GoPhish

Personalized Anti-Phishing Training

Team 56
Advised by: Kris Varhus (Faculty), Prekshi Vyas (TA)
Team: Ishir Vaidyanath, Vishakh Arora, Anjalee Narenthiren, Kasyap Chakravadhanula

Executive Summary

For the first time, technology can mimic human thought, speech, and even imagery so convincingly that most people struggle to differentiate between what is real and what is artificial. This advancement has paved the way for a new breed of security attacks, with phishing evolving into an increasingly sophisticated threat. AI agents have the capability to personalize messages, style, and content, crafting ultra-targeted phishing attacks that are more convincing than ever. As a result, individuals and organizations are at greater risk, with the potential damage of such attacks only increasing over time. Phishing scams specifically have become a significant cybersecurity threat, with increasingly sophisticated AI-generated tactics exploiting personalization to inflict substantial financial and data security damages.

GoPhish addresses this escalating threat by providing highly personalized phishing training simulations. Using AI and data retrieval techniques, our platform simulates elaborate phishing attacks, identifying human weaknesses that traditional methods often overlook. By dynamically adapting to organizational requirements, creating custom phishing scenarios, and providing accurate and powerful phishing analytics, GoPhish aims to bridge the gap left by existing generic and inflexible solutions. This proactive approach allows organizations to adapt to emerging threats and stay ahead of attackers who are continuously refining their tactics. Our goal is to create a more secure digital landscape by providing a robust, scalable solution that evolves with the constantly expanding universe of phishing scams.

The culmination of our efforts this semester is a functional MVP tested in a pilot program with Penn students, in addition to analytics features designed to measure effectiveness. Moving forward, we plan to enhance our platform by integrating more advanced data summarizations, creating custom data profiles for users, and running additional pilot tests with Penn faculty.

Value Proposition

GoPhish offers a unique approach to phishing prevention by enhancing traditional phishing training with AI-driven adaptability. Unlike existing tools, which often lack contextual relevance and accurate analytics, GoPhish generates targeted simulations based on real-time data, increasing efficiency. Our platform helps small to medium organizations protect their employee and customer data, ultimately reducing the risk of financial and reputational harm.

Key differentiators:

1. **Personalized training simulations:** GoPhish creates phishing scenarios for each organization with content specific to employees, increasing engagement and awareness.
    a. Our solution generates a fake website where users have the ability to enter PII, a differentiating feature that allows us to measure how deep a user actually goes within a phishing scam, leading to more personalized training options and insights.
2. **Integration with organization-specific data:** Our solution dynamically adapts to organizational data, ensuring relevance. Proprietary access to company data allows our

solution to create more advanced phishing scams than scammers themselves, better preparing individuals for any new threats that come their way. GoPhish is not an email filter - rather our goal is to train employees to detect phishing themselves, which we believe to be a more effective way of combating phishing rather than attempting to detect all new scams ourselves.

3. **Advanced analytics to track and improve campaign effectiveness:** GoPhish offers actionable insights into user behavior and areas for improvement by measuring how deep an individual goes into a scam, rather than just detecting the scam. Most solutions only record clicks onto scams rather than whether any PII was entered, resulting in false positives in reported statistics as a person might not have entered PII even if they did click on a scam.
4. **Scalability and adaptability for diverse industries:** A generalized approach to generating emails means our platform is designed to scale with the needs of organizations across several industries.

Stakeholders

1. **Primary Users:** Educational institutions, including faculty, students, and IT staff, who directly interact with the phishing simulation and training platform and have the largest responsibility to keep data safe.
2. **Decision-Makers:** IT security teams and university administrators who are responsible for selecting and implementing cybersecurity solutions.
3. **Development Partners:** Mentors and cybersecurity industry experts who provide feedback and guidance to refine the product.
4. **Beneficiaries:** Employees and students who benefit from reduced data compromise risk and improved phishing/cybersecurity awareness.

