

An Efficient MobileNet-Based Deep Learning Framework for Automated Dog Breed Classification

Mykhailo Paslavskyi^{1,*}, Ibrahim Ahmed Ibrahim Abd Elaal El Sayed², Lin Chih-Hao³, Amarilys González García⁴, Vinothini Kasinathan⁵, Kohila Malar Kalesamy⁶

¹Department of Computer Science, Ukrainian National Forestry University, Lviv, Lviv Oblast, Ukraine.

²Department of Biomedical Sciences, Zewail City of Science and Technology, October Gardens, Giza Governorate, Egypt.

³Department of Manufacturing, Taipei Tech Industries, Da'an District, Taipei, Taiwan.

⁴Department of Research and Development, Placental Histotherapy Center, Playa, Havana, Cuba.

^{5,6}School of Computing, Asia Pacific University of Technology and Innovation, Kuala Lumpur, Wilayah Persekutuan Kuala Lumpur, Malaysia.

mykhailo.paslavskyi@nltu.edu.ua¹, s-ibrahim.aibrahim@zewailcity.edu.eg², lin.chihhao@taipeitech.tw³, agonzalez@miyares-cao.cu⁴, vinothini@apu.edu.my⁵, kohila.malar@apu.edu.my⁶

*Corresponding author

Abstract: In computer vision, classifying dog breeds is a difficult problem because there are many breed classes, high similarity among related breeds, variations in pose and illumination conditions, and the presence of noisy, low-quality, or blurry images. In this paper, researchers propose a transfer learning approach using TensorFlow Hub to develop a large-scale multi-class image classification system for accurate dog breed identification. Researchers use pre-trained architectures like MobileNetV2 and EfficientNet as feature extractors since they can build rich, discriminative image representations at lower computational cost than training a deep neural network. The Kaggle Dog Breed Identification dataset has over 20,000 labelled photos of 120 dog breeds, which researchers used to train, validate, and assess the model. Image quality, data diversity, and overfitting are improved by resizing, pixel normalisation, and data augmentation through rotation, flipping, zooming, and shifting. The classification model is trained with categorical cross-entropy loss and an optimisation approach and assessed using Top-1 and Top-5 accuracy. Experimental studies reveal that transfer learning can classify visually related breeds well. The study also addresses class imbalance, intra-breed differences, background interference, and misclassification among breeds with similar physical traits. The proposed method shows that transfer learning can be used for large-scale dog breed recognition in pet identification systems, veterinary support, animal shelters, rescue services, mobile apps, edge devices, and real-time intelligent animal monitoring systems.

Keywords: Transfer Learning; Breed Recognition; Deep Learning (DL); Computer Vision; Data Augmentation; TensorFlow Hub; Class Imbalance; Background Interference; Veterinary Support.

Cite as: M. Paslavskyi, I. A. I. A. E. E. Sayed, L. Chih-Hao, A. G. García, V. Kasinathan, and K. M. Kalesamy, "An Efficient MobileNet-Based Deep Learning Framework for Automated Dog Breed Classification," *AVE Trends in Intelligent Computing Research*, vol. 1, no. 2, pp. 91–101, 2026.

Journal Homepage: <https://avepubs.com/user/journals/details/ATICR>

Received on: 21/08/2025, **Revised on:** 10/10/2025, **Accepted on:** 25/10/2025, **Published on:** 07/06/2026

DOI: <https://doi.org/10.64091/ATICR.2026.000341>

1. Introduction

Copyright © 2026 M. Paslavskyi *et al.*, licensed to AVE Trends Publishing Company. This is an open access article distributed under [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/), which permits copying, remixing, redistributing, transformation, and building upon the material in any medium so long as the original work is properly cited.

