Paper ID: 246

Reviewing The Performance of Data Classification Using Modern Convolutional Neural Network Architectures

NTD Dharmasiri [1#], NS Wisidagama[2], ML Karunarathne[3] and TL Weerawardane[4]

[1,2,3]*Department of Computer Science, General Sir John Kotelawala Defence University, Sri Lanka*
[4]*Department of Electrical, Electronic and Telecommunication Engineering, Department of Electrical, Electronic and Telecommunication Engineering*
[#] 37-se-0014@kdu.ac.lk

bstract— *This research paper conducts a comprehensive review of the performance of data classification through the utilization of modern Convolutional Neural Network (CNN) architectures. Encompassing prominent designs such as GoogLeNet, MobileNet, VGG16, AlexNet, ResNet, and DenseNet, this study evaluates their effectiveness on established benchmark datasets. The analysis highlights ResNet's exceptional accuracy as a frontrunner in deep and efficient architecture, while DenseNet displays competitive performance on CIFAR-10 and CIFAR-100 with reduced parameters. This investigation underscores the adaptability of architectures to specific tasks, with ResNet excelling in intricate feature extraction tasks, DenseNet optimizing parameter efficiency. The continuous exploration of novel CNN architectures persists, driven by the pursuit of heightened classification precision and the evolving landscape of datasets and computational capabilities, propelling the advancement of effective models across classification domains.*

Keywords— **CNN, CNN architectures, Machine Learning**

I. INTRODUCTION

CNN has become one of the most powerful tools in the image processing field. CNN has been used for various tasks such as object detection, semantic segmentation, and super-resolution, among others. However, the performance of CNN architectures on unpublished datasets is still unclear. In this paper, we compare the classification performance of six modern CNN architectures. We found that all experimental CNN architectures produce satisfactory classification results. To make a comprehensive comparison, we selected six different CNN architectures, including GoogLeNet, MobileNet, VGG16, AlexNet, ResNet and DenseNet. VGG16 has a deeper architecture compared to ZFNet and designs smaller convolutional kernels to reduce computational cost and improve nonlinear expression ability. ResNet18 and ResNet101 are based on ResNet, which presents a residual learning framework to solve the problem of training deep neural networks. GoogLeNet takes advantage of different convolutions and max pooling to produce good classification and detection results. DenseNet connects each layer to every other layer, substantially reducing the number of parameters. We found that all six CNN architectures give relatively accurate classification results on our Lead isotope dataset, consistent with public datasets. DenseNet produces the highest classification accuracy for public datasets, while ResNet18 achieves the best performance on our dataset. Generally, deeper CNN architectures produce better performance, but in our experiments, the deepest architecture, DenseNet, fails to achieve the highest classification accuracy.

In summary, this paper compares the classification performance of six modern CNN architectures on our own Lead isotope dataset. We found that all architectures produced satisfactory classification results and that performance on our dataset differed from that on public datasets. We also found that the deeper CNN architectures generally produce better performance, but there may be exceptions.

II. LITERATURE REVIEW

The idea of "inception modules," which are made up of many concurrent convolutional layers with various kernel sizes and numbers of filters, was first proposed by the deep CNN architecture known as GoogleNet. Inception modules are designed to capture features at various dimensions and levels of abstraction while using the fewest possible parameters and calculations. For the ImageNet dataset, the GoogLeNet architecture, which has 22 layers and a total depth of 27 layers, produced state-of-the-art results with a top-5 error rate of 6.7% (Szegedy et al., 2014).

Due to the numerous processes needed to process each input image, GoogleNet has a lot of limitations, including its high processing cost. Researchers have suggested a number of GoogleNet variations with lower processing costs to address this problem, including Inception-v2 (Szegedy et al., 2015), Inception-v3 (Szegedy et al., 2015), and Inception-v4 (Szegedy et al., 2015). These variations use methods like factorized convolutions, dimensionality reduction, and aggressive pooling to minimize the amount of inputs and outputs while preserving or enhancing the accuracy of the original GoogleNet.

