

# Cry Wolf: Toward an Experimentation Platform and Dataset for Human Factors in Cyber Security Analysis

William Roden

University of North Carolina, Wilmington  
Wilmington, North Carolina, USA  
roden011@gmail.com

Lucas Layman

University of North Carolina, Wilmington  
Wilmington, North Carolina, USA  
laymanl@uncw.edu

## ABSTRACT

Computer network defense is a partnership between automated systems and human cyber security analysts. The system behaviors, for example raising a high proportion of false alarms, likely impact cyber analyst performance. Experimentation in the analyst-system domain is challenging due to lack of access to security experts, the usability of attack datasets, and the training required to use security analysis tools. This paper describes Cry Wolf, an open source web application for user studies of cyber security analysis tasks. This paper also provides an open-access dataset of 73 true and false Intrusion Detection System (IDS) alarms derived from real-world examples of *impossible travel* scenarios. Cry Wolf and the impossible travel dataset were used in an experiment on the impact of IDS false alarm rate on analysts' abilities to correctly classify IDS alerts as true or false alarms. Results from that experiment are used to evaluate the quality of the dataset using difficulty and discrimination index measures drawn from classical test theory. Many alerts in the dataset provide good discrimination for participants' overall task performance.

## CCS CONCEPTS

• **Security and privacy** → **Usability in security and privacy**; *Intrusion detection systems*.

## KEYWORDS

cyber security, human factors, IDS

### ACM Reference Format:

William Roden and Lucas Layman. 2020. Cry Wolf: Toward an Experimentation Platform and Dataset for Human Factors in Cyber Security Analysis. In *2020 ACM Southeast Conference (ACMSE 2020)*, April 2–4, 2020, Tampa, FL, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3374135.3385301>

## 1 INTRODUCTION AND BACKGROUND

Computer network defense is a partnership between humans and machines. Machines are necessary to process voluminous network data, and humans must investigate suspicious behaviors identified on the network. Intrusion Detection Systems (IDSes) inspect network activity for known attack signatures, violations of user-configured rules, and statistical anomalies [6]. IDSes then alert human cyber analysts of suspicious activity. However, IDS systems

often produce false alarm rates above 95%, requiring cyber analysts to evaluate hundreds or thousands of false alarms [10]. This behavior may lead to vulnerabilities at the analyst-system interface; research shows that low true event probability and high false alarm rate reduce human operator performance [2, 3]. While researchers strive to improve the accuracy of IDSes, comparatively little research has been published on how the *behaviors* of IDSes and other semi-automated cyber security systems impact the performance of human cyber analysts.

Experimentation on cyber security analyst performance is challenging. Cyber security professionals are difficult to access due to the sensitivity of their work. Further, security-focused datasets that may be used to simulate real attacks in an experimental environment (e.g., [4]) are often raw network captures and system logs rather than cyber defense tool output, and training study participants on real IDS and network monitoring systems is intractable. Consequently, few quantitative experiments have been published on analyst performance in cyber security tasks (e.g., [1, 7]).

This paper presents two open-access resources that help address the need for controlled experimentation of cyber analyst performance. Section 2 describes an IDS alarm dataset derived from real true and false alarms from the University of North Carolina, Wilmington (UNCW)'s IDSes. Section 3 introduces the open source Cry Wolf web application wherein users evaluate alarms from the dataset, answer a questionnaire on user expertise, and reflect on the task. Finally, Section 4 presents an initial evaluation of the dataset's quality using results from a controlled experiment conducted using the Cry Wolf platform.

## 2 THE IMPOSSIBLE TRAVEL DATASET

The first resource is a dataset of simulated IDS alerts derived from real *impossible travel* alarms. Impossible travel alerts are triggered when a user authenticates from two geographic locations within a period where physical travel between the two locations is impossible, e.g., authentications from London and Moscow with a time between authentications of 30 minutes. Physical travel in this time frame is impossible, but the authentications may be legitimate through technical means such as a Virtual Private Network (VPN).

