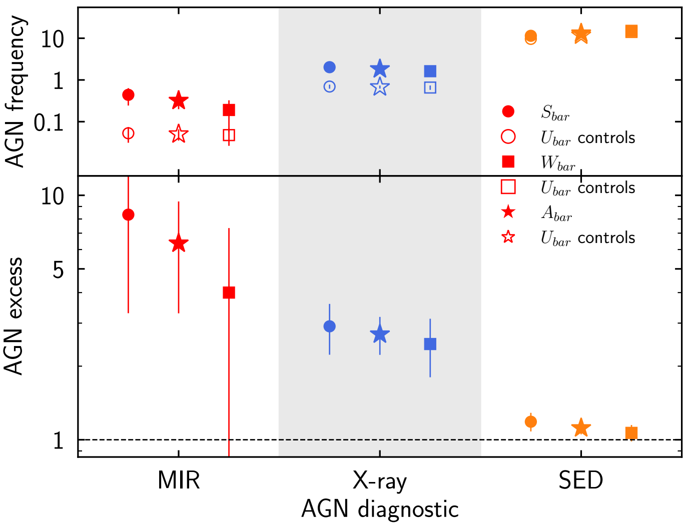

Fig. 10. AGN frequency and relative AGN excess in barred vs. unbarred disc galaxies. *Top panel*: Frequency of AGN (MIR: red, X-ray: blue, SED: orange) in $S_{\rm bar}$ (circles), $W_{\rm bar}$ (squares), $A_{\rm bar}$ (stars), and the corresponding $U_{\rm bar}$ control samples (empty symbols of the same type). Errors are calculated using binomial statistics. *Bottom panel*: Ratio of the AGN frequency in the barred samples to that in their corresponding $U_{\rm bar}$ controls (AGN excess).

and $A_{\rm bar}$, we observed a p-value$<$ 0.05 in 16.7% and 19.3% of the cases, respectively. This confirms that while the excess is present, the statistical significance of the SED-selected AGN result is marginal and relies on the precision provided by the large control sample.

Our finding of an enhanced AGN fraction in barred galaxies is broadly consistent with several previous observational studies. For instance, Galloway et al. (2015) found that barred galaxies are more likely to host an AGN compared to unbarred ones. Similarly, Alonso et al. (2018) reported a higher incidence of AGN in local barred galaxies. More recently, Garland et al. (2023, 2024), using a variety of AGN selection techniques, also found that barred galaxies are more likely to host AGN. Some studies, such as Lee et al. (2012) and Cheung et al. (2015), found a higher AGN frequency in barred than in unbarred galaxies, but this difference disappeared after controlling for host galaxy properties. Indeed, Cisternas et al. (2015) showed that differences in bar fractions between active and inactive galaxies can be suppressed after controlling for $M_\star$, suggesting that direct comparisons without careful matching can be misleading. Our results, based on meticulously matched control samples, reinforce the idea that bars do play a role in AGN triggering, even if the effect varies with AGN type and luminosity.

The varying strength of the AGN excess across our three selection methods warrants further discussion. For the MIR-selected sample, the very small number of objects (only seven AGN in $A_{\rm bar}$ galaxies) precludes a robust physical interpretation of its large excess; while tantalising, this result requires confirmation with much larger samples. For the more numerous X-ray-selected sample, the observed excess could be influenced by selection effects inherent to our methodology. Our analysis excludes, by necessity, edge-on galaxies where bars cannot be identified, which favours face-on systems. Since X-ray emission can be significantly obscured by the high column densities of the host galaxy's interstellar medium (in addition to the AGN torus), our sample may be systematically biased towards less obscured systems, a selection effect that may not be perfectly accounted for in our control matching. More broadly, it is well-established that different selection techniques are sensitive to distinct AGN populations and physical processes, which evolve on different timescales (Alexander & Hickox 2012). A full deconvolution of these effects is beyond the scope of this work but represents a crucial avenue for future investigation.

4.2. Bar–AGN connection using the relative and absolute AGN power

In this Section, we present the results on the bar-AGN connection using the continuous $f_{\rm AGN}$ and $L_{\rm disc}$ parameters, derived from SED fitting. Since the MIR AGN sample is limited to 15 instances, we focus only on X-ray and SED AGN.

4.2.1. The relative AGN power in barred and unbarred galaxies

We now investigate the bar-AGN connection by examining the distribution of relative AGN power, quantified by the $f_{\rm AGN}$ (3 − 30 μm) parameter. Figure 11 displays the normalised $f_{\rm AGN}$ distributions for X-ray and SED AGN hosted in $S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ galaxies, compared to their respective $U_{\rm bar}$ control samples. We used the two-sample Kolmogorov–Smirnov (KS) test (Hodges 1958) to evaluate whether the $f_{\rm AGN}$ distributions of barred AGN differ significantly from those of their unbarred controls (Fig. 11). The KS test compares the cumulative distributions of two samples; a small p-value (typically below a chosen significance level, here we take it to be 0.05) suggests that the two distributions are significantly different.

For the X-ray AGN, the $f_{\rm AGN}$ distributions in $S_{\rm bar}$ and $A_{\rm bar}$ samples differ from their $U_{\rm bar}$ controls (p-value $<$ 0.05). Specifically, the $U_{\rm bar}$ controls show a higher fraction of galaxies in the $f_{\rm AGN}$ range of 0.2−0.5. No significant difference is observed between the $W_{\rm bar}$ X-ray AGN and their controls (p-value =

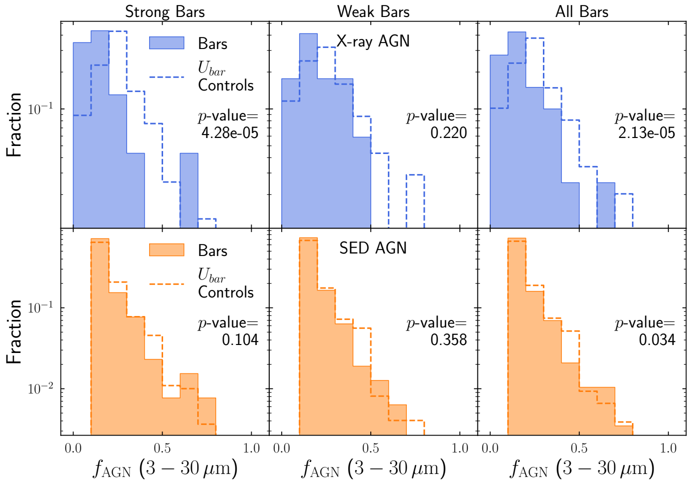

Fig. 11. Normalised distributions of $f_{\rm AGN}$ for X-ray (top panels) and SED (bottom panels) AGN. Distributions are shown for $S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ galaxies (filled histograms) compared to their respective $U_{\rm bar}$ control samples (dashed lines). The p-value from a KS test comparing the barred sample with its control is reported in each panel.

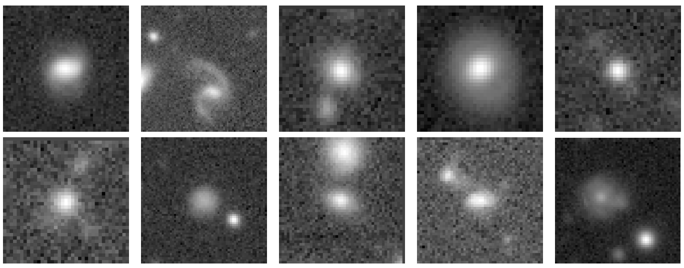

Fig. 12. HSC-SSP i-band image cutouts of the ten disc galaxies from the parent sample with the highest $f_{\rm AGN}$ values (all $f_{\rm AGN} > 0.7$). The cutout size and scaling are the same as in Fig. 3. The predicted class is indicated at the top of each panel.

