

ZERO-DAY MALWARE ROBUSTNESS IN ANDROID: A MULTI-VIEW DEEP LEARNING MODEL WITH INCREMENTAL FINE-TUNINGShaik Salma Begum^{1*}, Shaik Mahammad Riyaz¹, Rayala Bhargav Sai¹, Pagadala Naga Sri Lakshmi¹, Sureddy Surya¹¹ Dept. of CSE (AI&ML), S R Gudlavalleru Engineering College, Gudlavalleru- 521356, Andhra Pradesh, India
(shaiksalma.gec@gmail.com, mdriyazrafi786@gmail.com, rayalabhargavsai@gmail.com,
nagasrilakshmipagadala@gmail.com, sureddysurya123@gmail.com)Corresponding Author Email: shaiksalma.gec@gmail.com**ABSTRACT**

Android malware detection systems based on static analysis and machine learning often report high accuracy when evaluated on datasets collected within the same time period. However, their robustness under long-term malware evolution remains insufficiently studied. In this paper, a longitudinal evaluation of a multi-view deep neural network (MV-DNN) trained on the Drebin dataset (2012–2014) is performed and tested on CIC-AndMal2020 (2020), representing an eight-year temporal gap. Experimental results show severe performance degradation, where the Drebin-trained model achieves only 45.22% accuracy on CIC-AndMal2020, indicating poor generalization under temporal drift. To mitigate this issue, a lightweight transfer learning strategy is proposed that partially fine-tunes only the last layers of the pretrained MV-DNN using a small labeled calibration set of 1,000 CIC samples. The proposed approach significantly improves detection performance, achieving 87.11% accuracy with balanced precision and recall on the full set of CIC-AndMal2020 tests. The computational cost of the proposed adaptation relative to full-model retraining is quantified and discusses the statistical reliability and limitations of the reported improvements. These results demonstrate that limited target-domain supervision combined with partial fine-tuning can effectively restore malware detection performance under long-term dataset shift.

Keywords: Android Malware Detection, Calibration Subset, CIC-AndMal2020, Dataset Shift, Domain Adaptation, Drebin, Multi-View Deep Neural Network, Temporal Drift, Transfer Learning.

1. INTRODUCTION

The increasing popularity of the Android ecosystem has enabled smartphones to become one of the primary modes of communication, finance, healthcare, and enterprise services. However, the widespread use of the Android ecosystem to offer various smartphone-based services has made it a growing target for malicious apps. Malware on the Android platform has undergone tremendous transformation in the past ten years, as malware authors make use of sophisticated evasion techniques such as “code obfuscation” and “behavioral manipulation.” Android malware detection techniques are generally categorized into static analysis, dynamic analysis, and hybrid approaches [1]. Of these, static analysis is the most popular technique as it has low computational cost and is scalable, and is used to extract features like permissions, API calls, intents and opcode-level patterns of the application without running it. There have been several representative systems proposed such as Drebin [2], DroidAPIMiner [3], MaMaDroid [4] and MADAM [5] which are behavior-based. Application of deep learning to Android malware detection has also been recently shown, with automatic feature representation learning and high detection accuracy [6] [7].

Though these advances have been made, one of the critical dimensions that have not been explored is temporal robustness [8]. Most malware classifiers are trained and tested with a timeframe similar to that used in this paper. Malware

continually changes in the real world, and the statistical distribution of features over time [9]. This is known as temporal drift or dataset shift, and can result in drastic performance drops for more recent malware when using a model trained on older malware. Thus, it is crucial to understand the reliability of Android malware detection systems in long-term deployment, and evaluating cross-time generalization is vital [10].

To investigate this issue, this article performs a longitudinal evaluation of a multi-view deep neural network (MV-DNN) trained and tested on CIC-AndMal2020 [11], representing an eight-year temporal gap. Our baseline results reveal a severe degradation in detection capability, where the Drebin-trained model achieves only 45.22% accuracy and a ROC-AUC of 0.558 on CIC-AndMal2020 [12]. To mitigate this temporal performance drop, a lightweight transfer learning strategy is proposed by fine-tuning the final layers of the pretrained MV-DNN using a small labeled calibration subset from the target dataset. This adaptation approach significantly improves cross-dataset performance, achieving 87.11% accuracy and ROC-AUC of 0.932 on CIC-AndMal2020.

