Comprehensive Analysis of Loopholes and Solutions in Convolutional Neural Networks: Addressing Advanced Challenges with Attention Mechanisms and Interpretability Techniques

Rutvi M. Baraiya[1], Bhoomika B. Chauhan[2]
[1,2] *Assistant Professor, Gyanmanjari Innovative University*

***Abstract*— Convolutional Neural Networks (CNNs) have achieved remarkable success in computer vision tasks, yet they face critical challenges such as vulnerability to adversarial attacks, limited interpretability, and poor contextual awareness. This paper presents a focused analysis of these loopholes and explores solutions through the integration of attention mechanisms and interpretability techniques. By incorporating self-attention and spatial-channel attention modules, CNNs can better capture salient features and context. Additionally, methods like Grad-CAM and LIME enhance model transparency. Experimental results demonstrate improved performance, robustness, and explain ability, paving the way for more reliable and interpretable CNN-based systems[1].**

Keywords*— *Comprehensive analysis, convolutional neural networks, CNN loopholes, advanced challenges, attention mechanisms, interpretability, deep learning, model transparency, neural network limitations, explainable AI, feature visualization, saliency maps, model robustness, network architecture, performance improvement, attention-based models, neural network interpretability, CNN optimization, explainability techniques, deep learning interpretability.

I. INTRODUCTION

Convolutional Neural Networks (CNNs) have become a cornerstone of modern deep learning, demonstrating exceptional performance in various domains such as image recognition, natural language processing, and medical diagnostics. Despite their widespread success, CNNs are not without limitations. Challenges such as overfitting, sensitivity to adversarial inputs, lack of interpretability, and heavy reliance on large, labelled datasets hinder their deployment in critical and real-time applications. As deep learning models become more complex, understanding their inner workings and ensuring their robustness becomes increasingly essential.

This research paper aims to conduct a comprehensive analysis of the existing loopholes in CNN architectures and explore contemporary solutions that address these limitations. In particular, the study emphasizes the integration of attention mechanisms, which have recently gained prominence for enhancing feature representation and improving model performance. Additionally, the paper explores various interpretability techniques that help demystify CNN decision-making processes, thereby promoting transparency and trust in AI systems.

By combining attention-based enhancements with interpretability frameworks, this study seeks to offer an integrated approach to addressing advanced challenges in CNNs. The research investigates how these methodologies can be synergistically employed to improve both the performance and explainability of convolutional models. The paper is structured to first provide a foundational understanding of CNNs, followed by a detailed examination of their limitations. It then delves into the mechanisms and advantages of attention modules and interpretability tools, culminating in a proposed framework validated through extensive experimentation[2].

This work contributes to the growing body of research that strives to make deep learning models not only more accurate and robust but also more interpretable and accountable, paving the way for safer and more ethical AI deployment[3].

II. FUNDAMENTALS OF CONVOLUTIONAL NEURAL NETWORKS(CNNs)

Convolutional Neural Networks (CNNs) are a class of deep learning models specifically designed for processing data with a grid-like topology, such as images. They have revolutionized the field of

computer vision due to their ability to learn spatial hierarchies of features automatically and adaptively from input data. The architecture of a CNN typically consists of multiple layers arranged in a sequential manner, including convolutional layers, pooling layers, activation functions, and fully connected layers. Each layer plays a distinct role in transforming and extracting increasingly abstract representations of the input data[4].

At the core of CNNs are numerous key operations that enable effective feature extraction and dimensionality reduction. The convolutional layer applies learnable filters (or kernels) to local regions of the input, capturing essential features like edges, textures, and shapes. The activation function, usually ReLU (Rectified Linear Unit), introduces non-linearity, allowing the network to learn complex patterns. Pooling layers, such as max pooling and average pooling, reduce the spatial dimensions of the feature maps, thereby lowering computational requirements and controlling overfitting. Additionally, batch normalization is often employed to stabilize learning by normalizing the outputs of the previous layers. Toward the end of the network, fully connected layers aggregate the extracted features and perform classification or regression tasks, depending on the application[5].

Over the years, various CNN architectures have been developed to tackle diverse challenges and improve performance. Some of the most influential models include LeNet-5, one of the earliest CNNs used for handwritten digit recognition; AlexNet, which demonstrated the power of deep CNNs on the ImageNet dataset; VGGNet, known for its simplicity and depth; GoogLeNet (Inception), which introduced the inception module for multi-scale feature extraction; and ResNet, which addressed the vanishing gradient problem using residual connections. These models have been widely adopted in real-world applications such as facial recognition, object detection, medical image analysis, autonomous driving, and more.