The dataset contains *true alarms* where the impossible travel alert warrants further investigation for potential malicious activity, and *false alarms* where the alert is not cause for concern. The dataset contains 73 alerts in all: 30 true alarms and 43 false alarms and is available at <https://uncw-hfcs.github.io/ids-simulator-analysis/>. These particular numbers of alerts were needed for the experiment introduced in Section 4 and fully reported in [13]. Each alert contains the following data:

ACMSE 2020, April 2–4, 2020, Tampa, FL, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *2020 ACM Southeast Conference (ACMSE 2020)*, April 2–4, 2020, Tampa, FL, USA, <https://doi.org/10.1145/3374135.3385301>.

- **Cities of Authentication** – The two geographic locations from which the IDS detected an authentication.
- **Number of Successful Logins** – The number of successful authentications from each location in the past 24 hours.
- **Number of Failed Logins** – The number of failed logins from each location in the past 24 hours.
- **Source Provider** – The type of internet provider the authorizations came from at each location. Possible values are:
  - **Telecom** – traditional Internet Service Providers
  - **Mobile/cellular** – wireless carriers
  - **Hosting/server** – hosted service providers, e.g., VPNs, web hosts, and cloud computing
- **Time between Authentications** – The shortest time between authentications from the two cities in the past 24 hours. Reported in decimal hours. This the field that triggers an alarm in a real IDS.
- **VPN Confidence** – A percent likelihood that the user utilized a VPN.

Examples are provided in §2.2. These data fields were chosen based on the first author’s experience as a security operations analyst at UNCW. These fields were most often considered when deciding whether an impossible travel alert was normal from the university IDSes. The data values for events were based on real samples from UNCW’s IDSes with random noise added. Location pairs were generated from combinations of 12 notable cities.

## 2.1 Evaluating Alerts – the Security Playbook

A "Security Playbook" accompanies the dataset and guides how to use the data fields to evaluate whether an alert is a true alarm or false alarm. The guidance reflects how an expert familiar with IDS behavior and the network being protected would evaluate the alerts. The playbook may be viewed at <https://git.io/JvE0I>. This section summarizes the main elements of the playbook as articulated to the reader.

One of the cities in each alert corresponds to legitimate access. Every city has a *concern level* based on a history of attacks: Moscow and Beijing are "High" concern cities, North American cities are "Low" concern, and all others are "Medium" concern. However, users travel legitimately and hosted services may connect from other countries, thus location should not be the sole deciding factor.

*Time between authentications* is the field that triggered the alert. Authentications from different locations in a short time could be an indication of account compromise. The Security Playbook provides a table of typical travel times between all pairs of locations. The *ratio of successful logins to failed logins* from each location may also indicate malicious activity. More failed logins than successful logins could be an indication of password guessing, but legitimate users may fail to authenticate as well, e.g., by using an expired password.

The *source providers* help interpret the other information. A telecom provider implies the login is originating from a user physically at the location, whereas a mobile/cellular or hosting/server provider does not necessarily mean the user is physically present. There is nothing inherently safe or malicious about any type of source provider. The *VPN confidence* percentage is the likelihood that the authentication attempts were made via a VPN.

The Security Playbook contains a "Things to Keep in Mind" section with additional considerations. This section reminds evaluators that IDSes can be inaccurate and that a few, many, or all of the alerts may be false alarms. Another consideration is that users visiting countries with restrictive governments will often use a VPN to circumvent that nation’s firewall. Finally, the section states that it is not unusual for a mobile device to ping the country in which the mobile device is registered when the user is traveling abroad.

## 2.2 Scenarios

The impossible travel dataset covers seven scenarios encountered in real alerts from UNCW’s IDSes.

**2.2.1 True Impossible Travel.** The dataset contains 19 impossible travel alerts that are true alarms. An example is shown in Table 1. In these alerts, the *time between authentications* is less than the time required for a person to travel between the locations as provided in the Security Playbook, and the *source providers* are set to *telecom* to indicate that the persons attempting the authentications are in those physical locations.

**Table 1: Event #66 – A True Alarm with Password Guessing**

City of Authentication	Seattle	Moscow
# Successful Logins	4	11
# Failed Logins	1	3
Source Provider	Telecom	Telecom
Time between Authentications	1.75	
VPN Confidence	0%	

**2.2.2 Password Guessing.** The dataset contains six true alarms where the number of *failed logins* from one or both cities exceeds the number of *successful logins* from that city. In contrast, the ratio of failed-to-successful logins is less than 1.0 in the rest of the dataset.

