CNNs for Bulk Material Defect Detection

Akshay Aravindan, Harrison Greenwood, Aakriti Varshney
{akshay14, hgreenw2, aakritiv}@stanford.edu
CS230 Deep Learning, Stanford University

Stanford
Mechanical Eng.

Abstract

Detecting defects in bulk materials is a problem for industries worldwide, and is currently solved by time consuming, manual inspection. This project is aimed at developing a model to detect defects in bulk steel and plastic. Ideally, it would lay groundwork for bulk material inspection across all engineering materials. Transfer learning was applied to pretrained ResNet18 and ResNet50 Convolutional Neural Networks (CNNs) by retraining them with 12568 steel defect images. It was hypothesized that these networks would perform well on a small dataset of similar plastic defect images. A test accuracy of 87% was obtained with the ResNet18 model while an accuracy of 90.2% was obtained using the ResNet50 model on the steel dataset. However, while applying both these models on the plastic dataset, we only obtained a 60% accuracy, indicating that the initial hypothesis was not accurate.

Data

12,568 annotated, rectangular, grayscale images of steel plates with and without defects came from Severstal kaggle competition[1]. The competition was focused on segmentation, so annotations we modified for binary classification. The plastic test set came from Kolektor and features about 500 grayscale microscope images of plastic parts with and without defects.

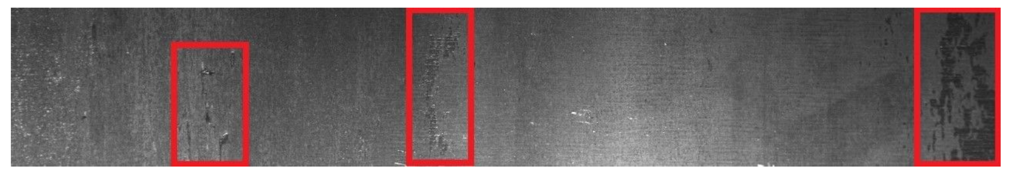

From the figure we see a sample image from the training dataset that has defects. (As isolated by the red box)

Data Preprocessing

Before training on the ResNet model, images were padded from 1600x256 to a square 1600x1600. Images were then rescaled to 512x512 and were flipped to achieve a defect-no_defect ratio closer to 50/50.

Model Architecture

The models used in this project are pretrained ResNet18 and ResNet50 CNNs. The last fully connected layer was changed to be in accordance with binary classification and cross entropy loss was used to validate results. A learning rate of 0.001 was used in all the cases for backpropagation.