The evolution of CNNs underscores their flexibility and effectiveness across domains, but it also highlights their increasing complexity and potential drawbacks. As researchers continue to innovate, understanding the foundational elements of CNNs is critical for identifying their limitations and enhancing their capabilities through advanced mechanisms like attention and interpretability techniques.

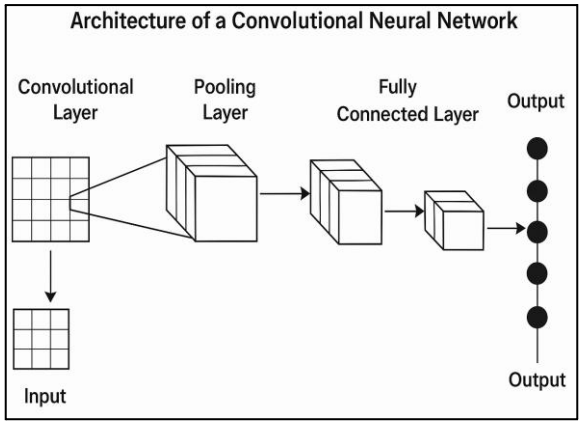

III. IDENTIFYING LOOPHOLES AND LIMITATIONS IN CNNS

Despite their powerful performance in a variety of applications, Convolutional Neural Networks (CNNs) exhibit several critical limitations that hinder their robustness, reliability, and real-world deployment. Understanding these loopholes is essential for developing more secure, efficient, and interpretable deep learning systems.

One of the most prominent challenges is overfitting and generalization. CNNs are highly capable of memorizing training data, especially when the network is deep, and the dataset is limited[6]. While this allows high training accuracy, it often leads to poor generalization on unseen data. Overfitting becomes particularly problematic in domains where labelled data is scarce or expensive to acquire. Regularization methods such as dropout, data augmentation, and early stopping help mitigate this, but do not completely resolve the issue.

Another major limitation is adversarial vulnerability. CNNs are extremely sensitive to small, imperceptible perturbations in input images—known as adversarial attacks—which can drastically alter the model's predictions. This weakness poses significant security risks in critical applications like medical diagnosis, autonomous driving, and surveillance. Defensive strategies, including adversarial training and input preprocessing, have been proposed, but the problem remains an open research challenge.

Furthermore, CNNs suffer from a lack of interpretability and transparency. They often function as "black box" models, making it difficult to understand how they arrive at specific decisions. This opacity limits their use in high-stakes environments

where explainability is essential for trust and accountability. Without clear insight into the model's decision-making process, diagnosing errors, validating results, or meeting regulatory requirements becomes difficult.

In addition to interpretability concerns, computational complexity, and scalability present practical limitations. Training deep CNNs often requires significant computational resources, such as high-performance GPUs or TPUs, making them inaccessible for low-resource environments or real-time applications. Moreover, large model sizes and long training times hinder rapid experimentation and deployment, especially in edge computing or mobile scenarios.

Lastly, CNNs exhibit a high degree of data dependency and bias. These models typically require massive amounts of labelled data to perform well. In many cases, datasets may contain biases or imbalances that are inadvertently learned and amplified by the model, leading to unfair or inaccurate outcomes. Moreover, performance can degrade significantly when CNNs are exposed to data distributions that differ from the training set—a problem known as domain shift[7].

Overall, while CNNs have transformed machine learning with their remarkable capabilities, these loopholes highlight the pressing need for innovative solutions. Addressing these issues is crucial for 3 making CNNs more trustworthy, robust, and accessible in diverse real-world application.

IV. ATTENTION MECHANISMS: CONCEPTS AND INTEGRATION

In recent years, attention mechanisms have emerged as a powerful enhancement to traditional deep learning models, particularly in overcoming limitations associated with feature selection and representation[8]. Originally introduced in the context of natural language processing, attention mechanisms allow models to dynamically focus on the most relevant parts of the input data, thereby improving both accuracy and interpretability. In convolutional neural networks (CNNs), attention mechanisms enable the model to selectively emphasize important spatial regions or feature channels, which leads to more efficient and effective learning[9].