**2.2.3 Edge Case Travel.** The dataset contains 15 false alarms where the *time between authentications* is within 20 minutes of the typical time between the two cities. The Security Playbook emphasizes the travel time table shows "typical times" and not minima or maxima. All of the cities are "low concern" and the *source providers* are either "telecom" or "mobile/cellular" to encourage focus on the times.

**2.2.4 Eurotrip.** The dataset contains six false alarms with the same features as the "edge case travel" scenarios except that both cities are in Europe. All European cities are of "medium concern", which may lead some evaluators to escalate these scenarios.

**2.2.5 Mobile.** The dataset contains eight false alarms of authentication from a mobile device. Table 2 shows an example. A mobile device may initially route through its home country when traveling abroad as mentioned in the Security Playbook. For all the mobile scenario alerts, the *source provider* is "mobile/cellular" from a city in the USA, and the other city’s *source provider* is "telecom". The idea is the user authenticates from a computer abroad while their mobile device is occasionally pinging their home wireless service.

**Table 2: Event #18 – A False Alarm from Mobile Usage**

City of Authentication	Miami	London
# Successful Logins	3	12
# Failed Logins	2	0
Source Provider	Mobile/Cellular	Telecom
Time between Authentications	0.90	
VPN Confidence	0%	

**2.2.6 VPN.** The dataset contains seven false alarms stemming from users employing VPN services in a location separate from the user's physical location. The *VPN confidence* for these alerts is >90%. One of the cities is "high concern" and has a *telecom* source provider to indicate the user is physically present. The other city is low or medium concern and has a *hosting/service* source provider that is intended to be the VPN service.

**2.2.7 Hosting/servers.** The dataset contains 12 false alarms that entail using a server or hosting service. These alerts involve only low or medium concern cities, and one of the *source providers* is "hosting/server" while the other is "telecom" to indicate the user is geographically situated in one place. The *VPN confidence* values are between 53-71% to distinguish these alarms from the higher confidence VPN scenario with high and medium concern cities.

### 3 THE CRY WOLF PLATFORM

The Cry Wolf web application provides a simulated environment for evaluating IDS alarms from the impossible travel dataset. Cry Wolf is written in the Flask micro-framework for Python. Source code and screenshots of Cry Wolf are available at <https://uncw-hfcs.github.io/ids-simulator/>. The webapp implements an experiment to evaluate the impact of IDS *false alarm rate (FAR)* on analyst performance in correctly classifying alerts as true or false alarms. The webapp provides the following experimentation structure:

- (1) A login page with experiment description and Institutional Review Board information. User logins are assigned by a human proctor, and participants is placed into a 50% or 86% FAR treatment group based on the login name.
- (2) A questionnaire that captures participant expertise in cyber security and networking for use in performance analysis.
- (3) A training page introducing the experiment's scenario, the Security Playbook, and five training alerts with rationale.
- (4) The main alert evaluation task, which displays a table of alerts and links to the alert details. The details are presented similarly to the examples shown in Tables 1–2.
- (5) A post-survey that administers the NASA Task Load Index questionnaire [9] and self-reflection questions on the IDS alert evaluation for use in analysis.

In part (4), the participant reviews the alert data against the Security Playbook and chooses to "escalate" or "don't escalate" the alert, which classifies the alert as a true or false alarm respectively. An analysis script compares participants' answers against an oracle of whether the alerts were true or false alarms, and generates a confusion matrix to derive classification performance measures of

sensitivity, specificity, and precision. The repository of analysis scripts is <https://github.com/uncw-hfcs/ids-simulator-analysis>.

The Cry Wolf platform has several benefits for user experimentation. Nearly all of the experimental procedure is automated in the application; the only intervention required by the experimenter is to obtain informed consent and assign a pre-generated login code. The platform captures precise data on when the participants view and classify each alert, enabling fine-grained analysis of time-on-task. Subjects can exit and resume the experiment as necessary using their assigned logins, and the web application can be deployed on a cloud server for maximum availability. Cry Wolf is open source, and experimenters knowledgeable of Python and web development can adapt it to variations of the experimental structure. Further details the Cry Wolf platform structure and implementation are provided in [13].

