

# Adversarial Campaign Mitigation via ROC-Centric Prognostics

Javier Echauz<sup>1</sup>, Keith Kenemer<sup>2</sup>, Sarfaraz Hussein<sup>3</sup>, Jay Dhaliwal<sup>4</sup>, Saurabh Shintre<sup>5</sup>,  
Slawomir Grzonkowski<sup>6</sup>, and Andrew Gardner<sup>7</sup>

<sup>1-3,7</sup> Symantec Corporation, Atlanta, GA, 30328, USA  
{Javier\_Echauz, Keith\_Kenemer, Sarfaraz\_Hussein, Andrew\_Gardner}@symantec.com

<sup>4,5</sup> Symantec Corporation, Mountain View, CA, 94043, USA  
{Jasjeet\_Dhaliwal, Saurab\_Shintre}@symantec.com

<sup>6</sup> Symantec Corporation, Dublin, Ireland  
Slawomir\_Grzonkowski@symantec.com

## ABSTRACT

Machine learning models are vulnerable to adversarial inputs that induce seemingly unjustifiable errors. As automated classifiers are increasingly used in industrial control systems and machinery, these adversarial errors could grow to be a serious problem. Despite numerous studies over the past few years, the field of adversarial ML is still considered alchemy, with no practical unbroken defenses demonstrated to date, leaving PHM practitioners with few meaningful ways of addressing the problem. We introduce turbidity detection as a practical superset of the adversarial input detection problem, coping with adversarial *campaigns* rather than statistically invisible one-offs. This perspective is coupled with ROC-theoretic design guidance that prescribes an inexpensive domain adaptation layer at the output of a deep learning model during an attack campaign. The result aims to approximate the Bayes optimal mitigation that ameliorates the detection models degraded health. A *proactively reactive* type of prognostics is achieved via Monte Carlo simulation of various adversarial campaign scenarios, by sampling from the models own turbidity distribution to quickly deploy the correct mitigation during a real-world campaign.

## 1. INTRODUCTION

A machine learning application often begins with a dataset of examples and the task is to find a classification model that will turn inputs into class-label predictions, while preserving some sense of minimum expected error. The learning problem is often *unrealizable*, so no perfect model exists that will have 0 generalization error (Shalev-Shwartz &

Ben-David, 2014). But less obviously, it is often possible to deterministically find input examples that *force* the model to misclassify (Szegedy et al., 2014). Machine learning (ML) models can be subjected to adversarially crafted small perturbations that purposely induce these errors, and they can seem unjustified or surprising to a human observer (e.g., a digital image of a school bus mistaken for a bird). As automated ML-based classifiers pervade across applications in transportation, medicine, finance, and cybersecurity, adversarial errors could grow to be a very serious problem. The danger is particularly acute in industrial control systems (ICS), industrial Internet of Things (IIoT), automation equipment, and factory robotics, where malfunctions can be life-threatening (e.g., steel mill furnace explosions, power grid crashes, etc.). Unfortunately, ICS attacks are on the rise, with increased vectors for malicious party access to critical infrastructure (National Cybersecurity and Communications Integration Center, 2017). Detection of attacks to cyberphysical systems (Yan et al., 2018), particularly as it relates to adversarial ML, is a growing area of concern that has been underserved in PHM literature.

Despite vigorous study over the past few years (see review in (Gilmer, Adams, Goodfellow, Andersen, & Dahl, 2018)), the field of adversarial ML is considered by researchers to be at a nascent stage (Evans, 2018), with no practical unbroken defenses demonstrated to date (attacks succeed with  $p > .25$ ) (Carlini & Wagner, 2017a), and still talks of an “arms race” between attackers and defenders (Goodfellow, McDaniel, & Papernot, 2018). This leaves PHM practitioners with few meaningful ways of addressing the problem. We have identified a fundamental flaw in the current interpretation of adversarial defenses, and offer an alternative practical reformulation of the problem that copes with population-level

---

Javier Echauz et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

*campaign mitigation* as opposed to individual input, case-by-case protection.

The defense side of adversarial ML has tried to answer a blend of two questions: (a) How to *robustify* a model (make it harder for an attacker to fool)? This has led to adversarial training, defensive distillation, feature squeezing, architecture modification, and minimax optimization (Madry, Makelov, Schmidt, Tsipras, & Vladu, 2017); and (b) What can be *measured* about adversarial inputs that is different from regular ones? This has led to input validators and adversarial detectors (Goodfellow et al., 2018). Generative adversarial networks (GANs) can synthesize adversarial examples which can then be used to retrain the classifier, however, this only helps insofar as it gets a classifier closer to Bayes optimality (it can also make things worse). By omissions in the current discourse, these methods have created the illusion that we could one day prebake a solution at training time that will protect a model against *one-or-few-off* adversarial inputs at deployment time. Our work suggests that the latter goal comes at a disproportionate price in expected error. Intuitively, if there was a way to accurately detect error-inducing inputs at runtime, then that same detector would have been used to augment or improve the training to begin with.

In the following sections, we will introduce *turbidity detection* as a different, ROC-centric way of thinking about adversarial example detection that fixes current widespread misinterpretations and leads to a practical mitigation. Our theory yields 3 previously unreported results: (i) unqualified use of an adversarial detector inverts ROC (harms); (ii) adversarial campaign pinches down ROC (harms); and (iii) conditions exist where the ROC can be repaired to at least a gracefully degraded state during the campaign. We propose a methodology for putting that into practice and show experimental results using image (digit recognition) and IIoT security (malware detection) data.

## 2. TURBIDITY DETECTION THEORY

Our first aim is to show that the unqualified use of an adversarial detector has deleterious effect on ROC. To that end, we will start in a seemingly restrictive setting: 1-dimensional input, uniform distribution over  $-10$  to  $10$ , binary output from binomial discrete-choice theory with logistic noise, equiprobable classes, and Bayes decision rule. However, our main results (ROC inversion, pinching, and repair) will not critically depend on these specific choices, retaining clarity of illustration without loss of generality.