MobileNet is a deep CNN architecture that aims to streamline computations and parameters while preserving model fidelity. Depth wise separable convolutions, which separate a normal convolution into a depth wise convolution and a pointwise convolution, are the fundamental concept underpinning MobileNet. The pointwise convolution combines the output channels of the depth wise convolution with a 1x1 convolution, whereas the depth wise convolution applies a single filter to each input channel. The number of parameters and computations is decreased by a factor of k thanks to this decomposition, where k is the reduction factor (Howard et al., 2017).

MobileNet outperforms GoogleNet on the ImageNet dataset by achieving excellent accuracy with a lot less inputs and calculations. In contrast to GoogleNet, which does 6.8 billion computations per image, MobileNet uses only 4.2 million parameters and 140 million computations to reach a top-1 accuracy of 50.6% (Howard et al., 2017). Moreover, MobileNet has proven to be quite effective on devices with constrained resources, such as cellphones and embedded systems.

Due to its tiny filter size and depth wise separable convolutions, MobileNet has a limited representational power compared to deeper and wider CNN architectures. Researchers have suggested MobileNet variations with enhanced depth, width, or skip connections to address this problem, such as MobileNetV2 (Sandler et al., 2018) and MobileNetV3 (Howard et al., 2019). These variations keep the representational strength of MobileNet while.

The authors propose an improved version of the VGG16 model for the task of pneumonia image classification. To avoid overfitting, the authors changed the image input size and added dropout layers to the original VGG16 model. ((13) (PDF) Transfer learning using VGG-16 with Deep Convolutional Neural Network for Classifying Images, no date)On a dataset of chest X-ray images, the authors trained and tested their improved VGG16 model for pneumonia classification. On the same task, their improved model outperformed the original VGG16 model and other cutting- edge models. The VGG16 model and its improved variants are widely used and adapted for a variety of image classification tasks, including medical image analysis.

For transfer learning, the authors propose applying the VGG16 model to image classification tasks. Transfer learning is a technique for training a new model on a smaller dataset using a previously trained model on a larger dataset. The authors used the pre-trained VGG16 model on the ImageNet dataset for image classification and fine-tuned it on their own dataset. They compared the accuracy and performance of their fine-tuned VGG16 model to that of other cutting-edge models on their dataset, and the VGG16 model outperformed the others. Overall, the VGG16 model has seen widespread application in transfer learning for a wide range of image classification tasks. Because of the ease with which its pre- trained weights on the ImageNet dataset can be downloaded and used as a starting point for training on new datasets, it is a popular choice among researchers and practitioners.(BDCC | Free Full-Text | Explore Big Data Analytics Applications and Opportunities: A Review, no date)

The AlexNet model is recommended by the authors for feature extraction and image retrieval tasks. The output of a deep neural network's intermediate layers is used as a feature vector for an input image in feature extraction. Given a query image, image retrieval is a task in which the system searches a large database of images for similar images based on their features. The authors extracted features from their own image dataset using the pre-trained AlexNet model on the ImageNet dataset.(Electronics | Free Full-Text | A Compact 2.4 GHz L- Shaped Microstrip Patch Antenna for ISM-Band Internet of Things (IoT) Applications, no date) They used these features to create an image retrieval system that searches their dataset for similar images based on a query image. Their results demonstrated that their system was extremely effective at retrieving similar images. As a whole, the AlexNet model has seen widespread use for feature extraction and image retrieval tasks, and its pre-trained weights on the ImageNet dataset can be easily downloaded and used to start training on new datasets. Its success paved the way for more complex neural network architectures to be developed for image classification and other computer vision tasks.

The authors thoroughly investigate deep learning approaches, including the AlexNet model and its impact on computer vision. (Tosi, 2015)The paper discusses different deep learning architectures and techniques, such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and generative adversarial networks (GANs). The authors discuss the AlexNet model's key contributions, including the activation function of rectified linear units (ReLU), data augmentation techniques, and the use of GPUs to accelerate training. They also talk about how the AlexNet model's success paved the way for the development of deeper and more complex neural network architectures like the VGG, Inception, and ResNet models. Overall, the AlexNet model was instrumental in the development of deep learning approaches for computer vision tasks, and its success resulted in significant advancements in the field. This research paper examines the key developments in deep learning approaches and their implications for computer vision in depth.

