Exploring the Universe with Artificial Intelligence
Proceedings IAU Symposium No. 397,
doi:10.1017/xxxxx

Modern Deep Learning Approaches for Galaxy Classification

Dimitrios Katsaros, Nasia Zacharia, Stamatia-Danae Kravariti, Dimitrios Papakostas

University of Thessaly, Greece
dkatsar@uth.gr, {nasiazacharia,sdkravariti,jim.papakostas}@gmail.com

Abstract. One of the significant problems in cosmology is the morphological classification of galaxies in order to gain insights about the origin and evolution of our Cosmos. Here, we investigate the classification accuracy of several modern convolutional neural network architectures (some of them pretrained on non-astronomical data and others trained from scratch to galaxy data), and of transformers with primary goal to understand their performance on conditions of fast training (i.e., in a few rounds). We experimented with Transfer Learning, and Vision Transformers (ViT). We contrasted their performance with that of traditional Convolutional Neural Networks (CNN) methods, i.e., AlexNet, and VGG. Our results show a 75% classification accuracy (by the ViT and VGG), beating the rest of the competitors. Moreover, one pretrained VGG model showed also excellent performance.

Keywords. Galaxy, Morphology Classification, Convolutional Neural Networks, Transformers, Pretrained Neural Networks, Deep Learning

1. Introduction

During the last two decades we have witnessed tremendous advances in computing power and data collection capabilities accelerating and transforming how science and technology progresses with astrophysics being significanlty impacted by this progress. The operation of new scientific instruments such as the JWST pose new challenges to our ability to extract knowledge from unprecedented vast volumes of collected data. The *Astroinformatics* discipline, which employs Deep Learning (DL), big data management and distributed computing for serving astronomy's goals, has emerged as a result of this wave of scientific evolution. In particular, the ability of deep learning algorithms, which are based on neural networks, to extract knowledge directly from raw data, without sophisticated preprocessing, along with their performance behaviour, which gets significant improvements with larger and larger datasets, make them ideal for being deployed in astroinformatics. Thus, neural networks have been extensively applied for gravitational wave detection Nousi et al. (2023), for light curve classification using the transformer-based Large Language Models Li et al. (2025), and even for detecting possible signs of technological activities from nearby stars Ma et al. (2023).

In cosmology, the creation of galaxy catalogs is a significant goal related to questions concerning the origin, age and evolution of the Cosmos. Morphological classification of galaxies is maybe the most basic information when creating these catalogs and comprises an important steps to study the formation, structure and evolution of galaxies. Starting from the famous Hubble classification system Hubble (1936) to more detailed classification systems e.g., de Vaucouleurs, the aim of these efforts is to classify galaxies into one of (sub)categories, usually in a probabilistic manner. Various methods have been applied/devised for achieving this goal; photometric-based, pure image-based, and so on. We focus here on computerized learning-based techniques; for instance Mukundan et al. (2024) uses a machine (but not deep) learning-based method for performing the classification, namely *k*-nearest neighbors. Barchi et al. (2020) investigated the performance of two somewhat outdated neural architectures, namely ResNet and GoogleNet Goodfellow et al. (1916) along with a decision tree and Support Vector Machine; Cao et al. (2024) used a vision transformer, but they worked on a five-class classification task, which apparently is a less challenging relatively, since the complexity of the task increases with the increase of the number of classes. Variawa et al. (2022) investigated transfer learning but for the purposes of comparing and combining crowd-sourcing to more traditional methods in galaxy classification.

In this article, we focus exclusively on deep learning algorithms, namely convolutional neural networks, transformer-based neural networks and transfer learning for performing morphological classification of galaxies, while deploying minimal image preprocessing. We wish to investigate a) whether these methods can reap significant performance gains with a few training rounds and b) also the impact of techniques such as early stopping, which further accelerates the training on the performance.

The rest of this paper is organized as follows: section 2 describes the deep learning neural architectures we have implemented and evaluated and section 3 describes the datasets used to evaluate these methods. Section 4 presents and explains the results obtained and finally, section 5 summarizes our findings and concludes the article.

2. Deep learning algorithms

We implemented and evaluated several convolutional neural networks (CNNs); some baseline ones, namely the *AlexNet* Goodfellow et al. (1916), more advanced ones, namely a *VGG* family member, some pretrained models and a transformer. *AlexNet* is comprised by eight layers: the first five are convolutional layers, some of them followed by max-pooling layers, and the last three are fully connected layers. Our *VGG* variant has a total of four convolutional layers with 3×3 kernel filters, two 2D pooling layers with 2×2 kernel filters, and two fully connected layers. The activation function throughout the hidden layer is ReLU, in the first fully connected layer, the activation function is ReLU while in the second layer the activation function is SoftMax. The pretrained models include *EfficientNetB3* and a pretrained *VGG-16* model. *EfficientNetB3* is a CNN that employs *Mobile Inverted Bottleneck Convolution* blocks, including depthwise separable convolutions with varying kernel sizes (3×3 and 5×5), and Squeeze-and-Excitation (SE) blocks for feature extraction; it has a $300 \times 300 \times 3$ input shape, a depth of 210 layers. *VGG16* has thirteen convolutional layers, five Max Pooling layers, and three Dense layers, which sum up to 21 layers, but it has only sixteen weight layers, i.e., learnable parameters layers. Our pretrained models were trained on the ImageNet dataset. Our deployed Vision Transformer (ViT) is the classic one reported in Dosovitskiy et al. (2022).

