IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

Detection of E-Banking Phishing Websites Using Associative Classification: A Review of Machine Learning, Feature Engineering, and Explainable Detection

Prof. Parag B. Patil[1], Dr. Sushma Zade[2], Gaurav Prabhakar Kalwaghe[3], Pratiksha Sanjay Jonwal[4], Khushi Devi Anand[5] and Himanshu Krushnarao Deshmukh[6]

[1]Project Guide, [2]Project Co-coordinator

[123456]Department of Information Technology

Sandip Institute of Technology and Research Centre, Nashik

[3]kalwaghegaurav1@gmail.com, [4]pratiksharajput6070@gmail.com

[5]devikhushi383@gmail.com, [6]hkdeshmukh2005@gmail.com

Abstract : *Electronic banking has made financial services available at any hour, but it has also given attackers a profitable target. Phishing websites imitate genuine banking portals to collect credentials, card details and one-time passwords. Blacklists and manual checks are reactive and can miss newly created sites, so researchers have turned to machine learning, which classifies a site from features of its URL, domain and page. Many such classifiers, however, give a label without a readable reason. Associative classification combines association rule mining with classification and produces IF–THEN rules that an analyst can inspect. This paper reviews the twenty-six sources from our project synopsis that bear on this problem and groups them into six themes: traditional detection, URL and feature-based detection, machine learning, rule induction and associative classification, advanced methods, and datasets. The sources show that feature-based learning is the dominant direction and that rule-based models offer transparency. The supplied sources also leave open how associative classification compares with other classifiers, how well any model copes with unseen sites, and how recall can be kept high without raising false alarms. We outline a conceptual framework for our project; it is a proposal and has not been implemented or validated.*

Key Words: E-banking security, phishing website detection, associative classification, association rule mining, machine learning, feature extraction, explainable artificial intelligence

1. INTRODUCTION

A. Background of Electronic Banking and Phishing Online fund transfers, bill payments, mobile banking and digital wallets let customers handle money without visiting a branch. The same convenience draws criminals, and phishing remains one of the most common ways to reach a customer's account. A phishing website copies the look and workflow of a bank's login or payment page so that a visitor types in a user name, password, card number or one-time password (OTP) [3], [11]. Attackers then use the stolen data for unauthorised transfers or identity theft. Dhamija et al. examined why such imitations work on users [20], and the broader background of cyber crime is covered in [10].

B. Importance of Phishing Website Detection A user rarely has the means to judge a page that looks correct, so detection has to happen before the credentials are typed.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

For a bank, a missed phishing site means direct financial loss for a customer, while a genuine page wrongly blocked erodes trust in the service. Detection of fraudulent banking sites therefore supports both customer protection and institutional reputation, which is the motivation stated in our synopsis.

C. Existing Technological Approaches The literature offers several families of methods. Antiphishing practice has relied on blacklists and heuristics [3], [16], [18]. Other work studies the URL itself [14], [17], collects a wider set of page and host features [1], [13], compares how a page looks [4], or classifies large collections of pages automatically [19]. Surveys and systematic reviews organise this field [5], [11], [16]. Fuzzy and rule-based models, such as [2] and [12], aim to give a decision that can be read.

D. Limitations of Traditional Detection Methods A list can only name sites that have already been reported, and many phishing pages are short-lived, so a list is always behind the attacker. Detecting a site that no one has seen before is the central difficulty, and it motivates the move to learned models [16], [19]. Learned models bring their own problem: a classifier may perform well on a benchmark yet give no account of its decision, which is a drawback when an analyst must justify blocking a bank's customers from a site or acting on an alert.