ResNet18 Architecture

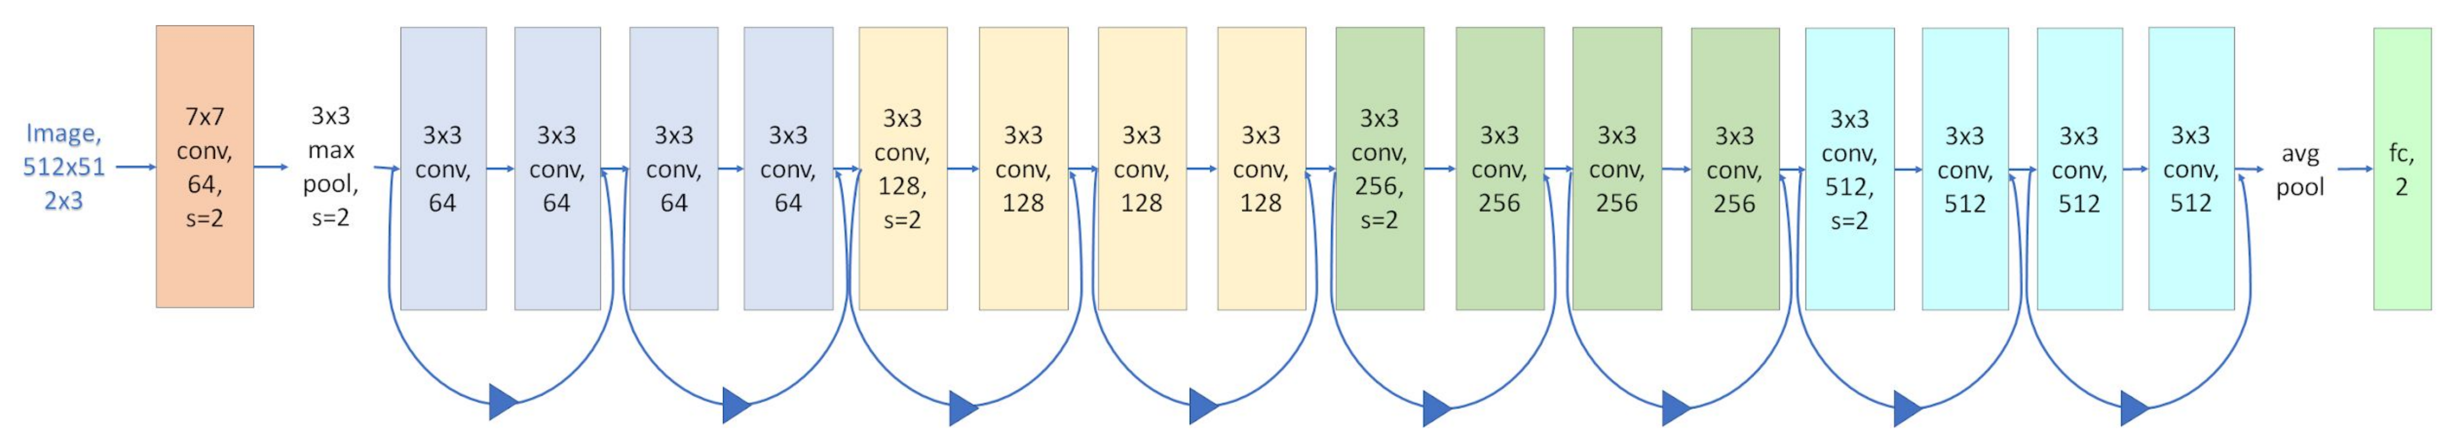

ResNet50 Architecture

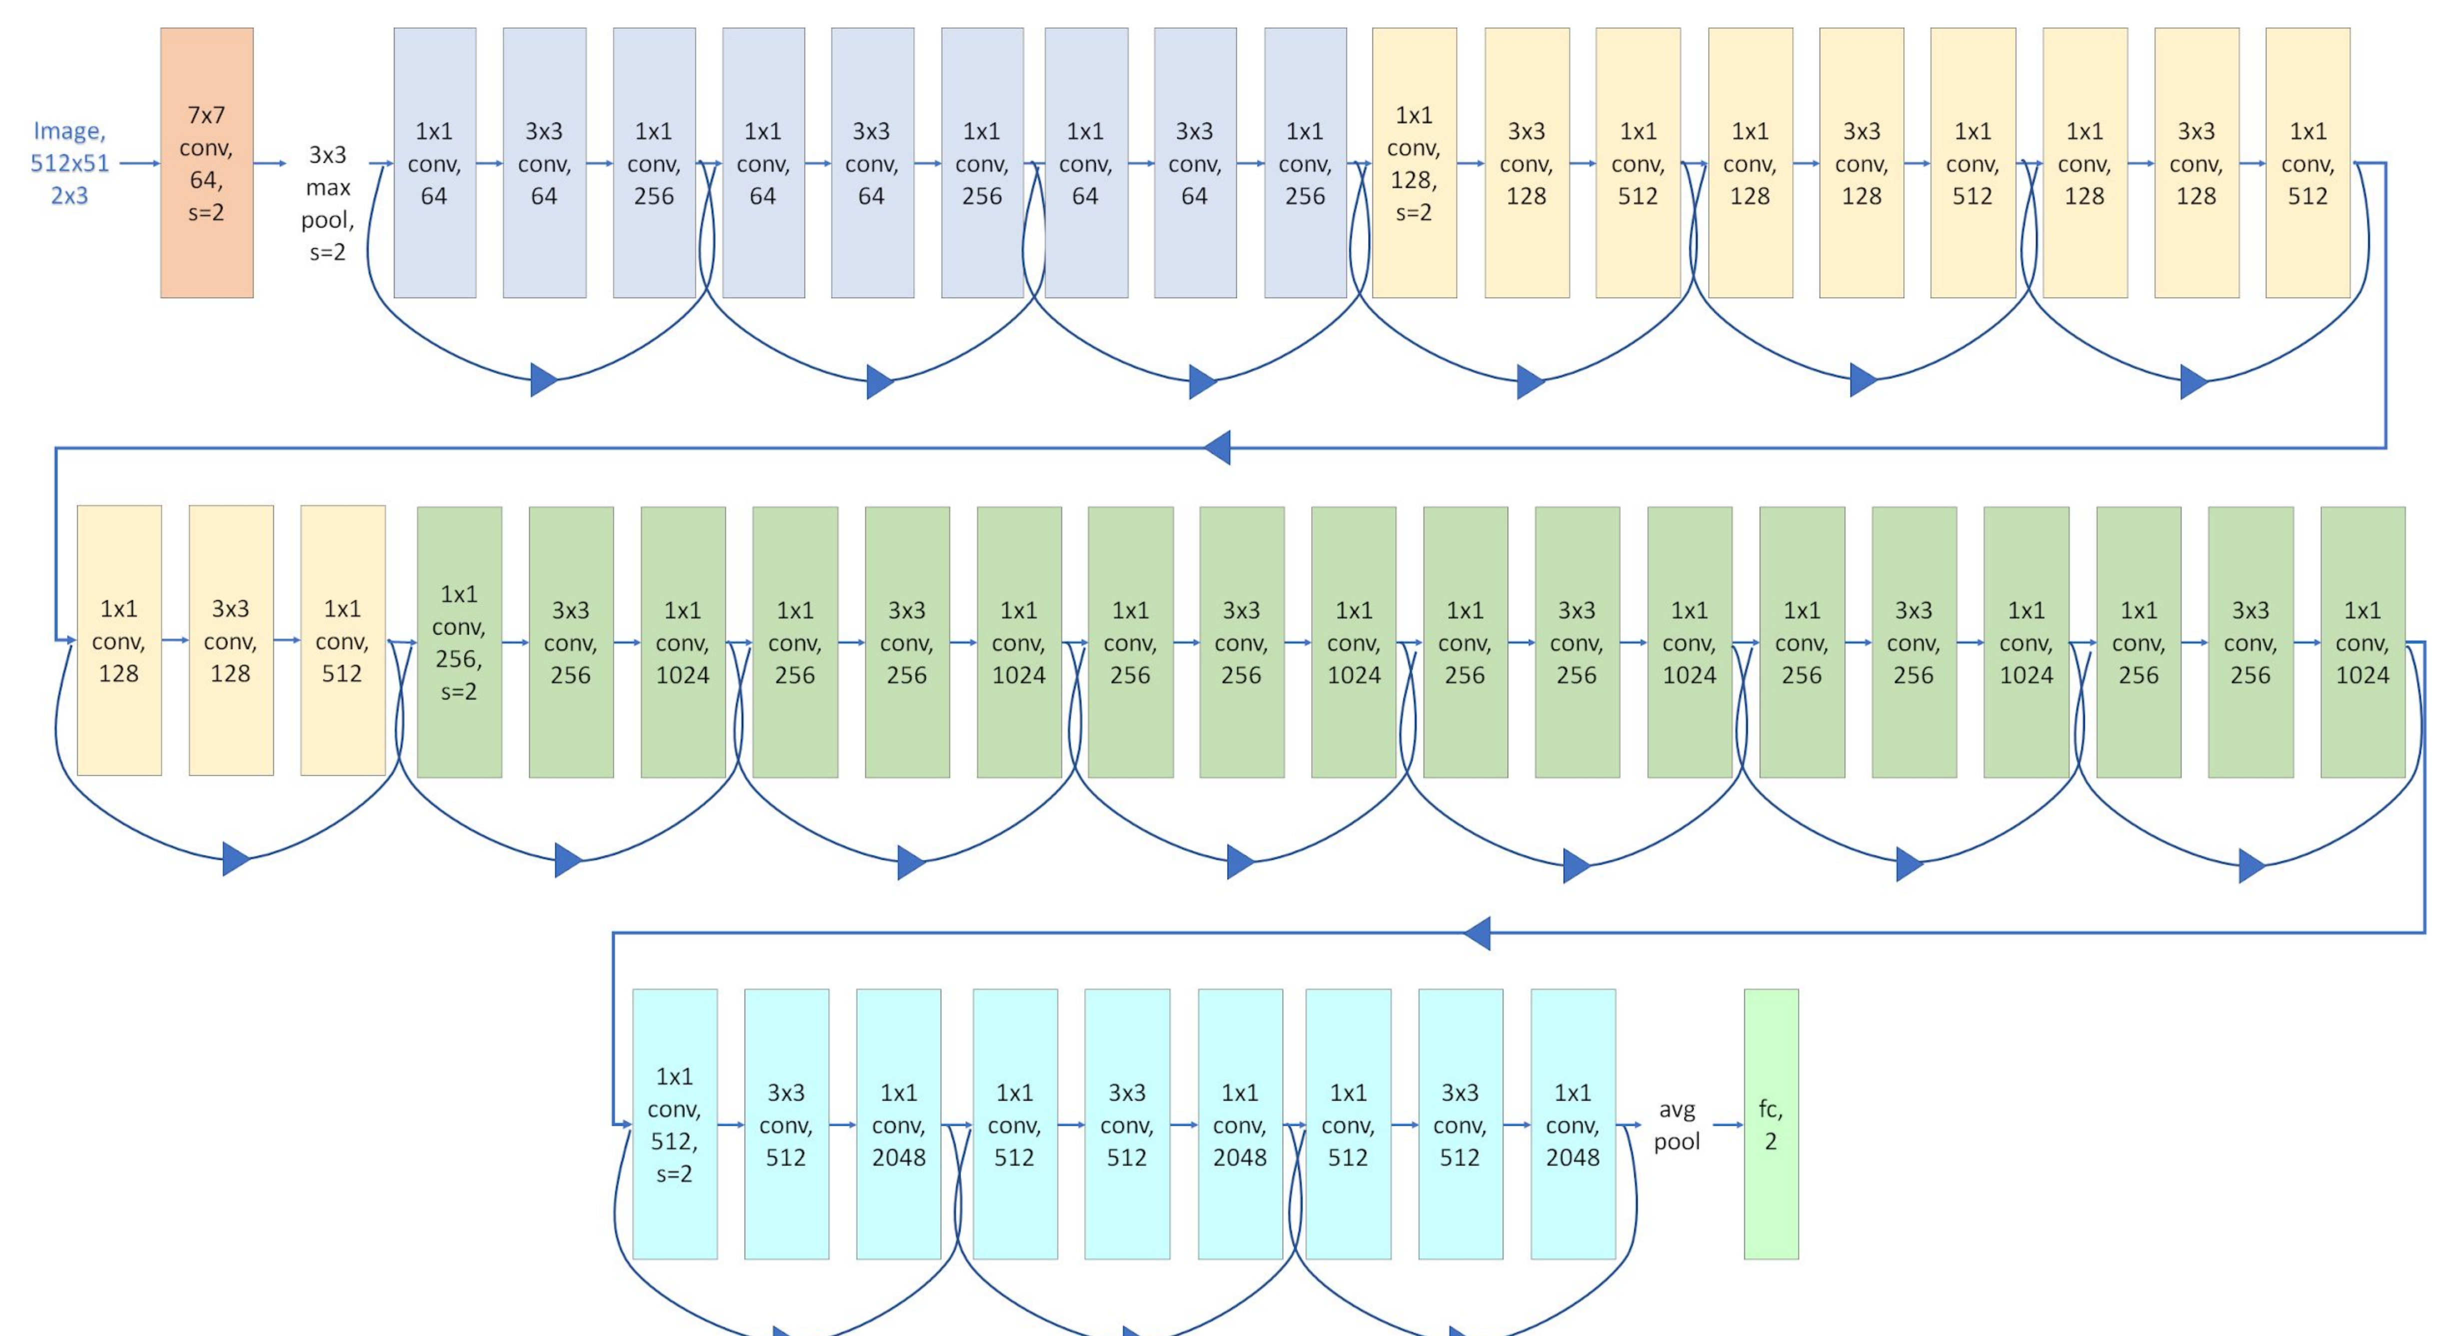

Discussion

Model (on steel)	Test Accuracy
ResNet18 (Last 5 layers unfrozen)	70.2%
ResNet18 (Last 16 layers unfrozen)	87.0%
ResNet101(Last 50 layers unfrozen)	70.5%
ResNet101 (Last 98 layers unfrozen)	90.6%
ResNet50 (Last 48 layers unfrozen)	90.2%

- Bias problem solved by increasing number of learnable parameters but not as much by making network deeper
- Variance problem was not solved even after increasing regularization parameter value
- Applying models on plastic however, yielded only 60.2% accuracy which proved our hypothesis wrong. Some additional work needs to be done to make the model perform better on all bulk materials.

Future

- Perform error analysis to understand sources of inaccuracy - particularly misclassification and compare to human performance.
- Understand reasoning for low performance on plastic dataset and improve accuracy on the data.
- Generalize model to other bulk materials to detect surface defects.

References

[1] "Severstal: Steel Defect Detection." *Kaggle*, https://www.kaggle.com/c/severstal-steel-defect-detection.
[2] "Kolektor Surface-Defect Dataset." *Downloads/KolektorSDD - Visual Cognitive Systems Laboratory*, https://www.vicos.si/Downloads/KolektorSDD.