ACIG
APPLIED CYBERSECURITY & INTERNET GOVERNANCE

NASK

From ATT&CK Technique to Remediation Priority: Bayesian Estimation of Weakness-Class Likelihood for Operational Cyber Defence

Emad Sherif | Faculty of Technology, Arts and Culture, De Montfort University, Leicester, UK | ORCID: 0000-0002-8450-338X

Abstract

The existing vulnerability prioritisation methods emphasise either technical severity, such as the Common Vulnerability Scoring System, or exploitation likelihood, such as the Exploit Prediction Scoring System, but neither helps defenders decide which weakness classes to remediate first after observing a specific MITRE ATT&CK technique. We propose a Bayesian co-occurrence model that estimates the conditional probability of Common Weakness Enumeration (CWE) classes, given an ATT&CK technique by integrating three publicly available datasets: the known exploited vulnerabilities (KEV) catalog, the National Vulnerability Database, and the European Repository of Cyber Incidents. A KEV-derived exploitation-frequency distribution is used as the prior, while incident-level technique, CWE, co-occurrences provide the conditioning signal. Using a temporal holdout split, we evaluate technique-conditioned CWE rankings against an unconditioned marginal baseline with Normalized Discounted Cumulative Gain (NDCG). Across the four evaluated techniques, the conditioned model improves mean NDCG@10 by 693% over the baseline. The findings indicate that technique-conditioned Bayesian estimation

Received: 31.03.2026

Accepted: 15.06.2026

Published: 24.07.2026

Cite this article as:
E. Sherif "From ATT&CK technique to remediation priority: Bayesian estimation of weakness-class likelihood for operational cyber defence," ACIG, vol. 5, no. 2, 2026, doi: 10.60097/ACIG/224570

Corresponding author:
Emad Sherif, Faculty of Technology, Arts and Culture, De Montfort University, Leicester, UK;
E-mail: emad.o.sherif@dmu.ac.uk
0000-0002-8450-338X

can generate interpretable and operationally useful CWE remediation priorities.

Keywords

Bayesian inference, cyber risk management, vulnerability prioritisation, MITRE ATT&CK, CISA KEV

1. Introduction

Effective cyber risk management requires organisations to translate threat observations into prioritised remediation actions. When an incident responder observes adversary behaviour consistent with a MITRE ATT&CK technique, such as Bypass User Account Control (BUAC T1548.002) or network denial of service (DoS T1498), an immediate operational question arises: which classes of software weakness should be remediated first to mitigate the observed behaviour? Although the ATT&CK framework provides a structured description of adversary tradecraft, the existing vulnerability prioritisation approaches do not directly answer this technique-conditioned remediation question [1].

The dominant operational framework for vulnerability scoring, the Common Vulnerability Scoring System (CVSS), was designed to communicate technical severity rather than estimate risk in a specific threat context. Official guidance from the National Vulnerability Database (NVD) states that CVSS is a measure of severity, not risk [2, 3]. Empirical work has further shown that CVSS alone has limited value for prioritising vulnerabilities by exploitation likelihood, particularly when compared with exploit-based or data-driven approaches [4]. By contrast, the Exploit Prediction Scoring System (EPSS) estimates the probability of exploitation in the wild and offers a more informative basis for prioritisation [5]. However, both CVSS- and EPSS-based approaches remain technique-agnostic wherein they rank vulnerabilities without conditioning on adversary behaviour observed in incidents [6, 7].

Moreover, prior research has sought to link vulnerabilities to ATT&CK techniques, using both expert-defined mappings and machine-learning (ML) approaches [8, 9]. These efforts provide important context on how vulnerabilities may enable adversary behaviour, but they largely stop at linkage or classification and do not estimate remediation priorities at the level of Common Weakness Enumeration (CWE) classes. At the same time, exploitation data shows that confirmed exploitation is concentrated in a relatively small number of recurring

weakness classes, suggesting that CWE provides a useful abstraction for operational prioritisation [10]. A marginal exploitation distribution *P (CWE)* is informative, but it does not answer the more relevant conditional question that if a specific ATT&CK technique has been observed, how should remediation priorities shift? To address this gap, we propose a Bayesian co-occurrence model that estimates *P (CWE | ATT&CK technique)* from incident-level technique–CWE associations derived from the European Repository of Cyber Incidents [11], regularised by an exploitation prior derived from the Cybersecurity and Infrastructure Security Agency's (CISA) known exploited vulnerabilities (KEV) catalog [12] and NVD [13]. The model yields a ranked list of weakness classes tailored to the observed technique, thereby providing a direct operational output for vulnerability managers and security operations centres.

Building on this gap, this study addresses three research questions as outlined below:

RQ1: Can the conditional probability distribution *P (CWE | ATT&CK technique)* be reliably estimated by integrating publicly available cyber incident, vulnerability, and exploitation datasets without requiring proprietary threat intelligence or manual expert elicitation?

RQ2: Does conditioning CWE rankings on an observed ATT&CK technique produce prioritisation outputs that measurably outperform the unconditioned marginal exploitation baseline *P (CWE)* under a temporal holdout evaluation design?

RQ3: How sensitive are the resulting CWE rankings and probability estimates to the choice of smoothing concentration parameter α, and which value provides the most appropriate balance between technique-specific signal and prior regularisation for this data regime?

Therefore, the main contributions are as follows: we propose a Bayesian co-occurrence model for estimating *P (CWE | ATT&CK technique)* using an exploitation-grounded prior derived from the KEV catalog; develop an integrated data pipeline combining 225,144 Common Vulnerabilities and Exposures (CVEs) records, 1551 KEV entries, and 1432 European Repository of Cyber (EuRepoC) incidents–techniques–CWE triples across 24 ATT&CK techniques; and demonstrate, under a temporal holdout design, that technique-conditioned ranking substantially outperforms the marginal baseline on the evaluated techniques.