Instead of asking where adversarial examples are “hiding” in high-dimensional input space, we focus on scalar decision score output axis (preactivation/logit or post-activation/pseudo-probability), where model-processed samples have to end up anyway, and where any decision confusion actually occurs. Fig. 1 shows a deep neural network taking input array

through convolutional and nonlinear activation layers, then dense layers reducing the output to a scalar decision score  $s$  (here logit). Consider a data-generating process (DGP) such that ground-truth bipolar labels  $y \in \{-1, +1\}$  come from adding the score to a *Logistic*(0, 1) symmetric noise  $\xi$  (whose scale and bias control class separability and class imbalance respectively), and taking sign:

$$y = \text{sign}(s + \xi) .$$

The symmetry of the noise about its 0 mean implies our DGP emits equiprobable labels:  $P(0) = P(1) = 0.5$ . In order to output a monotonic and correctly calibrated posterior probability  $P(1|s)$ , what the last-layer activation of the deep net “wants to be” is the CDF of the discrete-choice noise:

$$\begin{aligned} P(y = 1|s) &= \int_0^{\infty} p_{\xi}(s' - s) ds' \\ &= \int_{-s}^{\infty} p_{\xi}(s') ds' = 1 - \int_{-\infty}^{-s} p_{\xi}(s') ds' \\ &= 1 - CDF_{\xi}(-s) = CDF_{\xi}(s) . \end{aligned}$$

The logistic-distributed noise has logistic sigmoid CDF, agreeing with an output neuron (here with  $\mu = 0, c = 1$ ):

$$\xi \sim \text{Logistic}(\mu, c) \Rightarrow CDF_{\xi}(s) = \frac{1}{1 + e^{-\left(\frac{s-\mu}{c}\right)}} = \sigma\left(\frac{s-\mu}{c}\right) .$$

The Bayes-optimal decision rule (one yielding least probability of misclassification in our DGP) corresponds to the homogeneous halfspace (here semiaxis) obtained by thresholding the above posterior probability at 0.5, or directly thresholding the preactivation score:

$$\hat{y} = \text{sign}(s) . \quad (4)$$

Now we derive the ROC for this ideal detector in its regular environment. From Bayes theorem, the 0 (“clean”) vs-1 (“mal”) class-conditionals of the score are

$$\begin{aligned} p(s|0) &= F_{\xi}(-s)\mathcal{U}(s; -10, 10)/0.5 , \\ p(s|1) &= F_{\xi}(s)\mathcal{U}(s; -10, 10)/0.5 , \end{aligned}$$

(see Fig. 2(a) where  $F(\cdot)$  denotes CDF from now on,  $\mathcal{U}(s; -10, 10) = [10 \leq s \leq 10]/20$  is the uniform PDF, and that last Iverson bracket  $[\cdot]$  means indicator function: valued 1 when the event  $s$ -within-the-interval is true and 0 otherwise. In our 1D mathematical reference figures, the score is directly equal to the input:  $s = x$  (while in higher dimensions, it will be an inner product where coordinates can be explanatory variables, features, previous neural layers, etc.).




Figure 1. Deep neural network from input to scalar decision score  $s$  with discrete-choice true label  $y$ .

The marginal of scores is the uniform PDF (Fig. 2(b)), while the malicious class-posterior is (Fig. 2(c)):

$$P(1|s) = \frac{p(s|1) \cdot \frac{1}{2}}{p(s|0) \cdot \frac{1}{2} + p(s|1) \cdot \frac{1}{2}} = F_{\xi}(s). \quad (6)$$

Finally, monotonicity of the class-1 posterior allows us to obtain the ROC from a single sweep on the  $s$ -axis, yielding the parametric curve (Fig. 2(d)):

$$(FPr, TPr) = (1 - F_0(s), 1 - F_1(s)), \quad (7)$$

where

$$F_0(s) = \int_{-\infty}^s p(s'|0) ds', \quad F_1(s) = \int_{-\infty}^s p(s'|1) ds'.$$

## 2.1. Clarity and Turbidity Distributions

Unless classes are 100% separable in a generalization-preserving way relative to the DGP (input features, label noise, and their statistical relation), every model, including the Bayes-optimal one, experiences difficulty whenever it makes the wrong class prediction. We say that samples

that confuse the model, i.e., FPs and FNs, are **turbid** from the model's point of view, whereas all the other correctly-classified TNs and TPs are **clear**. We can think of every model that tackles the original 0-vs-1 problem as having an inherent dual problem: separating clear-vs-turbid (denoted e-vs-d as mnemonic for "easy"-vs-"difficult"), for which a different detector can be built. Since the model's confusion depends on its output threshold, by default we peg the associated turbidity detection concept to the maximum balanced-accuracy/Youden index threshold in the original detector (ROC operating point closest to upper-left corner).

Next we present the clarity and turbidity distributions for a DGP where there is 50-50% proportion of clear vs turbid samples (something that we will characterize as a *toxic* environment compared to the regular one where mistakes should be rare), and 50-50% proportion of clean vs mal within each. Obtain each conditional as a mixture of the truncated class-0 plus the truncated class-1 PDFs. For example, the left half of turbidity  $p(s|d)$  consists of the left tail of  $p(s|1)$  (= FNs) normalized by the area under it up to 0 (=  $F_1(0)$ ), while the right half has the right tail of  $p(s|0)$  (= FPs) normalized by the area under it from 0 onward (=  $1 - F_0(0) = F_1(0) = .0693$ ). The mixture of these 2 densities then gives the Laplace inflexed arch shape (purple in Fig. 2(e)) and is in fact the same as a truncated Laplace distribution:

