



GreenForensics: Deep Hybrid Edge-Cloud Detection and Forensics System for Battery-Performance-Balance Conscious Devices

By:

Mohit Sewak, Sanjay K. Sahay and Hemant Rathore

From the proceedings of
The Digital Forensic Research Conference
DFRWS APAC 2022
Sept 28-30, 2022

DFRWS is dedicated to the sharing of knowledge and ideas about digital forensics research. Ever since it organized the first open workshop devoted to digital forensics in 2001, DFRWS continues to bring academics and practitioners together in an informal environment.

As a non-profit, volunteer organization, DFRWS sponsors technical working groups, annual conferences and challenges to help drive the direction of research and development.

<https://dfrws.org>



Contents lists available at ScienceDirect

Forensic Science International: Digital Investigation

journal homepage: www.elsevier.com/locate/fsidi

DFRWS 2022 APAC - Proceedings of the Second Annual DFRWS APAC

GreenForensics: Deep hybrid edge-cloud detection and forensics system for battery-performance-balance conscious devices

Mohit Sewak^{a,*}, Sanjay K. Sahay^b, Hemant Rathore^b^a Security & Compliance Research, Microsoft R&D, India Pvt. Ltd., India^b Dept. of CS & IS, Goa Campus, BITS Pilani, India

ARTICLE INFO

Article history:

Keywords:

Efficient forensics
Battery-performance balance aware deep learning
Deep clustering
Hybrid edge

ABSTRACT

Motivated by the advancements made by the recently proposed DRo algorithm to uplift the performance of data scarce Deep Learning malware detector for edge, we propose an adaptive and efficient system for hybrid edge-cloud detection and forensics, named GreenForensics. The proposed adaptive enhancement, makes the system more suitable for devices with custom battery-performance optimization mandates like tablets and laptops. Further, the enhancements offer various discrete and continuous controls for influencing the detection coverage and model footprints in real time. To further enhance the detection efficiency and making the detection resilient to adversarial-attacks, the proposed system can work with adversarial-DL immune algorithms. In the experiments conducted, GreenForensics was able to significantly outperform even the best baseline deep architectures and improved the detection and forensics robustness by up to **100%** and performance by up to **40%**. This gains further significance as the incumbent baseline DL architecture had up to **6700%** higher neural inference complexity, and had its performance and robustness benchmarks had remained unchallenged for a long time.

© 2022 The Author(s). Published by Elsevier Ltd on behalf of DFRWS This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

As per the *Work Trend Index* report published by Microsoft (Wiseman, 2021), more than 73% of employees would like to continue to work in a hybrid manner post pandemic, and 66% of business leader are considering redesigning their work culture to enable this. This sudden and extreme shift to hybrid work cultures has led to the adoption of more personal and mobile devices like Windows laptops and tablets and raising the concerns on enabling similarly high level of security mechanism on these battery-performance conscious machines. Since such battery powered client devices can ill-afford the current advanced and computationally involved detection mechanisms as were designed for their more powerful and AC connected workstations and server counterparts, this adoption trend is also seen as an opportunity by many ransomware operators and malware developers to specifically target such devices. Therefore, the recent issue of Microsoft's *Digital Defence Report* (Hogan-Burney, 2021) observes that human-operated-ransomware (HumOR) operators are increasingly attacking devices belonging to this segment. In 2021, the individual

consumers were the highest engagement segment for the (13%) Ransomware as a Service (RaaS) operators, followed by financial and manufacturing sectors (at 12% each). Therefore, innovation for performant and efficient mechanisms to protect such battery-conscious Windows clients from ransomware and other advanced malware has increasingly become critical.

Recently Sewak et al. presented a novel and innovative algorithm named DRo (Sewak et al., 2021a) (for **Deep Router**), and proposed many end-to-end (e2e) Deep Learning (DL) architectures for implementing DRo under multiple unique scenarios; including some for malware detection (followed by forensics via vault) at the edge device itself. The promising experimental results indicated that DRo could revolutionize the performance of DL based security systems even under extremely data-scarce scenarios. The DRo algorithm was benchmarked using an Android based on-edge architecture (DRoID), and it was demonstrated that DRo could help design a performant and efficient malware detector even with low-information features extracted from sparsely labeled data. Further, it was demonstrated that the model could be tuned for varying degrees of performance (accuracy/recall) and false-positive-rates (FPR) balance mandates for the downstream DL model.