The remainder of this paper is organised as follows. Section 2 reviews related work, Section 3 presents the methodology, Section 4 reports the results, Section 5 discusses implications and limitations, and Section 6 concludes with directions for future research.

2. Related Work

2.1. Vulnerability Prioritisation and CVSS Limitations

The CVSS remains the most widely used standardised framework for communicating vulnerability severity [7]. Official documentation from both Forum of Incident Response and Security Teams (FIRST) and NVD makes clear, however, that CVSS is intended primarily as a measure of technical severity, not as a direct measure of risk. Its base metrics capture intrinsic characteristics, such as attack vector, attack complexity, privileges required, user interaction, and impacts on confidentiality, integrity, and availability [2, 3].

Prior works have questioned the value of CVSS as a standalone prioritisation mechanism for real-world exploitation. Allodi and Massacci [14] showed that prioritising remediation solely on the basis of high CVSS scores provides limited discrimination for exploited vulnerabilities and can approximate random patch selection in practice. Howland [7] extended this critique by arguing that CVSS is deeply embedded in vulnerability management workflows despite weak alignment with observed exploitation patterns. Similarly, Jacobs et al. [15] demonstrated that remediation strategies grounded in exploit prediction outperform severity-based heuristics when balancing efficiency and coverage, reinforcing the argument that CVSS alone is a weak basis for prioritising vulnerabilities by likely exploitation.

2.2. Exploitation Prediction and EPSS

The EPSS was developed to address the gap between severity-based scoring and the practical need to estimate the likelihood of exploitation [5]. EPSS is a data-driven machine-learning model that estimates the probability that a published CVEs entry will be exploited in the wild within the next 30 days, and it publishes daily probability scores and percentiles for CVEs through openly available data feeds [6]. The EPSS work positioned this as a more operationally useful basis for remediation prioritisation than severity alone, and subsequent iterations have continued to improve the model's predictive accuracy [5].

Recent research shows that EPSS outperformed CVSS-only classifiers on the same KEV-labelled dataset, and a composite Key Risk Indicator (KRI) that combined EPSS with CVSS severity and CWE prevalence improved impact-weighted remediation value further [4]. Nevertheless, both EPSS and derivative risk indicators remain technique-agnostic, wherein they estimate exploitation likelihood at the level of individual CVEs but do not condition prioritisation on the adversary behaviour observed in an incident. As a result, they do not answer the question that motivates the present study that given an observed ATT&CK technique, which weakness classes should be remediated first?

2.3. MITRE ATT&CK and Vulnerability Mapping

The MITRE ATT&CK framework has become the *de facto* standard for describing adversary behaviour, providing a structured knowledge base of tactics, techniques, and sub-techniques grounded in real-world observations [16]. A related body of work has sought to connect vulnerability data to ATT&CK in order to contextualise how specific vulnerabilities may enable adversary actions. Kuppa et al. [8] and Grigorescu et al. [9] have explored automated CVE-to-ATT&CK linkage using enriched metadata and machine-learning models. A closely related strand of work concerns the automated extraction of ATT&CK techniques from unstructured threat intelligence. The threat report ATT&CK mapper (TRAM) project was developed to reduce the cost and inconsistency of mapping threat reports to ATT&CK, and later iterations incorporated fine-tuned large language models to improve both speed and mapping quality [17]. The present study is positioned downstream of that literature wherein rather than asking how to map vulnerabilities or reports to ATT&CK, it assumes that relevant techniques have already been identified and instead asks how those observed techniques should alter CWE remediation priorities.

2.4. CWE Exploitation Landscape

The CWE framework provides a standard taxonomy of software and hardware weaknesses [18], and the annual CWE top 25 continues to show that a relatively small set of recurring weakness classes accounts for a large share of serious software risk [10]. The CWE programme has also recently introduced dedicated views, such as the top 10 KEV weaknesses, reinforcing the practical value of analysing exploitation at the weakness-class level rather than only at the level of individual CVEs [19].

A recent work done by Simonetto et al. [20] shows that large language model-assisted key-term extraction can improve transformer-based CVE-to-CWE mapping performance by up to 8.88% F1, highlighting ongoing progress in the automatic identification of weakness classes from vulnerability text. However, the literature mainly focuses on assigning CWE labels or characterising the marginal exploitation landscape; it does not estimate technique-conditioned probabilities of the form *P (CWE | ATT&CK technique)*, which is the gap addressed by the present work.

2.5. Probabilistic and Graph-Based ATT&CK–Vulnerability Linkage

A body of work has addressed the probabilistic and structural linkage between adversary behaviour and vulnerability data, though from perspectives that differ from the present study.

Graph-based approaches have sought to represent ATT&CK techniques, CVEs, and CWEs as nodes in a knowledge graph, with edges encoding structural or semantic relationships. Hemberg et al. [21] introduced BRON, a heterogeneous knowledge graph that links ATT&CK tactics, techniques, and procedures (TTPs) to CVEs and CWEs through formally defined relationships derived from MITRE's published mappings. BRON enables reachability queries over the attack-vulnerability graph, for example, identifying which CVEs are structurally connected to a given ATT&CK technique. Similarly, Li et al. [22] proposed AttacKG, a technique knowledge graph constructed from cyber threat intelligence reports using natural language processing (NLP), which captures procedural relationships between adversary actions and affected system components. These graph-based approaches are valuable for structural exploration of the ATT&CK-vulnerability space but do not estimate probability distributions over weakness classes conditioned on observed technique activity.

