AMERICAN JOURNAL OF
ICT AND INNOVATION (AJICTI)

VOLUME 1 ISSUE 1 (2026)

PUBLISHED BY
E-PALLI PUBLISHERS, DELAWARE, USA

American Journal of
ICT and Innovation (AJICTI)

Volume 1 Issue 1, Year 2026
https://journals.e-palli.com/home/index.php/ajicti

Establishing the Relationship Between Jacobian Sensitivity and Perturbation Propagation in Post-Training Quantized Convolutional Neural Networks

Keya Abel Juma[1*], Michael Kimwele[1], Jael S Wekesa[1]

Article Information

Received: May 14, 2026

Accepted: July 31, 2026

Published: September 07, 2026

Keywords

Affine–Jacobian Framework, Convolutional Neural Networks, Edge Artificial Intelligence, Jacobian Sensitivity, Perturbation Propagation, Post-Training Quantization

ABSTRACT

Post-training quantization (PTQ) is widely used to deploy convolutional neural networks (CNNs) on resource-constrained edge devices because it reduces model size, memory use and computational requirements without retraining. However, PTQ introduces numerical approximation errors whose effects may differ across architectures. This study examined the relationship between Jacobian sensitivity and perturbation propagation in post-training quantized CNNs using an Affine-Jacobian framework. Six architectures, MobileNetV2, InceptionV3, DenseNet121, VGG16, NASNetMobile and EfficientNetB0, were quantized from FP32 to INT8 using TensorFlow Lite without retraining. Local sensitivity was quantified using the Frobenius norm of the Jacobian matrix, while post-quantization robustness was evaluated through progressively increasing Gaussian perturbations on a Raspberry Pi edge platform. The architectures showed substantial differences in Jacobian sensitivity and perturbation response. MobileNetV2 recorded the largest mean Jacobian norm, whereas InceptionV3 showed the largest accuracy degradation at the highest reported noise level. Exploratory architecture-level analysis nevertheless indicated a positive association between Jacobian sensitivity and degradation, although the relationship was not perfectly monotonic and, with only six architectures, was not statistically conclusive at the 0.05 level. The findings therefore support the Affine-Jacobian framework as a promising sensitivity-based explanation of quantization robustness rather than as definitive proof of a universal predictor. The study provides a basis for sensitivity-aware model selection and further validation across additional architectures and deployment conditions.

INTRODUCTION

The rapid advancement of deep learning has transformed numerous application domains, including computer vision, healthcare, intelligent transportation systems, industrial automation and precision agriculture. Among deep learning architectures, convolutional neural networks (CNNs) have emerged as the dominant models for image classification and visual recognition because of their ability to automatically learn hierarchical feature representations from large volumes of data. According to Goodfellow *et al.* (2016), CNNs have consistently achieved superior performance over traditional machine learning algorithms owing to their capacity to learn increasingly abstract visual features through multiple convolutional layers.

The widespread adoption of Edge Artificial Intelligence has further increased the demand for efficient deep learning models capable of operating on resource-constrained devices with limited computational power, memory and energy resources. Shi *et al.* (2016) observed that moving intelligence from centralized cloud servers to edge devices significantly reduces communication latency, improves data privacy and enables real-time decision-making. On the other hand, deploying modern CNNs on embedded hardware remains challenging because most state-of-the-art models contain millions of parameters that require substantial computational and storage resources.

To address these computational constraints, post-training quantization has become one of the most widely adopted model compression techniques. Rather than retraining a neural network, PTQ converts pretrained floating-point parameters into lower-precision numerical representations, thereby reducing model size and accelerating inference while preserving acceptable predictive performance. Jacob *et al.* (2018) demonstrated that integer-only quantization substantially improves deployment efficiency on embedded hardware without modifying the original network architecture. Consequently, PTQ has become a standard optimization technique for Edge AI applications because of its simplicity and computational efficiency.

Despite these advantages, quantization inevitably introduces numerical perturbations resulting from parameter approximation. These perturbations propagate through successive network layers and may alter feature representations, decision boundaries and ultimately classification accuracy. Existing studies have largely focused on measuring performance degradation after quantization while paying relatively little attention to the mathematical mechanisms responsible for perturbation propagation. Banner *et al.* (2019) argued that the impact of quantization differs considerably across neural network architectures because some models exhibit substantially greater numerical sensitivity than others. That said, the fundamental mathematical factors governing these

[1] Jomo Kenyatta University of Agriculture and Technology, Kenya
[*] Corresponding author's email: oringo.james@ku.ac.ke

variations remain insufficiently understood.

