PAPER ID: 20260202016

Channel-wise Analysis of Ultra-Low Resolution Images for Resource-Constrained Image Classification

Shikha Badhani and Manju Bhardwaj
Department of Computer Science, Maitreyi College, University of Delhi, Delhi, India

Abstract: Image classification on resource-constrained systems needs compact image representations that balance prediction metrics with computing efficiency. This study offers a detailed evaluation of RGB, Red, Green, Blue, and Grayscale image representations using the Agricultural Multi-Class Dataset (AMCD). Images were resized to 16 × 16 and 32 × 32 pixels and tested with Random Forest, Linear Support Vector Machine, Logistic Regression, and Extreme Gradient Boosting. Additionally, three channel-wise ensemble strategies—Majority Voting, Weighted Majority Voting, and Stacking—were explored to see the benefits of combining individual color channels. The experimental results indicate that increasing the image resolution from 16 × 16 to 32 × 32 consistently boosts classification performance. Linear SVM achieved the best accuracy of 63.05% with RGB images, while Random Forest offered the best balance between classification performance and computing efficiency. Among the individual color channels, the Blue channel emerged as the strongest lightweight alternative, reducing the feature dimensionality by nearly threefold while maintaining competitive classification performance. Although the ensemble methods improved some single-channel representations, they did not surpass the RGB representation. This suggests that preserving full color information is more useful than combining decisions. The findings give practical guidelines for choosing image representations and classifiers for effective image classification in IoT, embedded, and edge computing settings.

Keywords: Channel analysis, IoT, Edge computing, Ultra-low-resolution image classification, Channel-wise ensemble, Machine learning.

INTRODUCTION:

The growing adoption of the Internet of Things (IoT), edge computing, and embedded vision systems has enabled image classification to be deployed directly on resource-constrained devices, reducing the reliance on cloud infrastructure and enabling real-time decision-making. However, these devices usually have limited processing power, memory, storage, and energy, making computational efficiency an important consideration when designing image classification algorithms (Shi et al., 2016; Satyanarayanan, 2017). In such environments, reducing the amount of image data without significantly affecting recognition performance can improve both processing speed and resource utilization.

One straightforward way to reduce computational requirements is to decrease image resolution and simplify color representation (Torralba, A., 2009). Although high-resolution RGB images contain rich visual information, they also increase feature dimensionality and computational cost. In contrast, ultra-low-resolution images and single-channel representations (Red, Green, Blue, or Grayscale) require considerably fewer features, making them attractive for deployment on resource-constrained and edge devices. However, the extent to which these compact representations preserve the information required for accurate classification has not been thoroughly investigated.

To explore this trade-off, this study uses the recently introduced Agricultural Multi-Class Dataset (AMCD), which consists of 5,405 images of size 512 × 512 pixels belonging to 77 classes of agricultural crops and flowers captured under diverse real-world conditions (Karim et al., 2026). Images are resized to 16 × 16 and 32 × 32 pixels and represented using RGB, individual color channels, and Grayscale images. Four machine learning classifiers are evaluated using stratified cross-validation to identify the most suitable classifier for ultra-low-resolution image classification. The best-performing classifier is subsequently used to train independent models on the Red, Green, and Blue channels, which are then combined using different ensemble strategies to examine whether channel-wise fusion improves classification performance. The proposed framework is evaluated in terms of classification accuracy, feature dimensionality, training time, and inference time, providing practical insights into the impact of image resolution, color representation, and classifier selection for efficient image classification on IoT, embedded, and edge computing platforms.

LITERATURE REVIEW

Image classification has become a fundamental task in computer vision because of its wide range of applications, including healthcare, agriculture, surveillance, industrial inspection, and autonomous systems. While deep learning has substantially improved classification accuracy, its computational requirements often make deployment difficult on resource-constrained platforms, such as IoT devices, embedded systems, and edge devices. These platforms typically operate with limited processing capability, memory, and energy, making computational efficiency an equally important consideration alongside predictive performance. As a result, lightweight machine learning techniques continue to play an important role in applications where rapid inference and efficient resource utilization are essential (Shi et al., 2016; Satyanarayanan, 2017).