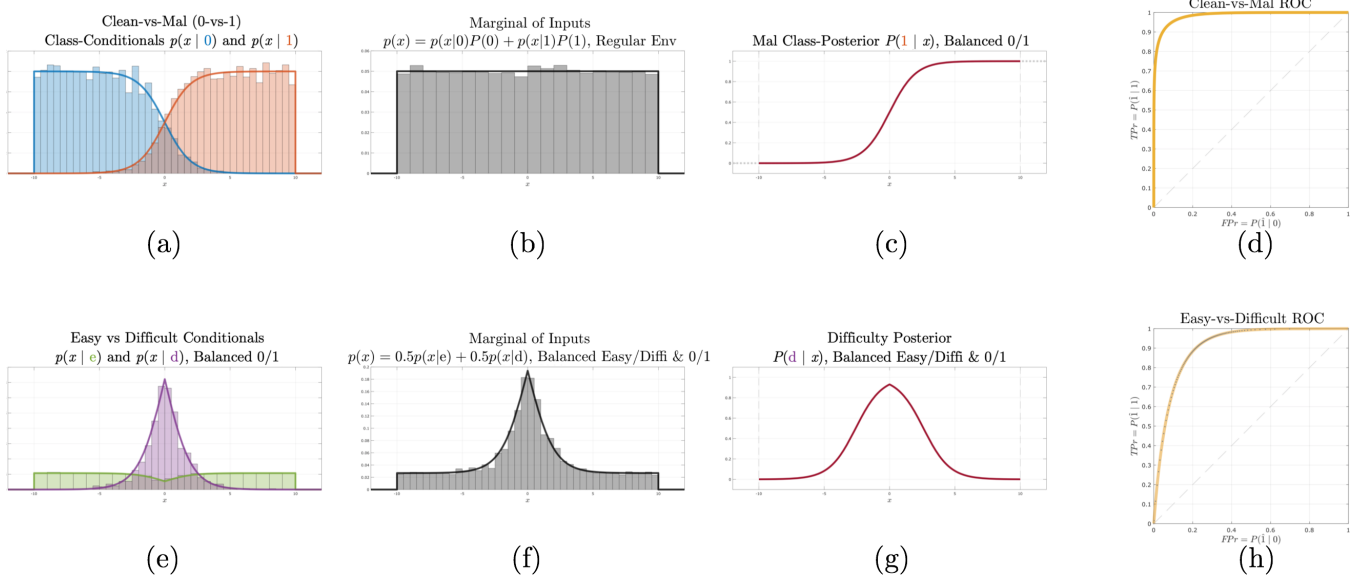


Figure 2. Top: Class-conditionals, marginal, posterior, and ROC of 0-vs-1 detector in regular environment. Bottom: Class-conditionals, marginal, posterior, and ROC of corresponding e-vs-d detector.

$$\begin{aligned}
 p(s|e) &= \frac{\frac{1}{2}p(s|0)[s < 0] + \frac{1}{2}p(s|1)[s \geq 0]}{F_0(0)}, \\
 p(s|d) &= \frac{\frac{1}{2}p(s|0)[s \geq 0] + \frac{1}{2}p(s|1)[s < 0]}{F_1(0)}.
 \end{aligned} \tag{9}$$

Obtain the marginal of scores from the equiprobable mixture of the clear and turbid distributions (or from total probability theorem; Fig. 2(f)), and the turbidity class-posterior as (Fig. 2(g)):

$$P(d|s) = \frac{p(s|d) \cdot \frac{1}{2}}{p(s|e) \cdot \frac{1}{2} + p(s|d) \cdot \frac{1}{2}}. \tag{10}$$

Note that this symmetric reverse-ogee arch is nonmonotonic. This implies that the theoretical ROC curve can no longer be obtained simply by sweeping a single threshold over the  $s$  domain; doing so would result in a suboptimal improper curve (under diagonal chance line). The most general method is to sweep a descending threshold on the vertical axis of the class-posterior, nonlinearly solve/root-find all critical  $s$  values where posterior intersects the threshold, then calculate area under class-conditionals over  $s$  regions so as to obtain the pair  $(FPr, TPr)$ . In effect, the ROC curve computation becomes multibranching, with number of connected segments dependent on number of intersections encountered during the sweep. Fig. 2(h) shows the exact ROC, using either the multibranching algorithm just described or an alternative monotonic version afforded by symmetry in this case. Luckily, when data scientists compute an empirical ROC (i.e., from a data sample), they automatically obtain a Monte Carlo es-

timate, so theoretical complications like the nonmonotonicity above are never encountered. However, the scores should be presented as the possibly nonmonotonic posteriors instead of as preactivations.

## 2.2. Relation to Adversarial Detection

The widely accepted *oracle* definition of adversarial examples (Evans, 2018) states that: (i) they are created with intent to deceive, (ii) they start from a seed example of say class A, correctly seen as class A by the model, and (iii) after perturbation they still behave like class A according to the oracle/ground-truth, yet they are now incorrectly seen as class B by the model. However, the goal of adversarial example detection (accurately determining at runtime whether an input is adversarial) has been widely misconstrued, leading to overfitting and/or invalidly-dichotomized detectors. If we insist we can detect a particular set of adversarial samples, then that same detector is bound to fail on a freshly created one operating in a regular environment. It will work if operated in a toxic environment, but then for a whole different reason as we'll see below.

A typical adversarial detection experiment starts from a dataset of regular samples, takes each instance in the dataset as a seed to which a transformation (e.g., from CleverHans library (Papernot et al., 2016)) is applied in order to create an adversarial counterpart, and then sees if the “regular-vs-adversarial” examples are discernible in some way (e.g., by showing differences in distributions or by building adversarial detectors and measuring their above-chance discrimination). By definition, *all adversarial examples are turbid*. Further, they can exist with “high model confidence”

(with  $P(0|x)$  or  $P(1|x)$  near 1). But exactly the same is true of natural, unforced errors. All regular FPs and FNs are turbid, and while most are associated with low confidence ( $P(0|x) = P(1|x) = 0.5$ ) near model’s decision boundary, high-confidence ones also arise. They happen as predicted even by the 1D DGP, just less frequently, consistent with the tapered-but-still-nonzero tails of the turbidity distribution. Thus, regular and adversarial samples can share the same domain.