One mathematical concept that offers considerable potential for explaining these variations is Jacobian sensitivity. The Jacobian matrix characterizes how small perturbations introduced at the input are transformed as they propagate through a neural network. Jakubovitz and Giryes (2018) demonstrated that the Jacobian provides an effective measure of local network sensitivity because it quantifies the degree to which input perturbations influence network outputs. Networks with relatively large Jacobian norms tend to amplify perturbations, whereas networks with smaller norms generally exhibit greater numerical stability. While Jacobian analysis has primarily been applied to adversarial robustness, its application to understanding quantization-induced perturbation propagation remains largely unexplored.

The theoretical basis for analyzing perturbation propagation is provided by matrix perturbation theory. Stewart and Sun (1990) explained that structured perturbations propagate according to the underlying mathematical properties of a system rather than behaving as random noise. This perspective suggests that quantization errors should be interpreted as deterministic perturbations whose effects depend on local network geometry. Within convolutional neural networks employing Rectified Linear Unit (ReLU) activation functions, these local geometric properties can be represented through piecewise affine mappings, making Jacobian analysis particularly suitable for investigating perturbation behavior following post-training quantization.

Building upon these theoretical foundations, this study proposes an Affine–Jacobian framework that models post-training quantization as a deterministic structured perturbation process. Rather than viewing quantization errors as independent stochastic noise, the framework relates local numerical sensitivity to Jacobian behavior within piecewise affine regions of convolutional neural networks. By integrating mathematical analysis with experimental evaluation across multiple CNN architectures, the study examines whether Jacobian sensitivity is associated with perturbation response and can serve as a useful indicator of numerical robustness following post-training quantization.

LITERATURE REVIEW

Theoretical Foundation

This study is anchored on three complementary theories that collectively explain the relationship between Jacobian sensitivity and perturbation propagation in post-training quantized convolutional neural networks. These are Piecewise Affine Neural Network Theory, Matrix Perturbation Theory and Sensitivity Analysis Theory. Together, these theories provide the mathematical foundation for understanding how deterministic perturbations introduced during quantization propagate through deep neural networks and influence prediction accuracy.

Piecewise Affine Neural Network Theory

Piecewise Affine (PWA) Neural Network Theory provides the primary theoretical basis for analyzing perturbation propagation in deep neural networks. According to Montúfar *et al.* (2014), convolutional neural networks employing Rectified Linear Unit (ReLU) activation functions partition the input space into numerous locally linear regions. Within each region, the network behaves as an affine transformation even though the complete network remains globally nonlinear. This mathematical property enables complex deep learning models to be analyzed using linear approximation techniques while preserving their predictive capabilities.

Building on this perspective, Hanin and Rolnick (2019) explained that the remarkable expressive power of deep neural networks results from the composition of many locally affine transformations across successive network layers. As the depth of a network increases, the number of affine regions grows exponentially, enabling CNNs to approximate highly complex decision boundaries. Even so, this increasing complexity also implies that local numerical perturbations may propagate differently across regions depending on the underlying affine mapping. Consequently, two networks exposed to identical perturbations may exhibit substantially different levels of numerical stability.

The local behavior of these affine regions can be quantified through the Jacobian matrix. Jakubovitz and Giryes (2018) demonstrated that the Jacobian provides a direct mathematical representation of local network sensitivity by measuring how small variations in the input affect network outputs. Networks characterized by relatively small Jacobian norms tend to attenuate perturbations, whereas those exhibiting larger Jacobian norms amplify perturbations, increasing their susceptibility to numerical instability. This study extends this theoretical perspective by investigating whether Jacobian sensitivity similarly governs perturbation propagation following post-training quantization rather than adversarial perturbations alone.

Matrix Perturbation Theory

Matrix Perturbation Theory explains how small changes introduced into mathematical systems influence their subsequent behavior. Stewart and Sun (1990) argued that structured perturbations do not propagate randomly but instead follow deterministic relationships governed by the mathematical properties of the underlying system. Consequently, the response of a system depends not only on the magnitude of the perturbation but also on its structural characteristics and the sensitivity of the transformation through which it propagates.