Motivated by the performance of DRo and some of its architectures that intrinsically support a vault based-mechanism for

* Corresponding author.

E-mail address: mohit.sewak@microsoft.com (M. Sewak).

subsequent forensics, we propose new and unique architectures and application areas for DRO. In particular, we propose multiple adaptive enhancements to the architecture of DRO and apply them to the Windows endpoint security domain. We further strengthen the vault mechanism to have adaptive coverage to support efficient forensics. We call the resulting system as GreenForensics. We take one of DRO's (hierarchical) architecture that has been proposed for low-powered edge devices, and adapt it to suit multitude of hybrid edge-cloud inference requirements. The adaptive enhancement proposed by us, makes DRO more suitable for use with custom battery-performance-balance mandate devices like Windows clients, tablets, and laptops; and simultaneously also provides opportunities for influencing performance and coverage balance in real time. This is achieved by offering various discrete and continuous controls options with varying complexity and coverage outcomes that takes the *Power Mode* mandate and battery drain status as inputs. Moreover, to further enhance the detection and forensics efficiency and to make the system resilient to adversarial-DL attacks (Rathore et al., 2021; Sewak et al., 2021b), we adapted DRO to also work with adversarial-attack immune classical Machine Learning (ML) algorithms like the Random Forest (RF) (Sewak et al., 2018). In line with the results claimed in the DRO's paper (of improving the robustness and accuracy of downstream DL detection by up to $\approx 67\%$ and 11% respectively (Sewak et al., 2021a)), GreenForensics offers similar advantages and further improves upon these benchmarks and has more controlled coverage of malware samples for subsequent forensics. On a popular and standardized malware and ransomware dataset (Nappa et al., 2015), GreenForensics was able to improve the detection robustness by up to **100%** (0% FPR across multiple experiments) and performance/area-under-curve statistics (AUC) by up to $\approx 40\%$ as compared to even the best baseline e2e DL configurations. This is significant, as the comparison was made with some of best baseline e2e DL architectures, which had up to $\approx 6700\%$ higher neural inference complexity as compared to GreenForensics, and their established benchmarks had remained unchallenged by any other DL architecture, technique, or model since a long time now.

Following are the key research contributions of this paper:

1. We propose a new architecture for DRO, named GreenForensics, which could balance between the edge detection:forensics coverage ratio and detection complexity (without compromising with the detection performance or robustness); and hence is extremely useful for hybrid edge-cloud detection.
2. The original DRO introduced hyperparameters to influence a desirable trade-off between accuracy and FPR; we propose further enhancements to extend such flexibility to balance between on-demand coverage and on-demand complexity. Moreover, we provide two unique types of controls, one discrete control and another continuous control for this purpose.
3. Where the primary aim of the original DRO proposal was to enhance the performance of a downstream DL classifier under scarce label scenario (as an alternative to Transfer Learning, but under more extreme conditions); we further demonstrate that DRO, despite being a highly-efficient and low-complexity algorithm, could (mostly) by itself surpass the benchmarks of most DL architectures, even e2e DL pipelines comprising of multiple DL models, all of which being many times larger and more complex than DRO; and does not need an assistance of existing e2e DL pipeline to achieve this objective. To this end, we combine DRO with exceedingly simple, efficient and adversarial-DL immune classical ML algorithms and obtain several benchmarks to support this claim.
4. Finally, we further established the universality and wide applicability of the DRO mechanism by implementing it on datasets

beyond Android malware, and beyond simplistic/non-intuitive features. We apply DRO on informative features, on which exceedingly high benchmarks already exists. Our experiments demonstrate the DRO continue to break all benchmarks established by any existing DL models as used in security, including e2e DL pipelines and several layered deep malware detection systems.