1.1 Research contributions

The contributions of this work are threefold. First, we present a longitudinal, cross-temporal evaluation of a multi-view deep neural network across an eight-year dataset gap (Drebin, 2012–2014, to CIC-AndMal2020, 2020), quantifying how severely static-feature Android malware detectors degrade

under realistic long-term deployment conditions (45.22% accuracy, ROC-AUC 0.558). Second, we propose a lightweight partial fine-tuning strategy that freezes the pretrained feature-extraction branches and adapts only the classification head using a small, balanced calibration subset of 1,000 target-domain samples, avoiding the cost and data requirements of full model retraining while mitigating catastrophic forgetting [13]. Third, this adaptation strategy is tested on the full CIC-AndMal2020 test set and shown to recover detection performance to 87.11 accuracy and 0.932 ROC-AUC, its computational and data efficiency compared with full retraining, its statistical reliability and its limitations are examined.

2. RELATED WORKS

2.1 Static analysis-based malware detection

Due to its efficiency and scalability, static analysis has been used to detect Android malware for a long time. A recent work by Drebin [2] showed that a comprehensive set of static features (permissions, API calls and network indicators) can lead to effective and interpretable detection via linear classifiers. Likewise, DroidAPIMiner [3] used API-level frequency features to boost the adversarial ability to withstand basic obfuscation techniques. Other static methods included using parallel classifiers [14], semantic feature engineering [15] and lightweight detection strategies that are able to resist obfuscation [16].

MaMaDroid [4] used behavioral modeling techniques to improve generalization over the purely frequency based representations of the API call sequences by using a Markov chain. A significant number of these methods depend critically on manually designed features, and are usually tested only in one-dataset environments, which prevents the use of them to gain insights into cross-dataset or temporal robustness.

2.2 Deep learning-based approaches

In order to minimize reliance on manual feature engineering, deep learning models have been suggested for automated representation learning. McLaughlin et al. [6] showed that deep neural networks were able to categorize Android malware by only analyzing raw static features. Later, multimodal deep architectures were studied, e.g., adversarial learning frameworks like DANDroid [17] and multimodal deep models that combine heterogeneous feature sources [18].

The classifiers of deep architectures like DroidDetector [7] and DeepRefiner [19] comprised multiple neural layers to model complex feature interactions and obtain better classification performance [20]. Classifier based on neural networks have also been used in other areas, like medical image analysis and voice-based harassment and fraudulent call protection and cyber threat detections [21].

2.3 Benchmark datasets and temporal challenges

Despite many other datasets being created, the Drebin data set continues to be one of the most commonly used data sets for static malware detection research [22]. Recently, to solve the problems of previous data sets, CIC-AndMal2020 [11] proposed a modern and systematically-generated benchmark. Lashkari et al. [23] reiterate the importance of consistently generating structured datasets and labelling for reproducible research.

However, there is a lack of cross-temporal and cross-dataset evaluations despite these datasets. Concept drift can cause degradation in the performance of a model based on an older dataset like Drebin when it is run on newer malware samples. In the field of machine learning, the issue of concept drift and distribution shift is well explored [10] and thus there is a need for adaptive strategies. In malware detection, it has been demonstrated that ignoring temporal ordering of samples can lead to an over-estimate of the performance [8, 9] and drift-aware and continual-learning detectors have been proposed to overcome the resulting decay [24, 25, 26].

2.4 Research gaps

So far, it has been shown that classical approaches and deep learning methods achieve good results on the same dataset. Focusing on limited aspects of temporal robustness, lightweight adaptation and transfer learning between legacy sets (e.g. Drebin [2]) and contemporary sets (e.g. CIC-AndMal2020 [11]) has not received much attention. To meet this need, architectures must have the ability to retain the knowledge gained through security and adjust to the changing nature of malware.

2.5 Comparative summary of prior work

Table 1 summarizes representative prior studies alongside the dataset, method, evaluation setting, key findings, and the specific limitation or research gap each leaves open, to situate the present work relative to the literature.

Table 1. Comparative summary of representative prior studies

Study	Dataset	Method / model	Evaluation setting	Key findings	Limitation / gap
Arp et al. [2] (Drebin)	Drebin (2012–2014)	Linear SVM on static features	Single-dataset, same period	High accuracy, interpretable, low cost	No cross-time / cross-dataset evaluation
Aafer et al. [3] (DroidAPIMiner)	Contemporary malware corpus	API-frequency features + KNN	Single-dataset	Robust to simple obfuscation	Not tested under temporal drift
Mariconti et al. [4] (MaMaDroid)	Contemporary malware corpus	Markov-chain API abstraction	Single-dataset	Better generalization than frequency features	Handcrafted; limited temporal analysis
McLaughlin et al. [6]	Raw APK opcodes	CNN on raw static features	Single-dataset	Automated representation learning improves accuracy	No temporal / cross-dataset study
Millar et al. [17] (DANDroid)	Contemporary malware corpus	Adversarial multi-view network	Single-dataset	Improved robustness to obfuscation	No long-term drift evaluation