Besides the learning algorithm itself, the way an image is represented has a significant influence on both classification performance and computational cost. Most image classification systems process images in the RGB color space by default, even though the three channels often contain redundant information. This has encouraged researchers to investigate alternative color representations that either improve discriminative capability or reduce computational complexity. Gowda and Yuan (2019), for instance, showed that combining information from different color spaces can improve image classification while maintaining a relatively compact model. Their findings suggest that color representation should be considered an integral part of the classification pipeline rather than merely a preprocessing choice. Interest has also grown in understanding the contribution of individual color channels. Instead of treating RGB as a single entity, recent studies have shown that each channel captures different visual characteristics. Yoo (2023) demonstrated that quantitative information extracted from individual Red, Green, and Blue channels can be effectively utilized in machine vision applications. Likewise, Gelir et al. (2025) reported that features extracted separately from RGB and HSV channels improved retinal disease classification, indicating that color-specific information can enhance visual recognition. Although these studies highlight the importance of color information, their primary objective is feature extraction for specific application domains rather than evaluating the effectiveness of individual color channels as compact image representations.

Reducing image resolution offers another practical way to decrease computational requirements. Smaller images contain fewer pixels, resulting in lower memory consumption, reduced storage requirements, and faster processing. These advantages are particularly valuable for IoT and edge computing applications, where hardware resources are inherently limited. Despite these benefits, most existing studies continue to rely on relatively high-resolution images or deep feature representations (Khayya et al., 2024; Gowda & Yuan, 2019). Comparatively little attention has been given to understanding how aggressively image resolution can be reduced before classification performance begins to deteriorate (Khayya et al., 2024).

Similarly, evaluations involving lightweight machine learning algorithms remain less common than deep learning approaches (Khayya et al., 2024) despite their suitability for deployment on resource-constrained platforms. Consequently, there is still limited evidence regarding how different color representations perform after substantial spatial downsampling or whether processing full RGB images provides sufficient benefits to justify the additional computational cost.

Although previous studies have demonstrated the importance of image resolution and color representation for image classification, their combined influence has not been systematically investigated using lightweight machine learning approaches. To address this gap, the present study uses the Agricultural Multi-Class Dataset (AMCD) (Karim et al., 2026) to compare RGB, individual Red, Green, Blue, and Grayscale representations at ultra-low resolutions using four widely adopted machine learning classifiers. In addition, a channel-wise ensemble is evaluated to determine whether combining predictions from individual color channels offers any additional benefit. By examining classification performance alongside computational efficiency, this study provides practical insights into selecting suitable image representations and classifiers for resource-constrained IoT, embedded, and edge computing applications.

METHODOLOGY

This study investigates the influence of image resolution and color representation on image classification performance under resource-constrained settings. RGB, individual Red, Green, Blue, and Grayscale representations are evaluated at ultra-low resolutions using lightweight machine learning classifiers to examine the trade-off between classification performance and computational efficiency. Based on the comparative analysis, the best-performing classifier is further used to investigate whether combining independently trained RGB channel models through ensemble learning improves classification performance. The overall workflow of the proposed framework is illustrated in Figure 1.