Classifying dog breeds is an important and challenging computer vision task because it requires the model to distinguish among many classes that often show only subtle visual differences [1]. In many general image classification tasks, objects in different classes have clearly distinguishable characteristics. However, when identifying a dog breed, body structure, coat color, facial features, ear shape, and overall appearance can be similar across several breeds [2]. At the same time, dogs of the same breed can look very different due to variations in age, stance, body size, coat length, grooming, lighting, camera angle, background and image quality [3]. These characteristics make dog breed recognition a fine-grained image classification problem, where the model needs to learn highly discriminative visual patterns rather than relying only on broad object-level features. Moreover, real-world environments often present additional challenges such as low resolution, blurriness, partial occlusion or poor illumination. Therefore, developing a precise and computationally efficient dog breed classification system is a useful benchmark for evaluating modern deep learning and computer vision techniques. Recent progress in machine learning and deep learning has greatly enhanced computer systems' ability to automatically learn useful representations from visual information and classify images into predefined categories [4]. Convolutional neural networks have been very successful at image recognition because they learn hierarchical representations from raw image pixels. Early layers detect simple visual information, such as edges, lines, colors, and textures. In contrast, deeper layers detect more complex structures, such as shapes, facial features, body patterns, and object-specific characteristics [5]. These hierarchical features help with dog breed identification because distinguishing between breeds usually requires a combination of local and global features.

Two breeds may have similar coat colors but differ in muzzle structure, ear carriage, facial proportions, fur texture or body shape. To make a final breed prediction, a successful classification model needs to extract and combine different kinds of visual information [6]. In this work, researchers present a large-scale multi-class image classification framework for automatically identifying 120 dog breeds using deep learning, transfer learning and TensorFlow Hub. Instead of training a deep neural network from scratch, the proposed approach uses pre-trained architectures such as MobileNetV2 and EfficientNet. These models have already learned useful visual representations from large-scale image datasets and can therefore be adapted to dog breed classification [7]. Transfer learning is especially appropriate for this task because collecting and labeling large, breed-specific image datasets is expensive and time-consuming. By reusing pre-trained features, the model can focus on learning only the breed-specific characteristics needed for fine-grained classification and significantly reduce training time, computational requirements, and the amount of task-specific training data [8]. The primary data source used for building and evaluating the model is the publicly accessible Kaggle Dog Breed Identification dataset. The dataset includes about 20,000 images in total, including over 10,000 labeled training images and roughly the same number of unlabeled test images. The labeled set includes 120 dog breeds, making it much harder than typical binary or small multi-class classification problems. The dataset contains dogs in different conditions, such as different backgrounds, poses, scales, viewpoints, lighting environments and image resolutions [9]. This diversity matters because a model used in practice should not simply memorize a limited set of training images. Instead, it should learn features that can generalise to unseen images of the same breeds. A major difficulty of the dataset is the low inter-class variation with high intra-class variation. High inter-class similarity means two different breeds can look very similar [10].

Closely related breeds may differ only in minor details involving facial structure, fur patterns, ear shape, body proportions or other localized characteristics [11]. High intra-class variation means dogs of the same breed can look very different from each other because of pose, age, grooming style, lighting, background or natural biological diversity. Thus, the classifier must learn sufficiently robust representations that identify common features within the same breed while remaining sensitive to small differences between breeds [12]. This makes dog breed identification a good application for studying the performance of transfer learning and advanced feature extraction techniques. Before training, researchers perform pre-processing operations to improve the consistency and quality of the input data. Researchers resize the images to dimensions compatible with the selected pre-trained architecture for efficient batch processing [13]. Researchers normalize pixel values according to the model used for each task to match the data distribution used during pre-training. Normalization stabilizes optimization and removes unwanted variance from differences in raw pixel intensity ranges. The breed labels are then converted into a machine-readable format for multi-class classification, such as one-hot encoded vectors, where each output position corresponds to one of the 120 possible dog breeds. The training pipeline also uses data augmentation to improve generalization and reduce overfitting [14]. Augmentation techniques can involve horizontal flipping, small rotations, zooming, cropping, shifting and other real-life geometric transformations. These operations produce slightly modified versions of the available training images and encourage the model to learn relevant breed-specific features rather than memorizing exact positions, backgrounds or orientations. For example, a dog photographed from a slightly different angle should still be recognized as the same breed [15].

