PUP Journal of Science and Technology
Volume 18 (2025), pp 1-17
doi: https://doi.org/10.70922/4cya0y10

Comparison of Harris Performance as Activation Function to Rectified Linear Unit (ReLU), Leaky ReLU, and Tanh in Convolutional Neural Network for Image Classification

Luther Villacruz and Maria Lyn Bernadette Mendoza

Department of Mathematics and Physics, College of Science, Adamson University, Manila, 1000, Metro Manila

Corresponding author: luther.villacruz@adamson.edu.ph

Received: 27 June 2025; Accepted: 04 November 2025; Available online: 04 Dec 2025

Abstract: Activation functions (AFs) are the building blocks of a deep neural network (DNN) to perform image classification effectively by handling nonlinear data and extracting complex features and patterns. This paper introduces a new activation function (AF) called "Harris" a piecewise and nonlinear nonmonotonic AF inspired from the field of photonics. The AF was integrated to a simple convolutional neural network (CNN) using Canadian Institute for Advanced Research (CIFAR-10) dataset to determine the model's performance in terms of accuracy in training and testing, image classification capability, and feature extraction. Harris was able to exceed the leaky Rectified Linear Unit (ReLU) and hyperbolic tangent function (tanh) target accuracies from α-values −0.80 to −1.00 in image classification, while the testing accuracies were able to exceed the target accuracies of ReLU from α-values −0.80 to −0.95. It was able to handle negative values solving the dead neuron problem and extract complex features through its feature maps which improve the F1-scores of the CNN model in image classification.

Keywords: *Harris, Activation Function, Convolutional Neural Network, ReLU*

1. INTRODUCTION

Deep learning is a field of artificial intelligence (AI) that utilizes deep neural networks (DNNs) that mimic biological neurons in performing regression and classification tasks by reducing the use of hand-tuning for problem-solving in recent times (Roy et al., 2022). According to Goodfellow et al. (2016), a DNN is a chain structure that is composed of layers also known as hidden layers, and each layer is composed of perceptrons or artificial neurons that are connected for allowing information processing. DNNs use forward and backward propagation to optimize the perceptrons' weight values. Activation functions (AFs) are the building blocks of deep learning and are known for their ability to introduce nonlinearity to a DNN, which suggests that the network can address real-world problems, handle nonlinear data and extract nonlinear and complex features that help the DNN to be accurate, reliable, and effective in mimicking the biological neurons (Nguyen et al., 2021). AFs are divided into 3 types, the linear activation function, nonlinear monotonic activation function, and nonlinear nonmonotonic activation function (Zhu et al., 2021).

Nonlinear monotonic AFs were the earliest known AFs for computer vision. However, they are prone to the vanishing gradient problem owing to the mapping of the input, which saturates to a certain interval, such as tanh (blue), which saturates the output from (−1, 1) in its positive ($x > 0$, green shade) and negative ($x < 0$, yellow shade) regions (Figure 1). Nonlinear nonmonotonic AFs significantly improve the network because these AFs address the vanishing gradient problem and allow better output mapping as it avoids saturation of information, swish (orange) as shown in Figure 1 is unbounded on the positive region but bounded on the negative region (Zhu et al., 2021).

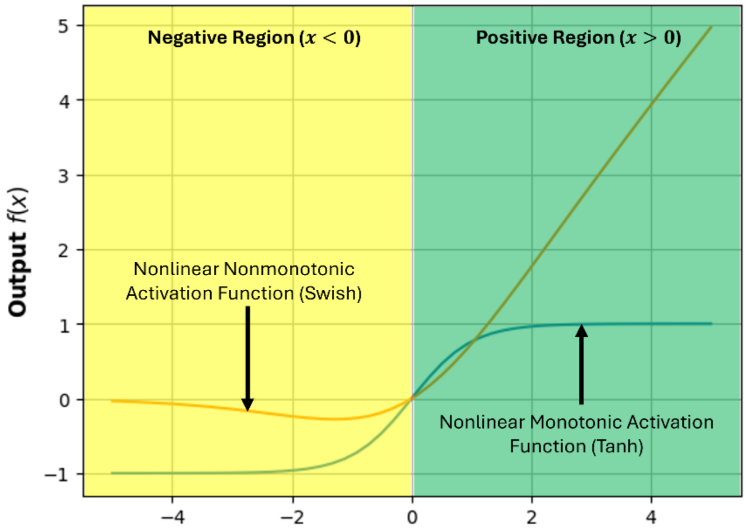

Figure 1. Examples of nonlinear monotonic (tanh) and nonlinear nonmonotonic (swish) activation functions and their behavior in the negative (yellow shade) and positive (green shade) regions.

A convolutional neural network (CNN) is a widely used DNN model because of its super resolution, which is important for CNN because it improves the representational ability of the model and its overall performance (Segawa et al., 2020; Zhu et al., 2021). The basic CNN architecture shown in Figure 2 can be divided into several parts. The base is typically composed of convolution layers that extract features using filters that are smaller tensors than the input and pooling layers, which downsize the extracted features to reduce the trainable parameters. The fully connected layers are responsible for the classification task based on the extracted features from the base.

This paper introduces a novel AF called Harris, a piecewise activation function with monotonic and non-monotonic characteristics inspired by earlier works on photonics by Fred Harris in his book section entitled "A Most Efficient Digital Filter: The Two-Path Recursive All-Pass Filter" which was published in 2007. This research compares the performance of the proposed AF with that of existing AFs, such as ReLU, Leaky ReLU, and tanh, using a CNN model for image classification and investigated the performance of Harris' α-values.