### 4 EVALUATION OF THE IMPOSSIBLE TRAVEL DATASET

This section presents an initial evaluation of the quality of the impossible travel dataset for use in controlled experimentation of cyber analyst performance. In Fall 2019, 51 individuals participated in an experiment using the Cry Wolf platform. The goal of that experiment was to evaluate the impact of IDS *false alarm rate* on analysts' abilities to *correctly classify IDS alerts as true or false alarms*. The participants were randomly assigned to one of two treatment groups: 25 participants were treated with a 50% false alarm rate (25 true and 25 false alarms), and 26 participants were treated with an 86% false alarm rate (seven true and 43 false alarms). Participants were shown alerts from the dataset and classified each as a true or false alarm. The experiment was not time-limited, and the median time to complete the classification of all 50 alerts was 15.6 minutes. The experiment, its initial findings, and threats to validity are fully reported in [13]. The remainder of this section examines the quality of the impossible travel dataset using data from that experiment.

#### 4.1 Measures of Dataset Quality

Classical test theory provides two measures to evaluate the quality of individual alerts: the alert's *difficulty index* and *discrimination index* [5]. These measures are traditionally used to evaluate the quality questions in psychometric and educational tests.

An alert's *difficulty index*,  $p$  is the *proportion* of participants who correctly classified an alert as a true or false alarm to the total number of people who classified the alert. Lord suggests a target of  $p \approx 0.85$  for items with a binary response to maximize validity of the overall test while accounting for random guessing [12]. An alert with  $p = 1.0$  indicates that the correct classification may be obvious to the participants, whereas an alert with a  $p \ll 0.85$  may indicate that the alert data is insufficient or misleading.

An item's *discrimination index*,  $D$ , measures the difference in responses to the item between high- and low-performers on the overall task. Each participant's number of correctly classified alerts is counted. Two groups are formed from the 27% of participants with the highest and lowest scores to minimize chance error [11]. To calculate an alert's  $D$ , subtract the number correct in the low group from the number correct in high group, then divide the difference by the size of the larger group. Values of  $D$  are in the range  $[-1.0, 1.0]$ .

**Table 3: Item Difficulty and Discrimination Index Statistics per False Alarm Rate (FAR) Treatment**

	50% FAR		86% FAR	
	$p$	$D$	$p$	$D$
mean	0.75	0.30	0.76	0.41
std	0.22	0.30	0.20	0.35
min	0.29	-0.29	0.16	-0.29
$Q_1$	0.60	0.00	0.62	0.14
$Q_2$ (median)	0.76	0.29	0.77	0.43
$Q_3$	0.95	0.57	0.92	0.71
max	1.00	0.86	1.00	1.00
participants	25		26	
true alarms	25		6	
false alarms	25		44	

Ebel states that items where  $D > 0.40$  are useful discriminators,  $0.20 \leq D \leq 0.40$  are in need of improvement, and  $D < 0.20$  are poor discriminators and should be eliminated or rewritten [8].

## 4.2 Analysis

Table 3 shows the item difficulty and discrimination index statistics calculated per treatment group. Item difficulty ( $p$ ) scores are similar across groups, which suggests the dataset is reliable across treatments. The mean and median discrimination index scores ( $D$ ) are in the range warranting improvement per Ebel [8]. The negative min  $D$  value shows that at least one alert is misleading because low performers answered it correctly more than high performers.

Table 4 bins the alerts according to their  $p$  and  $D$  scores. Approximately 40% of the alerts evaluated by each group have a good discrimination index per Ebel [8]. Approximately 30% of the alerts fall into the "too easy" category where many of the participants correctly classified the alert. In the Cry Wolf experiment, unlike in educational tests, such alerts are useful because they help build and reinforce the participants' notions of true and false alarms.

The alarms that are "too hard" are cause for concern as they may indicate that the Cry Wolf training exercises are insufficient or the alerts are unclear. Two *eurotrip* alerts have travel times that are nearly impossible but are false alarms – likely the participants were erring on the side of caution. Four alerts from the *mobile* scenario fell into the "too hard" category for the 50% FAR group. One possible explanation is that the participants did not read or remember the line near the end of the Security Playbook stating that mobile devices abroad often ping their home city first. The other alerts in the "too hard" group exhibit no obvious pattern. The

**Table 4: Aggregate Measures of Item Analysis. Values are the Number of Alerts**

	50% FAR	86% FAR
$D > 0.4$ (best)	24	27
$p \geq Q_3$ and $D \leq 0.4$ (too easy)	13	17
$p < Q_2$ and $D \leq 0.4$ (too hard)	7	2

$D$  score is sensitive to random noise given the small sizes of the high and low groups ( $n = 7$  in each group) for each treatment.