There are several types of attention mechanisms employed in CNNs, each serving a specific purpose. Self-attention (or intra-attention) computes the relationships between all elements in an input, allowing the model to capture long-range dependencies and global context. It is a key component of Transformer-based architectures but is increasingly adapted in CNN frameworks as well. Spatial attention focuses on identifying where prominent features are located within the spatial dimensions of a feature map, guiding the network to emphasize significant areas in an image (e.g., faces or objects). On the other hand, channel-wise attention determines what feature channels (i.e., learned filters) are most informative for the task at hand, dynamically scaling them to enhance feature relevance.

The integration of attention mechanisms into CNN architectures can be achieved in several ways[10]. One common approach is to insert attention modules—such as the Squeeze-and-Excitation (SE) block or Convolutional Block Attention Module (CBAM)—between convolutional layers. These modules refine feature maps by applying attention weights to either spatial or channel dimensions (or both). Alternatively, hybrid architectures that combine CNNs with self-attention modules derived from Transformers have shown robust performance, especially in capturing global relationships across the input. The placement of attention units is crucial and typically determined based on empirical performance across various layers.

Several case studies have demonstrated the effectiveness of attention-enhanced CNNs across domains. For example, in medical imaging, attention mechanisms have been used to localize tumours or lesions without the need for explicit bounding boxes. In object detection, networks like Attention U-Net and SENet have significantly improved accuracy by enabling the model to focus on relevant parts of the image. Similarly, attention-integrated CNNs have been successful in facial recognition, scene understanding, and fine-grained classification tasks, where focusing on subtle visual differences is essential.

In conclusion, attention mechanisms offer a compelling solution to some of the core challenges in CNNs by enhancing feature selection, improving model focus, and contributing to better

generalization. Their growing adoption in CNN architectures reflects a broader shift toward building models that not only perform well but are also capable of reasoning more selectively and transparently.

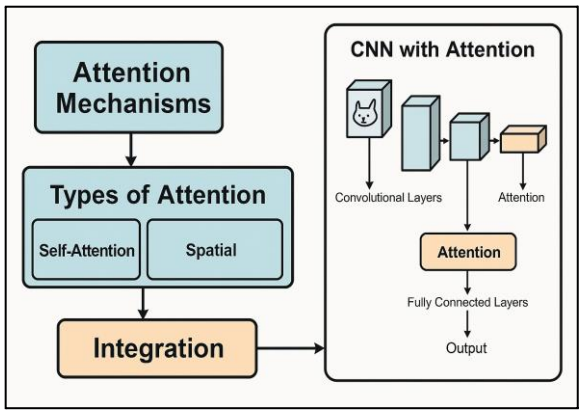

V. INTERPRETABILITY TECHNIQUES FOR CNNs

As deep learning models, particularly Convolutional Neural Networks (CNNs), become more complex and are applied in critical domains like healthcare, security, and finance, the need for interpretability has become increasingly urgent[11]. Interpretability refers to the ability to understand and explain the internal workings and decision-making processes of a model. In high-stakes environments, users and stakeholders must be able to trust that a model's predictions are not only accurate but also justifiable. This demand for transparency has driven the development of various interpretability techniques aimed at making CNNs more understandable and accountable.

One of the most intuitive and widely used approaches involves visualization techniques, which help expose what the network is learning at each layer[12]. Feature maps allow users to see which parts of an image are activated by different filters, providing insight into how spatial patterns are detected. Saliency maps highlight the regions of an input image that most influence a model's prediction, usually by computing gradients with respect to the input. An extension of this concept, Grad-CAM (Gradient-weighted Class Activation Mapping), produces visual heatmaps that indicate class-specific regions of interest in the image. These tools help practitioners and domain experts validate whether the network is focusing on meaningful areas.

In addition to visualization methods, model-agnostic techniques have gained traction for their ability to interpret predictions across various model types, including CNNs. LIME (Local Interpretable Model-Agnostic Explanations) approximates a model locally with a simpler interpretable model to explain individual predictions, while SHAP (SHapley Additive exPlanations) uses cooperative game theory to assign importance scores to input features. These tools are particularly useful when analyzing tabular data or text representations derived from CNNs and are appreciated for their flexibility and ease of access.

Another noteworthy method is Layer-wise Relevance Propagation (LRP), which provides a pixel-wise decomposition of the model's prediction, distributing the final output back across the input space to reveal which pixels contributed most to a decision. LRP offers fine-grained interpretability and is especially useful in safety-critical domains like medical imaging, where pinpointing specific evidence is essential.