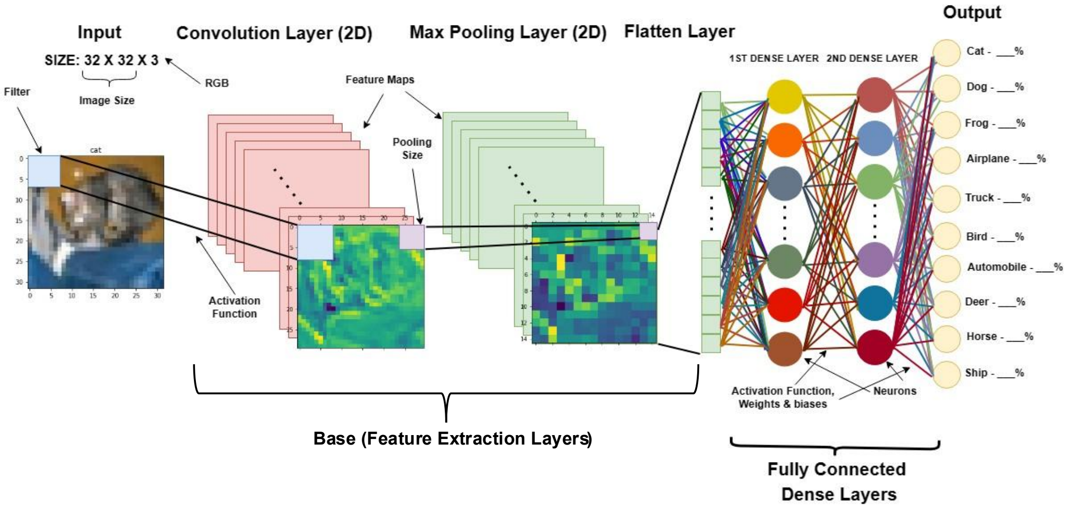

Figure 2. Single convolutional block neural network architecture.

2. METHODOLOGY

Harris Activation Function

Harris is a piecewise and nonlinear, nonmonotonic activation function defined by

$$f(x) = -n\ \text{arctan}\left[\frac{1+\alpha}{1-\alpha}\tan\left(M\frac{x}{n}\right)\right], \qquad M = -1 \qquad \text{(Eqn 1)}$$

where n is the normalizing factor which is a data-driven parameter, α represents the shape of the function as it is an adjustable parameter, and x is the input data that can be either positive or negative that represents the regions shown in Figure 1 in the x-axis. Table 1 summarizes the values used on the positive and negative regions of Harris. The n = 255 represents the maximum value of the data from CIFAR-10 (Canadian Institute for

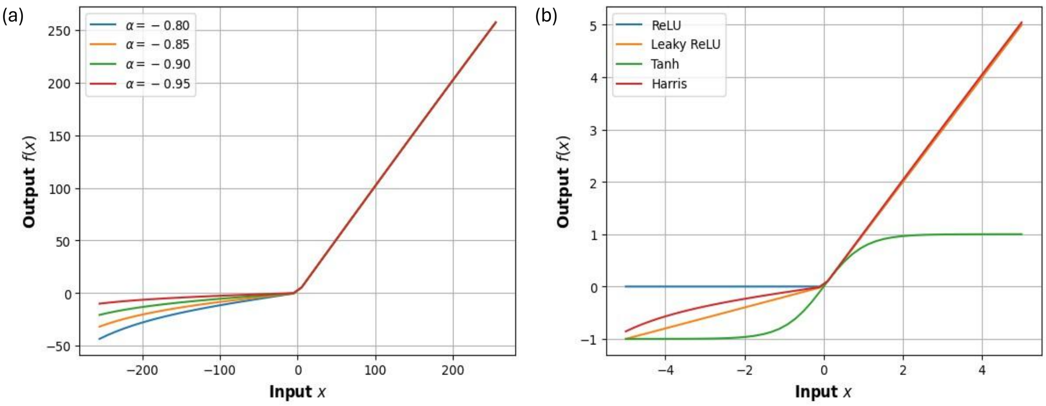

Figure 3. (a) Harris activation function with different α-values on the negative region, (b) Harris in comparison with other existing activation function in terms of graphical behavior.

Advanced Research), an established computer-vision dataset used for object recognition. The α-values when x > 0 will be 0.01 as it is the closest linear behavior on the positive region as shown in Figures 3a and 3b, while the α-values when x < 0 varies from −0.5 to −1.0 with −0.05 difference on each interval. Some of the values shown in Figure 3a are the values on the negative region. This will be the focus for the test.

The M-value was set to −1 as the positive values of M will flip the graphical behavior of the function as shown in Figure 4a which represents the original behavior of the function from Fred Harris which represents the order of the polynomial, while equation (1) represents the transfer function of phase response of the type-1 all-pass filter (Harris, 2007). Figure 4b shows that the M = -1 avoids the sudden shift of values and provides a smooth output mapping. M-value affects the periodicity of graphical behavior.

Table 1. Harris parameters as activation function.

Harris	$\boldsymbol{\alpha}$	$\boldsymbol{M}$	$\boldsymbol{n}$
Positive Region	0.01	-1	255
Negative Region	varies	-1	255

From the general function as shown in Eqn. 1, applying the piecewise, as described as one of the characteristics of an effective AF (Obla et al., 2020), and using the parameters

$$f(x) = \begin{cases} -255 \text{ arctan}\left[\dfrac{1+0.01}{1-0.01}\tan\left(-\dfrac{x}{255}\right)\right], & \text{for } x \geq 0 \\ -255 \text{ arctan}\left[\dfrac{1+\alpha}{1-\alpha}\tan\left(-\dfrac{x}{255}\right)\right], & \text{for } x < 0 \end{cases} \qquad \text{(Eqn 2)}$$

from Table 1 transforms the function given in Eqn. 2, which is the current form of the AF applied to the CNN model. In addition, shown in Figure 5 is the derivative of Eqn. 2. The bounded behavior of the derivative prevents Harris from mapping extreme values as an output during backpropagation. Thus, identifying the derivative of AF helps to determine if the DNN model has a possibility to suffer from exploding or vanishing gradient problems during backpropagation (Obla et al., 2020; Roy et al., 2022).