Study	Dataset	Method / model	Evaluation setting	Key findings	Limitation / gap
Kim et al. [18]	Multi-source static/dynamic features	Multimodal deep network	Single-dataset	Heterogeneous fusion improves accuracy	No cross-temporal generalization
Su et al. [11] (CIC-AndMal2020)	CIC-AndMal2020 (2020)	Benchmark construction	Dataset paper	Modern, systematically labeled benchmark	Not paired with legacy datasets for drift study
Gama et al. [10]	General ML drift literature	Concept-drift survey	Theoretical / survey	Formalizes drift and adaptation strategies	Not applied to Android malware detection

Compared with these studies, the present work is, to our knowledge, the first to explicitly quantify the accuracy and ROC-AUC drop across an eight-year Drebin-to-CIC-AndMal2020 gap and to pair that measurement with a lightweight, target-domain-efficient partial fine-tuning strategy rather than a full retraining pipeline, directly addressing the research gap identified above.

3. DATASETS

In order to test how robust Android malware detection could be in dealing with long-term malware evolution, we conduct our experiments by using two datasets collected at different times. This simulates a reasonable temporal generalization setting, i.e., training on historical malware data and testing on current malware data. Table 2 summarizes the two datasets.

3.1 Drebin dataset (2012–2014)

Among publicly available datasets, Drebin [2] remains one of the most frequently utilized. It contains Android applications collected during 2012–2014 and provides static feature representations extracted from application manifests and disassembled code. Drebin includes multiple feature categories, such as API calls, permissions, intents, and network-related strings.

3.2 CIC-AndMal2020 dataset (2020)

To evaluate temporal drift, we use CIC-AndMal2020 [11], [27], a modern Android malware dataset collected around 2020. This dataset is treated as the target domain. Since CIC-AndMal2020 is temporally separated from Drebin by approximately eight years, it represents a challenging and realistic scenario where a detector trained on older malware is applied to newer applications.

In our experiments, CIC-AndMal2020 contains 357,805 samples, including 162,181 benign applications and 195,624 malware samples. Of these, a balanced subset of 1,000 samples (500 benign, 500 malware) is randomly drawn to form the calibration set used exclusively for fine-tuning the classification head, and the remaining 356,805 samples (357,805 – 1,000) form the independent test set, which is never seen during fine-tuning and is used only for the final evaluation reported in Section 7. This calibration/test split is applied identically to the baseline (no-adaptation) and fine-tuned evaluations so that the two settings are directly comparable.

For CIC-AndMal2020, we utilize the preprocessed static feature version provided in parquet format. The feature columns are anonymized and stored as F0 to F214, where each

feature corresponds to one element in the aligned multi-view representation.

Table 2. Summary of datasets used in this study

Property	Drebin	CIC-AndMal2020
Year introduced	2014	2020
Source	Real-world Android apps	Systematically generated benchmark
Analysis type	Static analysis	Static analysis
Feature representation	215 static features	Aligned to 215 features
Usage in this work	Pretraining	Fine-tuning and evaluation
Calibration subset	–	1,000 (500 benign / 500 malware)
Test samples	–	356,805
Purpose	Base knowledge learning	Temporal adaptation and testing

4. PROPOSED METHODOLOGY

4.1 Problem definition

Let $D^s = \{(x_i^s, y_i^s)\}$ represent the source dataset (Drebin, 2012–2014) and $D^t = \{(x_i^t, y_i^t)\}$ represent the target dataset (CIC-AndMal2020, 2020). Each sample x is represented using static analysis features and the label $y \in \{0, 1\}$ indicates benign or malware.

The goal of this work is to evaluate temporal robustness under an eight-year time gap and to design a lightweight adaptation strategy that improves performance on D^t while reusing the pretrained model trained on D^s .

4.2 Multi-view feature representation

Android applications can be described through multiple complementary static views. Following prior work on static Android malware detection [2], [3], we represent each application using three views:

- Permissions view (x_p): a binary vector indicating requested Android permissions.
- API calls view (x_a): a vector representing API usage extracted from application code.
- Intent filters view (x_i): a vector indicating intent-related components and actions.

The final input representation is expressed in Eq. (1):

$$X = \{x_p, x_a, x_i\} \quad (1)$$

where $x_p \in R^{113}$, $x_a \in R^{74}$, and $x_i \in R^{28}$. In our experiments, the combined feature dimension is 215 (113 + 74 + 28).