Market Research

Phishing scams are one of the fastest-growing cybersecurity threats, costing businesses an estimated **$12.5 billion** globally in 2023 and surging by **58%** this year. Furthermore, while the losses have been increasing, the number of reported incidents received by the FBI has slightly decreased, reflecting how organizations are becoming increasingly unaware that they're even being scammed and that scams are becoming more harmful. The phishing simulation market is expected to grow at a double digit compound annual growth rate (CAGR) of **14%** between 2024 and 2030, reflecting increasing awareness and adoption of training solutions.

Surveys conducted during our research revealed that:

- **60%** of participants in a Harvard Business Review study were susceptible to AI-generated phishing attacks.
- Feedback from educational institutions highlighted gaps in current solutions, including lack of personalization, unreliable statistics, and cumbersome implementation.

Based on interviews with faculty and IT personnel at Penn Engineering, we identified a gap in phishing training tailored to organizational contexts. Highly complex organizations such as Penn use one person to manually write and send out emails, and don't have specific training or feedback measures These insights guided our product design and market entry strategy (see Appendix for survey data and testimonials).

Customer Segment

1. **Educational Institutions (primary):** Universities and colleges with high exposure to phishing threats due to extensive online interactions with third parties and reliance on digital communication. We will focus on this market first as they manage the most amount of data with the least amount of sophistication in security (compared to large companies managing the same volume of data).
2. **IT Security Teams:** Responsible for implementing training programs, maintaining organizational cybersecurity, and creating phishing simulations. We will target small to medium sized organizations, as larger companies tend to already have a robust security team in-house to tackle any security threats, which we found through our interviews with Penn faculty that have previously worked at Cisco and DARPA.

Secondary markets include nonprofit and public sector organizations looking for tailored cybersecurity training solutions.

Competition

Key competitors in the automated phishing simulation market include Proofpoint, KnowBe4, Ironscales, CultureAI, and Infosec. While these platforms provide phishing prevention solutions, they are focused on detecting user clicks and providing general training rather than understanding the depth of engagement with phishing scams. This limits their effectiveness in addressing the root causes of susceptibility to phishing.

GoPhish stands out in the competitive landscape by optimizing for two critical dimensions determined through our user interviews: **ease of use** and **personalization**.

- **Dynamic Content:** GoPhish generates highly tailored phishing simulations similar to real-world scenarios using public and private employee data, ensuring higher training relevance.
- **Comprehensive Feedback:** Beyond click detection, our solution assesses user interactions with phishing attempts, offering deeper insights into user behavior and vulnerabilities.
- **Ease of Implementation:** The platform is intuitive and integrates seamlessly with organizational systems, reducing administrative burdens.

Intellectual Property (IP)

Our proprietary approach includes:

1. **AI-Driven Simulation Algorithms:** Using LLMs fine-tuned for phishing email generation ensures realistic and dynamic phishing scenarios.
2. **Custom Data Integration Framework:** Organizations can input specific data and context, resulting in unique training scenarios aligned with their needs.
3. **Analytics Platform:** Offering in-depth insights into campaign performance and user behavior, enabling customers to measure and improve cybersecurity awareness over time.

Cost and Revenue Model

- **Core Technology (Approx. $40,000 – $60,000)**
    - AI Model Development ($20,000 – $30,000)
        - Fine-tuning open-source or licensed LLMs (e.g., GPT-3.5/BERT)
        - Minimal data labeling and prompt engineering
    - Dynamic Scenario Engine ($10,000 – $15,000)
        - Simple rule-based or lightweight ML-based scenario adaptation
    - Multi-Channel Simulation ($10,000 – $15,000)
        - Covers email and SMS phishing scenarios
        - Voice simulation initially deferred or done minimally
- **Supporting Infrastructure (Approx. $20,000 – $35,000)**
    - User Interface ($15,000 – $25,000)
        - Basic web-based dashboard for campaign setup and analytics
        - Initial, lean front-end development costs
    - Security & Pen Testing ($5,000 – $10,000)
        - Essential third-party penetration testing
        - Minimal compliance readiness audits
- **Total Initial Investment**
    - Approximately $60,000 – $95,000