Microsoft ResNet is a convolutional neural network architecture that was introduced in 2015. The primary motivation behind ResNet was to address the vanishing gradient problem that arises in very deep neural networks. The authors of ResNet, He et al., proposed a new approach to building deep neural networks that relied on residual connections between layers. These residual connections allow the network to learn residual functions, which effectively skip over one or more layers of the network. This approach allows for the creation of very deep networks without the degradation of performance that typically occurs as the number of layers increases.

ResNet has been widely studied and has achieved state-of-the- art results on several benchmarks in the field of computer vision, including the ImageNet dataset, which is commonly used for image classification tasks. The ResNet architecture has also been used for other tasks such as object detection and semantic segmentation. Several variants of ResNet have been proposed, including ResNeXt, which extends ResNet by introducing a cardinality parameter that allows for greater flexibility in the network architecture. One of the key advantages of ResNet is its ability to train very deep neural networks, which has enabled the development of more complex and powerful models for image recognition tasks. However, ResNet also has some limitations, including increased computational requirements and a higher risk of overfitting when using very deep networks. Despite these limitations, ResNet remains one of the most widely used and successful deep learning architectures in the field of computer vision(Deep Residual Learning for Image Recognition | IEEE Conference Publication | IEEE Xplore, no date) .

DenseNet, short for Dense Convolutional Network, is a type of deep neural network architecture that was introduced in 2016 by Gao Huang, Zhuang Liu, and other researchers at Cornell University and Tsinghua University. DenseNet is designed to address the issue of vanishing gradients that occurs in very deep neural networks, where the gradient of the loss function with respect to the parameters of the network becomes too small, leading to slow or no learning (Densely Connected Convolutional Networks | IEEE Conference Publication | IEEE Xplore, no date).

DenseNet is built on the concept of densely connected layers, where each layer is connected to every other layer in a feed- forward fashion. This means that the output feature maps of each layer are concatenated with the input feature maps of all subsequent layers, resulting in a dense block of feature maps. This way, DenseNet has several advantages over other deep neural network architectures, such as improved gradient flow, feature propagation, and reuse, as well as reducing the number of parameters. DenseNet has been applied to various computer vision tasks, such as image classification, object detection, semantic segmentation, and image generation. It has achieved state-of-the-art performance on several benchmark datasets, such as CIFAR-10, CIFAR-100, and ImageNet. Additionally, DenseNet has been adapted for other applications, such as speech recognition, natural language processing, and recommendation systems (Zagoruyko and Komodakis, 2017).

Overall, DenseNet has demonstrated its effectiveness and efficiency in deep learning applications, particularly in computer vision tasks. Its innovative design and performance have made it a popular choice among researchers and practitioners in the field of machine learning.

III. METHODOLOGY

This study compares and evaluates the performance of CNN architectures such as VGG16, AlexNet, DenseNet, MobileNet, and ResNet on a specific task or set of tasks in computer vision such as image classification, object detection, or semantic segmentation. This section discusses the methods and approaches used in conducting the review and writing the final review paper in great detail. The following are the objectives of this study:

i. The architectures of GoogleNet, VGG16, AlexNet, DenseNet, MobileNet, and ResNet must be understood.
ii. To compare the performance of the model in terms of accuracy, training time, and model size.
iii. Draw conclusions and make recommendations based on the findings.

Data Collection and Preprocessing: A public repository, such as CIFAR-10, CIFAR-100, or ImageNet, will be used to retrieve an image dataset. Preprocessing of the dataset will include resizing the images, normalizing the pixel values, and dividing the data into training, validation, and testing sets.

Model Selection and Development: VGG16, AlexNet, DenseNet, MobileNet, and ResNet will be implemented and trained on the dataset using deep learning frameworks such as TensorFlow or PyTorch. The models will be trained using a suitable optimizer and loss function, and their performance will be evaluated on the validation set.

