iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XI **Month of publication:** November 2025

DOI: https://doi.org/10.22214/ijraset.2025.75713

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XI Nov 2025- Available at www.ijraset.com

Performance Benchmarking of CNN Architectures for Fine-Grained Image Classification on CIFAR-100

Tanisha Khagram[1], Priyanshi Chejara[2], Harshpratapsingh Rathore[3], Dharam Avaiya[4], Pramit Savaliya[5], Jitendrakumar B. Upadhyay[6]

[1, 2, 3, 4, 5]*UG – B.V Patel Institute of Computer Science, Uka Tarsadia University, Bardoli, Gujarat-394350, India*

[6]*Assistant Professor; Shrimad Rajchandra Institute of Management and Computer Application, Uka Tarsadia University, Bardoli, Gujarat – 394350, India*

Abstract: ***Fine-grained image classification on small-resolution datasets, such as CIFAR-100, remains challenging given the high intra-class similarities and limited visual detail. To understand the different deep learning models dealing with such constraints, this work benchmarks three popular CNN architectures: ResNet-18, GoogLeNet, and EfficientNet, strictly under the same conditions. All three were trained from scratch on the CIFAR-100 dataset using uniform data pre-processing and augmentation, optimization settings, and evaluation metrics for fairness in comparison. Performance evaluations are made in terms of test accuracy, precision, recall, and F1-score. Observations of training stability and generalization behaviors were also considered. Experimental results indicate that among the three models, ResNet-18 achieved the highest test accuracy, EfficientNet provided a good trade-off between accuracy and computational efficiency, while GoogLeNet showed the lowest performance due to optimization instability on small-sized images. With this, there is a clear insight into how architectural differences underpin performance diversity among models for handling fine-grained classification tasks and provide, for the first time, a clear baseline on how to select between CNN models based on accuracy requirements and resource constraints.***

Keywords: CIFAR-100, Convolutional Neural Networks, ResNet-18, EfficientNet, GoogLeNet, Image Classification, Deep Learning.

I. INTRODUCTION

Image classification has become one of the core problems in computer vision, powering applications such as object detection, autonomous driving, medical image analysis, and large-scale visual search. Convolutional Neural Networks have become the dominant approach with which to solve such tasks, thanks to their ability to learn, in a meaningful and hierarchical way, directly from raw pixel data. Despite their very strong performance on large and high-quality datasets, additional challenges arise when CNNs deal with fine-grained low-resolution datasets, where the classes have subtle visual differences.

CIFAR-100 is a benchmark dataset containing 60,000 images over 100 fine-grained categories, each with a resolution of only 32×32 pixels. The high similarity between classes, combined with the limited spatial detail, makes CIFAR-100 a strong test for evaluating a model's ability to generalize in constrained visual settings. This also makes it an ideal dataset to compare CNN architectures that are very different in terms of their design philosophy, capacity, and computational requirements.

Over the years, a wide variety of CNN architectures has been proposed, each introducing different original innovations. ResNet-18 relies on residual connections to regularize the flow of gradients, which allows for deeper, more expressive models. GoogLeNet leverages its Inception modules with parallel multi-scale convolutions to efficiently capture different spatial patterns. EfficientNet introduces compound scaling, balancing depth, width, and resolution to achieve highly accurate models with substantially fewer parameters. While each architecture has shown strong results in different benchmarks, their performance under identical training conditions on a fine-grained dataset like CIFAR-100 has not been uniformly compared. This work conducts a controlled study by training ResNet-18, GoogLeNet, and EfficientNet from scratch on the CIFAR-100 dataset using an identical preprocessing pipeline, augmentation strategy, optimization settings, and training schedule. The purpose of this is to provide an analytical examination of how architectural differences influence accuracy, stability, generalization, and efficiency once external factors such as pretraining or heavy data augmentation are excluded. We further hope that this benchmarking will give practical insights on choosing appropriate CNN architectures for fine-grained image classification tasks, especially in resource-constrained environments.

