

Grounded but Misleading: Evaluating Semantic Alignment in AI-Generated Security Explanations

Heajun An¹, Connor Ng¹, Sandesh Sharma Dulal², Junghwan Kim², Jin-Hee Cho¹

¹Department of Computer Science, Virginia Tech, Blacksburg, Virginia, U.S.A.

²Department of Geography, Virginia Tech, Blacksburg, Virginia, U.S.A.

{heajun, connorng, jicho}@vt.edu, {sandesh339, junghwankim}@vt.edu

Abstract

Online scams increasingly leverage fluent and context-aware social-engineering strategies, creating growing demand for AI systems that explain why a message may be risky. However, explanations that cite detector-derived evidence may still semantically weaken or redirect the intended risk interpretation. We introduce **VEXA** (Verifying Semantic EXplanation Alignment), a controlled testbed for studying the gap between lexical grounding and semantic risk alignment in AI-generated scam-risk explanations. VEXA generates ungrounded, risk-aligned, and risk-diluting explanations by independently controlling evidence grounding and semantic framing. Across LLM-as-a-judge and human evaluations, explanations can remain comparatively grounded even when their interpretations weaken the detector’s intended risk assessment. In human evaluation, risk-diluting XAI-grounded explanations retained comparatively elevated Perceived Evidence Grounding scores (3.66 ± 1.02) despite lower Helpfulness (3.00 ± 1.41) and Reasoning Support (3.14 ± 1.05) scores. These findings provide controlled evidence of grounding-illusion effects and suggest that trustworthy explanation evaluation must verify not only whether evidence is cited, but also how it is interpreted.

Introduction

Online scams have become a pervasive security challenge across emails, SMSs, and social media, causing billions of dollars in annual losses (Federal Trade Commission 2025). This challenge is further intensified by generative AI, which enables highly fluent and context-aware scam messages at scale (Gressel, Pankajakshan, and Mirsky 2024). As AI-assisted scam detection systems become increasingly deployed, explanations are expected not only to provide binary warnings, but also to justify why a message may be risky.

Explainable Scam Detection. Recent transformer-based detectors have substantially improved scam and phishing detection (Uddin, Mahiuddin, and Chowdhury 2024). In parallel, eXplainable AI (XAI) methods such as SHAP (Lundberg and Lee 2017) and LIME (Ribeiro, Singh, and Guestrin 2016) identify influential scam-related cues. However, they remain largely token-level and model-centric, providing limited explanation for why such cues indicate deception or how users should interpret them in context (Miller 2019).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Dimension	Evaluation Question	Failure Mode
Lexical Grounding	Does the explanation reference detector-derived cues?	Cue citation without aligned interpretation
Semantic Alignment	Does the explanation preserve the detector’s intended risk direction when interpreting those cues?	Benign reinterpretation, uncertainty injection, risk dilution

Table 1: Lexical grounding versus semantic risk alignment.

From Lexical Grounding to Semantic Alignment. Large Language Models (LLMs) can transform detector-derived evidence into natural-language explanations. However, evidence citation alone does not guarantee detector-aligned interpretation. Even when explanations reference suspicious cues, they may weaken or redirect the detector’s intended risk assessment (Huang et al. 2025; Turpin et al. 2023).

Grounding Illusion. We define a *grounding illusion* as an explanation that (1) references detector-derived evidence (*lexical grounding*) while (2) weakening or redirecting the detector’s risk assessment (*semantic misalignment*). Such explanations may appear trustworthy yet induce false reassurance, leading users to miscalibrate trust and make unsafe security decisions. These risks may disproportionately affect digitally vulnerable populations, including older adults, users with low cybersecurity literacy, and those who rely heavily on AI-generated explanations.

Grounding Robustness and Evaluation Reliability. Recent LLM-as-a-judge frameworks enable scalable evaluation of open-ended outputs (Zheng et al. 2023), yet evaluators may over-trust evidence-grounded explanations despite weakened risk interpretation. Figure 1 illustrates this grounding illusion: explanations can preserve detector-derived evidence references while reducing perceived scam severity.

This raises a central research question:

Can evaluators distinguish semantically aligned explanations from evidence-grounded explanations that weaken the detector’s semantic risk assessment?

Motivated by this challenge, we introduce **VEXA** (Verifying Semantic EXplanation Alignment), a controlled framework




Figure 1: Illustration of grounded-looking but semantically misaligned scam-risk explanations.

for evaluating explanation robustness under risk-diluting grounding conditions. VEXA independently manipulates: (1) whether explanations reference detector-derived evidence, and (2) whether those cues are interpreted consistently with the detector’s intended risk assessment.

Specifically, VEXA generates: (1) *Ungrounded* explanations without detector-derived evidence, (2) *Risk-Aligned XAI-Grounded* explanations aligned with detector risk cues, and (3) *Risk-Diluting XAI-Grounded* explanations that preserve grounding while weakening semantic risk interpretation.

Using this setup, we compare LLM-as-a-judge and human evaluations under semantically distinct explanation conditions. Across both settings, explanations may still appear grounded even when their semantic interpretation weakens the detector’s intended risk assessment, revealing limitations of groundedness-based evaluation alone.

Key Contributions. (1) Lexical Grounding vs. Semantic Alignment: We formalize the distinction between lexical grounding and semantic risk alignment in AI-generated security explanations. **(2) VEXA Testbed:** We introduce VEXA, a controlled evaluation framework that independently manipulates evidence citation and semantic risk interpretation across scam-risk explanations. **(3) Controlled Human and LLM Evaluation:** Through controlled human and LLM-based evaluations, we show that evidence-grounded explanations may continue to receive elevated Perceived Evidence Grounding scores despite weakening the detector’s semantic risk assessment. We further introduce the *Grounding-Illusion Gap (GIG)* to measure divergence between perceived grounding and interpretation-oriented semantic quality. **(4) Groundedness Evaluation Limitation:** We show that groundedness-based evaluation may overestimate explanation reliability when evaluators focus primarily on evidence citation without verifying semantic alignment.

More broadly, this work examines how AI explanations can

appear trustworthy while semantically weakening safety-critical interpretations, raising important challenges for trustworthy and human-centered AI deployment.

Related Work

Scam-Risk Detection and Explainability. Transformer-based models have substantially improved scam and phishing detection across email, SMS, and social media (Uddin, Mahiuddin, and Chowdhury 2024). Many systems incorporate XAI methods such as SHAP (Lundberg and Lee 2017) and LIME (Ribeiro, Singh, and Guestrin 2016) to highlight influential tokens or suspicious patterns. However, these approaches remain largely token-level and model-centric, offering limited explanation of why cues such as urgency framing, impersonation, or suspicious links indicate deception (Miller 2019). Consequently, existing explanations provide limited support for user-facing risk interpretation.

Faithfulness, Grounding, and Semantic Alignment. LLMs can generate fluent explanations, but these may rely on generic heuristics or parametric knowledge rather than faithfully reflecting detector reasoning (Jacovi and Goldberg 2020; Turpin et al. 2023). Such unfaithful explanations can produce misleading rationales in safety-critical settings (Huang et al. 2025). Unlike hallucination, grounding illusions may preserve evidence citation while semantically distorting interpretation. Although XAI-derived evidence can improve grounding, evidence inclusion alone does not guarantee semantic risk alignment. The same cited cue may support either scam-risk warning or benign reinterpretation, enabling grounded-looking but semantically misleading explanations.

