

VU Research Portal

Are Machine Learning Models for Malware Detection Ready for Prime Time?

Cavallaro, Lorenzo; Kinder, Johannes; Pendlebury, Feargus; Pierazzi, Fabio; Massacci, Fabio; Bodden, Eric; Sabetta, Antonino

published in

IEEE Security and Privacy
2023

DOI (link to publisher)

[10.1109/MSEC.2023.3236543](https://doi.org/10.1109/MSEC.2023.3236543)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Cavallaro, L., Kinder, J., Pendlebury, F., Pierazzi, F., Massacci, F., Bodden, E., & Sabetta, A. (2023). Are Machine Learning Models for Malware Detection Ready for Prime Time? *IEEE Security and Privacy*, 21(2), 53-56. <https://doi.org/10.1109/MSEC.2023.3236543>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:


vuresearchportal.ub@vu.nl

Are Machine Learning Models for Malware Detection Ready for Prime Time?

Lorenzo Cavallaro  | University College London

Johannes Kinder  | Bundeswehr University

Fergus Pendlebury  | University College London

Fabio Pierazzi  | King's College London

We investigate why the performance of machine learning models for malware detection observed in a lab setting often cannot be reproduced in practice. We discuss how to set up experiments mimicking a practical deployment and how to measure the robustness of a model over time.

In academic research on malware detection, machine learning-based techniques have become the de facto standard. Over the years, many papers have been published on the topic, using everything in the machine learning toolbox, from traditional models^{1, 2} to more recent neural network architectures.^{3, 4, 5} With published accuracies and other performance numbers regularly close to the 100% mark, you could not be blamed for thinking that the problem is practically solved. Why would anyone be interested in continuing to use malware signatures and simple patterns?

Cannot Reproduce?

Yet, practitioners find themselves disillusioned when trying to put machine learning models from research into practice. In real-world deployments, machine learning-based malware classifiers are known to often be

unreliable and are not trusted to act as the main line of defense.

Where does this discrepancy come from? It would be too easy to shrug this off as expected technical challenges in technology transfer; we must get to the bottom of the issue or risk exacerbating the reproducibility crisis decried in machine learning-based science.⁶ So, what's the deal? If the same model performs so differently in production, then the lab settings that yield near-perfect performance must not be representative of how such models are deployed.

The typical research project on machine learning-based malware detection goes as follows: 1) obtain a malware dataset from an academic project or public sources; 2) collect a dataset of benign software, e.g., through scraping repositories and app stores; 3) engineer or learn features to represent the malicious and benign apps in your datasets; and 4) train a classifier while following common good

practice guidelines, for instance, to prevent overfitting.⁷ Unfortunately, this seemingly straightforward approach comes with several pitfalls that may prevent your results from translating to practice.^{8, 9}

The first and maybe the most obvious issue is that performance metrics are influenced by the relative sizes of the classes, i.e., the ratio of malware to benign software in the evaluation dataset. To illustrate, imagine a silly malware detection model that classifies everything as malware. If your dataset consists of 95% malware and 5% benign software, then the model's performance will actually be quite decent! Widely used performance metrics are precision, recall, and F_1 —the harmonic mean of the two. With precision being the fraction of true malware in everything the classifier detected as malware, we would get to a respectable 95% because only 5% of the detections would be false. Recall is the percentage of actual

malware detected, which would even be a perfect 100%. The F_1 then comes out at an impressive 97%, although the model is completely useless in practice. Now, this may be an extreme example, but the effect is gradual and gets worse the more the class ratio at testing time differs from the class ratio in the real world. This effect of misrepresenting the ratio between classes is also known as the *base rate fallacy*.¹⁰ In most practical settings, there is vastly more benign software than malware, making real datasets heavily imbalanced toward the benign class. A classifier that is good at detecting malware but misclassifies much benign software as malicious therefore has poor performance in practice. When tested on equal amounts of malware and benign software, and even on a majority of malware, its performance will be inflated.

A second issue is that, as a result of the dataset collection procedure, malicious and benign software datasets often end up being from separate time periods. It is not unusual for malware datasets to be several years old (the Drebin dataset¹ of Android malware from August 2010 to October 2012 is still commonly used today), while software scraped from app stores will be current. Because of this, they may differ in file formats, metadata, compiler versions, and the application programming interface (API) methods used. A machine learning model trained to optimally separate the two classes is then very likely to pick up on these spurious noncausal differences instead of just the actual malicious behavior. After all, if the easiest way to distinguish the old malware from the new benign applications is to look for the presence of some deprecated permission, then this is what the model will learn.

