Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing <u>scholarworks-group@umbc.edu</u> and telling us what having access to this work means to you and why it's important to you. Thank you.

Leveraging Uncertainty for Improved Static Malware Detection Under Extreme False Positive Constraints

Andre T. Nguyen^{1,2,3}, Edward Raff^{1,2,3}, Charles Nicholas³ and James Holt¹

¹Laboratory for Physical Sciences

²Booz Allen Hamilton

³University of Maryland, Baltimore County {andre,edraff}@lps.umd.edu, nicholas@umbc.edu, holt@lps.umd.edu

Abstract

The detection of malware is a critical task for the protection of computing environments. This task often requires extremely low false positive rates (FPR) of 0.01% or even lower, for which modern machine learning has no readily available tools. We introduce the first broad investigation of the use of uncertainty for malware detection across multiple datasets, models, and feature types. We show how ensembling and Bayesian treatments of machine learning methods for static malware detection allow for improved identification of model errors, uncovering of new malware families, and predictive performance under extreme false positive constraints. In particular, we improve the true positive rate (TPR) at an actual realized FPR of 1e-5 from an expected 0.69 for previous methods to 0.80 on the best performing model class on the Sophos industry scale dataset. We additionally demonstrate how previous works have used an evaluation protocol that can lead to misleading results.

1 Introduction

Classifying a new file as benign (safe to run) or malicious (not safe, termed "malware") is a current and growing issue. Malware already causes billions in damages [Anderson et al., 2019; Hyman, 2013], and with healthcare systems increasingly targeted has directly led to deaths [Eddy and Perlroth, 2021]. For years most anti-virus (AV) vendors have been seeing at least 2 million malicious new files per month [Spafford, 2014], and benign files on a network tend to outnumber malicious files at a ratio of 80:1 [Li et al., 2017]. This creates a common need for malware detection systems to operate with extremely low false positive rates. If false positives are too frequent, then analysts, IT, and support staff have to spend too much work on non-threats while simultaneously interrupting normal workflow. Even with this focus, Computer Incident Response Teams (CIRT) are often dealing with over 50% false positive rates and cite it as the main issue with current tooling [NISC, 2020].

The natural goal for AV style deployments of a malware detector is to maximize the true positive rate (TPR) for some maximally acceptable false positive rate (FPR). Production deployments are often concerned with FPRs of 0.1% at most, and preferably $\leq 0.001\%$. The issue of low FPR has been recognized since the very first research on machine learning based malware detection [Kephart *et al.*, 1995], yet surprisingly little work has been done to study how to maximize TPR@FPR. We present the first work addressing this gap by applying ideas from ensembling and Bayesian uncertainty estimation to a variety of common malware detection methods in use today on the two largest public corpora. We develop a number of contributions and previously unrealized insights:

1) All prior malware detection work we are aware of have evaluated TPR@FPR incorrectly or or did not specify their approach. The common error is to measure the TPR at the desired FPR on the test set, but this presupposes knowledge of the exact threshold to achieve the desired FPR. By instead estimating the threshold on a validation set, we show prior results have often misidentified their true TPR rates.

2) While the benefits of ensembling have long been known, it is often presumed that significant model diversity is required to obtain meaningful benefit. We show even moderately diverse or Bayesian approaches can significantly improve the TPR, especially for the low-FPR regimes needed for malware detection.

3) By using a Bayesian approach to estimate the epistemic and aleatoric uncertainty of a model on a given sample, we develop a new approach to thresholding a model's decision that can improve TPR and better approximate the desired FPR on unseen data.

4) Malware detection deployment requires detecting novel malware families, an intrinsically out-of-distribution task. We show how epistemic and aleatoric uncertainty relates to errors and novel malware families, allowing for faster detection of new threats.

The rest of this work is organized as follows. First we will review the related research to our work in section 2. Next we will detail the data, algorithms, and metrics used in all of our experiments in section 3. We present extensive experimental results organized by our major contributions in section 4, which show that prior TPR estimates could be off by $\geq 35\%$ relative error, that ensembles of limited diversity can raise TPR rates by $\geq 11\%$. Then we leverage uncertainty estimates to show a statistically significant improvement over the naive approach of thresholding models for TPR/FPR trade-offs, and that our uncertainty estimates are useful to malware analysts in identifying mispredicted files. Our conclusions are presented in section 5.

2 Related Work

The need for low FP rates has been paramount since the inception of machine learning malware detection research by Kephart et al. [1995]. Much of the history in this domain is focused on "signature" like tasks, where the goal was to recognize the set of already known malware, smaller than the total population of malware. This led to works that used the training data as part of the final evaluation data [Gavrilut et al., 2012; Fukushima et al., 2010]. This approach is not meaningful for determining TPR at any FPR due to overfitting, and is not tenable due to the now large and growing population of malware with more sophisticated obfuscation techniques. There is no agreed-upon threshold for exactly how low FPs should be, with most published work ranging from 0.1% down to $\leq 0.00002\%$ [Rafique and Caballero, 2013; Smutz and Stavrou, 2012; Pitsillidis et al., 2010; Anderson et al., 2016; Perdisci et al., 2008; Kolbitsch et al., 2009; Saxe and Berlin, 2015; Raff et al., 2019]. Beyond some works using training data at test time, all of these works evaluate their false positive rates on the test set, selecting the threshold from the test set that gives them the desired FPR, and then report the associated TPR. This is an understandable but incorrect approach, because the threshold is selected explicitly from the test set, when our goal is to test the ability of the model to achieve an FPR on unseen data. As far as we are aware, our work is the first in the malware detection space to identify this and propose selecting the threshold from a validation set, and then evaluate the precision of the FPR estimate in conjunction with the obtained TPR.

It is also worth noting that these cited prior works attempt to minimize FPR primarily by feature selection, engineering, or ML process pipeline choices that they believe will result in a more accurate model or be biased towards low FPR. Our approach is instead model agnostic, and seeks to better understand the nature of selecting thresholds to achieve TPR@FPR targets and improving it with uncertainty estimates. The only other work we are aware of that has this goal is from the related domain of spam detection by Yih et al. [2006] who propose two dataset re-weighting approaches, but also determine their success using the test set. Because they focus on data reweighting their approach is orthogonal to our own, and both could be used simultaneously. The closest other work we are aware of is [Eban et al., 2017] that develop differentiable approximations to AUC and Precision at a minimum Recall, but their approach does not apply to our situation because we instead need a maximum FP rate and specific points on the AUC curve. Further, our need for very low FPR is problematic for their setting as a mini-batch of data will be insufficient for estimating low FPR.