Augmentation can effectively increase the diversity of training samples without collecting additional labeled data [1]. However, apply transformations carefully, since overly aggressive augmentation may distort important anatomical features and degrade classification performance. At its heart, the proposed framework relies on transfer learning using architectures such as MobileNetV2 and EfficientNet. MobileNetV2 is particularly attractive for its good trade-off between prediction accuracy and computational cost [2]. Its design targets situations where memory, processing power and power consumption are at a premium. The MobileNetV2 architecture uses depthwise separable convolutions and inverted residuals to reduce computations and

parameters compared to classical convolutional networks [3]. Therefore, it is well suited for real-time prediction applications or deployment on smartphones, embedded systems and edge devices. The proposed framework primarily employs the pre-trained MobileNetV2 network as a feature extractor [4]. The original classification layer is removed/replaced, and new classification layers are added to predict the 120 dog breed classes. Another option for extracting more detailed image representations is to explore EfficientNet. The EfficientNet family is designed to systematically balance the network depth, width and input resolution to improve model performance. This allows EfficientNet models to achieve high image classification accuracy while using parameters efficiently [5]. Comparing MobileNetV2 and EfficientNet helps clarify the trade-off between accuracy and computational cost. A more powerful model may perform better on visually similar breeds, while a lightweight model may be more suitable when inference speed, memory usage and deployment efficiency are important requirements. TensorFlow Hub offers an easy path to use pre-trained models within the classification pipeline.

You can initialize the feature-extraction layers with weights learned in large-scale image-identification training and attach them to task-specific classification layers [6]. In the first phase of transfer learning, you can freeze most of the pre-trained backbone so the learned weights don't change while you train the new layers on the dog breed dataset. This lets the classifier learn the mapping from general-purpose image features to the 120 breed classes without changing the pre-trained representation right away. Then researchers can unfreeze some deeper layers and fine-tune them with a smaller learning rate [7]. Fine-tuning allows the network to specialize its high-level visual features to capture differences between dog breeds. The proposed classification model produces a probability distribution over the 120 output classes. The last layer usually uses the softmax activation function so the outputs can be interpreted as class probabilities [8]. Researchers use categorical cross-entropy as a loss function during training because it is appropriate for mutually exclusive multi-class classification problems. The loss quantifies the difference between the predicted probability distribution and the true breed label, and provides the optimisation signal needed to update the trainable model parameters [9]. Optimisation algorithms (e.g., Adam) can iteratively update the model weights via backpropagation. Choosing the right learning rate, batch size, number of epochs, dropout rate, and fine-tuning strategy is critical for stable convergence while minimizing overfitting. Split the labelled data into training and validation subsets so you can evaluate model performance on images not used for parameter optimization [10]. During training, you can also monitor both training loss and validation performance to see how the model converges and whether it is starting to overfit.

If needed, you can add methods such as early stopping and learning-rate adjustment. Finally, use the trained model on unseen test images to estimate its ability to generalize beyond the training data [11]. The performance is mainly measured by Top-1 and Top-5 accuracy. Top-1 accuracy is the percentage of times the predicted class with the highest probability matches the true breed. This is a strict evaluation criterion and is especially difficult when many dog breeds have similar visual features [12]. Top-5 accuracy considers a prediction correct if the actual breed is among the top 5 classes with the highest predicted probabilities. Top-5 evaluation is useful for fine-grained classification because the model may assign similar probabilities to several closely related breeds [13]. Therefore, the Top-1 and Top-5 performance difference can provide further insight into whether classification errors occur between visually similar categories rather than between completely unrelated breeds. An important part of the study is also the analysis of misclassification. For example, a model might confuse breeds that look alike, such as breeds with similar coat textures, head shapes, ear shapes, or body proportions. Examining these errors can reveal limitations in the learned feature representations [14]. Confusion patterns can hint whether more training samples, stronger augmentation, higher-resolution inputs, attention-based mechanisms, or more advanced architectures are needed. Error analysis can also indicate the effect of image background. If certain breeds tend to occur in similar contexts, for example, the model may inadvertently pick up contextual cues, rather than focusing purely on the animal. Thus, better dog localization in the image, or the inclusion of object detection and segmentation, could further improve performance [15]. Another challenge to consider is class imbalance.

