

PhishNet:

Measuring the Effectiveness of Phishing Awareness Training

Group Members: Samuel Castellanos, Devin Giles, Carlos Rey, Ahsan Sanders

CIS 4951: Capstone II

Prof. Masoud Sadjadi

4/15/2026

Abstract

Phishing attacks remain a major cybersecurity threat due to their focus on human behavior. This project presents PhishNet, a system that analyzes phishing susceptibility using simulated data in an ethical and controlled environment. The system integrates a fast API backend with a web based dashboard to compute and visualize metrics such as failure rate and reporting rate. Our simulated results show that 25% of users failed phishing attempts overall, with trained users at 13% and untrained users at 37%. These findings demonstrate that awareness training significantly reduces phishing susceptibility, while also highlighting the importance of structured evaluation methods in cybersecurity.

1.Introduction

Phishing attacks remain one of the most common and effective cybersecurity threats, primarily because they exploit human behavior rather than technical vulnerabilities. Attackers use deceptive emails and messages to trick users into revealing sensitive information, clicking malicious links, or downloading harmful content. Despite advances in security technologies like firewalls and intrusion detection systems, the human element continues to be a significant point of weakness. Industry reports, including Verizon's 2025 Data Breach Investigations Report, indicate that the human element is involved in a large percentage of security breaches, highlighting the importance of addressing user behavior in cybersecurity. As a result, phishing awareness training has become a widely adopted strategy for reducing user susceptibility. Organizations increasingly rely on simulated phishing campaigns and training programs to educate users and improve their ability to recognize and respond to phishing attempts. However, evaluating the effectiveness of such training remains a challenge. Simple metrics such as click rates or failure rates do not always provide a complete picture of user awareness, and results can vary depending on factors such as message difficulty and user context.

This project introduces PhishNet, a phishing awareness analysis system designed to evaluate and visualize user susceptibility to phishing attacks in a structured and ethical manner. PhishNet uses simulated phishing data to compute key metrics, including failure rate and reporting rate, and presents these results through an interactive dashboard. The system enables comparison between trained and untrained user groups, allowing for clear visualization of how awareness training impacts user behavior. The goal of this project is not to claim definitive real world effectiveness of phishing training, but rather to provide a framework for analyzing phishing awareness outcomes and communicating trends in a clear and accessible way. By combining data processing with visualization, PhishNet contributes to cybersecurity education and highlights the importance of user focused defenses in mitigating phishing risks.

2. Background

Phishing is a form of social engineering attack in which attackers attempt to trick users into revealing sensitive information or performing actions that compromise their security. These attacks commonly take the form of fraudulent emails, messages, or websites that appear to originate from trusted sources. Unlike many other cybersecurity threats that target technical vulnerabilities, phishing primarily exploits human decision making, making it particularly difficult to defend against using traditional security mechanisms alone. The effectiveness of phishing attacks is closely tied to psychological factors such as trust, urgency, and familiarity. Attackers often design messages that mimic legitimate messages and create a sense of urgency to encourage quick responses. As a result, even users with some level of technical knowledge can fall victim to well crafted phishing attempts. This human centered nature of phishing has made it one of the most persistent threats in modern cybersecurity.

To address this challenge, organizations have increasingly adopted phishing awareness training programs. These programs aim to educate users on how to recognize suspicious messages and respond appropriately. Training often includes simulated phishing campaigns, where users are exposed to realistic phishing scenarios in a controlled environment. The goal is to improve user awareness and reduce the likelihood of successful attacks over time.

Evaluating the effectiveness of phishing awareness training is an important but complex task. Common metrics include failure rate, which measures the percentage of users who fall for phishing attempts, and reporting rate, which measures how often users correctly identify and report suspicious messages. However, research from the National Institute of Standards and Technology (NIST) suggests that these metrics should be interpreted carefully, as phishing difficulty can vary depending on factors such as message content and user context. The NIST Phish Scale provides a framework for understanding how different phishing messages may affect user detection.

Given these challenges, there is a need for systems that can analyze and visualize phishing awareness outcomes in a structured and interpretable way. PhishNet builds on these concepts by using simulated data to compute key metrics and present them through a dashboard interface. This approach supports both technical analysis and educational understanding of phishing risks and user behavior.

3. Motivation

Despite widespread awareness of phishing as a cybersecurity threat, it continues to be one of the most successful attack methods in practice. Organizations invest heavily in security technologies and awareness programs, yet users still fall victim to phishing attempts at significant rates. This persistence highlights a fundamental challenge: while technical defenses have improved, the human element remains difficult to manage and evaluate effectively. Existing phishing awareness training programs aim to reduce user susceptibility, but their effectiveness is not always clear. Many organizations rely on basic metrics such as click rates or failure rates to assess user performance. However, these metrics can be misleading when considered in isolation, as they do not account for the varying difficulty of phishing messages or the context in which users interact with them. Research and industry guidance, including work from NIST, emphasize that phishing outcomes should be interpreted using multiple factors rather than a single metric.