The remaining the paper is organized as follows. In section 2 we cover the related work with several aspect of DRO and the problem we solve with GreenForensics. Next, we cover the mathematics of (original) DRO and our proposed architecture for GreenForensics in section 3. We cover details of an illustrative implementation of GreenForensics for a Windows based power-performance optimized device (WinDRO) in section 3.4. Then we provide the details on the dataset and baseline experiments in section 4, and next the benchmark experiments and their results in section 4.1. Finally, we provide conclusion in section 5.

2. Related work

In this section, we first describe some related work on the mathematical basis of DRO, which is built over recent advances in the field of Deep Clustering. Next, since GreenForensics at a high level does data sampling for training and inferencing at the edge vs. on the cloud, we produce some art related to these aspects.

2.1. Deep clustering

End-to-End Deep Learning for discrete representation learning (clustering and hashing) has is becoming popular in research fraternity and different algorithms have been proposed that have been proposed to this end. Such discrete representation could be for clustering or hashing. When Deep Learning is used for generating discrete representations like hashing or clustering, we in general refer to it as Deep Clustering (DC). Many of the conventional and popular *representation* learning algorithms like Gaussian Mixture Modeling (GMM) and K-Means are incapable of modelling non-linear boundary separation between groups/clusters. Others like the *kernel* (Xu et al., 2005; Kulis and Darrell, 2009) and *spectral* (Weiss et al., 2009) clustering-based techniques though can model arbitrary cluster boundaries, but can not efficiently scale to large datasets. Here is where DL, especially deep clustering can create a large impact, as it is capable of modeling and learning non-linear and complex separation boundaries and is extremely scalable and flexible. Many of the recent algorithms in this field are also capable of producing data embeddings along with the cluster representations.

Besides learning discrete data representation, the reason why deep clustering became popular as a semi-supervised pre-processing approach for DL algorithms is because it could generate embeddings that could be used to make e2e DL architectures. In this regard, the first DC algorithm that could simultaneously also generate embedding and clustering representations was Deep Embedding for Clustering (DEC) (Xie et al., 2016). Since this is an evolving field of research different ideas have gained popularity on how to make discrete representation learning more robust. In this regard, some approaches for generating adversarial robustness to different types of DL models exist (Jiang et al., 2017). In these approaches, the distribution of data is modeled by a *generative* algorithm, and then GMM models are used to represent the prior distributions for these models. This approach is atypical to data generation using Variational Auto Encoders (VAE) (Sewak et al., 2020). Leen (1995) proved that applying data augmentation to a DL classifier gives a similar effect as applying regularization to the

original cost function of a conventional machine learning (ML) model. Miyato et al. (2015), and Sajjadi et al. (2016) successfully adapted this approach to semi-supervised DL algorithms and demonstrated successful results. Self Augmentation Training (SAT) is inspired by this approach of regularization. Dosovitskiy et al. (2014) proposed to use data augmentation to model the invariance of learned representations for unsupervised algorithms like clustering. IMSAT takes a similar approach but applies invariance directly to the learned representation instead of applying it to the surrogate classes, and directly learns discrete representations (clusters), instead of learning continuous representations that are later converted to discrete class representations or clusters predictions. Further, it combines Information Maximization (IM) along with SAT to produce more robust cluster formations; a special coefficient λ is used to balance between entropy loss and augmentation loss for learning cluster representations.

2.2. Sampling for forensics and hybrid edge-cloud detection

Where the role of a malware classifier is to detect a malware sample, the role of a forensic sampler is to analyze which sample could or could not be detected well by a malware detector and hence sample them either to an automated detector or to a dedicated forensics team or a vault. Where there exists good research on designing effective and efficient malware classifiers, there is a dearth of research on automated forensic/detection sampler systems. Whatever work exists, use complex Deep Reinforcement Learning (DRL) algorithms and offer black-box solutions to one of the variables in the integrated problem scenario of training data sampling, associated inference data routing, hybrid mobile-edge detection, and model complexity-based routing. Some selected works in each of these problem areas are given as follows. Recently some suggestions have come from the research in Network Intrusion Detection System (NIDS), where Lopez et al. (Lopez-Martin et al., 2020) used DRL to train agents to learn to sample anomaly logs for training an anomaly detection system. Similarly, DRL solutions have been proposed by Xiaoyue et al. (Wan et al., 2017) for selective offloading of mobile-based inference to the cloud. Similarly, for the problem of overall detection complexity management, a DRL based solution to route appropriate inference candidates to one of the many available models of different complexity is proposed by Yoni et al. (Birman et al., 2020).