We don’t believe human intent is a distinguishing feature that can be measured either—a view hinted in (Carlini & Wagner, 2017a)—anymore than telling if the person who made the samples was left-handed from looking at the numerical input coordinates. So what is it in the standard adversarial detection experiment that is being detected? **Answer:** the observed difference between regular and adversarial conditions stems from the fact that *all* adversarial samples are turbid/difficult by definition, whereas in the regular environment turbid samples are rare. Turbid tends to concentrate while regular  $\approx$  clear spreads, thus second moments separate. A scientific animation clarifying this point can be seen in (Echaz, 2019).

These issues can be fixed by moving from a fortuitous “regular-vs-adversarial” dichotomy to the principled “clear-vs-turbid,” and by not blurring the line between detector and its intended deployment environment (Gilmer et al., 2018). Dropping intent and seed-of-origin out of the adversarial character makes the problem realistic and applicable to campaign mitigation.

### 2.3. ROC Inversion

We now show that a realistic adversarial (i.e., e-vs-d) detector cannot actionably help 0-vs-1 decision-making in a regular environment as it leads to **ROC inversion**. In the spirit of *reductio ad absurdum*, let the theoretical 1D adversarial detector  $A(x)$  in Figs. 2(g,h), Eq. (10) augment the probabilistic 0-vs-1 detector  $D(x)$  in Figs. 2(c,d), Eq. (6). Given any input  $x'$  at test time, if  $A(x')$  is accurately declaring that  $x'$  is adversarial then we would want to contradict the decision from  $D(x')$ . From the point of view of  $A(x)$ , the Bayes trigger to declare adversariness is  $A(x) > 0.5$ , which is equivalent to checking if input magnitude is within a critical cutoff:  $|x| < 2.597$  (also equal to the crossover points in Fig. 2(e)). The augmented detector becomes

$$\hat{y} = \text{sign}(x)\text{sign}(|x| - 2.597),$$

and the augmented-system posterior probability is:

$$P_{\text{aug}}(1|x) = F_{\xi}(x)[|x| > 2.597] + F_{\xi}(-x)[|x| \leq 2.597].$$

(Any function that reverses the  $D(x)$  decision within that in-

terval will work.) Fig. 3(a) shows this nonmonotonic posterior. Fig. 3(b) shows the exact ROC using multibranching algorithm. The original accuracy of  $F_0(0)P(0) + (1 - F_1(0))P(1) = 0.93$  goes down to 0.79. The original detector was already optimal for its intended regular environment, and overriding its decisions only makes it worse. Thus, *protection against one-off adversarial examples is a misguided design goal*.

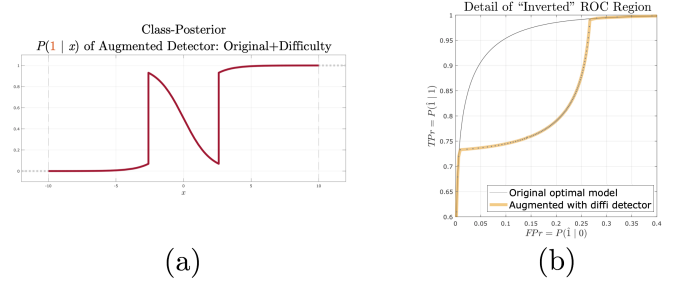


Figure 3. Posterior and ROC of the augmented original + turbidity detector leading to ROC inversion.

### 2.4. ROC Pinch-Down

We now show that operating the original 0-vs-1 detector in a toxic environment leads to an **ROC pinch-down**. Under adversarial campaign, class-conditionals can display abnormal concentrations around the original decision score threshold, a single crossover (as in Fig. 2(a)) can become multiple, class-posterior can turn nonmonotonic, and errors become frequent, making model performance plummet. Continuing with the balanced proportions of 50-50% e-vs-d and 50%-50% 0-vs-1 of Section 2.1, the class-conditional likelihoods that the original detector now has to confront are (Fig. 4(a)):

$$\begin{aligned} p_{\text{toxic}}(x|0) &= p_e(x)[x \geq 0] + p_d(x)[x < 0], \\ p_{\text{toxic}}(x|1) &= p_e(x)[x < 0] + p_d(x)[x \geq 0], \end{aligned}$$

where  $p_e$  and  $p_d$  are the conditionals in Eq. 9.

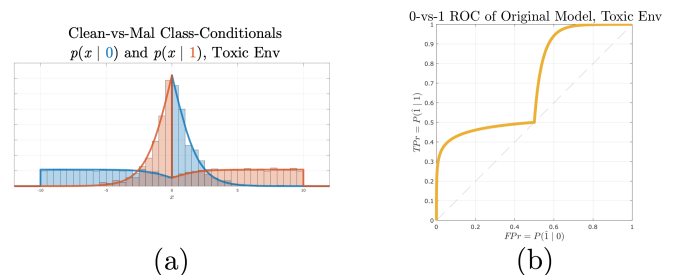


Figure 4. Class-conditionals in toxic environment leading to ROC pinch-down.

The marginal of inputs is identical to Fig. 2(f), just composed differently from the average of the above toxic conditionals.

The posterior (not shown) is the same one in Fig. 2(c) since the model remained naively unchanged during this toxic campaign. The exact ROC (Fig. 4(b)) can be obtained here from the monotonic sweep form in Eq. 7, except with  $F_{0,\text{toxic}}$  and  $F_{1,\text{toxic}}$ . So while augmenting the detector during normal operation was harmful, ignoring the problem during abnormal operation is also potentially worse. Thus, *protection against adversarial campaigns (not one-offs) is needed*.