LLM-as-a-Judge and Evaluation Reliability. LLM-as-a-judge frameworks provide scalable evaluation of open-ended outputs, including explanation quality (Zheng et al. 2023). Compared with lexical metrics, they better assess coherence and reasoning quality (Wang et al. 2023, 2025). However, prior work reports sensitivity to fluency, verbosity, positional bias, and self-preference effects (Zheng et al. 2023; Ye et al. 2025; Wataoka, Takahashi, and Ri 2024), raising concerns about whether evaluators can reliably distinguish evidential grounding from semantic risk alignment.

Human-Centered Security Explanations and Stylistic Framing. Usable security research shows that warnings and interventions must remain understandable and cognitively appropriate for users (Sasse, Brostoff, and Weirich 2001; Egelman, Cranor, and Hong 2008). Prior work further suggests that differences in risk perception and information processing influence the interpretation of scam cues and security explanations (Anawar et al. 2019; Cho, Çam, and Oltramari 2016; Tornblad et al. 2021). Motivated by these findings, we examine persona-adaptive stylistic framing as a secondary factor affecting sensitivity to semantically weakened explanations.

Unlike prior work on evidence citation or fluency, we study whether detector-derived evidence is interpreted consistently with the intended risk assessment. To examine this distinction, we introduce VEXA as a controlled framework for comparing grounded-looking and semantically aligned explanations.




Figure 2: Overview of the VEXA testbed.

Proposed Framework: VEXA

Controlled Explanation Generation Framework

VEXA studies explanation-layer semantic robustness rather than detector evasion. Given a fixed scam-risk detector, VEXA independently controls whether explanations cite detector-derived evidence and whether those cues preserve or weaken the detector’s intended risk interpretation. This separation enables controlled analysis of lexical grounding, semantic alignment, and grounding-illusion effects.

Figure 2 summarizes the VEXA testbed. VEXA extracts detector-based risk evidence from scam messages and separately applies persona-driven style constraints. A controlled prompt constructor then combines these signals to generate explanations under factorial grounding and semantic-alignment conditions. This design allows us to test when explanations appear evidence-grounded while semantically reducing perceived scam risk.

Given an input message, VEXA follows four stages:

- Scam Detection:** A frozen transformer-based detector predicts scam risk and provides a stable attribution source.
- XAI Evidence Extraction:** GradientSHAP identifies salient detector-derived evidence supporting the detector’s scam-risk prediction.
- Stylistic Conditioning:** Supportive or analytical framing modifies explanation style without changing detector-derived evidence.

- Controlled Explanation Generation:** An LLM generates: (a) *Ungrounded* explanations without detector-derived evidence, (b) *Risk-Aligned XAI-Grounded* explanations that preserve the detector’s intended risk interpretation, or (c) *Risk-Diluting XAI-Grounded* explanations that cite the same evidence while weakening or redirecting its risk interpretation.

Controlled Grounding and Semantic Alignment

Scam Detection. VEXA employs a DeBERTa-v3-based encoder (He, Gao, and Chen 2021) fine-tuned for binary scam classification. The detector is frozen after training to provide stable attribution signals during explanation generation, enabling VEXA to isolate explanation-layer robustness rather than detector evasion. Channel-specific markers (email, SMS, SNS) are prepended to preserve modality information.

XAI Evidence Extraction. VEXA applies GradientSHAP (Lundberg and Lee 2017) in embedding space using the Captum interpretability framework to identify salient detector-derived evidence. Attributions are computed with respect to the scam-class logit using zero embedding baselines and 20 attribution samples. Subword attributions are aggregated at the word level, and evidence phrases are ranked by absolute attribution magnitude. The top-10 salient tokens are retained after removing channel markers and non-informative tokens while preserving URLs, currency mentions, and emphatic punctuation. The resulting evidence phrases operationalize *lexical grounding* during explanation generation. However, lexical grounding alone does not guarantee faithful or semantically aligned interpretation, since explanations may still weaken or redirect the detector’s intended risk assessment.

Persona-Based Stylistic Conditioning. VEXA adopts stylized explanation framings inspired by prior work linking communication preferences and scam susceptibility (Ge et al. 2021; Tornblad et al. 2021; Sarno, Harris, and Black 2023). Rather than inferring participant personality traits directly, these framings are used solely to modulate explanation tone and structure under controlled experimental conditions.

Style Framing	Conscientiousness	Neuroticism	Agreeableness
Supportive-Contextual	Low	High	High
Analytical-Concise	High	Low	Low

Table 2: Stylized explanation framings used for controlled persona-adaptive prompting.

Guided by prior literature (Anawar et al. 2019; Cho, Çam, and Oltramari 2016; Tornblad et al. 2021), we use stylized high/low abstractions of conscientiousness, neuroticism, and agreeableness to modulate explanation framing. Supportive-contextual framings use calmer and more contextualized phrasing, whereas analytical-concise framings emphasize direct and evidence-focused reasoning.

We adopt a binary high/low abstraction to preserve control-

lability and interpretability while avoiding fine-grained psychometric modeling or user profiling. Persona conditioning affects linguistic framing only and does not alter detector-derived evidence or the underlying risk signal. This separation enables controlled analysis of how stylistic framing influences perceived groundedness, trust calibration, and sensitivity to semantically weakened explanations.

Controlled Explanation Generation. The final explanation is generated by GPT-5.2 (Singh et al. 2025) using the original message, extracted evidence, and condition-specific instructions. Across all settings, the underlying LLM and base prompting structure remain fixed, while only grounding conditions and stylistic framing vary. This setup enables systematic analysis of explanations that preserve grounding while diverging from the detector’s semantic risk assessment.

To study semantic misalignment, VEXA incorporates risk-diluting interpretation strategies summarized in Table 3. These manipulations preserve evidence references while weakening the detector’s semantic risk assessment.

Type	Example Interpretation
Benign reinterpretation	“routine account verification”
Uncertainty injection	“may simply reflect normal activity”
Risk normalization	“common business practice”
Contextual reframing	“standard security workflow”

Table 3: Representative semantic misalignment in VEXA.

These controlled manipulations enable evaluation of grounding-illusion effects, where explanations continue to appear evidence-grounded despite semantically weakening the detector’s intended risk assessment.

Experimental Setup

Data Collection. To evaluate explanation robustness across heterogeneous communication environments, we construct a multi-channel corpus spanning *email*, *SMS*, and *social media*. All datasets are publicly available and capture both classical and contemporary forms of online deception, including human-authored and AI-generated scam messages (Table 4). Because source datasets differ in label semantics, we harmonize positive labels as broader *scam-risk* rather than treating all positives as identical financial scams. Throughout this work, *scam-risk* includes phishing, impersonation, suspicious links, urgency-based persuasion, spam, and related social-engineering-oriented deception.

