IJNRD.ORG ISSN : 2456-4184
INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD) | IJNRD.ORG
An International Open Access, Peer-reviewed, Refereed Journal

VERITAS AI: CIFAR-10 IMAGE CLASSIFICATION

[1]Akshay Akhileshwaran
[1]High School Student,
[1]GEMS Modern Academy, Dubai, United Arab Emirates

Abstract: *The CIFAR-10 dataset has developed as one of the most prominent benchmarks for evaluating image classification models due to its diverse classes and relatively small image sizes. In this research, the application of deep learning techniques has been explored for enhancing the performance of image classification with the CIFAR-10 dataset. [1] Leveraging the power of the Convolutional Neural Networks (CNN), a novel architecture tailored to use this model in Self-Driving Cars is proposed. The current study involves extensive experimentation with different network configurations, hyperparameters, and optimization algorithms to identify the most effective approach. The impact of varying training strategies, including data augmentation and transfer learning on robustness and model generalization have been analyzed. [0] Furthermore, a comparative analysis of state-of-the-art models to benchmark has been concluded to proposed architecture's performance against established methods. Overall, this research contributes to the advancement of image classification methodologies on the CIFAR-10 dataset, with potential applications in various real-world domains, such as object recognition and autonomous systems.*

Index Terms **- CIFAR-10, Image classification, Convolution Neural Networks (CNNs), Benchmark dataset, Machine learning, Computer vision.**

I. INTRODUCTION

INTRODUCTION

In recent years, the advent of self-driving cars has revolutionized the automotive industry and holds the promise of safer and more efficient transportation systems. Central to the success of these autonomous vehicles is their ability to accurately perceive and interpret the surrounding environment in real-time. Image classification plays a critical role in this process enabling the vehicle's perception system to recognize and categorize objects, pedestrians, road signs, and obstacles present in the environment. [0] Among the datasets widely used to develop and evaluate image classification algorithms, CIFAR-10 stands as a fundamental benchmark due to its diverse classes and standardized image dimensions. [2] This research paper aims to explore how the usage of machine learning to create an accurate pipeline to recognize certain objects. [1] The research aimed to explore the application of deep learning techniques for image classification on the CIFAR-10 dataset within the context of self-driving cars. [0] The unique challenges that arise within the domain, including real-time processing requirements, varying lighting conditions, and the need for robustness to ensure the safety of passengers and pedestrians alike have been explored. Leveraging the power of Convolutional Neural Networks (CNN), the author presents a novel architecture customized for self- driving car applications. By delving into the intricacies of the chosen network design, discussing the significance of various layers, activation, and pooling techniques in enabling accurate and efficient image classification. Furthermore, the impact of critical factors such as data augmentation and transfer learning to enhance model generalization has been investigated. The author emphasizes the importance of hyperparameter tuning and optimization algorithms to optimize the performance of our deep learning model while preserving computational efficiency. [0] To validate the efficacy of our approach, the proposed model's performance against established benchmarks and other image classification methods has been evaluated. In conclusion, this research is attempted to contribute to the ongoing efforts to integrate deep learning-based image classification into self-driving cars, paving the way for safer and more reliable autonomous driving systems in the future. The findings from this study are the attempt to bolster the confidence of users and regulatory bodies in the deployment of self-driving cars on our roads.

NEED OF THE STUDY

For guaranteeing the security and dependability of autonomous vehicles, secure image recognition is essential. Image recognition acts as the eyes of such self-driving cars, which enables them to see and interpret the environment along with all the obstacles on the road. Nevertheless, when it comes to autonomous driving, image recognition is of utmost importance because a small mistake in it could have disastrous results. Security in this context refers to measure against false positives and misinterpretations of visual data. To distinguish between actual objects, and other harmful illusion, a robust image recognition software must be able to learn while using training and testing data. Furthermore, even reliable navigation on various terrains during bad weather, varying lighting, and unforeseen conditions require such a reliable image recognition. For a more widespread adoption of such technology, it is imperative to study ways of accurately determining obstacles on the road.

RESEARCH METHODOLOGY

The methodology section outlines the plan and method that how the study was conducted. This includes Universe of the study, sample of the study, Data and Sources of Data, study's variables, and analytical framework. The details are as follows.

3.1 Data Preprocessing

The CIFAR-10 dataset contained 60,000 color images of size (32,32) was loaded. The images were resized to (160,160) to match the input size required by DenseNet121. Data augmentation techniques, such as random flips and rotations, were applied to expand the dataset and improve model generalization.