Even when all 120 breeds are represented, the number and diversity of available images can vary across categories. A model trained without addressing imbalance can become biased toward classes represented by more examples or less diverse images. This problem can be mitigated by techniques such as class weighting, balanced sampling, target augmentation and additional data collection [1]. Therefore, the evaluation must measure not only overall accuracy but also breed-level performance to assess whether the classifier is consistent across classes. Automatic dog breed identification has a practical significance beyond a benchmark image classification problem. A good model could help veterinary clinics provide breed information early when an animal's breed is unknown, or records on mixed breeds are incomplete [2]. Automated classification could help animal shelters and rescue organisations better organise animal records and support adoption listings. Pet identification apps could offer owners breed predictions from photos, and animal welfare organizations could use similar technologies for database management and large-scale visual analysis [3]. Breed recognition can also support law enforcement, research and animal monitoring under certain circumstances. These systems should be treated as decision-support tools rather than replacements for professional veterinary or biological assessment, especially when breed determination has medical, legal or welfare implications [4]. The real-world implementation requires computational efficiency. A very accurate model might require expensive hardware, which is not suitable for mobile or edge-based applications. Therefore, MobileNetV2 is a strong candidate for practical deployment because its lightweight structure reduces memory consumption and inference time.

TensorFlow models can be transformed into an optimized form and executed on mobile devices, embedded platforms and other resource-limited devices [5]. This enables breed recognition to be performed closer to the user, rather than relying entirely on cloud-based computation. Edge deployment can also minimize network dependency and improve response time, making the system better suited to real-time applications [6]. Overall, the proposed work shows the applicability of deep learning and transfer learning for large-scale fine-grained dog breed classification. By combining structured preprocessing, data augmentation, pre-trained feature extraction, task-specific classification layers, and proper optimization techniques, the system can learn meaningful distinctions among 120 visually related dog breeds [7]. MobileNetV2 provides a lightweight solution, while EfficientNet offers an opportunity to test whether greater representational capacity improves recognition performance. TensorFlow Hub lets you easily access powerful pre-trained models and makes the workflow more reproducible and scalable [8]. The findings highlight both the potential and the remaining challenges of automatic breed recognition, particularly those associated with visually similar breeds, class imbalance, variable image quality, and intra-class diversity. Future enhancements could include more advanced augmentation techniques, attention mechanisms, ensemble learning, fine-tuning strategies, explainable artificial intelligence techniques, object localization, and larger datasets with more diverse real-world images [9]. The overall framework provides a practical foundation for developing efficient image recognition systems. It shows how transfer learning can be applied to complex fine-grained classification problems while reducing training time, computational cost, and dependence on extremely large, domain-specific labeled datasets [10].

2. Literature Review

Machine learning and deep learning techniques have made remarkable progress in computer vision and enabled reliable classification of visual information into predefined categories. Dog breed classification is one of the most challenging fine-grained classification problems among different image identification challenges. This is because different dog breeds usually have only subtle differences in their physical characteristics, including body structure, coat patterns, facial features, and size. At the same time, images of the same breed can vary widely because of differences in posture, age, lighting conditions, background, camera angle and image quality. Thus, accurate dog breed recognition has become a key benchmark for the performance of modern image recognition systems and has practical uses in pet identification, veterinary diagnostics, animal management and automated monitoring systems [2]. Recent work shows that convolutional neural networks (CNNs) and transfer learning techniques are highly effective for solving difficult image classification problems. CNN-based models can automatically extract hierarchical visual features from raw images without manual feature engineering. In particular, transfer learning lets researchers leverage pre-trained deep learning models that have learned general visual representations from large-scale datasets and adapt them to specific applications such as dog breed identification. This approach has lower computational cost, shorter training time and better classification performance than traditional machine learning methods [1]; [7]. In this paper, researchers propose a robust multi-class classification model to identify 120 dog breeds from images using deep learning-based transfer learning techniques. The problem is challenging because of low inter-class variance, where visually similar breeds share similar features, and high intra-class variance, where images of the same breed can differ greatly due to environmental and biological variations.

These characteristics make dog breed recognition an ideal benchmark for assessing advanced feature extraction approaches and state-of-the-art deep learning architectures. Prior studies have explored various CNN-based methods, hybrid deep learning models, and feature fusion techniques for enhancing the accuracy and robustness of dog breed classification systems [3]; [4]. This study uses the Kaggle Dog Breed Identification competition dataset. Researchers use more than 10,000 labeled images for training and an equal number of unlabeled images for testing [10]. The proposed workflow extracts relevant image features, builds a multi-class classification model using labeled examples, and evaluates the model's performance on unseen test images. The proposed method uses previously learned image representations and adapts them through transfer learning to identify breed-specific characteristics. This provides an efficient and scalable approach toward solving fine-grained real-world image classification problems in situations where large quantities of labeled data may not always be available [2]; [8]. The proposed architecture shows that deep learning methods are well suited to complex, unstructured image data and emphasizes the importance of high classification accuracy for practical use. Correctly identifying dog breeds from photographs can support many areas, including veterinary clinics, pet adoption websites, animal advocacy groups, biological research and law enforcement applications. Similarly, computer vision studies have shown that state-of-the-art detection and classification frameworks can extract key information from complex visual environments [5]; [9]. This research presents a novel method to automatically classify dog breeds using deep learning and computer vision techniques, eliminating the need for human expert intervention.