4.3 Multi-view deep learning model

We design a multi-view deep neural network architecture where each view is processed by an independent branch. Each branch consists of two fully connected layers, as given in Eq. (2):

$$h_p = f_p(x_p), h_a = f_a(x_a), h_i = f_i(x_i) \quad (2)$$

Each branch output is then fused using concatenation, as given in Eq. (3):

$$h = [h_p \parallel h_a \parallel h_i] \quad (3)$$

Finally, the fused representation is passed through a classification head, as given in Eq. (4):

$$\hat{Q} = \sigma(g(h)) \quad (4)$$

Here $\sigma(\cdot)$ represents the sigmoid activation function and $\hat{Q}$ represents the probability of malware. The MV-DNN is trained using the binary cross-entropy loss given in Eq. (5). The overall workflow of the proposed methodology is shown in Figure 1.

$$G = -\left(\frac{1}{C}\right) \sum_{v=1}^C [Q_v \log(\hat{Q}_v) + (1 - Q_v) \log(1 - \hat{Q}_v)] \quad (5)$$




Figure 1. Workflow of the proposed methodology

4.4 Baseline temporal evaluation (Drebin → CIC)

To quantify temporal degradation, we first train the MV-DNN on Drebin and evaluate it directly on CIC-AndMal2020 without any adaptation. Although the model achieves high performance on Drebin (approximately 98% accuracy), the accuracy drops to 45.22% on CIC-AndMal2020. This confirms a severe temporal dataset shift where the pretrained model fails to generalize to modern malware.

4.5 Lightweight domain adaptation via partial fine-tuning

To mitigate temporal drift, we propose a lightweight adaptation strategy based on partial fine-tuning, which builds on the transferability of features learned by pretrained networks [28]. Instead of retraining the entire MV-DNN, we reuse the Drebin-trained model and adapt only the final layers

using a small labeled calibration subset from CIC-AndMal2020.

Let θ represent all parameters of the pretrained MV-DNN. We divide the parameters into two groups, as given in Eq. (6):

$$\theta = \{\theta_{\text{base}}, \theta_{\text{head}}\} \quad (6)$$

where θ_{base} corresponds to the feature extraction branches and fusion layers, and θ_{head} corresponds to the final classification layers. During adaptation, we freeze θ_{base} and fine-tune only θ_{head} using the calibration set, as given in Eq. (7):

$$\theta_{\text{base}} \leftarrow \text{fixed}, \theta_{\text{head}} \leftarrow \underset{\theta_{\text{head}}}{\text{argmin}} L(D_{\text{cal}}^t) \quad (7)$$

In our experiments, the calibration set contains 1,000 labeled samples (500 malware and 500 benign). We then progressively unfreeze the last layers of the network and fine-tune with a low learning rate to improve generalization while avoiding catastrophic forgetting.

4.6 Final evaluation

After adaptation, the fine-tuned model is evaluated on the full CIC-AndMal2020 dataset. The proposed strategy significantly restores detection performance and improves the accuracy to 87.11%, demonstrating that lightweight fine-tuning can effectively address long-term malware evolution.

5. MULTI-VIEW DEEP NEURAL NETWORK ARCHITECTURE

5.1 Input layer

The proposed model employs three parallel input layers corresponding to heterogeneous static feature groups extracted from Android applications:

- Permissions input: $x_p \in R^{113}$
- API calls input: $x_a \in R^{74}$
- Intents input: $x_i \in R^{28}$

Each input layer receives a binary feature vector representing the presence or absence of specific static attributes. These layers do not contain trainable parameters.

5.2 View-specific feature learning branches

Each feature group is processed independently through a dedicated dense transformation branch to learn specialized latent representations.

a) Permission branch: Dense layer with 128 neurons, ReLU activation, 14,592 trainable parameters, followed by Dropout [29] for regularization, as given in Eq. (8):

$$h_p = \text{Dropout}(\text{ReLU}(W_p \cdot x_p + b_p)) \quad (8)$$

b) API branch: Dense layer with 128 neurons, ReLU activation, 9,472 trainable parameters, followed by Dropout, as given in Eq. (9):

$$h_a = \text{Dropout}(\text{ReLU}(W_a \cdot x_a + b_a)) \quad (9)$$

c) Intent branch: Dense layer with 64 neurons, ReLU activation, 1,856 trainable parameters, followed by Dropout, as given in Eq. (10):

$$h_i = \text{Dropout}(\text{ReLU}(W_i \cdot x_i + b_i)) \quad (10)$$

The independent branches enable the network to extract view-specific representations, thereby reducing cross-view interference.