E. Motivation, Purpose and Scope of the Review

Associative classification mines rules that link feature values to a class and then uses a selected set of those rules as the classifier [6]. Its rules read as IF–THEN statements, which suits a security setting. We do not claim that it is better than other classifiers; part of the purpose of this review is to see what the available evidence says. The review draws on the reference list of our project synopsis only; no outside source has been added. A caution is needed about depth: the synopsis supplies bibliographic details but not the full text of each work, so each study is described at the level its title, venue and our synopsis support, and no accuracy figure, dataset size or ranking is quoted. Three questions guide the review. RQ1: which features and methods do the sources use to detect phishing websites? RQ2: what do they indicate about rule-based or interpretable detection? RQ3: which gaps remain for an associative-classification detector of e-banking phishing? The sources are grouped into the themes of Fig. 1. The topic also has social relevance. Our synopsis links the project to safer digital banking, resilient digital infrastructure and reduced cybercrime (SDG 8, 9 and 16); these are aims of the project, not outcomes shown by this review. The rest of the paper explains key concepts (Section II), reviews the

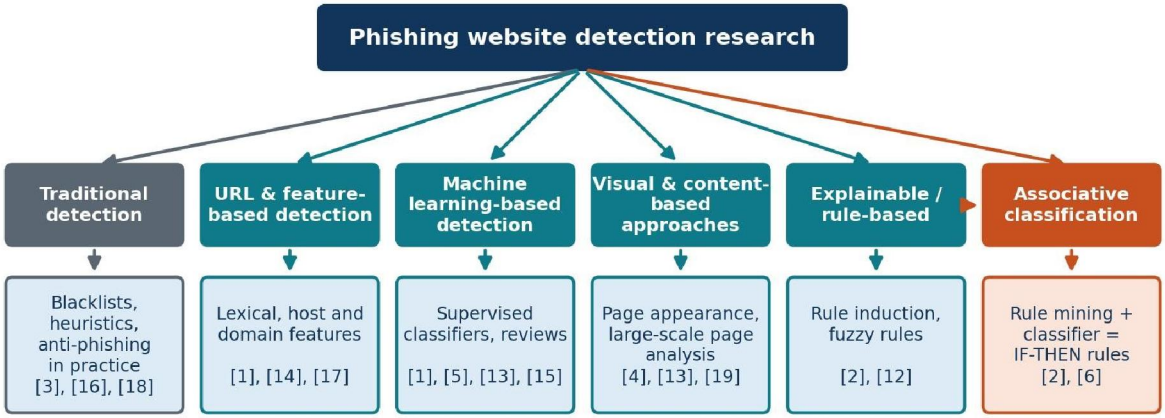


Fig. 1. Research taxonomy for phishing website detection. Reference numbers follow the project synopsis; the grouping of sources into themes is our interpretation, and some sources appear under more than one theme.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

literature (III), compares studies (IV), discusses them critically (V), lists gaps (VI), presents a conceptual framework (VII), outlines an evaluation plan (VIII) and concludes (IX).

II. BACKGROUND AND KEY CONCEPTS

A. E-Banking Phishing Threats

In this paper a phishing website is a page that impersonates a bank in order to obtain confidential data. Systematic reviews of the field catalogue the forms phishing takes [11], and the network-security concepts behind secure sessions, such as certificates and encrypted transport, are treated in standard texts [9]. Human factors matter as well, because a convincing page can defeat a careful user [20].

B. Conventional Phishing Detection Techniques Blacklists compare a requested address against known malicious ones. Heuristic and rule-based checks look for warning signs in the URL or page [3], [18]. Visual-similarity methods compare the rendered appearance of a suspect page with that of a legitimate one [4]. These methods are easy to understand, but the first two depend on prior knowledge of attacks, and the last depends on a reference page being available for comparison.

C. Machine Learning-Based Detection A supervised learner is trained on websites labelled as phishing or legitimate and then predicts the label of a new site [7]. Studies in the synopsis apply this idea to extracted website features [1], [13], to the URL alone [14], [17], and to the wider problem as a general detection task [15]. Systematic reviews of machine learning for phishing give the broader picture [5]. Deep learning, a further branch of the field [8], is mentioned here for completeness; our project does not use it.