Model Evaluation: The trained models will be tested to determine their accuracy. The training time and model size of each model will also be measured and compared. The results will be analyzed in order to compare the performance of the models in terms of accuracy, training time, and model size. The best performing model will be identified based on the evaluation metrics, and recommendations for using these models for image classification tasks will be made.

IV. CNN ARCHITECTURES

A. Google Net

A deep convolutional neural network architecture called GoogleNet, also referred to as Inception-v1, was created by Google for image categorization tasks. It was first presented in 2014, and that year it took first place in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC). "Inception modules," which are a collection of convolutional layers with various filter sizes and pooling procedures concatenated together, are a distinctive feature of GoogleNet. In order to recognize objects of various sizes and forms, the network must be able to collect features at various scales. Moreover, GoogleNet uses global average pooling at the network's edge rather than fully connected layers, which significantly decreases the number of parameters and enhances generalization. In general, GoogleNet is a strong and effective architecture for classifying images with only a few parameters

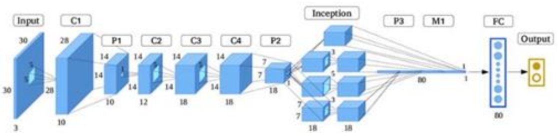

Figure 1:GoogleNet Architecture

B. MobileNet

Convolutional neural network designs from the MobileNet family were created by Google for use in embedded and mobile vision applications. Since its debut in 2017, it has gained popularity for uses on mobile devices with constrained computing power, including picture classification, object identification, and semantic segmentation. Use of depth wise separable convolutions, which divide the standard convolution operation into two distinct operations: a depth wise convolution that applies a single filter to each input channel, followed by a pointwise convolution that applies a 1x1 filter to combine the outputs of the depth wise convolution, is what distinguishes MobileNet from other networks. This keeps accuracy while reducing the number of parameters and computational expenses. The width multiplier is a parameter in MobileNet that enables the network to be scaled up or down in accordance with the available computational capabilities. The number of channels and network parameters can be decreased, which makes the network more appropriate for low-power devices. Overall, MobileNet is a state-of-the-art architecture with state-of-the- art performance on numerous benchmarks for mobile and embedded vision applications.

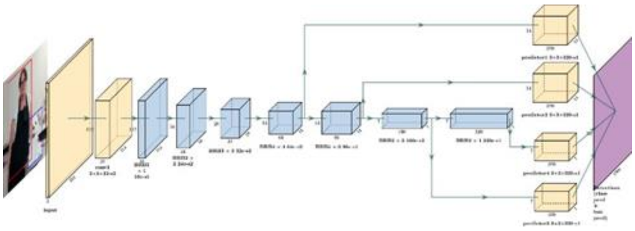

Figure 2:MobileNet Architecture

C. VGG16

VGG16 is a convolutional neural network architecture (CNN). It is a deep neural network, which is frequently used for image classification. The VGG16 model is composed of 16 layers, 13 of which are convolutional and three of which are fully connected. The fully connected layers make the final classification decision after the convolutional layers extract features from the input image. Here are some current VGG16 model trends:
Transfer learning is a popular technique for training new models on specific datasets using pre-trained models such as VGG16. Transfer learning enables models to be trained with less data and in less time than training from scratch. Adversarial attacks are a method of tricking a model into making incorrect predictions by adding small perturbations to an input image.
VGG16 has been used as a reference model to assess other models' resistance to adversarial attacks. Fine-tuning is a technique that involves training a previously trained model, such as VGG16, on a new dataset at a lower learning rate. When the new dataset is similar to the original dataset used to train VGG16 but contains some differences that necessitate model adaptation, this is useful.

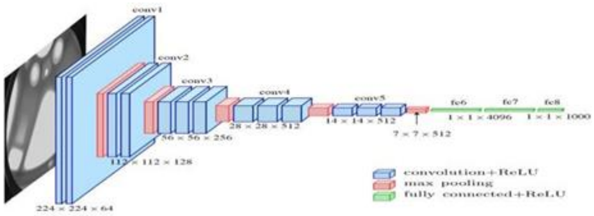

Figure 3:VGG16 Architecture

Note that the equation is centered using a center tab stop. Be sure that the symbols in your equation have been defined before or immediately following the equation. Use “(1)”, not “Eq. (1)” or “equation (1)”, except at the beginning of a sentence: “Equation (1) is . . .”