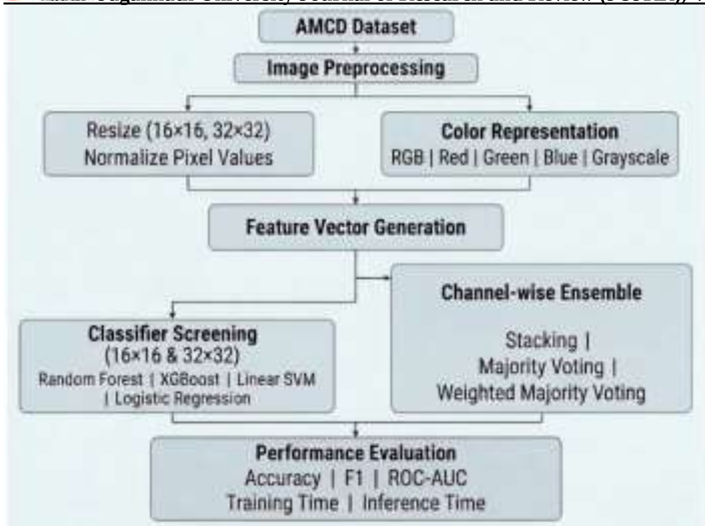

Fig. 1 – Experimental Workflow

3.1. Dataset

Experiments were performed using the Agricultural Multi-Class Dataset (AMCD) (Karim et al., 2026), which contains 5,405 JPEG images belonging to 77 subclasses distributed across four categories: flowers (1,765 images; 25 subclasses), vegetables (1,610 images; 23 subclasses), fruits (1,400 images; 20 subclasses), and crops/grains (630 images; 9 subclasses). The images were captured under diverse environmental conditions. The variations in illumination, viewpoint, object orientation, and background provide a challenging benchmark for evaluating image classification algorithms. Although originally developed for agricultural image classification, the dataset is used here as a benchmark to evaluate compact image representations for resource-constrained image classification.

3.2. Image Preprocessing and Color Representation

All images were resized to 16 × 16 and 32 × 32 pixels prior to feature extraction. These resolutions were chosen because they significantly reduce the amount of image data that must be processed, making them well suited for resource-constrained environments such as IoT devices, embedded systems, and edge platforms, where memory and computational resources are limited. Each resized image was represented in five different forms: the original RGB image, the individual Red, Green, and Blue channels, and a Grayscale image. While the RGB representation retains complete color information, the individual channels and grayscale image provide more compact representations, allowing the effect of color information on classification performance and computational efficiency to be evaluated. Pixel intensities were normalized to the range [0, 1] before feature extraction to ensure a consistent numerical scale across all experiments.

3.3. Feature Representation

Each processed image was transformed into a one-dimensional feature vector (d) by flattening its pixel values as given in Eq. (1). Consequently, the feature dimensionality depends on both the image resolution and the selected color representation.

For an image of dimensions $M \times N$,

$$d = \begin{cases} 3MN, & RGB\ Image \\ MN, & R, G, B, or\ Grayscale\ Image \end{cases} \quad (1)$$

Thus, RGB images contain three times more features than the corresponding single-channel representations, resulting in higher memory requirements and computational cost.

3.4. Classification Models

Four supervised machine learning classifiers were evaluated in this study, namely Random Forest [RF] (Breiman, 2001), Linear Support Vector Machine [LSVM] (Cortes & Vapnik, 1995), Logistic Regression [LR] (Cox, 1958), and Extreme Gradient Boosting [XGBoost] (Chen & Guestrin, 2016). These algorithms were selected because they represent different learning paradigms while remaining computationally efficient compared with deep learning models. Random Forest and XGBoost are ensemble learning methods capable of modelling complex decision boundaries, whereas Linear SVM and Logistic Regression provide efficient linear classification models that perform well on high-dimensional feature spaces. To study the effect of complementary information across color channels, the classifier demonstrating the best overall performance was used to develop a channel-wise ensemble. Separate classifiers were trained using the Red, Green, and Blue image representations, and three ensemble strategies—Majority Voting, Weighted Majority Voting, and Stacking—were investigated to examine whether combining channel-specific predictions improves classification performance. Majority Voting determines the final class based on the prediction receiving the highest number of votes from the three independently trained channel-specific classifiers. Weighted Majority Voting extends this approach by assigning different weights to the predictions of the individual classifiers before determining the final class (Kuncheva, 2004). In this study, the weights were derived from the classification accuracy of the corresponding channel-specific classifiers. Stacking combines the predictions of the base classifiers through a meta-classifier that learns how to integrate their outputs (Wolpert, 1992). In this work, Linear Support Vector Machine was used as the meta-classifier.

