Comparing State-of-the-Art Convolutional Neural Network Performance in Cancer Prediction Tasks

Cael Elmore, ‘25

IPHS 300 AI for the Humanities (Spring 2023) Prof Elkins and Chun, Kenyon College

Abstract

Convolutional neural networks (CNNs) have gained significant popularity in recent years, showcasing their potential in various fields. This study explores the performance of CNN models, including ResNet-50, EfficientNet-B7, and a bespoke model, in predicting lung and colon cancer from histopathological scans. Through transfer learning, these pretrained models exhibit remarkable predictive capabilities with high accuracy rates in classifying tissue types. The analysis reveals strengths and weaknesses in each model. The bespoke model outperforms others in colon cancer prediction, while the EfficientNet-B7 model excels in distinguishing benign and malignant lung tissues. The ResNet-50 model demonstrates consistent accuracy without confusing different lung cancer types. Considering the findings, a combination of models is recommended for optimal performance in cancer prediction tasks. Furthermore, despite training on unrelated image data, the pretrained models showcase impressive adaptability. Future advancements could involve the creation of a medical image database to further enhance CNN performance in classification and diagnosis tasks. As AI continues to advance in healthcare, these AI tools are expected to improve and become more integrated into the healthcare system.

Introduction

Over the course of the last decade, convolutional neural networks, often referred to as CNNs, have seen a resurgence in popularity. In 2012, a convolutional neural network called AlexNet won the ImageNet 2012 Challenge—a competition for image classification among different softwares—with a top-5 error over 10 percentage points better than its nearest competitor. This victory, due to its dominating fashion, sparked interest in the potential for the field of deep learning, not only in the technology sphere, but in the world at large.
In 2015, Microsoft engineers created ResNet, a convolution neural network using a residual learning framework. This new framework allowed for the creation of networks over 100 layers deep without falling prey to the vanishing gradient problem. While AlexNet, an eight-layer network, had a top-5 error of 15.3%, the deeper ResNet had an error rate of 3.57%.
A few years later, Google researchers took the advancements created by the ResNet engineers even further. In short, they created a principled method that allowed for neural networks to scale their depth, width, and resolution optimally under a given computational constraint, usually measured in terms of floating point operations per second. As an extension of this work, they also used a grid search of different neural architectures to create a family of models known as the EfficientNet family. EfficientNet-B7, the most robust model of the EfficientNet family, achieved state-of-the-art image classification while being over 8x smaller and 6x faster than the previous best model.
Now, these models have been adapted by data scientists and researchers for many different image classification tasks. While models like ResNet-50 v2 and EfficientNet-B7 were trained on an image set containing images like strawberries, desks, keyboards, and other common nouns, the feature extraction done by these models can be abstracted to other tasks. This process is known as transfer learning. This paper examines the performance of these models in predicting lung and colon cancer from different histopathological scans. These pretrained models, with only a few added pooling and fully connected layers, can achieve stunning performance metrics on this task.

Methodology

To conduct this analysis, three different convolutional neural network models were extracted from Kaggle. All three used the Lung and Colon Cancer Histopathological Image Dataset (LC25000) as the training and testing data for their project.The dataset contains images of benign colon tissue, colon adenocarcinoma, benign lung tissue, lung adenocarcinoma, and lung squamous cell carcinoma. Two models took advantage of transfer learning; one used a slightly modified version of the ResNet-50 v2 model, while the other used a modified version of the aforementioned EfficientNet-B7 architecture. The third model is a bespoke 21-layer model, created by Abdallah Wagih Ibrahim, which uses an 80-10-10 train-validation-test split on the 25,000 images in the lung and colon cancer dataset. The performance metrics of each model were compared to see which models performed the best under given circumstances.