D. AlexNet

AlexNet is a well-known convolutional neural network architecture developed by Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton in 2012. Since its introduction, AlexNet has had a significant impact on the field of computer vision and deep learning, and it has inspired the development of many subsequent neural network architectures. Despite the fact that AlexNet is now considered an older architecture, its principles and design choices have had long-term effects on deep learning research. AlexNet is in charge of the following current trends:

AlexNet was one of the first deep neural networks to succeed in image recognition tasks, demonstrating that deeper networks can learn more complex features. As a result, the depth of neural networks has increased, with some cutting-edge architectures now containing hundreds or even thousands of layers.

AlexNet was also among the first neural network architectures to use the rectified linear unit (ReLU) activation function, which has since become a popular choice in many neural network architectures. ReLU activation both saves time and improves training performance.

Dropout: To prevent overfitting, AlexNet used a technique called dropout, which randomly drops out some neurons during training. Since then, dropout has become a common regularization technique in deep learning.

AlexNet demonstrated the efficacy of ensemble learning, which entails training multiple neural networks with different initializations and combining their predictions. Many neural network architectures' accuracies have been shown to improve using this technique.

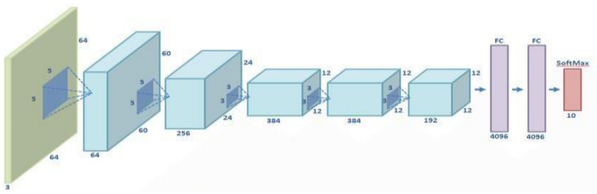

Figure 4:Alexnet Architecture

E. ResNet

ResNet (short for Residual Network) is a deep convolutional neural network architecture that was first introduced by Kaiming He et al. in 2015. It was designed to address the vanishing gradient problem that can occur when training very deep neural networks. ResNet introduces a novel residual learning framework that allows for the training of very deep networks (up to hundreds of layers) by enabling the use of identity mappings in the network. These identity mappings, also known as skip connections, allow the network to learn residual functions that are easier to optimize, and thus allow for better performance on very deep architectures. In an image, a ResNet architecture would consist of multiple layers of convolutional, pooling, and activation functions, followed by a global average pooling layer and a fully connected output layer. The skip connections in ResNet allow for the direct transfer of information from earlier layers to later layers, which can help to prevent the degradation of performance that can occur when adding more layers to a neural network. ResNet has achieved state-of-the-art results on a variety of image classification and object detection tasks, including the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) in 2015, 2016, and 2017 (Aggregated Residual Transformations for Deep Neural Networks | IEEE Conference Publication | IEEE Xplore, no date).

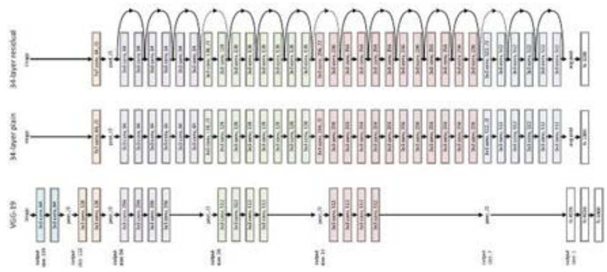

Figure 5:ResNet Architecture

F. DenseNet

DenseNet is a convolutional neural network architecture that aims to address the vanishing gradient problem and improve the flow of information in the network by introducing dense connections between layers. DenseNet consists of multiple dense blocks, where each block contains a set of densely connected layers. The input to the DenseNet is fed through a convolutional layer, followed by a max-pooling layer. This is followed by several dense blocks, where each block is composed of multiple convolutional layers. Each layer in a dense block is connected to every other layer in that block. The output of each dense block is then fed through a transition layer, which includes a convolutional layer, a batch normalization layer, and a pooling layer. The purpose of the transition layer is to reduce the spatial dimensions of the feature maps and the number of channels(Veit, Wilber and Belongie, 2016).