A number of prior works have investigated diverse ensembles of different kinds of algorithms to improve malware detection accuracy [Ye *et al.*, 2010; Liu and Wang, 2016; Kang *et al.*, 2016; Menahem *et al.*, 2009; Khasawneh *et al.*, 2015], following the common wisdom that averaging over diverse and uncorrelated predictors improves accuracy [Wolpert, 1992; Breiman, 1996; Jacobs *et al.*, 1991]. As far as we are aware, we are the first to study the performance of small ensembles of low-diversity (i.e, different runs of the same algorithm) and identify their especially large impact on TPR when needing extremely low FPR. This is important for malware detection as a diverse ensemble often necessitates algorithms that are too slow for deployment, and high compute throughput is critical to practical utility in this domain.

Building upon the use of ensembles, the modeling of uncertainty for decision making is notably missing from the current machine learning for malware literature. An exception is the approach of Backes and Nauman [2017] which introduces new classes for uncertain decisions in the context of a simple Bayesian logistic regression model applied to the Drebin Android malware dataset [Arp et al., 2013; Spreitzenbarth et al., 2013]. Understanding when a machine learning model is uncertain about its prediction is critical in high risk applications such as malware detection. When an automated malware detection algorithm is uncertain about a sample, the uncertainty estimate could be used to flag the sample for analysis by a more computationally expensive algorithm or for review by a human. Our work is the first we are aware to study how modeling uncertainty can be used to improve TPR@FPR scores, and to aid analysts by showing new samples with high uncertainty are more likely to be novel malware families.

3 Methods

We provide details about the data and machine learning models used in our experiments. The majority of the existing research in machine learning applied to malware detection has focused on the automation of static malware analysis, where a file is analyzed without being run. We will also focus on the static analysis problem.

3.1 Data

Due to the need to estimate low FPR rates, we use the two largest available corpora for malware detection. These are the EMBER2018 and Sophos AI SOREL-20M datasets. We note that both of these datasets focus on low FPR evaluation, but make the same error in evaluation. Our first results in section 3 will show the relative magnitudes of the errors.

We use the EMBER2018 dataset which consists of portable executable files (PE files) scanned by VirusTotal in or before 2018 [Anderson and Roth, 2018]. The dataset contains 600,000 labeled training samples and 200,000 labeled testing samples, with an equal number of malicious and benign samples in both sets. The malicious samples are also labeled by malware family using AVClass [Sebastián *et al.*]. All of the testing samples. EMBER2018 includes vectorized features for each sample encoding general file information, header information, byte histograms, byte-entropy histograms, and string

information [Anderson and Roth, 2018]. While the 1.1TB of raw PE files are not available as part of EMBER2018, they can be downloaded via VirusTotal. We note that the EM-BER2018 dataset was designed to be more challenging for machine learning algorithms to classify than the original EM-BER2017 dataset.

We also use the recent Sophos AI SOREL-20M dataset, consisting of 20 million files [Harang and Rudd, 2020]. The much larger number of data points in the Sophos dataset is advertised as "industry scale" and allows for the exploration of FPR constraints much smaller than allowed by EMBER2018. In particular, the test set size for Sophos consists of 1,360,622 malicious samples and 2,834,441 benign samples. As part of the Sophos dataset release, two baseline models are provided, including a feed-forward neural network (FFNN) and a LightGBM (LGBM) gradient-boosed decision tree model. Five versions of each of the models are pre-trained using different random seeds on the Sophos data using the same featurization as EMBER2018.

3.2 Models

EMBER2018: We apply three models to the EMBER2018 dataset that each rely on different types of features. The first model we apply is a Bayesian deep learning model based on the MalConv model of Raff *et al.* [2017], a convolutional neural network for malware detection that operates on the raw byte sequences of files. We will refer to this model as Bayesian MalConv (BMC). As exact Bayesian inference cannot be done for deep neural networks such as MalConv, approximate inference methods need to be used.

Gal and Ghahramani [2016] introduced an easy to implement approach to variational inference in Bayesian neural networks. In particular, they showed that a neural network with dropout, a technique commonly used to reduce overfitting in neural networks by randomly dropping units during training [Hinton et al., 2012; Srivastava et al., 2014], applied before every weight layer is equivalent to an approximation of a deep Gaussian process [Damianou and Lawrence, 2013], and that training with dropout effectively performs variational inference for the deep Gaussian process model. The posterior distribution can be sampled from by leaving dropout on at test time. For Bayesian MalConv, we follow this approach and apply dropout before each fully connected layer of the MalConv model, with a dropout probability of p = 0.1. We use the Adam optimizer [Kingma and Ba, 2014] to train the model, and we produce 16 samples at evaluation time using multiple forward passes on the trained model with dropout left on.

The second model we apply is a Bayesian logistic regression (BLR) model which takes as input the binary presence of 94,225 byte 8-grams [Raff *et al.*, 2019, 2018] selected using LASSO from the one million most common byte 8-grams. Dropout is used in a similar manner to Bayesian MalConv, with a dropout probability of p = 0.1 and 16 samples at evaluation time.

The third model we apply is an ensemble of 16 Light Gradient Boosting Machine (LGBM) models [Ke *et al.*] trained with different random seeds on the EMBER features as described in Anderson and Roth [2018].

Sophos: We apply two models to the Sophos dataset that both rely on the EMBER2018 featurization. The first is an ensemble of 5 feed-forward neural network (FFNN) models as described in Harang and Rudd [2020], a simplified version of the model from Rudd et al., trained using different random seeds. The second is an ensemble of 5 LGBM models trained using different random seeds. We use the publicly available pre-trained models provided with the Sophos dataset as our ensemble members for both models. While 5 ensemble members may seem small, Ovadia et al. [2019] found that increasing ensemble sizes beyond 5 has diminishing returns with respect to the quality of the uncertainty estimates, so an ensemble size of 5 may be sufficient. Our results will show that not only are they sufficient, but for our goal of low FPR they can be significantly more effective than has been previously reported.