0.22), suggesting that their $f_{\rm AGN}$ distributions are consistent with being drawn from the same parent population. For SED-selected AGN, there are no significant differences in the $f_{\rm AGN}$ distributions between the strongly and weakly barred galaxy types ($S_{\rm bar}$, $W_{\rm bar}$) and their corresponding $U_{\rm bar}$ controls (p-value > 0.05). On the contrary, in the $A_{\rm bar}$ case, we cannot rule out the hypothesis that the two distributions are drawn from the same parent distribution, since we obtained a p-value < 0.05.

A striking feature across all bar samples and AGN types is the near absence of galaxies with $f_{\rm AGN} > 0.8$. The opposite trend is observed in studies focusing on major mergers (e.g. LM24), which often host such AGN-dominated systems. This scarcity suggests that bar-driven secular processes are unlikely to be the primary drivers of the most AGN-dominated phases. To further explore this, Fig. 12 shows image cutouts of the ten galaxies in our entire disc galaxy sample (barred and unbarred) with the highest $f_{\rm AGN}$ values (all $f_{\rm AGN} > 0.7$). None of these galaxies exhibit clear bar features, while approximately half show signs of ongoing interaction or disturbed morphologies (i.e., merger features). This observation lends further support to the idea that major mergers, rather than bars, are predominantly responsible for triggering AGN with very high $f_{\rm AGN}$.

To quantify the relationship between bar presence and $f_{\rm AGN}$, we calculated the bar fraction ($f_{\rm bar}$) as a function of $f_{\rm AGN}$. This was done by dividing the number of barred AGN ($S_{\rm bar}$, $W_{\rm bar}$, or $A_{\rm bar}$) by the total number of AGN (barred and unbarred) in bins of $f_{\rm AGN}$. For the X-ray AGN, we used four equally spaced bins between $f_{\rm AGN} = 0$ and 1. For the SED AGN, given the larger number statistics, we adopted six bins: five equally spaced between $f_{\rm AGN} = 0$ and 0.7, and one final bin spanning $0.7 < f_{\rm AGN} \leq 1$[3]. The bar fraction is defined as $f_{\rm bar} = N_{\rm bar,\,class}/(N_{\rm bar,\,class} + N_{\rm unbar})$[4]. We used bootstrapping with resampling (1000 samples for each population) for calculating $f_{\rm bar}$ in each $f_{\rm AGN}$ bin.

Figure 13 presents the median $f_{\rm bar}$–$f_{\rm AGN}$ relations for X-ray and SED AGN, for strong, weak, and all bars. For both AGN selections, $f_{\rm bar}$ remains largely constant as a function of $f_{\rm AGN}$, fluctuating between ∼20–50% depending on the bar type. Only for the $S_{\rm bar}$ sample, we notice a mild declining trend in $f_{\rm bar}$ with $f_{\rm AGN}$, which is likely due to the poor number statistics. For the most dominant AGN ($f_{\rm AGN} > 0.75$), we observe larger statistical uncertainties due to the smaller numbers of galaxies. As shown in Fig. 7, this near-absence of bars at $f_{\rm AGN} > 0.75$ is consistent across all bar classes and both AGN types.

This behaviour contrasts with the merger fraction–AGN fraction ($f_{\rm merg}$–$f_{\rm AGN}$) relation presented for the first time in our previous work (LM24). In Fig. 13, we overlay the best-fit parametrisations of the $f_{\rm merg}$–$f_{\rm AGN}$ relation from LM24 for MIR, X-ray, and SED AGN. The details of how these relations were parametrised and fitted are provided in Appendix C. These overlaid lines show that $f_{\rm merg}$ is flat up to $f_{\rm AGN} \approx 0.8$, followed by a rapid increase for $f_{\rm AGN} > 0.8$. This opposing trend between $f_{\rm bar}$ and $f_{\rm merg}$ at high $f_{\rm AGN}$ suggests that while both bars and mergers can fuel AGN with low to intermediate $f_{\rm AGN}$, major mergers become the dominant (if not the sole) mechanism for triggering the most AGN-dominated systems ($f_{\rm AGN} > 0.8$), where bar-driven fuelling appears inefficient or the bars themselves may be disrupted.

4.2.2. The absolute AGN power in barred and unbarred galaxies

In addition to relative AGN power, we analysed the distributions of absolute AGN accretion disc luminosity ($L_{\rm disc}$) for barred and unbarred galaxies. Figure 14 presents the normalised $L_{\rm disc}$ distributions for X-ray and SED-selected AGN hosted in $S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ galaxies, compared to their $U_{\rm bar}$ controls. We used the two-sample KS test to assess the statistical significance of any observed differences. For X-ray AGN, the $L_{\rm disc}$ distributions in barred galaxies ($S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$) are all statistically different from their respective $U_{\rm bar}$ controls (p-values < 0.05 in all cases). Figure 14 shows that the $U_{\rm bar}$ controls tend to host a larger fraction of high luminosity X-ray AGN ($L_{\rm disc} > 10^{44.5}$ erg s^{-1}) compared to barred galaxies. Conversely, barred galaxies show a relative excess of X-ray AGN at intermediate luminosities ($L_{\rm disc} \approx 10^{43}$–$10^{44.5}$ erg s^{-1}). For SED-selected AGN, however, there are no significant differences in the $L_{\rm disc}$ distributions between $S_{\rm bar}$, $W_{\rm bar}$ and their $U_{\rm bar}$ controls (p-values > 0.05), suggesting that bar presence does not strongly influence the typical luminosities of SED-selected AGN in our sample. When the two samples are combined, $A_{\rm bar}$, a mild difference with the unbarred controls (p-value = 0.014) emerges.

To further explore how bar presence relates to AGN luminosity, we calculated $f_{\rm bar}$ as a function of $L_{\rm disc}$. The results are shown in Fig. 15 for X-ray and SED AGN, separately for

[3] We varied the number and width of the bins, but the overall trends remained unchanged.

[4] $S_{\rm bar}$ and $W_{\rm bar}$ are treated independently: when computing one, the other is excluded from the corresponding sample.

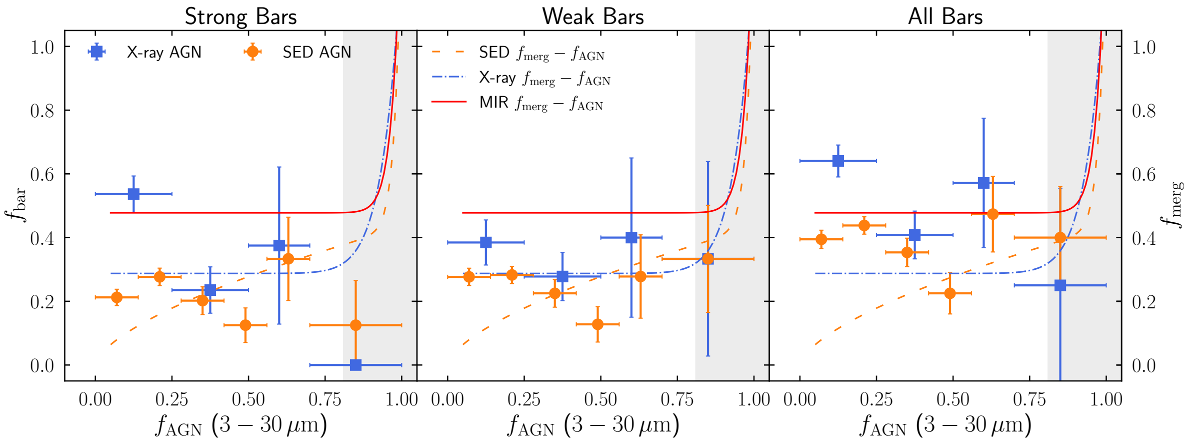

Fig. 13. Bar fraction ($f_{\rm bar}$) as a function of $f_{\rm AGN}$ for X-ray AGN (blue squares) and SED AGN (orange circles). Results are shown for $S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ separately. The $f_{\rm bar}$ is calculated in $f_{\rm AGN}$ bins using bootstrapping with resampling. The values reported are the median $f_{\rm bar}$ (with 1σ standard deviations as uncertainties); the x-axis error bars indicate bin width. Overlaid are the best-fit merger fraction ($f_{\rm merg}$)–$f_{\rm AGN}$ relations from LM24 for the MIR (solid red line), X-ray (dotted-dashed blue line), and SED (dashed orange line) AGN, parametrised as described in Appendix C. In each subplot, the grey-shaded area highlights the region where $f_{\rm merg}$ sharply rises as a function of $f_{\rm AGN}$.