3. Background and the proposed architecture

In this section we provide a brief on the DRo algorithm (in section 3.1) cover the mathematical basis of (original) DRo (in section 3.2). Next, we cover an illustrative architecture for implementing a GreenForensics based system for hybrid edge-cloud inference (in section 3.3) and compare it with that of one of the architectures (for edge devices) of DRo. Finally, we cover the formulation of the WinDRo solution that applies the GreenForensics architecture to a battery-performance optimized connected Windows client (tablet/laptop) for hybrid edge-cloud detection and routing (in section 3.4).

3.1. About DRo

Sewak et al. (2021a) used the ideas of Deep Clustering to solve a novel problem faced by most DL application domains, i.e. learning robust detection under scarce label data scenarios. In this regard, they further modified the deep clustering algorithms to form DRo (Deep Router) such that the cluster assignment directly ends in providing routing suggestions for data samples based on the cluster-separation-margin which is indirectly an indication of the

associated noise level in the sample. The DRo mechanism proves useful for effective, noise-free training of a (label data affine) DL classifier. Further, the same routing suggestions were used for selectively routing inference candidates to a classifier trained by DRo sampled data. Since the training and inferencing are both controlled by the same model, DRo enabled downstream DL discriminators to effectively learn meaningful discriminative features with little representative data. In doing so, the downstream DL classifiers consistently outperformed even more sophisticated models trained on the same input data. With an end-objectively to influence the recall and robustness of the downstream discriminator, Sewak et al. introduced another hyperparameter, μ , that strike a balance between the absolute and conditionally entropy losses to alter the routing suggestions. We cover the related mathematical details on DRo in section 3.2.

3.2. Mathematical basis of DRo

Representation Learning and unsupervised learning like clustering are difficult tasks as the data associations are not defined. But acquiring labeled data is expensive, and sometimes even infeasible. Also, for many tasks, the labels are not determined apriori. Hence, discrete representation learning, like hashing and clustering, from unlabeled is a critical task. Classical Machine Learning (ML) offers many popular clustering algorithms like the k-means (Hartigan and Wong, 1979), hierarchical clustering (Johnson, 1967) etc. For this purpose; but these algorithms do not scale to large datasets and complex non-linear patterns. These are the areas where Deep Learning (DL) has invariably replaced many classical ML solutions. DL offers complex algorithms to model sophisticated non-linear patterns in the data. This also makes DL over-fit to the training data easily, and in the process lose the ability to identify meaningful domain representations from noise. In unsupervised learning, since the target is not specified, the problem is unconstrained, which further exacerbates this problem. Therefore, such an algorithm needs optimal regularization to perform under different scenarios. In DL invariance of different types can be introduced to make it immune to slight alterations in patterns (Leen, 1995) and subsequently regularize the learning. One type of invariance is local invariance, which uses Self Augmented Training (SAT). Such augmented learning in the simplest form could be generated by random perturbations or strategic adversarial perturbations. Perturbations-based SAT generate local perturbations from original data samples and in the loss function tries to minimize the loss between surrogate classes and their predicted representation to make the model invariant to such local perturbations.