The third issue is also due to inconsistent dates in the ordering of samples. The canonical way to train and evaluate a model is to

first randomly split the dataset into a training set and a test set.⁹ The model is trained on the training set and evaluated against the test set. However, a model's performance on randomly chosen training and testing sets is not necessarily representative of its performance in practice. Because the partitioning is random, the training set can contain samples that would not have been available to learn from at the time it would have first encountered the samples in the test set. For instance, a randomly chosen test set can contain older versions of malware contained in the training set. This way, the classifier is allowed to learn from future knowledge that would not be available in realistic settings.

In other domains, such as natural language processing, this effect is not present or at least much less pronounced. There, past and future samples can be mixed because the underlying data distribution from which the samples are drawn is mostly stable, as human language changes only slowly. In the malware detection context, the distribution of the samples changes relatively quickly. In fact, this phenomenon is common in security settings and fueled in part by human attackers acting in an adversarial fashion.^{8,11,12} Not only do newer malware samples use new API methods, their authors also actively try to circumvent detection and mitigation mechanisms and adapt their strategies accordingly. The resulting change in the distribution is known as *dataset shift* or *adversarial drift*, depending on the context. It makes any performance assessments partial and leaves us in uncharted territory when trying to understand the actual performance a system will deliver in practice.

We Can Do Better

If we wish to estimate more accurately the performance of a malware classifier in a practical deployment, we have to model the experiments

after reality. Of course, nothing can replace observing the performance of a system in production, but the best we can do in a lab setting is use the data we have to simulate how the system would have fared in practice against the malware that appeared in the past, replaying any dataset shifts that occurred.

First, we must ensure that the class ratio between our benign and malicious classes at test time mimics the one seen at deployment time. This means that researchers must have a reasonable understanding of what to expect "in the wild": Will 50% of the applications seen be malicious or, rather, 10%? In fact, this kind of understanding of the real-world environment is not just necessary for configuring the experimental setup but also of great benefit when designing the approach in the first place. This lends yet more importance to an open exchange between practitioners and academics in the field. Do note, however, that one is free to choose different class ratios at training time. While a realistic class ratio at training time will minimize the overall error rate (the fraction of misclassifications over all samples), other ratios can be used to move the decision boundary and hence trade off the likelihood of false positives against that of false negatives.⁸

Second, malware and goodware have to come from the same time frame; otherwise, the classifier will pick up noncausal artifacts as discriminative features. This is no easy task, as datasets of malware are treated as trade secrets by IT security firms, and benign software may not be readily available from past time periods or at all; this is true in particular for paid apps in closed source environments.

Notably, the easy availability of standard datasets has been one of the drivers in applications of machine learning to domains such as image recognition and natural language processing. If we are serious about advancing

research in machine learning for security, we need to make high-quality datasets publicly available. Only then can new generations of Ph.D. students easily try out new ideas and compare approaches objectively.

The AndroZoo project¹³ is a great example of such an effort. It provides access to over 20 million Android applications, both benign and malicious, from a time span of over 10 years and continues to be continuously updated. For the Windows platform, the Sophos-ReversingLabs (SOREL)¹⁴ dataset contains pre-extracted features and metadata for 20 million portable executable files, half benign and half malicious. Unfortunately, it is typically not allowed to freely redistribute Windows binaries, so SOREL has full binary files only for malware.

Having a dataset of time stamped malware and benign applications also allows us to avoid the third pitfall mentioned in the preceding: if the machine learning model should be used for proactive malware detection, i.e., detecting new malware as it comes in, then your training data should not be newer than your test data. For instance, this is a typical setting for a malware detector deployed at an app store that scans newly uploaded applications. Known malware may effectively be detected using signatures, so classifier performance on older specimens may be less relevant in practice. We can model this setup in an experiment through a sliding window of test data, using only applications with an earlier timestamp as training data.⁸ Interestingly, this also allows us to evaluate how a classifier trained at a certain time in the past would have performed on the new malware as it comes in.

Performance Degradation

Once deployed, a trained malware classifier will degrade in performance over time if it is not regularly retrained.^{11, 12} This is a result of the

forementioned dataset shift due to malware and, in fact, all applications changing over time as malicious strategies and available APIs change.