5.3 Fusion layer

The learned embeddings are fused using feature concatenation, $h = [h_p \parallel h_a \parallel h_i]$ (Eq. (3)). The resulting fused representation has dimension $128 + 128 + 64 = 320$. This fusion layer integrates complementary structural and behavioral signals from all feature groups.

5.4 Classification head

The fused representation is passed into a fully connected layer consisting of a Dense layer with 128 neurons, a ReLU activation, 41,088 trainable parameters, followed by Dropout, as given in Eq. (11):

$$h_f = \text{Dropout}(\text{ReLU}(W_f h + b_f)) \quad (11)$$

This layer captures cross-view interactions and higher-level abstractions.

5.5 Output layer

The final prediction layer consists of a Dense layer with 1 neuron and a sigmoid activation, with 129 trainable parameters, as given in Eq. (12):

$$\hat{Q} = \sigma(W_o h_f + b_o) \quad (12)$$

where $\hat{Q} \in [0,1]$ represents the probability that the input application is malicious. The complete architecture is shown in Figure 2.




Figure 2. Model architecture

6. EXPERIMENTAL SETUP

6.1 Implementation details

All programs used Python as their programming language and TensorFlow/Keras as their model development framework, also NumPy and Scikit-learn for data preprocessing and evaluation tasks. Google Colab was used for training and

evaluation because of the extensive size of the CIC-AndMal2020 dataset.

Reproducibility details: Hardware/GPU configuration—GPU/TPU model and VRAM used in Google Colab, e.g., NVIDIA T4/A100, or CPU. Software versions—Python [version], TensorFlow, NumPy, Scikit-learn, pandas/pyarrow. Random seed—a fixed seed of was used for data splitting, weight initialization, and calibration-subset sampling. Data protocol—the Drebin set was split into training/validation ([80]/[20]%) for base-model training; the 1,000-sample calibration subset was itself split into fine-tuning/validation ([80]/[20]%) to monitor for overfitting during adaptation, with the remaining 356,805 samples held out strictly for final testing. Number of runs—each configuration (baseline and fine-tuned) was repeated across [20] independent runs with different random seeds (see Section 7.5). Hyperparameter selection—the learning-rate range (10^{-4} to 5×10^{-4}) and epoch range (15–25) were selected stopping when validation loss on the held-out calibration split stopped improving for [25] consecutive epochs.

6.2 Baseline model training (Drebin)

The proposed MV-DNN is initially trained on the Drebin dataset using the 215-dimensional static feature representation. The model performs binary classification (benign / malware) optimized with Adam [30] and binary cross-entropy loss.

6.3 Cross-dataset evaluation (Drebin → CIC)

To assess temporal generalization, the Drebin-trained model is directly evaluated on CIC-AndMal2020 without adaptation. CIC samples are mapped to the fixed 215-feature space. Performance is measured using the confusion matrix and other metrics.

6.4 Calibration and fine-tuning

A balanced calibration subset of 1,000 samples (500 benign, 500 malware) is randomly selected from CIC-AndMal2020. This subset is used exclusively for adaptation.

Fine-tuning is performed in a staged manner: early feature extraction layers are frozen, and only higher layers are updated. Subsequently, selected layers are progressively unfrozen and trained with a reduced learning rate to adapt to modern malware characteristics while mitigating catastrophic forgetting.

6.5 Training configuration

Fine-tuning uses the Adam optimizer with a learning rate between 10^{-4} and 5×10^{-4} , a batch size of 32, and 15–25 epochs. An output layer with binary cross-entropy loss is employed.

6.6 Evaluation protocol

The final performance is reported on the complete CIC-AndMal2020 test set using accuracy, precision, recall, F1-score, and other evaluation metrics.

7. RESULTS AND DISCUSSIONS

7.1 Baseline temporal transfer (Drebin → CIC)

We first evaluate the Drebin-trained MV-DNN on CIC-AndMal2020 without any adaptation. Although the model achieves approximately 98% accuracy on Drebin, its performance drops sharply on CIC-AndMal2020 due to the eight-year temporal gap. The baseline transfer accuracy is 45.22% (Figure 3), indicating severe temporal dataset shift.




Figure 3. Baseline confusion matrix (Drebin → CIC)

7.2 Performance of the proposed adaptation

To mitigate temporal drift, we apply a lightweight fine-tuning strategy using a small labeled calibration set from CIC-AndMal2020. The adapted model significantly improves detection performance on the target dataset. The final fine-tuned model achieves 87.11% accuracy on the full CIC-AndMal2020 test set (Figure 4). The ROC-AUC improves to 0.932, demonstrating strong separability between benign and malware samples after adaptation.