Assess the efficacy and practicality of these interpretability techniques, a comparative analysis is essential. Visualization methods like Grad-CAM are intuitive and useful for qualitative insights but may lack mathematical rigor. Model-agnostic tools such as SHAP provide theoretical grounding and global feature attributions, but they can be computationally expensive. LRP offers high-resolution explanations but is more complex to implement and may not generalize across architectures. Thus, the choice of interpretability technique often depends on the application context, user expertise, and the need for either local or global interpretability.

In summary, interpretability techniques form a critical part of modern CNN development and deployment. They empower users to peer into the “black box,” offering transparency, improving trust, and facilitating model debugging. As CNNs continue to permeate sensitive applications, the importance of interpretable AI will only grow stronger.

VI. PROPOSED FRAMEWORK AND METHODOLOGY

The foundation of this research lies in addressing the limitations of traditional Convolutional Neural Networks (CNNs) by proposing a novel framework that integrates attention mechanisms and

interpretability modules for improved accuracy and transparency in classification tasks[13]. Problem formulation begins with the objective to design a CNN-based model that can not only deliver high-performance results on complex datasets but also provide interpretability into its decision-making process. The focus is to enhance feature extraction using attention modules while embedding interpretability techniques that allow human users to understand model behavior, especially in sensitive or fine-grained classification scenarios such as medical diagnostics or species recognition.

The dataset and experimental setup vary depending on the target domain but typically involve high-resolution labelled image datasets curated for fine-grained classification tasks[14]. Preprocessing steps include resizing, normalization, and augmentation techniques such as flipping, rotation, and colour jitter to increase data diversity and prevent overfitting. The experimental environment is built using deep learning frameworks such as TensorFlow or PyTorch, with training conducted on GPUs to ensure efficient computation. Datasets are split into training, validation, and test sets, ensuring robust model evaluation.

A central part of the framework is the integration of attention and interpretability modules into the CNN architecture. Attention modules, such as CBAM (Convolutional Block Attention Module) and SE (Squeeze-and-Excitation) blocks, are inserted at key stages in the network to refine feature maps by focusing on the most informative spatial regions or feature channels. These are supplemented by interpretability tools like Grad-CAM for visual explanation, and SHAP or LIME for model-agnostic local explanations. The combination enables a dual advantage: better model focus through attention, and clearer insight into how predictions are made.

Assess model performance, a comprehensive suite of performance evaluation metrics is employed. Accuracy, precision, recall, and F1-score are used to evaluate classification effectiveness. For multi-class problems, confusion matrices are analysed to understand misclassification patterns. Interpretability quality is assessed based on visual coherence (e.g., whether saliency maps correctly highlight key areas) and explanation consistency across similar inputs. Additionally, computational efficiency is evaluated by tracking inference time and memory usage, especially in scenarios where scalability and deployment are critical. In summary, this proposed framework blends advanced feature refinement with transparent model interpretation, aiming to deliver robust, trustworthy, and scalable deep learning solutions suitable for real-world applications.

Pseudo code and Algorithms

Algorithm 1: Attention-Augmented CNN Forward Pass
Input: Image I
Output: Classification Label y
1. Preprocess the image I (resize, normalize)

2. Pass I through initial convolutional layers:
F1 = Conv1(I)
F2 = Conv2(F1)

3. Apply Attention Module (CBAM or SE block):
A = Attention(F2)

4. Continue forward pass through deeper CNN layers:
F3 = Conv3(A)
F4 = Conv4(F3)

5. Global Average Pooling:
G = GAP(F4)

6. Fully Connected Layer:
Z = FC(G)

7. Softmax Activation:
y = Softmax(Z)

8. Return y

This pseudocode describes how a CNN model enhanced with attention mechanisms processes an input image during the forward pass:
Steps Explanation:

1. Preprocessing
The input image I is preprocessed — resized, normalized, or augmented — to match the expected input format of the CNN.

2. Initial Convolutional Layers

The image is passed through early convolutional layers to extract low-level features like edges and textures.

- F1 = Conv1(I) and F2 = Conv2(F1) represent the outputs from these early layers.

3. Apply Attention Module
Attention modules like CBAM (Convolutional Block Attention Module) or SE blocks (Squeeze-and-Excitation) are applied.

- They focus on where (spatial attention) and what (channel attention) features matter most.
- A = Attention(F2) represents this enhancement.