3.3 Uncertainty Estimation

The Bayesian framework allows for the principled modeling of uncertainty in machine learning and decision making. Within this framework, probabilities represent degrees of belief as opposed to the frequentist interpretation of probabilities as long run frequencies [Liu and Wasserman, 2014]. Bayesian inference uses Bayes' Theorem to update beliefs (that are represented in the form of probability distributions) when new data is observed.

In the context of machine learning, a Bayesian update takes the following form where θ represents model parameters, D represents the data, and M represents the model class: $\mathbb{P}(\theta|D,M) = \frac{\mathbb{P}(D|\theta,M)\mathbb{P}(\theta|M)}{\mathbb{P}(D|M)}$, where $\mathbb{P}(\theta|D,M)$ is the posterior belief about the model parameters given the data, $\mathbb{P}(D|\theta, M)$ is the likelihood of the data given the model parameters, $\mathbb{P}(\theta|M)$ is the prior belief about model parameters, and $\mathbb{P}(D|\theta,M)$ is the marginal likelihood or evidence. Bayesian inference is usually intractable due to the integrals involved, unless the prior distribution is conjugate to the likelihood distribution. Unfortunately, conjugate priors exist for only exponential family distributions [Murphy, 2012] and so can't be directly applied to complex models like Bayesian deep neural networks.

As exact Bayesian inference cannot be done for Bayesian deep learning models, approximate inference methods need to be used. Given sufficient compute time, Markov Chain Monte Carlo (MCMC) methods can be used to sample from the posterior [Neal, 1995]. Unfortunately, common sampling based approaches are difficult to scale to problems in the malware space where practical dataset sizes are measured in terabytes because they require gradient computations over the entire dataset. As MCMC is hard to scale in practice, variational inference is often used instead which converts the integration problem into an optimization problem where the posterior is approximated using a simpler variational distribution [Blei et al., 2017]. Variational inference for neural networks was first introduced in the early nineteen nineties [Hinton and van Camp, 1993], and Graves [2011] revived interest in variational inference for neural networks by introducing a stochastic variational method for inference in neural networks.

We note that complicated Bayesian inference is not necessarily needed to provide useful uncertainty estimates. Lakshminarayanan *et al.* [2017] introduce an alternative that trains an ensemble of randomly initialized models. These deep ensembles have been shown to produce competitive uncertainty estimates [Ovadia *et al.*, 2019; Ashukha *et al.*, 2020] because they are able to explore different modes in function space [Fort *et al.*, 2019]. Wilson and Izmailov [2020] argue that deep ensembles are not a competing approach to Bayesian deep learning but rather are an effective approach for Bayesian model averaging.

Two kinds of uncertainty can be distinguished [Gal, 2016]. Aleatoric uncertainty is caused by inherent noise and stochasticity in the data. More training data will not help to reduce this kind of uncertainty. Epistemic uncertainty on the other hand is caused by a lack of similar training data. In regions lacking training data, different model parameter settings that produce diverse or potentially conflicting predictions can be comparably likely under the posterior.

For classification tasks where epistemic and aleatoric uncertainty don't need to be differentiated, uncertainty can be measured using the predictive distribution entropy:

$$H[\mathbb{P}(y|x,D)] = -\sum_{y \in C} \mathbb{P}(y|x,D) \log \mathbb{P}(y|x,D)$$

Aleatoric uncertainty is measured using expected entropy:

$$u_{alea} = \mathbb{E}_{\mathbb{P}(\theta|D)} H[\mathbb{P}(y|x,\theta)]$$

Mutual information is used to measure epistemic uncertainty:

$$u_{epis} = I(\theta, y | D, x) = H[\mathbb{P}(y | x, D)] - \mathbb{E}_{\mathbb{P}(\theta | D)} H[\mathbb{P}(y | x, \theta)]$$

Monte Carlo estimates obtained by sampling from the posterior can be used to approximate the terms of these equations for our Bayesian models [Smith and Gal, 2018]. In particular, $\mathbb{P}(y|x, D) \approx \frac{1}{T} \sum_{i=1}^{T} \mathbb{P}(y|x, \theta_i)$ and $\mathbb{E}_{\mathbb{P}(\theta|D)} H[\mathbb{P}(y|x, \theta)] \approx \frac{1}{T} \sum_{i=1}^{T} H[\mathbb{P}(y|x, \theta_i)]$ where the θ_i are samples from the posterior over models and T is the number of samples.

For our ensemble based models which are not explicitly Bayesian (because each ensemble member receives the same weight) but Bayesian inspired, uncertainties can be computed in a similar way where the θ_i are no longer samples from a posterior, but instead multiple independent trainings of a model with T different random seeds.

3.4 Classification Metrics

We use multiple metrics to evaluate and compare approaches. Accuracy is defined as the percent of correct predictions made. Area under the receiver operating characteristic curve (AUC) is the probability that the classifier will rank a randomly selected malicious file higher in probability to be malicious than a randomly selected benign file. The true positive rate (TPR) is defined as the number of true positives over the sum of true positives and false negatives. The false positive rate (FPR) is defined as the number of false positives over the sum of false positives and true negatives.

An important contribution of our work is to recognize that the TPR obtained at any given FPR on the test set is not the actual measure of interest in malware detection, but an overfit measure due to the implicit assumption that the correct decision threshold is known at test time. The threshold must be estimated during training or validation, and then applied to the test set. This means we have a target maximum FPR T_{FPR} that we wish to obtain, and a separate *actualized FPR* that is obtained on the test set. In order to capture the tradeoff between TPR and actualized FPR constraint satisfaction, we define the following combined metric Equation 1 where T_{FPR} is the desired maximum FPR.

$$C = TPR - \frac{\max(\text{actualized } FPR - T_{FPR}, 0)}{T_{FPR}} \quad (1)$$

This metric captures that we have a desired TPR, but penalizes the score based on the degree of violation of the FPR. This is done by a division so that the magnitude of the violation's impact grows in proportion to the target FPR shrinking. This matches the nature of desiring low FPR itself. For example, 90% TPR at a target FPR of 0.1% is still quite good if the actualized FPR is 0.11% (C = 0.8), but is unacceptably bad if the target FPR was 0.01% (C = -9.1).