IMSAT (Hu et al., 2017) (refer section 2.1 for a discussion on other deep clustering algorithms), uses perturbation based SAT. Besides using the SAT penalties to provide regularization, IMSAT also use the (Regularized) Information Maximization (RIM) criteria (Krause et al., 2010), (Bridle et al., 1992) for identifying cluster boundaries. These two criteria are balanced in the total loss function of clustering by a regularization coefficient λ . This can be represented as in equation (1), where R_{SAT} represents the regularization loss for SAT and R_{IM} the loss due to cluster representation learning (negative information gain).

$$Loss_{total} = R_{SAT} + \lambda R_{IM} \quad (1)$$

3.2.1. R_{SAT} loss

Assume that $T: \mathbf{X} \rightarrow \mathbf{X}$ is the transformation function that generates the local perturbations (data augmentation) to ensure invariance. If p is a small perturbation that does not alter the

discrete representation of the data sample in the given context, then in the simplest form, the transformation function, T , function could be expressed as $T(x) = x + p$. SAT against such small, local perturbations, p , generates local-invariance in the learned representations, and hence pushes that the cluster-separation boundaries to the low sample density regions. Such cluster-separation boundaries are desired as these abide by the low-density-separation principle. For a sample with feature vector x , with surrogate label y , the probability $p_{\theta}(y_m|x)$ of the learned representation y_m in an M class discrete representation learning problem ($m \in [1, \dots, M]$), using a function parameterized over hyperparameter vector $\hat{\theta}$, the SAT regularization penalties, $\mathcal{R}_{SAT}(\theta; x, T(x))$, is given in equation (2).

$$\mathcal{R}_{SAT}(\theta; x, T(x)) = - \sum_{m=1}^M \sum_{y_m=0}^{V_m-1} p_{\theta}(y_m|x) \log p_{\theta}(y_m|T(x)) \quad (2)$$

The SAT regularization loss for the entire batch \mathbf{X} of data (where $x \in \mathbf{X}$) is the average of the individual sample level SAT loss $\mathcal{R}_{SAT}(\theta; x, T(x))$. This is given as in equation (3).

$$\mathcal{R}_{SAT}(\theta; T) = \frac{1}{N} \sum_{n=1}^N \mathcal{R}_{SAT}(\theta; x_n, T(x_n)) \quad (3)$$

3.2.2. R_{IM} loss

The Regularized Information Maximization algorithm minimizes the objective as in equation (4). In this, $x \in \mathbf{X}$ are the data samples, and $Y \in \mathbf{Y} = \{1, \dots, M\}$ are the cluster predictions (for M class clusters assignment problem). \mathcal{R}_{IM} is the RIM regularization penalty, and $\mathbf{I}(\mathbf{X}; \mathbf{Y})$ is Mutual Information (MI) between surrogate classes of the data samples and the cluster assignment representations.

If $\mathbf{I}(\mathbf{X}; \mathbf{Y})$ represents the Mutual Information (MI) between surrogate classes (y_m) of the data sample $x(x \in \mathbf{X})$ and the cluster assignment representations, $y_k[0, \dots, k] \in K$, the R_{IM} loss could be given as in equation eq: rim-objective.

$$\mathcal{R}_{IM} = -\mathbf{I}(\mathbf{X}; \mathbf{Y}) \quad (4)$$

The RIM's cluster probability prediction function could be expressed as $p_{\theta}(y_1, \dots, y_m|x)$ by the associated DNN architecture. Assuming that the data dimensions are conditionally independent, the joint probability could be expressed as a product of individual conditional probabilities as in equation (5).

$$p_{\theta}(y_1, \dots, y_m|x) = \prod_{m=1}^M p_{\theta}(y_m|x) \quad (5)$$

If the marginal-entropy (ME) $H(Y)$ of a classes in a data is expressed as in the equation (6), and the conditional-entropy (CE) (equation (7)) of the cluster-assignments of the data conditioned on the input features is $H(Y|X)$, then the Mutual Information Gain ($-R_{IM}$) due to clustering is as a difference between ME and CE ($-(H(Y) - H(Y|X))$) (Cover and Thomas, 2006); where the entropy function is given in equation (8).