D. Association Rule Mining and Associative Classification Association rule mining finds combinations of attribute values that occur together often. A rule is judged by its support, the share of records containing it, and its confidence, the share of those records that also have the stated outcome [6]. Associative classification restricts the right-hand side to a class label, keeps rules that pass chosen support and confidence thresholds, and uses them to label new records [6]. The model is therefore a list of readable rules. An illustrative rule, not taken from any experiment of ours, would be: IF the URL contains suspicious characteristics AND relevant security indicators are abnormal, THEN classify the website as potentially phishing.

E. Website Feature Extraction and Selection

Features come from several layers of a website. URL features include length, structural complexity, the use of an IP address in place of a domain name, special characters, the number of subdomains and redirection [14], [17]. Host and security features include HTTPS usage, SSL/TLS certificate information, domain age and registration details, and DNS and WHOIS records [1], [13]. Page features include hyperlinks and page structure, and the synopsis also lists HTML and JavaScript characteristics [13], [19]. Not every feature helps equally, so feature selection removes irrelevant ones, lowers computation and can limit overfitting, as our synopsis notes. The library support for such pre-processing is described in [24] and [26]. Table II summarises the feature groups and the sources in which each appears.

TABLE II. WEBSITE FEATURE GROUPS AND SOURCES

Feature group	Examples	Sources
URL	Length, IP address, special characters, subdomain count, redirection, URL shortening	[14], [17]
Host and security	HTTPS usage, SSL/TLS certificate, domain age, DNS and WHOIS records	[1], [13]
Page and content	Hyperlinks, page structure, HTML and JavaScript characteristics	[13], [19]
Visual	Rendered page appearance compared with a legitimate page	[4]

F. Performance Evaluation Metrics

Accuracy is the share of correct predictions. Precision is the share of sites flagged as phishing that really are phishing; recall is the share of phishing sites that are found. The F1score combines the two, and the confusion matrix lists the four outcomes. The false-positive rate (FPR) is the share of legitimate sites wrongly flagged, and the false-negative rate (FNR) is the share of phishing sites missed [26]. In banking, a missed site (low recall, high FNR) exposes a customer

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

directly, whereas a high FPR blocks genuine services and teaches users to ignore warnings. Accuracy alone can hide either problem when classes are unequal, so all of these measures should be reported together. With TP, FP, TN and FN denoting true positives, false positives, true negatives and false negatives (phishing is the positive class), the measures are:

$$\text{Precision} = TP / (TP + FP) \quad (1)$$

$$\text{Recall} = TP / (TP + FN) \quad (2)$$

$$F1 = 2 \times \text{Precision} \times \text{Recall} / (\text{Precision} + \text{Recall}) \quad (3)$$

$$FPR = FP / (FP + TN), \quad FNR = FN / (FN + TP) \quad (4)$$

For a rule A ⇒ C over N training records, support is count(A and C) / N (5) and confidence is count(A and C) / count(A)
(6) [6].

III. LITERATURE REVIEW

A. Traditional and Rule-Based Phishing Detection Yue and Wang study anti-phishing measures as they operate in the wild [3], and Garera et al. present a framework for detecting and measuring phishing attacks [18]. Khonji et al. survey the detection literature [16]. Taken together, these works describe how detection was first built from lists, heuristics and measurement, which is the baseline against which later learned models are judged. From the details available to us we cannot state the exact checks used in each, so the comparison in Section IV stays at the level of approach. Dhamija et al. [20] add a human-factors angle: if users misjudge convincing pages, detection that does not rely on the user's own judgement becomes more valuable.

B. URL and Website Feature-Based Detection Le et al. argue in their title that URL names "say it all", and their PhishDef work focuses on the URL as the evidence [17]. Sahingoz et al. likewise classify phishing from URLs with machine learning [14]. Xiang et al. take a broader view with CANTINA+, a feature-rich framework that draws on more than the address [13]. The common thread is that a site can be judged without waiting for it to appear on a list. The difference is cost and reach: URL-only methods are light and fast but see less, while richer feature sets see more but need more data to be collected about each site, some of it only available after fetching the page or querying a registry.