4 Experiments and Discussion

Now that we have discussed the methods of our work and the metrics by which they will be examined, we will show empirical results demonstrating our primary contributions: 1) Evaluating test-set performance thresholds from the test set leads to misleading results at lower FPR, 2) Simple non-diverse ensembles can dramatically improve TPR at any given FPR rate, 3) we can further improve TPR@FPR by explicitly modeling Bayesian uncertainty estimates into our decision process, and 4) these uncertainty estimates have practical benefits to application by showing that errors and previously unseen malware families have uncertainty distributions that place more weight on higher uncertainties. For each of these we will include the empirical results on the EMBER2018 and the Sophos 2020 corpora, and include additional discussion and nuance to how these relate to practical deployment.

4.1 Misleading Evaluation

A currently accepted practice for evaluating malware detection models under FPR constraints is to report the test set ROC curve. Once the test set ROC curve is produced, the desired FPR rates from the curve are selected to show their associated TPR. This is misleading as in practice the test set is not available when choosing the decision threshold, causing this evaluation procedure to be invalid. Instead, we must recognize that there are a priori target FPRs that are the FP rates that we desire from the model, and the actualized FPRs which are what is obtained on the test (read, "production") data. Selecting the threshold from the test set hides that the target and actualized FPRs are different, especially for low FPRs that require large amounts of data to estimate. The valid approach to this scenario when evaluating a classifier at different FPRs is to select the thresholds using a validation set. Once the thresholds are selected that obtain the target FPRs, they can be applied to the test set to obtain the actualized FPRs and their associated TPRs. We show the impact this has on the entire TPR/FPR curve in Figure 1 which shows the absolute relative error in TPR for a given actualized FPR. Depending on the model and dataset, the resulting TPR for any



Figure 1: Absolute relative error in TPR when using the invalid evaluation protocol, for three different model and dataset combinations. A valid evaluation protocol will use a validation set ROC curve to select a threshold given a desired FPR. The evaluation protocol that is currently the norm in the malware detection literature is invalid because it uses the test set ROC curve, which is never available in practice, to select a threshold. The use of the invalid evaluation protocol can lead to over a 30 percent relative error in TPR.

actualized FPR can change by over 30%, and the relative error generally increases as the FPR decreases. This is expected because low FPRs naturally require more data to estimate: if you want an FPR of 1:1,000 and you want 1,000 FPRs to estimate the threshold from you would expect to need $1,000^2 = 1$ million examples.

We note that the Sophos FFNN model seems to be particularly robust with the lowest error in Figure 1. This is in part a testament to the FFNN approach, but more broadly a function on the magnitude of the Sophos dataset. With 2.5 million samples in the validation set and 4.2 million in the test set, the corpus is large enough to mitigate the impact of some inappropriate practices. To demonstrate the impact the validation set can have, we show the same results in Figure 2 when only the validation set used to select the threshold is reduced by various orders of magnitude.

One can clearly see that as the validation set size decreases, the ability to estimate the FPR decreases. This causes more errors and a "shortening" of the curves as it becomes impossible to estimate lower desired FPR rates. This last point is important as some prior works have reported FPRs lower than what their dataset could accurately estimate. If the test set size times the desired FPR is less than 100 samples, it is unlikely the TPR@FPR reported will be an accurate estimate (e.g., as done in [Anderson *et al.*, 2016]).

We note that this distinction between invalid vs. valid approaches is *not* a critique on the evaluation of entire ROC curves. The fundamental distinction is whether we care about



Figure 2: Absolute relative error in TPR when using the invalid evaluation protocol at various levels of subsampling of the validation set.

the entire ROC curve, or only specific points from the ROC curve. If you care about the entire ROC curve, evaluating the ROC on the test set is valid and appropriate. But because malware detection is concerned with particular points from the ROC curve, it becomes necessary to evaluate if the approach can hit its desired location on the curve (i.e., a specific FPR in production). There are also valid scenarios to consider just the ROC curve as a whole for malware analysis and its associated AUC, as it represents a metric of ability to rank that is applicable to other scenarios within malware detection and analysis. Our critique is for just those concerned with AV-like deployments that aim for low FPRs specifically.

4.2 Ensembles

We have now shown that the correct approach to developing a ROC curve when one wishes to evaluate specific points on the curve is to select the threshold from a validation set rather than the test set. We will apply this to the results of this section to show that creating an ensemble of randomly seeded models can improve the obtained TPR at almost any actualized FPR, especially under extreme FPR constraints. Figure 3 shows the ROC curves for individual models as well as for the ensemble consisting of those individual models. The Sophos trained FFNN ensemble notably performs significantly better than any individual member of the ensemble, with the gap in performance widening as FPR becomes smaller.

The fact that these are all the same type of model, but with different random seeds at initialization, is an important and not previously recognized phenomena. Classical wisdom is that ensembles should maximize diversity to reduce correlation of predictions, and thus maximize accuracy. But in our scenario malware detection models are designed to be lightweight in model size, low latency, and high throughput, so that the AV system does not interrupt the users of a computing system. A classically diverse model with different types of algorithms or features, as done in prior work in



Figure 3: Sophos ensembles that take the average of predictions from randomly seeded models can lead to significant TPR gains under extreme FPR constraints, compared to individual models.

this space, ends up including approaches that are many orders of magnitude larger and slower than the lighter weight approaches we study in this work. Because we can use multiple versions of the same type of model with the same features, we can maintain the high throughput, low latency & size requirements while obtaining these large gains in TPR.

Table 1 compares the accuracy, AUC, and AUC @ $\leq 0.1\%$ FPR achieved by ensembles to the average of those achieved by individual ensemble members. In all cases, the ensemble has better performance than the expected performance of individual ensemble members, even though they are using an ensemble of low diversity. Of particular importance is the performance of each model in the low FPR domain has a greater relative improvement (median 11% improvement) than one may have anticipated looking at more standard metrics like Accuracy and AUC (median improvements of 0.6% and 0.4% respectively). The only exception to this is the Logistic Regression approaches which have difficulty operating at the extremely low FPR ranges, which we will see repeated.