$$H(Y) \equiv h(p_{\theta}(y)) = h\left(\sum_{i=1}^N p_{\theta}(y|x)\right) \quad (6)$$

$$H(Y|X) \equiv \frac{1}{N} \sum_{i=1}^N h(p_{\theta}(y|x)) \quad (7)$$

$$h(p(y)) \equiv - \sum_{y'} p(y') \log p(y') \quad (8)$$

3.2.3. L_{DRo}

An enhanced the Marginal Entropy $H(Y)$ could enforce the cluster sizes to be uniform. Whereas a diminished Conditional Entropy $H(Y|X)$ could enforce unambiguous and low-noise cluster assignments (Bridle et al., 1992). To influence this balance desirably, an additional hyperparameter $\mu \in \mathbb{Z}$, was added to further tune the Conditional Entropy in DRo. This mechanism offered added advantage of controlling the coverage across different routes of DRo, thus enabling different modes like the *Vault Mode* and *Hierarchical Mode* settings. The modified loss function of the DRo algorithm could be given as in equation (9).

$$Loss_{DRo} = \mathcal{R}_{SAT} - \lambda((H(Y) - \mu H(Y|X))) \quad (9)$$

Experimentally the coefficient $\mu \in$ (Sewak et al., 2021a, 2021b) provides a decent range for default settings, and $\mu \in$ (Wiseman, 2021; Weiss et al., 2009) provides a decent range to influence recall and FPR across a wide range of applications.

3.3. Proposed architecture for GreenForensics

The original DRo (Sewak et al., 2021a) offers multiple architectures. The prominent ones of those being a single-layer architecture, mostly suitable for an unconnected edge device with on-device vault provision, and a second one for slightly more powerful edge devices, that can use multiple layers of detection (e.g., as in Advanced Threat Protection (ATP)) which might be sparingly connected to online services for dynamic analysis, feedback, and model updates. Both architectures are remarkably useful for mobile devices that are mostly configured of extreme power efficiency and in-frequent network usage. The hierarchical (multi-routing) DRo architecture however could be configured for a variety of usage. Therefore, we take this second architecture as a base, and modify it to propose an architecture for more powerful clients (e.g., tablets, and laptops), which though value battery efficiency, but largely desires an optimal battery-performance balance (as is available in Windows 11), and are more often connected. The architecture for the proposed GreenForensics system is given in Fig. 2 and for comparison, the architecture of Hierarchical multi-routing DRo is given in Fig. 1. The following are the key difference between these architectures:

- As compared to the hierarchical DRo, GreenForensics transfers the subsequent detection routing beyond the first layer *selectively* to the cloud infrastructure.
- Both the architectures use multiple DRo models, but where the hierarchical DRo uses these in a tree based hierarchical section, and all the models could have similar model complexity and footprint, GreenForensics selective choose one of the multiple models for edge-detection routing, and each model has a varying level of neural complexity and footprint.
- Instead of assisting a DL classifier and forming an e2e DL architecture with shared features, GreenForensics use a classical ML algorithm (e.g., Random Forest) for 2 main reasons; first to make the systems more robust to adversarial-DL attacks, and second to further lower the model neural complexity and footprints, and make it ideal for devices which do not have a powerful/dedicated GPU.

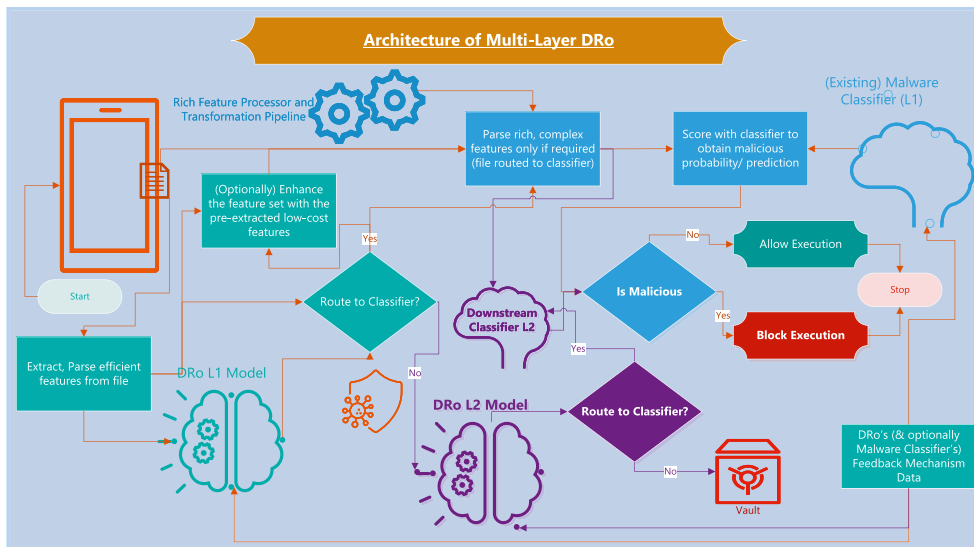


Fig. 1. Architecture of hierarchical (multi-routing) DRo.