3. Dataset and evaluation setting

The Galaxy Zoo (GZ) project invited volunteers to classify galaxies from the Sloan Digital Sky Survey based on the morphology of the given color images Lintott et al. (2008). Subsequently, a Kaggle competition centered around the GZ was developed. This paper uses the dataset from the GZ project available on the Kaggle platform, which has been used by other studies as well, e.g., Cao et al. (2024); Variawa et al. (2022). In particular, our dataset contains 61578 images of JPG format for training, and 79975 images of JPG format for testing purposes. There is a solution file of training images that shows the probability distribution for each image. We performed minimal preprocessing, simply to reduce the size of each image. The original shape of each image is (424×424) while taking advantage of the fact that the point of interest is at the center, we just crop and downsampled the images. The final shape of each image became (60×60).

We used as loss function the Binary Cross Entropy and the Adam optimizer with learning rate equal to 0.001.

4. Classification results

Our first observation is that the loss for all our architectures reduces exponentially fast (neither linearly, nor step-wise). We show it only for AlexNet (right plot of Figure 1). This is an encouraging fact, since it implies that we can get very good performance in the first few rounds of training.

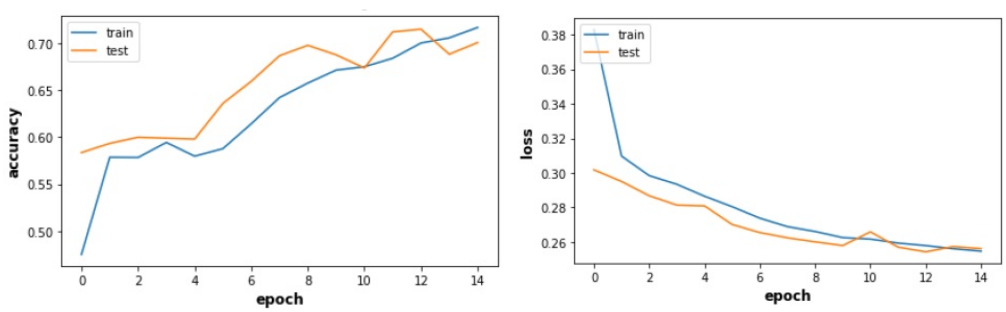

Figure 1: AlexNets' performance: (left) accuracy (right) loss.

Our second observation concerns issues of overfitting. As it can be seen in Figure 2, the technique of *early stopping* is very beneficial by improving the test accuracy and on the other hand improving convergence.

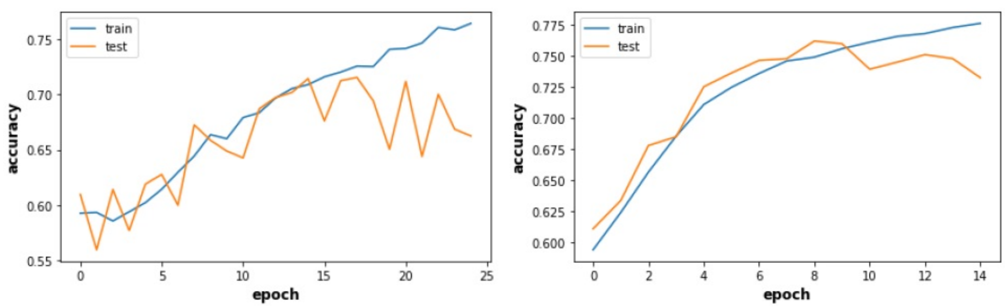

Figure 2: VGG's accuracy (left) without early stopping, (right) with early stopping.