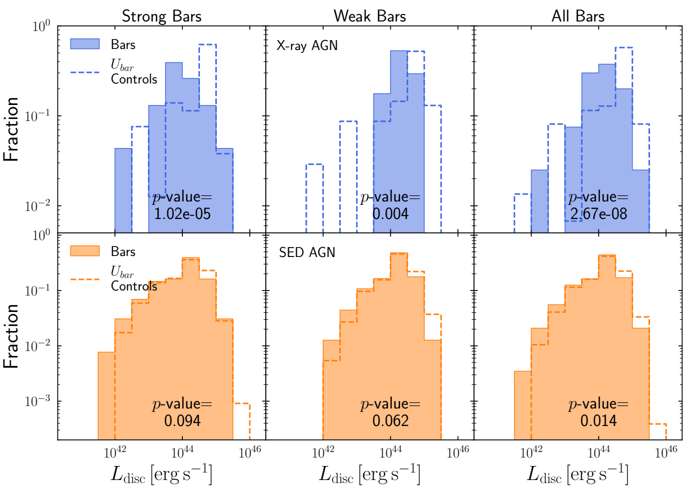

Fig. 14. Normalised distributions of AGN accretion disc luminosity ($L_{\rm disc}$) in X-ray (top panels) and SED (bottom panels) AGN. Distributions are shown for $S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ (filled histograms) compared to their respective $U_{\rm bar}$ controls (dashed lines). The p-value from a KS test comparing the barred sample with its control is reported in each panel.

$S_{\rm bar}$, $W_{\rm bar}$, and $A_{\rm bar}$ classifications. The $f_{\rm bar}$ was calculated in $L_{\rm disc}$ bins (six for X-ray AGN and seven for SED AGN), logarithmically spaced between 10^{42} and 10^{45} erg s^{-1} and an additional bin for AGN brighter than 10^{45} erg s^{-1}. For SED AGN, for $S_{\rm bar}$ and $A_{\rm bar}$, a clear trend emerges: $f_{\rm bar}$ generally decreases with increasing $L_{\rm disc}$: from ≈40–50% at $L_{\rm disc} = 10^{42}$ erg s^{-1}, to ≈20–30% at $L_{\rm disc} = 10^{45.5}$ erg s^{-1}. Similarly, X-ray AGN in the same bar types show a $f_{\rm bar} \approx 50\%$ at $L_{\rm disc} \leq 10^{44}$ erg s^{-1}, dropping to $f_{\rm bar} \approx 30\%$ for brighter AGN. For $W_{\rm bar}$ SED AGN, $f_{\rm bar}$ roughly stays constant at 30%, while X-ray AGN show a larger scatter due to the low number statistics. Notably, very few or no barred galaxies are found hosting the most luminous AGN ($L_{\rm disc} \approx 10^{46}$ erg s^{-1}) in our sample.

This observed decrease in bar fraction at high AGN luminosities is opposite to the behaviour of the merger-AGN relation, where the merger fraction typically increases with AGN luminosity (e.g. Treister et al. 2012; Glikman et al. 2015, LM24). We overlay the $f_{\rm merg}$–$L_{\rm disc}$ parametrised relation in Fig. 15 (we report the details of the parametrisation in Appendix C). Our findings suggest that while bars may effectively trigger and fuel AGN up to intermediate luminosities, they are not efficient mechanisms for sustaining the most powerful AGN. Several studies support the idea that bars are more commonly associated with lower-to-moderate luminosity AGN activity. For instance, Knapen et al. (2000) and Laine et al. (2002) found connections between bars and Seyfert nuclei, often indicative of moderate accretion. Similarly, in their sample of X-ray AGN, Koss et al. (2011) found that mergers are prevalent among the most luminous AGN, while bars (and other secular processes) might predominantly power moderate luminosity AGN. Marels et al. (2025) showed that the fraction of barred luminous AGN is at most 25%. Furthermore, Cisternas et al. (2013) found no significant correlation between bar strength and SMBH fuelling, suggesting that the extent of the bar-driven inflow is not directly connected with the degree of ongoing accretion (i.e. AGN luminosity). The decline we observe at high $L_{\rm disc}$ might imply either that the bars cannot channel material at sufficiently high rates to power the brightest AGN, or that the energetic feedback from such luminous AGN could disrupt or destroy the bar structure over time (e.g. Zee et al. 2023).

The observed $f_{\rm bar}$–$L_{\rm disc}$ trend is in excellent agreement with the emerging paradigm of a dichotomy in AGN fuelling. In this scenario, the most luminous AGN, such as quasars, are predominantly triggered by violent galaxy major mergers (e.g. Ramos Almeida et al. 2011, 2012; Treister et al. 2012; Euclid Collaboration: La Marca et al. 2026). Conversely, less luminous, Seyfert-like activity has been linked to internal secular processes, among which bars are a viable path (e.g. Cisternas et al. 2011).

5. The effect of mergers on the continuous analysis

In our analysis in Sect. 4 above, we deliberately excluded galaxies classified as major mergers in LM24 to isolate the influence of bars on AGN fuelling in relatively undisturbed disc galaxies. However, it is important to consider how this exclusion might affect our conclusions, particularly regarding the trends observed in the $f_{\rm bar} - f_{\rm AGN}$ and $f_{\rm bar} - L_{\rm disc}$ relations at high AGN dominance or power, where mergers are known to be significant. In this section, we explore two key aspects: first, whether any of the most dominant or luminous AGN, previously excluded as

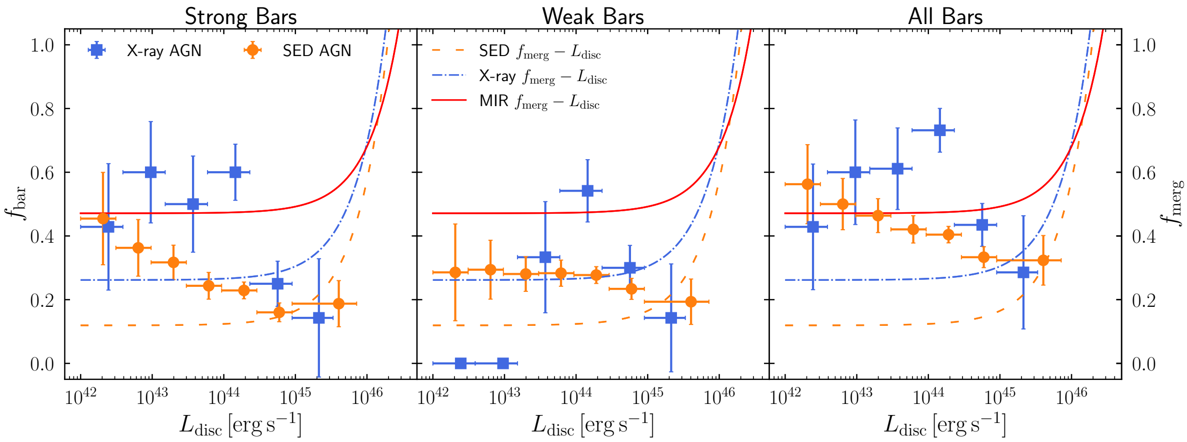

Fig. 15. Bar fraction as a function of the AGN total disc accretion luminosity (L_{disc}) for X-ray (blue lines) and SED (orange circles) AGN. The $f_{\rm bar}$ is calculated for the $S_{\rm bar}$ and $W_{\rm bar}$ individually (left and middle panels, respectively), and for the total $A_{\rm bar}$ sample (right panel). The error bars are the same as in Fig. 13. Overlaid are the best-fit $f_{\rm merg}$–$L_{\rm disc}$ relations from LM24 for the MIR (solid red line), X-ray (dotted-dashed blue line), and SED (dashed orange line) AGN, parametrised as described in Appendix C.