To support controlled evaluation across communication contexts, we apply stratified sampling to select balanced scam and ham distributions per channel. This design captures both channel-specific deception patterns and shared social-engineering cues such as urgency framing, impersonation, suspicious links, and reward-based persuasion (Table 5). Duplicate messages are removed before splitting, and a stratified 80/10/10 train/validation/test partition is applied by channel and label to reduce leakage across data splits.

Preprocessing and Controlled Evaluation Filtering. Mes-

Dataset	Channel	Reference
Ling-Spam Dataset	Email	(Sakkis et al. 2003)
Enron-Spam Dataset	Email	(Metsis, Androutsopoulos, and Paliouras 2006)
Human-LLM Phishing AI-Generated Email Dataset	Email	(Greco et al. 2024)
Super SMS Dataset	SMS	(Opapa, Modesti, and Golightly 2025)
UCI SMS Spam Collection	SMS	(Salman, Ikram, and Kaafar 2024)
NUS SMS Corpus	SMS	(Almeida and Hidalgo 2011)
SpamHunter	SMS	(Chen and Kan 2013)
Social Honeypot Dataset	SNS	(Tang et al. 2022)
UTKML Twitter Spam Dataset	SNS	(Lee, Eoff, and Caverlee 2011)
		(Kaggle 2020)

Table 4: Public scam-risk datasets used in VEXA.

Channel	Original Size	Spam / Ham	Sampled Size	Spam / Ham
Email	38,062	18,143 / 19,919	35,000	17,500 / 17,500
SMS	67,008	26,178 / 40,830	35,000	17,500 / 17,500
SNS	573,343	241,379 / 331,964	35,000	17,500 / 17,500

Table 5: Dataset statistics before and after stratified sampling.

sages are standardized into a unified format for scam detection and explanation generation. We prepend channel-specific tokens (<Email>, <SMS>, <SNS>) to preserve modality information. Email messages include subject and body fields, whereas SMS and SNS messages include only body text.

To isolate explanation-layer semantic robustness from detector misclassification, we retain only correctly classified scam-risk instances when generating explanation conditions. This design allows VEXA to evaluate grounding and semantic-alignment failures while holding the detector’s underlying risk assessment fixed. We additionally remove non-English messages and truncate extremely long messages to 1,500 tokens to reduce boilerplate noise.

Detection Performance and Attribution Stability. Table 6 summarizes the performance of evaluated transformer-based scam detectors. DeBERTa-v3-base (He, Gao, and Chen 2021) achieves the highest Macro F1 on both validation and test sets and is therefore selected for all VEXA experiments. Hyperparameters are optimized using Optuna (Akiba et al. 2019), and the selected model is retrained with early stopping to provide stable attribution signals during explanation generation.

Rather than focusing solely on classification accuracy, the detector in VEXA serves as a stable evidence source for controlled evaluation of lexical grounding, semantic alignment, and grounding-illusion effects across explanation conditions.

Model	Val (Macro F1)	Test (Macro F1)
ALBERT-base (Lan et al. 2019)	0.93149	0.93320
DistilRoBERTa-base (Sanh et al. 2019)	0.93342	0.92875
DeBERTa-v3-base (He, Gao, and Chen 2021)	0.93856	0.93875
DeBERTa-v3-small (He, Gao, and Chen 2021)	0.93666	0.92941
RoBERTa-base (Liu et al. 2019)	0.93551	0.93675
Electra-base (Clark et al. 2020)	0.93038	0.92905

Table 6: Performance of transformer-based detectors.

Grounding and Semantic-Alignment Evaluation

Evaluation Samples and Analysis Protocol. LLM-as-a-judge and human evaluations were conducted on 18 scam-risk message instances spanning email, SMS, and social-media contexts. Unless otherwise stated, LLM-as-a-judge statistics are computed at the explanation level, whereas human-evaluation results use participant-item ratings after excluding incomplete responses. Descriptive tables report marginal condition means, while inferential comparisons use participant-level paired aggregates across matched conditions.

Since VEXA studies explanation robustness, not detector performance alone, the evaluation protocol isolates grounding and semantic-alignment effects under controlled settings. This setup tests whether evaluators distinguish semantically aligned explanations from evidence-grounded explanations that weaken the detector’s semantic risk assessment.

Threat Model and Semantic Misalignment Attacks. We study explanation-layer semantic attacks rather than detector evasion. Specifically, we model an adversary that seeks to reduce the perceived severity of a correct warning through semantically misleading explanations.

The attack surface is restricted to the downstream explanation-generation layer, where the LLM reinterprets detector-derived evidence using benign, uncertain, or risk-diluting semantic framing while leaving the detector and attribution signals unchanged. Under these conditions, the detector may correctly identify risk-relevant evidence while the explanation layer generates plausible but semantically weakened interpretations that downplay the intended risk assessment.

Such failures may arise from prompt injection attacks, over-cautious alignment policies, user-support-oriented tone calibration, or unintended semantic drift during LLM generation. In human-facing AI systems, these failures may distort trust calibration by encouraging overreliance on evidence-grounded explanations despite weakened risk interpretation.

Accordingly, the risk-diluting grounding conditions are designed as controlled stress tests for explanation robustness and evaluation reliability rather than a specific exploit pipeline. This setup further enables analysis of grounding-

illusion effects in safety-critical explanation systems.

Table 7 illustrates how explanations may preserve detector-derived evidence references while semantically weakening the detector’s semantic risk assessment.

Component	Example
Message	“Your [BANK] account has been suspended. Please update your account information immediately to avoid account lock.”
Detector Evidence	“[BANK] account,” “suspended,” “update,” “account locked”
Risk-Aligned Explanation	“The message is risky because it uses account-suspension urgency and immediate verification requests, which are common phishing tactics used to obtain financial credentials.”
Risk-Diluting Explanation	“The message may reflect a routine banking security procedure. Temporary suspension and account-update requests are commonly used in fraud-prevention workflows.”

Table 7: Illustrative grounding-illusion example under risk-diluting semantic framing.

Controlled Evaluation Conditions. We investigate how lexical grounding and semantic framing influence explanation robustness and evaluation reliability under both human and LLM-based assessment settings. Rather than optimizing explanation quality alone, we frame the evaluation setup as a controlled robustness probe examining whether evaluators distinguish semantically aligned explanations from grounded-looking but semantically misleading ones.

We manipulate two experimental factors: (1) **semantic condition**, and (2) **evidence/style condition**. Explanations are generated from a stratified 10% subset of scam-risk messages held constant across all conditions.

Semantic Condition. The semantic condition controls whether explanations preserve or weaken detector risk interpretation:

- **Risk-Aligned Interpretation:** Explanations semantically align with the detector’s intended reasoning and interpret salient scam indicators as risk-supporting evidence.
- **Risk-Diluting Interpretation:** Explanations preserve lexical grounding while weakening or redirecting the intended scam-risk interpretation through benign reinterpretation, uncertainty injection, contextual reframing, or risk normalization. Generated risk-diluting explanations are manually reviewed to verify that detector-derived evidence references are preserved while semantic risk interpretation is weakened or redirected.

Evidence and Style Condition. The evidence/style condition controls how explanations reference detector-derived evidence and stylistic framing:

- **No-XAI:** Explanations are generated without detector-derived evidence. Under risk-diluting conditions, this setting tests whether evaluators infer *Perceived Evidence Grounding* from fluent message-specific references alone.