Probabilistic attack graph models represent an earlier and related tradition. Sheyner et al. [23] introduced automated generation of attack graphs with associated probability annotations, enabling the computation of path probabilities through chains of exploitable vulnerabilities. Subsequent work formalised the computation of attack likelihood in enterprise networks under models, such as k-zero-day safety [24]. While these models quantify exploitation probability across multi-step attack paths, they operate at the level of network topology and individual CVEs and do not address the technique-conditioned weakness-class prioritisation problem.

Kim et al. [25] proposed BAN, a Bayesian network model for predicting advanced persistent threat (APT) attack progression sequences using the ATT&CK framework. BAN models conditional dependencies between ATT&CK techniques across attack stages, making it a probabilistic model that involves ATT&CK conditioning. However, BAN addresses the question of which technique is likely to follow a previously observed technique, a next-step prediction problem, rather than asking which weakness classes are most likely to underlie an observed technique. The output of BAN is a probability distribution over ATT&CK techniques; the output of the present model is a probability distribution over CWE weakness classes intended to guide remediation prioritisation.

Taken together, these prior approaches either model structural linkage without probability estimation (BRON and AttacKG), compute path-level exploit probabilities across network topology (probabilistic attack graphs), or model technique-to-technique transition probabilities (BAN). The present work is, to our knowledge, the first to estimate the conditional probability as a remediation prior, using incident-level co-occurrence data regularised by an empirical exploitation baseline. This positions the proposed model as a complement to the above approaches rather than a replacement for them, wherein BRON or AttacKG could, in principle, provide the structural linkage from which co-occurrence counts are derived, while the Bayesian model converts those counts into actionable ranked probability estimates for CWE remediation.

2.6. Bayesian Methods in Cyber Risk

Bayesian methods have long been recognised as valuable for cyber risk analysis because they provide a principled way to reason under uncertainty, combine sparse empirical data with expert judgement, and update beliefs as new evidence becomes available [26]. The authors advocate this broader quantitative perspective explicitly, including the use of Bayesian examples for cybersecurity problems with limited data.

Within the cybersecurity literature, Bayesian networks have been applied to model uncertainty in enterprise security analysis and near real-time attack reasoning [27], and they have also been used to extend formal risk frameworks, such as factor analysis of information risk (FAIR), with more flexible probabilistic structure and inference [28]. More recent work continues this trend by applying Bayesian statistical updating to dynamic cyber-threat quantification and prioritisation [29]. However, the existing literature primarily

addresses risk assessment, threat propagation, or exploit prediction; it does not, to our knowledge, address Bayesian smoothing for estimating conditional probabilities over structured security ontologies, such as ATT&CK techniques and CWE classes. Hence, the present work contributes specifically at that intersection by using Bayesian smoothing to estimate technique-conditioned weakness distributions for remediation prioritisation.

3. Methodology

The methodology comprises four stages: (1) data integration, (2) construction of a marginal exploitation prior, (3) estimation of technique-conditioned CWE probabilities using Bayesian smoothing, and (4) evaluation under a temporal holdout design. The data sources are the NVD, the KEV catalogue, and the EuRepoC, which, respectively, provide vulnerability metadata, confirmed exploitation evidence, and incident-level ATT&CK annotations.

3.1. Dataset Integration

We use the full NVD CVE dataset, comprising 225,144 CVE records with specific CWE annotations covering vulnerabilities disclosed between 1999 and 2025 [13]. Records labelled with non-specific weakness categories (NVD-CWE-Other and NVD-CWE-noinfo) are excluded, leaving 707 unique specific CWE types. For each retained record, we extract the CVE identifier, vulnerability description, CVSS base score, severity band, and mapped CWE identifier.

The KEV catalogue is used as the source of confirmed exploitation evidence. At the time of the experiment, the fetched KEV snapshot contained 1551 CVE entries [12]. Joining KEV to NVD by CVE identifier produced 1217 matched records (78.5%), with the remaining unmatched cases largely reflecting CVEs without specific NVD CWE annotations or entries added to KEV before NVD enrichment was complete. The matched $KEV \cap NVD$ set forms the exploitation corpus used to estimate the baseline prior over CWE classes.

The KEV-derived prior was constructed by joining KEV entries to NVD on CVE identifier and retaining only matched CVEs with specific CWE annotations; records labelled as NVD-CWE-Other or NVD-CWE-noinfo were excluded from the prior estimation.

EuRepoC provides structured, expert-reviewed cyber-incident data [11]. EuRepoC incidents were linked to CWE labels through a pre-processed mapping pipeline developed prior to the analysis.

In that pipeline, each incident record was associated with one or more CWE classes derived from the linked vulnerability information available in the local dataset. Where multiple CWE labels were associated with the same incident, all such labels were retained. The resulting dataset was then normalised to incident–technique–CWE triples and deduplicated at the level of (incident description, ATT&CK technique, CWE), so that repeated records referring to the same incident–technique–weakness combination contributed only one co-occurrence observation. The final dataset consists of 1432 unique records, spanning 24 ATT&CK techniques and 42 CWE classes between 2011 and 2025 [30].

3.2. Baseline Prior

The marginal exploitation-frequency prior *P (CWE)* is computed from the KEV-confirmed exploited CVEs matched to NVD CWE annotations:

$$P(CWE = c) = \frac{count(c\ in\ KEV \cap NVD)}{\Sigma_{c'} count(c'\ inKWEV \cap NVD)} \quad (1)$$

This prior directly replicates the exploitation-frequency baseline, providing methodological continuity and enabling direct comparison. The prior spans 128 CWE types; CWE-787 (Out-of-bounds Write) is the most prevalent at probability P = 0.1035 (126 KEV-confirmed CVEs), followed by CWE-78 (OS Command Injection, P = 0.0748) and CWE-416 (Use-After-Free, P = 0.0731). Table 1 presents the top 15 CWEs in the prior.