Dataset and Convolutional Neural Network Model

The CNN model was trained using the CIFAR-10 dataset which consists of 60,000 Red, Blue, Green (RGB) images with 32×32 resolution categorized into 10 different classes

(a)

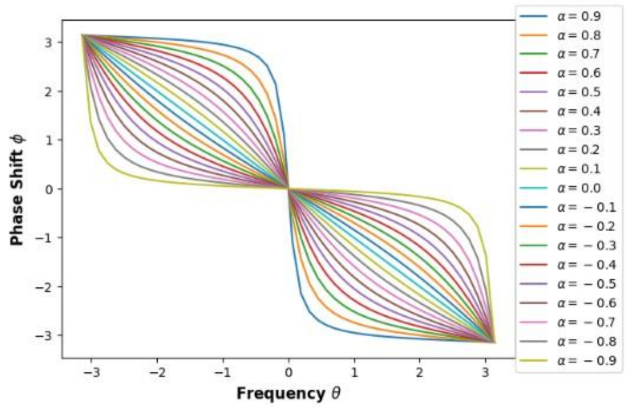

(b)

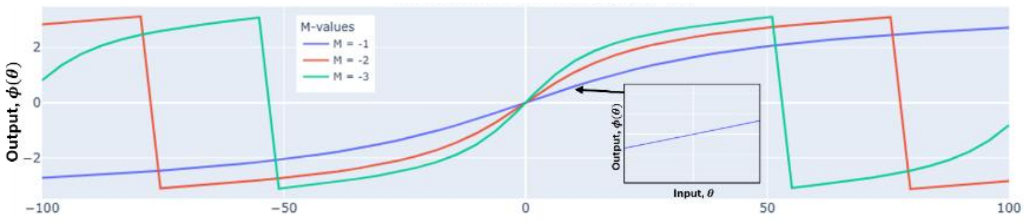

Figure 4. (a) Harris when the M-values is positive and (b) the behavior in different M-values.

(airplane, automobile, bird, cat, deer, dog, frog, horse, ship, truck) (Krizhevsky., 2009). The 50,000 RGB images were used for the training while 10,000 RGB images were used for testing.

Figure 2 showed the CNN architecture which is composed of an (i) input layer with dimensions 32×32×3 that conforms to the dimension of the images in the dataset, (ii) one convolution layer, (iii) one max pooling layer, and (iv) two fully connected tuned using the hyperband under different AFs with their corresponding training and test accuracies which will also be the target accuracies for the Harris under different α-values. For this paper, 3

AFs were involved for the comparison in terms of the performance of the Harris shown in Figure 3b that includes leaky ReLU, ReLU, and tanh.

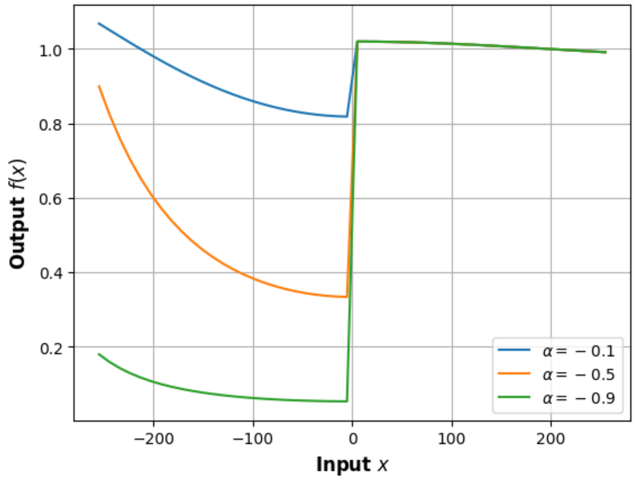

Figure 5. First-order derivative of the Harris activation function.

The experimentation involved 3 trials with the same CNN model and set of α-values on the negative region as shown in Table 1. This means that 11 α-values were involved per trial which produced training and testing accuracies. CNN models that have been trained were saved and used to produce feature maps and image classification in a set of images that the CNN model were seen for the first time.

Model Evaluation and Statistical Analysis

The results for the 3 trials for each α-values underwent a one-way analysis of variance (ANOVA) to determine if each α-values had significant difference within the obtained training and testing accuracies with a level of significance at 5%. Accuracies for both training and testing were compared to other AFs as target training and testing accuracies (Table 2) to determine which of the Harris' α-values exceed the accuracies from leaky ReLU, ReLU, and tanh. The F1-score provided the classification performance of CNN model containing different AFs in classifying images from different categories. The confusion matrix helped identify which categories does the misclassified images fall under during the testing. The feature maps were used to identify the complexity of features the CNN models learns during the training.

Table 2. Hyperparameters in each layer with their corresponding activation functions and their accuracy.

	Convolution Layer		Max Pooling Layer	Fully Connected Layers		Target Accuracies	
Activation Functions	Number of Filters	Filter Size	Pooling Size	1st Dense Layer Neurons	2nd Dense Layer Neurons	Training Accuracy	Testing Accuracy
Leaky ReLU	40	2 × 2	2 × 2	140	140	90.10%	61.38%
ReLU	40	2 × 2	2 × 2	140	140	94.12%	61.33%
Tanh	40	2 × 2	2 × 2	140	140	59.67%	53.26%

3. RESULTS AND DISCUSSION

Training and Testing Accuracies

The training and testing accuracies of the Harris across different α-values is summarized in Figure 6. The training accuracies in Figure 6a with α-values of −0.85 (93.20%), −0.90 (93.30%), −0.95 (93.34%), and −1.0 (93.76%) highlighted by the red box produced high average accuracies with −1.0 having the highest average accuracy. The graphical trend of each trial with the average was shown in Figure 7b. Testing accuracy is important as it evaluates the model's capability to predict in the images it sees for the first time. The α-values −0.8 (62.27%), −0.85 (62.27%), −0.90 (62.28%), and −0.95 (61.98%) provided a high average accuracy highlighted in red box (Figure 6b) with -0.9 having the highest average accuracy. In Figure 7d, the closeness of the average testing accuracies from α-values −0.8 to −0.95 can also be observed.

In comparison with other AFs, leaky ReLU, ReLU, and tanh as shown in Figures 7a and 7c respectively, the α-values −0.80 to −0.95 consistently surpass the target accuracies of leaky ReLU and tanh while in Table 2 only ReLU's training accuracy was not surpassed by any of the α-values. This means that α-values from −0.8 to −0.95 are consistent in terms of producing accuracies that exceeds leaky ReLU and tanh for both the training and testing.

