

Galaxy Classification with an Attentional Convolutional Neural Network

Jai Misquith

Received August 9th, 2022

Accepted December 20th, 2022

Electronic access January 15th, 2023

While machine learning has been used to automate and improve processes in galaxy classification, advanced AI techniques like attention have yet to be implemented, despite significant success in its application in other fields. In this paper, we incorporate attention into a convolutional neural network designed to classify 10 different types of galaxies using 22,000 69x69 color images of galaxies. Our network uses a modified version of an attention network designed for CIFAR-10 images, which allows us to look at detailed statistics and heatmaps post-training and testing. We found that our network achieved a testing accuracy of 83.28% over 158 epochs. The advanced data allowed us to also realize that the network has very high capacity and could utilize the attention even better than it does given more training data and computational power. We conclude that attention is as applicable in astronomy as it is in many other fields, given the high testing accuracy, and we believe that our research will advise future networks in the realm of galaxy classification and possibly other stellar or planetary classification networks as well.

Introduction

Classifying galaxies has long been important to the field of astronomy. By knowing a galaxy's type, we can derive a plethora of information about it very easily without having to investigate further. For example, a galaxy's type is very closely correlated with its star formation rate and dynamical properties¹. Additionally, galaxy morphology provides a significant amount of information on the formation and evolution of galaxies, furthering our general knowledge of the universe. However, there are billions upon billions of potential galaxies to be classified, making it impossible for humans to classify all known galaxies. Fortunately, The classification of galaxies, when looked at properly, is a perfect example of what machine learning can be used to achieve. Humans are extremely adept at taking one look at a galaxy and picking out important features to accurately predict the type. This combination of relative simplicity and overwhelming population creates the perfect environment for a neural network to thrive. Machine learning has already been widely utilized in the field of astronomy to accurately classify galaxies². Current state-of-the-art models use Bayesian convolutional neural networks³. However, despite all preexisting experimentation, no one has attempted to incorporate a relatively new concept: attention.

To better study galaxy classification networks, this paper seeks to incorporate attention to help improve state-of-the-art results, advise future galaxy classification architectures, and highlight the contrast between galaxy image features used by humans and features used by machines. The heat maps generated from the attention model could provide valuable in-

sight as to how machines make their decisions when classifying a galaxy. It could also provide clues into how to better design galaxy classification architectures, as well as highlight any interesting differences between human and machine classification. This could end up helping human classification of galaxies when the lines between two categories get blurry. When applied to convolution neural networks (CNNs), attention forces the network to focus on the features of an image that will be most helpful in making the prediction. The research our model was based on, Jetley et al. (2018)⁴, details the strengths and application of one version of attention in a traditional classification CNN. Jetley et al. (2018) takes a traditional classification CNN and adds attention modules that take in lower-level features from the network at various layers, then trains the attention modules to focus on the important features in those lower-level layers, producing this idea of attention. It also showcases some examples of heatmaps the model can generate, as well as some testing and training data when their sample model was used on CIFAR-10 images. Ever since being popularized, attention has been used in other novel ways. In An et al. (2021), attention was used to construct a multiscale CNN using attention to classify lung nodules and breast cancer images⁵. Galaxy classification often involves picking out specific features among seemingly uninteresting images, and An et al. (2021) proves the ability of attention to classify homogenous groups of objects. Additionally, in Vázquez et al. (2020), attention was even used to model sentence pairs, an application similarly as abstract as galaxy morphology⁶. Given the wide variety of fields that attention has already been successfully applied in, there is compelling ev-