This perspective is particularly relevant to post-training quantization because the conversion of floating-point parameters into low-precision integer representations introduces structured numerical perturbations rather than arbitrary random noise. Higham (2002) explained that perturbation analysis enables researchers to distinguish between stable systems, where perturbations remain bounded and unstable systems, where small numerical changes are progressively amplified. Within CNNs, this distinction implies that perturbation propagation depends largely on the local mathematical properties of the network, providing the theoretical basis for modelling quantization errors as deterministic structured perturbations.

The present study applies Matrix Perturbation Theory to examine the proposition that perturbation behavior following post-training quantization is related to local network sensitivity. Rather than interpreting quantization errors as independent stochastic events, the proposed Affine–Jacobian framework treats them as structured numerical changes whose effects may be investigated using Jacobian analysis.

Sensitivity Analysis Theory

Sensitivity Analysis Theory provides the mathematical framework for evaluating how variations in system inputs influence corresponding outputs. Saltelli *et al.* (2008) described sensitivity analysis as a systematic approach for quantifying the contribution of input variations to observed changes in model behavior. Within deep neural networks, this approach enables researchers to identify regions where predictions are either robust or highly susceptible to perturbations.

Applying this principle to deep learning, Novak *et al.* (2018) demonstrated that the Jacobian norm provides an effective quantitative measure of local sensitivity because it captures the responsiveness of network outputs to infinitesimal input perturbations. Smaller Jacobian norms indicate greater robustness and improved numerical stability, whereas larger norms signify increased perturbation amplification and reduced resilience to numerical disturbances. These characteristics make Jacobian sensitivity an appropriate metric for evaluating perturbation propagation following model quantization.

The relevance of Sensitivity Analysis Theory to this study lies in its ability to establish a measurable relationship between local network sensitivity and perturbation propagation. By combining sensitivity analysis with Piecewise Affine Neural Network Theory and Matrix Perturbation Theory, this study develops an integrated theoretical framework capable of explaining why different CNN architectures exhibit varying degrees of robustness after post-training quantization.

Empirical Literature Review

Post-Training Quantization in Convolutional Neural Networks

Post-training quantization has emerged as one of the most practical techniques for compressing deep neural networks without requiring additional model training. Jacob *et al.* (2018) demonstrated that converting pretrained floating-point models into low-precision integer representations substantially reduces memory consumption and computational complexity while maintaining acceptable predictive performance. Their work established the practical foundation for integer-only inference, making PTQ an important optimization technique for deploying CNNs on embedded and edge computing devices.

Subsequent research has shown that the effects of quantization differ considerably across network architectures. Banner *et al.* (2019) observed that some CNNs maintain high predictive accuracy following quantization, whereas others experience substantial performance degradation despite undergoing identical quantization procedures. These findings suggest that numerical robustness depends on intrinsic architectural properties rather than on the quantization algorithm alone.

Expanding this line of inquiry, Nagel *et al.* (2021) demonstrated that quantization performance is strongly influenced by activation distributions, parameter characteristics and layer-specific numerical behavior. Their findings indicate that understanding internal network sensitivity is essential for improving quantization performance. Even though these studies explain factors affecting quantization accuracy, they provide limited mathematical insight into how perturbations propagate through deep neural networks after quantization.

Jacobian Sensitivity and Perturbation Propagation

The Jacobian matrix has become an important analytical tool for investigating the stability of deep neural networks. Jakubovitz and Giryes (2018) demonstrated that Jacobian sensitivity provides a reliable indicator of network robustness because it quantifies how small perturbations introduced at the input influence the resulting network outputs. Their findings showed that networks exhibiting lower Jacobian norms generally demonstrate greater resilience to perturbations than networks characterized by higher sensitivity.

Similarly, Novak *et al.* (2018) reported that controlling Jacobian norms contributes to improved generalization and increased numerical stability in deep learning models. Their study established a direct relationship between local sensitivity and network robustness, suggesting that Jacobian analysis can provide valuable insight into the behavior of perturbations during inference. That said, their investigation focused primarily on model generalization rather than quantization-induced perturbations.