Figure 4. Confusion matrix after proposed fine-tuning on CIC-AndMal2020

This confusion matrix displays the results of the fine-tuned model evaluation conducted on CIC-AndMal2020. The results indicate that the proposed approach successfully restores malware detection performance while maintaining a low false-positive rate.

7.3 Overall comparison

Table 3 summarizes the baseline and proposed performance on CIC-AndMal2020. Compared to the baseline transfer evaluation, the proposed approach improves accuracy from 45.22% to 87.11%, and increases ROC-AUC from 0.558 to 0.932. These results confirm that temporal drift is a critical challenge for Android malware detection, and that lightweight fine-tuning can effectively adapt a pretrained model to modern malware distributions. Figures 5 and 6 compare the two settings in terms of per-metric performance and ROC curves, respectively, and Figure 7 summarizes the impact of the temporal gap and the proposed adaptation.

Table 3. Baseline vs. proposed performance on CIC-AndMal2020

Metric	Baseline	Proposed
Accuracy	45.22%	87.11%
Precision (malware)	15.0%	91.17%
Recall (malware)	0.0%	84.64%
F1-score (malware)	0.0%	87.78%
ROC-AUC	0.558	0.932




Figure 5. Comparative accuracy analysis: baseline Drebin model vs. fine-tuned MV-DNN on modern CIC-AndMal2020 samples




Figure 6. Performance enhancement of the proposed model over the baseline, achieving a ROC-AUC of 0.932




Figure 7. Impact of temporal gap and proposed adaptation

7.4 Computational complexity analysis

Table 4 summarizes the trainable-parameter budget of the MV-DNN and its partitioning under the proposed partial fine-tuning strategy. The network contains 67,137 trainable parameters in total. Full retraining would require gradient computation and weight updates across all 67,137 parameters on the full CIC-AndMal2020 training split, plus careful replay of Drebin data to avoid catastrophic forgetting. In contrast, the proposed strategy freezes the three view-specific branches (25,920 parameters, $\approx 38.6\%$ of the network) and updates only the classification head and output layer (41,217 parameters, $\approx 61.4\%$ of the network) using the 1,000-sample calibration set. Because the frozen branches require no gradient computation with respect to their own weights, adaptation needs orders of magnitude fewer labeled target-domain samples than would be required to retrain the full network reliably.

Table 4. Parameter budget and complexity of the proposed partial fine-tuning strategy

Component	Trainable parameters	Status under fine-tuning
Permission branch (Dense-128)	14,592	Frozen (θ_{base})
API-call branch (Dense-128)	9,472	Frozen (θ_{base})
Intent branch (Dense-64)	1,856	Frozen (θ_{base})
Fusion (concatenation)	0	N/A (no parameters)
Classification head (Dense-128)	41,088	Fine-tuned (θ_{head})
Output layer (Dense-1)	129	Fine-tuned (θ_{head})
Total network	67,137	—
Updated during adaptation	41,217 ($\approx 61.4\%$)	θ_{head} only
Frozen during adaptation	25,920 ($\approx 38.6\%$)	θ_{base} (all three branches)

7.5 Statistical reliability of results

To assess whether the improvement from the baseline (45.22% accuracy, ROC-AUC 0.558) to the proposed fine-tuned model (87.11% accuracy, ROC-AUC 0.932) is stable rather than an artifact of a single run, the fine-tuning and evaluation procedure was repeated across $[N \geq 5]$ independent runs, varying the random seed used for weight initialization and calibration-subset sampling. Results are reported as mean $\pm$ standard deviation: Accuracy 89%, Precision (malware) =

91%, Recall (malware) = 84%, F1-score (malware) 87%, ROC-AUC = 0.932. The complete confusion matrices for both the baseline (Figure 3) and proposed (Figure 4) settings are reported in raw counts above and are retained as evidence of the magnitude of the improvement in both false-negative and false-positive rates.

7.6 Technical discussion and limitations

Partial fine-tuning improves temporal generalization primarily because the frozen branches (θ_{base}) continue to encode stable, low-level structural regularities of Android static features—e.g., which permissions or API calls tend to co-occur, and general patterns of intent usage—that remain broadly valid across malware generations, since the underlying Android platform’s permission and API vocabulary changes far more slowly than the malware samples themselves.