3.5. Experimental Setup

All experiments were implemented in Python 3.x using the Scikit-learn library. The abovementioned four machine learning classifiers were evaluated using 3-fold stratified cross-validation, ensuring that the class distribution was preserved in each fold. Performance was reported as the average across the three folds to provide a reliable estimate of classifier performance. The framework was evaluated using three performance metrics- Accuracy, macro F1-score and macro ROC-AUC (Area under ROC curve). To assess the computational efficiency of algorithms, the training time and inference time was also recorded.

RESULTS AND DISCUSSION

This section presents the results of experiments performed on the different image representations. The performance of the abovementioned machine learning classifiers is first compared on RGB images to assess the algorithms on the basis of their performance and computation timings. RF and LSVM were selected as they performed better than the other two classifiers. The selected classifiers were then used for classification of the converted low-resolution RGB images, the individual Red, Green, and Blue channels, and Grayscale image for each image in the given dataset. The best-performing LSVM and RF classifiers over individual Red, Green and Blue channels were then used to generate the channel-based ensemble.

4.1. *Selection of Classifier*

The performance of the four classifiers was assessed using feature vectors obtained from RGB image format for the 16 × 16 and 32 × 32 images (See Sec. 3.2). Table 1 lists the performance metrics and the training/inference timings for each classifier on the basis of RGB image format.

Table 1. Performance comparison of classifiers using RGB images.

Image Resolution	Classifier	Accuracy	F1-score	ROC—AUC	Training Time (s)	Inference Time (s)
16×16	LR	0.54	0.54	0.90	27.7	0.0
	LSVM	0.57	**0.57**	0.92	26.3	4.6
	RF	**0.59**	**0.57**	0.91	10.6	0.3
	XGBoost	0.56	0.56	**0.95**	528.5	0.1
32×32	LR	0.57	0.57	0.94	77.6	0.0
	LSVM	**0.63**	**0.63**	**0.95**	115.0	14.1
	RF	0.61	0.60	0.92	24.5	0.2
	XGBoost	0.59	0.59	**0.95**	3746.4	0.2

*Highest values of performance metrics are highlighted in bold

It may be noted that increasing the image resolution from 16 × 16 to 32 × 32 consistently improved the performance of all classifiers. But this improvement was accompanied by a fourfold increase in feature dimensionality. For RGB images, the number of input features increased from 768 at 16 × 16 resolution to 3072 at 32 × 32, providing richer image information but also increasing the computational cost of classification.

As can be seen from Table 1, Linear SVM was the best performer with the highest classification accuracy (63%), F1-score (63%) and ROC-AUC (95%) at the higher resolution, Random Forest was the best performer in terms of accuracy (59%) and F1-score (57%) on lower resolution. Also, it delivered comparable performance on higher resolution image with substantially lower training and inference times. Although XGBoost achieved the highest ROC-AUC, its training time increased considerably with image resolution, exceeding 3,700 seconds at 32 × 32, making it less suitable for resource-constrained deployment. Logistic Regression remained computationally efficient but consistently produced the lowest classification accuracy.

4.2. *Color-channels for Image Identification*

To understand the contribution of individual color channels, each classifier was evaluated using Red, Green, Blue, and Grayscale images. Table 2 summarizes the performance metrics and training/inference timings for the best-performing classifier for each image representation.

Table 2. Best-performing classifier for each color channel.

Image Res.	Best Classifier	Color	Acc.	F1-score	ROC-AUC	Training Time (s)	Inference Time (s)
16×16	RF	RGB	**0.59**	**0.57**	**0.91**	10.6	0.3
		R	0.52	0.51	0.89	6	0.2
		G	0.52	0.51	0.88	6	0.2
		B	0.54	0.53	0.9	8	0.3
		GS	0.5	0.5	0.88	6	0.2
32×32	LSVM	RGB	**0.63**	**0.63**	**0.95**	115	14.1
		R	0.54	0.55	0.92	26	4.2
		G	0.54	0.54	0.91	31	4.7
		B	0.56	0.56	0.92	26	4.2