In addition, recent studies have shown that the effectiveness of phishing training is not always consistent across different environments. Some training approaches may produce measurable improvements in controlled settings, while others show limited impact in real world organizational contexts. This uncertainty makes it important to develop tools that focus not only on training itself, but also on how training outcomes are analyzed and communicated.

The motivation behind PhishNet is to address this gap by providing a structured and transparent way to evaluate phishing awareness outcomes. Instead of attempting to implement or test a specific training program, PhishNet focuses on analyzing user behavior using simulated data and presenting the results through an interactive dashboard. By emphasizing visualization, comparison, and interpretation, the system helps make phishing risk more understandable and highlights the importance of evidence based evaluation in cybersecurity.

4. Design

The PhishNet system is designed as a full stack application that integrates a backend data processing layer with a frontend visualization dashboard. The architecture follows a client-server

model, where the backend is responsible for computing phishing awareness metrics and the frontend presents these results in an accessible and interactive format. This separation of concerns allows the system to remain modular, scalable, and easy to maintain.

The backend of PhishNet is implemented using the FastAPI framework in Python. It serves as the core data processing component of the system. The backend processes simulated phishing datasets and computes key metrics such as total users, failure rate, and reporting rate. These metrics are exposed through a RESTful API endpoint, allowing the frontend to request data dynamically. The backend also supports filtering based on training status, enabling comparisons between trained and untrained user groups. This design allows for efficient data retrieval and supports real-time interaction with the frontend.

The frontend is implemented as a web based dashboard using HTML, CSS, and JavaScript. It utilizes the Chart.js library to visualize phishing metrics through graphical representations such as bar charts. The dashboard retrieves data from the backend API and dynamically updates key performance indicators, including failure rates and reporting rates. Users can interact with the system through filtering options, allowing them to view results for all users, trained users, or untrained users. This interactive design enhances usability and makes the system accessible to both technical and non-technical audiences.

The overall data flow begins with the frontend sending a request to the backend API. The backend processes the request, computes the relevant metrics, and returns the results in JSON format. The frontend then parses this data and updates both the numerical metrics and visual charts displayed on the dashboard. This real time interaction enables users to quickly interpret phishing awareness outcomes and compare trends across different user groups.

By combining a structured backend with an interactive frontend, PhishNet provides a cohesive system for analyzing and visualizing phishing susceptibility. The design emphasizes clarity, usability, and modularity, ensuring that the system can effectively support both analysis and educational objectives.

5. Evaluation

5.1 Experimental Setup

The evaluation of the PhishNet system is conducted using simulated phishing datasets derived from publicly available cybersecurity research and industry reporting. These datasets model user interactions with phishing attempts and categorize users based on training status, allowing for controlled comparisons between trained and untrained groups. Because the system relies on simulated data rather than real phishing campaigns, the evaluation remains ethical, reproducible, and appropriate for an academic environment.

The backend processes these simulated user interactions and computes key phishing awareness metrics, including failure rate, reporting rate, and overall user outcomes. The results are then visualized through the frontend dashboard, allowing for clear interpretation of trends and comparisons across user groups.

5.2 Metrics and Benchmarks

The primary metric used in this evaluation is the phishing failure rate, defined as the percentage of users who fail to correctly identify a phishing attempt. This metric directly reflects user susceptibility to phishing attacks. However, failure rate alone is not sufficient to fully capture user awareness.

To provide a more comprehensive evaluation, the system also considers reporting rate and comparative performance between trained and untrained users. This approach aligns with guidance from the National Institute of Standards and Technology (NIST), which emphasizes that phishing outcomes should be interpreted in context rather than relying on a single metric.

To contextualize results, PhishNet compares simulated outcomes with industry benchmarks. Reports such as Verizon's 2025 Data Breach Investigations Report highlight the continued importance of the human element in cybersecurity incidents. Additionally, phishing awareness benchmarking studies indicate that untrained users often exhibit failure rates in the range of 30–35%, while trained users typically show significantly lower rates after sustained training.

5.3 Results Overview

The results produced by PhishNet demonstrate a clear difference between trained and untrained users within the simulated environment. Overall, approximately 25% of users failed phishing attempts. When segmented by training status, trained users exhibited a failure rate of approximately 13%, while untrained users showed a significantly higher failure rate of approximately 37%.

These results are consistent with general trends observed in both academic research and industry reports, where repeated training and awareness efforts are associated with reduced phishing susceptibility. The dashboard visualization further enhances interpretation by presenting these differences in a clear and accessible format, allowing users to quickly compare performance across groups.

5.4 Discussion and Interpretation

The evaluation results suggest that phishing awareness training can reduce user susceptibility within a controlled, simulated environment. The clear difference between trained and untrained users demonstrates that the system is capable of modeling and visualizing meaningful behavioral trends.

However, it is important to interpret these results with caution. Existing research indicates that the effectiveness of phishing training can vary depending on factors such as training design, user population, and real world context. Some studies suggest that commonly used awareness approaches may not always produce consistent improvements in practice.

For this reason, PhishNet should be understood as an analytical and educational tool rather than a definitive measure of real world training effectiveness. Its primary contribution is to provide a structured framework for comparing outcomes, visualizing trends, and supporting evidence based discussions about phishing risk and user behavior.