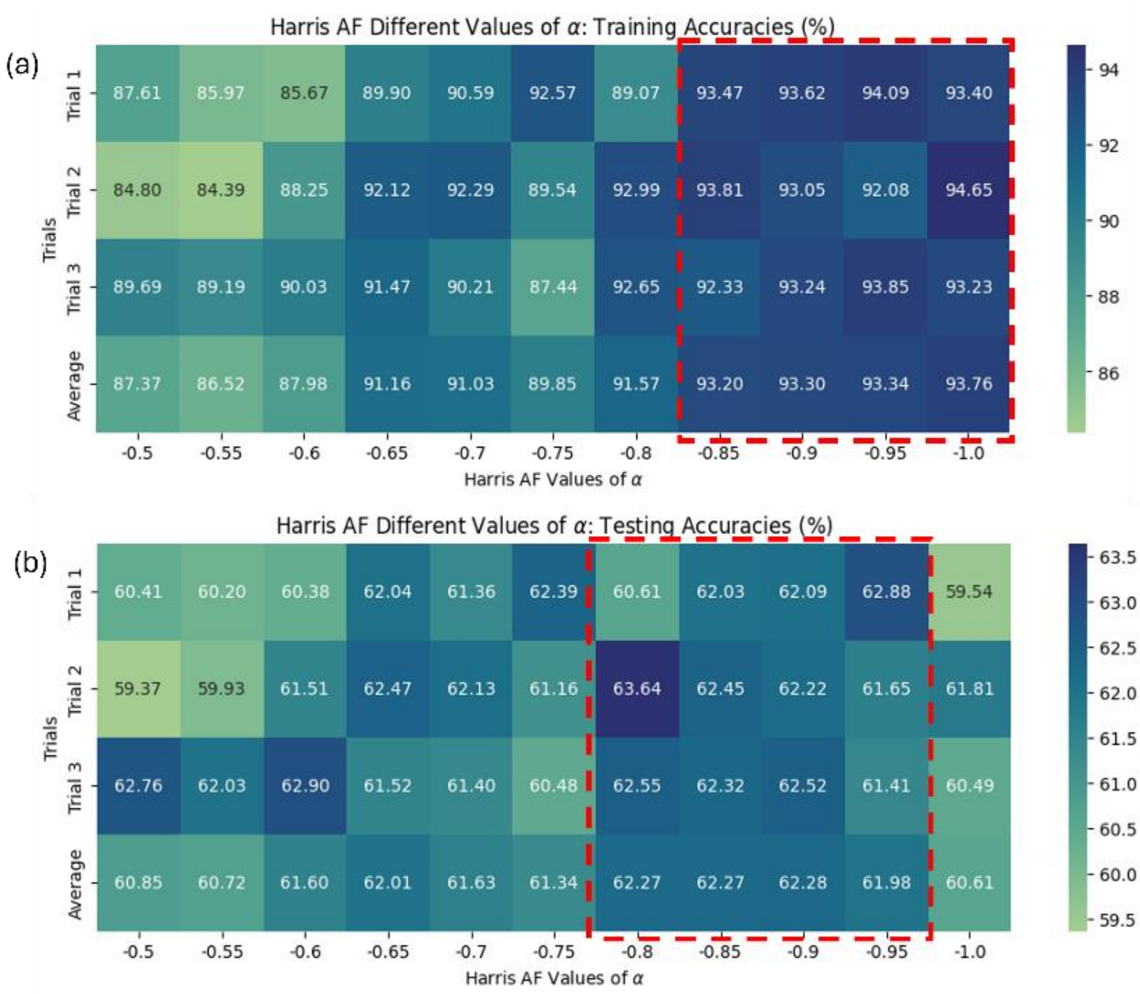

Figure 6. Heatmap summary of accuracy values from the (a) training and (b) testing from each trial with their average using different Harris α-values.

Interestingly, the trend of the average accuracy for training and testing in Figures 7a and 7c is consistent from 0.5 to −1.0. The accuracy tends to increase in both training and testing, but with the exception of the testing accuracy with an α-value of −1.0 which has a drop of testing accuracy. Correlating this using Figure 3a when the α-value increases, Harris reaches a zero output wherein the AF could mimic the ReLU. However, due to the difference of their derivatives, the backpropagation process affects the weight values that can affect the overall training of the CNN model. When Harris α-value is −1.0 as shown in Figures 7a and 7c, target accuracy for ReLU in both the training and testing cannot exceed. This is similar for other α-values with their training accuracy as shown in Figure 7a.

Image Classification

Table 3 showed the F1-scores that assessed the ability of the CNN model to classify and correctly label images from the dataset. Based on the table, Harris was able to get high accuracies as compared to other AFs in some categories. Harris got a high F1-score for the automobile in α-value −0.80 (76%), bird in α-values −0.85 (53%) and −0.95 (53%), cat in α-value −0.80 (46%), frog in α-values −0.80 (71%) and −0.95 (71%), horse in α-values −0.80 (69%) and −0.90 (69%), and truck in α-value −0.80 (71%). On the other hand, categories such as the airplane, deer, dog, and ship someα-values were able to reach the F1-score of the other AF.