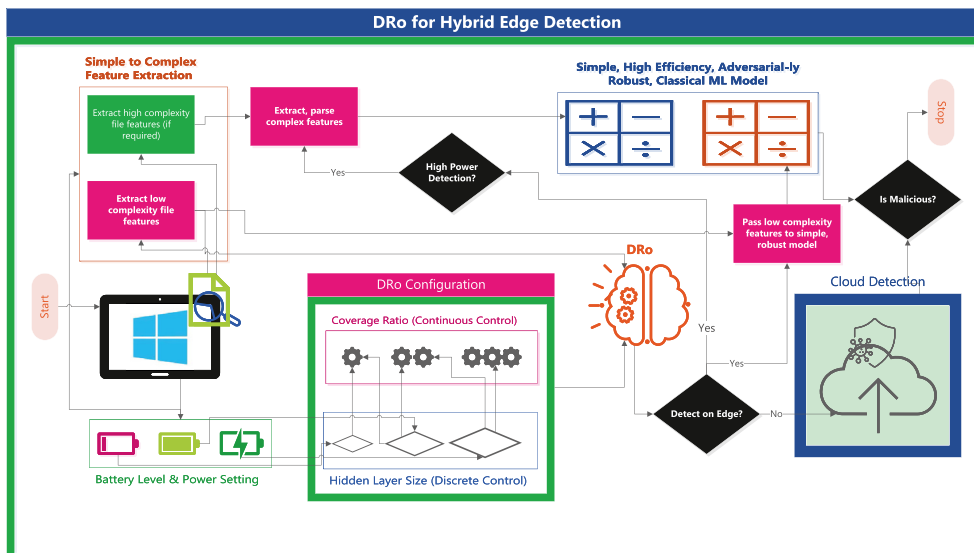


Fig. 2. Architecture of GreenForensics

3.4. Adaptive formulations for WinDRo

We use the GreenForensics architecture in the Windows context, to design a hybrid edge-cloud detection system, named WinDRo. The WinDRo system is based on GreenForensics architecture (section 3), is trained on Windows PE dataset (section 4), and the different DRo configurations with varying discrete and continuous battery-performance optimization control dimensions (section 4.1). This solution is illustrated in Fig. 3. For this, we take 2 inputs, one is the *Performance Mode* set by the user for their device (e.g., Windows Laptop), second is the battery level of the device. Where the first adaptive control input is purely discrete, the second option is largely based on continuous inputs but may use discrete configurations. Based on the *Performance Mode*, a suitable (pre-compiled) architecture of the DRo model is loaded into the model. All of these architecture are single layer ones to offer high efficiency and low model footprint (a function of trainable parameters and computations), and differs only in the size of the single hidden

layer. Next, based on the battery level, an equally performant configuration, but with varying coverage ratio is activated. Based on the configuration parameters (λ, μ), the neuron connection weights may also change, and hence these are also modified without re-compiling the model or changing its architecture. For lower battery levels, a configuration with lower edge-detection coverage/forensics is loaded and vice-versa.

4. Dataset and baseline models

We use the malware from the overly popular Malicia dataset (Nappa et al., 2015). One reason for choosing this highly popular dataset is that many benchmarks exist on this dataset, and these are so high that it is considered futile to break. The most recent and the highest benchmark claimed on this dataset using e2e DL architectures is by Sewak et al. (2018). They obtained an accuracy of 99.21% at an FPR of 0.2% while using a 3-layer Auto-Encoder, coupled with a 4-layer MLP-DNN. The features used were the frequency vector of

