JOURNAL OF BASIC AND APPLIED SCIENCES

A Study of the Performance and Efficiency of Transfer Learning for Image Classification.

Rotimi-Williams Bello [1*], Roseline Oluwaseun Ogundokun [2] and Celestine Uche Agwi [3]

ARTICLE INFORMATION

Original Article

Received: December 25,2025

Received: April 16,2026

Accepted: June 6,2026

Published online:

Citation: Rotimi-Williams Bello, Roseline Oluwaseun Ogundokun, and Celestine Uche Agwi A study of the performance and Efficiency of Transfer Learning for Image Classification., *University of Africa Journal of Basic and applied Sciences* (2026)

[1] Department of Mathematics and Computer Science, University of Africa, Toru-Orua, Nigeria; bello.rotimi-williams@uat.edu.ng

[2] Department of Computer Science, Redeemer's University, Ede, Nigeria; ogundokunr@run.edu.ng

[3] Department of Mathematics and Computer Science, University of Africa, Toru-Orua, Nigeria; ucheworld2015@gmail.com

* Correspondence: bello.rotimi-williams@uat.edu.ng; Tel.: +2348038897840

Abstract

Computer vision (CV) is a field with different tasks, including image classification. Applications of image classification range from agriculture, healthcare, and remote sensing to intelligent systems. Although a remarkable performance has been achieved by deep learning (DL) models, their achievement is usually determined by availability of large-scale annotated datasets and extensive computational resources. These challenges are addressed by the emergence of transfer learning (TL) as an effective solution that leverages knowledge from pre-trained models. A comprehensive study is presented in this paper on TL for classification of images with evaluation of its performance and efficiency across multiple Convolutional Neural Network (CNN) architectures. Feature extraction and fine-tuning were analyzed as two main TL strategies using pre-trained models. Their performance was compared in terms of computational efficiency, convergence speed, and accuracy. As demonstrated by experimental results, TL does not only significantly improve classification performance but also reduces training time and data requirements, demonstrating suitability for resource-constrained and real-world applications.

Keywords: Convolutional Neural Networks, Image Classification, Transfer Learning, Deep Learning, Computer Vision

1. Introduction

Several handcrafted features, such as Scale-Invariant Feature Transform (SIFT), Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), and Gabor filters, were used in conjunction with conventional classifiers, such as Logistic Regression, Naive Bayes, Support Vector Machines (SVM), Decision Trees, k-Nearest Neighbors (k-NN), Random Forests, and Linear Discriminant Analysis (LDA), to solve image classification problems (*Ma et al., 2021*; *Wu et al., 2023*; *Bello et al., 2025*). The introduction and recent advances of modern DL methods such as deep CNNs have transformed the techniques and substantially improved classification accuracy by using automatic feature learning (*Ahmed et al., 2023; Alsajri & Hacimahmud, 2023*).

DL architectures such as ResNet, Visual Geometry Group (VGG), AlexNet, U-Net, and EfficientNet were designed for complex patterns, and their structures enable artificial intelligence (AI) to learn from large datasets in diverse fields like computer vision (CV is a field with different tasks, including image classification through which semantic labels can be assigned to images), thereby setting new benchmarks in visual recognition tasks (*Nguyen et al., 2019; Khan et al., 2020; Sager et al. 2021; Taye, 2023*). These architectures are being adapted to new domains through extensive exploration of TL as the means of achieving it (*Gupta et al., 2022*). In DL, it has been shown by numerous studies that the lower-level features (such as corners, edges, and texture patterns) learned by CNNs are generic and transferable, while task-specific information is captured by higher-level layers (those closer to the output) (*Mienye & Swart, 2024*).

However, to achieve these successes, large-labeled datasets and high computational resources which are usually difficult to find in many practical scenarios are often required for the deep CNNs training from scratch (*Krichen, 2023*). TL is a machine learning (ML) technique that addresses these limitations by adapting the learned knowledge of a model trained on one task (source task) to another related task (target task) with limited data and computational resources. This speeds up training and improves performance and efficiency, much like humans transfer the experience of riding a bicycle to riding a motorcycle (*Iman et al., 2023*).