The goal is to improve classification accuracy while reducing training time and computational requirements compared with more complex, resource-intensive approaches. Researchers build an efficient classification system using transfer learning with MobileNetV2. Researchers chose MobileNetV2 for its strong feature-extraction ability, lightweight architecture and suitability for real-time applications and deployment on low-resource devices. It uses depthwise separable convolution operations that greatly reduce parameters and computational complexity while maintaining effective image representation [1]; [7]. Transfer

learning improves model efficiency by reusing knowledge learned from previously trained neural networks. This way, the model doesn't have to learn all visual features from scratch; it retains general image recognition ability and fine-tunes high-level features to identify dog breeds. This method enables the model to perform better with fewer training samples and lower computational requirements. Several studies have shown the effectiveness of transfer learning-based CNN models for animal image classification and other fine-grained recognition tasks [8]; [3]. The dataset used in this work comprises about 20,000 images of 120 dog breeds. This provides a diverse set of images for training and evaluation. The images come from different sources and were manually enhanced to improve the classification model's robustness and generalization.

Different preprocessing techniques, such as image resizing, normalization and augmentation, enhance image quality and diversity during training. These techniques help the model learn breed-specific features while decreasing the likelihood of overfitting due to limited training examples [10]; [4]. Combining MobileNetV2 with transfer learning provides a scalable, efficient solution for automatic dog breed classification. Compared with deeper, more computationally expensive networks, MobileNetV2 offers a strong trade-off between accuracy and efficiency for real-world deployment scenarios such as mobile applications, embedded devices, and real-time identification systems. Comparative studies of deep learning algorithms show that optimized CNN architectures can achieve competitive classification performance with lower computational requirements [3]. The experimental results show that the proposed approach is robust and efficient and can replace more complex models that require significant computational resources. The system provides useful insights for solving fine-grained classification problems with visually similar categories and variable image conditions. Future improvements could include incorporating attention mechanisms, ensemble learning approaches, advanced object detection methods (e.g., YOLO-based frameworks), and larger datasets to improve classification accuracy and generalization performance further [5]; [6]. Overall, this work demonstrates the effectiveness of deep learning-based transfer learning for large-scale dog breed classification. The proposed MobileNetV2-based framework offers an efficient, accurate and practical solution for dog breed classification from images by reducing computational complexity and training time. In the subsequent sections, researchers discuss related work, methodology, experimental evaluation and directions for future research in detail.

3. Methodology

3.1. Dog Breed Classification Workflow

3.1.1. Data Collection

Researchers used the Dog Breed Identification Dataset from Kaggle. It is a good choice for a supervised learning task because it contains more than 20, 000 dog photos of 120 breeds (Figure 1).




Figure 1: Dataset breeds

The dataset has 10,222 labeled photos for training and 10,357 unlabeled images for testing. All images are in .jpg format, and a CSV file contains metadata such as breed labels and image IDs. Researchers downloaded the dataset from Kaggle (either

manually or using the platform's API). Researchers checked that all photos matched the CSV entries and that there were no duplicates, corrupted images, or missing files. After sorting training photos by breed labels, researchers split the dataset into training, validation, and testing subsets:

- **Data Preprocessing:** The data pretreatment steps ensured the raw image data was properly prepared for model training. Key steps included scaling images, partitioning the dataset, normalising pixel values, batching data, and converting images to tensors.
- **Tensor Conversion:** Tensors were invented for images, and tensors are multidimensional arrays. They're the key to getting picture data into a form that machine learning models can easily understand and process.
- **Data Splitting:** Researchers create three dataset subsets: the testing dataset, the validation dataset, and the training dataset. Researchers can thoroughly examine the model's performance by training it on one dataset and evaluating it on another. By clearly separating training and evaluation, researchers can more accurately assess the model's effectiveness.
- **Pixel Value Normalisation:** Researchers normalized the photos' pixel values to the ranges 0 to 255 and 0 to 1. By scaling all pixel values to the same range, this normalisation step improves model convergence and helps train better models.
- **Resizing:** To keep things consistent, researchers resized all images to 224x224 pixels. Since CNNs are at the heart of most deep learning models, it makes sense to standardize image dimensions for consistent processing.
- **Batching Data:** Eventually, the data was split into 32 batches. This makes training faster and more effective, especially when using smaller batch sizes. Batching allows the model to process data in small batches, making training more efficient and effective (Figure 2).