		GS	0.52	0.53	0.9	28	4.3

R: Red, G: Green, B: Blue, GS: Grayscale

It is clear from Table 2 that RGB consistently outperformed all single-channel representations, highlighting the importance of preserving complete color information. Among the individual channels, the Blue channel achieved the highest classification accuracy at both resolutions, followed by the Red and Green channels, while Grayscale format reported the weakest performance. Although the individual color channels resulted in lower classification accuracy as compared to RGB format, they required substantially fewer input features. At 16 × 16 resolution, RGB images generated 768 features, whereas each individual color channel contained only 256 features. Similarly, at 32 × 32 resolution, the feature dimensionality was reduced from 3072 to 1024 features in case of individual color channel feature representation. Despite this threefold reduction in feature size, the Blue channel achieved an accuracy of 56%, only about 7 percentage points lower than the RGB representation (63%), demonstrating that competitive classification performance can be achieved using considerably fewer features. Thus, *RGB image format holds a richer set of information as compared to individual Red, Green and Blue channels or the corresponding Grayscale image.*

4.3. *Channel-based Ensemble*

To investigate whether combining independently trained Red, Green, and Blue channel models could improve classification performance, three decision-level fusion strategies were evaluated, namely, Majority Voting, Weighted Majority Voting and Stacking. The results obtained on channel-based ensembles using these strategies are presented in Table 3.

Table 3. Performance comparison of channel-wise ensemble methods.

Resolution	Ensemble Method	Acc.	F1-score	ROC-AUC	Training Time (s)	Inference Time (s)
16×16	Majority Voting	0.51	0.51	0.9	45	7
	Weighted Majority Voting	0.52	0.52	0.78	31	5
	Stacking	0.52	0.53	0.93	35	5
32×32	Majority Voting	0.56	0.57	0.93	108	18
	Weighted Majority Voting	0.56	0.57	0.81	82	14
	Stacking	0.55	0.59	0.95	95	15

The three ensemble methods produced comparable classification performance, with accuracies ranging from 51% to 56%. Majority Voting and Weighted Majority Voting achieved the highest accuracy at 32 × 32, whereas Stacking produced the highest F1-score and ROC-AUC, indicating better ranking ability and class separability. However, none of the ensemble methods exceeded the performance of the RGB representation.

The limited improvement obtained through ensemble learning suggests that the discriminative information contained in the Red, Green, and Blue channels is largely overlapping. Consequently, combining independently trained channel-specific models provides only a modest benefit over the individual channels and cannot recover the richer information available in the original RGB image. Among the evaluated ensemble approaches, Weighted Majority Voting achieved a favorable balance between predictive performance and computational cost, while Stacking required additional training without a corresponding improvement in classification accuracy.

Overall, the results demonstrate that RGB remains the most informative representation for ultra-low-resolution image classification. However, reducing the feature dimensionality by nearly threefold through the use of a single color channel resulted in only a moderate reduction in classification accuracy, with the Blue channel providing the strongest performance among the compact representations. These findings suggest that significant feature reduction can be achieved while retaining much of the discriminative information required for classification, making single-channel representations attractive for deployment on resource-constrained IoT and edge devices.

It is important to note that the objective of this study differs from that of the original AMCD work, which focused on maximizing classification accuracy using deep learning models trained on high-resolution (512 × 512) images and reported substantially higher performance (Karim et al., 2026). In contrast, the present work investigates how far image resolution and feature dimensionality can be reduced while maintaining acceptable classification performance using lightweight machine learning models. For many IoT, embedded, and edge computing applications, reducing memory usage, computational cost, and inference time is often more important than achieving the highest possible classification accuracy.