C. Machine Learning Approaches for Phishing Detection Mohammad et al. classify phishing website features with an extreme learning machine [1], and Basnet et al. present a machine learning approach to detecting phishing attacks [15]. Whittaker et al. address automatic classification of phishing pages at large scale [19]. Umer et al. provide a systematic review of machine learning for phishing website detection [5], and the review in [11] concentrates on the nature of phishing attacks rather than on detection models. These studies show a steady move from hand-written rules to trained models, but they report results on their own data and settings, which is why direct ranking of one against another is unsafe (Section IV). A further point is that the older studies [1], [15] and the newer review [5] span more than a decade, so the choice of classifier, the available data and the style of attack are not constant across them. Comparing these works as if they were run on the same footing would therefore be misleading.

D. Associative Classification and Rule-Induction Techniques

Thabtah and Peebles propose a model based on the induction of rules for phishing website detection [2]. This is the closest study in our list to our project, because its output is a rule set that a person can read, and it shares with associative classification the idea of learning rules from labelled sites. The general method of association rule mining, and the way rules can be turned into a classifier, is described in the datamining text [6]. Aburrous et al. use fuzzy techniques for an intelligent phishing detection system [12], which also yields rule-like decisions but handles uncertain feature values with fuzzy logic rather than fixed thresholds. We found no study in the supplied list that applies associative classification itself to e-banking phishing and compares it with other classifiers under the same conditions; this is a statement about our list only, not about the literature as a whole.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

E. Deep Learning, Visual Similarity and Other Advanced Methods

Jain and Gupta analyse visual-similarity-based approaches, which judge a site by how closely it looks like a legitimate page [4]. Deep learning is covered in general terms in [8], although none of the synopsis sources on phishing is a deeplearning study as far as their titles indicate. These advanced methods may raise detection power, but they are less transparent than rules and can be expensive to compute, a point that supports keeping a rule-based model for the explanation role. A hybrid arrangement, in which a rule model gives the reason and another component handles hard cases, is conceivable, but none of the sources we hold evaluates it, so we treat it only as a possible future direction.

TABLE I. COMPARISON OF SELECTED STUDIES FROM THE SYNOPSIS (BASED ON BIBLIOGRAPHIC INFORMATION; LIMITATIONS ARE OUR INTERPRETATION)

Ref.	Approach	Features / techniques	Main contribution	Limitations / cautions	Relevance to our project
[1]	Machine learning (extreme learning machine)	Phishing website features	Classifies sites from a set of website features	Black-box model; results not extracted	Shows the feature-based starting point
[2]	Rule induction	Website features; learned rule set	Readable rule-based phishing model	Details need full-text check	Closest rule-based comparator
[3]	Anti-phishing in practice	Deployed measures	Describes anti-phishing in the wild	Detail needs full-text check	Baseline of existing defences
[4]	Visual similarity analysis	Page appearance	Analyses visual-similarity approaches	Needs a reference page; costly	Complements feature rules
[5]	Systematic review	Machine learning methods	Reviews ML for phishing detection	Review; no new model	Map of ML options and baselines

F. Phishing Datasets and Evaluation Practices The synopsis names four dataset sources: a Kaggle phishing website detection dataset [27], the UCI phishing websites data set [28], the PhishTank community feed [29] and the OpenPhish feed [30]. Feeds such as [29] and [30] list reported phishing URLs, so they supply the phishing side of a sample and need a separate source of legitimate sites; the entries give no further detail on collection dates or labelling. The bibliographic details we hold for these entries are minimal (no authors, URLs or access dates), so dataset size, age and balance cannot be assessed from the synopsis, and we flag this as a verification need.