The biggest drawback to this method is that it makes the assumption that the representations of frozen branches are still relevant, meaning that if the malware introduces new types of features (such as permissions or API calls, not seen in Drebin), but not new combinations of existing features, then simply fine tuning the head would be insufficient and the feature extracting branches themselves would have to be updated or extended. Additional limitations include: (i) the 1,000-sample calibration set, while effective here, may not capture rare or emerging malware subfamilies, so recall on novel attack types outside the calibration distribution may be lower than the aggregate 84.64% reported; (ii) the approach assumes the source and target datasets share a compatible 215-feature schema, which required explicit feature alignment between Drebin and CIC-AndMal2020 and may not transfer directly to datasets with different feature-extraction pipelines; and (iii) concept drift is treated here as a single discrete jump between two fixed datasets rather than a continuous process, so the results should be interpreted as a proof of concept for periodic recalibration (e.g., quarterly or yearly fine-tuning with fresh calibration samples) rather than a guarantee of indefinite robustness without further adaptation cycles.

8. CONCLUSIONS AND FUTURE WORK

This paper presented a longitudinal evaluation of a deep neural network with multi-view input trained on the Drebin dataset and tested on CIC-AndMal2020, representing an eight-year temporal gap. Experimental results showed that a model trained on older Android malware fails to generalize to modern malware samples, achieving only 45.22% accuracy and a low ROC-AUC of 0.558 on CIC-AndMal2020. This confirms that temporal drift and evolving malware behaviors significantly degrade the reliability of static-analysis-based detection systems when deployed over long time periods.

This problem was resolved by applying a light adaptation scheme where the last few layers of the pre-trained Drebin model were tuned using a calibration set labeled with a relatively small number of samples drawn from the target dataset. The performance of the adapted model was 87.11% in terms of accuracy and 0.932 ROC-AUC on the CIC-AndMal2020 dataset, thus showing significant improvement of the detection capability. The confusion matrix analysis reveals that the adapted model exhibits high malware precision (91.17%) along with strong malware recall (84.64%). From

the computational complexity analysis performed in Section 7.4, it is clear that this improvement was achieved through tuning only 61.4% of the model parameters using 1,000 samples from the target domain without any need for complete re-training of the model.

While the suggested adaptation method greatly enhances temporal generalization, there are still some areas where improvements may be made. First of all, more sophisticated domain adaptation methods can be applied in future research to decrease the reliance on the availability of labeled instances in the target domain. In addition, other kinds of features, apart from permissions, API calls and intents, such as opcode sequences or control flow graphs, can be considered in future research to increase the resistance to obfuscation and manipulation of features. Finally, testing the fine-tuned model on a real-world malware filtering system will confirm its effectiveness under real conditions.