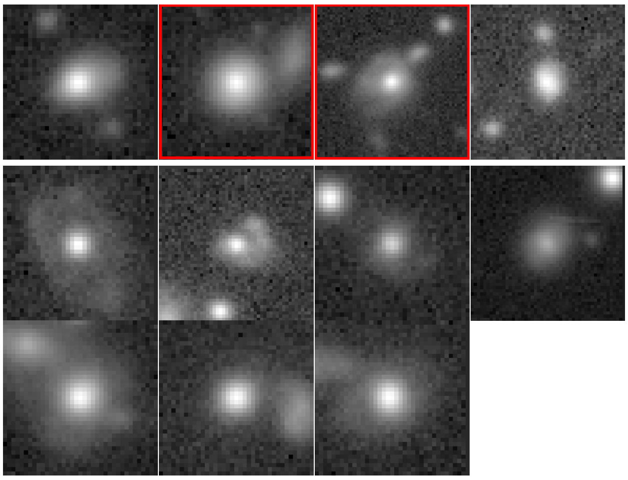

Fig. 16. HSC-SSP i-band cutouts of the most dominant ($f_{\rm AGN} > 0.7$, top row) and luminous AGN ($L_{\rm disc} > 10^{44.5}$ erg s^{-1}, middle and bottom rows) in our sample, previously classified as major mergers in LM24. The red boxes mark the galaxies in both samples. Scaling matches Fig. 3.

mergers, might also host bars; and second, how the continuous relations change when mergers are re-included in the sample.

First, we quantified the overlap between our `Zoobot` bar classifications and the merger classifications from LM24. Within the parent sample matched between both catalogues, we found that 226 $A_{\rm bar}$ galaxies (102 $S_{\rm bar}$, 124 $W_{\rm bar}$) and 224 $U_{\rm bar}$ discs were classified as major mergers. These figures correspond to a 10% increase in both samples. We then examined the AGN properties of these 'barred mergers'. We found that none of the 226 barred mergers host AGN with $f_{\rm AGN} > 0.7$. Similarly, nearly no mergers hosting an AGN with $L_{\rm disc} > 10^{44.5}$ erg s^{-1} were labelled as barred. This indicates that even when considering galaxies that are both barred and undergoing a major merger, such systems do not contribute to the population of extremely dominant or luminous AGN in our sample. Figure 16, which displays the i-band images of all merger-classified galaxies (barred or unbarred by `Zoobot`) hosting $f_{\rm AGN} > 0.7$ or $L_{\rm disc} > 10^{44.5}$ erg s^{-1}, further supports this: none of these dominant or luminous AGN hosts exhibit clear bar structures, while most of them show merging features. This confirms that excluding major mergers did not inadvertently remove a significant population of barred galaxies hosting extremely powerful AGN; such systems appear to be genuinely rare or absent.

Next, we re-evaluated the continuous relations – $f_{\rm bar}$ as a function of $f_{\rm AGN}$ and $L_{\rm disc}$ – by including the previously excluded mergers in our analysis. For this test, we recalculated $f_{\rm bar}$ in $f_{\rm AGN}$ and $L_{\rm disc}$ bins, as done in Sect. 4.2. The revised $f_{\rm bar}$–$f_{\rm AGN}$ and $f_{\rm bar}$–$L_{\rm disc}$ relations are presented in Appendix D, as Fig. D.1. This experiment revealed that the overall trends presented in Figures 13 and 15 remain qualitatively unchanged. To conclude, these tests robustly indicate that our primary conclusions regarding the inefficiency of bars in fuelling the most dominant (high $f_{\rm AGN}$) and most luminous (high $L_{\rm disc}$) AGN are not significantly biased by the initial exclusion of mergers. The dearth of bars in these extreme AGN regimes appears to be a genuine feature, with major mergers being the more likely drivers of such activity.

6. Summary and conclusions

We investigated the role of galactic bars in fuelling AGN in disc-dominated galaxies up to $z \approx 0.8$. Our analysis utilised a large, multi-wavelength catalogue from which we selected galaxies based on HSC-SSP i-band imaging (LM24). We employed a deep learning model (`Zoobot`, Walmsley et al. 2023), fine-tuned on Galaxy Zoo classifications, to identify strongly barred, weakly barred, and unbarred disc galaxies. Active galactic nuclei were selected via three independent methods (MIR colours, X-ray detections, and SED fitting with `CIGALE`). Moreover, their properties – specifically the AGN fraction ($f_{\rm AGN}$(3–30 µm)) and accretion disc luminosity ($L_{\rm disc}$) – were quantified. The impact of bars was assessed by comparing AGN incidence and properties in barred galaxies against carefully constructed redshift-, stellar mass-, and colour-matched unbarred control samples. We explicitly excluded major mergers from our primary analysis to isolate bar-driven secular processes, with their impact separately assessed. Our main findings can be summarised as follows:

1. Barred discs exhibit a higher fraction of AGN than their unbarred control counterparts across all three AGN selections. The AGN excess is more pronounced and robust for MIR AGN (factor of ~6) and X-ray AGN (factor of ~2.7). For the more numerous SED AGN population, the excess is more modest (factor of ~1.1), and its statistical significance

depends on the control sample size. This consistent trend across multiple independent tracers suggests that galactic bars might contribute to the triggering of AGN activity.
2. The fraction of AGN hosted in barred galaxies ($f_{\rm bar}$) remains largely flat as a function of relative AGN power ($f_{\rm AGN}$), fluctuating between ∼20−50%. We find a significant scarcity of barred galaxies, whether strongly or weakly, hosting AGN with $f_{\rm AGN} > 0.75$.
3. Similarly, $f_{\rm bar}$ is flat or mildly decreases with increasing $L_{\rm disc}$. While bars are associated with AGN up to intermediate luminosities (peaking around $L_{\rm disc} \approx 10^{44}\,{\rm erg\,s^{-1}}$ for X-ray AGN), there is a sharp decline in $f_{\rm bar}$ at higher luminosities, with very few or no barred galaxies found hosting the most luminous AGN ($L_{\rm disc} > 10^{45}\,{\rm erg\,s^{-1}}$).

Therefore, our work helps resolve the long-standing debate on the role of bars by defining their operational boundaries: they appear to be a viable, though not dominant, mechanism for sustaining low AGN activity but are inefficient at fuelling moderately to highly luminous as well as dominant phases of black hole growth. Our analysis, including the reintroduction of major mergers, confirms that these extreme AGN are preferentially hosted in unbarred merging galaxies, and rarely, if ever, in clearly barred disc galaxies in our sample. In conclusion, our findings align with the previously proposed dual-mode picture of AGN fuelling: internal processes, such as bars, sustain a population of low-luminosity AGN, while major mergers are the primary drivers of the most powerful, high-luminosity AGN.