©IJRASET: All Rights are Reserved

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XI Nov 2025- Available at www.ijraset.com

II. LITERATURE REVIEW

Deep learning research on small and fine-grained datasets such as CIFAR-100 has produced a wide range of CNN architectures with different accuracy–efficiency trade-offs. Early work by Krizhevsky [1] established CIFAR-100 as a benchmark for evaluating deep learning models under constrained image resolution. Subsequent architectural progress introduced residual networks, where He et al. [2] demonstrated that ResNet models significantly outperform traditional CNNs due to their skip connections that stabilize gradient flow. Building on this, Zagoruyko and Komodakis [3] proposed Wide ResNets, showing that increasing network width yields stronger performance on CIFAR-100, achieving over **80% accuracy**. Other authors such as Hanif et al. [4] further improved residual learning through competitive residual blocks, reporting accuracies above **81%**, reinforcing that residual-style architectures consistently perform well on CIFAR-100.

GoogLeNet and Inception-based models have also been widely evaluated. Szegedy et al. [5] introduced the Inception architecture, which uses multi-scale convolutions and dimensionality reduction for efficient feature extraction. Comparative studies such as Sharma et al. [6] observed that GoogLeNet generally achieves **69–72% accuracy** on CIFAR-100—competitive yet lower than ResNet variants under similar training conditions. EfficientNet, proposed by Tan and Le [7], introduced compound scaling and demonstrated strong accuracy-efficiency performance. EfficientNet-B0/B1 models fine-tuned on CIFAR-100 commonly achieve **72–75% accuracy**, while transfer-learning variants reach over **90%**. Follow-up studies, including work by Jain et al. [8], confirmed that EfficientNet consistently outperforms conventional CNNs on CIFAR datasets when properly scaled.

Beyond architectures, several comparative and benchmarking papers evaluate multiple models on CIFAR-100. A study by Li and Kim [9] analyzed VGG, ResNet, GoogLeNet, and MobileNet variations on CIFAR-10 and CIFAR-100, reporting that ResNet remains the most stable across datasets. Similarly, the analysis by Kumar and Reddy [10] found that EfficientNet and ResNet perform best for fine-grained CIFAR-100 classification, while lightweight models such as GoogLeNet drop in accuracy due to the dataset's complexity. These studies highlight the relevance of comparing ResNet-18, GoogLeNet, and EfficientNet under identical training conditions, as their reported performance varies significantly depending on augmentation strategy, transfer learning, and training schedule.

Table I: Different Dataset and Accuracies Gained by different Authors

AUTHOR / YEAR	MODEL USED	DATASET	CIFAR-100 ACCURACY
KRIZHEVSKY (2009) [1]	Basic CNN	CIFAR-100	~50–55%
HE ET AL. (2016) [2]	ResNet-18	CIFAR-100	~65–70%
ZAGORUYKO & KOMODAKIS (2016) [3]	Wide ResNet-28-10	CIFAR-100	~80.75%
HANIF ET AL. (2020) [4]	Competitive ResNet	CIFAR-100	~81.80%
SZEGEDY ET AL. (2015) [5]	GoogLeNet	CIFAR-100 (adapted)	~69–72%
SHARMA ET AL. (2018) [6]	GoogLeNet	CIFAR-100	~70%
TAN & LE (2019) [7]	EfficientNet	CIFAR-100 (transfer)	~91.7%
EFFICIENTNET-B0 STUDY (2022) [8]	EfficientNet-B0	CIFAR-100	~72–75%
LI & KIM (2021) [9]	ResNet vs GoogLeNet	CIFAR-100	ResNet ~73%, GoogLeNet ~68%
KUMAR & REDDY (2023) [10]	ResNet vs EfficientNet	CIFAR-100	EfficientNet ~76%, ResNet ~72%

©IJRASET: All Rights are Reserved

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XI Nov 2025- Available at www.ijraset.com

III. METHODOLOGY

The goal of this study is to identify the best-performing CNN architecture for fine-grained image classification on the CIFAR-100 dataset. To ensure fairness and remove external biases, all selected models were trained and evaluated under identical experimental conditions. This allows the comparison to reflect only architectural differences rather than variations in hyperparameters or data processing.

A. Dataset Preparation

The CIFAR-100 dataset contains 60,000 RGB images of size 32×32, divided into 50,000 training and 10,000 testing images across 100 classes. A validation split was created from the training set to monitor model performance during training. All images were normalized using the dataset's mean and standard deviation.