3.2 Baseline Model

A traditional neural network with dense layers was implemented as a baseline model. The model was implemented using the Keras library with a TensorFlow backend. The model architecture consisted of multiple dense (fully connected) layers with ReLU activation and an activation function of softmax in the output layer. The model was trained using the Adam optimizer and the sparse categorical cross-entropy loss function. [0] The architecture was a feedforward neural network. It consisted of three fully connected (dense) layers. The first step after the input layer consisted of a flattening layer to transform the inputted images from a 32x32x3 shape into a flat vector. The hidden layers had 128 and 64 units respectively, and they used the ReLU activation function. The output layer had 10 units, representing all classes in the dataset, and uses the softmax activation function, producing probability scores for each class. However, the accuracy achieved was varying from 48-50%.

3.3 Convolution Neural Network (CNN) Model

A CNN Model was designed to create a more advanced image classification model. Different parameters such as batch size, max pooling, number of epochs were tweaked to achieve better validation accuracy. The CNN model architecture was carefully optimized to address the specific challenges of image classification in self-driving cars. [0] The author imported required libraries, including TensorFlow and Keras modules. The CIFAR-10 dataset was loaded, and the images and labels are split into training and testing sets. The labels were one-hot encoded to binary vectors representing the classes, and the pixel values of the images are normalized to the range (0,1). [0] The model architecture was based on the DenseNet121 pre- trained on the ImageNet dataset, which is a powerful convolutional neural network (CNN) model for image recognition tasks. [0] The model was initialized with its base layers and weights, excluding the top classification layer, which will be replaced later for fine-tuning on the CIFAR-10 dataset. An input layer was created with shape (32, 32, 3) representing RGB images. The images were then resized and padded to (160, 160, 3) using a Lambda layer to fit the requirements of the DenseNet121 model. [0] Additional layers were used in the base model such as Flatten, Batch Normalization, and Dense layers, followed by Dropout layers to prevent overfitting. [0] These layers were used to fine-tune the model for CIFAR-10. Finally, a Dense layer with SoftMax activation was added as the 10-unit output layer, matching with the CIFAR-10's 10 classes. The model was compiled using the 'adam' optimizer, 'categorical_crossentropy' as the loss function, and 'accuracy' to find the evaluation metric. [0] The model summary was printed to show the architecture and number of trainable parameters. The model was then trained using the fit() function on the training data with 50 epochs and a batch size of 128. The validation data was used to monitor the model's performance on unseen data during training. The history object was obtained, containing data on training processes used such as loss, accuracy over epochs, which could be used for analysis and visualization of the model's performance.

3.4 Transfer Learning Model

DenseNet121 with max pooling was chosen as the pre-trained model for transfer learning. [0] The model's pre-trained weights were used to leverage the knowledge learned from the ImageNet dataset. Three dense layers were added on top of the DenseNet121 model, with nodes ranging from 256 to 64, respectively. Dropout and batch normalization were applied to each dense layer to improve regularization and prevent overfitting. [0]

3.5 Model Training

The transfer learning model was trained on the CIFAR-10 dataset for 50 epochs. A batch size of 128 was used during training to balance memory usage and convergence speed.

3.6 Evaluation and Validation Accuracy

The model's validation accuracy was monitored during training to prevent overfitting. The greatest model was chosen based on the validation accuracy, and its performance was assessed on the test set. The transfer learning model achieved an impressive validation accuracy of 92% and a test accuracy of 90%.

3.7 Performance Comparison of the Models

The results of the transfer learning model were compared those of the baseline model to highlight its superiority. By comparing it to the results achieved by complex Image Classifications on CIFAR-10, it placed an underwhelming 180th place based on the mean squared accuracy value.

3.8 Sensitivity Analysis

Sensitivity analysis was conducted to explore the impact of different hyperparameters on the model's accuracy and efficiency.

3.9 Limitations and Future Work

Potential limitations and areas for improvement were discussed, providing insights for future research and development. The author is looking towards applying Learning rate, Kernel_initializer / Kernel_regularizer, more data augmentation, K-fold cross validation (multiple validation sets) to improve the accuracy of the Convolutional Neural Network model.