5.5 Limitations

Several limitations affect the interpretation of this evaluation. First, the use of simulated data means that the system cannot fully capture real world conditions such as user distraction,

time pressure, or highly personalized phishing attacks. These factors can significantly influence user behavior in practice.

Second, benchmark comparisons are approximate and should not be interpreted as exact targets. Industry datasets are collected under different conditions and may not directly align with the simulation environment used in this project.

Finally, phishing failure rate alone does not fully represent user resilience. As noted in NIST guidance, additional factors such as reporting behavior and message difficulty should be considered when evaluating phishing awareness outcomes.

5.6 Summary of Evaluation

Overall, the evaluation demonstrates that PhishNet can effectively model, compute, and visualize phishing awareness outcomes using simulated data. The results show a consistent and interpretable difference between trained and untrained users, supported by both numerical metrics and dashboard visualizations.

While the system does not claim to replicate real world conditions, it provides a valuable framework for analyzing phishing susceptibility, communicating risk trends, and supporting cybersecurity education through structured evaluation.

6. Related Work

Phishing awareness and user susceptibility have been widely studied in both academic research and industry practice. A consistent finding across the literature is that phishing is not solely a technical issue but also a human centered problem. Industry reports such as Verizon's 2025 Data Breach Investigations Report highlight that the human element continues to play a significant role in security breaches, reinforcing the importance of evaluating user behavior in cybersecurity.

One important area of related work focuses on how phishing outcomes should be measured. The National Institute of Standards and Technology (NIST) developed the Phish Scale, which emphasizes that phishing detection difficulty varies depending on message

characteristics and user context. This framework highlights the limitations of relying solely on simple metrics such as click rates or failure rates and encourages a more nuanced interpretation of phishing outcomes. This perspective is directly relevant to PhishNet, which also evaluates user susceptibility using multiple metrics and emphasizes contextual interpretation.

Another body of research examines the effectiveness of phishing awareness training programs. Studies such as those conducted in healthcare environments have shown that users can remain vulnerable even after exposure to training, indicating that awareness alone does not eliminate risk. More recent reviews of cybersecurity training methods suggest that while many studies report positive outcomes, the evidence base is often limited by small sample sizes, non experimental designs, and inconsistent methodologies.

Recent empirical research has further highlighted the complexity of phishing training effectiveness. Some studies have found that commonly used training approaches may not produce significant improvements in real world settings. These findings suggest that organizations should not assume that training programs automatically reduce phishing risk, but instead should carefully evaluate outcomes using structured and evidence based methods.

In addition to academic research, industry benchmarking reports provide valuable context for phishing awareness performance. Large scale datasets from security awareness platforms indicate that untrained users often exhibit relatively high phishing failure rates, while sustained training can lead to measurable improvements over time. Although such benchmarks should be interpreted cautiously, they provide useful reference points for evaluating systems like PhishNet.

PhishNet builds upon these areas of work by providing a simulation based platform for analyzing phishing awareness outcomes. Rather than attempting to implement or validate a specific training program, the system focuses on measuring, visualizing, and interpreting user behavior. This positions PhishNet as a complementary tool within the broader field of phishing awareness research, supporting both analysis and education.

7. References

Verizon. *2025 Data Breach Investigations Report*. 2025.

https://www.verizon.com/business/resources/T16f/reports/2025_dbir_data_breach_investigations_report.pdf

Dawkins et al. *NIST Phish Scale User Guide (NIST TN 2276)*. National Institute of Standards and Technology, 2023.

<https://nvlpubs.nist.gov/nistpubs/TechnicalNotes/NIST.TN.2276.pdf>

Steves, Greene, and Theofanos. “Categorizing Human Phishing Detection Difficulty: A Phish Scale.” *Journal of Cybersecurity*, 2020.

<https://academic.oup.com/cybersecurity/article/6/1/tyaa009/5905453>

Gordon et al. “Evaluation of a Mandatory Phishing Training Program for High Risk Employees at-a U.S. Healthcare System.” *JAMIA*, 2019.

<https://academic.oup.com/jamia/article-abstract/26/6/547/5376646>

Prümmer, van Steen, and van den Berg. “A Systematic Review of Current Cybersecurity Training Methods.” *Computers & Security*, 2024.

<https://www.sciencedirect.com/science/article/pii/S0167404823004959>

Marshall and Sturman. “Exploring the Evidence for Email Phishing Training: A Scoping Review.” *Computers & Security*, 2024.

<https://www.sciencedirect.com/science/article/pii/S0167404823006053>

Ho et al. *Understanding the Efficacy of Phishing Training in Practice*. IEEE Symposium on Security and Privacy, 2025.

https://people.cs.uchicago.edu/~grantho/papers/oakland2025_phishing_training.pdf

KnowBe4. *2025 Phishing by Industry Benchmarking Report* / related press summary, 2025.

<https://www.knowbe4.com/press/knowbe4-report-reveals-security-training-reduces-global-phishing-click-rates-by-86>