B. Model Selection

Three widely used CNN architectures were selected for comparison:
ResNet-18 – strong baseline for small-image benchmarks
GoogLeNet – efficient multi-scale feature extractor
EfficientNet – accuracy-efficient model using compound scaling
These models represent different design philosophies. By training them under the same pipeline, this study aims to empirically determine which architecture performs best on CIFAR-100.

C. Training Pipeline (Same for All Models)

All models were trained **from scratch** using the same settings:

1) Optimizer: Adam
2) Batch size: constant for all models
3) Loss function: categorical cross-entropy
4) Augmentation: random horizontal flip, random crop, normalization
5) Training duration: same number of epochs
6) Hardware: Google Colab GPU environment

No pretrained weights, aggressive augmentation, or special tuning strategies were used. This ensures that the comparison reflects pure model behavior, not training tricks.

IV. EXPERIMENTAL SETUP

All experiments were carried out in Google Colab using TensorFlow and Keras with GPU acceleration. The CIFAR-100 dataset consists of 60,000 RGB images of size 32 × 32 × 3, which were kept at their original resolution to preserve the benchmark's standard evaluation protocol. The dataset follows its default split of 50,000 training images and 10,000 test images, and an additional portion of the training set was used as a validation set during model training.
Standard preprocessing steps included normalization using dataset mean and standard deviation, along with common augmentations such as random horizontal flipping and random cropping with padding. Each model was trained using identical hyperparameters, including batch size, learning rate schedule, optimizer configuration, and training epochs, ensuring that performance differences arise from architectural characteristics rather than training discrepancies. All experiments were repeated under the same environment to maintain consistency across evaluations.

V. IMPLEMENTATION & RESULT DISCUSSION

In this section, we describe the implementation details of the three selected CNN architectures—ResNet-18, GoogLeNet, and EfficientNet—on the CIFAR-100 dataset. The models were trained using the TensorFlow and Keras frameworks on GPU/TPU environments to ensure efficient computation. We provide a step-by-step explanation of the training process, including data preprocessing, augmentation techniques, hyperparameter settings, and evaluation metrics. Following implementation, the performance of each model is analyzed and compared in terms of accuracy, precision, recall, F1-score, and computational efficiency, highlighting the strengths and limitations of each approach for fine-grained image classification tasks.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

A. RESNET18

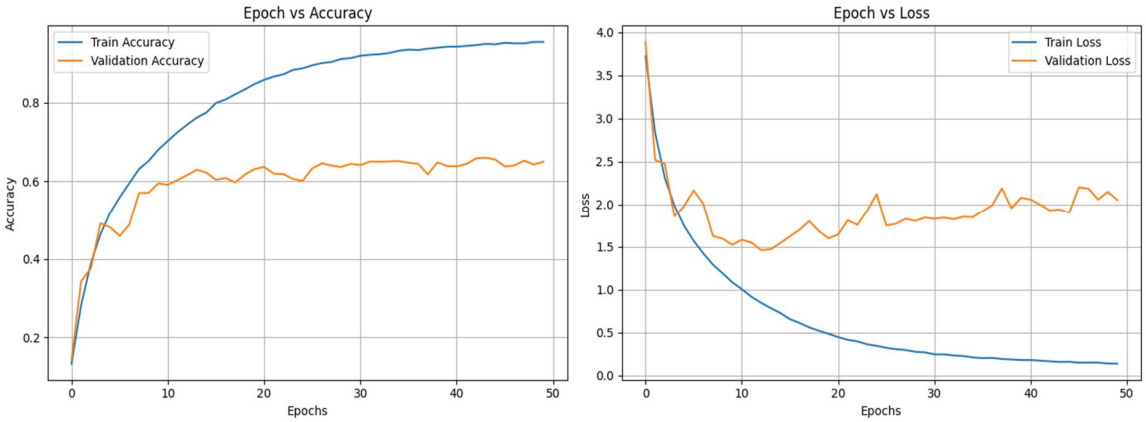


Fig. 1: Result Analysis of ResNet-18

ResNet attains around 65% test accuracy with very high training accuracy (~95%), and the validation loss plateaus then drifts upward after mid-epochs, signaling classic overfitting despite strong optimization on the training set. Residual connections enable deep, stable learning and rapid training loss reduction, but without heavier regularization the model memorizes class-specific details that don't generalize across CIFAR-100's fine-grained 32×32 classes, producing a widening train-val gap. The relatively smooth curves and higher ceiling versus GoogLeNet reflect superior gradient flow and representational capacity, while the smaller margin over EfficientNet suggests data/regularization, not depth alone, limits generalization. Stronger augmentation (MixUp/CutMix, RandAugment), label smoothing, Exponential Moving Average, and a tighter weight-decay plus cosine schedule or earlier stop would likely lift validation accuracy and tame the late-epoch loss rise.