Future work should incorporate detailed bar morphologies (e.g. bar length) and host galaxy star formation to clarify bar-driven AGN fuelling conditions. Larger, deeper galaxy and AGN samples are crucial for robustly probing these relationships across wider parameter ranges, especially at high AGN power, and for studying their evolution with redshift, a task for which our current sample lacks sufficient statistical power. The forthcoming ESA *Euclid* mission data release (Euclid Collaboration: Mellier et al. 2025c) will provide vast, high-quality imaging, identifying hundreds of thousands of barred systems (Euclid Collaboration: Huertas-Company et al. 2026). This unprecedented dataset will enable a more detailed and statistically powerful investigation of bar-driven AGN fuelling and its evolution across cosmic time, decisively advancing our understanding of secular processes in galaxy evolution, particularly at higher redshifts.

Data availability

Catalogues containing the galaxy morphological classification are available at the CDS via `https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/707/A152`, and from the Zenodo repository `https://doi.org/10.5281/zenodo.18299075`

Acknowledgements. We thank the anonymous referee for their valuable feedback. The Hyper Suprime-Cam (HSC) collaboration includes the astronomical communities of Japan and Taiwan, and Princeton University. The HSC instrumentation and software were developed by the National Astronomical Observatory of Japan (NAOJ), the Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU), the University of Tokyo, the High Energy Accelerator Research Organization (KEK), the Academia Sinica Institute for Astronomy and Astrophysics in Taiwan (ASIAA), and Princeton University. Funding was contributed by the FIRST program from Japanese Cabinet Office, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), the Japan Society for the Promotion of Science (JSPS), Japan Science and Technology Agency (JST), the Toray Science Foundation, NAOJ, Kavli IPMU, KEK, ASIAA, and Princeton University. Based on data collected at the Subaru Telescope and retrieved from the HSC data archive system, which is operated by Subaru Telescope and Astronomy Data Center at National Astronomical Observatory of Japan. Based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs 177.A-3016, 177.A-3017 and 177.A-3018, and on data products produced by Target/OmegaCEN, INAF-OACN, INAF-OAPD and the KiDS production team, on behalf of the KiDS consortium. OmegaCEN and the KiDS production team acknowledge support by NOVA and NWO-M grants. Members of INAF-OAPD and INAF-OACN also acknowledge the support from the Department of Physics & Astronomy of the University of Padova, and of the Department of Physics of Univ. Federico II (Naples). This publication makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration. This work is based on data from eROSITA, the soft X-ray instrument aboard SRG, a joint Russian-German science mission supported by the Russian Space Agency (Roskosmos), in the interests of the Russian Academy of Sciences represented by its Space Research Institute (IKI), and the Deutsches Zentrum für Luft- und Raumfahrt (DLR). This publication is part of the project 'Clash of the titans: deciphering the enigmatic role of cosmic collisions' (with project number VI.Vidi.193.113) of the research programme Vidi which is (partly) financed by the Dutch Research Council (NWO). The training and testing of the CNN were carried out on the Dutch National Supercomputer (Snellius). We thank SURF (`www.surf.nl`) for the support in using the National Supercomputer Snellius. We thank the Center for Information Technology of the University of Groningen for support and access to the Hábrók high performance computing cluster.

References

Aguerri, J. A. L., Méndez-Abreu, J., & Corsini, E. M. 2009, A&A, 495, 491
Aihara, H., Arimoto, N., Armstrong, R., et al. 2018, PASJ, 70, S4
Aihara, H., AlSayyad, Y., Ando, M., et al. 2022, PASJ, 74, 247
Alexander, D. M., & Hickox, R. C. 2012, New Astron. Rev., 56, 93
Alonso, S., Coldwell, G., Duplancic, F., Mesa, V., & Lambas, D. G. 2018, A&A, 618, A149
Athanassoula, E. 2003, MNRAS, 341, 1179
Barazza, F. D., Jogee, S., & Marinova, I. 2008, ApJ, 675, 1194
Boquien, M., Burgarella, D., Roehlly, Y., et al. 2019, A&A, 622, A103
Brunner, H., Liu, T., Lamer, G., et al. 2022, A&A, 661, A1
Bruzual, G., & Charlot, S. 2003, MNRAS, 344, 1000
Burgarella, D., Buat, V., & Iglesias-Páramo, J. 2005, MNRAS, 360, 1413
Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, ApJ, 533, 682
Chabrier, G. 2003, PASP, 115, 763
Cheung, E., Athanassoula, E., Masters, K. L., et al. 2013, ApJ, 779, 162
Cheung, E., Trump, J. R., Athanassoula, E., et al. 2015, MNRAS, 447, 506
Cisternas, M., Jahnke, K., Bongiorno, A., et al. 2011, ApJ, 741, L11
Cisternas, M., Gadotti, D. A., Knapen, J. H., et al. 2013, ApJ, 776, 50
Cisternas, M., Sheth, K., Salvato, M., et al. 2015, ApJ, 802, 137
de Jong, J. T. A., Kuijken, K., Applegate, D., et al. 2013, The Messenger, 154, 44
de Sá-Freitas, C., Fragkoudi, F., Gadotti, D. A., et al. 2023, A&A, 671, A8
de Sá-Freitas, C., Gadotti, D. A., Fragkoudi, F., et al. 2025, A&A, 698, A5
Di Matteo, T., Springel, V., & Hernquist, L. 2005, Nature, 433, 604
Draine, B. T., Aniano, G., Krause, O., et al. 2014, ApJ, 780, 172
Driver, S. P., Hill, D. T., Kelvin, L. S., et al. 2011, MNRAS, 413, 971
Edge, A., Sutherland, W., Kuijken, K., et al. 2013, The Messenger, 154, 32
Erwin, P. 2018, MNRAS, 474, 5372
Eskridge, P. B., Frogel, J. A., Pogge, R. W., et al. 2000, AJ, 119, 536
Euclid Collaboration (Huertas-Company, M., et al.) 2026, A&A, in press, `http://dx.doi.org/10.1051/0004-6361/202554606`
Euclid Collaboration (La Marca, A., et al.) 2026, A&A, in press, `http://dx.doi.org/10.1051/0004-6361/202554579`
Euclid Collaboration (Mellier, Y., et al.) 2025c, A&A, 697, A1
Galloway, M. A., Willett, K. W., Fortson, L. F., et al. 2015, MNRAS, 448, 3442
Garland, I. L., Fahey, M. J., Simmons, B. D., et al. 2023, MNRAS, 522, 211
Garland, I. L., Walmsley, M., Silcock, M. S., et al. 2024, MNRAS, 532, 2320
Géron, T., Smethurst, R. J., Lintott, C., et al. 2021, MNRAS, 507, 4389
Glikman, E., Simmons, B., Mailly, M., et al. 2015, ApJ, 806, 218
Guo, Y., Jogee, S., Wise, E., et al. 2025, ApJ, 985, 181
Hart, R. E., Bamford, S. P., Willett, K. W., et al. 2016, MNRAS, 461, 3663
Heckman, T. M., & Best, P. N. 2014, ARA&A, 52, 589
Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19
Hodges, J. L. 1958, Arkiv for Matematik, 3, 469
Holwerda, B. W., Robertson, C., Cook, K., et al. 2024, PASA, 41, e115
Hopkins, P. F., & Quataert, E. 2010, MNRAS, 407, 1529
Hopkins, P. F., Hernquist, L., Cox, T. J., et al. 2006, ApJS, 163, 1
Hopkins, P. F., Hernquist, L., Cox, T. J., & Kereš, D. 2008, ApJS, 175, 356
Kauffmann, G., Heckman, T. M., Tremonti, C., et al. 2003, MNRAS, 346, 1055
Knapen, J. H., Shlosman, I., & Peletier, R. F. 2000, ApJ, 529, 93