A setup as described previously, and detailed by Pendlebury et al.,⁸ can evaluate how the performance of classifiers of interest would have degraded in the past after some deployment date. Figure 1 shows the results for the Drebin classifier trained on a dataset containing apps up to December 2016. The classifier was tested 24 times against the new malware appearing each month, to cover a period of two years up to December 2018. The lines show how the various metrics decline, with the F_1 score on the malware class being the most commonly used metric for this type of work. Depending on the activity of individual malware campaigns, performance can rise temporarily, but the overall downward trend is clearly visible. From the plot, we can also see that all performance metrics are much higher on the majority class; that is, the classifier continues to do a relatively good job in identifying most benign software. This is expected in an imbalanced setting and reaffirms the importance of picking the class ratios correctly. This type of plot lets us visualize the performance over time. For easier comparison, we can also capture the time-related decay in a single number: the area under time is the area under the F_1 curve and provides a measure of the classifiers' robustness against dataset shifts.

A classifier has to be maintained just like any other detection mechanism. For a machine learning classifier, this essentially amounts to retraining, possibly incrementally. Designing a retraining strategy that trades off costs for labeling, training, and, possibly, model distribution against performance degradation is a subtle problem. We can equip classifiers with a rejection option to allow them to defer classification of a particular sample, due

to it being an outlier with respect to the known classes.¹¹ This allows the system to identify drift and trigger retraining.

Looking Forward

There is still much research to be done in the area despite the host of available publications. We now have a better understanding of the performance of machine learning models for malware detection and related security tasks. With the appropriate precautions described previously and fleshed out in recent work,^{6, 8, 9} the deployment and maintenance of machine learning-based classifiers is promising yet still a delicate matter.⁷ For instance, the right abstractions to approximate program semantics, detecting and dealing with dataset shifts,¹¹ the costs of labeling in retraining and risks in online learning strategies, and defenses from realizable adversarial

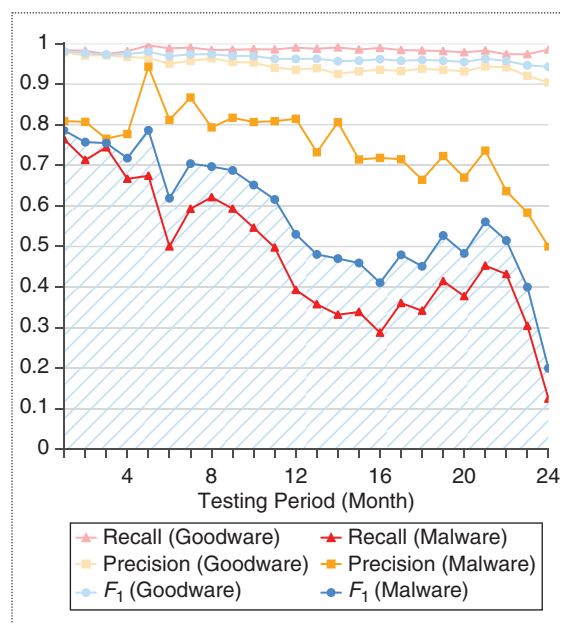


Figure 1. The performance decay in several metrics of a malware detection model trained at month 1 when evaluated against the new malware and benign software of each month for 24 months (adapted from Pendlebury et al.⁸). Performance numbers are shown both for detecting malware and detecting benign software (goodware, in light colors), respectively. The shaded area under the F_1 curve for malware detection is the area under time, measuring robustness against decay.

attacks¹⁵ are just some of the open research questions our community must deal with going forward. ■