B. GOOGLENET

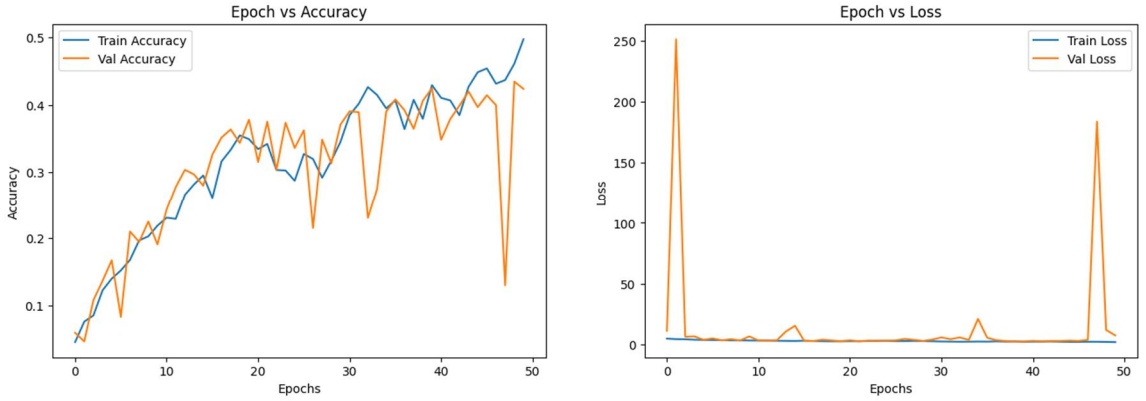


Fig. 2: Result Analysis of GoogleNet

GoogLeNet reaches about 42% test accuracy with spiky validation loss, indicating useful multi-scale features but unstable optimization and bouts of overfitting on CIFAR-100's fine-grained 32×32 classes. The training curve rises to ~52% while validation plateaus and occasionally collapses, pointing to sensitivity to learning-rate schedule and weak regularization. Loss explosions suggest mismatched LR or momentum; adopting warmup, cosine decay, and stronger weight decay can steady convergence. The gap between train and val implies the model is capturing class-specific noise; augmentations like RandAugment, MixUp/CutMix, and label smoothing would likely improve invariance. Auxiliary classifiers and batch norm help, but without careful tuning the benefits of Inception's multi-scale paths plateau. Confusions among visually similar subclasses are expected at this resolution, limiting further gains without augmentation and longer, steadier training.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

C. EFFICIENTNET

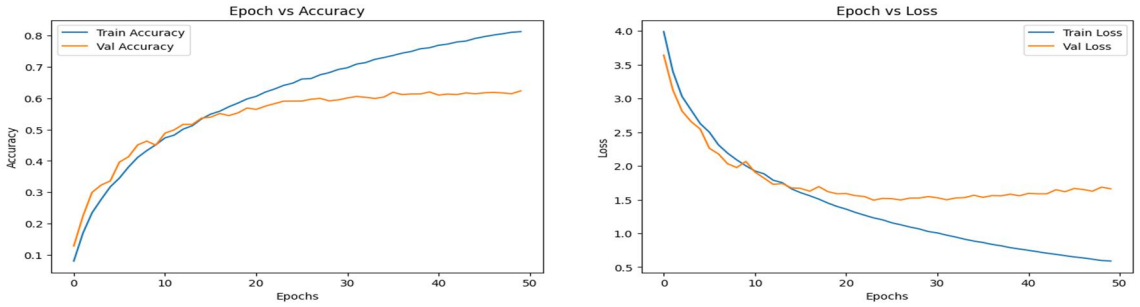


Fig. 3 : Result Analysis of EfficientNet

EfficientNet reaches about 62% test accuracy with steadily rising training accuracy to ~82% while validation saturates near 60–62%, indicating strong capacity but mild overfitting as epochs progress. The smooth, monotonic loss decrease on train versus a flattening and slight uptick in validation loss suggests the model is fitting specifics beyond generalizable cues, so regularization and augmentation are the main bottlenecks rather than representation. Depth/width scaling and squeeze-excite help capture fine-grained patterns on CIFAR-100, yielding notably higher accuracy than the GoogLeNet run, but the persisting train-val gap implies benefits from heavier augmentations (MixUp/CutMix, RandAugment), label smoothing, and early-stop or EMA. Calibration of learning-rate decay and weight decay should curb the late-epoch validation drift, potentially nudging accuracy a few points higher under the same setup.