4.3 Uncertainty Based Threshold Adjustments

While ensembling predictions by taking an average leads to improved results, there is more information within ensembles that can be leveraged to further ameliorate performance under extreme FPR constraints. In particular, estimates can be computed for epistemic (u_{epis}) and aleatoric (u_{alea}) uncertainty. We introduce a simple threshold adjustment approach that leverages data point specific uncertainty to locally adjust the threshold based on the uncertainty. We explore three uncertainty aware local threshold adjustments:

Table 1: Accuracy and AUC for each model and dataset combination. Ensembles are compared to the expected performance of their components. Best results are shown in **bold**.

Dataset	Model	Accuracy	AUC	$\mathrm{AUC}_{\leq 0.1\% FPR}$	
EMBER	Bayesian MalConv	91.64	97.47	0.04079	
	MalConv	90.88	97.05	0.03288	
	Bayesian Log. Reg.	94.72	98.15	0.0	
	Log. Reg.	94.15	97.32	0.0	
	LightGBM Ensemble	93.98	98.62	0.06054	
	LightGBM	93.88	98.55	0.05433	
Sophos	FFNN Ensemble	98.81	99.83	0.09274	
	FFNN	98.56	99.75	0.08990	
	LightGBM Ensemble	86.10	98.41	0.04459	
	LightGBM	85.47	98.05	0.03637	

$$\hat{y}_{adj} = \hat{y} + \alpha_1 \cdot u_{epis} + \alpha_2 \cdot u_{alea} \tag{2}$$

$$\hat{y}_{adj} = \hat{y} + \alpha_1 \cdot \exp(\alpha_3 \cdot u_{epis}) + \alpha_2 \cdot \exp(\alpha_4 \cdot u_{alea})$$
(3)

$$\begin{aligned}
\hat{y}_{adj} &= \hat{y} + \mathbb{I}[\hat{y} > \alpha_0](\alpha_1 \cdot u_{epis} + \alpha_2 \cdot u_{alea}) + \\
& \mathbb{I}[\hat{y} \le \alpha_0](\alpha_3 \cdot u_{epis} + \alpha_4 \cdot u_{alea})
\end{aligned} \tag{4}$$

where $\mathbb{1}[\cdot]$ is the indicator function, \hat{y} is the original ensemble prediction for a data point, u_{epis} is the epistemic uncertainty (mutual information) for a data point's prediction, u_{alea} is the aleatoric uncertainty (expected entropy) for a data point's prediction, the α_i are learned scaling factors, and \hat{y}_{adj} is the uncertainty adjusted prediction for a data point. The scaling factors are learned by iteratively optimizing each α_i to maximize TPR given a desired FPR, where after each scaling factor adjustment a new global adjustment is computed.

Because TPR@FPR is not a differentiable metric, we use a gradient free approach to altering the weights α . In particular, we use a coordinate descent style approach where we take turns optimizing the individual α_i values while holding all others fixed, repeating the process until convergence. This is feasible thanks to the convex behavior that occurs with respect to the TPR scores. If α_i is set too large, then the associated feature (e.g, u_{epis} or u_{alea}) becomes the only effective factor by overshadowing all other components, but is not sufficient on its own to make meaningful predictions, resulting in a low TPR when selecting the target FPR. If the weight is too small then the associated feature has no impact, and the result is unchanged. This creates the two "low points," and a weight that results in a higher score (hopefully) exists between the two extrema. This can then be selected with high precision by using a golden search by treating the extrema as brackets on a solution. We use Brent's method to solve this because it allows faster searches by approximating the problem with a parabola when possible, and switching to golden search in the worst case, allowing it to solve the optimization quickly.1

This gives us an approach to directly optimize our target metric even though it is non-differentiable, and to do so with high precision in just a few minutes of computation. All optimization occurs on the validation set. While the Sophos

¹In our experience each call to Brent's method takes less than 100 optimization steps.

Table 2: Sophos comparison of the standard global adjustment (labeled as (g)) to the uncertainty aware local adjustments from Equation 2 (labeled as (g+l)) and Equation 3 (labeled as (g+lv2)) and Equation 4 (labeled as (g+lv3)). Best combined score (TPR penalized for over-runs on target FPR) shown in **bold**.

Target FPR	Test Perf.	LGBM				FFNN				
		(g)	(g+l)	(g+lv2)	(g+lv3)	(g)	(g+l)	(g+lv2)	(g+lv3)	
1%	TPR	8.060E-01	8.125E-01	8.150E-01	8.137E-01	9.779E-01	9.779E-01	9.779E-01	9.779E-01	
	FPR	1.175E-02	1.123E-02	1.125E-02	1.129E-02	8.664E-03	8.663E-03	8.666E-03	8.665E-03	
	Comb.	6.315E-01	6.899E-01	6.904E-01	6.845E-01	9.779E-01	9.779E-01	9.779E-01	9.779E-01	
0.1%	TPR	5.264E-01	5.318E-01	5.343E-01	5.342E-01	9.440E-01	9.471E-01	9.450E-01	9.450E-01	
	FPR	1.493E-03	1.088E-03	9.699E-04	9.681E-04	9.695E-04	9.473E-04	1.024E-03	1.022E-03	
	Comb.	3.338E-02	4.434E-01	5.343E-01	5.342E-01	9.440E-01	9.471E-01	9.208E-01	9.233E-01	
0.01%	TPR	2.296E-01	2.352E-01	2.371E-01	2.278E-01	9.017E-01	9.037E-01	9.086E-01	9.086E-01	
	FPR	4.339E-05	5.751E-05	5.786E-05	4.375E-05	8.855E-05	9.032E-05	8.961E-05	9.102E-05	
	Comb.	2.296E-01	2.352E-01	2.371E-01	2.278E-01	9.017E-01	9.037E-01	9.086E-01	9.086E-01	
0.001%	TPR	9.940E-02	1.007E-01	1.075E-01	1.007E-01	8.022E-01	8.046E-01	8.043E-01	8.041E-01	
	FPR	3.881E-06	4.586E-06	4.939E-06	4.586E-06	4.234E-06	5.292E-06	5.998E-06	4.586E-06	
	Comb.	9.940E-02	1.007E-01	1.075E-01	1.007E-01	8.022E-01	8.046E-01	8.043E-01	8.041E-01	

dataset comes with a validation set, EMBER2018 does not, so we create a validation set for EMBER2018 using five fold cross-validation. The α_i are optimized in an alternating manner for Equation 2 and in a randomized sequential manner for Equation 3 and Equation 4. Brent's method is used as the optimizer [Brent, 1972], with a bracketing interval of [-100, 100] for Equation 2, of [-10, 10] for Equation 3, and of [0, 1] for Equation 4 with [-0.1, 0.1] for α_0 . A ROC curve is then computed using the locally adjusted \hat{y}_{adj} to obtain the final global threshold. For all methods, when fitting the global and local threshold adjustments using the validation set, a target FPR of 0.9 times the actual desired FPR is used in order to conservatively satisfy the constraint given differences in the test set data.