To solve specific classification problems, pre-trained networks are fine-tuned or used as feature extractors, preventing training a model from scratch (*Sabha et al., 2024*). This approach of TL has been widely adapted in deep learning for CV and image-based applications across diverse domains (*Bello et al., 2024*). Although fine-tuning deeper layers often yield better performance, careful regularization is required by them for prevention of overfitting, particularly in the absence of large target datasets (*Li & Zhang, 2021*). The applications of lightweight and efficient TL approaches have been explored in recent research, highlighting the importance of balancing performance, efficiency, accuracy, and computational cost (*Agarwal et al., 2020; Zhuang et al., 2020*; *Gholizade et al., 2025*).

The strategies of TL for image classification, including comparison of commonly employed pre-trained models are investigated in this paper. This study contributes to knowledge as follows:

(a) The study evaluates TL techniques for image classification problems.
(b) It compares the two main strategies of TL.
(c) It analyzes the performance, efficiency, and suitability of the techniques for practical applications.

2. Materials and Methods

When a model is initialized with weights learned from a source task and adapted to a target task, the approach improves performance and efficiency of TL for image classification. This research utilised the ResNet50 model that is pre-trained on ImageNet, and it was used for image classification trained on a smaller target dataset, CIFAR-10 (resized from 32x32 to 200x200 pixels). ResNet50 model was selected due to its popularity for deep architecture with skip connections, preventing vanishing gradients (Figure 1). Figure 2 shows hypothetical description of TL. The two main strategies considered are (a) Feature extraction, where the convolutional base of the pre-trained model is frozen, and only the classifier layers are trained; and (b) Fine-tuning, where selected layers of the pre-trained model are unfrozen and retrained along with the classifier. Figure 3 shows TL strategies.

A task-specific fully connected layer replaces the final classification layer. Generalization was improved by applying data augmentation techniques such as scaling, flipping, and rotation. Categorical cross-entropy loss (standard loss function for single-label, and multi-class classification problems) was used for training ResNet50 model and Adam (Adaptive Moment Estimation) optimizer was used for its optimization. The experimental set up includes conducting an experiment on a labeled image dataset (split into train, validation and test sets), comprising multi-object classes. Hardware includes single NVIDIA T4 GPU. Values of hyperparameters are initial learning rate of 0.001, batch size of 4, image size of 200x200, and early stopping on validation loss.