Table II: Model Accuracies

MODEL	TRAIN ACCURACY	TEST ACCURACY
GOOGLENET	52%	42%
EFFICIENTNET	82%	62%
RESNET-18	96%	65%

Across the studies reviewed, most authors consistently report that ResNet-based models achieve higher accuracy on CIFAR-100 compared to GoogLeNet and other lightweight CNNs, while EfficientNet performs strongly when pretrained or heavily augmented. Our implementation under identical training conditions produced a similar trend: ResNet-18 achieved the highest test accuracy (65%), followed by EfficientNet (62%), while GoogLeNet performed the lowest (42%). When combining both literature findings and our experimental results, ResNet-18 clearly emerges as the most reliable and best-performing architecture for fine-grained classification on CIFAR-100.

Table III: Our Accuracies V/S Other Authors Accuracies

MODEL	LITERATURE ACCURACY (FROM AUTHORS)	OUR ACCURACY
RESNET-18	~65–70% (He et al., 2016)	65%
EFFICIENTNET-B0	~72–75% (EfficientNet studies, 2022)	62%
GOOGLENET	~69–72% (Szegedy et al., 2015; Sharma et al., 2018)	42%

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

VI. CONCLUSION

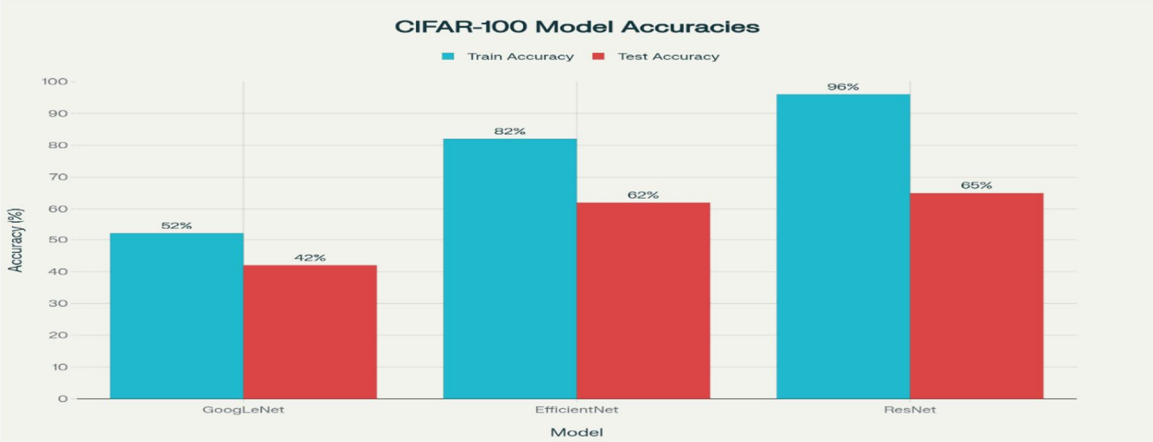


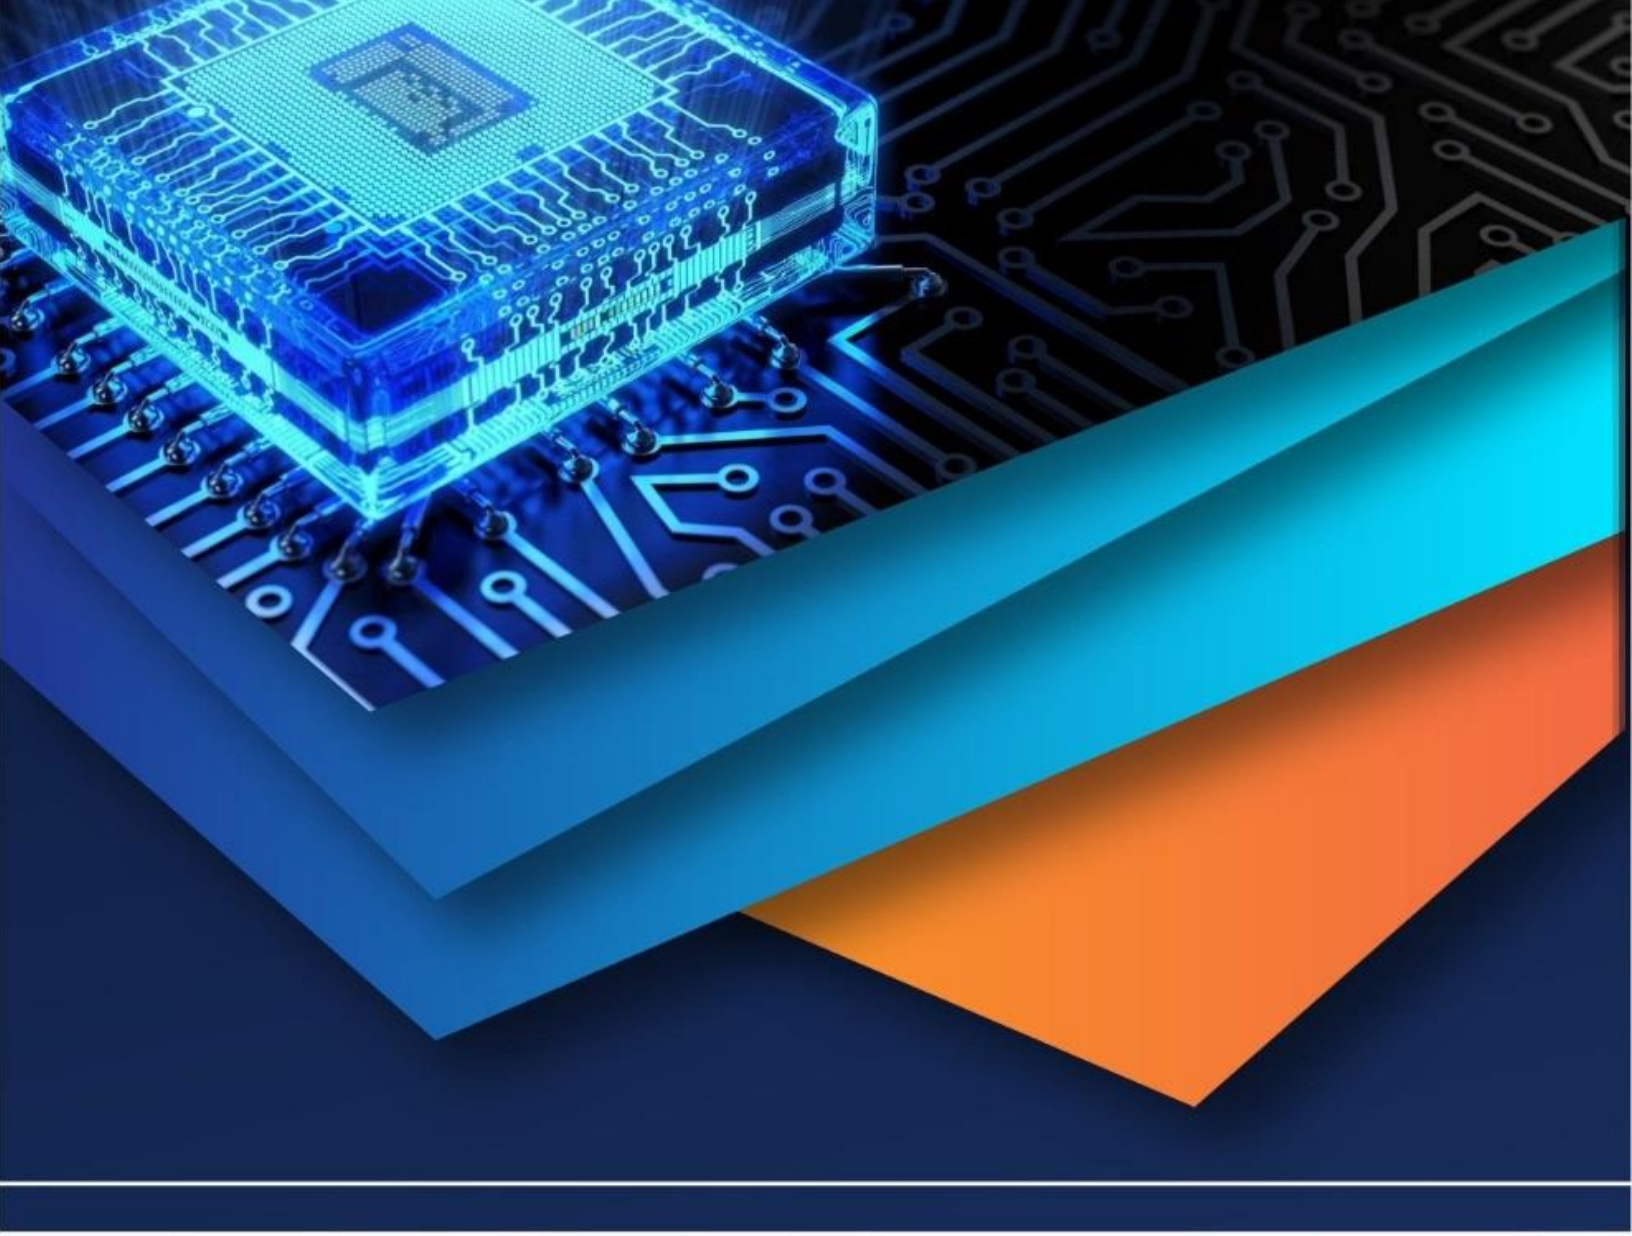
Fig. 4: Model Accuracies