Figure 2: Batch images

3.1.2. Input and Output Shapes

The model input shape was [None, IMG_SIZE, IMG_SIZE, 3], where IMG_SIZE is 224 pixels and specifies the scaled image size. The 3 in the shape indicates the number of RGB color channels, allowing the model to handle color images. The number of dog breeds in the dataset defines the output shape; each breed is a unique output class.

4. Results and Discussions

4.1. Model Architecture

Researchers used the pre-trained MobileNetV2 model from TensorFlow Hub as our baseline because it performs well on image inference. MobileNetV2 is a lightweight deep neural network designed for fast image classification. So, it is suitable for dog breed classification. The model has the following layers.

4.2. Feature Extraction Layer

The Feature Extraction Layer is the first layer of the model, a pre-trained model that researchers downloaded from TensorFlow Hub. It recognises general patterns in images, such as edges, forms, and textures. Because MobileNetV2's weights are fixed during training, it acts as a feature extractor rather than learning new features.

4.3. Classification Layer

The second dense layer has several units equal to the total number of dog breeds, forming the final output. A fully connected layer with a softmax activation function then converts this output into a probability distribution over all breed categories.

4.4. Model Compilation

The model architecture was defined and compiled in the following way:

- **Loss Function:** For multi-class classification problems, where each image is assigned to a single class among multiple distinct categories, researchers selected categorical cross-entropy as the loss function. This choice is ideal when the objective is to estimate the probability distribution across multiple classes, ensuring each image is classified into exactly one category.
- **Optimiser:** Researchers chose the Adam optimiser because it performs well for training deep learning models, especially by adapting the learning rate during training.
- **Evaluation Metrics:** Researchers chose accuracy as the evaluation metric because it is the most straightforward and reliable measure of performance for this task. Accuracy is the ratio of correctly predicted cases to the total number of predictions made.
- **Callbacks:** By providing tools to monitor and control the training process under different conditions, callbacks help improve model training. In this study, researchers used TensorBoard and Early Stopping as the two main callbacks.

4.5. TensorBoard Callback

Researchers monitored the model's performance during training using the TensorBoard callback. TensorBoard provides real-time visualisation of accuracy, training loss, validation loss, and other metrics. Researchers recorded the accuracy, training loss, validation loss, and other metrics during training to monitor the model's performance. This gave us insights into how the model learned and generalised on the validation data. For example, TensorBoard visualisations could identify slow convergence, underfitting, or overfitting.

4.6. Early Stopping Callback

To limit training time and prevent overfitting, researchers use the Early Stopping callback. This approach monitors a specific metric, such as validation loss, to determine if it remains unchanged for a set number of consecutive epochs (referred to as 'patience'). If no improvement occurs during this period, training halts. By stopping training early, researchers prevent the model from learning further once it has reached its maximum performance on the training data. In addition, because this callback stops training after convergence, it saves computational resources.

4.7. Model Building

Researchers used TensorFlow and TensorFlow Hub to train the model in the Dog Breed Classification task. Researchers applied transfer learning using a pre-trained model for feature extraction. A detailed description of the model-building process appears below.

4.8. Training the Model

Researchers trained the model for 40 epochs on the training and validation datasets. Each epoch updated the weights based on training-data feedback while evaluating generalisation on the validation set. Researchers implemented Early Stopping to monitor the validation loss and halt training once no further improvement was observed. Researchers also used TensorBoard to track metrics such as loss and accuracy. TensorBoard provides real-time visualisations to help optimise training and prevent

overfitting. Model Training: Researchers trained the model for 40 epochs on the training and validation datasets. Each epoch updated the model's weights on the training data while evaluating generalization on the validation set. Researchers implemented Early Stopping to monitor the validation loss and halt training once no further improvement was observed (Figure 3).