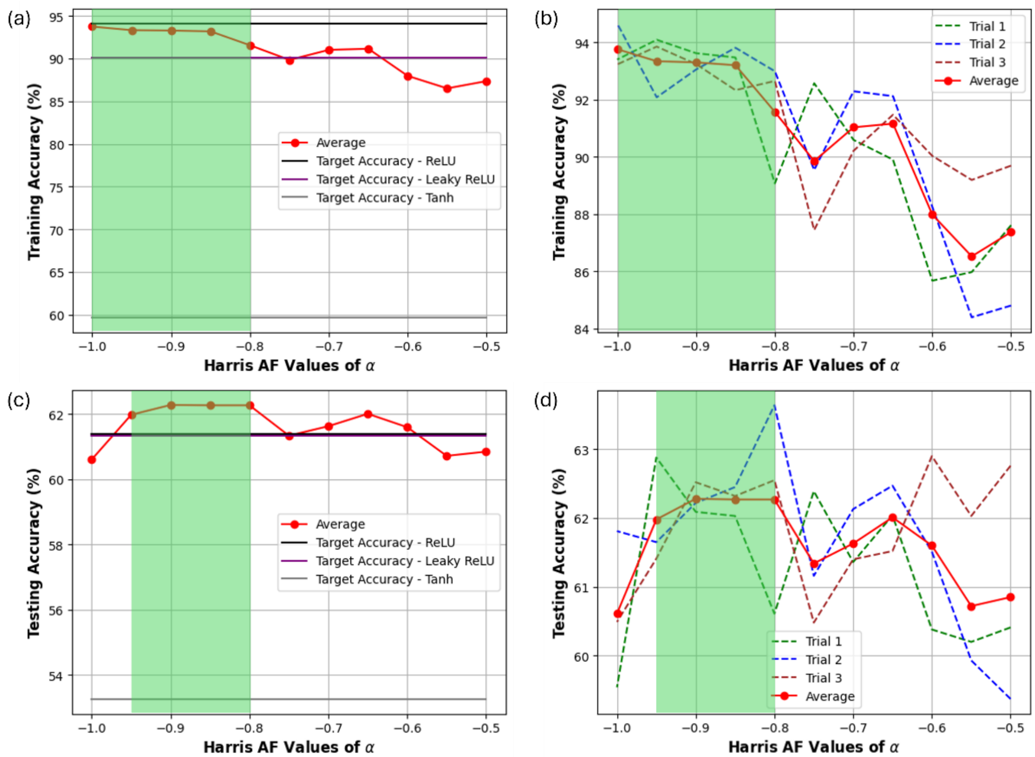

Figure 7. Graphical representation of the average accuracies using different Harris α-values in comparison with the target accuracies for (a) training and (b) testing, and the trend for each trial together with the average for the Harris α-values for (b) training and (d) testing.

Interestingly, by looking at the confusion matrix of Harris at some α-values shown in Figure 8 for the misclassified images, which identifies on what categories they were classified instead of their actual categories, there were similarities on how Harris misclassified the images from CIFAR-10. For the airplane, Harris misclassified most of the images as either a bird or a ship. For the automobile, it misclassified most of the images as either a truck or a ship. For the bird, it misclassified most of the images as either a cat, deer, or a dog. For the cat, it misclassified most of the images as a dog. For the ship, it misclassified most of the images as an airplane. For the truck, it misclassified most of the images as either an airplane or an automobile. While there are categories that still varies like the deer, dog, and frog where the classification varies depending on the Harris α-value.

Comparing some of the misclassified images to the other AFs as shown in Figures 9 and 10, it can be observed that Harris was able to classify most of the sample images in different α-values. For example, the α-value −0.95 misclassified 4 out of 21 images shown in Figure 10d similar to leaky ReLU and ReLU as shown in Figures 9a and 10b, respectively. Overall, Harris was able to compete with other AFs in terms of the F1-score and even exceed in some of the categories.

Table 3. F1-scores of different activation functions.

				Harris α-value			
Categories	Leaky ReLU	ReLU	Tanh	-0.80	-0.85	-0.90	-0.95
Airplane	69%	68%	53%	66%	67%	66%	68%
Automobile	71%	74%	63%	76%	75%	74%	71%
Bird	49%	50%	44%	52%	53%	50%	53%
Cat	39%	45%	35%	46%	45%	44%	45%
Deer	58%	52%	46%	58%	56%	57%	56%
Dog	51%	50%	40%	48%	49%	50%	50%
Frog	70%	70%	63%	71%	69%	68%	71%
Horse	68%	68%	60%	69%	65%	69%	68%
Ship	74%	74%	58%	73%	73%	74%	73%
Truck	65%	69%	64%	71%	68%	70%	68%