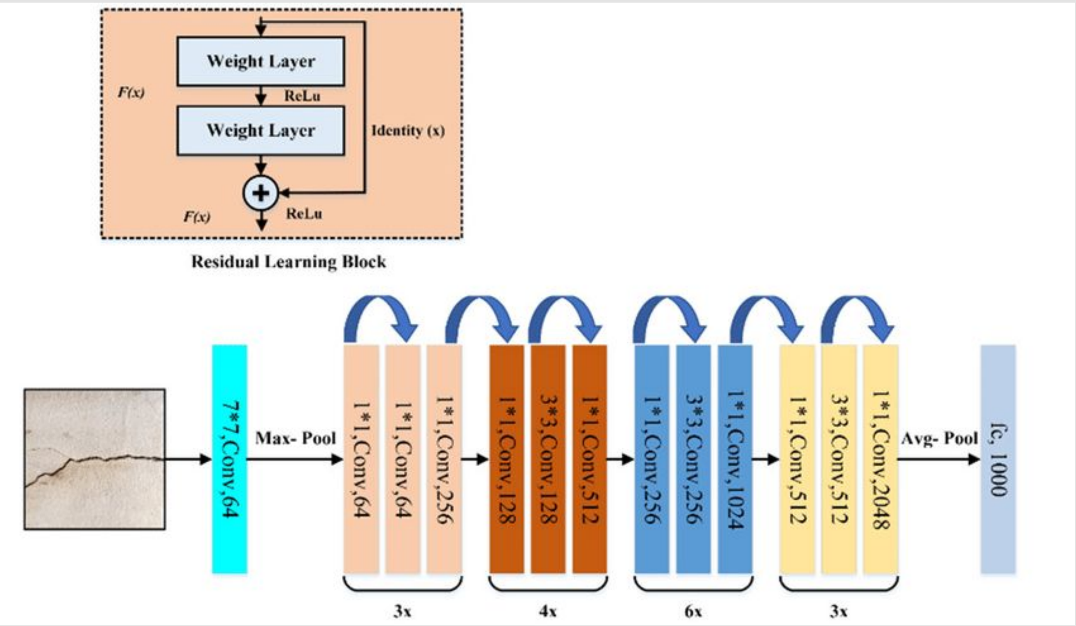

General architecture of ResNet-50, in the context of concrete crack detection. The arrows represent identity shortcut connections that allow the model to skip layers, allowing for models hundreds of layers deep without diminished performance.

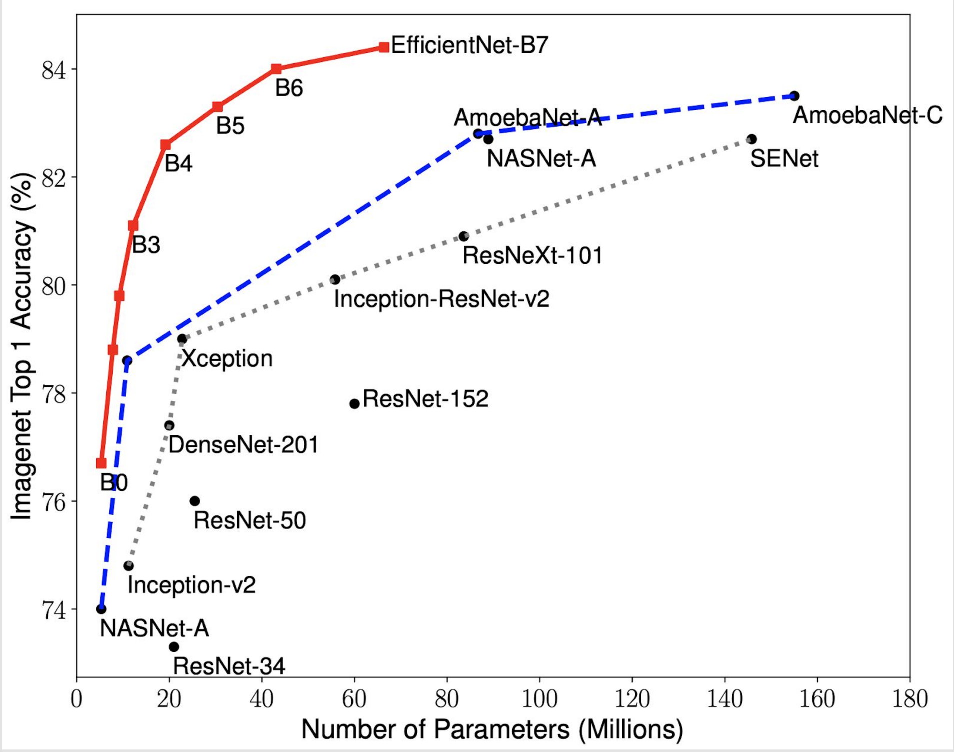

A graph detailing how the EfficientNet family of models compares to other state-of-the-art image classification models. More often than not, EfficientNet architectures achieve higher performance with less parameters than most other models in existence.