Figure 3: Flow diagram

Researchers also used TensorBoard to track metrics such as loss and accuracy. TensorBoard provides real-time visualisations to help optimise training and prevent overfitting.

4.9. Model Evaluation

After 40 epochs of training, researchers test the model on the validation dataset. The model made predictions fairly close to the true labels, as shown by the validation loss of 0.0102, which is relatively low (Figure 4).




Figure 4: TensorBoard loss

Also, it correctly classified 99.89% of the validation images, as shown by the validation accuracy of 0.9989. By these measures, the model achieves high accuracy and low error, making it suitable for dog breed classification. The evaluation confirms the model's strength by showing it generalises well to unseen data (Figure 5).




Figure 5: TensorBoard accuracy

After extensive training, the model's validation loss is 0.0102 and validation accuracy is 99.89%. The model shows very little discrimination error between dog breeds. The high accuracy suggests it generalises well to new data, making it reliable for practical applications. The model predicts well, and the reduced validation loss supports the precision. Overall, the model does well, suggesting it can be used in automated dog breed classification systems (Figure 6).




Figure 6: Model prediction probabilities

The results showed that after 40 epochs of training, the model performed very well in dog breed classification. With a validation loss of 0.0102 and a validation accuracy of 99.89%, the model trained to classify 120 dog breeds with low error. This result strongly implies that the model learned the relevant features for breed categorisation. While the model performs well, it can still be improved. Exploring additional data augmentation methods can increase training-sample diversity and improve generalization. Investigating alternative model architectures or training approaches could further optimise performance. In addition, a more comprehensive and diverse dataset could help the model recognize rare dog breeds and handle complex scenarios more effectively. While the model performs well, it could still be improved. Applying a wider range of data augmentation techniques can enhance the diversity of the training set, potentially improving generalisation. Trying different model architectures or training approaches may also help optimise performance. Lastly, a larger, more diverse dataset will help the model recognise rare dog breeds and better handle challenging cases. Overall, the model's performance is encouraging for real-world dog breed classification tasks and could be used for automatic breed identification in applications such as mobile apps and animal shelter systems.

5. Conclusion and Future Work

Researchers developed a deep learning model for dog breed classification using TensorFlow Hub's MobileNetV2 architecture. After 40 epochs of training on a dataset of more than 20,000 images, the model achieved a validation accuracy of 99.89% and a validation loss of 0.0102. These results show that the model classifies dog breeds effectively and generalises well to new data. Using TensorFlow's rich toolkit, TensorBoard for real-time monitoring, and Early Stopping for regularisation helped avoid overfitting. Performance is still strong, although it could improve with more varied datasets and more experimentation with different deep learning architectures. This paper presents a transfer learning approach to image classification and applies it to automatic dog breed identification. With further improvements, the proposed model can be applied in veterinary clinics, animal shelters and mobile applications. Model improvements could include training on larger, more diverse datasets to make it more robust across breeds and different image conditions (i.e., different lighting, background and quality variations). Exploring more advanced deep learning architectures, such as ViTs and EfficientNet variants, may also improve classification accuracy. Lastly, applying explainable AI techniques could increase transparency by providing insights into how the model makes decisions and making it more interpretable. Another interesting direction for future work is deploying the model on edge devices or mobile applications for real-time breed classification. Optimising the model for low-latency, resource-constrained environments will further improve its applicability in real-world settings. Last but not least, multimodal learning that combines textual metadata, such as breed descriptions, with image classification may further improve prediction accuracy and provide additional contextual information. By improving these aspects, the proposed model can offer a more comprehensive and practical approach to automatic dog breed classification.

Acknowledgement: The authors gratefully acknowledge the support and contributions of the Ukrainian National Forestry University, Zewail City of Science and Technology, Taipei Tech Industries, Placental Histoherapy Center, and Asia Pacific University of Technology and Innovation. Their valuable support and cooperation contributed to the successful completion of this research.

Data Availability Statement: The research contains data related to an efficient MobileNet-based deep learning framework for automated dog breed classification. The dataset and associated information used in this study are available to the authors and can be provided upon reasonable request.

Funding Statement: No external funding was received by the authors for conducting this research.

Conflicts of Interest Statement: The authors declare that there are no conflicts of interest regarding this study.