References

1. D. Arp, M. Spreitzenbarth, M. Hubner, H. Gascon, and K. Rieck, "DREBIN: Effective and explainable detection of android malware in your pocket," in *Proc. Annu. Netw. Distrib. Syst. Secur. Symp. (NDSS)*, 2014, pp. 1–15, doi: 10.14722/NDSS.2014.23247.
2. L. Onwuzurike, E. Mariconti, P. Andriotis, E. D. Cristofaro, G. J. Ross, and G. Stringhini, "MaMaDroid: Detecting android malware by building Markov chains of behavioral models (Extended Version)," *ACM Trans. Privacy Security*, vol. 22, no. 2, pp. 14:1–14:34, May 2019, doi: 10.1145/3313391.
3. K. Grosse, N. Papernot, P. Manoharan, M. Backes, and P. McDaniel, "Adversarial examples for malware detection," in *Proc. Eur. Symp. Res. Comput. Secur. (ESORICS)*. Springer-Verlag, 2017, pp. 62–79, doi: 10.1007/978-3-319-66399-9_4.
4. E. Raff, J. Barker, J. Sylvester, R. Brandon, B. Catanzaro, and C. K. Nicholas, "Malware detection by eating a whole EXE," in *Proc. AAAI Workshop Artif. Intell. Cyber Secur.*, 2018, pp. 268–276.
5. K. Pei, Z. Xuan, J. Yang, S. Jana, and B. Ray, "Trex: Learning execution semantics from micro-traces for binary similarity," 2020, *arXiv:2012.08680*.
6. S. Kapoor and A. Narayanan, "Leakage and the reproducibility crisis in ML-based science," 2022, *arXiv:2207.07048*.
7. J. Saxe and H. Sanders, *Malware Data Science: Attack Detection and Attribution*. San Francisco, CA, USA: No Starch Press, 2018.
8. F. Pendlebury, F. Pierazzi, R. Jordaney, J. Kinder, and L. Cavallaro, "TESSERACT: Eliminating experimental bias in malware classification across space and time," in *Proc. USENIX Secur. Symp.*, 2019, pp. 729–746.
9. D. Arp et al., "Dos and don'ts of machine learning in computer security," in *Proc. USENIX Secur. Symp.*, 2022. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity22/presentation/arp>
10. S. Axelsson, "The base-rate fallacy and the difficulty of intrusion detection," *ACM Trans. Inf. Syst. Security (TISSEC)*, vol. 3, no. 3, pp. 186–205, Aug. 2000, doi: 10.1145/357830.357849.
11. F. Barbero, F. Pendlebury, F. Pierazzi, and L. Cavallaro, "Transcending TRANSCEND: Revisiting malware classification in the presence of concept drift," in *Proc. IEEE Symp. Secur. Privacy (S&P)*, 2022, pp. 805–823, doi: 10.1109/SP46214.2022.9833659.
12. J. G. Moreno-Torres, T. Raeder, R. Alaíz-Rodríguez, N. V. Chawla, and F. Herrera, "A unifying view on dataset shift in classification," *Pattern Recognit.*, vol. 45, no. 1, pp. 521–530, 2012, doi: 10.1016/j.patcog.2011.06.019.
13. K. Allix, T. F. Bissyandé, J. Klein, and Y. Le Traon, "AndroZoo: Collecting millions of Android apps for the research community," in *Proc. Int. Conf. Mining Softw. Repositories (MSR)*, ACM, Jan. 2016, pp. 468–471.
14. R. Harang and E. M. Rudd, "SOREL-20M: A large scale benchmark dataset for malicious PE detection," 2020, *arXiv:2012.07634*.
15. F. Pierazzi, F. Pendlebury, J. Cortellazzi, and L. Cavallaro, "Intriguing properties of adversarial ML attacks in the problem space," in *Proc. IEEE Symp. Secur. Privacy (S&P)*, 2020, pp. 1332–1349, doi: 10.1109/SP40000.2020.00073.

Lorenzo Cavallaro is a full professor of computer science at University College London, WC1E 6BT London, U.K., where he leads the Systems Security Research Lab. His research focuses on understanding and improving the effectiveness of machine learning (ML) methods for systems security in the presence of adversaries. In particular, he and his lab investigating the intertwined relationships

of program analysis and machine learning and the implications they have toward realizing trustworthy ML for systems security. Cavallaro received a Ph.D. in computer science from the University of Milan, Italy. Contact him at l.cavallaro@ucl.ac.uk.

Johannes Kinder is a full professor of computer science at Bundeswehr University Munich, 85577 Neubiberg, Germany, where he is a member of the Cyber Defense Research Institute and the Institute for Systems Security. His research focuses on program analysis, comprehension, and transformation for security. Kinder received a PhD in computer science from TU Darmstadt, Germany. He is a member of Association for Computing Machinery and the German GI. Contact him at johannes.kinder@unibw.de.

Fergus Pendlebury is a visiting scholar at University College London, WC1E 6BT London, U.K. His research interests include developing machine learning techniques for hostile environments where adversaries seek to evade detection. Pendlebury received a Ph.D. in computer science from Royal Holloway, University of London, U.K. Contact him at fpindlebury@gmail.com.

Fabio Pierazzi is a lecturer (assistant professor) of computer science and the deputy head of the Cybersecurity group, Department of Informatics, King's College London, WC2R 2LS London, U.K. His research interests include systems security and machine learning, with a particular emphasis on settings in which attackers adapt quickly to new defenses. Pierazzi received a Ph.D. from the University of Modena, Italy. He is a Member of IEEE. Contact him at fabio.pierazzi@kcl.ac.uk.