#### 2.4.1. Asymmetric Toxic Environments

At the 1:1 ratio of e-vs-d samples under both 0 and 1 classes, the characteristic ROC seagull (Fig. 4(b)) has curve pinned at the chance line (50% accuracy). However, at other ratios of turbidity proportions under each class, conditionals become asymmetric and the pinch point moves somewhere else. Class-0 e:d ratio controls the horizontal axis (FPr), while class-1 e:d ratio independently controls the vertical axis ( $TPR = 1 - FNr$ ). This means that if adversarial campaign actors could not only add samples but also subtract from the environment seen by the model, they would be able to place the pinch-down point anywhere on the ROC plane! But they would have to be oracles themselves, for example, to force the model to be always wrong in the future would pin operating point at the bottom-right corner (something we cant do ourselves with imperfect knowledge). Fig. 5 shows a toxic formulation where class-0 samples are regular (i.e., no adversarial FPs), with their natural  $F_0(0) : F_1(0)$  proportion of clear to turbid, whereas class-1 samples have an unnatural 37.5%-62.5% proportion.

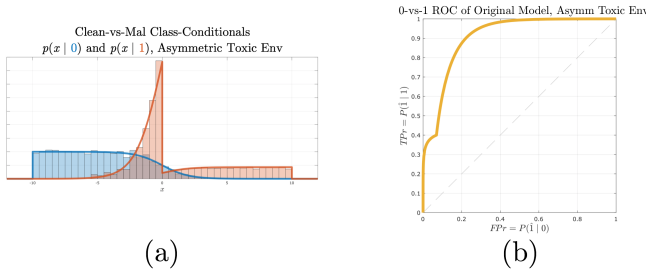


Figure 5. Class-conditionals in asymmetric toxic environment, moving ROC pinch-down somewhere else.

#### 2.5. Mitigation/Repair of the Degraded ROC

In its intended regular environment, the original model can adapt to changes in maliciousness imbalance (0-vs-1 prevalences) by simply sliding its operating point along the intact, class prevalence-agnostic ROC curve. However, in the adversarially toxic environment it is no longer enough to simply adjust a threshold to match the environment; a fundamentally different detection problem must be solved. In order to “unpinch” the ROC to the best available shape given the adversarial campaign, we should obey the new posterior:

$$P_{\text{toxic}}(1|x) = \frac{p_{\text{toxic}}(x|1)P(1)}{p_{\text{toxic}}(x|0)P(0) + p_{\text{toxic}}(x|1)P(1)}. \quad (15)$$

This will typically be nonmonotonic (Fig. 6(a)). The exactly repaired ROC is obtained from the multibranch algorithm as shown in Fig. 6(b).

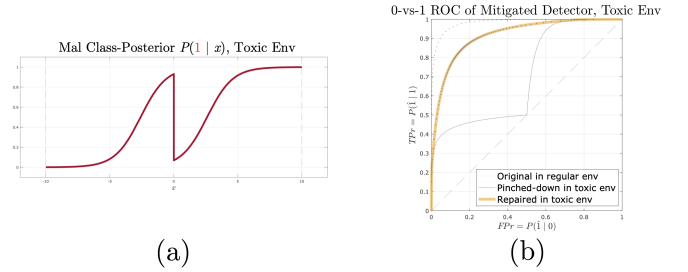


Figure 6. New posterior and mitigated/repaired ROC.

The new optimal maximum-a-posteriori Bayes classifier implements *decision reversals* relative to the original one. Reversals occur only within the decision score intervals where the new heights of 0-vs-1 conditional likelihoods have swapped their dominance, due to the new concentration of turbid/difficult samples in the environment. Thus, *the mitigated detector is gracefully (rather than catastrophically) degraded, restoring acceptable error rates and adaptability to maliciousness imbalance*.

#### 2.5.1. When Repair Isn’t Possible

In some cases it isn’t really possible to “unpinch” the ROC because the curve morphs into a seamless one with no dent (as if in Fig. 5(b) the pinched point fused into the left branch), e.g., with ratio of e-vs-d samples still at 1:1 but with malicious class prevalences falling outside of the interval  $F_1(0) < P(1) < F_0(0)$ . The curve is still depressed compared to the original regular one due to 50% of samples being turbid, and only detection threshold remains as a potential adjustment.

We have also uncovered **adversarial covariate shift** as another condition where ROC repair isn’t possible. This would make score class-conditionals and marginal more turbid while keeping the posterior intact. For example

$$\begin{aligned} p_{\text{toxic}}(x|0) &= (1 - F_{\xi}(s)) \cdot (0.5p(x|e) + 0.5p(x|d)) / 0.5, \\ p_{\text{toxic}}(x|1) &= F_{\xi}(s) \cdot (0.5p(x|e) + 0.5p(x|d)) / 0.5. \end{aligned}$$

The reader can verify that the posterior  $P(1|x)$  is exactly recovered as  $F_{\xi}(x)$ , the CDF of the label noise (also true for other imbalanced 0-vs-1 priors). However, it seems unrealistic that an adversary could shape the conditionals in this fashion as it would require omnipotent control of the environment

beyond merely adding adversarial samples to the regular one seen by the model.

## 2.6. Generalization to Suboptimal Models and Higher Dimensions

We have systematically charted an atlas documenting how the above 1D reference theory is impacted when the model is sub-optimal instead of Bayes-optimal (via mistuned bias and/or misaligned weights), and in higher-dimensional input space, where the decision score  $s$  is taken to be the preactivation  $\mathbf{x}^\top \mathbf{w}$ , i.e., the (possibly augmented) dot product at the output layer of a probabilistic binary classifier. Due to space restrictions, we only mention that all ROC inversion, pinch-down, and mitigation repair results remain qualitatively identical. The score is still 1D; the only difference is that with independent components of  $\mathbf{x}$ , all distributions become windowed/tapered, and ROCs get “dumber”/shallower from the CLT centrality effect of marginal  $p(\mathbf{x})$ , which makes samples appear close to decision boundary more frequently. Injecting correlation structure in components of  $\mathbf{x}$  also weakens separability, but the main ROC results hold. Further, nothing above prevents decision scores from being computed by *non*-neural network models. Thus, the described campaign effects and mitigation apply to any decision-making component that exposes its scores to attackers, including ensembles of decision trees widely prevalent in industrial settings.