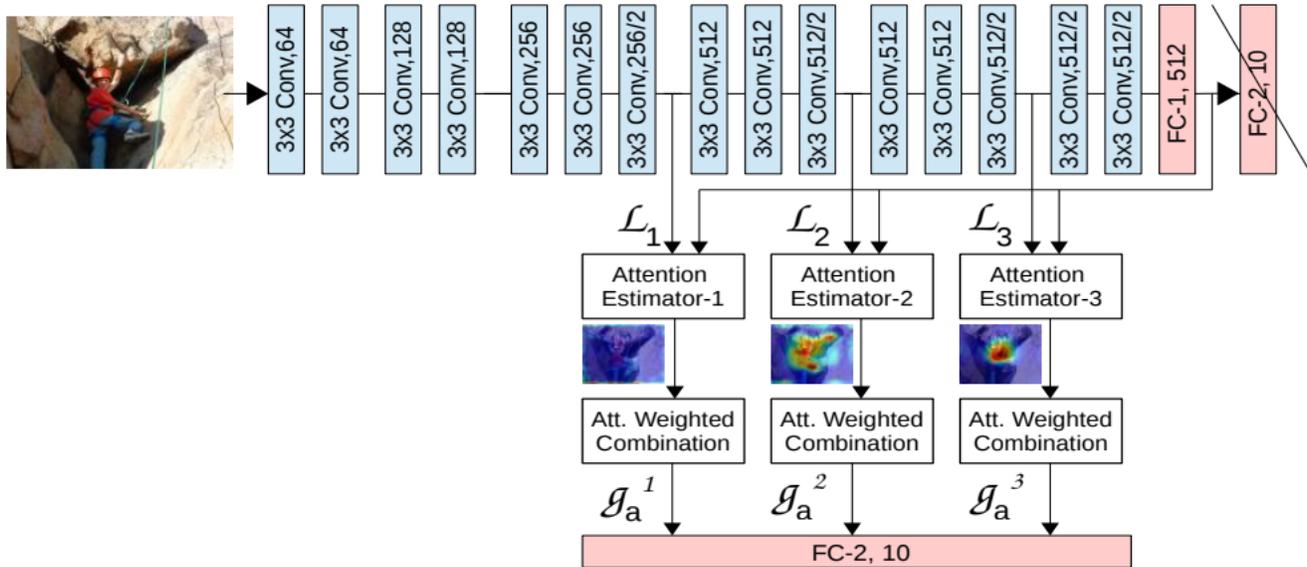


Fig. 1 The layout of the attentional convolutional neural network used.

idence that it should be applied experimentally for the task of galaxy classification. The network we used was a modified version of an attentional CNN designed for CIFAR-10 images. We then trained the model on 22 thousand images of galaxies across 10 classes for 158 epochs. We observed the various heatmaps and graphs produced by the model and found that it achieved a testing accuracy of 83.28%. Through analyzing the model’s output, we concluded that the model had a far greater capacity than we first thought. Its ability to achieve such a high accuracy with such limited data in such a short timespan suggests that attention could be the next breakthrough in not just galaxy classification, but astronomical classification as a whole.

Method

In traditional convolutional classification neural networks, images pass through multiple layers of convolution to pull out and distribute the important features of the image. Eventually, the image gets passed through a fully connected layer to turn the image into a class prediction, and this ultimately represents the output for one image. An attentional network uses attention estimators at regularly spaced intervals along the convolution to have the network relay what exactly it finds useful at each given time. A new fully connected layer takes in the information from the attention estimators to make its prediction, thus forcing it to use only features from the attention estimators to make the final prediction. This means that the

attention estimators are forced to narrow in on the more relevant features, because the neural network can only use them for predictions.

The data we used is the Galaxy 10 DECaLS Dataset⁷, which consists of around 22 thousand images of galaxies spread roughly equally into 10 different classifications. These images were in color and came as 69x69x3 pixels.

The images were modified to instead be 32x32 using PyTorch, as it significantly cut down on processing time and saved memory space. We used 19 thousand images for training and 3 thousand images for testing.

The layers of the network consist of 6 convolutional “blocks” (multiple similar convolution layers packaged together) that are connected with ReLU functions and MaxPooling2D layers to add a level of non-linearity. All convolutional layers use a 3x3 filter, and eventually lead into 2 fully connected layers to turn the data back into an image, then turn that image into a prediction.