3.3. Bayesian Conditional Probability Model

The co-occurrence count *n (t, c)* is defined as the number of unique incident–technique–CWE triples in the training partition, where technique *t* is associated with CWE class *c*. The Bayesian posterior is computed as follows [25]:

$$P\ (CWE = c \mid Tech = t) \propto n\ (t, c) + \alpha\ P\ (CWE = c), \quad (2)$$

where α is a smoothing concentration parameter. The posterior is normalised over all CWE types in the vocabulary (union of the KEV prior and observed incident CWEs). This formulation encodes the following semantics: when technique-specific incident evidence is abundant, the posterior is dominated by *n (t, c)* and the prior has negligible influence; when evidence is sparse, the posterior degrades gracefully to the KEV exploitation baseline.

Table 1. Top 15 CWE types in the KEV exploitation prior P (CWE).

Rank	CWE	Weakness type	P (CWE)	KEV count
1	CWE-787	Out-of-bounds Write	0.1035	126
2	CWE-78	OS command Injection	0.0748	91
3	CWE-416	Use-after-free	0.0731	89
4	CWE-22	Path traversal	0.0600	73
5	CWE-20	Improper input Validation	0.0493	60
6	CWE-94	Code injection	0.0427	52
7	CWE-502	Deserialisation	0.0411	50
8	CWE-119	Buffer overflow	0.0362	44
9	CWE-843	Type confusion	0.0296	36
10	CWE-287	Improper authentication	0.0279	34
11	CWE-306	Missing authentication	0.0238	29
12	CWE-77	Shell injection	0.0230	28
13	CWE-79	Cross-site Scripting	0.0230	28
14	CWE-284	Improper access control	0.0205	25
15	CWE-89	SQL injection	0.0181	22

Setting $\alpha = 1.0$, assigns one total unit of prior mass distributed according to the KEV-derived prior, is the chosen value, justified by the sensitivity ablation in Section 4.4.

The prior $P(CWE = c)$ is specifically the exploitation-frequency distribution from $KEV \cap NVD$ rather than a uniform prior. This encodes the domain knowledge that, in the absence of technique-specific signal, a vulnerability manager should prioritise CWEs that are most frequently confirmed as exploited globally, a principled default that is empirically grounded rather than uninformative.

3.4. Evaluation Design

Incident records are sorted chronologically by incident date and split at the 80th percentile [8], yielding a training partition (incidents prior to 3 January 2024) and a test partition (incidents from 3 January 2024 onwards). This temporal design ensures that the model is evaluated on its ability to generalise to future incidents, the operationally relevant setting, rather than on randomly sampled records that may share temporal context with training data.

For each technique *t* in the test partition, the ground-truth relevance set is defined as the set of CWE classes observed in test incidents associated with that technique. We evaluate ranked predictions using three standard ranking metrics. *Precision@k* measures the proportion of relevant CWE classes among the top-*k* predictions and captures actionable precision under a fixed remediation budget [31]. Mean Reciprocal Rank (MRR) measures the reciprocal rank of the first correct prediction and captures how quickly the model surfaces at least one relevant weakness class [32]. Normalized Discounted Cumulative Gain (*NDCG@k*) evaluates ranking quality while penalising relevant classes placed lower in the ranking [31]. NDCG is appropriate when early ranked results matter most to the user, although in this design relevance is binary unless frequency-based gains are introduced. All metrics are computed at $k \in \{5,10,20\}$" and compared against the unconditioned baseline ranking induced by $P(CWE)$.

To quantify uncertainty in the technique-conditioned posterior estimates, we computed 95% equal-tailed credible intervals for representative posterior probabilities. Treating each technique-specific posterior as Dirichlet with parameters $a_{t,c} = n(t,c) + \alpha P(c)$, the marginal posterior for each individual class probability was evaluated as Beta($a_{t,c}$,$A_t - a_{t,c}$), where "A_t is the sum of the posterior parameters over the CWE vocabulary.

4. Results

4.1. Co-occurrence Matrix

The training partition yields 90 unique (technique, CWE) co-occurrence pairs across 21 techniques and 42 CWE types. The distribution is skewed, wherein T1548.002 (BUAC) contributes 455 co-occurrences with CWE-434, and T1548 (Abuse Elevation Control Mechanism) contributes 127 and 126 co-occurrences with CWE-290 and CWE-306, respectively. Most other techniques have single-digit or zero counts for most CWE types, reflecting the realistic sparsity of expert-annotated incident-CWE linkages. This sparsity is the primary motivation for the Bayesian smoothing approach, whereby the model must produce meaningful predictions even when direct co-occurrence evidence is limited [8].

4.2. Conditional Probability Distributions

The conditional probability displays *P (CWE | Technique)* for all 21 modelled techniques across the 20 most-discriminating CWE types. Three behavioural regimes are visible, each reflecting

a distinct relationship between evidence density and posterior concentration [20].

Four techniques produce near-complete probability concentration on a single CWE class. T1548.002 (BUAC) assigns $P = 0.998$ to CWE-434 (unrestricted file upload), representing a 62× uplift over the baseline probability of 0.016. This association is domain-correct wherein UAC bypass commonly exploits DLL hijacking through unsigned or loosely validated file loading, precisely the exploitation pathway described by CWE-434. T1204 (User Execution) concentrates on CWE-306 (Missing Authentication, $P = 0.947$), consistent with user-executed payloads exploiting unauthenticated application endpoints. T1491.002 assigns $P = 0.955$ to CWE-416 (Use-After-Free), and T1601 (Modify System Image) assigns $P = 0.875$ to CWE-494 (Download of Code without Integrity Check), the latter representing a particularly compelling result because CWE-494 has negligible baseline probability (effectively absent from the global KEV prior), making this prediction entirely technique-specific rather than prior-driven.