More recently, Gholami *et al.* (2022) argued that mathematical analyses of quantized neural networks should extend beyond empirical accuracy measurements to include theoretical models capable of explaining numerical behavior. Whereas the authors recognized the importance of sensitivity analysis in quantization research, they did not establish a mathematical relationship linking Jacobian sensitivity to perturbation propagation. This limitation provides an important motivation for the present study.

Research Gap

While the previous studies have significantly advanced the understanding of post-training quantization and neural network robustness, several important knowledge gaps remain. Existing research has primarily concentrated on evaluating classification accuracy, computational efficiency and compression performance after quantization, with relatively little emphasis on the mathematical mechanisms governing perturbation propagation. Zhang *et al.* (2023) noted that much of the current literature remains performance-oriented, leaving the theoretical explanation of quantization-induced perturbations insufficiently developed.

In addition, studies investigating Jacobian sensitivity have largely focused on adversarial robustness and generalization rather than deterministic perturbations introduced during post-training quantization. Consequently, there is limited theoretical understanding of whether local Jacobian sensitivity can reliably predict perturbation propagation across different CNN architectures. This study addresses this gap by proposing the Affine–Jacobian framework, which integrates Piecewise Affine Neural Network Theory, Matrix Perturbation Theory and Sensitivity Analysis Theory to establish the mathematical relationship between Jacobian sensitivity and perturbation propagation in post-training quantized convolutional neural networks.

MATERIALS AND METHODS

Research Design and Model Quantization

The study adopted a computational and experimental design to examine architecture-level sensitivity and perturbation response after post-training quantization. Six pretrained CNN architectures, MobileNetV2, InceptionV3, DenseNet121, VGG16, NASNetMobile and EfficientNetB0, were converted from FP32 to INT8 using TensorFlow Lite without retraining. Experimental validation was conducted on a Raspberry Pi edge computing platform so that the observed behavior reflected resource-constrained deployment conditions.

Sensitivity and Perturbation Evaluation

Local network sensitivity was quantified using the Frobenius norm of the Jacobian matrix. Robustness after quantization was then evaluated by applying progressively increasing Gaussian perturbations to the INT8 models and observing changes in classification accuracy. The Gaussian perturbations were used as a controlled post-quantization stress test of how each architecture responds to numerical disturbance; they were not treated as identical to the deterministic rounding error introduced by quantization itself. This distinction permits comparison of sensitivity across architectures without equating stochastic test noise with quantization error.

Architecture-Level Relationship Analysis

The relationship between mean Jacobian sensitivity and total reported accuracy degradation was evaluated at the architecture level. In addition to descriptive comparison, exploratory Pearson and Spearman associations were calculated from the six reported architecture-level summaries. Because the analysis contains only six architectures, the coefficients are interpreted as evidence of trend and effect direction rather than as definitive population-level inference.

RESULTS AND DISCUSSION

Overview of the Results

This study examined the relationship between Jacobian sensitivity and post-quantization perturbation response in convolutional neural networks (CNNs). The proposed Affine–Jacobian framework hypothesized that architectures with greater local Jacobian sensitivity would generally exhibit greater response to numerical disturbance. Six CNN architectures—MobileNetV2, InceptionV3, DenseNet121, VGG16, NASNetMobile and EfficientNetB0—were evaluated using Jacobian sensitivity analysis and progressive Gaussian perturbation experiments after INT8 quantization. The findings presented in this section are derived from the computational experiments reported in the parent study and are interpreted as architecture-level evidence rather than definitive proof of a universal predictive relationship.

Jacobian Sensitivity Across Convolutional Neural Network Architectures

The first objective of the analysis was to determine whether the selected CNN architectures exhibited different levels of Jacobian sensitivity after post-training quantization. The experimental results revealed considerable variation in the Frobenius norms of the Jacobian matrices across the evaluated models, indicating substantial architecture-level differences in local numerical sensitivity. MobileNetV2 recorded the largest Jacobian norm, whereas EfficientNetB0 exhibited the smallest Jacobian norm among the six architectures investigated.

These findings support the theoretical proposition advanced in the literature that Jacobian norms provide a quantitative measure of local network stability. According to Jakubovitz and Giryes (2018), larger Jacobian norms indicate that small perturbations introduced at the input are amplified as they propagate through successive network layers. On the other hand, smaller Jacobian norms indicate greater numerical stability because the network suppresses perturbation growth rather than amplifying it. The observed variation among the evaluated CNN architectures therefore suggests that network topology plays an important role in determining sensitivity to post-training quantization.