What makes this network unique is the inclusion of attention. The model has attention estimators after convolution 7, 10, and 13. These estimators will attempt to pinpoint exactly what the model thinks is most important at each of these steps. Once the image is passed through the first fully connected layer, however, the model passes the image through the attention estimators and uses only the data garnered from them to make its prediction. [Fig 1] PyTorch was used to build and modify the network to fit these specifications. The network uses a batch size of 128, initial learning rate of 0.1 which was halved after every 25 epochs, and 300 total epochs. It was

Table 1

Class	Precision	Recall	F1-Score
Disturbed Galaxies	0.95	0.92	0.93
Merging Galaxies	0.98	0.99	0.98
Round Smooth Galaxies	0.97	0.98	0.98
In-between Round Smooth Galaxies	0.93	0.88	0.90
Cigar Shapes Smooth Galaxies	0.98	0.99	0.98
Barred Spiral Galaxies	1.00	1.00	1.00
Unbarred Tight Spiral Galaxies	0.99	0.99	0.99
Unbarred Loose Spiral Galaxies	0.97	0.93	0.95
Edge-on Galaxies without Bulge	0.95	0.96	0.95
Edge-on Galaxies with Bulge	0.98	0.97	0.97

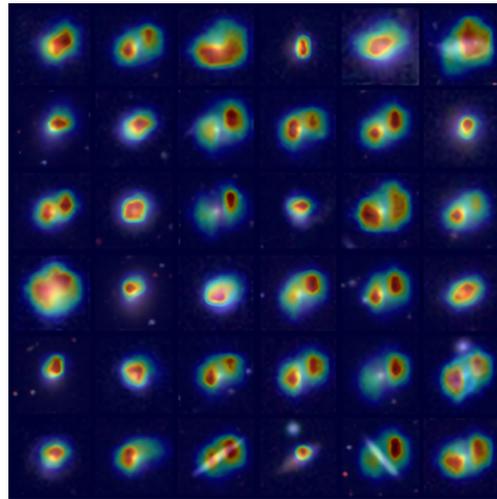


Fig. 2 The attention heatmaps for a sample of galaxies from epoch 42.

trained on a supercomputer with an Intel Xeon Platinum processor and 192 gigabytes of ram. The model was initially set at the stated 300 epochs, but while it was training, it began to produce a training accuracy of 100% at around epoch 50. After further noticing that the testing accuracy was not changing substantially past epoch 50, the decision was eventually made to halt training at epoch 158. Training continued to epoch 158 to ensure that the model did cap at its final accuracy value and ensure that we did not miss any spikes in accuracy right after epoch 50. Results were ultimately analyzed using Tensorboard.

Results

With our model, we achieved an eventual testing accuracy of 83.28% after training the neural network for 158 epochs. The model was able to fully learn the data in what seemed to be

only 50 epochs, and the heatmaps show that after those 50 epochs, the model relied less on utilizing attention because of the high capacity of the network. The network began to overfit to the data and began relying less on utilizing attention, which suggests future modifications such as including more training data or reducing the free parameters of the network.

Discussion

The model ultimately produced a testing accuracy of 83.28% after 158 epochs, and as previously stated, began producing a training accuracy of 100% starting at epoch 50. See Figure 2 for the model's heatmaps at roughly epoch 50.

The ability of the model to essentially fully train in 1/6 the time we expected proves the effectiveness of utilizing attention in a galaxy classification model. This has proved that the model has learned our data much quicker than expected

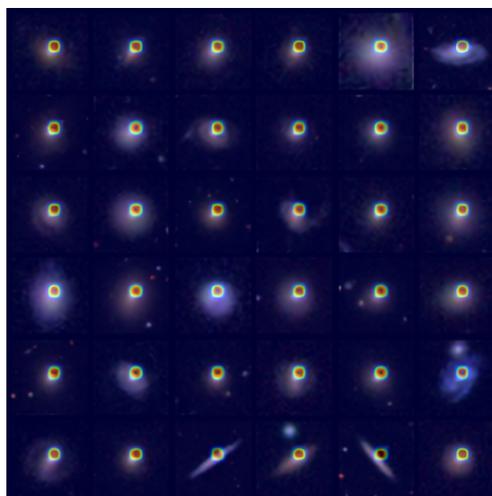


Fig. 3 The attention heatmaps for a sample of galaxies from epoch 67 onwards.