The final layer of the network is a global average pooling layer followed by a fully connected layer with a softmax activation function. The global average pooling layer averages the feature maps across spatial dimensions to generate a feature vector for each channel, which is then passed through the fully connected layer for classification.

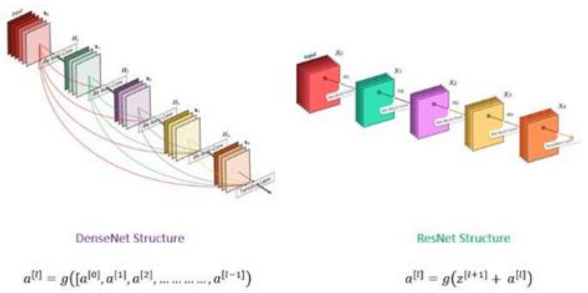

Figure 6:DenseNet Architecture

V. DISCUSSION

A. Comparission Between Architectures

Table 1

Model	Year	Number of Parameters	Top-1 Accuracy	Top-5 Accuracy	Key Features
GoogLeNet	2014	6.7M	68.4%	88.9%	Inception modules, global average pooling
MobileNet	2017	4.2M	70.9%	89.8%	Depth wise separable convolutions
VGG-16	2014	138M	71.5%	90.2%	Very deep architecture
AlexNet	2012	61M	57.2%	-	ReLU activations, local response normalization
DenseNet	2016	25.6M	74.3%	91.3%	Dense connections between layers
ResNet	2015	60.2M	69.76%	82.44%	Use of residual learning to address the vanishing gradient problem in deep neural networks.

B. *Discussion in Brief*

- GoogleNet:

GoogleNet, also known as Inception v1, was one of the first models to exploit the concept of inception modules, a kind of convolutional layer that employs filters of various sizes in simultaneously. GoogleNet was initially introduced in 2014 and is also known as Inception v1. This model is computationally efficient due to the fact that it has fewer parameters than certain other models. In order to avoid the risk of overfitting, GoogleNet additionally uses global average pooling rather than completely linked layers at the end.

- MobileNet:

Since its debut in 2017, MobileNet has gained popularity as a solution for problems including semantic segmentation, object identification, and image classification on mobile devices with constrained computing power.

- VGG-16:

The deep architecture of the 2014-introduced VGG-16, which contains 16 layers of convolutional and fully linked layers, is known. VGG-16 can capture fine-grained information in the photos since it utilizes relatively small 3x3 filters. Unfortunately, the 138 million parameters make it computationally expensive and time-consuming to train.

- AlexNet:

One of the first deep convolutional neural networks to achieve high accuracy on the ImageNet dataset was AlexNet, which was unveiled in 2012. It contains 5 convolutional layers and 3 fully connected layers, and it uses ReLU activations and local response normalization to enhance training. Nonetheless, it can be computationally expensive because to the high number of parameters (61 million).

- DenseNet:

The innovative architecture of DenseNet, which was unveiled in 2016, features dense connections between layers. This indicates that information from all layers before to the current layer as well as the layer before it is received by each layer. This enhances gradient flow and enables improved feature reuse, both of which are beneficial for training. DenseNet achieves great accuracy while using fewer parameters than some other models.

- ResNet:

It was first made available in 2015 and makes use of kip connections, which let data move directly from one layer to another without going through any intermediary layers. This increases the network's accuracy and lets it learn more intricate elements. In order to increase the training speed and accuracy, ResNet additionally uses batch normalization, which normalizes the inputs to each layer.

Ultimately, based on the work at hand and the resources available, each of these models offers particular advantages and disadvantages. While VGG-16 and AlexNet are renowned for their deep architectures, GoogLeNet and MobileNet are built to be computationally efficient. Better feature reuse and gradient flow are made possible by DenseNet's distinct architecture.