CONCLUSION

This study focused on the use of ultra low resolution images for resource constrained systems such as IoT, edge devices, and embedded applications. Experiments were performed using the Agricultural Multi-Class Dataset (AMCD), the images in the dataset were converted to ultra-low resolution of 16 × 16 and 32 × 32. Each resized image was converted to feature vectors in five formats: the original RGB image, the individual Red, Green, and Blue channels, and a Grayscale image. Experimentation indicates improved classification performance on higher resolution as compared to low resolution. Performance of four machine learning algorithms- Linear Regression, Linear SVM, Random Forest and XGBoost was assessed on this dataset. Linear SVM (LSVM) achieved the highest accuracy, while Random Forest provided the best trade-off between accuracy and training timings.

Features generated from RGB images outperformed those for grayscale and single channel (Red, Green, Blue) inputs, showing the importance of preserving full color information. Information conveyed by Blue channel, although low-performer as compared to RGB, emerged as the closest lightweight alternative, offering a good compromise between reduced feature set and performance.

Ensemble methods combining the three individual Red, Green and Blue channels produced competitive results, but they still fell short of the RGB representation. Thus, we can safely conclude that the original RGB images already contain the most discriminative information for this dataset.

Overall, this study highlights that a highest possible image resolution and image representation in RGB format is more conducive for classification performance than use of Grayscale images or ensemble learning on individual color channels. This insight is valuable for designing efficient image classification systems in environments with limited computational resources.

Future work will explore additional datasets and use of feature selection strategies to further improve classification performance while maintaining computational efficiency. Adaptive image representations can further improve classification performance for resource-constrained applications.

REFERENCES

[1]. Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

[2]. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). ACM. https://doi.org/10.1145/2939672.2939785

[3]. Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning, 20*(3), 273–297. https://doi.org/10.1007/BF00994018

[4]. Cox, D. R. (1958). The regression analysis of binary sequences. *Journal of the Royal Statistical Society: Series B (Methodological), 20*(2), 215–242. https://doi.org/10.1111/j.2517-6161.1958.tb00292.x

[5]. Gelir, F., Akan, T., Carmichael, O. T., Bhuiyan, M. S., Conrad, S. A., Vanchiere, J. A., Kevil, C. G., & Bhuiyan, M. A. N. (2025). Topological feature extraction from multi-color channels for pattern recognition: An application to fundus image analysis. *Journal of Imaging Informatics in Medicine*. Advance online publication. https://doi.org/10.1007/s10278-025-01791-1

[6]. Gowda, S. N., & Yuan, C. (2019). ColorNet: Investigating the importance of color spaces for image classification. In C. Jawahar, H. Li, G. Mori, & K. Schindler (Eds.), Computer Vision – ACCV 2018 (Lecture Notes in Computer Science, Vol. 11364, pp. 581–596). Springer. https://doi.org/10.1007/978-3-030-20870-7_36

[7]. Karim, M. A., Tashrif, M. T. A., Hossain, M. S., Ahmed, M. K., Raha, M. T. H., & Sara, U. (2026). AMCD: A multi-domain agricultural crop and flower image dataset for deep learning-based classification. *Data in Brief*, 59, 112974. https://doi.org/10.1016/j.dib.2026.112974

[8]. Khayya, E. K., El Oirrak, A., & Datsi, T. (2024). A survey on RGB images classification using convolutional neural network (CNN) architectures: Applications and challenges. In 2024 International Conference on Circuit, Systems and Communication (ICCSC) (pp. 1–8). IEEE. https://doi.org/10.1109/ICCSC62074.2024.10617086

[9]. Kuncheva, L. I. (2004). *Combining pattern classifiers: Methods and algorithms*. John Wiley & Sons. https://doi.org/10.1002/9781118914564

[10]. Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30–39. https://doi.org/10.1109/MC.2017.9

[11]. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198

[12]. Torralba, A. (2009). How many pixels make an image? Visual Neuroscience, 26(1), 123–131. https://doi.org/10.1017/S0952523808080930

[13]. Wolpert, D. H. (1992). *Stacked generalization*. *Neural Networks*, *5*(2), 241–259. https://doi.org/10.1016/S0893-6080(05)80023-1

[14]. Yoo, W. S. (2023). Image-based quantification of color and its machine vision applications. *Technologies*, 11(2), 49. https://doi.org/10.3390/technologies11020049