and is able to cap out at around 83%. [Fig 2] In fact, going any further than those 50 epochs shows that the network stops utilizing the attention, since the first attention estimator goes completely blank from epoch 67 onward. [Fig 3] The third estimator only looks at the very center for that same range, and does not update its attention any further. The model can now memorize the data with very few pixels, making many predictions using attention that are homogenous across different types of galaxies. This signals that the capacity is too high, and that the model has begun to overfit the data. The overfitting is also the explanation as to why the model seems to only look at the centers near the end of all 158 epochs. What we should really be looking at are the heatmaps before epoch 50, where the model looks in slightly different areas for each galaxy. [Fig 2] For example, for some galaxies with a significant amount of light surrounding them, the model attempts to almost disregard the center, and only analyzes that light.

The fact that the model learns the data so efficiently means a similar model could be applied for astrophysical objects other than galaxies. Stars, supernovas, and even black holes could be identified and classified in a similar way as our model, and we would expect the resulting output to be just as monumental.

An easy way to improve this model would be to incorporate a few data augmentation techniques to increase the amount of data we have. It has already been proven that the model learns the images incredibly quickly with just the limited dataset that was used, so if there was more, it could achieve an even higher accuracy than now. The model could also be improved by implementing regularization or normalization. Batch normalization could stabilize the learning process and reduce the number of training epochs necessary, reducing the chance of overfitting. Similarly, adding dropout layers could force the net-

work to learn different ways to classify a given galaxy instead of relying on the same weights each time, again reducing the chance of overfitting.

Our research has contributed to the overall progress of astronomical classification as a field, as we have proven that attention allows galaxy classification models to achieve a high accuracy at a fast speed, providing an insight into galaxy classification that was previously unknown. The model also learns the images very quickly, meaning a dataset of more images can be used without fear of increasing runtime by too much, and thus achieve a higher accuracy. Our research should also be an example of how attention can be used to great effect in the field of astronomy and could persuade future models in the field to incorporate attention as well, boosting overall accuracy and speed across the board.

Conclusion

We sought to test the application of attention in a galaxy classification network, and our model achieved a testing accuracy of 83.28% from 158 epochs. We present evidence via the statistics and heatmaps that the model has a capacity far greater than what we have tested. From these results, we can conclude that incorporating attention into a galaxy classification CNN works very well and could prove useful for future research in both this field and astronomical classification as a whole. It is likely that attention can be utilized in various other fields of astronomy, such as identifying black holes or distant supernovas, and could provide the boost that some models need to make major breakthroughs. Through our research, we have expanded the field of galaxy classification, and we hope that our results will pave the way for the application of attention

models in future astronomical research.

References

- 1 B. Lee, M. Giavalisco, C. Williams, Y. Guo, J. Lotz, A. Wel, H. Ferguson, S. Faber, A. Koekemoer, N. Grogin, D. Kocevski, C. Conselice, S. Wuyts, A. Dekel, J. Kartaltepe and E. Bell, *The Astrophysical Journal*, **774**, 47.
- 2 X.-P. Zhu, J.-M. Dai, C.-J. Bian, Y. Chen, S. Chen and C. Hu, *Astrophysics and Space Science*, **364**, year.
- 3 M. Walmsley, C. Lintott, T. Géron, S. Kruk, C. Krawczyk, K. Willett, S. Bamford, L. Kelvin, L. Fortson, Y. Gal, W. Keel, K. Masters, V. Mehta, B. Simmons, R. Smethurst, L. Smith, E. Baeten and C. Macmillan, *Monthly Notices of the Royal Astronomical Society*, **509**, 3966–3988.
- 4 S. Jetley, N. Lord, N. Lee and P. Torr, *Accessed*.
- 5 F. An, X. Li and X. Ma, *Oxidative Medicine and Cellular Longevity*, 1–12.
- 6 R. Vázquez, A. Raganato, M. Creutz and J. Tiedemann, *Computational Linguistics*, **46**, 387–424.
- 7 *Galaxy10 decals dataset. Galaxy10 DECals Dataset - astroNN 1.1.dev0 documentation*, <https://astronn.readthedocs.io/en/latest/galaxy10.html>, Retrieved December 11, 2022, from.