The experimental observations are consistent with the theoretical assumptions of Piecewise Affine Neural Network Theory. Since convolutional neural networks contain numerous locally affine regions, differences in architecture can produce different local Jacobian geometries. Consequently, comparable numerical disturbances may propagate differently depending on the local affine transformations encountered during inference. The observed architecture-level variation therefore provides empirical evidence consistent with the

Table 1: Comparison of Theoretical Jacobian Sensitivity and Observed Accuracy after PTQ.

CNN Architecture	Mean Jacobian Norm	Predicted Stability	FP32 Accuracy (%)	INT8 Accuracy (%)	Accuracy Difference (%)
MobileNetV2	1367.48	Highly Sensitive	92.07	91.44	−0.63
InceptionV3	541.04	Sensitive	92.85	93.14	+0.28
DenseNet121	456.11	Moderately Stable	90.00	90.00	0.00
VGG16	226.50	Moderately Stable	90.00	90.00	0.00
NASNetMobile	77.08	Stable	90.00	90.00	0.00
EfficientNetB0	39.12	Stable	90.00	90.00	0.00

Source: Researcher (2026).

Affine–Jacobian framework proposed in this study, while broader validation is required to establish its generality.

Perturbation Propagation Following Post-Training Quantization

The second stage examined how the quantized CNN architectures responded to controlled perturbations after post-training quantization. Progressively increasing Gaussian perturbations were applied to the INT8 models and the resulting changes in predictive performance were observed. The experiment was designed as a sensitivity stress test rather than as a direct simulation of the deterministic rounding error generated by PTQ. The models differed markedly in their response as perturbation magnitude increased.

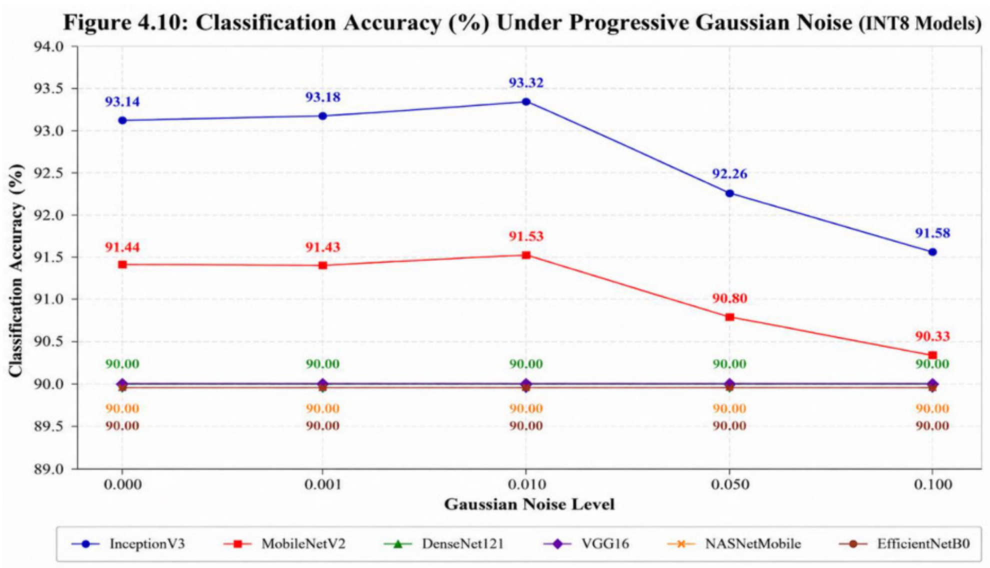

Figure 1: Classification Accuracy (%) Under Progressive Gaussian Noise in INT8 Models.
Source: Researcher (2026)

Among the evaluated models, InceptionV3 recorded the largest total accuracy degradation at the highest reported noise level (1.55%), followed by MobileNetV2 (1.11%). DenseNet121, VGG16 and EfficientNetB0 showed no reported degradation at that level, while the very small increase reported for NASNetMobile was practically negligible. The results therefore show that post-quantization perturbation response differs across architectures. Importantly, the ordering is not perfectly identical to the ordering of Jacobian norms, so the evidence should be interpreted as an overall sensitivity trend rather than a strictly monotonic relationship.