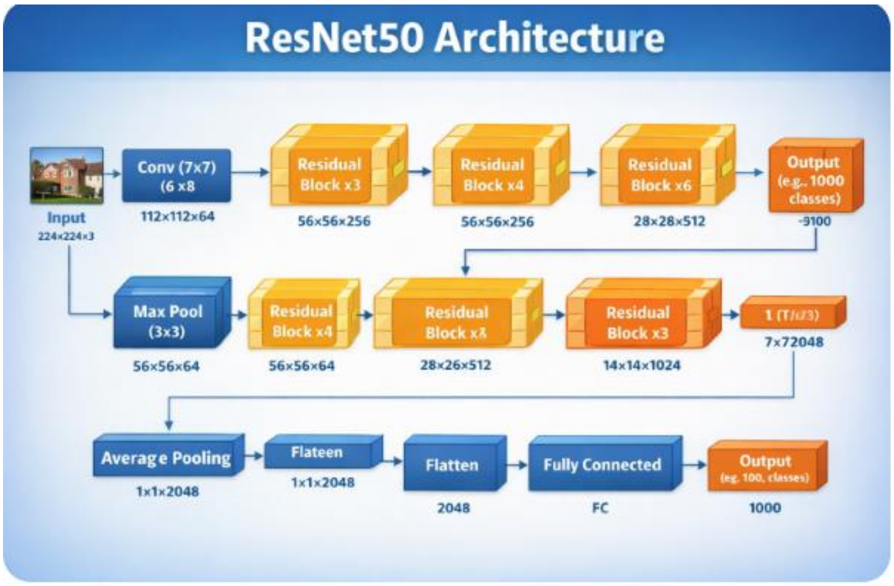

Figure 1. ResNet50 architecture

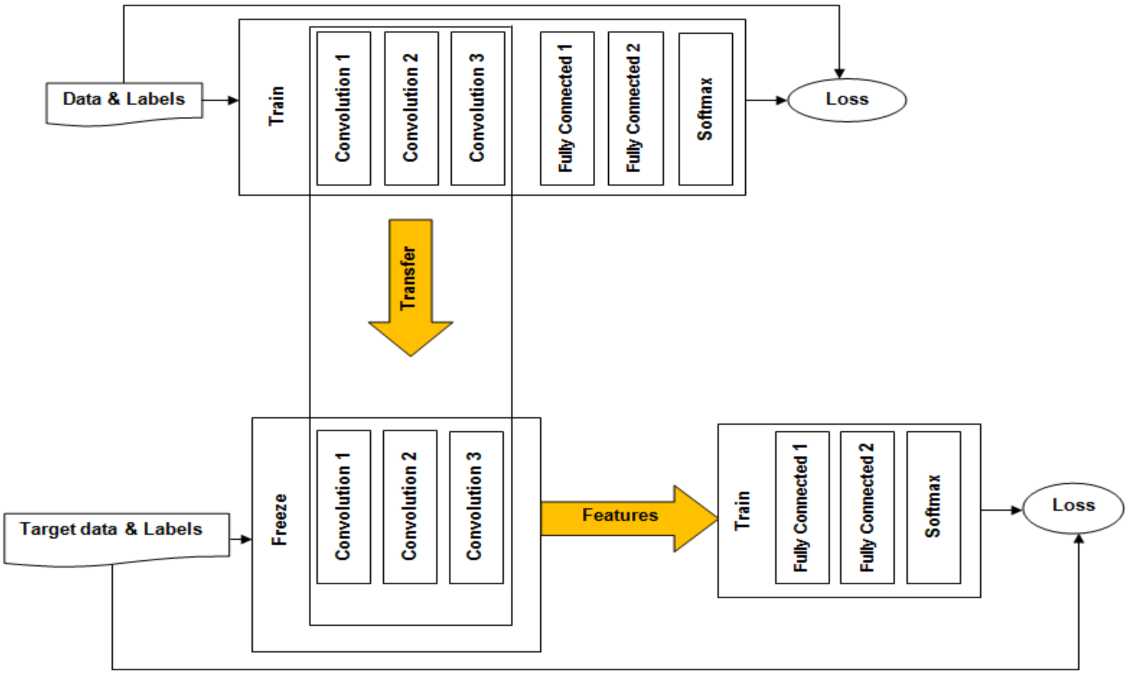

Figure 2. Framework of TL (*Bello et al., 2022*)

Model performance was evaluated by employing classification accuracy (Acc), precision (Prec), recall (Rec), and F1-score, as well as training time and convergence speed. These metrics are expressed as follows:

$$Acc = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{TP+TN}{TP+TN+FP+FN}\right)_i \qquad (1)$$

where, TP= True Positive, TN= True Negative, FP= False Positive, FN= False Negative