---

2. Monthly Operating Expenses

- **Technology Infrastructure (Approx. $1,500 – $2,500/Month)**
    - Compute Resources ($1,000 – $1,500)
        - Cloud GPU/CPU instances for AI inference
        - Ability to scale up for model retraining and down for normal ops
    - Data Storage ($500 – $1,000)
        - Encrypted user data and phishing simulation logs
        - Basic S3/Blob/Cloud Storage solutions
- **Business Operations (Approx. $5,500 – $8,500/Month)**

    - Threat Intelligence ($1,500 – $2,500)
        - Subscription to basic intelligence feeds and dark web monitoring
        - Possibly utilizing open-source or shared data sources
    - Customer Support ($4,000 – $6,000)
        - One or two support specialists plus minimal AI chatbot handling FAQs
        - Coverage of basic administrative overhead
- **Total Monthly Operating Expenses**
    - Approximately $7,000 – $11,000

---

3. Revenue Model

- **Subscription Tiers**
    - Starter: $1.00/user/month
        - Basic simulations and limited templates
        - Suited for SMBs (<100 employees)
    - Pro: $3.00/user/month
        - AI-driven campaigns and basic risk profiling
        - Aimed at mid-market companies (100–1,000 employees)
    - Enterprise: $5.00/user/month
        - Custom domain setup, SLA guarantees
        - Ideal for larger organizations
- **Discounts**
    - Annual Commitment: 15% discount
    - Non-Profit Organizations: 30% discount
- **Additional Revenue Streams**
    - Custom Phishing Templates: $0.50–$0.80 per template
    - CISO Advisory Services: $200/hour for strategic consulting
    - Industry Benchmarking Reports: $8,000/year for peer comparisons

---

4. Market Performance (Illustrative)

- **Year 1 Target: $600,000**
    - Possible composition: 700 Pro-tier users + 10 Enterprise contracts (each ~500 users)
    - Potential upsell from custom templates or advisory services
- **Year 3 Target: $3.0 – $4.0 Million**
    - Driven by channel partnerships, expanded template marketplaces, deeper compliance modules

---

5. Annual Compliance Investment

- **SOC 2 Readiness: $10,000 – $15,000**
  - Involves hiring a consultant or using readiness tools
  - Basic controls and documentation rather than full certification
- **ISO 27001 (Partial / Gap Audit): $8,000 – $12,000**
  - Phased approach to eventually reach full certification
  - Full ISO 27001 might cost $20k+ down the line
- **Ongoing Security Monitoring: $5,000 – $8,000**
  - Intrusion detection, SIEM tools (e.g., Splunk, Datadog)
- **Total Annual Compliance**
  - Approximately $23,000 – $35,000

---

6. Breakeven Analysis

- **Fixed Costs (Year 1)**
  - Initial Development: $77,500 (midpoint of $60k–$95k)
  - Annual Compliance: $29,000 (midpoint of $23k–$35k)
  - Total Fixed Costs: $106,500
- **Monthly Operating Costs**
  - Technology + Business Ops: $9,000 (approx. midpoint of $7k–$11k)
- **Variable Cost per User**
  - Approximately $0.40 per user/month
  - Covers AI inference, data storage, incremental support
- **Average Revenue per User (ARPU)**
  - Blend of Starter ($1), Pro ($3), and Enterprise ($5)
  - Weighted example: 50% Pro, 40% Starter, 10% Enterprise
  - Approx. $2.50 gross ARPU
  - Minus 15% discount for annual commitments/non-profits → $2.13 net
- **Monthly Contribution Margin per User**
  - $2.13 – $0.40 = $1.73
- **Breakeven User Count**
  - Monthly Fixed Costs = $106,500 / 12 = ~$8,900
  - $8,900 / $1.73 ≈ 5,150 users

With around 5,150 active users (paying a net $2.13 each), the platform can cover monthly operating expenses plus amortized fixed costs.