- **XAI-Only:** Explanations are grounded in detector-derived evidence from GradientSHAP without stylistic adaptation.
- **XAI + Supportive Style:** Grounded explanations incorporate supportive and contextualized framing, emphasizing calmer and more reassurance-oriented phrasing.
- **XAI + Analytical Style:** Grounded explanations incorporate concise and analytical framing, emphasizing direct and evidence-focused reasoning.

This factorial design enables controlled analysis of grounding-illusion effects, including whether evidence-grounded explanations continue to appear trustworthy despite semantically weakening the detector’s semantic risk assessment.

Statistical Analysis and Grounding-Illusion Contrasts.

We conduct paired comparisons for key contrasts associated with the grounding-illusion framework. For the human evaluation, participant-item ratings are aggregated at the participant-condition level before inferential testing. We report paired two-sided *t*-tests together with 95% confidence intervals, Holm-corrected *p*-values, and Cohen’s *d_z* effect sizes.

We evaluate two contrast categories: (1) *semantic degradation contrasts* comparing Risk-Aligned and Risk-Diluting conditions, and (2) *grounding-illusion contrasts* examining whether *Perceived Evidence Grounding* remains elevated under risk-diluting conditions relative to interpretation-oriented dimensions. To characterize this divergence, we define the *Grounding-Illusion Gap* (GIG) as:

$$\text{GIG} = \text{PEG} - \text{SQ}, \quad (1)$$

where PEG is Perceived Evidence Grounding and SQ denotes Semantic Quality measured by the average of Helpfulness, Appropriateness, and Reasoning Support. Larger values indicate explanations that continue to appear comparatively grounded despite semantically weakened risk interpretation.

LLM-as-a-Judge Evaluation Framework. To evaluate explanation robustness under both risk-aligned and risk-diluting conditions, we employ an *LLM-as-a-judge* framework (Zheng et al. 2023). Because evidence citation alone does not capture higher-level interpretive properties such as contextual appropriateness, semantic alignment, or explanatory plausibility, we complement the controlled generation setup with both automated and human evaluation.

Like the experimental design, explanations are evaluated across: (1) **semantic condition** (Risk-Aligned vs. Risk-Diluting), and (2) **evidence/style condition** (No-XAI, XAI-only, XAI + supportive style, and XAI + analytical framing).

The **evaluation rubric** is grounded in prior work on human-centered XAI, trustworthy explanations, and usable security (Miller 2019; Jacovi and Goldberg 2020; Sasse, Brostoff, and Weirich 2001; Turpin et al. 2023): (1) **Helpfulness:** Measures support for understanding scam risk (Miller 2019). (2) **Clarity:** Evaluates readability and interpretive accessibility (Miller 2019). (3) **Perceived Evidence Grounding:** Measures explicit reference to message-specific evidence, including grounded-looking but semantically misleading explanations (Jacovi and Goldberg 2020; Turpin et al. 2023). (4) **Appropriateness:** Evaluates contextual suitability and persona-informed framing (Sasse, Brostoff, and Weirich 2001). (5)

Reasoning Support: Evaluates whether explanations connect evidence to the detector’s semantic risk assessment rather than weakening or diluting it (Turpin et al. 2023).

Each explanation is evaluated using the structured 5-point Likert rubric shown in Table 8. To reduce self-evaluation bias, GPT-5.2 (Singh et al. 2025) is used as the explanation generator, whereas GPT-5.4 (Singh et al. 2025) and Gemini-2.5-Pro (Comanici et al. 2025) serve as independent judge models. All evaluations use deterministic decoding settings (temperature = 0). Judges receive the original message and the generated explanation text, but are not provided with detector-derived evidence tokens or condition labels. XAI-grounded conditions incorporate extracted evidence phrases, whereas risk-diluting conditions instruct the generator to preserve lexical grounding while semantically weakening or redirecting the intended risk interpretation.

Because judges do not observe detector-derived attribution signals directly, *Perceived Evidence Grounding* measures perceived message-specific grounding rather than attribution faithfulness. Consequently, evaluators may assign elevated grounding scores to explanations that appear evidence-based while weakening the detector’s intended risk assessment.

The controlled evaluation set consists of 18 scam-risk message instances balanced across communication channels, with 9 assigned to risk-aligned variants and 9 to risk-diluting variants. Using the same evaluation pool across human and LLM-based assessments enables controlled comparison of sensitivity to semantic misalignment under identical conditions.

Importantly, we treat LLM-based judges not only as evaluation tools, but also as objects of analysis. This setup tests whether automated evaluators distinguish evidence-grounded explanations from semantically aligned ones.

Human Evaluation of Grounding Illusions. To complement the LLM-as-a-judge evaluation, we conducted a human study examining how users assess explanation quality under varying grounding and semantic-alignment conditions. In particular, the study evaluates whether evidence-grounded explanations remain trustworthy even when their semantic interpretation weakens the detector’s intended risk assessment.

(a) Participants. We recruited adult participants fluent in English and familiar with everyday online communication environments such as messaging, email, and social media. Participants were recruited through university-affiliated mailing lists and online communication channels. All responses were collected anonymously, and participation was voluntary with informed consent obtained prior to the study.

Participants who completed fewer than the assigned evaluation items were excluded from analysis. The resulting human-evaluation dataset included $N = 28$ complete responses. The study protocol was approved by the institutional IRB.

(b) Controlled Evaluation Conditions. The evaluation pool consisted of 18 scam-risk message instances balanced across communication channels, with 9 assigned to risk-aligned explanation variants and 9 assigned to risk-diluting variants. Consistent with the overall experimental design, we manipulated: (1) **semantic condition** (Risk-Aligned vs. Risk-

Criterion	Evaluation Question and Scale
Helpfulness (1–5 Likert)	<i>Question:</i> To what extent does the explanation help the user understand why the message is risky? 1 = Not helpful; fails to explain the risk or is misleading. 2 = Slightly helpful; mentions risk but lacks clarity or relevance. 3 = Moderately helpful; provides some useful insight but remains generic or incomplete. 4 = Mostly helpful; clearly explains the risk with minor limitations. 5 = Very helpful; provides clear, specific, and informative reasoning that strongly aids understanding.
Clarity (1–5 Likert)	<i>Question:</i> How easy is the explanation to read and understand? 1 = Very unclear; difficult to understand due to poor structure or wording. 2 = Somewhat unclear; partially understandable but contains confusing phrasing. 3 = Moderately clear; generally understandable but lacks smoothness or precision. 4 = Mostly clear; easy to understand with minor issues. 5 = Very clear; well-structured, fluent, and immediately understandable.
Perceived Evidence Grounding (1–5 Likert)	<i>Question:</i> To what extent does the explanation reference specific cues in the message rather than relying on generic reasoning? 1 = Not grounded; explanation is generic or unsupported by the message. 2 = Weakly grounded; minimal reference to message-specific content. 3 = Moderately grounded; some relevant cues mentioned but incomplete. 4 = Mostly grounded; references key cues with minor omissions. 5 = Highly grounded; explicitly references multiple salient message-specific cues.
Appropriateness (1–5 Likert)	<i>Question:</i> How appropriate is the explanation in relation to the message content and context? 1 = Not appropriate; explanation is irrelevant or poorly matched to the context. 2 = Slightly appropriate; limited relevance to the message. 3 = Moderately appropriate; generally relevant but lacks precision. 4 = Mostly appropriate; well-aligned with minor issues. 5 = Highly appropriate; clearly relevant and well-matched to the message context.
Reasoning Support (1–5 Likert)	<i>Question:</i> To what extent does the explanation provide explicit reasoning connecting message cues to the detector’s intended risk assessment? 1 = No reasoning; only states the conclusion without explanation. 2 = Weak reasoning; minimal or vague justification provided. 3 = Moderate reasoning; some reasoning present but lacks depth. 4 = Strong reasoning; clear justification with minor gaps. 5 = Very strong reasoning; detailed and coherent reasoning connecting evidence to the intended interpretation.