## 3. PREEMPTIVE DOMAIN ADAPTATION

The theoretical results in the previous section can be put into PHM practice by monitoring estimates of the decision score class-conditional distributions in order to declare if and when an adversarial campaign is in effect, repair the degraded ROC during campaign, restore the original model after campaign subsides, and improve readiness for future attacks via simulation. Assume a well-trained classifier has been deployed in its originally-intended threat environment where errors are rare (e.g.,  $> 95\%$  hit at  $< 0.1\%$  FP rates). A health management methodology can track 0-vs-1 conditional score histograms (and optionally error rates), from which class-conditionals curves are kernel-density estimated (KDE) as smooth functions. This still requires ground-truth label estimates; in cybersecurity they are obtained after some lag ranging from sub-seconds (with access to cloud-based reputation, etc.) to days (offline endpoints with sporadic live updates, air-gapped IIoT devices, etc.). There is also a way to detect without ground truth, by introspectively looking at whether too many decision scores are falling in a low-confidence interval, but this is bypassed if attackers actively inject only high-confidence samples.

Under adversarial campaign, class-conditionals may develop multimodality, with multiple crossover points that misfit the original decision rules, making model performance plummet

for some period of time. This condition can be declared from observation of empirical ROC pinch-down or abnormally large error rates. (In the low-confidence campaign case, the model can track its own decision scores falling in an interval near decision threshold at higher-than-historical rates, suggesting adversarial manipulation since it would be rare to see that in the regular environment.) At that time, an equiprobable class-posterior function implementing Eq. (15) is transmitted to the endpoint model to be used as a post-transformation layer of the original decision scores. The output of this function can then be thresholded to obtain a desired mitigated ROC operating point. In effect, this is an inexpensive statistical domain adaptation that reverses decisions where it makes sense to do so. The same logic can be applied in reverse to restore the original model when campaign has subsided.

The methodology above is still 100% reactive defense. Our investigation suggests that a 100% proactive defense (where the model is hardened at training time against all future one-off adversarial samples in the regular environment) is mathematically impossible. Thus, prognostics in the usual sense of predicting remaining life until failure, to do something about it before it occurs, is outside the scope of our work. However, we introduce a *proactively reactive* compromise. It precomputes the optimal response to each of several plausible adversarial attack scenarios, via Monte Carlo simulation drawing from the model’s own turbidity distribution, and stores that information as a look-up table to quickly deploy the correct mitigation during a real-world campaign. A final health maintenance modality is to put the decision modification layer into effect continuously/ prophylactically without waiting to detect that a campaign has begun (in which case the last layer calculation automatically yields simply an identity function). This way, as machine operating conditions change even gradually, the method is already there to mitigate possibly harmful effects while any persistent shifting is investigated.

## 4. EXPERIMENTS WITH REAL-WORLD DATA

This section verifies the main ROC inversion, pinching, and repair results using real-world data with corresponding attacks against a deep neural network classifier in two application areas: digit recognition and IIoT malware detection.

### 4.1. Digit Recognition

The standard MNIST benchmark dataset was used, containing 60,000 grayscale  $28 \times 28$ px images of handwritten digits. The deep convolutional neural network trained in (Dhaliwal & Shintre, 2018), whose first layers are visualized in Fig. 1, achieved over 99% accuracy on a holdout split of the data. A stratified random sample of 2400 images was taken to equally represent all digits. We adversarially generated 2400 FPs and 2400 FNs using the Carlini-Wagner algorithm (Carlini

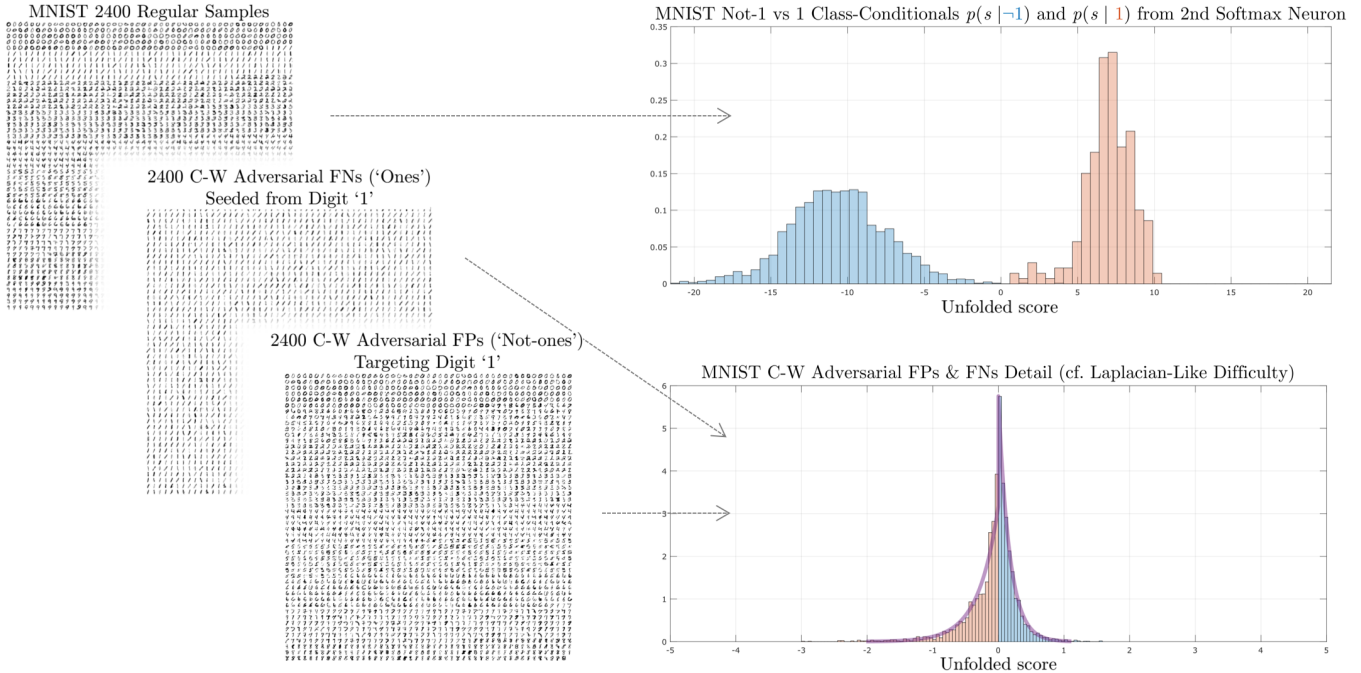


Figure 7. Regular not-1 vs 1 conditional histograms and balanced turbidity histogram from adversarial FPs & FNs.