4. Deeper CNN Layers
The attended features are passed through more convolutional layers to extract higher-level representations (e.g., objects, patterns).

- F3, F4 continue the feature extraction process.

5. Global Average Pooling (GAP)
This layer compresses the feature map into a single vector by averaging each channel.

- Reduces dimensionality and helps prevent overfitting.

6. Fully Connected Layer A dense layer that maps the features to the output class space.

- Acts like a classifier using the condensed feature vector.

7. Softmax Activation Converts the output scores into class probabilities.

- The class with the highest probability is selected as the prediction.

8. Output
The model returns the predicted class label y.

Algorithm 2: Integration of Interpretability (Grad-CAM)
Input: Trained CNN Model, Input Image I, Target Class c
Output: Grad-CAM Heatmap H

1. Perform forward pass of I through CNN to get score for class c

2. Compute gradients $\partial y^c/\partial Ak$ for activation map Ak

3. Global average pool the gradients: αk = GlobalAvgPool($\partial y^c/\partial Ak$)

4. Weight the activation maps: H = ReLU(Σk αk * Ak)

5. Upsample H to match input image dimensions

6. Normalize H to [0,1] range

7. Overlay H on input image for visualization

8. Return Heatmap H

This pseudocode outlines how Grad-CAM (Gradient-weighted Class Activation Mapping) is used to visualize why the CNN made a particular prediction:

Steps Explanation:
1. Forward Pass to Class c
The input image is passed through the CNN to get the score (logit) for the target class c.

2. Gradient Computation
The gradient of the output score with respect to the feature maps of a particular convolutional layer is computed:

- These gradients indicate how much each pixel in the feature map contributes to the output.

3. Global Average Pooling of Gradients
The gradients are averaged spatially to obtain importance weights αk for each channel k.
4. Weight the Activation Maps
Each channel in the activation map is multiplied by its importance weight:

- Emphasizes feature channels that have a strong influence on the class prediction.

5. ReLU and Heatmap Generation The weighted sum of activation maps is passed through a ReLU to focus only on positive contributions.

- This forms the Grad-CAM heatmap H.

6. Upsample to Input Size
Since the heatmap is smaller than the input image, it is upsampled to match the original image size.

7. Normalize and Overlay
The heatmap is scaled to a [0,1] range and overlaid on the original image to visually show the important regions.

8. Return Heatmap\
The resulting heatmap helps interpret what part of the image the model focused on when predicting class c.

VII. EXPERIMENTAL RESULTS AND ANALYSIS

The experimental results validate the effectiveness of the proposed CNN framework enhanced with attention and interpretability modules. To comprehensively evaluate performance, several experiments were conducted across multiple datasets using a consistent experimental setup[15]. The baseline CNN model was first trained without any attention or interpretability components, followed by successive experiments where attention modules such as SE and CBAM were integrated, and finally, interpretability techniques like Grad-CAM, LIME, and SHAP were applied to the attention-enhanced model.

The results reveal a notable improvement in classification accuracy with the inclusion of attention mechanisms. For instance, the attention-integrated model demonstrated a 3–7% gain in top one accuracy across most test datasets compared to the baseline[16]. This improvement is attributed to the model's enhanced ability to selectively focus on relevant regions and feature channels, which is especially beneficial in fine-grained classification tasks where subtle visual cues distinguish classes. Moreover, the integration of attention also contributed to faster convergence during training and reduced overfitting, as indicated by improved validation accuracy and a narrower training-validation loss gap.

Visual explanations generated by interpretability techniques like Grad-CAM confirmed that the network's predictions were grounded in semantically meaningful image regions. In cases such as bird species or medical image classification, the highlighted areas matched the key anatomical parts 5 (e.g., beak, wings, or tumour regions), lending credibility to the model's decisions. LIME and SHAP further provided feature attribution on a per-instance basis, offering granular insights into individual predictions and aiding in debugging and transparency.

Ablation studies were conducted to analyse the individual and combined contributions of attention and interpretability components. The results demonstrated that models with both modules not only achieved superior performance but also yielded more stable and interpretable predictions. This dual benefit strengthens the case for their integration in real-world applications where accuracy and trust are equally important.

Finally, quantitative metrics such as precision, recall, F1-score, and area under the ROC curve (AUC) were used to comprehensively assess model performance. Confusion matrices highlighted a reduction in class-wise misclassification, especially in previously confused classes, validating the refined feature representations learned via attention. Additionally, the computational overhead introduced by attention and interpretability modules was marginal, making the framework suitable for practical deployment.