Table 8: LLM-as-a-Judge evaluation criteria and scales. Perceived Evidence Grounding measures message-specific evidence citation, whereas Reasoning Support evaluates whether such evidence is connected to the detector’s semantic risk assessment.

Diluting), and (2) **evidence/style condition** (No-XAI, XAI-only, and grounded explanations with stylistic framing).

For grounded conditions, style-adaptive framing used supportive-contextual and analytical-concise prompting strategies. These framings were inspired by prior work on communication preferences and scam susceptibility, but were used only to modulate explanation style under controlled conditions rather than model participant personality traits.

Because stylistic-condition coverage was imbalanced across responses, grounded explanations with stylistic framing were aggregated into a single *XAI + Style* category for analysis. Consequently, style-related findings should be interpreted as exploratory rather than direct comparisons between supportive and analytical framing strategies. The primary human evaluation, therefore, compares three explanation settings: (1) *No XAI*, (2) *XAI Only*, and (3) *XAI + Style*.

Participants evaluated a counterbalanced subset of 9 explanation items, viewing only one explanation variant per item to reduce fatigue and repeated exposure effects. After excluding incomplete responses, the final dataset contained 252

participant-item evaluations. Condition assignment followed a Latin Square counterbalancing design to maintain balanced exposure across semantic conditions while minimizing repeated exposure to identical scam-risk content (Table 9).

Instance	Group A	Group B	Group C
Item 1	Ungrounded	Risk-Aligned	Risk-Diluting
Item 2	Risk-Aligned	Risk-Diluting	Ungrounded
Item 3	Risk-Diluting	Ungrounded	Risk-Aligned

Table 9: Illustration of the Latin Square counterbalancing design used in the human evaluation.

(c) Measures and Evaluation Procedure. Participants evaluated explanations using the same five dimensions employed in the LLM-as-a-judge framework: Helpfulness, Clarity, Perceived Evidence Grounding, Appropriateness, and Reasoning Support. Before evaluation, participants were provided with concise definitions and scoring guidelines for each dimension. Optional hint descriptions were also available during the study interface for additional clarification.

For descriptive statistics, we report means and standard deviations over participant-item ratings. For inferential testing, ratings were first averaged within participant $\times$ condition cells; paired contrasts were computed only for participants with observations in both compared conditions.

Each session lasted approximately 30 min. Participants evaluated only the original messages and generated explanations; detector-derived evidence tokens and condition labels were hidden. This setup enables controlled analysis of whether users infer trustworthiness and evidence grounding from semantically weakened yet evidence-grounded explanations under risk-diluting conditions. Consistent with the grounding-illusion framework, we compare human and LLM-based evaluations to examine whether evaluators distinguish lexical grounding from semantically aligned risk interpretation. The protocol reduced uncontrolled variation through standardized rubrics, hidden condition labels, and counterbalanced semantic-condition exposure.

Experimental Results and Analyses

LLM-as-a-Judge Evaluation Results. Tables 10 and 11 summarize explanation-quality evaluations from GPT-5.4 (Singh et al. 2025) and Gemini-2.5-Pro (Comanici et al. 2025) under risk-aligned and risk-diluting conditions.

(a) Ceiling Effects under Risk-Aligned Conditions. Under risk-aligned conditions, both judge models exhibit strong ceiling effects, assigning near-perfect scores across most evaluation dimensions regardless of grounding or stylistic configuration. GPT-5.4 produces scores near the upper bound across settings, whereas Gemini-2.5-Pro assigns almost uniformly perfect scores under semantically aligned conditions.

These findings suggest that when detector-derived evidence is interpreted consistently with the detector’s intended risk assessment, both evaluators assign highly favorable scores across dimensions. At the same time, the observed ceiling effects indicate limited score discrimination under semantically aligned conditions, particularly for Gemini-2.5-Pro.

(b) Grounding Illusions under Risk-Diluting Conditions. Grounding illusions are concerning in human-facing AI systems because explanations that preserve evidence-looking references while weakening semantic risk interpretation may induce false reassurance despite appearing trustworthy. Under risk-diluting conditions, grounded explanations exhibit substantial degradation in Helpfulness, Appropriateness, and Reasoning Support across both judge models despite continuing to reference detector-derived evidence.

For example, GPT-5.4’s risk-diluting *XAI-Only* condition drops sharply in Helpfulness (1.00 ± 0.00) and Appropriateness (1.22 ± 0.44), while retaining moderate Perceived Evidence Grounding (3.44 ± 1.01). Similarly, Gemini-2.5-Pro’s risk-diluting *Analytical Style* condition maintains elevated Perceived Evidence Grounding (4.62 ± 1.06) despite declines in Helpfulness and Appropriateness (both 1.50 ± 1.41).

This divergence reflects a substantial *grounding-illusion gap*, where explanations continue to appear grounded despite weakened risk interpretation. Across risk-diluting grounded

conditions, Perceived Evidence Grounding remains higher than interpretation-oriented dimensions, suggesting that automated evaluators may over-trust evidence-grounded explanations even when semantic alignment deteriorates.

Notably, even risk-diluting *No-XAI* explanations occasionally receive elevated Perceived Evidence Grounding scores, indicating that LLM judges may infer grounding from fluent message-specific references without explicit detector-derived evidence. These findings suggest that perceived grounding and attribution faithfulness are not interchangeable.

(c) Style-Adaptive Framing Effects. Style-adaptive framing introduced comparatively modest variation under risk-diluting conditions. Supportive framing did not substantially mitigate semantic degradation caused by risk-diluting reinterpretation, suggesting that stylistic reassurance alone is insufficient to preserve semantically aligned risk communication. From a trustworthy AI perspective, this highlights the importance of distinguishing supportive communication style from faithful semantic interpretation. Explanations may remain fluent and contextually appropriate while still weakening the detector’s intended risk assessment.

(d) Cross-Model Evaluation Consistency. Although GPT-5.4 and Gemini-2.5-Pro differed in score variance under risk-aligned conditions, both evaluators exhibited the same qualitative grounding-illusion pattern under risk-diluting settings. In particular, both models maintained elevated Perceived Evidence Grounding despite severe degradation in interpretation-oriented dimensions. This consistency suggests that grounding-illusion effects are not isolated to a single evaluator architecture and raises concerns about evaluation reliability in human-facing AI systems, where evidence-grounded explanations may appear trustworthy despite misleading interpretation. Broader evaluation across additional judge models is needed to assess the generality of these patterns.