Techniques with genuine exploitation ambiguity correctly produce bimodal posteriors. T1548 (the parent technique of T1548.002) distributes approximately equally between CWE-290 (Spoofing, $P = 0.500$) and CWE-306 (Missing Authentication, $P = 0.496$), reflecting the two dominant sub-technique pathways for elevation control abuse. Similarly, T1087 (Account Discovery) splits between CWE-290 and CWE-451 at approximately $P = 0.482$ each. These bimodal distributions are the correct Bayesian response when training evidence equally supports two exploitation pathways and should not be interpreted as model failure.

T1498 (network DoS) distributes mass approximately equally across five CWE classes: CWE-78 ($P = 0.199$), CWE-416 ($P = 0.199$), CWE-94 ($P = 0.199$), CWE-77 ($P = 0.199$), and CWE-406 ($P = 0.199$). Critically, the baseline prior assigns CWE-787 (Out-of-bounds Write) the highest probability ($P = 0.103$), while the conditioned model correctly suppresses CWE-787 in favour of DoS-relevant weakness classes. This demonstrates the model actively overriding the global exploitation prior with technique-specific evidence, which is the central claim of this work.

4.3. Evaluation against Marginal Baseline

Table 2 presents full evaluation results across all four techniques with sufficient test observations. The temporal split at 3 January 2024 yields test partitions of 65 records for T1548, 196 for T1548.002, 1 for T1601, and 25 for T1498.

Table 2. Evaluation results: marginal baseline vs. technique-conditioned model (α = 1.0).

Technique	N	NDCG@5 base	NDCG@5 cond	NDCG@10 base	NDCG@10 cond	NDCG@20 base	NDCG@20 cond	MRR base	MRR cond
T1548	65	0.000	1.000	0.000	1.000	0.171	1.000	0.091	1.000
T1548.002	196	0.000	1.000	0.000	1.000	0.242	1.000	0.063	1.000
T1601	1	0.000	1.000	0.000	1.000	0.000	1.000	0.013	1.000
T1498	25	0.384	1.000	0.504	1.000	0.595	1.000	0.500	1.000
Mean	**-**	**0.096**	**1.000**	**0.126**	**1.000**	**0.252**	**1.000**	**0.167**	**1.000**

The conditioned model achieves *NDCG@10* = 1.000 for all four techniques. Mean *NDCG@10* improvement over the marginal baseline is +693% (from 0.126 to 1.000). Mean *Precision@5* increases from 0.100 to 0.450 (+350%), and mean *MRR* improves from 0.167 to 1.000 (+500%), meaning the model places the first correct CWE at rank 1 for every evaluated technique. T1498 is the most informative result wherein with a non-trivial baseline of *NDCG@10* = 0.504 (the baseline performs at chance-plus for this technique because DoS-relevant CWEs partially overlap with the global exploitation prior), the conditioned model still achieves perfect *NDCG@10* = 1.000 (+0.496), demonstrating improvement beyond a non-trivial baseline.

Posterior credible intervals clarify the distinction between well-supported and data-sparse techniques. For the highest-support case, T1548.002, the top-ranked class CWE-434 has posterior mean 0.9984 with a 95% credible interval of [0.9941, 0.99996], indicating very strong concentration. By contrast, T1498 produces a more diffused yet still informative posterior, with top-ranked CWE-78 assigned mean probability of 0.1989 and 95% credible interval of [0.1304, 0.2777]. For sparse techniques, uncertainty is substantially larger where T1601 assigns mean probability 0.8891 to CWE-494 but with a wider 95% credible interval of [0.6309, 0.9969], while T1119 assigns mean probability of 0.5306 to CWE-22 with interval [0.0329, 0.9815]. These results indicate that posterior concentration is high when co-occurrence evidence is abundant, but confidence in the precise probability estimate is necessarily lower for under-represented techniques.

4.4. Smoothing Parameter Sensitivity Ablation

Table 3 reports a three-point sensitivity analysis over $\alpha \in \{0.1, 0.5, 1.0, 5.0, 10.0\}$, showing the probability assigned to the

Table 3. Smoothing parameter ablation: P (top CWE | Technique) and distribution entropy *H* across α ∈ {0.1, 0.5, 1.0}. *NDCG*@10 = 1.000 for all techniques at all α values.

Technique	CWE	α = 0.1	α = 0.5	α = 1.0	α = 0.1	α = 0.5	α = 1.0
		P (CWE)	*P (CWE)*	*P (CWE)*	Entropy *H*	Entropy *H*	Entropy *H*
T1548.002	CWE-434	0.9998	0.9989	0.9978	0.003	0.013	0.024
T1548	CWE-290	0.502	0.501	0.500	0.698	0.713	0.730
T1498	CWE-78	0.200	0.199	0.199	1.618	1.647	1.678
T1601	CWE-494	0.986	0.933	0.875	0.127	0.497	0.850

top-ranked CWE class and the Shannon entropy of the corresponding conditional distribution for each evaluated technique [20]. Across the four techniques shown here, *NDCG@10* remains equal to 1.000 for all tested values of *α*, indicating that the rank ordering of CWE classes is stable within this range.

Three observations emerge from this ablation. First, ranking performance appears insensitive to α over the tested range, since *NDCG@10* remains unchanged for all four techniques. This suggests that the principal conclusions about relative CWE ordering do not depend strongly on the exact smoothing value within this interval.

Second, for data-rich techniques, such as T1548.002 (n = 455), the top-ranked CWE probability changes only slightly, from 0.9998 at α = 0.1 to 0.9978 at α = 1.0, with only a small increase in entropy. This pattern is consistent with the likelihood dominating the prior when co-occurrence evidence is abundant.