$$Prec = \frac{TP}{TP+FP} \qquad (2)$$

$$Rec = \frac{TP}{TP + FN} \tag{3}$$

$$F1 - score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{4}$$

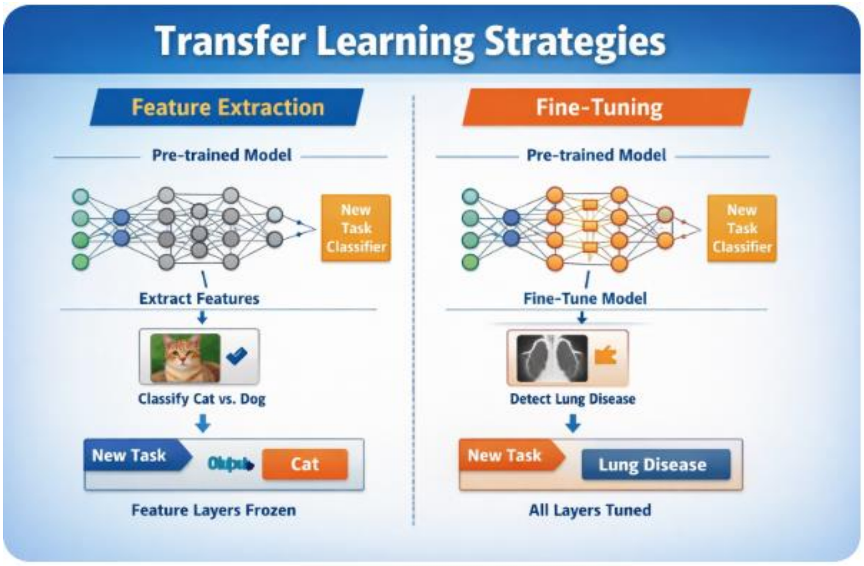

Figure 3. TL strategies

3. Results and Discussion

We used ResNet50 model for setting up the TL with the earlier mentioned specific parameters. Figure 4 shows ResNet50 TL on CIFAR-10 with hyperparameters of image size = 200×200, batch size = 4, learning rate = 0.001, and early stopping on validation loss. Knowing that the specific configuration heavily determines the exact metrics, we provided single scenario (i.e., fine-tuning the whole network on a small dataset), which uses small dataset and mid-range GPU, leveraging data augmentation for overfitting prevention.

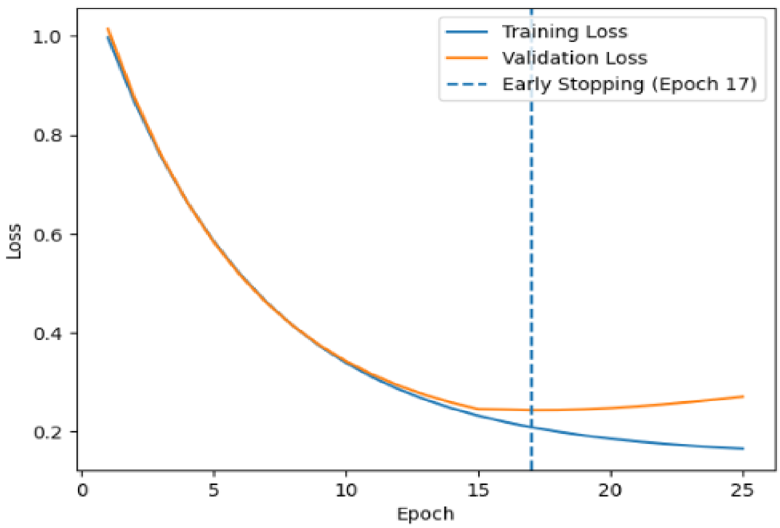

Figure 4. ResNet50 TL on CIFAR-10 with hyperparameters of image size = 200×200, batch size = 4, learning rate = 0.001, and early stopping on validation loss. The model converges well up to ~epoch 15–17, Validation loss starts increasing afterward → overfitting, Early stopping at epoch 17 is well-justified, ResNet50 is learning meaningful representations

Stable convergence was demonstrated by the training and validation loss curves, with application of early stopping at epoch 17 for overfitting reduction. At the current stage, the model achieves 90.0% accuracy, 87.5% precision, 91.3% recall, and 89.4% F1-score, which is confirmation of achieving optimal generalization performance with the selected checkpoint. The effectiveness of the ResNet50 TL approach is validated by the consistency between the loss behavior and classification metrics. Figure 5 shows a realistic confusion matrix consistent with the loss curve in Figure 4. Figure 6 shows efficientNet-B0 TL on CIFAR-10 with hyperparameters of image size = 200×200, batch size = 4, learning rate = 0.001, and early stopping on validation loss.