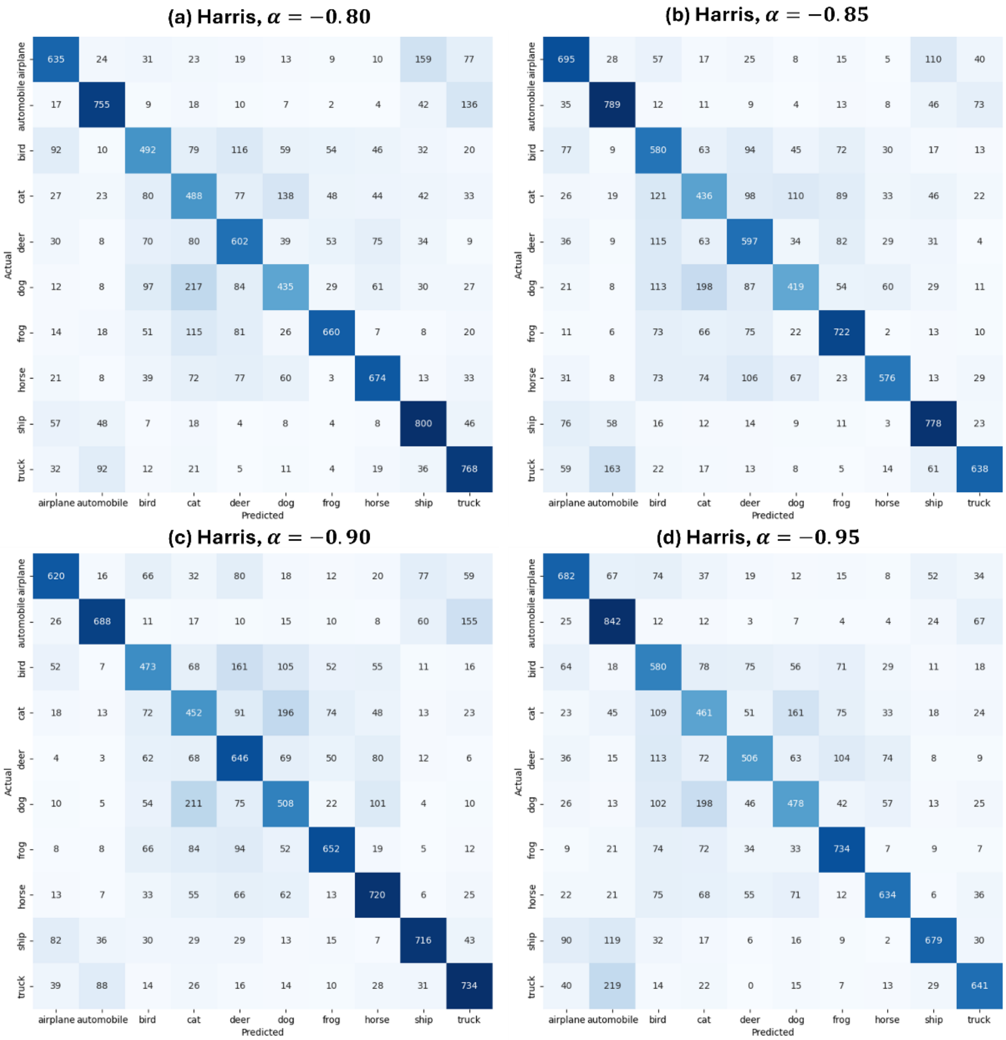

Figure 8. Confusion matrix of some of the α-values Harris activation function with high accuracies (a) $\alpha = -0.80$, (b) $\alpha = -0.85$, (c) $\alpha = -0.90$, and (d) $\alpha = -0.95$.

Feature Maps

Feature maps are essential that the AF were able to extract patterns and features that can help improve the CNN model during training. The CNN model architecture shown in Figure 2 has only one convolution layer, which extracts general features. As convolution layers increased, the feature extraction becomes more detailed (Pattanayak., 2023). In addition, a neural network with one hidden layer can only approximate a function while more than two hidden layers can learn complex representations like an automatic feature engineering (Nguyen et al., 2021).

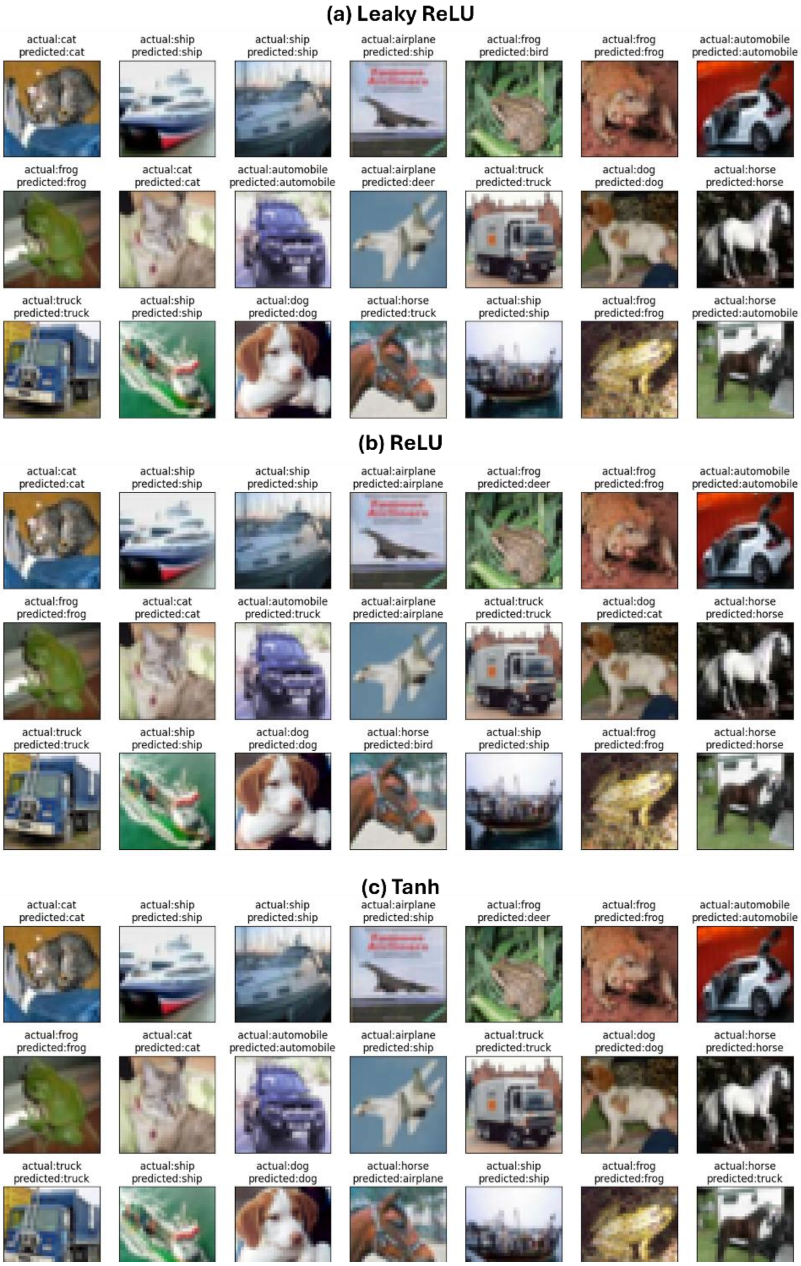

Figure 9. Sample predictions of (a) Leaky ReLU, (b) ReLU, and (c) tanh from theCIFAR-10 images.

In Figure 11a, leaky ReLU, extracts smooth feature maps that are almost similar while some features differ in terms of texture like features 1, 6, 7, and 8. Vanishing gradient problem cannot be seen in the feature maps of leaky ReLU since the CIFAR-10 images were normalized so that tanh will not be able to saturate. The feature maps of leaky ReLU suggest that during backpropagation the filter values of the convolution layer has close values, the reason why some features having almost similar in terms of texture and shape during feature extraction in the convolution layer like features 37 and 38.

For the ReLU shown in Figure 11b, vanishing gradient problem can be observed in most of the features as it can’t handle negative values and only accepts positive values during the training of the CNN model; but it is not clear if the backpropagation during the weight update can also cause vanishing gradient, given that the images were normalized.