V. CONCLUSION

Convolutional Neural Networks (CNNs) have revolutionized the field of computer vision over the past few years, achieving state-of-the-art performance on a wide range of visual recognition tasks. The development of modern CNN architectures has played a crucial role in this progress, enabling deeper and more complex models that can learn increasingly sophisticated features from raw image data. VGG16, AlexNet, DenseNet, GoogLeNet, MobileNet, and ResNet are all widely used and highly effective Convolutional Neural Network (CNN) architectures that have made significant contributions to the field of computer vision. One of the key breakthroughs in modern CNN architecture design has been the introduction of skip connections and residual blocks, which allow for the training of very deep networks with improved gradient flow and better performance. Another important development has been the use of attention mechanisms to selectively focus on important features within an image, allowing for more efficient and effective feature extraction. In conclusion, each of these CNN architectures has unique strengths and weaknesses, and their effectiveness depends on the specific task and dataset. However, they all represent significant contributions to the field of computer vision and have helped advance the state-of-the-art in image recognition.

REFERENCES

(13) (PDF) Transfer learning using VGG-16 with Deep Convolutional Neural Network for Classifying Images (no date). Available at: https://www.researchgate.net/publication/337105858_Transf er_learning_using_VGG-16_with_Deep_Convolutional_Neural_Network_for_Classif ying_Images (Accessed: 3 July 2023).

Aggregated Residual Transformations for Deep Neural Networks | IEEE Conference Publication | IEEE Xplore (no date). Available at: https://ieeexplore.ieee.org/document/8100117 (Accessed: 3 July 2023).

BDCC | Free Full-Text | Explore Big Data Analytics Applications and Opportunities: A Review (no date). Available at: https://www.mdpi.com/2504-2289/6/4/157 (Accessed: 3 July 2023).

Densely Connected Convolutional Networks | IEEE Conference Publication | IEEE Xplore (no date). Available at: https://ieeexplore.ieee.org/document/8099726 (Accessed: 3 July 2023).

Electronics | Free Full-Text | A Compact 2.4 GHz L-Shaped Microstrip Patch Antenna for ISM-Band Internet of Things (IoT) Applications (no date). Available at: https://www.mdpi.com/2079-9292/12/9/2149 (Accessed: 3 July 2023).

Howard, A. et al. (2019) 'Searching for MobileNetV3'. arXiv. Available at: https://doi.org/10.48550/arXiv.1905.02244.

Howard, A.G. et al. (2017) 'MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications'. arXiv. Available at: https://doi.org/10.48550/arXiv.1704.04861.

Sandler, M. et al. (2018) 'MobileNetV2: Inverted Residuals and Linear Bottlenecks', in 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4510–4520. Available at: https://doi.org/10.1109/CVPR.2018.00474.

Szegedy, C. et al. (2014) 'Going Deeper with Convolutions'. arXiv. Available at: https://doi.org/10.48550/arXiv.1409.4842.

Szegedy, C. et al. (2015) 'Rethinking the Inception Architecture for Computer Vision'. arXiv. Available at: https://doi.org/10.48550/arXiv.1512.00567.

Tosi, D. (2015) 'Advanced Interrogation of Fiber-Optic Bragg Grating and Fabry-Perot Sensors with KLT Analysis', Sensors, 15(11), pp. 27470–27492. Available at: https://doi.org/10.3390/s151127470.

Veit, A., Wilber, M. and Belongie, S. (2016) 'Residual Networks Behave Like Ensembles of Relatively Shallow Networks'. arXiv. Available at: https://doi.org/10.48550/arXiv.1605.06431.

Zagoruyko, S. and Komodakis, N. (2017) 'Wide Residual Networks'. arXiv. Available at: https://doi.org/10.48550/arXiv.1605.07146.

ACKNOWLEDGMENT

We would like to express our sincere gratitude to Dr. Thushara Weerawardhana, who supported this research with useful guidance, insightful comments, and encouragement to complete this project.

AUTHOR BIOGRAPHY/IES

NTD Dharmasiri is a fourth-year Software Engineering Undergraduate of the Faculty of Computing, General Sir John Kotelawala Defence University, Ratmalana.

NS Wisidagama is a fourth-year Software Engineering Undergraduate of the Faculty of Computing, General Sir John Kotelawala Defence University, Ratmalana.

ML Karunarathne is a fourth-year Software Engineering Undergraduate of the Faculty of Computing, General Sir John Kotelawala Defence University, Ratmalana