Overall, the difficulty and discrimination indices suggest that the impossible travel dataset is useful to discriminate between high- and low- performers. Further qualitative investigation on why the "too hard" alerts were mis-classified is warranted.

## 5 CONCLUSION AND FUTURE WORK

This paper presents a dataset of simulated IDS alarms, introduces the Cry Wolf platform for controlled experiments of cyber analyst performance, and provides an initial evaluation of the dataset's quality. The impossible travel dataset and Cry Wolf platform are open source, and the evaluation shows promising results for discriminating between high and low performers.

In future work, the impossible travel dataset will be expanded to include new alarm types. The Cry Wolf platform will be used in Summer 2020 to study the effects of *anchoring bias* [14] on cyber analyst performance. These studies are part of a research plan to investigate factors impacting cyber analyst performance. The long term goals of this research are to develop new training methodologies for coping with imperfect cyber defense systems, provide guidance for tool developers on user interface considerations, and develop defense strategies against analyst-machine vulnerabilities.

## REFERENCES

- [1] N. Ben-Asher and C. Gonzalez. 2015. Effects of Cyber Security Knowledge on Attack Detection. *Computers in Human Behavior* 48 (7 2015), 51–61. <https://doi.org/10.1016/j.chb.2015.01.039>
- [2] J. P. Bliss and M. C. Dunn. 2000. Behavioural Implications of Alarm Mistrust as a Function of Task Workload. *Ergonomics* 43, 9 (9 2000), 1283–1300. <https://doi.org/10.1080/001401300421743>
- [3] S. Breznitz. 1984. *Cry Wolf: The Psychology of False Alarms*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- [4] Canadian Institute for Cybersecurity. 2019. Datasets. <https://www.unb.ca/cic/datasets/index.html>
- [5] L. Crocker and J. Algina. 1986. *Introduction to Classical and Modern Test Theory*. Wadsworth/Thomson Learning, Belmont, CA, USA.
- [6] D. E. Denning. 1987. An Intrusion-Detection Model. *IEEE Transactions on Software Engineering* 13, 2 (1987), 222–232. <https://apps.dtic.mil/dtic/tr/fulltext/u2/a484998.pdf>
- [7] V. Dutt, Y.-S. Ahn, N. Ben-Asher, and C. Gonzalez. 2012. Modeling the Effects of Base-Rates on Cyber Threat Detection Performance. In *Proceedings of the 11th International Conference on Cognitive Modeling (ICCM 2012)*. Universitätsverlag der TU, Berlin, Berlin, Germany, 88–93.
- [8] R. L. Ebel. 1979. *Essentials of Educational Measurement* (3rd ed.). Prentice-Hall, Inc., Englewood Cliffs, NJ.
- [9] S. G. Hart and L. E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Advances in Psychology* 52, C (1 1988), 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- [10] K. Julisch. 2003. Clustering Intrusion Detection Alarms to Support Root Cause Analysis. *ACM Transactions on Information and System Security* 6, 4 (2003), 443–471. <https://doi.org/10.1145/950191.950192>
- [11] T. L. Kelley. 1939. The Selection of Upper and Lower Groups for the Validation of Test Items. *Journal of Educational Psychology* 30, 1 (1 1939), 17–24. <https://doi.org/10.1037/h0057123>
- [12] F. M. Lord. 1952. The Relation of the Reliability of Multiple-Choice Tests to the Distribution of Item Difficulties. *Psychometrika* 17, 2 (1952), 181–194. <https://doi.org/10.1007/BF02288781>
- [13] W. T. Roden. 2019. *An Empirical Study of Factors Impacting Cyber Security Analyst Performance in the Use of Intrusion Detection Systems*. Master's thesis. University of North Carolina, Wilmington, Wilmington, NC.
- [14] A. Tversky and D. Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 4157 (9 1974), 1124–31. <https://doi.org/10.1126/science.185.4157.1124>