Koss, M., Mushotzky, R., Veilleux, S., et al. 2011, ApJ, 739, 57
Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, Adv. Neural Inf. Process. Syst., 25
Kruk, S. J., Lintott, C. J., Bamford, S. P., et al. 2018, MNRAS, 473, 4731
Kuijken, K., Heymans, C., Dvornik, A., et al. 2019, A&A, 625, A2
La Marca, A., Margalef-Bentabol, B., Wang, L., et al. 2024, A&A, 690, A326
Laine, S., Shlosman, I., Knapen, J. H., & Peletier, R. F. 2002, ApJ, 567, 97
Le Conte, Z. A., Gadotti, D. A., Ferreira, L., et al. 2024, MNRAS, 530, 1984
Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, Proc. IEEE, 86, 2278
Lee, G.-H., Woo, J.-H., Lee, M. G., et al. 2012, ApJ, 750, 141
Lin, L.-H., Wang, H.-H., Hsieh, P.-Y., et al. 2013, ApJ, 771, 8
Liu, Z., Mao, H., Wu, C. Y., et al. 2022, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 11976
Małek, K., Buat, V., Roehlly, Y., et al. 2018, A&A, 620, A50
Marconi, A., Risaliti, G., Gilli, R., et al. 2004, MNRAS, 351, 169
Marels, V., Mesa, V., Jaque Arancibia, M., et al. 2025, arXiv e-prints [arXiv:2505.23958]
Martin, G., Kaviraj, S., Volonteri, M., et al. 2018, MNRAS, 476, 2801
Masters, K. L., Nichol, R. C., Hoyle, B., et al. 2011, MNRAS, 411, 2026
Masters, K. L., Nichol, R. C., Haynes, M. P., et al. 2012, MNRAS, 424, 2180
Méndez-Abreu, J., Ruiz-Lara, T., Sánchez-Menguiano, L., et al. 2017, A&A, 598, A32
Menéndez-Delmestre, K., Sheth, K., Schinnerer, E., Jarrett, T. H., & Scoville, N. Z. 2007, ApJ, 657, 790
Nair, P. B., & Abraham, R. G. 2010, ApJS, 186, 427
Noll, S., Burgarella, D., Giovannoli, E., et al. 2009, A&A, 507, 1793
Predehl, P., Andritschke, R., Arefiev, V., et al. 2021, A&A, 647, A1
Ramos Almeida, C., Tadhunter, C. N., Inskip, K. J., et al. 2011, MNRAS, 410, 1550
Ramos Almeida, C., Bessiere, P. S., Tadhunter, C. N., et al. 2012, MNRAS, 419, 687
Romeo, A. B., & Fathi, Y. 2015, MNRAS, 451, 3107
Romeo, A. B., & Fathi, Y. 2016, MNRAS, 460, 2360
Sakamoto, K., Okumura, S. K., Ishizuki, S., & Scoville, N. Z. 1999, ApJ, 525, 691
Salvato, M., Wolf, J., Dwelly, T., et al. 2022, A&A, 661, A3
Sellwood, J. A. 2014, Rev. Mod. Phys., 86, 1
Shlosman, I., Frank, J., & Begelman, M. C. 1989, Nature, 338, 45
Silva-Lima, L. A., Martins, L. P., Coelho, P. R. T., & Gadotti, D. A. 2022, A&A, 661, A105
Skibba, R. A., Masters, K. L., Nichol, R. C., et al. 2012, MNRAS, 423, 1485
Smethurst, R. J., Beckmann, R. S., Simmons, B. D., et al. 2023, MNRAS, 527, 10855
Somerville, R. S., & Davé, R. 2015, ARA&A, 53, 51
Stalevski, M., Fritz, J., Baes, M., Nakos, T., & Popovic, L. C. 2012, MNRAS, 420, 2756
Stalevski, M., Ricci, C., Ueda, Y., et al. 2016, MNRAS, 458, 2288
Stern, D., Assef, R. J., Benford, D. J., et al. 2012, ApJ, 753, 30
Treister, E., Schawinski, K., Urry, C. M., & Simmons, B. D. 2012, ApJ, 758, L39
Valiante, E., Smith, M. W. L., Eales, S., et al. 2016, MNRAS, 462, 3146
Walmsley, M., Lintott, C., Géron, T., et al. 2022, MNRAS, 509, 3966
Walmsley, M., Allen, C., Aussel, B., et al. 2023, J. Open Source Software, 8, 5312
Willett, K. W., Lintott, C. J., Bamford, S. P., et al. 2013, MNRAS, 435, 2835
Willett, K. W., Galloway, M. A., Bamford, S. P., et al. 2017, MNRAS, 464, 4176
Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, ApJ, 140, 1868
Wright, A. H., Hildebrandt, H., Kuijken, K., et al. 2019, A&A, 632, A34
Yang, G., Brandt, W. N., Vito, F., et al. 2018, MNRAS, 475, 1887
Yang, G., Boquien, M., Buat, V., et al. 2020, MNRAS, 491, 740
Yang, G., Boquien, M., Brandt, W. N., et al. 2022, ApJ, 927, 192
Yu, S.-Y., Kalinova, V., Colombo, D., et al. 2022, A&A, 666, A175
Zee, W.-B. G., Paudel, S., Moon, J.-S., & Yoon, S.-J. 2023, ApJ, 949, 91

Appendix A: Classification examples and additional confusion matrices

To visually illustrate the morphological characteristics of the galaxy classes defined by our fine-tuned `Zoobot` model and selection criteria (Sect. 3.1.2), Fig. A.1 presents random examples of galaxies classified as strongly barred ($S_{\rm bar}$), weakly barred ($W_{\rm bar}$), and unbarred discs ($U_{\rm bar}$). The figure displays HSC-SSP i-band image cutouts for these three classes, divided into examples of galaxies at lower redshifts ($z < 0.5$, top three rows) and higher redshifts ($z \geq 0.5$, bottom three rows).

Observing Fig. A.1, $S_{\rm bar}$ galaxies typically exhibit prominent, elongated bar structures at their centres, often connecting directly to spiral arms. $W_{\rm bar}$ galaxies display subtler bar-like features. These can include less elongated or lower contrast central structures, sometimes appearing as oval distortions or thickened inner discs that are suggestive of a bar but lack the definitive, sharp features of strong bars. Distinguishing these from unbarred galaxies or galaxies with prominent bulges can be challenging, reflecting the lower purity of this class as discussed in Sect. 3.1.2. The $U_{\rm bar}$ galaxies generally lack any clear, centrally dominant linear bar structure. They often show smooth central light profiles or distinct spiral arm patterns originating closer to the galactic nucleus without an intervening bar.

The examples at different redshifts demonstrate the challenge of morphological classification with increasing distance. Features become less distinct and more pixelated at higher z, which underscores the importance of carefully constructing the training sample for studies extending beyond the local Universe. Nevertheless, the `Zoobot` model, trained on GZ classifications which themselves span a range of redshifts, is capable of identifying these features across the redshift range of our study.

In this section, we also demonstrate how the definition of barred and unbarred galaxies we adopted affects the model's performance. Specifically, in Fig. A.2, we report the confusion matrices we get for the test set, using the $A_{\rm bar}$ and $U_{\rm bar}$ definitions given in Sect. 3.1.2 and the smooth and edge-on definitions in Sect. 3.1. As both matrices demonstrate, confusion between the $U_{\rm bar}$ class and the smooth class has considerably decreased, now falling below 10%.

Appendix B: SED AGN selection

In our previous work (LM24), we defined SED AGN as galaxies with an AGN fraction $f_{\rm AGN} \geq 0.05$, based on the MIR (3–30 μm) contribution derived from `CIGALE` SED fitting. While this threshold captures galaxies with non-negligible AGN contributions, we found that a more conservative selection improves the robustness of AGN identification. To this end, we performed a test on the initial sample of SED AGN (i.e. those with $f_{\rm AGN} \geq 0.05$) by re-running `CIGALE` with the AGN module turned off. We then compared the resulting reduced chi-square values (χ^2_ν) between the original fit (with AGN) and the new fit (without AGN). We found that galaxies with $f_{\rm AGN} \geq 0.1$ typically showed a significant degradation in fit quality when the AGN component was excluded, with a median increase in χ^2_ν exceeding 15%. In contrast, galaxies with $0.05 \leq f_{\rm AGN} < 0.1$ could be adequately fitted without an AGN component, showing only marginal or negligible increases in χ^2_ν.

