Dogon Skies, Machine Eyes: Celestial Image Classification with CNN and Dense Network

Aïssata Diop

Smith College

Introduction

Inspired by the rich celestial traditions of Dogon astronomy and the growing capabilities of machine learning, this project examines how neural network architectures shape our ability to classify astronomical images. The goal is to identify an efficient and reliable approach for distinguishing cosmic objects. Using supervised learning, we compare two models: a Convolutional Neural Network (CNN), designed to capture local spatial patterns and visual features such as edges, and a Dense Network, which treats each pixel independently without spatial context. Both models are trained on a binary classification task to differentiate between black holes and galaxies, allowing us to evaluate how architectural design influences performance in interpreting the universe through data.

Methodology

Model performance using a train/validation/test split to ensure robust generalization was evaluated. Both the Convolutional Neural Network (CNN) and Dense Network are trained using cross-entropy loss, with dropout layers incorporated to reduce overfitting. To compare architectures, we analyze accuracy, computational efficiency, and generalization behavior through learning curves. This framework allows us to assess how well each model learns from data and adapts to unseen examples, ultimately identifying the more effective approach for classifying astronomical images.

CAPACITY

Dense Network

- Fully connected layers
- Input: flattened image (no spatial structure)
- Learns global relationships between pixels

CNN

- Convolutional layers (3×3 kernels)
- Feature maps: 32 → 64 channels
- Captures edges, textures, and spatial hierarchies

VARIANCE

- Dense Network exhibits higher variance, reflected in larger gaps between training and validation performance
- CNN benefits from weight sharing and spatial structure, leading to more stable and consistent predictions across datasets

OPTIMIZATION

- Models are optimized using the Adam optimizer with a fixed learning rate
- Adaptive gradient scaling accelerates convergence, particularly in early training stages
- Cross-entropy loss guides parameter updates for binary classification