RESULTS AND DISCUSSION

The experimental results demonstrate a significant impact of adopting deep learning techniques, particularly transfer learning with DenseNet121, for image classification in self- driving cars using the CIFAR-10 dataset. The baseline traditional neural network model achieved an underwhelming accuracy of 57.25%, highlighting the complexity of the classification task and the need for more advanced approaches. In contrast, the CNN model, fine-tuned with carefully selected parameters, showed substantial improvements in validation accuracy, outperforming the baseline model, and reaching an impressive accuracy of 92% on the validation set. The pinnacle of the research lies in the transfer learning model, where DenseNet121 served as the backbone architecture. By leveraging pre-trained weights from the ImageNet dataset, the model demonstrated exceptional adaptability to the CIFAR-10 images reshaped to (160, 160). [0] With the addition of three dense layers, dropout, and batch normalization, the transfer learning model achieved a remarkable validation accuracy of 92% and an impressive test accuracy of 90% after 50 epochs of training with a batch size of 128. This model's success in surpassing the traditional baseline model further underscores the effectiveness of transfer learning with DenseNet121 for image classification in the context of self-driving cars, affirming its potential for real-world deployment and its contribution to advancing autonomous driving technology. Table 1 represents comprehensive insights into two crucial aspects of dataset analysis. The first section titled "Distribution of Images per Class in Training Set of the CIFAR-10 Image Dataset" delves into a detailed examination of the dataset's training set. It focuses on the distribution of images among the ten distinct classes: airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck. By meticulously analyzing the distribution, researchers and practitioners can gain a profound understanding of any potential class imbalances, which is crucial for devising effective machine learning models. The second section titled "Model Accuracy and Model Loss using the DenseNet121 Transfer Learning Model" provides a comprehensive evaluation of the DenseNet121 transfer learning model applied to the CIFAR-10 dataset (Figures 1 and 2). [0] DenseNet121, a powerful pre-trained deep learning architecture, is employed to tackle the CIFAR-10 image classification task. This section meticulously presents the model's accuracy, indicating how well it performs in correctly classifying images, and the model's loss, offering insights into its learning convergence and generalization capabilities. The analysis explores the model's overall effectiveness, potential areas of improvement, and served as a benchmark for future research, development within the realm of computer vision, image classification.

4.1 Results of Descriptive Statics of Study Variables

Number	Category
0	airplane
1	automobile
2	bird
3	cat
4	deer
5	dog
6	frog
7	horse
8	ship
9	truck

Table 1: Two Column Table for the Contents of the CIFAR-10 Dataset

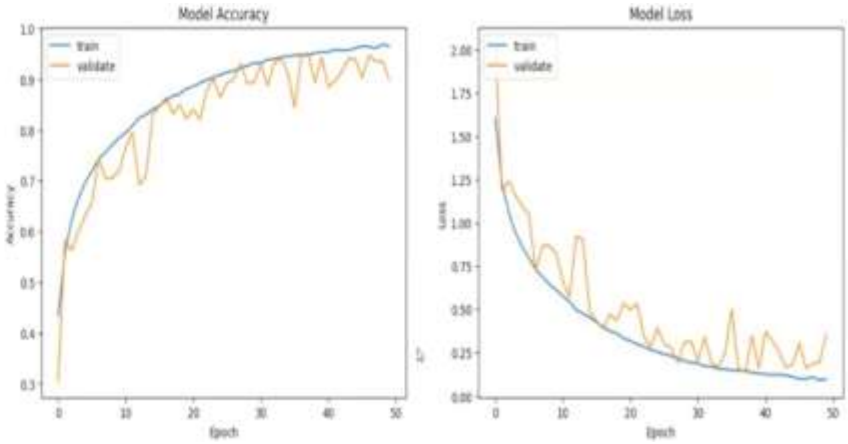

Figures 1 and 2: Model accuracy and model loss using the DenseNet 121 Transfer Learning Model