Human Evaluation Findings. Table 12 summarizes human evaluation results across risk-aligned and risk-diluting conditions. Overall, human evaluators exhibit clear sensitivity to semantic degradation under risk-diluting conditions, assigning lower scores across most interpretation-oriented dimensions.

(a) Sensitivity to Semantic Misalignment. Under risk-aligned conditions, all explanation settings receive relatively high scores across Helpfulness, Clarity, Appropriateness, and Reasoning Support, generally within the 4.0–4.4 range. This suggests that semantically aligned explanations are perceived as coherent, informative, and contextually appropriate regardless of explicit grounding configuration. In contrast, risk-diluting conditions produce substantial degradation across interpretation-oriented dimensions. For example, risk-diluting *No-XAI* explanations show clear declines in Helpfulness (2.68 ± 1.29), Appropriateness (3.16 ± 1.34), and Reasoning Support (3.08 ± 1.21) relative to their risk-aligned counterparts. Similar degradation patterns appear across risk-diluting *XAI-Only* and *XAI + Persona* conditions.

These findings suggest that human evaluators remain broadly sensitive to semantically weakened or misleading explanations, particularly in dimensions associated with explanatory usefulness and contextual appropriateness.

Semantic Condition	Evidence/Style Condition	Helpfulness	Clarity	Perceived Evidence Grounding	Appropriateness	Reasoning Support
Risk-Aligned	No-XAI	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	4.78 $\pm$ 0.44	5.00 $\pm$ 0.00	4.78 $\pm$ 0.44
	XAI-Only	4.89 $\pm$ 0.33	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	4.89 $\pm$ 0.33
	Supportive Style	4.89 $\pm$ 0.33	4.89 $\pm$ 0.33	4.89 $\pm$ 0.33	5.00 $\pm$ 0.00	4.89 $\pm$ 0.33
	Analytical Style	4.78 $\pm$ 0.44	4.78 $\pm$ 0.44	4.78 $\pm$ 0.44	4.78 $\pm$ 0.44	4.78 $\pm$ 0.44
Risk-Diluting	No-XAI	3.89 $\pm$ 0.33	3.22 $\pm$ 0.44	4.89 $\pm$ 0.33	4.11 $\pm$ 0.33	4.44 $\pm$ 0.53
	XAI-Only	1.00 $\pm$ 0.00	4.11 $\pm$ 0.33	3.44 $\pm$ 1.01	1.22 $\pm$ 0.44	2.33 $\pm$ 0.71
	Supportive Style	1.00 $\pm$ 0.00	4.22 $\pm$ 0.44	3.33 $\pm$ 1.00	1.33 $\pm$ 0.50	2.11 $\pm$ 0.33
	Analytical Style	1.11 $\pm$ 0.33	4.11 $\pm$ 0.33	3.56 $\pm$ 1.01	1.56 $\pm$ 0.73	2.44 $\pm$ 0.73

Table 10: Comparison of explanation quality metrics across semantic and evidence/style conditions using GPT-5.4 as the judge model (Mean $\pm$ SD over 9 message-level explanations per condition).

Semantic Condition	Evidence/Style Condition	Helpfulness	Clarity	Perceived Evidence Grounding	Appropriateness	Reasoning Support
Risk-Aligned	No-XAI	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	4.88 $\pm$ 0.35	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00
	XAI-Only	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00
	Supportive Style	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00
	Analytical Style	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00	5.00 $\pm$ 0.00
Risk-Diluting	No-XAI	4.22 $\pm$ 0.97	1.67 $\pm$ 0.50	5.00 $\pm$ 0.00	3.78 $\pm$ 0.97	5.00 $\pm$ 0.00
	XAI-Only	1.50 $\pm$ 1.41	2.25 $\pm$ 0.89	4.50 $\pm$ 0.93	1.50 $\pm$ 1.41	1.50 $\pm$ 1.41
	Supportive Style	1.00 $\pm$ 0.00	3.50 $\pm$ 1.51	4.25 $\pm$ 1.39	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00
	Analytical Style	1.50 $\pm$ 1.41	3.00 $\pm$ 0.76	4.62 $\pm$ 1.06	1.50 $\pm$ 1.41	2.50 $\pm$ 2.07

Table 11: Comparison of explanation quality metrics across semantic and evidence/style conditions using Gemini-2.5-Pro as the judge model (Mean $\pm$ SD over 9 message-level explanations per condition).

(b) Grounding Illusions in Human Evaluation. Despite overall degradation under risk-diluting conditions, Perceived Evidence Grounding remained comparatively resilient across grounded settings. Notably, risk-diluting *XAI-Only* explanations achieved the highest grounding score among risk-diluting conditions (3.66 ± 1.02), despite lower Helpfulness, Appropriateness, and Reasoning Support. This pattern reflects a human-side *grounding-illusion gap*: explanations may appear grounded even when they weaken the detector’s intended risk interpretation. Moreover, semantically aligned *XAI-grounded* explanations did not consistently receive higher grounding scores than ungrounded explanations, suggesting that perceived grounding does not directly track attribution-based grounding or semantic faithfulness.

Human evaluators showed a smaller within-condition grounding-illusion gap than LLM judges. Humans used a narrower scoring range and assigned more moderate scores under risk-diluting conditions, whereas LLM judges often assigned near-minimum scores under semantic misalignment. Thus, automated evaluators appeared more extreme, while human grounding perceptions remained comparatively resilient despite weaker reasoning support.

(c) Statistical Evidence for the Grounding-Illusion Framework. Table 13 reports selected paired comparisons from the human evaluation. Risk-diluting *XAI-Only* explanations show significant degradation in Helpfulness and Appropriateness, whereas Perceived Evidence Grounding exhibits smaller and statistically non-significant degradation after Holm correc-

tion. Moreover, under risk-diluting *XAI-Only* conditions, Perceived Evidence Grounding remains significantly higher than Reasoning Support, consistent with the proposed grounding-illusion gap. Note that descriptive statistics use participant-item ratings, whereas paired tests use participant-condition aggregates; therefore, marginal mean differences may differ from the reported contrasts.

(d) Persona-Adaptive Framing Effects. Persona-adaptive framing introduces comparatively modest variation in human evaluation outcomes. Across both risk-aligned and risk-diluting settings, the *XAI + Persona* condition does not consistently outperform the corresponding *XAI-Only* condition across interpretation-oriented metrics. Under risk-diluting conditions, it often received slightly lower Helpfulness and Reasoning Support scores, suggesting that stylistic framing alone cannot overcome semantic degradation when detector-derived evidence is misleadingly interpreted.

These findings suggest that contextually appropriate, reassuring explanations may weaken risk assessment while perceived grounding remains comparatively resilient.