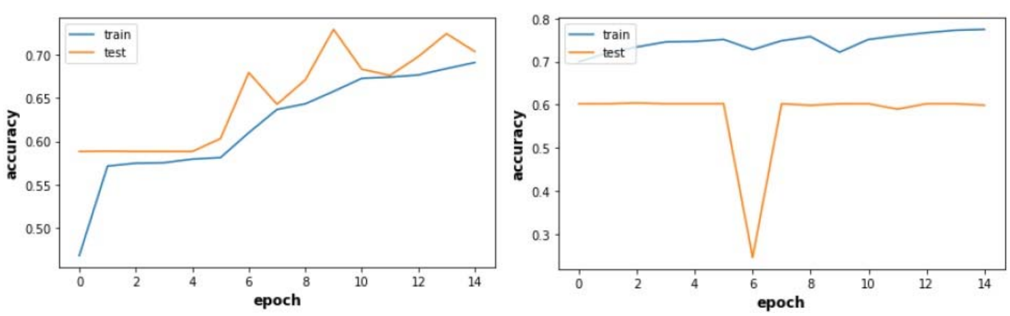

Figure 3: Pretrained models accuracy: (left) VGG16 (right) EfficientNet-B3.

As far as the accuracy is concerned, we observe in Figure 1(left), and Figures 2 – Figure 4 that with only a dozen training rounds, the examined methods are able to achieve around 75% accuracy, and the percentage keeps growing for some methods in subsequent rounds. Among the competitors, the Vision Transformer, achieves this performance from the very first few rounds.

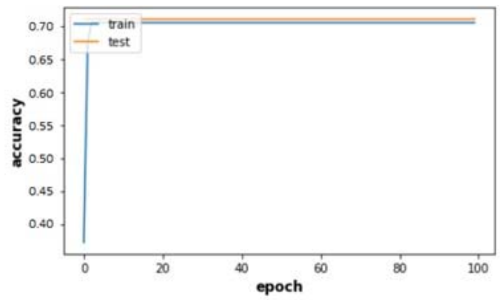

Figure 4: Vision Transformer's accuracy.

5. Discussion

Neural networks are very effective mechanisms for addressing important problems in astronomy and astrophysics. It is known however that neural networks need huge amounts of data in order to achieve significant performance. Here we investigated the question whether their deployment can reap performance gains with minimal training in the context of morphological galaxy classification. Our results strongly suggest that this is feasible.

The surprising and unexpected result is the performance of the pretrained VGG architecture, which – even though it has been trained on a dataset of different nature – showed almost identical performance to that of the top performers. This maybe implies that *transfer learning*, which can alleviate the burden of tremendous dataset sizes in astronomy/astrophysics, needs to be thoroughly investigated. We plan to extend our investigation along the lines of federated learning, so as to allow for training using geographically dispersed dataset repositories.

References

Barchi, P., de Carvalho, R., Rosa, R., Sautter, R., Soares-Santos, M., Marques, B., Clua, E., Goncalves, T., de Sa-Freitas, C., & Moura, T. 2020, Machine and deep learning applied to galaxy morphology - A comparative study. *Astronomy and Computing*, 30, 100334.

Cao, J., Xu, T., Deng, Y., Deng, L., Yang, M., Liu, Z., & Zhou, W. 2024, Galaxy morphology classification based on Convolutional vision Transformer (CvT). *Astronomy and Astrophysics*, 683, A42.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. 2022,. An image is worth 16*x*16 words: Transformers for image recognition at scale. Available at: https://arxiv.org/abs/2010.11929.

Goodfellow, I., Bengio, Y., & Courville, A. 1916,. *Deep Learning*. The MIT Press.

Hubble, E. 1936,. *Realm of the Nebulae*. Yale University Press.

Li, Y.-Y., Bai, Y., Wang, C., Qu, M., Lu, Z., Soria, R., & Liu, J. 2025, Deep learning and methods based on Large Language Models applied to stellar light curve classification. *Intelligent Computing*, 4, 0110.

Lintott, C., Schawinski, K., Slosar, A., Land, K., Bamford, S., Thomas, D., Raddick, M., Nichol, R., Szalay, A., Andreescu, D., Murray, P., & Vandenberg, J. 2008, Galaxy Zoo: Morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. *Monthly Notices of the Royal Astronomical Society*, 389(3), 1179–1189.

Ma, P., Ng, C., Rizk, L., Croft, S., Siemion, A., Brzycki, B., Czech, D., Drew, J., Gajjar, V., Hoang, J., Isaacson, H., Lebofsky, M., MacMahon, D., de Pater, I., Price, D., Sheikh, S. Z., & Worden, S. P. 2023, A deep-learning search for technosignatures from 820 nearby stars. *Nature Astronomy*, 7, 492–502.

Mukundan, K., Nair, P., Bailin, J., & Li, W. 2024, Automating galaxy morphology classification using k-nearest neighbours and non-parametric statistics. *Monthly Notices of the Royal Astronomical Society*, 533(1), 292–312.

Nousi, P., Koloniari, A., Passalis, N., Iosif, P., Stergioulas, N., & Tefas, A. 2023, Deep residual networks for gravitational wave detection. *Physical Review D*, 108(2), 024022–1–024022–10.

Variawa, M., van Zyl, T., & Woolway, M. 2022, Transfer learning and deep metric learning for automated galaxy morphology representation. *IEEE Access*, 10, 19539–19550.