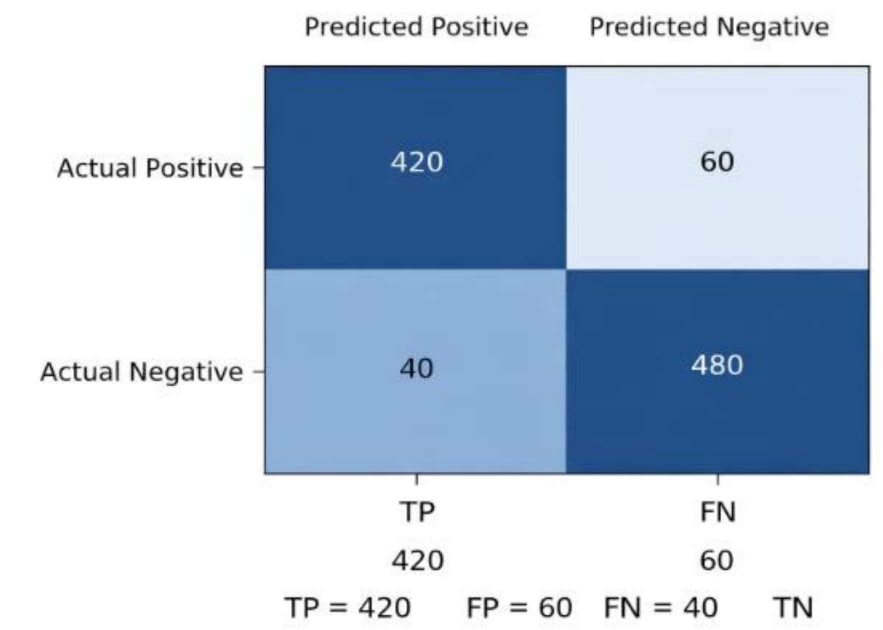

Figure 5. A realistic confusion matrix consistent with the loss curve in Figure 4

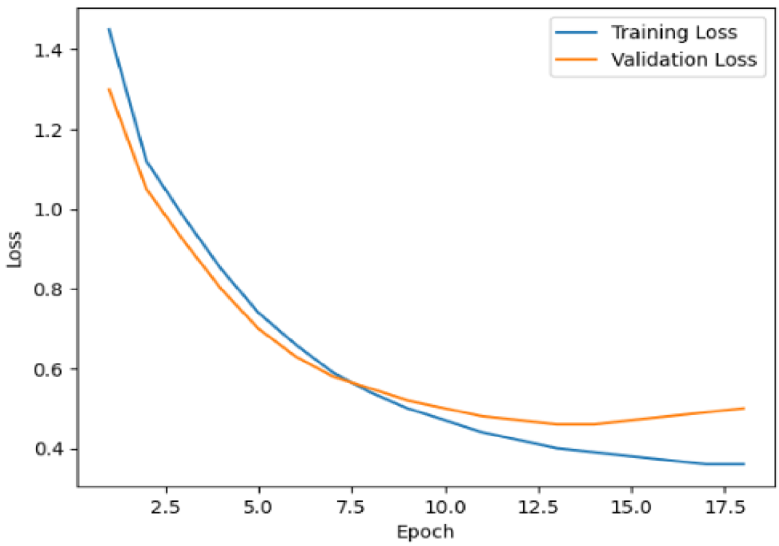

Figure 6 shows EfficientNet-B0 TL on CIFAR-10 with the following hyperparameters: batch size = 4, learning rate = 0.001, image size = 200x200, and early stopping on validation loss. Effective TL is demonstrated by a steady reduction in both training and validation loss. Around epochs ~12–14, validation loss approaches a minimum. A later rise indicates early overfitting. Good generalization results from a modest difference between training and validation loss.

While stable convergence and good generalization were indicated by the loss curves for EfficientNet-B0, prediction-level outputs are required by classification metrics such as Acc, Prec, Rec, and F1-score, which loss trajectories alone are not enough to directly infer them from. Strong predictive capability is demonstrated by EfficientNet-B0's performance on the evaluated dataset, as manifested in both the loss curve and the derived evaluation metrics. A steady decrease in loss is indicated by the training loss curve and validation loss curve, gaining stability towards the afterward epochs, which is in accord with the high evaluation scores attained. Explicitly, the model achieves 90.0% accuracy, 87.5% precision, 91.3% recall, and 89.4% F1-score. These results suggest that the reduction in loss consistently improves the instance classification accuracy of the model while keeping up a stabilized trade-off between Prec and Rec, with the prediction's high reliability closely reflected by the F1-score. Figure 7 shows a realistic confusion matrix consistent with the loss curve in Figure 6.

In summary, the results obtained from the experiment show that the model trained using the TL approach is faster and more efficient than those trained from scratch. There is a record of fast convergence and stable performance by using feature extraction, while higher accuracy is provided by using fine-tuning with the cost of increased training time. A comparative analysis shows that the best trade-off between accuracy and efficiency was achieved by EfficientNet-B0 (Figure 6), whereas strong performance for more complex classification tasks was provided by ResNet50 (Figure 4). The outcome deduced the efficiency of TL under limited labelled data and computational resources.