Lastly, for tanh in Figure 11c, the saturation is a problem during the training that caused most of the features to have similar textures, which suggests that most of the values of the filter suffers from the saturation during backpropagation. It learns similar features which reduced its complexity during training.

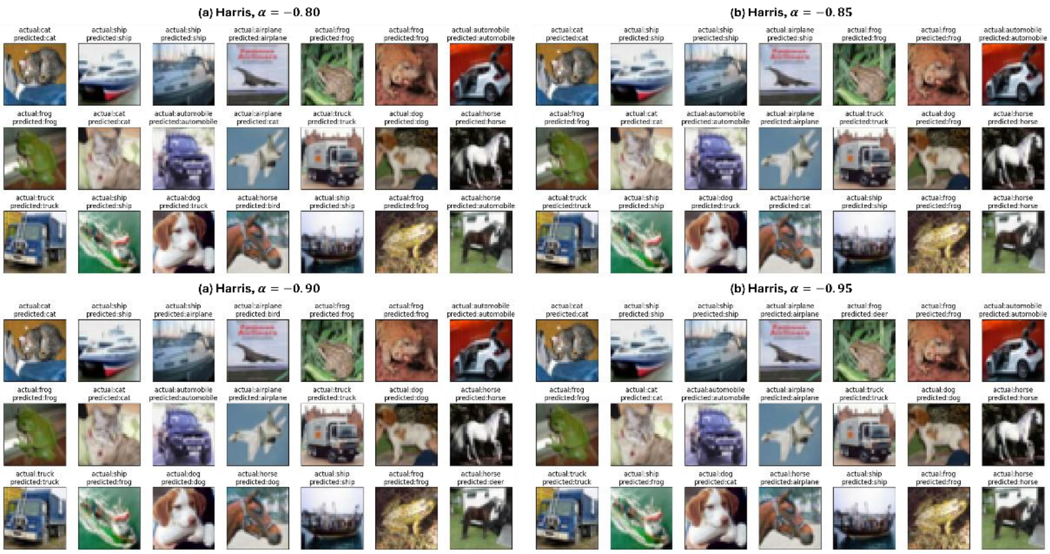

Figure 10. Sample predictions of some of the α -values Harris with high accuracies performance (a) α = −0.80, (b) α = −0.85, (c) α = −0.90, and (d) α = −0.95.

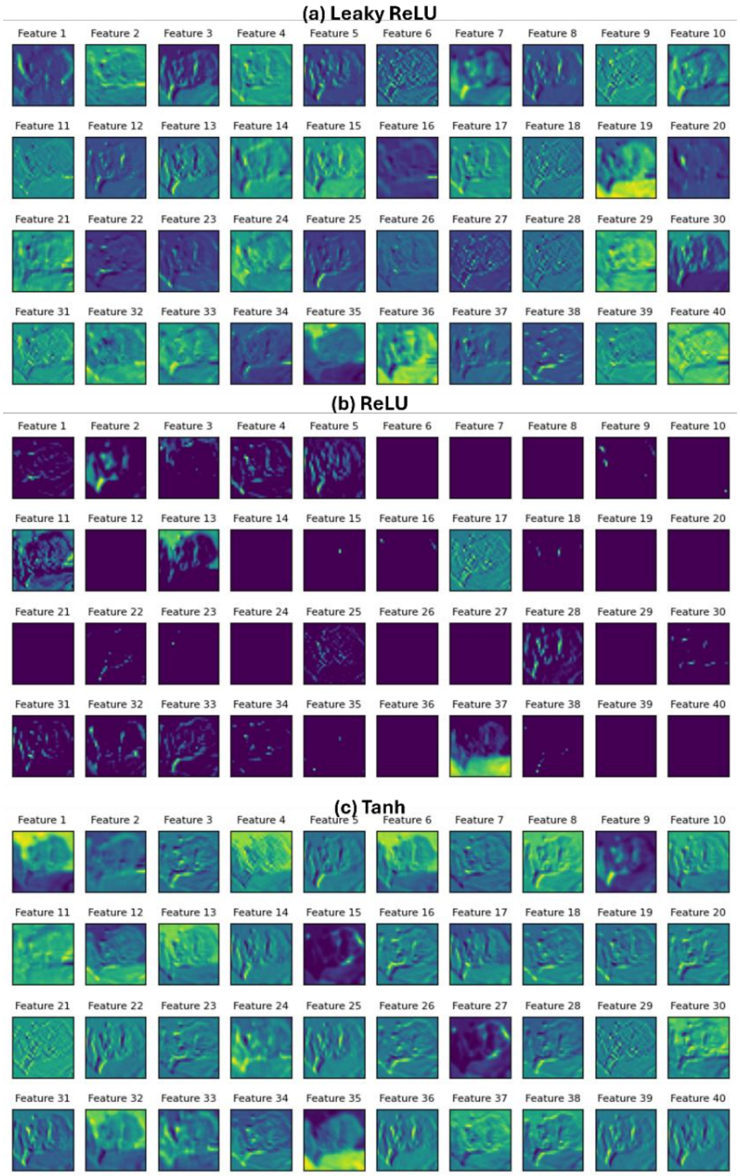

Figure 11. Feature maps of (a) Leaky ReLU, (b) ReLU, and (c) tanh from the CIFAR-10 images.

Harris was able to overcome these problems as shown in Figure 12. For the α -value −0.85, different features, textures, and shapes, Harris has extracted from the convolution layer. For example, feature 5 detects the edges of the cat image, feature 3 emphasized the platform where the cat is placed, and feature 37 focused on the shape of the cat that can aid the learning in terms of the position and geometry of the cat. The dead neuron problem is also addressed by Harris which means that it was able to handle negative values for both the forward propagation and backpropagation. The derivative of Harris shown in Figure 5 is essential, as it will affect the backpropagation process and determine if the model will suffer vanishing or exploding gradient problems affecting the learning of the DNN (Nguyen et al., 2021).