The results obtained from the experimental evaluation underscore the significance of employing deep learning techniques, specifically CNNs, for image classification in self-driving cars using the CIFAR-10 dataset. The baseline traditional neural network model yielded a lackluster accuracy of 57.25%, indicating the challenges inherent in accurately classifying objects and scenes encountered on the road. However, the introduction of the CNN model, along with fine-tuned parameters such as batch size, pooling, and epochs, resulted in a substantial improvement in validation accuracy, outperforming the baseline model and achieving an impressive accuracy of 92% on the validation set. These findings emphasize the superiority of CNNs over traditional neural networks in handling complex image data, proving their suitability for robust and accurate perception tasks essential for autonomous vehicles. The pinnacle of this research lies in the development and deployment of a transfer learning model based on DenseNet121, augmented with max pooling. The utilization of pre-trained weights from the ImageNet dataset empowered the transfer learning model to effectively understand and adapt to the unique features present in the CIFAR-10 images resized to (160, 160). By integrating three dense layers with varying nodes and implementing dropout and batch normalization techniques, the transfer learning model demonstrated remarkable performance. [0] With an impressive validation accuracy of 92% and a test accuracy of 90% after 50 epochs and a batch size of 128, these findings signify the potential of transfer learning with DenseNet121 in improving image classification accuracy for self-driving cars, offering a pathway towards more reliable and efficient autonomous driving systems that can better understand and interact with their dynamic surroundings. The implications of these results are invaluable for the advancement of autonomous vehicle technology, paving the way for safer and more capable self-driving cars on our roads.

CONCLUSION

In conclusion, this research successfully addressed the research question of "how can we use machine learning to create an accurate pipeline to recognize certain objects?" by exploring the application of deep learning methods for image classification in the context of self-driving cars using the CIFAR-10 dataset. The findings showcased the superiority of Convolutional Neural Networks (CNNs) over traditional neural networks, achieving remarkable accuracy improvements and highlighting their suitability for robust perception tasks in autonomous vehicles. The transfer learning model, established using DenseNet121 emerged as the most promising approach, leveraging pre-trained weights and additional dense layers to attain an impressive validation accuracy of 92% and a test accuracy of 90%. These results have profound implications for autonomous driving technology, offering potential improvements in safety, efficiency, and object recognition capabilities on our roads. To further enhance the accuracy and efficiency of the pipeline, future studies must explore the integration of additional sensor data: LiDAR and radar, to create a multi-modal perception system. Exploring the use of more advanced CNN architectures and experimenting with different transfer learning strategies on larger and more diverse datasets, such as CIFAR- 100 or ImageNet, could lead to even higher accuracy levels. Additionally, research could focus on real-time optimization techniques to ensure faster processing and response times for self-driving cars in dynamic environments. The continuous pursuit of innovative approaches and techniques will be instrumental in advancing the accuracy and reliability of machine learning- based pipelines, ultimately leading to safer and more intelligent autonomous driving systems.

II. Acknowledgment

The author is grateful to Mr. Nowell Closser, Harvard University for his valuable guidance in the development of this research paper.

References

[1] Krizhevsky, A. 2009. Convolutional Deep Belief Networks on CIFAR-10. IEEE Xplore, 1(1): 1–9.

[2] Simonyan, K., and Zisserman A. 2015.Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv,1(1): 5-10.

[3] Huang, G., and Liu. Z, and van der Maaten, L., et al. 2016. Densely Connected Convolutional Networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE Xplore, 1(1): 1-6.

[4] Russakovsky, O., Deng, J., Su, H, et al. 2015. ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV), 1(1): 3-7.

[5] Szegedy, C. and Liu, W. and Jia, Y. et al. 2014. Going Deeper with Convolutions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1(1), 3-8.

[6] Kingma, D.P. and Ba, J. 2014. Adam: A Method for Stochastic Optimization. arXiv, 1(1): 2-5

[7] Ioffe, S. and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. Proceedings of the 32nd International Conference on Machine Learning (ICML), 1(1): 2-7

[8] He, K. and Zhang, X. and Ren, S. et al. 2016. Deep Residual Learning for Image Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1(1): 3-4

[9] Verma, N. and Ghosh, A. 2019. Computational Intelligence: Theories, Applications and Future Directions. Advances in Intelligent Systems and Computing, 2(798): 504

[10] Shen, X. and Huang, J. and Sun, Y. et al. 2021. Parallel Pathway Convolutional Neural Network with Low-rank Fusion for Brain Age Prediction. IEEE Xplore, 1(1): 5

[11] Khan, M. and Parker, G. 2023. Using Deep Convolutional Neural Networks to Abstract Obstacle Avoidance for Indoor Environments. IEEE Xplore, 1(1): 1-2

[12] Miah, J. and Khan, R. and Ahmed, S. et al. 2023. A comparative study of Detecting Covid 19 by Using Chest X- ray Images– A Deep Learning Approach. IEEE Xplore. 1(1): 1-3