IV. COMPARATIVE ANALYSIS

Table I compares eleven sources from the synopsis. The “approach”, “features” and “contribution” columns restate what each work’s title and venue indicate; the “limitations” and “relevance” columns are our own reading and should not be taken as findings reported by the authors. Several observations follow from the table. First, the studies divide by how much evidence they use about a site: the address alone [14], [17], a wider set of host and page features [1], [13], [19], or appearance [4]. Second, rule-producing methods [2], [12] stand out as the ones whose decisions can be read, while the model in [1] is a trained network whose reasoning is not visible to a user. Third, new sites remain the weak point of list-based practice [3], [16], [18], and this is the reason the learned and rule-based methods exist at all. Fourth, every learned method depends on the quality of its data and features, so a result is tied to the dataset it was obtained on. Fifth, the sources do not report errors in a common way, so false positives and false negatives cannot be compared across them from the information we hold. This last point is important for a bank, and it argues for evaluating our own model against baselines under one protocol rather than quoting figures from different papers.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

V. CRITICAL ANALYSIS AND DISCUSSION

A. Common Strengths of Existing Approaches Feature-based learning can judge a site that no list has recorded, and it can use cheap signals such as URL structure [14], [17]. Surveys and reviews give a stable map of the field [5], [16]. Rule-based and fuzzy models contribute readable decisions [2], [12].

B. Recurring Technical Limitations

Three issues recur in the sources we hold. Methods depend on the features chosen, and the best feature set may differ between datasets. Richer features need more collection effort per site [13]. Results are usually given for one setting, so conclusions about generalisation are weak [5].

C. Trade-Offs Between Performance and Interpretability Complex models may detect more, but they explain less; rule-based models explain more but may need many rules to cover varied attacks [6]. We treat the size of this trade-off as an open empirical question for our own experiments, because the sources do not measure it for associative classification.

D. Dataset Quality, Generalisation and Evolving Attacks Phishing pages are short-lived, so a dataset ages quickly, and a model trained on old records may see little of current attack style [16], [19]. Feeds [29], [30] help refresh the phishing side, but class balance and the choice of legitimate sites are decisions that affect results. Duplicate or near-duplicate records between training and test sets can inflate scores, which is why duplicate removal is part of our plan.

Ref.	Approach	Features / techniques	Main contribution	Limitations / cautions	Relevance to our project
[12]	Fuzzy techniques	Fuzzy rules over indicators	Intelligent detection with fuzzy logic	Rule design and tuning effort	Alternative rule-like approach
[13]	Feature-rich ML framework	URL, host and page features	Broader evidence than URL alone	More data to collect per site	Source of feature ideas
[14]	ML from URLs	URL-based features	Detection from the address only	Sees less than full-page methods	Supports lightweight URL features
[16]	Literature survey	Detection techniques	Surveys detection methods	Older survey; may miss recent work	Background and gap evidence
[17]	URL-name analysis	URL names	Treats the URL as main evidence	Limited to URL-level evidence	Supports URL feature group
[19]	Large-scale classification	Page-level features	Automatic classification of many pages	Large data and infrastructure needs	Shows scalability concerns

E. Practical Challenges in Real-World E-Banking Security A deployed checker must obtain domain, DNS and certificate data quickly, and some of this metadata may be missing or blocked [13]. The tool must also keep false alarms low, since users stop trusting a warning that is often wrong, and it must be refreshed as attacks change.

F. Potential Benefits and Limits of Associative Classification

Possible benefits are human-readable IF–THEN rules, predictions that can be explained with the rules that fired, and the discovery of relationships between features [2], [6]. Possible limits are a very large number of rules (rule explosion), redundant rules, sensitivity to the support and confidence thresholds, difficulty with imbalanced classes, and weaker behaviour on sites unlike the training data. These are expectations drawn from the method's design as described in [6] and from the general machine-learning literature [7]; they are not measured results for phishing, and our experiments will have to test them.