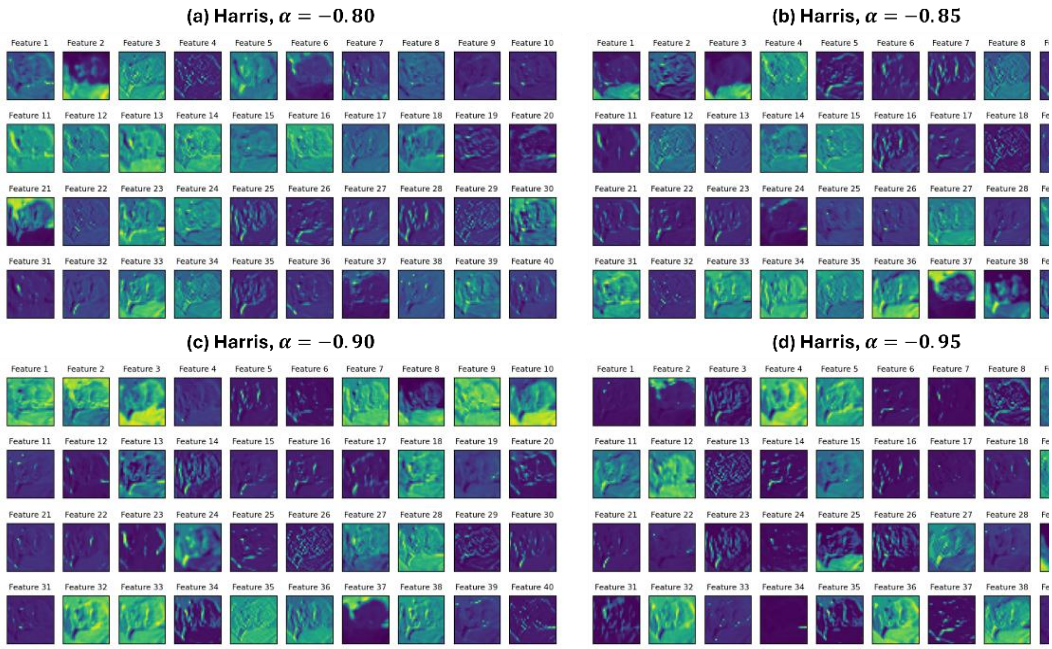

Figure 12. Feature maps of some of the α -values Harris activation function with high accuracies performance (a) $\alpha = -0.80$, (b) $\alpha = -0.85$, (c) $\alpha = -0.90$, and (d) $\alpha = -0.95$.

Statistical Analysis

The one-way ANOVA shows that for the training accuracy of the α-values from −0.80 to −1.00 obtained a p-value 0.000108 which is less than the 5% level of significance, and an F-value of 6.66 which is above the F-critical 2.30 which indicates that there is a significant difference in training accuracies in different α-values for Harris. The results supported the idea that the feature maps learned by the CNN model affects the model training when using Harris as an AF at different α-values.

For the testing accuracy, the one-way ANOVA obtained a p-value 0.38 which is greater than 5% level of significance and an F-value of 1.15 which is less than the F-critical 2.30. This means that there is no significant difference in the testing accuracy across Harris' α-values. This means that the CNN models trained and saved using Harris as an AF that it classifies images almost similarly as it sees it for the first time. This supports how Harris misclassify images as shown in the confusion matrix in Figure 8 wherein in some images it tends to have similar predicted categories which caused the testing accuracy to have no significant difference across the α-values.

4. CONCLUSION

Harris was able to outperform the leaky ReLU and the tanh in α -values in both the training and testing accuracies, that is α-values from -0.80 to -0.95 provided consistent accuracies with an average accuracy of 93.20%, 93.30%, and 93.34%, respectively for the training while α-values from -0.80 to -0.95 have average accuracies of 62.27%, 62.27%, 62.28%, and 61.98%, respectively for the testing. Harris was able to produce high F1-scores and was able to surpass the other AFs in most of the categories such as automobile, bird, cat, frog, horse, and truck at certain α-values. Harris, as an AF, was able to address the vanishing gradient problem and dead neuron problem as its feature map produces complex textures, shapes, and features compared to leaky ReLU, ReLU, and tanh. Based on the results from the one-way ANOVA, Harris α-values training accuracy showed a significant difference, while no significant difference in terms of the testing accuracy.

5. ACKNOWLEDGEMENT

We would like to acknowledge the support of the Mathematics and Physics Department and Center for Research and Development of Adamson University.

6. REFERENCES

Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning* (Vol. 1, No. 2). Cambridge: MIT press.

Harris, F. (2007). *Streamlining Digital Signal Processing: A Tricks of the Trade Guidebook* (pp. 85–104).

Krizhevsky, A., & Hinton, G. (2009). Learning multiple layers of features from tiny images.

Nguyen, A., Pham, K., Ngo, D., Ngo, T., & Pham, L. (2021). An analysis of state-of-the-art activation functions for supervised deep neural network. In *2021 International conference on system science and engineering (ICSSE)* (pp. 215–220). IEEE.

Obla, S., Gong, X., Aloufi, A., Hu, P., & Takabi, D. (2020). Effective activation functions for homomorphic evaluation of deep neural networks. *IEEE access*, *8*, 153098-153112.

Pattanayak, Santanu. (2023). Pro deep learning with tensorflow 2.0. *Pro Deep Learning with TensorFlow*, *2*.

Roy, S. K., Manna, S., Dubey, S. R., & Chaudhuri, B. B. (2022). Lisht: Non-parametric linearly scaled hyperbolic tangent activation function for neural networks. In *International Conference on Computer Vision and Image Processing* (pp. 462–476). Springer.

Segawa, R., Hayashi, H., & Fujii, S. (2020). Proposal of new activation function in deep image prior. *IEEJ Transactions on Electrical and Electronic Engineering, 15*, 1248–1249.

Zhu, H., Zeng, H., Liu, J., & Zhang, X. (2021). Logish: A new nonlinear nonmonotonic activation function for convolutional neural network. *Neurocomputing, 458*, 490–499.