We briefly note that we had also tried optimizing the uncertainty scaling factors and thresholds jointly using a CMA-ES gradient-free optimization approach [Hansen and Ostermeier]. Unfortunately, the obtained solutions were not precise enough given the extremely small FPR constraints, leading us to use the iterative optimization of each variable using Brent's method.

The results comparing our new uncertainty augmented local adjustments Equation 2 (g+l) and Equation 3 (g+lv2) and Equation 4 (g+lv3) against the naive approach (g) are provided in Table 2 for the Sophos dataset. Additional results for EMBER2018 are provided in Appendix Table 3. Bolded are the methods that performed best by the combined score Equation 1 which penalizes going over the target FPR. We note that across datasets, models, and target constraints, the inclusion of uncertainty based local adjustments (g+l and g+lv2 and g+lv3) improves over the standard use of only a global threshold adjustment (g). All three approaches are statistically significant in their improvement (Wilcoxon-signed rank test p-values of 0.02, 0.01, and 0.025 respectively). In almost all cases if (g+lv2) is the best performer, (g+l) and (g+lv3) are the second and third best. Similarly, when (g+lv2) is not the best, it is usually still better than (g).

4.4 Uncertainty on Errors and New AV Classes

Our local threshold adjustments using epistemic and aleatoric uncertainty estimates show improved TPR for extremely low target FPRs. We further investigate how this is possible, and in doing so show that these uncertainty estimates provide an additional benefit to practical application. We observe that



Figure 4: A comparison of uncertainty distributions for the EM-BER2018 models at test time between malware families seen and unseen during training.

the errors of our models are correlated with both uncertainty measures. This means we can use the uncertainty measures not just as a means of adjusting the threshold, but as a sourcing tool for analysts. The data with the highest uncertainty scores are the most likely to be mispredicted, and thus guide the analysts to the samples where their time is best spent. Additional details can be found in the Appendix, and next we focus on an extended version of this result, in that uncertainties are also correlated with a sample being novel malware or not.

Of the 200,000 files in the EMBER2018 test set, 363 belong to new malware families that were not present in the train set (we note that all the test set are new files that did not exist prior, as the train/test split is a split in time). Figure 4a and Figure 4b show that the Bayesian MalConv and LightGBM ensemble uncertainty distributions for test samples from malware families seen during training place most of their weight on lower uncertainty values, whereas the uncertainty distributions for novel families not seen during training place most of their weight on higher uncertainty values. We however found that the Bayesian logistic regression model mostly does not exhibit this behavior, likely due to the simplicity of the model class which limits the extent to which predictions can disagree. Overall, these results suggest that for some models, uncertainties can be leveraged for the detection of new, out of training distribution, malware families.

5 Conclusions

We have provided evidence that uncertainty estimation using ensembling and Bayesian methods can lead to significant improvements in machine learning based malware detection systems. In particular, improvements were especially large under extreme false positive rate constraints which are common in deployed, production scenarios. Local uncertainty based threshold adjustments were shown to lead to higher TPRs while satisfying desired FPR maximums. We additionally demonstrated how previous works have used an evaluation protocol that can lead to misleading results, and how uncertainty can be used to better detect model errors and new malware families.

Obtaining uncertainties has an inherent additional computational cost at prediction time which may limit use in resource limited deployed contexts. However, recent advances such as BatchEnsemble [Wen *et al.*, 2020] have introduced new methods to avoid the computational and memory costs of naive ensembles.

We are currently working with professional malware analysts and teams that believe this approach may benefit them in production environments based on the evidence this work has provided. Future work includes leveraging uncertainty estimates to decide when to run more expensive malware analysis algorithms and techniques such as dynamic analysis, exploring and explaining malware specific drivers of uncertainty, and evaluating these methods over a long period of time in production.

References

- Hyrum S. Anderson and Phil Roth. EMBER: An Open Dataset for Training Static PE Malware Machine Learning Models. 2018.
- Blake Anderson, David Mcgrew, and Subharthi Paul. Discovering Human and Machine Readable Descriptions of Malware Families. *The AAAI-16 workshop on Artificial Intelligence for Cyber Sercurity (AICS)*, pages 150–156, 2016.
- Ross Anderson, Chris Barton, Rainer Boehme, Richard Clayton, Michel J G Van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the Changing Cost of Cybercrime. *Workshop on the Economics of Information Security* (*WEIS*), 2019.
- Daniel Arp, Michael Spreitzenbarth, H Malte, Hugo Gascon, and Konrad Rieck. Technical Report IFI-TB-2013-02 D REBIN : Efficient and Explainable Detection of Android Malware in Your Pocket. 2013.
- Arsenii Ashukha, Alexander Lyzhov, Dmitry Molchanov, and Dmitry Vetrov HSE. Pitfalls of In-Domain Uncertainty Estimation and Ensembling in Deep Learning. In International Conference on Learning Representations (ICLR), 2020.
- Michael Backes and Mohammad Nauman. LUNA: Quantifying and Leveraging Uncertainty in Android Malware Analysis through Bayesian Machine Learning. *Proceedings -*2nd IEEE European Symposium on Security and Privacy, EuroS and P 2017, pages 204–217, 2017.
- David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.