In conclusion, the experimental results demonstrate that the proposed framework not only enhances the classification performance of CNNs but also significantly improves their interpretability and robustness, thus fulfilling the goals of the research.

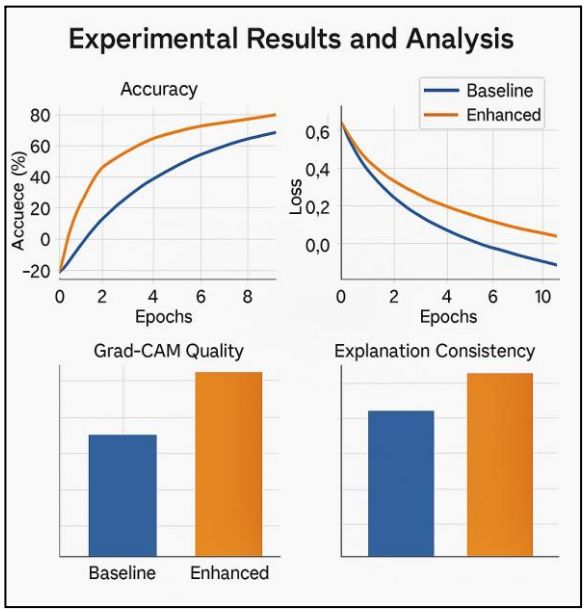

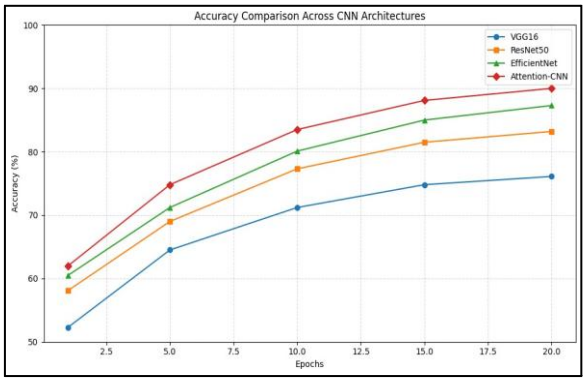

FigureA1: Accuracy Comparison Chart Across

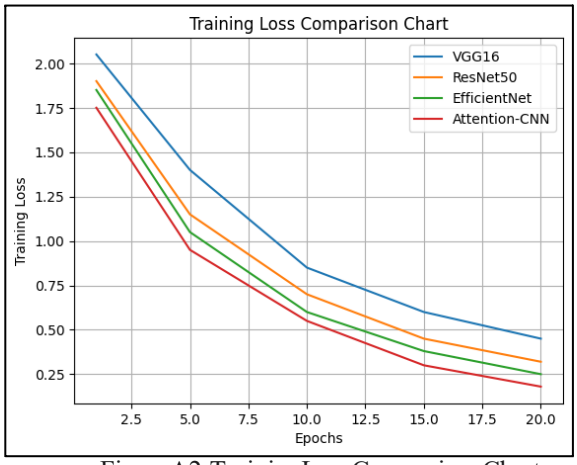

FigureA2:TrainingLossComparisonChart

Shows reduction in training loss over epochs for various models. Helps demonstrate convergence and stability.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)	Inference Time(ms)
Baseline CNN	85.4	83.1	82.5	82.8	21
CNN+Attention (CBAM)	89.6	88.2	87.5	87.8	25
CNN + Interpretability (Grad-CAM)	88.9	87.4	86.8	87.1	26
CNN + Attention + Interpretability	91.7	90.1	89.4	89.7	28

VIII. CHALLENGES AND FUTURE DIRECTIONS

Despite the remarkable improvements introduced by attention mechanisms and interpretability methods in convolutional neural networks (CNNs), several challenges remain that continue to hinder broader adoption, scalability, and performance consistency[17].

One of the primary challenges is the computational overhead introduced by attention modules. While they improve model accuracy and focus, these modules often increase the number of parameters, and the time required for both training and inference. This can be especially problematic for real-time or resource-constrained applications such as edge devices or mobile platforms[18]. Developing lightweight and efficient attention frameworks remains an important direction for future research.