This result indicates that $f_{\rm AGN} \geq 0.1$ provides a more secure threshold for identifying galaxies where the AGN contributes meaningfully to the infrared emission. Therefore, we adopt this stricter criterion in the present analysis.

Table C.1. Best-fit parameters of Eq. C.1 for X-ray and MIR AGN, and of Eq. C.2 for SED AGN.

AGN type	α	c	
X-ray	$16.2^{+4.87}_{-3.52}$	$0.287^{+0.009}_{-0.009}$	
MIR	$35.0^{+43.1}_{-9.46}$	$0.478^{+0.022}_{-0.023}$	
AGN type	**α**	**β**	**γ**
SED	$47.2^{+29.1}_{-14.0}$	$0.432^{+0.019}_{-0.018}$	$0.637^{+0.041}_{-0.030}$

Notes. Each parameter is estimated using a resampling process. The values reported are the median values of the 10 000 best-fit parameters and their 25^{th}–75^{th} percentile ranges.

Appendix C: Parametrisation of the merger fraction - AGN fraction and $L_{\rm disc}$ relations from LM24

In Sect. 4.2.1, we compare the bar fraction–AGN fraction ($f_{\rm bar}$–$f_{\rm AGN}$) relation derived in this work with the merger fraction–AGN fraction ($f_{\rm merg}$–$f_{\rm AGN}$) relation from our previous study (LM24). The overlaid lines in Fig. 13 representing the $f_{\rm merg}$–$f_{\rm AGN}$ trends were derived from LM24 as follows.

The $f_{\rm merg}$–$f_{\rm AGN}$ relation showed distinct behaviours for different AGN types and $f_{\rm AGN}$ regimes. For MIR and X-ray selected AGN, the data were best described by a simple power-law function of the form:

$$y = x^\alpha + c, \tag{C.1}$$

where y represents $f_{\rm merg}$, x is $f_{\rm AGN}$, and α and c are free parameters. For SED AGN, which exhibited a more complex trend, with a mild rising $f_{\rm merg}$ in the low $f_{\rm AGN}$ regime, a combination of two power laws was used:

$$y = x^\alpha + \beta \cdot x^\gamma, \tag{C.2}$$

where α, β, γ are free parameters.

To obtain robust parameter estimates, we employed a sampling method. First, the number of $f_{\rm AGN}$ bins used to calculate $f_{\rm merg}$ was varied (between six and 20 for X-ray and SED AGN, and up to 15 for MIR AGN due to smaller statistics). The merger fraction and its uncertainty were calculated for each bin. To account for these uncertainties, assumed to follow a Gaussian distribution, a bootstrapping method was applied by varying the $f_{\rm merg}$ values within their errors. The best-fit parameters for Equations C.1 and C.2 were then found for each bootstrapped dataset. This entire process was repeated 10 000 times. The final best-fit parameters reported in Table C.1 are the median values from these 10 000 fits, with uncertainties corresponding to their 25th and 75th percentile ranges. These parametrised relations are overlaid in Fig. 13.

For the parametrisation of the $f_{\rm merg}$–$L_{\rm disc}$ relation, we followed the same methodology, fitting the following linear function for all AGN types analysed:

$$f_{\rm merg} = \frac{1}{a}\left(\frac{L_{\rm disc}}{10^{42}\,{\rm erg\,s^{-1}}}\right) + c, \tag{C.3}$$

where a and c are free parameters. We calculated the best-fit parameters following the same methodology adopted for the $f_{\rm merg}$–$f_{\rm AGN}$ relation and reported the results in Table D.1, for each AGN type. These parametrised relations are overlaid in Fig. 15.

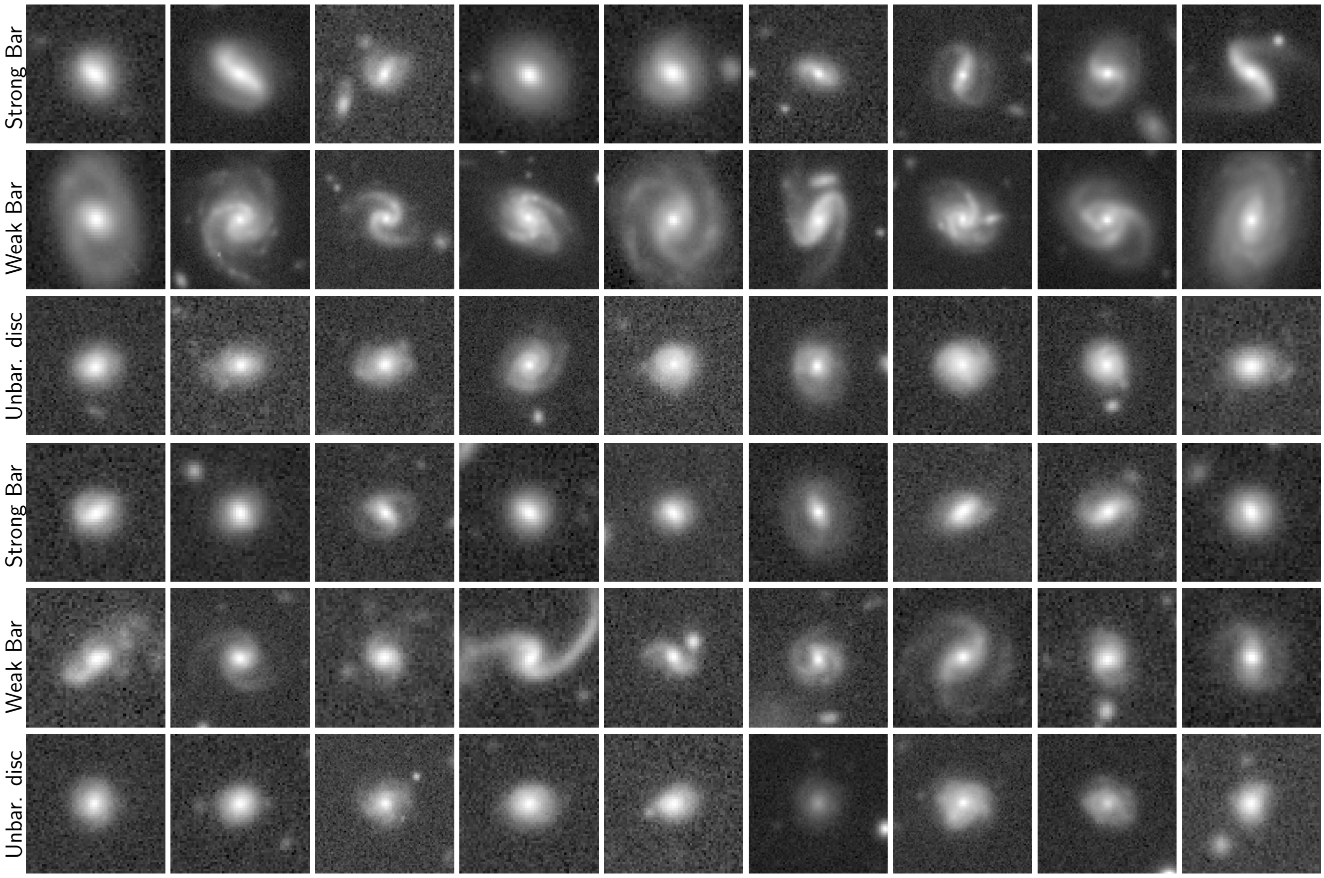

Fig. A.1. Random examples of barred and unbarred disc galaxies identified by the `Zoobot` model we trained. Top three rows: Examples at $z < 0.5$. Lower three rows: Examples at $z \geq 0.5$. Scaling matches Fig. 3.