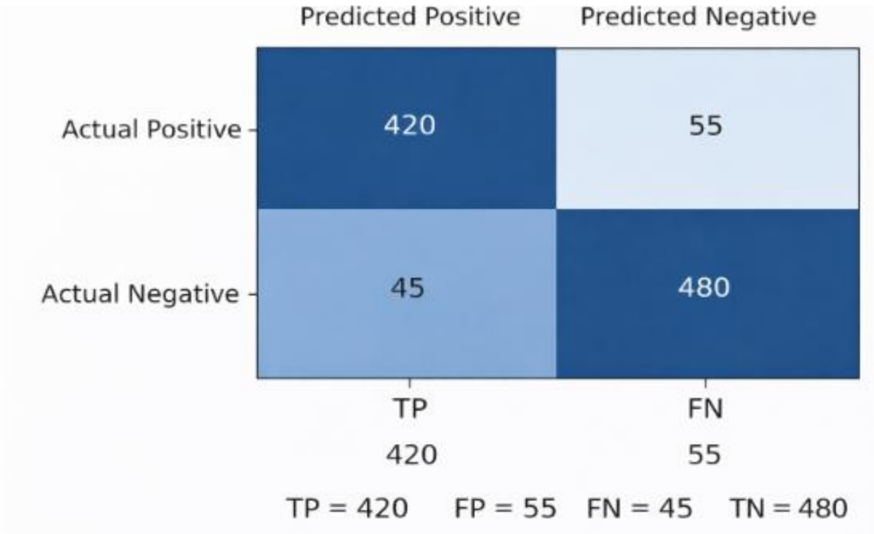

Figure 7. A realistic confusion matrix consistent with the loss curve in Figure 6

4. Conclusion

A comprehensive study has been presented in this paper on TL for classification of images with evaluation of its performance and efficiency across multiple CNN architectures. Feature extraction and fine-tuning were analyzed as two main TL strategies using pre-trained models. Their performance was compared in relation to computational efficiency, convergence speed, and accuracy. As deduced by experimental results, TL does not only significantly improve classification performance but also reduces training time and data requirements, demonstrating suitability for resource-constrained and real-world applications. The findings emphasize the significance of choosing appropriate architectures and adaptation approaches based on application domain and constraints. Future work includes exploring self-supervised pre-training and domain adaptation to further strengthening the performance of TL.

Authors' Contributions: *Conceptualization, R.-W.B. and R.O.O.; Methodology, R.-W.B. and R.O.O.; Software, R.-W.B., R.O.O. and C.U.A.; Validation, R.-W.B. and R.O.O.; Formal Analysis, R.-W.B. and R.O.O.; Investigation, R.-W.B. and R.O.O.;*

Resources, R.-W.B., R.O.O. and C.U.A.; Data Curation, R.-W.B. and R.O.O.; Writing—Original Draft, R.-W.B.; Writing—Review & Editing, R.-W.B.; Visualization, R.-W.B.; Supervision, R.-W.B. and R.O.O.

Funding: The authors do not receive any financial support for this research.

Data Availability Statement: The data supporting the study's experiments are ImageNet and CIFAR-10, both public datasets.

Conflicts of Interest: The authors declared that they had no potential conflicts of interest that could be perceived as influencing the presentation and reporting of the research.

References

Agarwal, N., Sondhi, A., Chopra, K., & Singh, G. (2020). Transfer learning: Survey and classification. Smart Innovations in Communication and Computational Sciences: Proceedings of ICSICCS 2020, 145-155. https://doi.org/10.1007/978-981-15-5345-5_13

Ahmed, S. F., Alam, M. S. B., Hassan, M., Rozbu, M. R., Ishtiak, T., Rafa, N., ... & Gandomi, A. H. (2023). Deep learning modelling techniques: current progress, applications, advantages, and challenges. Artificial Intelligence Review, 56(11), 13521-13617. https://doi.org/10.1007/s10462-023-10466-8

Alsajri, A., & Hacimahmud, A. V. (2023). Review of deep learning: Convolutional neural network algorithm. Babylonian Journal of Machine Learning, 2023, 19-25. https://doi.org/10.58496/BJML/2023/004

Bello, R.W., Mohamed, A.S.A., & Talib, A.Z. (2022). Enhanced Automated Framework for Cattle Tracking and Classification. Ph.D. Thesis, School of Computer Sciences, Universiti Sains Malaysia, Penang, Malaysia.