- Richard P Brent. Algorithms for Minimization Without Derivatives. 1972.
- Andreas C Damianou and Neil D Lawrence. Deep Gaussian Processes. 31, 2013.
- Elad Eban, Mariano Schain, Alan Mackey, Ariel Gordon, Ryan Rifkin, and Gal Elidan. Scalable Learning of Non-Decomposable Objectives. In AISTATS, volume 54 of Proceedings of Machine Learning Research, pages 832–840, 2017.
- By Melissa Eddy and Nicole Perlroth. Cyber Attack Suspected in German Woman's Death. *The New York Times*, Sept. 2021.
- Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv*, pages 1–15, 2019.
- Yoshiro Fukushima, Akihiro Sakai, Yoshiaki Hori, and Kouichi Sakurai. A behavior based malware detection scheme for avoiding false positive. In *IEEE Workshop on Secure Network Protocols*, pages 79–84, 2010.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. 33rd International Conference on Machine Learning, ICML 2016, 3:1651–1660, 2016.
- Yarin Gal. Uncertainty in Deep Learning. PhD thesis, University of Cambridge, 2016.
- Dragos Gavrilut, Razvan Benchea, and Cristina Vatamanu. Optimized Zero False Positives Perceptron Training for Malware Detection. In 2012 14th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, pages 247–253, 2012.
- Alex Graves. Practical variational inference for neural networks. Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011, NIPS 2011, pages 1–9, 2011.
- Nikolaus Hansen and Andreas Ostermeier. Completely Derandomized Self-Adaptation in Evolution Strategies. Technical report.
- Richard Harang and Ethan M. Rudd. SOREL-20M: A Large Scale Benchmark Dataset for Malicious PE Detection. 2020.
- Geoffrey E. Hinton and Drew van Camp. Keeping neural networks simple by minimizing the description length of the weights. pages 5–13, 1993.
- Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. pages 1–18, 2012.
- Paul Hyman. Cybercrime: it's serious, but exactly how serious? *Communications of the ACM*, 56(3):18–20, 2013.
- Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive Mixtures of Local Experts. *Neural Computation*, 3(1):79–87, 1991.

- Chanhyun Kang, Noseong Park, B Aditya Prakash, Edoardo Serra, and V S Subrahmanian. Ensemble Models for Datadriven Prediction of Malware Infections. In *Proceedings* of the Ninth ACM International Conference on Web Search and Data Mining, pages 583–592, 2016.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Light-GBM: A Highly Efficient Gradient Boosting Decision Tree. Technical report.
- Jeffrey O Kephart, Gregory B Sorkin, William C Arnold, David M Chess, Gerald J Tesauro, and Steve R White. Biologically Inspired Defenses Against Computer Viruses. In *IJCAI*, pages 985–996, 1995.
- Khaled N. Khasawneh, Meltem Ozsoy, Caleb Donovick, Nael Abu-Ghazaleh, and Dmitry Ponomarev. Ensemble Learning for Low-Level Hardware-Supported Malware Detection. In *Research in Attacks, Intrusions, and Defenses*. 2015.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. pages 1–15, 2014.
- Clemens Kolbitsch, Paolo Milani Comparetti, Christopher Kruegel, Engin Kirda, Xiaoyong Zhou, and XiaoFeng Wang. Effective and Efficient Malware Detection at the End Host. In *USENIX Security Symposium*. USENIX Association, 2009.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In Advances in Neural Information Processing Systems, 2017.
- Bo Li, Kevin Roundy, Chris Gates, and Yevgeniy Vorobeychik. Large-Scale Identification of Malicious Singleton Files. In *7TH ACM Conference on Data and Application Security and Privacy*, 2017.
- Liu Liu and Baosheng Wang. Malware classification using gray-scale images and ensemble learning. In *International Conference on Systems and Informatics*, 2016.
- Han Liu and Larry Wasserman. Bayesian Inference. In *Statistical Machine Learning*, chapter 12, pages 299–351. In Preparation., 2014.
- Eitan Menahem, Asaf Shabtai, Lior Rokach, and Yuval Elovici. Improving Malware Detection by Applying Multiinducer Ensemble. *Comput. Stat. Data Anal.*, 53(4):1483– 1494, 2009.
- Kevin P Murphy. *Machine learning: a probabilistic perspective (adaptive computation and machine learning series)*. The MIT Press, 2012.
- Radford M. Neal. *Bayesian Learning for Neural Networks*. PhD thesis, University of Toronto, 1995.
- NISC. NISC Survey Results. Technical report, Neustar International Security Council, 2020.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D Sculley, Sebastian Nowozin, Joshua V Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can You Trust Your Model's Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift. (NeurIPS), 2019.

- Roberto Perdisci, Andrea Lanzi, and Wenke Lee. McBoost: Boosting Scalability in Malware Collection and Analysis Using Statistical Classification of Executables. In Annual Computer Security Applications Conference (AC-SAC), 2008.
- Andreas Pitsillidis, Kirill Levchenko, Christian Kreibich, Chris Kanich, Geoffrey M. Voelker, Vern Paxson, Nicholas Weaver, and Stefan Savage. Botnet judo: Fighting spam with itself. In *Symposium on Network and Distributed System Security (NDSS)*, 2010.
- Edward Raff, Jon Barker, Jared Sylvester, Robert Brandon, Bryan Catanzaro, and Charles Nicholas. Malware Detection by Eating a Whole EXE. 2017.
- Edward Raff, Richard Zak, Russell Cox, Jared Sylvester, Paul Yacci, Rebecca Ward, Anna Tracy, Mark McLean, and Charles Nicholas. An investigation of byte n-gram features for malware classification. *Journal of Computer Virology and Hacking Techniques*, 14(1), 2018.
- Edward Raff, William Fleming, Richard Zak, Hyrum Anderson, Bill Finlayson, Charles K. Nicholas, and Mark Mclean. KiloGrams: Very Large N-Grams for Malware Classification. In *Proceedings of KDD 2019 Workshop on Learning and Mining for Cybersecurity (LEMINCS'19)*, 2019.
- M Zubair Rafique and Juan Caballero. FIRMA: Malware Clustering and Network Signature Generation with Mixed Network Behaviors. In *Proceedings of the 16th International Symposium on Research in Attacks, Intrusions, and Defenses - Volume 8145*, RAID 2013, pages 144–163, New York, NY, USA, 2013. Springer-Verlag New York, Inc.
- Ethan M Rudd, Felipe N Ducau, Cody Wild, Konstantin Berlin, and Richard Harang. *ALOHA: Auxiliary Loss Optimization for Hypothesis Augmentation*.
- Joshua Saxe and Konstantin Berlin. Deep neural network based malware detection using two dimensional binary program features. In 2015 10th International Conference on Malicious and Unwanted Software (MALWARE), 2015.
- Marcos Sebastián, Richard Rivera, Platon Kotzias, and Juan Caballero. AVCLASS: A Tool for Massive Malware Labeling. Technical report.
- Lewis Smith and Yarin Gal. Understanding measures of uncertainty for adversarial example detection. 34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018, 2:560–569, 2018.
- Charles Smutz and Angelos Stavrou. Malicious PDF detection using metadata and structural features. In *Proceedings* of the 28th Annual Computer Security Applications Conference, pages 239–248, 2012.
- Eugene C. Spafford. Is Anti-virus Really Dead? *Computers & Security*, 44:iv, 2014.
- Michael Spreitzenbarth, Felix Freiling, Florian Echtler, Thomas Schreck, and Johannes Hoffmann. Mobilesandbox: Having a deeper look into Android applications. *Proceedings of the ACM Symposium on Applied Computing*, pages 1808–1815, 2013.

- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal* of Machine Learning Research, 15:1929–1958, 2014.
- Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: An alternative approach to efficient ensemble and lifelong learning. arXiv:2002.06715, pages 1–20, 2020.
- Andrew Gordon Wilson and Pavel Izmailov. Bayesian Deep Learning and a Probabilistic Perspective of Generalization. (3), 2020.
- David H. Wolpert. Stacked generalization. *Neural networks*, 5:241–259, 1992.
- Yanfang Ye, Tao Li, Yong Chen, and Qingshan Jiang. Automatic Malware Categorization Using Cluster Ensemble. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 95–104, 2010.
- Scott Wen-tau Yih, Joshua Goodman, and Geoff Hulten. Learning at Low False Positive Rates. In *Proceedings of the 3rd Conference on Email and Anti-Spam.* CEAS, 2006.

A Reproducibility

The Sophos dataset and pre-trained models can be found here: https://github.com/sophos-ai/SOREL-20M

The EMBER2018 dataset and a description of the EM-BER2018 featurization method can be found here: https: //github.com/elastic/ember

B Additional Results and Figures

Results comparing our new uncertainty augmented local adjustments Equation 2 (g+l) and Equation 3 (g+lv2) and Equation 4 (g+lv3) against the naive approach (g) are provided in Table 3 for the EMBER2018 dataset.

The only case where (g) performed best is when using the Bayesian Logistic Regression (BLR) model on the EM-BER2018 corpus at a target FPR of 0.01%. In this one case we have pushed the model beyond what it is capable of achieving, and all three methods perform poorly - by happenstance the global threshold's degenerate solution of claiming that there is no malware receives a better score due to our uncertainty approaches failing to meet the FPR goal, which has a high penalty. However, we would argue our uncertainty based approaches are still preferable in this scenario because the degenerate model (g) is equivalent to having no anti-virus installed.

Figure 5a, Figure 5b, and Figure 5c on the EMBER2018 dataset and Figure 6a and Figure 6b on the Sophos data show that the uncertainty distributions for test samples that the models ultimately got wrong place most of their weight on higher uncertainties. Consistently, the uncertainty distribution for test samples that a model ultimately got right places most of its weight on lower uncertainties. This suggests that overall system performance can be improved by leveraging uncertainty and flagging high uncertainty predictions for further processing and review. This explains the success of our approach, which can learn to use the uncertainty terms as a kind of additional offset. The more we want to lower the FPR rate, the less we should trust the model's outputs if uncertainty is high.



Figure 5: A comparison of uncertainty distributions for all three EM-BER2018 models at test time between samples predicted correctly and incorrectly.

Table 3: EMBER2018 comparison of the standard global adjustment (labeled as (g)) to the uncertainty aware local adjustments from Equation 2 (labeled as (g+l)) and Equation 3 (labeled as (g+lv2)) and Equation 4 (labeled as (g+lv3)). Best combined score (TPR penalized for over-runs on target FPR) shown in **bold**.

	Target FPR=1%			Target FPR=0.1%			Target FPR=0.01%		
Method	TPR	FPR	Comb.	TPR	FPR	Comb.	TPR	FPR	Comb.
BMC (g)	7.602E-01	1.177E-02	5.832E-01	4.998E-01	8.100E-04	4.998E-01	2.422E-01	8.000E-05	2.422E-01
BMC (g+l)	7.617E-01	1.217E-02	5.447E-01	4.998E-01	8.300E-04	4.998E-01	2.431E-01	9.000E-05	2.431E-01
BMC (g+lv2)	7.594E-01	1.166E-02	5.934E-01	5.016E-01	8.200E-04	5.016E-01	2.434E-01	9.000E-05	2.434E-01
BMC (g+lv3)	7.605E-01	1.173E-02	5.875E-01	5.023E-01	8.200E-04	5.023E-01	2.434E-01	9.000E-05	2.434E-01
BLR (g)	7.778E-01	9.550E-03	7.778E-01	0.000E+00	0.000E+00	0.000E+00	0.000E+00	0.000E+00	0.000E+00
BLR (g+l)	7.781E-01	9.560E-03	7.781E-01	5.977E-02	7.500E-04	5.977E-02	1.000E-04	8.600E-04	-7.600E+00
BLR (g+lv2)	7.781E-01	9.560E-03	7.781E-01	5.248E-02	7.800E-04	5.248E-02	9.000E-05	7.900E-04	-6.900E+00
BLR (g+lv3)	7.779E-01	9.290E-03	7.779E-01	5.261E-02	8.200E-04	5.261E-02	5.840E-03	1.200E-04	-1.942E-01
LGBM (g)	8.805E-01	1.805E-02	7.547E-02	6.954E-01	1.550E-03	1.454E-01	4.888E-01	8.000E-05	4.888E-01
LGBM (g+l)	8.680E-01	1.494E-02	3.740E-01	6.892E-01	1.390E-03	2.992E-01	5.142E-01	9.000E-05	5.142E-01
LGBM (g+lv2)	8.693E-01	1.488E-02	3.813E-01	6.917E-01	1.430E-03	2.617E-01	5.142E-01	9.000E-05	5.142E-01
LGBM (g+lv3)	8.727E-01	1.512E-02	3.607E-01	6.890E-01	1.380E-03	3.090E-01	5.136E-01	9.000E-05	5.136E-01



(b) Sophos, LGBM Ensemble

Figure 6: A comparison of uncertainty distributions for an ensemble of Sophos models at test time on Sophos data between samples predicted correctly and incorrectly.