
Identifying AGN host galaxies with convolutional neural networks

Ziting Guo

Yale-NUS College
#01-220, 16 College Avenue West
Singapore 138527
guo.ziting@aya.yale.edu

John F. Wu *

Space Telescope Science Institute
3700 San Martin Drive
Baltimore, MD 21218
jowu@stsci.edu

Chelsea E. Sharon

Yale-NUS College
#01-220, 16 College Avenue West
Singapore 138527
chelsea.sharon@yale-nus.edu.sg

Abstract

Active galactic nuclei (AGN) are supermassive black holes with luminous accretion disks found in some galaxies, and are thought to play an important role in galaxy evolution. However, traditional optical spectroscopy for identifying AGN requires time-intensive observations. We train a convolutional neural network (CNN) to distinguish AGN host galaxies from non-active galaxies using a sample of 210,000 Sloan Digital Sky Survey galaxies. We evaluate the CNN on 33,000 galaxies that are spectrally classified as composites, and find correlations between galaxy appearances and their CNN classifications, which hint at evolutionary processes that effect both galaxy morphology and AGN activity. With the advent of the Vera C. Rubin Observatory, *Nancy Grace Roman Space Telescope*, and other wide-field imaging telescopes, deep learning methods will be instrumental for quickly and reliably shortlisting AGN samples for future analyses.

1 Introduction

Nearly all galaxies are theorized to contain a supermassive black hole at the galactic center, but only in certain galaxies do interstellar gas and dust accumulate around the supermassive black hole to form a bright, disk-shaped structure called active galactic nucleus (AGN). Most galaxy evolution models suggest that AGN regulates the rate of star formation in a galaxy.

As AGN are structurally small relative to their host galaxies, the identification of AGN traditionally relies on expensive spectroscopic analysis. The BPT diagram (Baldwin et al., 1981) plots the $[\text{OIII}]/\text{H}\beta$ emission line flux ratio against that of $[\text{NII}]/\text{H}\alpha$, to distinguish AGN from non-active galaxies. In between active and non-active galaxies are the spectral composites with intermediate emission-line ratios, such as ordinary AGN whose spectra may have been contaminated by nearby star-forming regions (Ho & Filippenko, 1993). While essential for AGN identification, spectroscopy takes $\sim 10^3$ times more time than imaging. It may be possible to identify AGN based on galaxy images alone, given AGN's connection to star formation. Galaxy evolution is theorized to depend strongly on the co-evolution between the central supermassive black hole and the rest of the galaxy,

*Secondary affiliation: Department of Physics & Astronomy, Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218

and AGN presence may therefore be linked to galaxy morphology. Kauffmann et al. (2003)’s study on galaxies from the Sloan Digital Sky Survey (SDSS) reveals patterns in the physical properties of AGN host galaxies, showing promise for an image-based AGN identification method.

Astronomers have already begun employing data-driven methods to search for AGN in wide-field imaging surveys. For example, Chang et al. (2021) identify various classes of AGN candidates using catalogues of optical and near-infrared photometry. Another work found that optically obscured (Type I) and unobscured (Type II) AGN can be distinguished from each other using features derived from X-ray observations (Falocco et al., 2022). Some AGN are variable, and Pasquet-Itam, J. & Pasquet, J. (2018) train a convolutional neural network to identify quasar candidates from photometric time-series data. Finally, Holwerda et al. (2021) have trained a CNN to predict galaxy spectra directly from image cutouts (Wu & Peek, 2020), and have estimated the BPT classification of a galaxy by extracting lines flux ratios from the CNN-derived spectrum. We present results that are most similar to this last work, except that we directly estimate BPT classifications from image cutouts.

In this work, we show that a CNN can detect the presence of AGN from its host galaxy’s appearance. We evaluate our trained CNN on galaxies spectroscopically labelled as composite, in order to determine whether they are more likely to be star-forming non-active objects, or AGN hosts. The results support our physical intuition: smaller, bluer composite systems are more often predicted to be star-forming, while redder bulge-dominated composite systems are more often predicted to host AGN. Using only broadband imaging, our method can be used to efficiently prioritize AGN candidates for follow-up spectroscopy in future imaging-only sky surveys, which will capture more galaxy images than previous sky surveys combined.