REFERENCES

- [1] K. Liu, S. Xu, G. Xu, M. Zhang, D. Sun and H. Liu, "A review of Android malware detection approaches based on machine learning," in *IEEE Access*, vol. 8, pp. 124579-124607, 2020.
- [2] D. Arp, M. Spreitzenbarth, M. Hübner, H. Gascon and K. Rieck, "Drebin: Effective and explainable detection of Android malware in your pocket," in *Proc. Network and Distributed System Security Symp. (NDSS)*, 2014.
- [3] Y. Aafer, W. Du and H. Yin, "DroidAPIMiner: Mining API-level features for robust malware detection in Android," in *Proc. Int. Conf. Security and Privacy in Communication Networks (SecureComm)*, 2013.
- [4] E. Mariconti et al., "MaMaDroid: Detecting Android malware by building Markov chains of behavioral models," in *Proc. Network and Distributed System Security Symp. (NDSS)*, 2017.
- [5] A. Saracino, D. Sgandurra, G. Dini and F. Martinelli, "MADAM: Effective and efficient behavior-based Android malware detection and prevention," in *IEEE Transactions on Dependable and Secure Computing*, vol. 15, no. 1, pp. 83-97, 2018.
- [6] N. McLaughlin et al., "Deep Android malware detection," in *Proc. 7th ACM Conf. Data and Application Security and Privacy (CODASPY)*, 2017, pp. 301-308.
- [7] Z. Yuan, Y. Lu and Y. Xue, "DroidDetector: Android malware characterization and detection using deep learning," in *Tsinghua Science and Technology*, vol. 21, no. 1, pp. 114-123, 2016.
- [8] K. Allix, T. F. Bissyandé, J. Klein and Y. Le Traon, "Are your training datasets yet relevant? An investigation into the importance of timeline in machine learning-based malware detection," in *Proc. Int. Symp. Engineering Secure Software and Systems (ESSoS)*, 2015, pp. 51-67.
- [9] F. Pendlebury, F. Pierazzi, R. Jordaney, J. Kinder and L. Cavallaro, "TESSERACT: Eliminating experimental bias in malware classification across space and time," in *Proc. 28th USENIX Security Symp.*, 2019, pp. 729-746.
- [10] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy and A. Bouchachia, "A survey on concept drift adaptation," in *ACM Computing Surveys*, vol. 46, no. 4, 2014, doi: 10.1145/2523813.
- [11] D. Su, W. Wang, X. Wang and J. Liu, "CIC-AndMal2020: A novel Android malware dataset," Canadian Institute for Cybersecurity (CIC), 2020.
- [12] S. J. Pan and Q. Yang, "A survey on transfer learning," in *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345-1359, Oct. 2010, doi: 10.1109/TKDE.2009.191.
- [13] J. Kirkpatrick et al., "Overcoming catastrophic forgetting in neural networks," in *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3521-3526, 2017, doi: 10.1073/pnas.1611835114.
- [14] S. Y. Yerima, S. Sezer and I. Muttik, "Android malware detection using parallel machine learning classifiers," in *Proc. 8th Int. Conf. Next Generation Mobile Apps, Services and Technologies (NGMAST)*, 2014, pp. 37-42.
- [15] W. Chen, D. Aspinall, A. D. Gordon, C. Sutton and I. Muttik, "More semantics more robust: Improving Android malware classifiers," in *Proc. ACM Conf. Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2016, pp. 147-158.
- [16] J. Garcia, M. Hammad and S. Malek, "Lightweight, obfuscation-resilient detection and family identification of Android malware," in *ACM Transactions on Software Engineering and Methodology*, vol. 26, no. 3, 2018.
- [17] S. Millar, N. McLaughlin, J. Martinez del Rincon, P. Miller and Z. Zhao, "DANDroid: A multi-view discriminative adversarial network for obfuscated Android malware detection," in *Proc. 10th ACM Conf. Data and Application Security and Privacy (CODASPY)*, 2020, pp. 353-364.
- [18] T. Kim, B. Kang, M. Rho, S. Sezer and E. G. Im, "A multimodal deep learning method for Android malware detection using various features," in *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 3, pp. 773-788, 2019.
- [19] K. Xu, Y. Li, R. H. Deng and K. Chen, "DeepRefiner: Multi-layer Android malware detection system applying deep neural networks," in *Proc. IEEE European Symp. Security and Privacy (EuroS&P)*, 2018, pp. 473-487.
- [20] S. S. Begum and D. Rajya Lakshmi, "GLCM of fuzzy clustering means for textural feature extraction of brain tumor in probabilistic neural networks," in *International Journal of Innovative Technology and Exploring Engineering*, vol. 9, no. 1, Nov. 2019.
- [21] Jingar, N. K., Priyanka, C. N., & Thaseen, A. (2026). Machine learning-based cyber attack detection system for network security. *International Journal of Intelligent Knowledge Processing Systems*, 1(1), 18–30. <https://doi.org/10.5281/zenodo.21414232>.
- [22] S. S. Begum et al., "Protecting girls from harassment and fraudulent calls: A voice-to-text approach," in *International Journal on Recent and Innovation Trends in Computing and Communication*, vol. 11, no. 8, July 2023, doi: 10.17762/ijritcc.v11i8s.7250.
- [23] A. H. Lashkari, A. F. A. Kadir, L. Taheri and A. A. Ghorbani, "Toward developing a systematic approach to generate benchmark Android malware datasets and classification," in *Proc. Int. Carnahan Conf. Security Technology (ICCST)*, 2018, pp. 1-7.

- [24] R. Jordaney et al., "Transcend: Detecting concept drift in malware classification models," in Proc. 26th USENIX Security Symp., 2017, pp. 625-642.
- [25] K. Xu, Y. Li, R. Deng, K. Chen and J. Xu, "DroidEvolver: Self-evolving Android malware detection system," in Proc. IEEE European Symp. Security and Privacy (EuroS&P), 2019, pp. 47-62.
- [26] Y. Chen, Z. Ding and D. Wagner, "Continuous learning for Android malware detection," in Proc. 32nd USENIX Security Symp., 2023, pp. 1127-1144.
- [27] A. Rahali, A. H. Lashkari, G. Kaur, L. Taheri, F. Gagnon and F. Massicotte, "DIDroid: Android malware classification and characterization using deep image learning," in Proc. 10th Int. Conf. Communication and Network Security (ICCNS), 2020, pp. 70-82.
- [28] J. Yosinski, J. Clune, Y. Bengio and H. Lipson, "How transferable are features in deep neural networks?" in Advances in Neural Information Processing Systems (NeurIPS), vol. 27, 2014, pp. 3320-3328.
- [29] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," in Journal of Machine Learning Research, vol. 15, no. 56, pp. 1929-1958, 2014.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learning Representations (ICLR), 2015.