The observed differences are consistent with the broader premise of Matrix Perturbation Theory that a system's response to disturbance depends on its intrinsic mathematical properties. In this experiment, controlled Gaussian perturbations applied after quantization produced architecture-specific responses, indicating that

Table 2: Overall Accuracy Degradation Under Progressive Gaussian Noise.

CNN Architecture	Baseline Accuracy (%)	Accuracy at Noise = 0.100 (%)	Total Accuracy Degradation (%)	Robustness Rank
DenseNet121	90.00	90.00	0.00	1
VGG16	90.00	90.00	0.00	1
EfficientNetB0	90.00	90.00	0.00	1
NASNetMobile	90.00	90.00	−0.00*	2
MobileNetV2	91.44	90.33	1.11	5
InceptionV3	93.14	91.58	1.55	6

*Source: Researcher (2026). *The slight increase for NASNetMobile (0.0049%) was reported as practically negligible in the thesis.*

numerical robustness is influenced by network structure and local sensitivity. Because the applied Gaussian noise is a stochastic stress test rather than the deterministic PTQ rounding error itself, the results support comparison of robustness across architectures but should not be interpreted as direct proof that quantization error propagates identically to Gaussian noise.

Relationship Between Jacobian Sensitivity and Perturbation Propagation

The principal objective was to determine whether Jacobian sensitivity is associated with post-quantization perturbation response. Architecture-level comparison showed a generally positive pattern: models with larger mean Jacobian norms tended to exhibit greater accuracy degradation, whereas several models with smaller norms remained stable under the reported perturbation level. The pattern was not perfectly monotonic, however, because InceptionV3 showed greater degradation than MobileNetV2 despite having a substantially smaller mean Jacobian norm.

An exploratory analysis of the six architecture-level summary values yielded a Pearson correlation of $r = 0.672$ ($p = 0.144$) between mean Jacobian norm and total accuracy degradation. The corresponding Spearman rank correlation was rho = 0.759 ($p = 0.080$). Both coefficients indicate a positive moderate-to-strong association in direction, but neither reaches the conventional 0.05 significance threshold with only six architectures. Accordingly, the results support Jacobian sensitivity as a potentially useful indicator of robustness, while additional architectures and repeated observations are required before it can be regarded as a statistically established predictor.

The overall correspondence between Jacobian sensitivity and perturbation response is therefore supportive, but not conclusive, evidence for the proposed Affine-Jacobian framework. The framework offers a plausible mathematical explanation for architecture-dependent robustness by linking local network geometry to sensitivity. The observed exception in the ordering of MobileNetV2 and InceptionV3 also indicates that Jacobian magnitude is unlikely to be the only factor governing robustness and that activation distributions, layer characteristics and other architectural properties may contribute.

Evaluation of the Affine-Jacobian Framework

A key contribution of the study was the empirical evaluation of the proposed Affine-Jacobian framework across six CNN architectures. The framework predicted that greater local Jacobian sensitivity would generally be associated with greater perturbation amplification. The architecture-level results broadly followed this direction, although the relationship was not exact across all models and the small number of architectures limits statistical certainty.

These findings suggest that deterministic sensitivity analysis can complement conventional empirical accuracy measurements when studying PTQ robustness. Higham (2002) emphasized the value of perturbation analysis for understanding structured numerical disturbances. In the present study, Jacobian analysis provides a mathematically grounded lens for interpreting architecture-dependent sensitivity, while the Gaussian experiments serve as a separate controlled stress test. Further work should directly measure layer-wise quantization error and its propagation to determine how closely those deterministic errors follow the sensitivity relationships predicted by the framework.

Discussion of Findings

The findings contribute to the literature on post-training quantization by showing that architecture-level numerical sensitivity is relevant to perturbation robustness. Previous PTQ studies have largely emphasized accuracy and efficiency outcomes; the present analysis adds a sensitivity-based explanation for part of the observed variation among architectures. The positive association between Jacobian norm and degradation supports this interpretation, but the non-monotonic ordering of

MobileNetV2 and InceptionV3 shows that perturbation response is not determined by Jacobian magnitude alone. Second, the study extends Jacobian sensitivity analysis from its established use in adversarial robustness and generalization to the evaluation of quantized models under controlled numerical perturbation. The architecture-level results indicate that Jacobian analysis may be useful for screening models for sensitivity before Edge AI deployment. However, because the exploratory correlations are based on six architectures and are not statistically significant at the 0.05 level, the framework should be considered promising rather than fully validated.