Ethics and Consent Statement: Ethical approval was obtained, and informed consent was received from the organization and all individual participants involved in the data collection.

References

1. P. Borwonginn, K. Thongkanchorn, S. Kanchanapreechakorn, and W. Kusakunniran, "Breakthrough Conventional Based Approach for Dog Breed Classification Using CNN with Transfer Learning," *11th International Conference on Information Technology and Electrical Engineering (ICITEE)*, Pattaya, Thailand, 2019.
2. A. Varshney, A. Katiyar, A. K. Singh, and S. S. Chauhan, "Dog Breed Classification Using Deep Learning," *2021 International Conference on Intelligent Technologies (CONIT)*, Hubli, India, 2021.

3. B. Valarmathi, N. S. Gupta, G. Prakash, R. H. Reddy, S. Saravanan, and P. Shanmugasundaram, "Hybrid Deep Learning Algorithms for Dog Breed Identification—A Comparative Analysis," in *IEEE Access*, vol. 11, no. 7, pp. 77228–77239, 2023.
4. C. Wang, J. Wang, Q. Du, and X. Yang, "Dog Breed Classification Based on Deep Learning," *2020 13th International Symposium on Computational Intelligence and Design (ISCID)*, Hangzhou, China, 2020.
5. J. Du, "Understanding of object detection based on CNN family and YOLO," *Journal of Physics: Conference Series*, vol. 1004, no. 1, pp. 1–9, 2018.
6. X. Hu, S. Zhu, and T. Peng, "Hierarchical attention vision transformer for fine-grained visual classification," *Journal of Visual Communication and Image Representation*, vol. 91, no. 3, p. 103755, 2023.
7. N. Pannurat, K. Eiamsaard, C. Suthanma, and A. Banharnsakun, "Machine learning techniques for supporting dog grooming services," *Results in Control and Optimization*, vol. 12, no. 9, pp. 1–9, 2023.
8. N. Towpunwong and N. Sae-Bae, "Dog Breed Classification and Identification Using Convolutional Neural Networks," *ECTI Transactions on Computer and Information Technology (ECTI-CIT)*, vol. 17, no. 4, pp. 554–563, 2023.
9. C. Cao, B. Wang, W. Zhang, X. Zeng, X. Yan, Z. Feng, Y. Liu, and Z. Wu, "An improved faster R-CNN for small object detection," in *IEEE Access*, vol. 7, no. 8, pp. 106838–106846, 2019.
10. M. Uno, X. H. Han, and Y. W. Chen, "Comprehensive Study of Multiple CNNs Fusion for Fine-Grained Dog Breed Categorization," *2018 IEEE International Symposium on Multimedia (ISM)*, Taichung, Taiwan, 2018.
11. O. O. Martins, C. C. Oosthuizen, and D. A. Desai, "Evaluating unified training optimisations for MobileNetV2: Efficiency-accuracy trade-offs in fine-grained dog breed classification," *Discover Applied Sciences*, vol. 7, no. 10, pp. 1–45, 2025.
12. P. Borwarnginn, W. Kusakunniran, S. Karnjanapreechakorn, and K. Thongkanchorn, "Knowing your dog breed: Identifying a dog breed with deep learning," *Int. J. Autom. Comput.*, vol. 18, no. 11, pp. 45–54, 2021.
13. D. N. Zou, S. H. Zhang, T. J. Mu, and M. Zhang, "A new dataset of dog breed images and a benchmark for fine-grained classification," *Computational Visual Media*, vol. 6, no. 4, pp. 477–487, 2020.
14. R. O. Sinnott, F. Wu, and W. Chen, "A mobile application for dog breed detection and recognition based on deep learning," in *Proc. 5th IEEE/ACM Int. Conf. Big Data Computing, Applications and Technologies (BDCAT)*, Zurich, Switzerland, 2018.
15. Z. C. Zhang, Z. D. Chen, Y. Wang, X. Luo, and X. S. Xu, "A vision transformer for fine-grained classification by reducing noise and enhancing discriminative information," *Pattern Recognition*, vol. 145, no. 1, p. 109979, 2024.

Publisher's Note: The publisher remains impartial concerning jurisdictional claims in published maps and institutional affiliations. Responsibility for the content rests entirely with the authors and does not necessarily reflect the publisher's perspectives.