Results

CNN Predictions

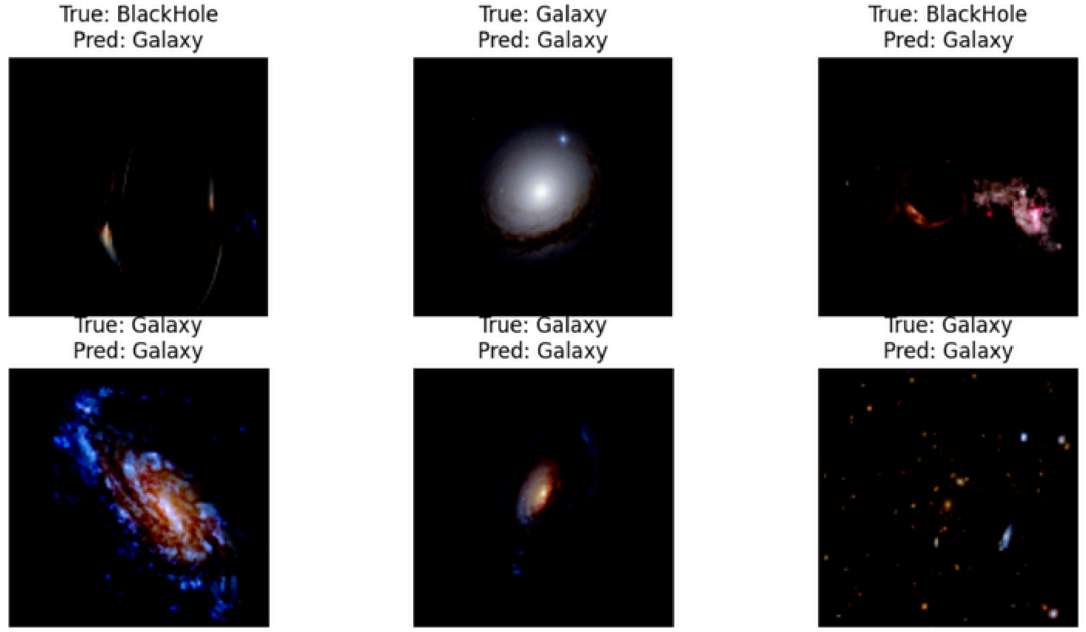

Dense Predictions

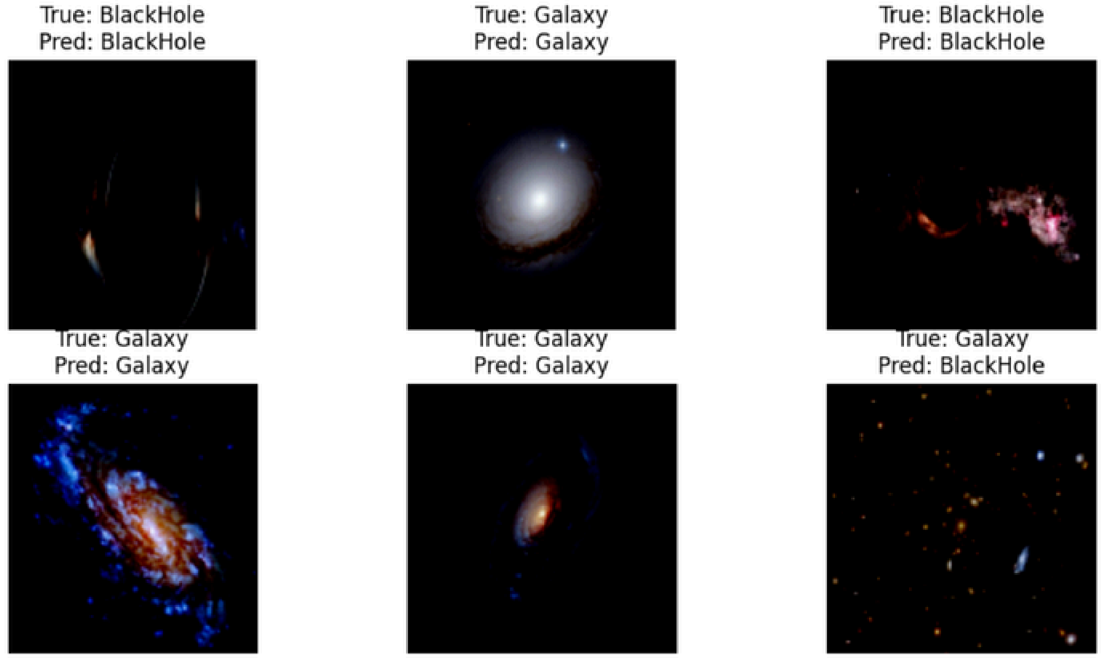

The CNN demonstrates stronger and more consistent performance in classifying astronomical images, correctly identifying key spatial features such as structure and brightness patterns. In contrast, the Dense Network shows more variability in its predictions, with several misclassifications likely due to its lack of spatial awareness. Learning curves further highlight this difference, as the CNN achieves higher validation accuracy and more stable convergence compared to the Dense model.

Conclusion

Loss & Accuracy - Dense

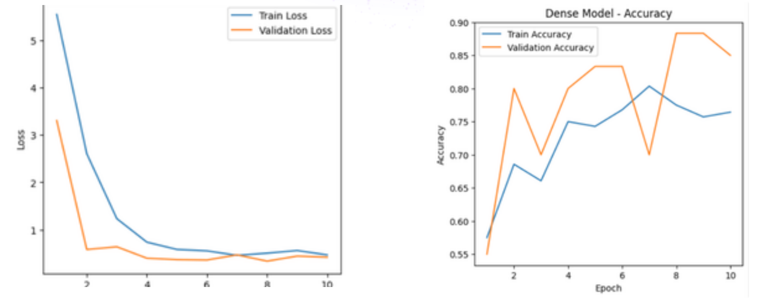

Loss & Accuracy - CNN

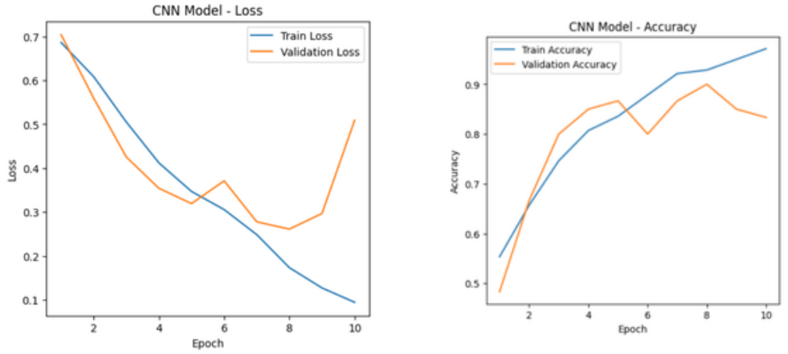

CNN beats Dense Network!

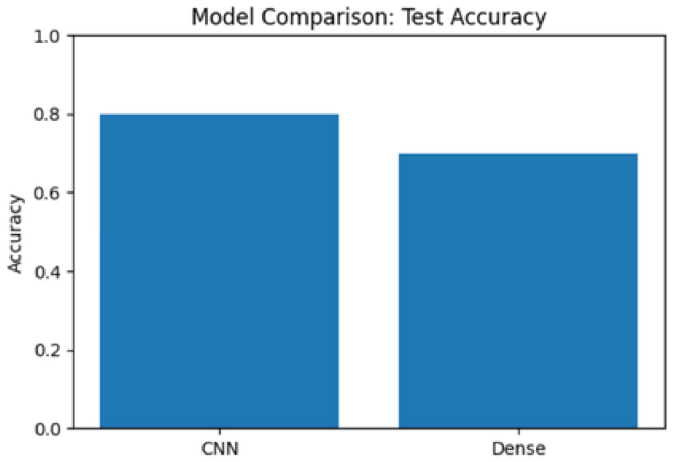

The CNN outperforms the Dense Network in both accuracy and generalization. While the CNN begins to slightly overfit after several epochs, it remains more robust due to its ability to capture spatial hierarchies through convolution and weight sharing. The Dense Network, lacks this inductive bias, shows higher variance and reduced sensitivity to important visual features. These results emphasize the importance of architecture choice when applying machine learning to astronomical image classification. Combining human curiosity about the cosmos with machine perception can lead to more effective ways of interpreting the universe.