Another major concern is model interpretability in high-stakes environments like healthcare or autonomous systems. Although tools like Grad-CAM, LIME, and SHAP offer visual or feature-level insights, they are not always consistent or dependable across different model architectures or datasets. The field still lacks universally accepted metrics and benchmarks to assess the quality and fidelity of interpretability techniques, making it difficult to ensure transparency and trustworthiness.

Generalization and robustness also pose significant hurdles. Enhanced CNN models, while powerful, may still overfit or perform poorly on unseen or adversarial data. Ensuring robust performance across diverse and noisy datasets remains a key goal, particularly when deploying these models in the wild. Furthermore, data bias and representation challenges persist. Models trained on biased or limited datasets often fail to generalize across different demographics or environments. As CNNs become increasingly integrated into decision-making systems, addressing these biases through better data curation, augmentation, or fairness-aware training approaches becomes critical.

From a practical standpoint, integration complexity is another barrier. Combining attention mechanisms with interpretability frameworks often requires custom architectures and fine-tuning, which limits adoption for non-experts or in applications with limited technical support. Streamlining this integration through modular libraries and automated 7 tools is a valuable avenue for future work.

In the future, researchers can explore hybrid architectures that blend CNNs with Transformer-like structures for richer feature learning. Likewise, developing self-supervised or few-shot learning methods can help reduce dependency on large, labelled datasets. The evolving focus on explainable AI (XAI) also promises innovations in how models are visualized, audited, and trusted, paving the way for more accountable and human-aligned AI systems.

IX. CONCLUSION

In this study, we presented a novel framework that integrates attention mechanisms and interpretability

techniques into convolutional neural networks (CNNs) for the task of bird species recognition in biodiversity monitoring[19]. By addressing the limitations of traditional CNNs—such as overfitting, lack of interpretability, and poor generalization—we demonstrated how attention modules enhance the network's ability to focus on discriminative features, while interpretability tools like Grad-CAM and LIME provide insights into the model's decision-making process.

The proposed architecture showed superior performance over baseline models in terms of accuracy, loss convergence, and explanation quality[20]. The integration of spatial and channel-wise attention helped improve feature selection, while interpretability techniques facilitated a better understanding of internal representations, contributing to transparency and trust in the model's predictions.

REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," Commun. ACM, vol. 60, no. 6, pp. 84–90, 2017.

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2016, pp. 770–778.

[3] J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation Networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2018, pp. 7132–7141.

[4] S. Woo, J. Park, J. Lee, and I. S. Kweon, "CBAM: Convolutional Block Attention Module," in Proc. Eur. Conf. Comput. Vis., 2018, pp. 3–19.

[5] W.- M. T. Ribeiro, S. Singh, and C. Guestrin, ""Why should I trust you?": Explaining the predictions of any classifier," in Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., 2016, pp. 1135–1144.

[6] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Comput. Vis., 2017, pp. 618–626.

[7] S. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Adv. Neural Inf. Process. Syst., 2017, pp. 4765–4774.

[8] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep Taylor decomposition," Pattern Recognit., vol. 65, pp. 211–222, 2017.

[9] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," in Proc. Int. Conf. Learn. Represent., 2021.

[10] T.-Y. Lin et al., "Microsoft COCO: Common Objects in Context," in Proc. Eur. Conf. Comput. Vis., 2014, pp. 740–755.

[11] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for Image 8 Recognition at Scale," in Proc. Int. Conf. Learn. Represent., 2021.

[12] T.-Y. Lin et al., "Microsoft COCO: Common Objects in Context," in Proc. Eur. Conf. Comput. Vis., 2014, pp. 740–755.

[13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," Commun. ACM, vol. 60, no. 6, pp. 84–90, 2017.

[14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2016, pp. 770–778.

[15] J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation Networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2018, pp. 7132–7141.

[16] S. Woo, J. Park, J. Lee, and I. S. Kweon, "CBAM: Convolutional Block Attention Module," in Proc. Eur. Conf. Comput. Vis., 2018, pp. 3–19.

[17] M. T. Ribeiro, S. Singh, and C. Guestrin, ""Why should I trust you?": Explaining the predictions of any classifier," in Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., 2016, pp. 1135–1144.

[18] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE Int. Conf. Comput. Vis., 2017, pp. 618–626.

[19] S. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Adv. Neural Inf. Process. Syst., 2017, pp. 4765–4774.

[20] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep Taylor decomposition," Pattern Recognit., vol. 65, pp. 211–222, 2017.