Finally, the Affine-Jacobian framework has practical implications for Edge Artificial Intelligence because sensitivity measures can be incorporated into model evaluation alongside conventional accuracy, latency and memory criteria. Developers may use Jacobian sensitivity as one diagnostic indicator when selecting models for quantization, while researchers can test whether layer-wise sensitivity-aware quantization improves robustness across larger model sets and hardware platforms.

CONCLUSION

This study examined whether Jacobian sensitivity is associated with perturbation response in six post-training quantized CNN architectures. The models were converted from FP32 to INT8, Jacobian sensitivity was measured using the Frobenius norm and robustness was assessed through progressive Gaussian perturbations applied after quantization. The results showed substantial architectural variation. MobileNetV2 recorded the largest mean Jacobian norm, while InceptionV3 showed the greatest reported accuracy degradation at the highest noise level. Overall, higher Jacobian sensitivity tended to be associated with greater degradation, although the relationship was not perfectly monotonic. Exploratory Pearson ($r = 0.672$, $p = 0.144$) and Spearman ($\rho = 0.759$, $p = 0.080$) analyses indicated a positive trend but were not statistically conclusive with only six architectures. The Affine-Jacobian framework should therefore be regarded as a promising explanatory approach requiring broader validation. Developers should use Jacobian sensitivity as one diagnostic criterion in model selection, while future studies should directly measure layer-wise quantization error across more architectures, quantization schemes and edge hardware platforms.

REFERENCES

Banner, R., Nahshan, Y., Hoffer, E., & Soudry, D. (2019). Post training 4-bit quantization of convolutional networks for rapid-deployment. *Advances in Neural Information Processing Systems, 32*, 7948–7956.

Creswell, J. W., & Creswell, J. D. (2023). *Research design: Qualitative, quantitative and mixed methods approach* (6th ed.). Sage Publications.

Gholami, A., Kim, S., Dong, Z., Yao, Z., Mahoney, M. W., & Keutzer, K. (2022). *A survey of quantization methods for efficient neural network inference. In L. O. Hall, K. Bowyer, D. Goldgof, P. Phillips, & J. Sarkar* (Eds.), Low-Power Computer Vision: Improve the Efficiency of Artificial Intelligence (pp. 291–326). Chapman and Hall/CRC.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.

Hanin, B., & Rolnick, D. (2019). Complexity of linear regions in deep networks. *Proceedings of the 36th International Conference on Machine Learning, 97*, 2596–2604.

Higham, N. J. (2002). *Accuracy and stability of numerical algorithms* (2nd ed.). Society for Industrial and Applied Mathematics.

Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., & Kalenichenko, D. (2018). Quantization and training of neural networks for efficient integer-arithmetic-only inference. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2704–2713.

Jakubovitz, D., & Giryes, R. (2018). Improving DNN robustness to adversarial attacks using Jacobian regularization. *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 514–529.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). *Deep learning. Nature, 521*(7553), 436–444.

Montúfar, G. F., Pascanu, R., Cho, K., & Bengio, Y. (2014). On the number of linear regions of deep neural networks. *Advances in Neural Information Processing Systems, 27*, 2924–2932.

Nagel, M., Amjad, R. A., van Baalen, M., Louizos, C., & Blankevoort, T. (2021). Up or down? Adaptive rounding for post-training quantization. *Proceedings of the 38th International Conference on Machine Learning, 139*, 7197–7206.

Novak, R., Bahri, Y., Abolafia, D. A., Pennington, J., & Sohl-Dickstein, J. (2018). Sensitivity and generalization in neural networks: An empirical study. *International Conference on Learning Representations* (ICLR).

Saltelli, A., Ratto, M. Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). *Global sensitivity analysis: The primer. John Wiley & Sons.*

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.

Stewart, G. W., & Sun, J.-G. (1990). Matrix perturbation theory. Academic Press.

Zhang, X., Liu, Y., Wang, H., & Chen, Z. (2023). A comprehensive review of post-training quantization for deep neural networks. *ACM Computing Surveys, 56*(2), 1–36.