G. Threats to the Validity of This Review The review is limited to the reference list of our synopsis and was not extended by a database search, so relevant studies may be missing. The sources were read at the level of their bibliographic details and the synopsis, not through full text, so details of method and results may differ from our short descriptions. Several bibliographic entries need verification before submission. Our groupings and the limitations in Table I are interpretations, and the reader should treat them as such.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

VI. RESEARCH GAPS

The following gaps emerge from the supplied sources. Each is stated for our list only.

Gap 1: comparative evidence. Among the sources we hold, [2] is the nearest rule-based study, and none of the sources we hold compares associative classification with other phishing classifiers under one protocol. This matters because our claim of interpretability must be balanced against detection quality.

Gap 2: unseen and short-lived sites. List-based practice is reactive [3], [16], and the learned methods in [14], [19] address new sites but, from the information we hold, do not tell us how a rule-based model degrades as attacks change. We will examine performance on records held out from training.

Gap 3: recall versus false alarms. The sources do not report FPR and FNR in a common form, so the balance cannot be compared. Our project will report both.

Gap 4: feature reliability. Different studies use different feature sets [1], [13], [14], [17], and some features depend on registry or DNS data that may be unavailable [13]. Which features are both informative and obtainable in real time remains to be established.

Gap 5: auditable rules. Rule-producing work exists [2], [12], but how a bank analyst would audit, prune and update a mined rule set is not addressed in the sources we hold.

Gap 6: dataset quality. Dataset descriptions in the synopsis are thin [27]–[30]. Imbalance, age and duplicate records are therefore unknown and are a risk for any evaluation. Table IV links each gap to its sources and to what our project will do about it.

TABLE IV. RESEARCH GAPS AND THEIR EVIDENCE BASE

Gap	Supporting sources	Relevance to our project
1 Comparison	[2], [6]	Compare associative classification with baselines on one protocol
2 Unseen sites	[3], [14], [16], [19]	Hold out later-collected records; report recall
3 Recall vs. FPR	[26] (metrics)	Report FPR and FNR with every result
4 Features	[1], [13], [14], [17]	Select features that are informative and obtainable
5 Auditable rules	[2], [12]	Prune and review the mined rule set
6 Dataset quality	[27]–[30]	Document origin, age, balance and duplicates

VII. PROPOSED CONCEPTUAL FRAMEWORK

As shown in Fig. 2, the framework follows the block diagram in our synopsis and has nine stages. This is a proposed design. It has not been implemented or experimentally validated.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

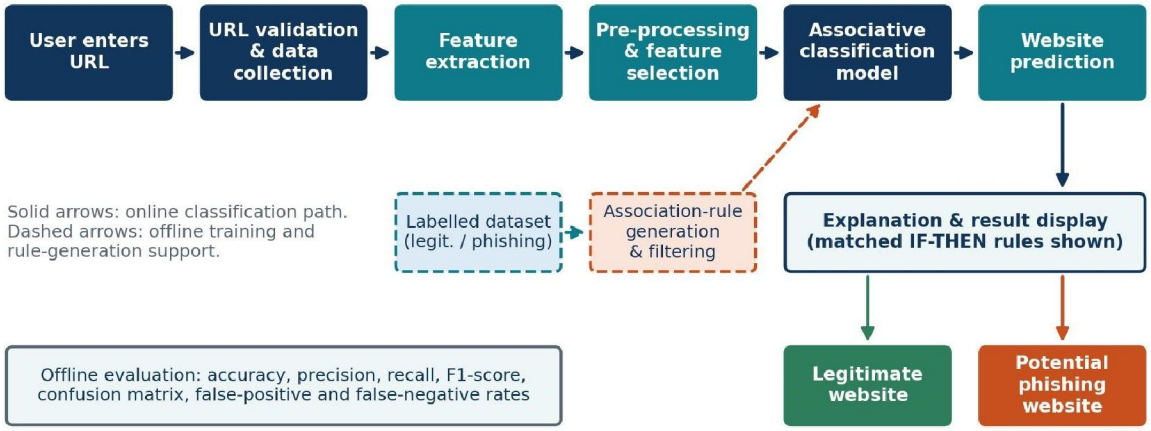


Fig. 2. Proposed e-banking phishing detection workflow. Dashed elements support offline training and rule generation. The framework is a proposed design that has not been implemented or validated.