& Wagner, 2017b). The 10-class problem was dichotomized into classes ‘not-1’ vs ‘1’ by unfolding the preactivation decision scores as

$$s = \frac{1}{2} (s_1 - \max\{s_0, s_2, s_3, s_4, s_5, s_6, s_7, s_8, s_9\}),$$

where  $s_1$  is the preactivation score at the neuron for class ‘1’ (2nd indicator in softmax layer). That leaves 2160 regular instances of class ‘not-1’ and 240 of ‘1’—a class-prior imbalance of 9:1. Fig. 7 shows the pdf-normalized regular conditional histograms (top) and the turbidity distribution from the adversarial FNs and FPs (bottom; additionally color-coded by not-1 vs 1 classes). As predicted by the theory, the latter distribution has a Laplace-like inflex concentration around the score decision-crossing point (cf. purple in Fig. 2(e)). For clarity, it is shown with balanced not-1s vs 1s within the turbid condition; the regular environment would have 9 times more FPs than FNs while toxic ones can be manipulated. In this potentially overfit “99%” accuracy case, we cannot display an empirical turbidity distribution with only the natural FPs and FNs because there were only 2 and 0 cases, respectively.

Aided by 2160 of the adversarially discovered FPs and 240 of the FNs to mimic the natural 9:1 class lopsidedness in the regular environment, we estimated decision-reversal interval as  $[-1.6, 1.6]$  (graphically from intersection of empirical conditional histograms, much like green vs purple curves in Fig. 2(e)), and generated the empirical ROC over the 2400 regular




Figure 8. KDE-smoothed class conditionals and corresponding equiprobable posterior for use in repair.

samples. Fig. 9(b) inset confirms that the resulting ROC is inverted. It appears to be small harm but there is actually almost an order-of-magnitude larger FP rate; this difference is critical in the field.

The two class-conditional PDFs were estimated using Gaussian kernel and bandwidth  $B$ . This is equivalent to fitting a Gaussian mixture distribution where means equal all individual data points, and covariances equal the shared constant  $B^2$ . During a very toxic adversarial campaign where half of all samples become turbid, KDE-smoothed PDFs from aggregated data at the desired e-vs-d proportions reveal 3 crossovers (Fig. 8(a)). The corresponding equiprobable posterior crosses the 0.5 threshold 3 times (Fig. 8(b)). Operating the original unmitigated detector yields the pinched-down ROC (Fig. 9(c)). In contrast, passing the scores thru the mitigated posterior function yields the repaired ROC (Fig. 9(d)). Fig. 9 confirms the inversion, pinch-down, and repair as predicted by the theory. We note however that the empirical nature of the construction makes ROCs look like staircases and it’s impossible to discern convexity of the inversion from con-



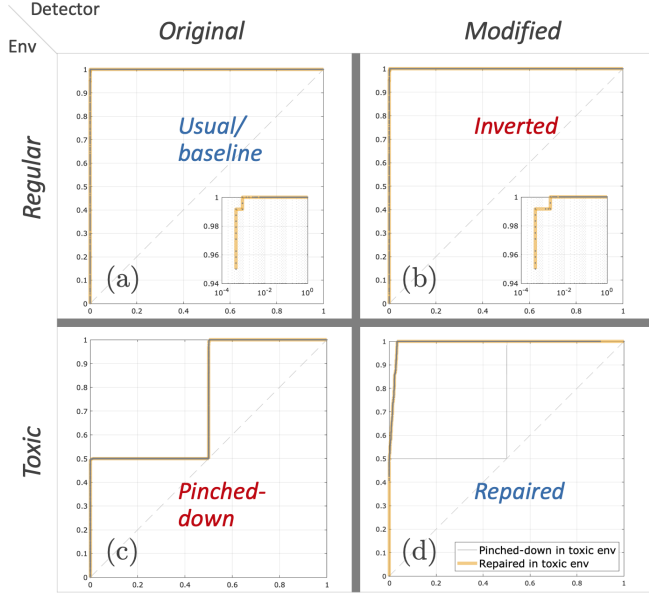


Figure 9. Regular, inverted, pinched-down & repaired ROCs in adversarial campaign against MNIST detector.

cavity of the pinch-down (something we know only from the theory).

#### 4.2. Malware Detection from Raw Bytes

We additionally verified the adversarial campaign health management framework using a pre-production model intended for an IIoT “ICSP Neural” USB scanning device. It contains a deep network with {embedding, 4 convolutional, 3 dense, softmax} layers trained on half a million raw executable files (aggregated from a mix of clean and malicious customer submissions and vendor feeds). This type of network is fed integers in  $[0,255]$  representing bytes of a file zero-padded or cropped to length 700,000 (as if it were a wide image that is only 1 pixel tall). The test dataset consisted of 2000 clean and 2000 malicious files sampled from a time split spanning 1 month after the training date. The model was made purposely suboptimal (with regular ROC curve far from upper-left corner) in order to better observe the manifestations of the theory.