Discussion: Trustworthy Grounding and Semantic Alignment

(a) Evidence Citation Is Not Semantic Alignment. The primary insight of this work is the empirical characterization of the *grounding illusion*. Under risk-diluting conditions, explanations preserve grounded-looking evidence references

Semantic Condition	Evidence/Style Condition	Helpfulness	Clarity	Perceived Evidence Grounding	Appropriateness	Reasoning Support
Risk-Aligned (RA)	No-XAI	4.16 $\pm$ 0.93	4.42 $\pm$ 0.84	3.96 $\pm$ 0.85	4.13 $\pm$ 0.92	4.09 $\pm$ 0.82
	XAI-Only	4.27 $\pm$ 0.80	4.27 $\pm$ 0.77	3.89 $\pm$ 0.99	4.14 $\pm$ 0.86	3.92 $\pm$ 0.89
	XAI + Persona	4.07 $\pm$ 0.89	4.24 $\pm$ 0.69	3.74 $\pm$ 0.94	4.38 $\pm$ 0.62	4.05 $\pm$ 0.82
Risk-Diluting (RD)	No-XAI	2.68 $\pm$ 1.29	3.24 $\pm$ 1.09	3.03 $\pm$ 1.12	3.16 $\pm$ 1.34	3.08 $\pm$ 1.21
	XAI-Only	3.00 $\pm$ 1.41	3.50 $\pm$ 0.97	3.66 $\pm$ 1.02	3.22 $\pm$ 1.27	3.14 $\pm$ 1.05
	XAI + Persona	2.71 $\pm$ 1.27	3.49 $\pm$ 1.05	3.22 $\pm$ 1.04	3.17 $\pm$ 1.12	2.90 $\pm$ 1.07

Table 12: Mean $\pm$ SD over participant-item ratings; $N = 28$ participants and 252 total participant-item evaluations.

Contrast	Metric	Mean Diff.	95% CI	Holm p	d_z
RA vs. RD	Helpfulness	-1.41	[-2.10, -0.72]	.0067	-0.91
XAI-Only	Perceived Evidence Grounding	-0.39	[-0.94, 0.16]	.638	-0.31
	Appropriateness	1.08	[-1.58, -0.57]	.0047	-0.95
RD XAI-Only	Perceived Evidence Grounding – Reasoning Support	0.50	[0.29, 0.71]	< .001	0.96

Table 13: Selected paired comparisons from the human evaluation. RA and RD denote Risk-Aligned and Risk-Diluting conditions, respectively.

while weakening the detector’s semantic risk assessment. Across both automated and human evaluations, risk-diluting grounded explanations retain comparatively elevated Perceived Evidence Grounding scores despite degradation in interpretation-oriented dimensions such as Helpfulness, Appropriateness, and Reasoning Support. These findings suggest that surface-level evidence citation alone is insufficient for trustworthy explanation evaluation.

(b) Perceived Grounding versus Attribution Faithfulness. Perceived Evidence Grounding does not necessarily correspond to attribution-based grounding or semantic faithfulness. Even risk-diluting *No-XAI* explanations occasionally receive elevated grounding scores, suggesting that both humans and LLM judges may infer grounding from fluent message-specific references alone. At the same time, LLM-based judges exhibit larger grounding-illusion gaps under risk-diluting conditions, whereas human evaluators show a smaller within-condition divergence between perceived grounding and semantic quality.

(c) Evaluation Reliability and Trust Calibration. The grounding-illusion gap highlights an important reliability challenge for human-facing AI systems. Automated evaluators may remain overly sensitive to grounded-looking evidence references even when semantic interpretation deteriorates, potentially overestimating explanation trustworthiness. For security-critical systems, lexical grounding alone should not be treated as sufficient evidence of trustworthy explanation behavior. Grounded-looking but semantically misleading

Evaluator	Perceived Evidence Grounding	Semantic Average	Illusion Gap
Human	3.66	3.12	0.54
GPT-5.4	3.44	1.52	1.92
Gemini-2.5-Pro	4.50	1.50	3.00

Table 14: Grounding-illusion gaps across evaluators under the Risk-Diluting XAI-Only condition.

explanations may distort user trust calibration by encouraging over-reliance on AI-assisted security assessments.

Accordingly, future explanation systems may require auditing not only for evidence citation fidelity but also for semantic risk preservation and trust calibration. More broadly, our findings suggest that lexical grounding, semantic alignment, and attribution faithfulness should be evaluated as distinct properties rather than interchangeable quality proxies.

These concerns may be especially consequential for digitally vulnerable populations, including users with limited cybersecurity literacy or familiarity with AI-assisted decision systems, who may rely more heavily on AI-generated explanations during security-related decision making.

Conclusions and Future Work

This work shows that evidence citation alone does not ensure semantic risk alignment in AI-assisted scam-risk explanations. Across human and automated evaluations, risk-diluting explanations retained comparatively high Perceived Evidence Grounding despite weaker interpretation-oriented quality, revealing grounding-illusion effects. Trustworthy evaluation should therefore assess lexical grounding, semantic alignment, and attribution faithfulness as distinct properties.

Our evaluation was limited to 18 scam-risk messages and a relatively small university-affiliated participant sample, constraining generalizability across messages, populations, and cybersecurity-literacy levels. Risk-diluting explanations were also generated under controlled rather than deployed conditions. Future work should examine grounding illusions across broader populations, real-world prompt-injection and conversational-assistance settings, detector uncertainty, trust calibration, and security decision making.

Ethics Statement

This study involves human evaluation and publicly available scam-risk datasets spanning email, SMS, and social media. All human-study responses were collected anonymously, and participation was voluntary. Participants provided informed consent prior to participation, and the study protocol was approved by the institutional IRB.

The scam-risk corpora used in this study were obtained from publicly available research datasets. Example messages included in the paper were manually reviewed to avoid unnecessary disclosure of sensitive or identifying information.

Because VEXA includes prompts for generating risk-diluting explanations, releasing full adversarial prompts or raw scam messages could create dual-use risks. Accordingly, we plan to release only sanitized examples, aggregate evaluation data, and controlled prompt templates that prevent direct reuse for scam facilitation. Public artifacts will remove URLs, phone numbers, account identifiers, and other sensitive information.

More broadly, this work aims to improve the robustness of AI-assisted security explanations by highlighting the limitations of relying solely on evidence citation without verifying semantic alignment and risk interpretation.

We avoid interpreting style-adaptive framing conditions as normative or clinical claims about individual users or psychological vulnerability. These stylized framings are used solely to examine how explanation style influences evaluation robustness under controlled semantic conditions. We do not interpret these stylized framings as direct indicators of real-world psychological vulnerability or user personality.

Acknowledgements

This research was supported by the Commonwealth Cyber Initiative (CCI) Southwest Virginia (SWVA) Innovation: Ideation to Commercialization program and the Cyber Innovation Scholars Program.