Some design considerations follow from the review. The synopsis aims for a prediction within a few seconds; this is a target, not a measured figure, and it depends on how quickly domain, DNS and certificate data can be retrieved [13]. Rules with high support and confidence will be retained, but the thresholds are tuning choices that interact with class balance, so they will be reported rather than fixed silently [6]. When no rule fires for a site, the system needs a default behaviour, for example a default class or an explicit "unknown" result; which is safer for a bank is a design question for the experiments. The literature informs the design in four ways. The feature groups follow the URL and host features in [1], [13], [14] and [17]. The rule output follows the rule-based direction of [2] and [12]. Visual comparison [4] is left out of scope because it needs reference pages and more computation. Datasets are chosen from the four named sources [27]–[30]. Fig. 3 shows the associative classification step itself. An illustrative explanation might read "classified as phishing because IP address in URL = yes, HTTPS = no and domain age is short". This is only an example of the output format, not a result.

(1) The user enters a banking website URL. (2) The URL is checked for validity and website information is collected, including domain, HTTPS, certificate and DNS data [9], [13]. (3) Features are extracted from the URL, host and page layers described in Section II-E [1], [14], [17]. (4) Data are cleaned, duplicates and missing values are handled, and features are selected [24], [26]. (5) A labelled dataset from the sources in [27]–[30] supplies training records. (6) Association rules are mined and filtered by support and confidence, then used as the associative classifier (Fig. 3) [2], [6]. (7) The website is classified as legitimate or potentially phishing. (8) The prediction is explained by the rules that apply to it. (9) The result is displayed with a warning or a safe message, and the model is evaluated offline.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

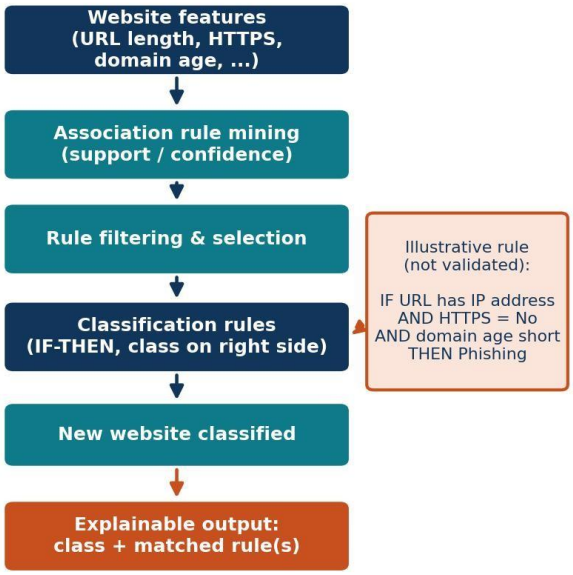


Fig. 3. Conceptual associative classification process. The rule shown is illustrative and has not been validated.