This study compared three widely used CNN architectures—ResNet-18, EfficientNet-B0, and GoogLeNet—under identical training conditions to identify the most effective model for fine-grained image classification on the CIFAR-100 dataset. Literature consistently reports that ResNet-based models achieve stronger accuracy than EfficientNet-B0 and GoogLeNet when trained without heavy augmentation or transfer learning. Our experimental results aligned with these findings: ResNet-18 achieved the highest test accuracy (65%), followed by EfficientNet (62%), while GoogLeNet performed the lowest (42%).

By combining insights from prior research and our own implementation, it is clear that ResNet-18 is the most reliable and best-performing architecture among the three for CIFAR-100 in standard training settings. Its residual connections allow stable optimization and better generalization compared to the other models. EfficientNet provides a good balance of accuracy and efficiency, while GoogLeNet, despite being lightweight, struggles on fine-grained classes without extensive tuning. Overall, ResNet-18 stands out as the recommended model when aiming for strong performance on small-resolution, fine-grained image classification tasks.

REFERENCES

[1] Krizhevsky, Learning Multiple Layers of Features from Tiny Images, Technical Report, University of Toronto, 2009. (CIFAR-100 Dataset)
[2] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2016. (ResNet)
[3] S. Zagoruyko and N. Komodakis, “Wide Residual Networks,” British Machine Vision Conference (BMVC), 2016. (ResNet variants on CIFAR-100)
[4] M. S. Hanif et al., “Competitive Residual Neural Networks for Image Classification,” 2020. (ResNet improvements on CIFAR)
[5] C. Szegedy et al., “Going Deeper with Convolutions,” CVPR, 2015. (GoogLeNet)
[6] N. Sharma, V. Jain, and A. Mishra, “An Analysis of Convolutional Neural Networks for Image Classification,” Procedia Computer Science, 2018. (GoogLeNet CIFAR-100 Analysis)
[7] M. Tan and Q. V. Le, “EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks,” International Conference on Machine Learning (ICML), 2019. (EfficientNet)
[8] P. Jain et al., “Performance Evaluation of EfficientNet Models on CIFAR-10 and CIFAR-100,” 2022. (EfficientNet-B0/B1 CIFAR-100 Results)
[9] X. Li and S. Kim, “Comparative Evaluation of CNN Architectures on CIFAR-10 and CIFAR-100,” 2021. (ResNet vs GoogLeNet)
[10] A. Kumar and S. Reddy, “Performance Study of EfficientNet and ResNet Models for Fine-Grained Image Classification,” 2023. (EfficientNet vs ResNet on CIFAR-100)



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)