2 Methodology

We select galaxy images from the SDSS Data Release 8 (DR8) spectroscopic catalog (Brinchmann et al., 2004; Tremonti et al., 2004a; Kauffmann et al., 2003). We download *gri*-band 160×160 -pixel SDSS image cutouts at a pixel scale of $0.262''$ from the Legacy Survey website (Dey et al., 2019). We require that galaxies in our sample be similar in size in final images by restricting galaxies to spectroscopic redshift $0.02 < z < 0.3$. Following the selection criteria from Kauffmann et al. (2003), each of the four emission lines, $H\alpha$, $[NII]$, $H\beta$, $[OIII]$, must have a non-zero flux with signal-to-noise ratio $(S/N) > 3$. Our ground truth labels of whether a galaxy contains AGN were obtained by BPT-based spectroscopic analysis in Kauffmann et al. (2003); we note that these classifications are model-dependent and only represent an estimate of the ground truth.

Composite and unclassifiable galaxies are excluded from training and validation. Of the remaining 210,332 galaxies, $\sim 21\%$ have AGN. Following a 80-20 train-validation split, we augment images in the training set to improve our model’s robustness. A good CNN should be able to make accurate predictions despite scatter in the central galaxy’s location, size and viewing angle. We stochastically apply an augmentation pipeline consisting of rotation, reflection, warping, and resizing to galaxy images in the training set, for which the augmentation degrees were optimized in preliminary experiments.

Our baseline model is *resnet18*, an 18-layer residual CNN, with ReLU replaced by Mish (Misra, 2019). With a batch size of 64, and *ranger* (Wright, 2019), a combination between *rectified Adam* (Liu et al., 2019) and *lookahead* (Zhang et al., 2019), as optimizer, the model learns the relationship between galaxy morphology and AGN presence by minimizing the cross-entropy loss between its predictions and ground truth labels. Learning rates are scheduled by *fastai*’s default learning rate scheduler, *1cycle* policy (Smith, 2018). These hyperparameter choices have shown to speed up convergence and improve accuracy in other astrophysical computer vision tasks (Wu et al., 2022).

3 Results

3.1 CNN classifications on active and non-active galaxies

We report that *resnet18* classifies active and non-active galaxies with 0.894 ± 0.001 accuracy, 0.738 ± 0.020 precision, 0.765 ± 0.046 recall, and 0.750 ± 0.012 F1-score. Uncertainties are determined by repeating the methodology using five different random seeds. Our experiments confirm




Figure 1: Spectroscopically classified AGN hosts and non-active galaxies from the validation set (*left*), and composites (*right*), as defined in Kauffmann et al. (2003). Panel (a), *left* shows ground truth active galaxies which have been classified as active by the CNN, and Panel (b), *left* shows correctly classified non-active galaxies. Panels (a) and (b), *right* are composite galaxies which have been classified as active and non-active respectively. Panel (c), *left and right* shows ambiguous classifications, for which CNN predicts equal probability for that it is active and that it is non-active. The small, reddish ambiguous objects in Panel (c) have features intermediate between active and non-active classifications.

the usefulness of morphological information in identifying active galaxies, and in turn, the connection between galaxy morphology and AGN presence.

Examples of typical active and non-active galaxies as identified by CNN are displayed in Figure 1, left. Panels (a) and (b), left show correctly classified active and non-active galaxies respectively. There is a clear distinction in morphology between what the model perceives as a typical AGN host and what as a typical non-active galaxy. Correctly classified AGN host galaxies in Figure 1(a) are redder and bigger (both in apparent and physical sizes) with dominant central bulges, which is consistent with Kauffmann et al. (2003)’s observation that nearly all dust-unobscured AGN reside in massive galaxies with little star formation. On the other hand, galaxies correctly classified as non-active are smaller star-forming galaxies with blue bulges. Among objects with highest-confidence CNN predictions, i.e. correct classifications with 10% least cross entropy losses, correctly classified active galaxies have median color index $g - i = 1.450$ and Petrosian r -band half-light radius $R_{50} = 3.195$ arcsec, whereas correctly classified non-active galaxies have median $g - i = 0.698$ and $R_{50} = 2.778$ arcsec. These visual patterns are consistent with the black hole mass - luminosity relation (Woo & Urry, 2002) and the mass-metallicity relation (Tremonti et al., 2004b).

Despite the presence of unrelated objects and artefacts in galaxy images in Figure 1, the CNN learns to make its predictions for the central galaxies, invariant to other unrelated objects in the field of view.