One of the simplest adversarial attacks for binaries is to append a crafted payload at the end of the file (Kuppa, Grzonkowski, & LeKhac, 2018). The brute-force algorithm appends random chunks until the model flips its decision (to within some trial count tolerance). A high-confidence campaign was defined as a set of new binaries bypassing the model with pseudo-probability output above 0.97. Drawing seeds from the size-4000 test set, 512 adversarial FPs and 524 adversarial FNs were created this way. The empirical class-conditional histograms in the toxic environment are characterized (when pegged to the regular minimum-balanced-error




Figure 10. Regular, pinched-down & repaired ROCs in special adversarial campaign against malware detector.

score threshold 2.2) by 357 unforced + 512 forced FPs, plus 498 unforced + 524 forced FNs, totaling 1891 errors and thus a 62.5%-37.5% clear-to-turbid ratio. The clean (class 0) conditional is estimated from

$$p_{\text{toxic}}(s|0) \approx \frac{1}{n_0} \sum_{i=1}^{n_0} \frac{1}{B_0} K \left( \frac{s - s_i^{(0)}}{B_0} \right),$$

where  $K(\cdot)$  is the standard Gaussian kernel,  $s_i^{(0)}$  are the unfolded decision scores (2nd-unit softmax preactivation minus 1st) under class 0, and  $B_0 = 0.687$  is the bandwidth from Silverman’s estimate. The malware (class 1) conditional was similarly obtained with  $B_1 = 0.928$ . Fig. 10 shows the devastating effect of this campaign on ROC and how much could be mitigated. Passing the original models decision scores through the KDE-formed posterior brings the whole curve back to at least a gracefully degraded state.

#### 5. CONCLUSIONS

The common misunderstanding surrounding what to do about adversarial inputs that fool detectors can be cleared by fixing the “regular-vs-adversarial” dichotomy and by recognizing the difference between one-off/per-trial basis protection vs adversarial campaign mitigation. Our investigation suggests that universal pre-hardening defenses are impossible without paying a price in accuracy of the original model operating in its regular environment.

We introduced turbidity detection, campaign mitigation, and

preemptive domain adaptation as conceptual frameworks leading to practicable detector health management solutions. The theory yielded previously unreported results about ROC inversion, pinch-down, and repair in the context of adversarial threats to deep neural networks increasingly used in industry, and also applicable to widely prevalent ensembles of decision trees.

It should be understood that our method is not a panacea to shield or empower a model; what it does is optimally mitigate the damage (dramatically so for some ROC operating points) caused by adversarial toxicity that the original model wasn't designed to tackle on its own.

## REFERENCES

- Carlini, N., & Wagner, D. (2017a). Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security* (p. 3-14).
- Carlini, N., & Wagner, D. (2017b). Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy* (p. 39-57).
- Dhaliwal, J., & Shintre, S. (2018). Gradient similarity: An explainable approach to detect adversarial attacks against deep learning. *arXiv preprint arXiv:1806.10707*.
- Echaz, J. (2019). *10 ML Gotchas for AI/ML Security!* <https://medium.com/ai-ml-at-symantec/10-ml-gotchas-for-ai-ml-security-8e7e79ae615>.
- Evans, D. (2018). Keynote presentation: Is "adversarial examples" an Adversarial Example? In *1st Deep Learning and Security Workshop, San Francisco*.
- Gilmer, J., Adams, R. P., Goodfellow, I., Andersen, D., & Dahl, G. E. (2018). Motivating the rules of the game for adversarial example research. *arXiv preprint arXiv:1807.06732*.
- Goodfellow, I., McDaniel, P., & Papernot, N. (2018). Making machine learning robust against adversarial inputs. *Communications of the ACM*, 61(7), 56-66.
- Kuppa, A., Grzonkowski, S., & LeKhac, N. A. (2018). Enabling trust in deep learning models: A digital forensics case study. In *17th IEEE International Conference On Trust, Security and Privacy* (p. 1250-1255).
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2017). Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- National Cybersecurity and Communications Integration Center. (2017). *ICS Cert Monitor*. [https://ics-cert.us-cert.gov/sites/default/files/Monitors/ICS-CERT\\_Monitor\\_Nov-Dec2017\\_S508C.pdf](https://ics-cert.us-cert.gov/sites/default/files/Monitors/ICS-CERT_Monitor_Nov-Dec2017_S508C.pdf).
- Papernot, N., Faghri, F., Carlini, N., Goodfellow, I., Feinman, R., Kurakin, A., ... others (2016). Technical report on the cleverhans v2. 1.0 adversarial examples library. *arXiv preprint arXiv:1610.00768*.
- Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge University Press.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2014). Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199v4*.
- Yan, W., Mestha, L., John, J., Holzhauer, D., Abbaszadeh, M., & McKinley, M. (2018). Cyberattack Detection for Cyber Physical Systems Security—A Preliminary Study. In *Proceedings of the annual conference of the phm society* (Vol. 10).

## BIOGRAPHIES



**Javier Echaz** is a Sr Machine Learning Research Scientist and Technical Director at Symantec Corporation. He obtained a PhD at the interface of electrical, computer, and biomedical engineering from the Georgia Institute of Technology.



**Keith Kenemer** is a Principal Research Engineer at Symantec Corporation. He obtained an MSEE degree at the Georgia Institute of Technology.



**Sarfaraz Hussein** is a Research Scientist and Engineer at Symantec Corporation. He obtained a PhD in Computer Vision / Medical Imaging at the University of Central Florida.



**Jay Dhaliwal** is a Sr Machine Learning Research Engineer at Symantec Corporation. He obtained an MS in Computer Science from the University of Massachusetts Amherst.



**Saurabh Shintre** is a Sr Principal Research Engineer at Symantec Corporation. He obtained a PhD in Electrical & Computer Engineering from Carnegie Mellon University.



**Slawomir Grzonkowski** is a Sr Principal Research Engineer at Symantec Corporation. He obtained a PhD in Computer Science from the National University of Ireland.



**Andrew Gardner** is a Senior Technical Director at Symantec Corporation. He obtained a PhD in Electrical Computer Engineering from the Georgia Institute of Technology.