Third, for sparse techniques, such as T1601, the posterior is more sensitive to the smoothing parameter. The top-ranked CWE probability decreases from 0.986 at α = 0.1 to 0.875 at α = 1.0, while entropy rises from 0.127 to 0.850. This indicates that, under limited evidence, stronger smoothing produces more diffused and conservative posterior estimates even when the rank ordering remains unchanged. In practical terms, this means that *α* has relatively little effect on which CWE class is ranked first in the present data, but it does influence how confidently ranking is expressed.

On that basis, α = 1.0 appears to be a reasonable regularisation choice for the present study. Under the model formulation used here, α = 1.0 assigns one total unit of prior mass distributed according to the KEV-derived exploitation prior rather than implementing classical add-one smoothing. This setting provides moderate

regularisation without overwhelming the technique-specific co-occurrence signal.

5. Discussion

The practical value of the conditioned model is most evident when the observed ATT&CK technique materially changes the ranking of weakness classes relative to the marginal exploitation prior. For example, when the model is conditioned on T1548.002 (BUAC), probability mass shifts away from globally dominant weakness classes and towards CWE-434, which is strongly represented in the technique-specific training data. In operational terms, this type of shift could help vulnerability managers narrow their initial remediation focus to a smaller subset of weakness categories that appear more relevant to the observed behaviour than the unconditioned baseline would suggest. At the same time, this output should be interpreted as a decision-support signal rather than as a deterministic prescription, whereby it reflects empirical co-occurrence patterns in the available data and does not, by itself, establish a causal mapping between the observed technique and a single underlying weakness class.

A different but equally important pattern appears for T1498 (network DoS), where the conditioned model distributes probability across several CWE classes rather than concentrating on one. In practice, such a diffused posterior may still be useful, because it indicates that the available evidence does not strongly support a single remediation pathway. In this case, the conditioned distribution suppresses some globally common weakness classes and reallocates probability towards a set of classes that are more closely associated with the observed technique in the incident data. For a defender, that type of result may support a broader, technique-aware remediation strategy rather than an overly narrow focus on the globally most exploited weaknesses. However, because the evidence for some techniques remains sparse, such diffused posteriors may reflect either genuine ambiguity or limited data density.

More broadly, the findings suggest that technique-conditioned CWE ranking may complement the existing vulnerability prioritisation methods rather than replace them. Prior work on EPSS, CVSS, and composite KRIs has focused primarily on estimating exploitation likelihood or remediation value at the level of individual CVEs [4]. The present work addresses a different question that given that a particular ATT&CK technique has been observed, which classes of weakness appear most relevant to investigate or remediate first? In

that sense, the contribution of this work is not a direct substitute for CVE-level exploit prediction but an additional behavioural layer that may help contextualise remediation decisions once an incident has been characterised in ATT&CK terms. The observed improvements over the marginal baseline therefore suggest that conditioning on adversary behaviour can add information beyond an unconditioned exploitation prior, at least for the subset of techniques evaluated here.

Some limitations constrain the interpretation of these results. The most important is the sparsity of the incident–technique–CWE matrix. After de-duplication, the EuRepoC-derived dataset contains 1432 unique triples across 24 techniques, but most technique–CWE pairs are represented by very few observations. Quantitative evaluation is available for only four techniques, and support across those techniques is highly uneven. As a result, the empirical findings are best understood as evidence of feasibility and initial promise rather than as a comprehensive assessment across the full ATT&CK technique space. Increasing linkage density through additional curated cyber threat intelligence sources would likely improve both model stability and evaluation coverage.

This sparsity is also uneven across techniques. In the present dataset, technique-level support ranges from 1 to 623 observations, with a median of 8 and a pronounced long-tail distribution. The top three techniques account for approximately 74.7% of all training observations, the top four for 81.7%, and the top five for 85.7%, indicating that a small subset of techniques dominates the empirical signal available to the model. At the other end of the distribution, 7 of the 27 modelled techniques have only 2 or fewer observations, 11 of 27 have 5 or fewer, and 16 of 27 have 10 or fewer. This imbalance introduces the possibility of representation bias wherein high-support techniques are likely to produce sharper and more reliable posterior estimates, whereas low-support techniques are more strongly influenced by the prior and are therefore less certain, even when their top-ranked CWE appears plausible.

Moreover, the conditioned model achieves very strong ranking scores on the evaluated techniques, including *NDCG@10* = 1.000 in all four cases. While encouraging, this result should be interpreted cautiously. The test partitions are small, the number of relevant CWE classes per technique is often limited, and one of the evaluated techniques is represented by a very sparse test case. Under such conditions, perfect ranking scores are more attainable than they would be in a larger and more heterogeneous evaluation setting.

Among the evaluated cases, T1498 remains the most informative because it combines a non-trivial baseline with a broader set of relevant CWE classes, making the gain from conditioning more meaningful than in the sharply concentrated cases.

The posterior credible intervals reinforce this interpretation. Data-rich techniques exhibit tightly concentrated top-CWE probabilities; for example, T1548.002 assigns a posterior mean of 0.9984 to CWE-434 with a 95% credible interval of [0.9941, 0.99996], whereas sparse techniques show much wider intervals, such as T1601 to CWE-494 with a mean of 0.8891 and an interval of [0.6309, 0.9969], and T1119 to CWE-22 with a mean of 0.5306 and an interval of [0.0329, 0.9815]. This pattern indicates that stable ranking behaviour does not imply equal confidence in posterior calibration across all techniques, and that low-support techniques should be interpreted more cautiously.

Furthermore, these results suggest that technique-conditioned Bayesian estimation is a promising direction for threat-informed vulnerability management but not yet a complete solution. Its current value lies in showing that incident-level ATT&CK observations can, in some cases, shift weakness-class rankings substantially relative to a global exploitation prior. Future work should therefore focus on expanding the incident corpus, strengthening the incident-to-CWE linkage process, incorporating the CWE hierarchy, and exploring how technique-conditioned weakness relevance might be integrated with CVE-level exploitability scores in a broader remediation workflow.