3.2 Using the trained CNN to classify composite galaxies

One of the key advantages of a machine learning based AGN identification method is to classify composite galaxies that traditional spectroscopic analyses fail to differentiate. Examples of composite galaxies likely to be active and non-active respectively, as well as those ambiguous to the CNN, are displayed in Figure 1, right.

We evaluate our trained CNN on the dataset of spectroscopic composites, while keeping in mind that composite galaxies are not represented at all in our training set. Unsurprisingly, composite galaxies classified as AGN by the CNN in Figure 1 (b) greatly resemble the typical AGN in Figure 1(a). Similarly, both typical non-active galaxies, and the composite galaxies identified to be non-active by the CNN (Figure 1 (c) and (d) respectively), lack dominant central bulges. Compared to predicting active and non-active galaxies, our CNN is less confident in predicting composite galaxies, suggesting that spectroscopically determined composites are also morphological intermediates.

4 Discussion




Figure 2: CNN predictions on galaxies spectroscopically classified as composites are shown on the BPT diagram in the main panel. Normalized histograms of composite galaxies’ $[\text{O III}]/\text{H}\beta$ and $[\text{N II}]/\text{H}\alpha$ line ratios are shown in the right and top panels respectively. Orange dots and lines represent spectroscopic composites the CNN identifies as active, whereas blue ones are classified as non-active. The central locus is bound by the two parabolas in Kewley et al. (2001) and Kauffmann et al. (2003).

Our work is the first application of CNN to identify AGN directly from optical images. We benchmark our CNN’s performance against a baseline *xgboost* model trained on *gri* photometric data that is similar to the approach by Chang et al. (2021). Again, we obtain the relevant photometric information from SDSS. Our *resnet18* trained on optical images outperforms the *xgboost* model in all chosen metrics (the best *xgboost* achieved 0.881 ± 0.001 accuracy, 0.726 ± 0.005 precision, 0.694 ± 0.007 recall, 0.710 ± 0.003 F1-score). *xgboost* performs on par with the CNN only when we include full *ugriz* photometry and Petrosian radius. This comparison indicates that galaxy size, in addition to color, is also important for identifying AGN in our SDSS sample.

In Figure 2, we show CNN classifications of spectroscopic composite galaxies in a BPT diagram. A $2d$ Kolmogorov-Smirnov test reveals that CNN-predicted active and non-active galaxies have different distributions ($p < 10^{-3}$), confirming our expectation that objects classified as active lie closer to the right-hand side of the diagram (the region of spectroscopically defined AGN), and objects classified

as non-active lie closer to the left-hand side. We note that $[\text{NII}]/\text{H}\alpha$ is also a tracer of gas-phase metallicity, and that the spectroscopic “ground truth” classifications are also metallicity-dependent. It has also been shown that metallicity can be predicted directly from galaxy image cutouts (Wu & Boada, 2019), which supports our thesis that galaxies’ morphologies are critical for identifying their physical properties.

Our deep learning approach can efficiently select AGN host galaxies from future imaging-only sky surveys. The Vera C. Rubin Observatory will go deeper and have higher resolution than SDSS. It will also detect many time-varying AGN at optical wavelengths. The *Nancy Grace Roman Space Telescope* will have exquisite resolution, and survey many AGN host galaxies out to higher redshifts. It will have tremendous low-surface brightness sensitivity, enabling better morphological characterization of extended galaxy disks, tidal debris, and other interstellar gas and dust. Using a comprehensive, CNN-selected sample of AGN candidates, we astronomers will be able to robustly study the co-evolution of galaxies and their supermassive black holes.

We have demonstrated that a CNN can identify AGN in low-redshift galaxies solely from optical imaging. We note that AGN hosts and non-active galaxies can overlap in the space of line flux ratios, particularly in the composite regime, but our machine learning technique hints at a way to separate probable AGN hosts from non-active galaxies on the basis of galaxy morphology. Future wide-field imaging surveys, such as the Legacy Survey of Space and Time by the Vera C. Rubin Observatory, and the High Latitude Wide Area Imaging Survey by the *Nancy Grace Roman Space Telescope*, can benefit from efficient selection of AGN candidates (or other interesting phenomena) by using CNNs.

Broader Impact

Like most deep learning projects, the process of training our model, required substantial computational resources and contributed to global warming.

For the positive impact, by keeping our code open-source, we have continued to build on open science with archival data sets.