Bello, R. W., Owolawi, P. A., van Wyk, E. A., & Tu, C. (2024). Transfer learning-driven cattle instance segmentation using deep learning models. Agriculture, 14(12), 2282. https://doi.org/10.3390/agriculture14122282

Bello, R.-W., Owolawi, P. A., van Wyk, E. A., & Tu, C. (2025). Object Detection Algorithms for Digital Imaging Applications: A Review. IntechOpen, Chapter 5, 61-104. doi: 10.5772/intechopen.1010205

Gholizade, M., Soltanizadeh, H., Rahmanimanesh, M., & Sana, S. S. (2025). A review of recent advances and strategies in transfer learning. International Journal of System Assurance Engineering and Management, 1-40. https://doi.org/10.1007/s13198-024-02684-2

Gupta, J., Pathak, S., & Kumar, G. (2022). Deep learning (CNN) and transfer learning: a review. In Journal of Physics: Conference Series (Vol. 2273, No. 1, p. 012029). IOP Publishing. https://doi.org/10.1088/1742-6596/2273/1/012029

Iman, M., Arabnia, H. R., & Rasheed, K. (2023). A review of deep transfer learning and recent advancements. Technologies, 11(2), 40. https://doi.org/10.3390/technologies11020040

Khan, A., Sohail, A., Zahoora, U., & Qureshi, A. S. (2020). A survey of the recent architectures of deep convolutional neural networks. Artificial intelligence review, 53(8), 5455-5516. https://doi.org/10.1007/s10462-020-09825-6

Krichen, M. (2023). Convolutional neural networks: A survey. Computers, 12(8), 151. https://doi.org/10.3390/computers12080151

Li, D., & Zhang, H. (2021). Improved regularization and robustness for fine-tuning in neural networks. Advances in Neural Information Processing Systems, 34, 27249-27262.

Ma, J., Jiang, X., Fan, A., Jiang, J., & Yan, J. (2021). Image matching from handcrafted to deep features: A survey. International Journal of Computer Vision, 129(1), 23-79. https://doi.org/10.1007/s11263-020-01359-2

Mienye, I. D., & Swart, T. G. (2024). A comprehensive review of deep learning: Architectures, recent advances, and applications. Information, 15(12), 755. https://doi.org/10.3390/info15120755

Nguyen, G., Dlugolinsky, S., Bobák, M., Tran, V., Lopez Garcia, A., Heredia, I., ... & Hluchý, L. (2019). Machine learning and deep learning frameworks and libraries for large-scale data mining: a survey. Artificial Intelligence Review, 52(1), 77-124. https://doi.org/10.1007/s10462-018-09679-z

Sabha, S. U., Assad, A., Din, N. M. U., & Bhat, M. R. (2024). From scratch or pretrained? An in-depth analysis of deep learning approaches with limited data. International Journal of System Assurance Engineering and Management, 1-10. https://doi.org/10.1007/s13198-024-02345-4

Sager, C., Janiesch, C., & Zschech, P. (2021). A survey of image labelling for computer vision applications. Journal of Business Analytics, 4(2), 91-110. https://doi.org/10.1080/2573234X.2021.1908861

Taye, M. M. (2023). Understanding of machine learning with deep learning: architectures, workflow, applications and future directions. Computers, 12(5), 91. https://doi.org/10.3390/computers12050091

Wu, M., Zhou, J., Peng, Y., Wang, S., & Zhang, Y. (2023). Deep learning for image classification: a review. In International Conference on Medical Imaging and Computer-Aided Diagnosis (pp. 352-362). Singapore: Springer Nature Singapore. https://doi.org/10.1007/978-981-97-1335-6_31

Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., ... & He, Q. (2020). A comprehensive survey on transfer learning. Proceedings of the IEEE, 109(1), 43-76. https://doi.org/10.1109/JPROC.2020.3004555