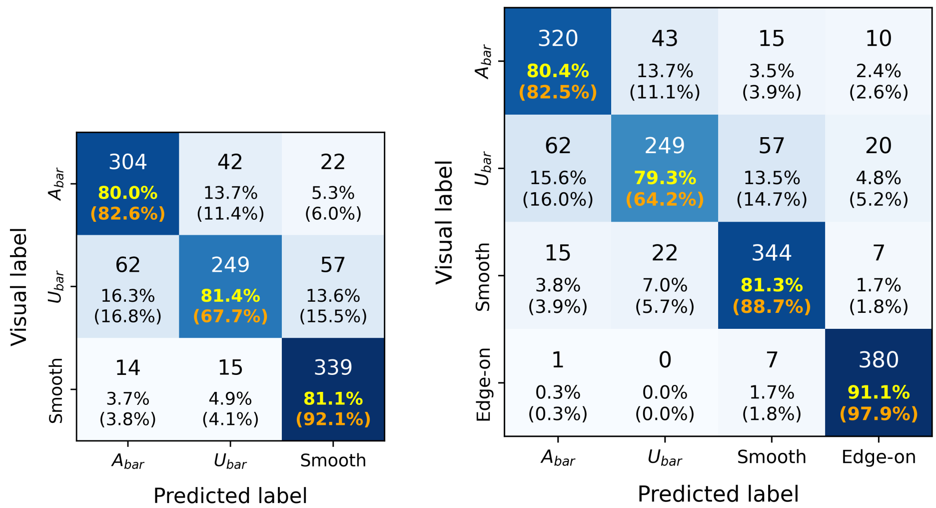

Fig. A.2. Confusion matrix for the `Zoobot` model, colour-coded relative to the total number of galaxies in the test set. Each cell shows the raw galaxy counts, the ratio over the number of predictions per class as a percentage (second line), and the ratio over the number of labels per class (in brackets). Along the diagonal, we highlight the precision of each class (yellow boldface text) and the recall (orange text). *Left panel*: Confusion matrix considering the $A_{\rm bar}$, $U_{\rm bar}$, and Smooth classes. *Right panel*: edge-on class.

Appendix D: $f_{\rm bar}$–$f_{\rm AGN}$ and $f_{\rm bar}$–$L_{\rm disc}$ including major mergers

Figure D.1 illustrates the $f_{\rm bar}$–$f_{\rm AGN}$ and $f_{\rm bar}$–$L_{\rm disc}$ relations presented in Sect. 4.2, recomputed after adding back the major mergers excluded in the main analysis. Comparing Fig. D.1 with Figures 13 and 15 (which excluded mergers entirely) reveals a good agreement, with both trends largely unaffected. The bar fraction still remains relatively flat or mildly decreases with increasing $f_{\rm AGN}$ for both AGN types, regardless of the bar strength. Similarly, $f_{\rm bar}$ continues to show a peak at intermediate luminosities ($L_{\rm disc} \approx 10^{44}\,{\rm erg\,s^{-1}}$) for X-ray AGN, followed by a decrease at higher luminosities. On the other hand, SED AGN show a monotonic decline towards higher luminosities.

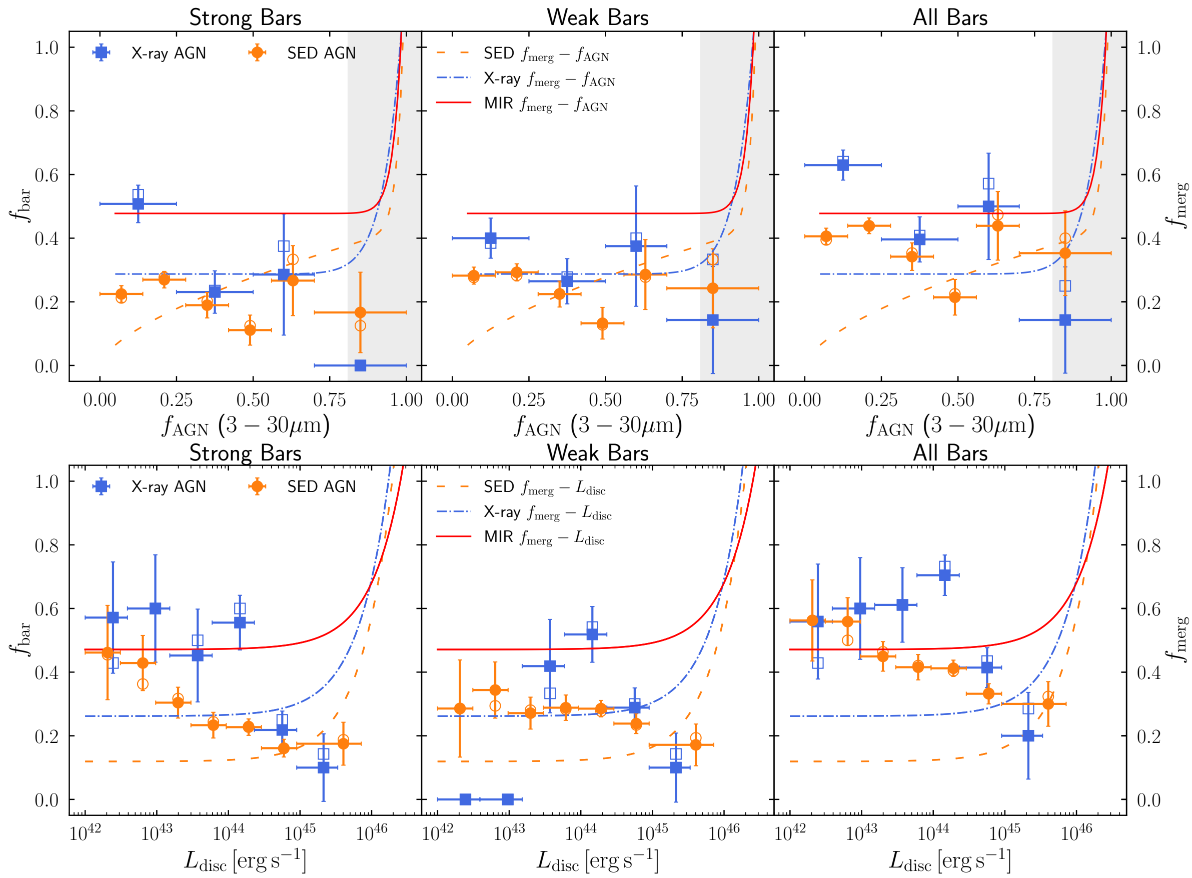

Fig. D.1. Bar fraction, $f_{\rm bar}$, as a function of $f_{\rm AGN}$(3–30 μm) (*top*) and $L_{\rm disc}$ (*bottom*) revised after including major mergers in the analysis. Markers and lines follow the same style as Figures 13 and 15. Empty symbols show the results presented in Figures 13 and 15.

Table D.1. Best-fit parameters of Eq. C.3 for MIR, X-ray, and SED AGN.

AGN type	a	c
MIR	$47.7^{+44.1}_{-8.3} \times 10^3$	$0.47^{+0.03}_{-0.03}$
X-ray	$23.2^{+2.3}_{-2.9} \times 10^3$	$0.26^{+0.01}_{-0.01}$
SED	$21.9^{+2.4}_{-6.0} \times 10^3$	$0.119^{+0.002}_{-0.003}$

Notes. Each parameter is estimated using a resampling process. The values reported are the median values of the 10 000 best-fit parameters and their 25^{th}–75^{th} percentile ranges.