6. Conclusions

We introduced a Bayesian co-occurrence model for technique-conditioned vulnerability prioritisation, estimating *P (CWE | ATT&CK technique)* by integrating NVD records, KEV catalogue, and EuRepoC incident data. By using an exploitation-frequency distribution derived from KEV as the prior, the model combines a globally grounded baseline of observed exploitation with incident-level evidence about adversary behaviour. In doing so, the paper addresses a question that is not answered directly by conventional severity-based or exploit-likelihood-based prioritisation methods which is how remediation priorities at the weakness-class level may shift once a specific ATT&CK technique has been observed.

Within the scope of the dataset and evaluation design, the results suggest that conditioning on ATT&CK technique can alter CWE rankings relative to an unconditioned exploitation prior. For the

evaluated techniques, the conditioned model outperformed the marginal baseline on the selected ranking metrics, and the sensitivity analysis indicated that the relative ordering of CWE classes was stable across the tested range of smoothing values. The findings are encouraging, but they should be interpreted cautiously. The quantitative evaluation covers only a small number of techniques, support across those techniques is uneven, and the learned distributions reflect empirical associations in the available data rather than independently validated causal mappings between techniques and weakness classes.

Accordingly, the principal conclusion of this work is not that technique-conditioned CWE prioritisation has been definitively established, but that it appears to be a feasible and potentially useful direction for threat-informed vulnerability management. The proposed formulation offers an interpretable way to combine observed adversary behaviour with an exploitation-grounded prior, while also providing a basis for future extensions involving richer incident corpora, time-bounded priors, hierarchy-aware CWE modelling, and integration with CVE-level exploitability scores. With such extensions, technique-conditioned Bayesian estimation may become a useful complement to the existing vulnerability prioritisation workflows rather than a replacement for them.

Funding

No financial support was provided for the conduct of this research.

Data Availability

To ensure reproducibility and support the practical implementation of our findings, all project resources have been made publicly accessible at: https://zenodo.org/records/19334720.

References

[1] MITRE. (2026). *Enterprise techniques*, MITRE ATT&CK. [Online]. Available: https://attack.mitre.org/techniques/enterprise/. [Accessed: Mar. 09, 2026].

[2] Forum of Incident Response and Security Teams (FIRST). (2005). *Introduction to CVSS*. [Online]. Available: https://www.first.org/cvss/v1/intro. [Accessed: Mar. 20, 2026].

[3] NIST. (2026). *Vulnerability metrics*. National Vulnerability Database (NVD). [Online]. Available: https://nvd.nist.gov/vuln-metrics/cvss. [Accessed: Mar. 27, 2026].

[4] E. Sherif, I. Yevseyeva, V. Basto-Fernandes, A. Cook, "Bridging the gap between security metrics and key risk indicators: An empirical framework for vulnerability prioritization," *arXiv paper: 2603.12450*, 2026, doi: 10.48550/arXiv.2603.12450.

[5] Forum of Incident Response and Security Teams (FIRST). (2025). *The exploit prediction scoring system.* [Online]. Available: https://www.first.org/epss/. [Accessed: Dec. 02, 2025].

[6] J. Jacobs, S. Romanosky, B. Edwards, I. Adjerid, M. Roytman, "Exploit prediction scoring system (EPSS)," *Digital Threats: Research and Practice*, vol. 2, no. 3, Art. No. 20, pp. 1–17, 2021, doi: 10.1145/3436242.

[7] H. Howland, "CVSS: Ubiquitous and broken," *Digital Threats: Research and Practice*, vol. 4, no. 1, Art. No. 1, pp. 1–12, Mar. 2023, doi: 10.1145/3491263.

[8] A. Kuppa, L. Aouad, N.-A. Le-Khac, "Linking CVE's to MITRE ATT&CK techniques," in *Proceedings of the 16th International Conference on Availability, Reliability and Security (ARES 2021)*, pp. 1-12, Vienna, Austria, 2021, doi: 10.1145/3465481.3465758.

[9] O. Grigorescu, A. Nica, M. Dascalu, R. Rughinis, "CVE2ATT&CK: BERT-based mapping of CVEs to MITRE ATT&CK techniques," *Algorithms*, vol. 15, no. 9, Art. no. 314, pp. 1-22, 2022, doi: 10.3390/a15090314.

[10] MITRE. (2025). *Top 25 most dangerous software weaknesses*. [Online]. Available: https://cwe.mitre.org/top25/archive/2025/2025_cwe_top25.html. [Accessed: Feb. 18, 2026].

[11] European Repository of Cyber Incidents. (2025). *EuRepoC database*. [Online]. Available: https://eurepoc.eu/database/. [Accessed: Mar. 12, 2026].

[12] Cybersecurity and Infrastructure Security Agency (CISA). (2025). *Known exploited vulnerabilities catalog*. [Online]. Available: https://www.cisa.gov/known-exploited-vulnerabilities-catalog. [Accessed: Dec. 03, 2025].

[13] NIST. (2026). *NVD dashboard*, National Vulnerability Database (NVD). [Online]. Available: https://nvd.nist.gov/general/nvd-dashboard. [Accessed: Mar. 20, 2026].

[14] L. Allodi, F. Massacci, "Comparing vulnerability severity and exploits using case-control studies," *ACM Transactions on Information and System Security*, vol. 17, no. 1, Art. no. 1, pp. 1–20, 2014, doi: 10.1145/2630069.

[15] J. Jacobs, S. Romanosky, I. Adjerid, W. Baker, "Improving vulnerability remediation through better exploit prediction," *Journal of Cybersecurity*, vol. 6, no. 1, Art. no. tyaa015, pp. 1-12, 2020, doi: 10.1093/cybsec/tyaa015.