References

- Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; and Koyama, M. 2019. Optuna: A Next-Generation Hyperparameter Optimization Framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2623–2631.
- Almeida, T. A.; and Hidalgo, J. M. G. 2011. SMS Spam Collection Data Set—UCI Machine Learning Repository.
- Anawar, S.; Kunasegaran, D. L.; Mas’ud, M. Z.; Zakaria, N. A.; et al. 2019. Analysis of Phishing Susceptibility in a Workplace: A Big-Five Personality Perspective. *Journal of Engineering Science and Technology*, 14(5): 2865–2882.
- Chen, T.; and Kan, M.-Y. 2013. Creating a Live, Public Short Message Service Corpus: The NUS SMS Corpus. *Language Resources and Evaluation*, 47(2): 299–335.
- Cho, J.-H.; Çam, H.; and Oltramari, A. 2016. Effect of Personality Traits on Trust and Risk to Phishing Vulnerability: Modeling and Analysis. *2016 IEEE CogSIMA*.
- Clark, K.; Luong, M.-T.; Le, Q. V.; and Manning, C. D. 2020. ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. *arXiv preprint arXiv:2003.10555*.
- Comanici, G.; Bieber, E.; Schaekermann, M.; Pasupat, I.; Sachdeva, N.; Dhillon, I.; Blistein, M.; Ram, O.; Zhang, D.; Rosen, E.; et al. 2025. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. *arXiv preprint arXiv:2507.06261*.
- Egelman, S.; Cranor, L. F.; and Hong, J. 2008. You’ve Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1065–1074.
- Federal Trade Commission. 2025. New FTC Data Show Big Jump in Reported Losses to Fraud: \$12.5 Billion in 2024.
- Ge, Y.; Lu, L.; Cui, X.; Chen, Z.; and Qu, W. 2021. How Personal Characteristics Impact Phishing Susceptibility: The Mediating Role of Mail Processing. *Applied Ergonomics*, 97: 103526.
- Greco, F.; Desolda, G.; Esposito, A.; Carelli, A.; et al. 2024. David versus Goliath: Can Machine Learning Detect LLM-Generated Text? A Case Study in the Detection of Phishing Emails. In *Proceedings of ITASEC 2024*, volume 3731.
- Gressel, G.; Pankajakshan, R.; and Mirsky, Y. 2024. Discussion Paper: Exploiting LLMs for Scam Automation: A Looming Threat. In *Proceedings of the 3rd ACM Workshop on the Security Implications of Deepfakes and Cheapfakes*, 20–24.
- He, P.; Gao, J.; and Chen, W. 2021. DeBERTaV3: Improving DeBERTa Using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing. *arXiv preprint arXiv:2111.09543*.
- Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; et al. 2025. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Transactions on Information Systems*, 43(2): 1–55.
- Jacovi, A.; and Goldberg, Y. 2020. Towards Faithfully Interpretable NLP Systems: How Should We Define and Evaluate Faithfulness? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4198–4205.
- Kaggle. 2020. UtkML’s Twitter Spam Detection Competition. Accessed: 2025-11-03.
- Lan, Z.; Chen, M.; Goodman, S.; Gimpel, K.; Sharma, P.; and Soricut, R. 2019. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. *CoRR*, abs/1909.11942.
- Lee, K.; Eoff, B.; and Caverlee, J. 2011. Seven Months with the Devils: A Long-Term Study of Content Polluters on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 5, 185–192.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.

- Lundberg, S. M.; and Lee, S.-I. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, 4768–4777. Long Beach, CA, USA: Curran Associates Inc.
- Metsis, V.; Androutsopoulos, I.; and Paliouras, G. 2006. Spam Filtering with Naive Bayes—Which Naive Bayes? In *CEAS*, volume 17, 28–69. Mountain View, CA.
- Miller, T. 2019. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artificial Intelligence*, 267: 1–38.
- Opara, C.; Modesti, P.; and Golightly, L. 2025. Evaluating Spam Filters and Stylometric Detection of AI-Generated Phishing Emails. *Expert Systems with Applications*, 276: 127044.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. “Why Should I Trust You?” Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.
- Sakkis, G.; Androutsopoulos, I.; Paliouras, G.; Karkaletsis, V.; Spyropoulos, C. D.; and Stamatopoulos, P. 2003. A Memory-Based Approach to Anti-Spam Filtering for Mailing Lists. *Information Retrieval*, 6(1): 49–73.
- Salman, M.; Ikram, M.; and Kaafar, M. A. 2024. Investigating Evasive Techniques in SMS Spam Filtering: A Comparative Analysis of Machine Learning Models. *IEEE Access*, 12: 24306–24324.
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2019. Distil-BERT: A Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. *arXiv*, abs/1910.01108.
- Sarno, D. M.; Harris, M. W.; and Black, J. 2023. Which Phish Is Captured in the Net? Understanding Phishing Susceptibility and Individual Differences. *Applied Cognitive Psychology*.
- Sasse, M. A.; Brostoff, S.; and Weirich, D. 2001. Transforming the “Weakest Link”: A Human–Computer Interaction Approach to Usable and Effective Security. *BT Technology Journal*, 19(3): 122–131.
- Singh, A.; Fry, A.; Perelman, A.; Tart, A.; Ganesh, A.; El-Kishky, A.; McLaughlin, A.; Low, A.; Ostrow, A.; Ananthram, A.; et al. 2025. OpenAI GPT-5 System Card. *arXiv preprint arXiv:2601.03267*.
- Tang, S.; Mi, X.; Li, Y.; Wang, X.; and Chen, K. 2022. Clues in Tweets: Twitter-Guided Discovery and Analysis of SMS Spam. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2751–2764.
- Tornblad, M. K.; Jones, K. S.; Namin, A. S.; and Choi, J. 2021. Characteristics That Predict Phishing Susceptibility: A Review. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 65. SAGE Publications.
- Turpin, M.; Michael, J.; Perez, E.; and Bowman, S. R. 2023. Language Models Don’t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting. In *Proceedings of the 37th NeurIPS*. Curran Associates Inc.
- Uddin, M. A.; Mahiuddin, M.; and Chowdhury, A. A. 2024. Explainable Email Spam Detection: A Transformer-Based Language Modeling Approach. In *2024 ICCIT*.
- Wang, J.; Liang, Y.; Meng, F.; Sun, Z.; Shi, H.; Li, Z.; Xu, J.; Qu, J.; and Zhou, J. 2023. Is ChatGPT a Good NLG Evaluator? A Preliminary Study. In *Proceedings of the 4th New Frontiers in Summarization Workshop*, 1–11.
- Wang, Y.; Yuan, J.; Chuang, Y.-N.; Wang, Z.; Liu, Y.; Cusick, M.; Kulkarni, P.; Ji, Z.; Ibrahim, Y.; and Hu, X. 2025. DHP Benchmark: Are LLMs Good NLG Evaluators? In *Findings of the Association for Computational Linguistics: NAACL 2025*, 8079–8094.
- Wataoka, K.; Takahashi, T.; and Ri, R. 2024. Self-Preference Bias in LLM-as-a-Judge. In *NeurIPS Safe Generative AI Workshop 2024*.
- Ye, J.; Wang, Y.; Huang, Y.; Chen, D.; Zhang, Q.; Moniz, N.; Gao, T.; Geyer, W.; Huang, C.; Chen, P.-Y.; et al. 2025. Justice or Prejudice? Quantifying Biases in LLM-as-a-Judge. In *International Conference on Learning Representations (ICLR)*, volume 2025, 102351–102390.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. *Advances in Neural Information Processing Systems*, 36: 46595–46623.