Citations

[1] "Cost of Phishing Study" - Ponemon Institute
https://ualr.edu/itservices/files/2016/10/Ponemon_Institute_Cost_of_Phishing.pdf

[2] "Phishing Simulation Overview" - CanIPhish https://caniphish.com/phishing-simulation

[3] "Key Factors in Phishing Simulation Costs" - usecure.io
https://blog.usecure.io/4-key-factors-that-affect-the-cost-of-phishing-simulations

[4] "Phishing Simulator Market Analysis" - Coherent Market Insights
https://www.coherentmarketinsights.com/industry-reports/phishing-simulator-market

[5] "Business Cybersecurity Cost Breakdown" - Mamori.io
https://www.mamori.io/blog/cost-breakdown-to-fully-cyber-secure-your-business

[6] "Pricing Structure" - CanIPhish https://caniphish.com/pricing

[7] "Service Pricing" - PhishingBox https://www.phishingbox.com/pricing

[8] "Security Analysis Report" - USENIX
https://www.usenix.org/system/files/usenixsecurity23-brunken.pdf

[9] "Security Awareness Training Costs" - CanIPhish Blog
https://caniphish.com/blog/how-much-does-security-awareness-training-cost

[10] "Calculating Organizational Phishing Costs" - PhishLabs
https://www.phishlabs.com/blog/how-and-why-you-should-calculate-your-organizations-cost-of-phishing

[11] "Top 10 Phishing Costs" - HoxHunt
https://hoxhunt.com/blog/what-are-the-top-10-costs-of-phishing

[12] "Security Awareness Training Pricing" - TitanHQ
https://www.titanhq.com/security-awareness-training/security-awareness-training-pricing/

Appendix

Testimonials

"Can't put a lot of trust in metrics we're getting from our system right now" - Kris Varhus, Chief Information Officer, Penn Engineering

“Technology like Chat GPT enables deep fake emails which are indistinguishable from a legitimate email” - Jonathan Smith, CIS professor and DARPA program manager

“There is one person that drafts these mock emails” - Kris Varhus discussing how Penn conducts phishing simulations currently

Market Research/Survey data:

FBI 2023 crime report, Zscaler 2024 phishing report, SSL Store, HBR Study, GIS Market Report

Pilot solution screenshots:

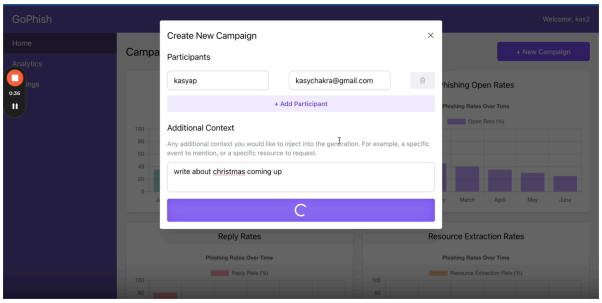

Creating a new phishing campaign on the home page

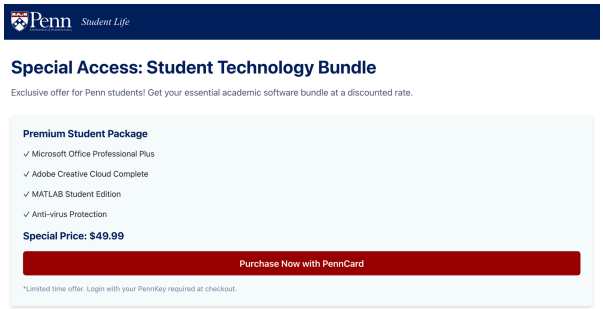

The fake website sent in phishing emails, modeled after a Penn webpage

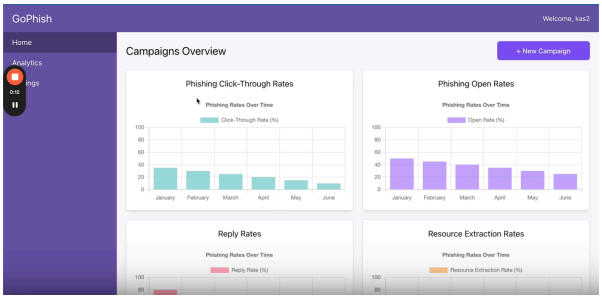

Summary statistics page

Campaign specific metrics