[16] B.E. Strom, A. Applebaum, D.P. Miller, K.C. Nickels, A.G. Pennington, et al., *MITRE ATT&CK: Design and Philosophy*. McLean, VA: MITRE Corporation, Rep. MP180360R1, 2020. [Online]. Available: https://attack.mitre.org/docs/ATTACK_Design_and_Philosophy_March_2020.pdf. [Accessed: Mar. 12, 2026].

[17] J. Baker, R. Struse. (2021). *TRAM: Advancing research into automated TTP identification in threat report*s, Medium, MITRE Engenuity. [Online]. Available: https://medium.com/mitre-engenuity/tram-advancing-research-into-automated-ttp-identification-in-threat-reports-2d868fecc791. [Accessed: Mar. 12, 2026].

[18] MITRE. (2025). *Common weakness enumeration*, the MITRE Corporation. [Online]. Available: https://cwe.mitre.org/. [Accessed: Dec. 03, 2025].

[19] MITRE. (2025). *2025 CWE top 10 KEV weaknesses list insights*, The MITRE Corporation. [Online]. Available: https://cwe.mitre.org/top25/archive/2025/2025_kev_insights.html. [Accessed: Mar. 20, 2026].

[20] S. Simonetto, R. Oostveen, T. van Ede, P. Bosch, W. Jonker, "What matters most in vulnerabilities? Key term extraction for CVE-to-CWE mapping with LLMs," in *Proceedings of Cryptology and Network Security: 24th International Conference, CANS 2025*, Osaka, 2025, Y. Kim, A. Miyaji, M. Tibouchi, Eds., *Lecture Notes in Computer Science*, vol. 16351. Singapore: Springer, 2025, pp. 467–492, doi: 10.1007/978-981-95-4434-9_22.

[21] E. Hemberg, J. Kelly, M. Shlapentokh-Rothman, B. Reinstadler, K. Xu, et al., "Linking threat tactics, techniques, and patterns with defensive weaknesses, vulnerabilities and affected platform configurations for cyber hunting," *arXiv paper: 2010.00533*, 2021, doi: 10.48550/arXiv.2010.00533.

[22] Z. Li, J. Zeng, Y. Chen, Z. Liang, "AttacKG: Constructing technique knowledge graph from cyber threat intelligence reports," in *Computer Security – ESORICS 2022,* Part I, V. Atluri, R. Di Pietro, C. D. Jensen, W. Meng, Eds., *Lecture Notes in Computer Science*, vol. 13554. Cham: Springer, 2022, pp. 589–609, doi: 10.1007/978-3-031-17140-6_29.

[23] O. Sheyner, J. Haines, S. Jha, R. Lippmann, J.M. Wing, "Automated generation and analysis of attack graphs," in *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, Berkeley, CA, 2002, Los Alamitos, CA, IEEE Computer Society, 2002, pp. 273–284, doi: 10.1109/SECPRI.2002.1004377.

[24] L. Wang, S. Jajodia, A. Singhal, S. Noel, "k-Zero day safety: Measuring the security risk of networks against unknown attacks," in *Computer Security – ESORICS 2010,* D. Gritzalis, B. Preneel, M. Theoharidou, Eds., *Lecture Notes in Computer Science, vol. 6345*. Berlin: Springer, 2010, pp. 573–587, doi: 10.1007/978-3-642-15497-3_35.

[25] Y. Kim, I. Lee, H. Kwon, K. Lee, J. Yoon, "BAN: Predicting APT attack based on Bayesian network with MITRE ATT&CK framework," *IEEE Access*, vol. 11, pp. 91949–91968, 2023, doi: 10.1109/ACCESS.2023.3306593.

[26] D.W. Hubbard, R. Seiersen, *How to Measure Anything in Cybersecurity Risk*, 2nd ed. Hoboken, NJ: Wiley, 2023.

[27] P. Xie, J.H. Li, X. Ou, P. Liu, R. Levy, "Using Bayesian networks for cyber security analysis," in *Proceedings of the 2010 IEEE/IFIP International Conference on Dependable Systems and Networks* (DSN), Chicago, IL, 2010, Piscataway, NJ: IEEE, pp. 211–220, doi: 10.1109/DSN.2010.5544924.

[28] J. Wang, M. Neil, N. Fenton, "A Bayesian network approach for cybersecurity risk assessment implementing and extending the FAIR model," *Computers & Security*, vol. 89, Art. no. 101659, pp. 1-20, 2020, doi: 10.1016/j.cose.2019.101659.

[29] S. Thevaratnam, Z. Rezaeifar, "Quantifying cyber threat using Bayesian statistical analysis," *International Journal of Information Security*, vol. 25, no. 2, Art. No. 43, 2026, doi: 10.1007/s10207-026-01220-6.

[30] E. Sherif, I. Yevseyeva, V. Basto-Fernandes, A. Cook. (2026). *Task-specific dataset for empirical analysis of exploited CWE types and ATT&CK technique prevalence*, Zenodo. [Online]. Available: https://zenodo.org/records/19109336. [Accessed: Mar. 30, 2026].

[31] M.R. Rahman, B. Wroblewski, Q. Matthews, B. Morgan, T. Menzies, et al., "Mining temporal attack patterns from cyberthreat intelligence reports," *Knowledge and Information Systems*, vol. 67, no. 10, pp. 8941–8981, 2025, doi: 10.1007/s10115-025-02491-6.

[32] I. Branescu, O. Grigorescu, M. Dascalu, "Automated mapping of common vulnerabilities and exposures to MITRE ATT&CK tactics," *Information*, vol. 15, no. 4, Art. no. 214, pp. 1-19, 2024, doi: 10.3390/info15040214.