Results

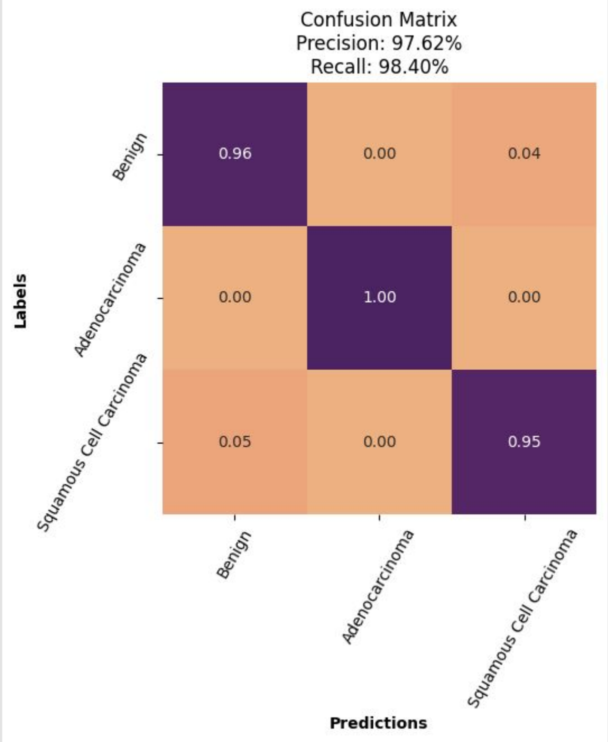

Results of ResNet-50-based model

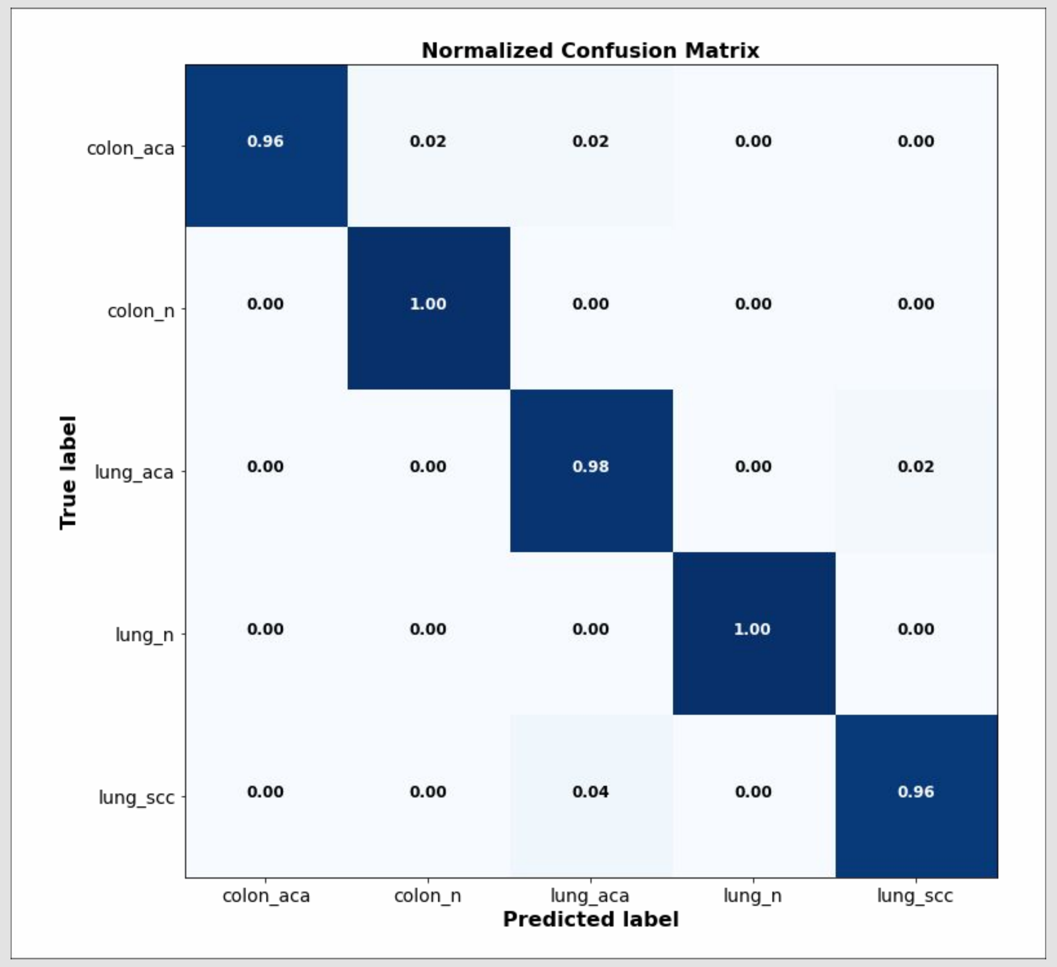

Results of EfficientNet-B7 model

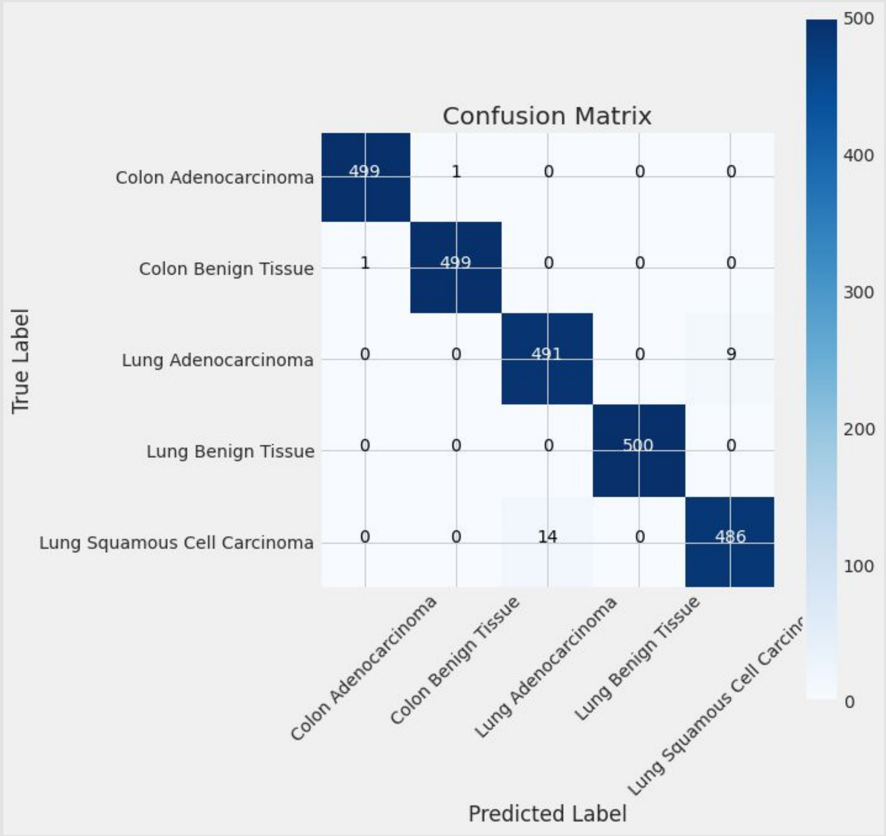

Results of bespoke model

Analysis and Conclusion

All three models did an excellent job classifying the different image scans correctly. Each model labeled each type of tissue correctly at least 95% of the time. As a result, on their own, each model would do a respectable job predicting these types of cancers. However, each model had their own strengths and weaknesses. First off, the ResNet-50 model did not include the colon scans in the training data, and as such did not predict on colon scans. The bespoke model did the best job of the two models that predicted colon cancer, as it classified the colon scans correctly 99.8% of the time, while the EfficientNet model labeled them correctly 98% of the time. When it comes to lung cancer, the picture is not so clear. Both the bespoke model and the EfficientNet model did the best job labeling benign lung tissue correctly, with both having a 100% accuracy rate. After that point, both models occasionally confused lung adenocarcinoma with lung squamous cell carcinoma. The bespoke model labeled 98.2% of adenocarcinoma samples correctly, while labeling 97.2% of squamous carcinoma samples correctly. EfficientNet had accuracy scores of 98% for the former and 96% for the latter. However, the ResNet-50 model never confused adenocarcinoma and squamous cell carcinoma samples. ResNet-50 labeled 100% of adenocarcinoma samples correctly, and correctly differentiated between benign tissue and squamous cell carcinoma 95.5% of the time, labeling 4% of benign tissues squamous carcinoma, and 5% of squamous cell carcinoma benign.
In terms of a best model, there is not a clear winner out of the three. Clearly, the bespoke model and the EfficientNet model did the best job labeling lung tissue benign or malignant, while the bespoke model did the best job labeling colon tissues correctly. The ResNet-50 model occasionally confused benign tissue with squamous carcinoma, but never confused the two types of lung cancer like the other models did. The author suggests an amalgamation of the models for best performance: use the bespoke model for colon tissue scans, the EfficientNet-B7 model to determine whether lung tissue is benign or malignant, and the ResNet-50 model to differentiate between malignant lung scans.
Overall, the bespoke model, despite having many less layers than the pretrained models, performed incredibly well. The ResNet-50 and EfficientNet-B7 models, despite being trained on data that has little in common with medical scans, did a great job transferring its feature learning to a new field. In the future, the creation of an ImageNet-esque image set of medical scans could potentially train large convolutional models to exceed current performance in classification and diagnosis. Many roadblocks to such a database exist, namely the thorny legal area of medical ethics and data access. Recent investments show that AI will only continue expanding in the field of healthcare, so naturally, the AI tools used in hospitals will only continue improving and permeating through the healthcare system at large.

Acknowledgements and References

For a full list of acknowledgements and references, please visit this link: https://tinyurl.com/ElmoreCNNProject