Acknowledgements and Disclosure

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

Code for this project can be found at <https://github.com/cherryquinnlg/agn-convnets>.

References

- Baldwin, J. A., Phillips, M. M., & Terlevich, R. 1981, *PASP*, 93, 5, doi: 10.1086/130766
- Brinchmann, J., Charlot, S., White, S. D. M., et al. 2004, *MNRAS*, 351, 1151, doi: 10.1111/j.1365-2966.2004.07881.x
- Chang, Y.-Y., Hsieh, B.-C., Wang, W.-H., et al. 2021, *ApJ*, 920, 68, doi: 10.3847/1538-4357/ac167c
- Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *AJ*, 157, 168, doi: 10.3847/1538-3881/ab089d
- Falocco, S., Carrera, F. J., & Larsson, J. 2022, *MNRAS*, 510, 161, doi: 10.1093/mnras/stab3435

Ho, L. C., & Filippenko, A. V. 1993, *Astrophysics and Space Science*, 205, 19, doi: 10.1007/BF00657952

Holwerda, B. W., Wu, J. F., Keel, W. C., et al. 2021, *ApJ*, 914, 142, doi: 10.3847/1538-4357/abffcc

Kauffmann, G., Heckman, T. M., Tremonti, C., et al. 2003, *Monthly Notices of the Royal Astronomical Society*, 346, 1055, doi: 10.1111/j.1365-2966.2003.07154.x

Kewley, L. J., Dopita, M. A., Sutherland, R. S., Heisler, C. A., & Trevena, J. 2001, *ApJ*, 556, 121, doi: 10.1086/321545

Liu, L., Jiang, H., He, P., et al. 2019, arXiv e-prints, arXiv:1908.03265. <https://arxiv.org/abs/1908.03265>

Misra, D. 2019, arXiv e-prints, arXiv:1908.08681. <https://arxiv.org/abs/1908.08681>

Pasquet-Itam, J., & Pasquet, J. 2018, *A&A*, 611, A97, doi: 10.1051/0004-6361/201731106

Smith, L. N. 2018, arXiv e-prints, arXiv:1803.09820. <https://arxiv.org/abs/1803.09820>

Tremonti, C. A., Heckman, T. M., Kauffmann, G., et al. 2004a, *ApJ*, 613, 898, doi: 10.1086/423264

—. 2004b, *ApJ*, 613, 898, doi: 10.1086/423264

Woo, J.-H., & Urry, C. M. 2002, *ApJ*, 579, 530, doi: 10.1086/342878

Wright, L. 2019, New Deep Learning Optimizer, Ranger: Synergistic combination of RAdam LookAhead for the best of both, Medium. <https://lessw.medium.com/new-deep-learning-optimizer-ranger-synergistic-combination-of-radaml-lookahead-for-the-best-of-2dc83f79a48d>

Wu, J. F., & Boada, S. 2019, *Monthly Notices of the Royal Astronomical Society*, 484, 4683, doi: 10.1093/mnras/stz333

Wu, J. F., & Peek, J. E. G. 2020, arXiv e-prints, arXiv:2009.12318. <https://arxiv.org/abs/2009.12318>

Wu, J. F., Peek, J. E. G., Tollerud, E. J., et al. 2022, *ApJ*, 927, 121, doi: 10.3847/1538-4357/ac4eea

Zhang, M., Lucas, J., Ba, J., & Hinton, G. E. 2019, *Advances in Neural Information Processing Systems*, 32

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? **[Yes]** See Section 2.
- Did you include the license to the code and datasets? **[No]** The code and the data are proprietary.
- Did you include the license to the code and datasets? **[N/A]**

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]**
 - (b) Did you describe the limitations of your work? **[Yes]**
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]**
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[N/A]**
 - (b) Did you include complete proofs of all theoretical results? **[N/A]**
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** Code will be made public in a GitHub link upon acceptance of the paper.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See Section 2.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]** See Section 3. Retrained the model with different random seeds.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** 10 epochs on one Tesla V100-SXM2-32GB GPU.Nvidia V100-32GB per training session.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]** See Section2 and Section 4.
 - (b) Did you mention the license of the assets? **[Yes]** See Section 4.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[Yes]** Code will be made public in a GitHub link upon acceptance of the paper.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[Yes]** All data and code in this project are open-source, and we made sure to accredit their authors.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**

- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]