VIII. EVALUATION PLAN AND FUTURE DIRECTIONS

We plan to select one labelled dataset from the sources in [27]–[30] and to record its origin, date and class balance. Cleaning will remove duplicates, handle missing values and encode categories [24]. The synopsis proposes an 80%/20% training and testing split; we will keep the test set untouched during rule and threshold selection, and use cross-validation inside the training set for those choices, to limit data leakage. Support and confidence thresholds will be varied and reported. If the project scope allows, baseline classifiers available in scikit-learn [26] will be trained on the same split for a fair comparison. We will report accuracy, precision, recall, F1-score, the confusion matrix, FPR and FNR, and we will examine class imbalance, performance on records collected later than the training data, and the number, length and redundancy of mined rules. The usefulness of the rules will be judged by whether a reviewer can follow why a site was flagged. To keep the work reproducible, we will fix random seeds, record the versions of Python and the libraries used [26], and store the thresholds and the final rule set with each experiment. The synopsis lists Python, pandas, scikitlearn, MLxtend (for Apriori and FP-Growth rule mining) and a Flask web interface as the planned tools; whether MLxtend alone gives a complete associative classifier, or whether the classification step must be coded separately, is a point we still have to confirm. No result is reported in this paper. Table III summarises the plan.

TABLE III. SUMMARY OF THE PLANNED EVALUATION

Dataset	One labelled set from [27]–[30]; origin, date and class balance recorded
Cleaning	Remove duplicates, handle missing values, encode categories [24]
Split	80% training / 20% testing; test set untouched during tuning; cross-validation for thresholds
Baselines	Scikit-learn classifiers on the same split, if scope allows [26]
Metrics	Accuracy, precision, recall, F1-score, confusion matrix, FPR, FNR
Risks checked	Class imbalance, data leakage, unseen and later-collected sites
Rule quality	Rule count, length, redundancy and readability for a reviewer

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

Future work follows from the gaps: a controlled comparison of associative classification with other classifiers (Gap 1), rule pruning for large rule sets (Gap 5), and periodic retraining on fresh phishing feeds (Gaps 2 and 6). Integration into a browser extension or banking application, noted in the synopsis, would depend on first obtaining validated results.

IX. CONCLUSION

Phishing against online banking is a continuing threat, and the sources reviewed show a clear shift from lists and heuristics towards models learned from website features. Rule-producing methods [2], [12] address the need for decisions that people can read, and associative classification [6] is a natural way to obtain such rules, although the sources do not show that it beats other classifiers. The main limits are dependence on dataset and features, weak handling of unseen sites, uneven reporting of false positives and false negatives, and thin dataset documentation. Our next steps are to implement the framework in Section VII, run the evaluation in Section VIII, and verify the bibliographic entries and full texts before the final submission. Only then can any claim about performance be made.

REFERENCES

[1] Rami M. Mohammad, Fadi Thabtah, Lee McCluskey, "Phishing Websites Features Classification Based on Extreme Learning Machine," International Journal of Computer Applications, Vol. 44, No. 16, pp. 1–7, April 2012.
[2] Fadi Thabtah, David Peebles, "A New Machine Learning Model Based on Induction of Rules for Phishing Website Detection," Applied Soft Computing Journal, Elsevier, Vol. 68, pp. 667–674, July 2018.
[3] Chuan Yue, Haining Wang, "Anti-Phishing in the Wild," IEEE Internet Computing, Vol. 13, No. 5, pp. 21–28, September–October 2009.
[4] A. K. Jain, B. B. Gupta, "Phishing Detection: Analysis of Visual Similarity Based Approaches," Security and Communication Networks, Vol. 2017, Article ID 5421046, pp. 1–20, 2017.
[5] Mohammad A. Umer, Muhammad Naveed, et al., "Phishing Website Detection Using Machine Learning: A Systematic Review," IEEE Access, Vol. 11, pp. 41012–41039, 2023.
[6] Jiawei Han, Micheline Kamber, Jian Pei, "Data Mining: Concepts and Techniques," Morgan Kaufmann Publishers, 3rd Edition, pp. 1–703, 2011.
[7] Tom M. Mitchell, "Machine Learning," McGraw-Hill Education, New York, pp. 1–432, 1997.
[8] Ian Goodfellow, Yoshua Bengio, Aaron Courville, "Deep Learning," MIT Press, Cambridge, pp. 1–800, 2016.
[9] William Stallings, "Network Security Essentials: Applications and Standards," Pearson Education, 7th Edition, pp. 1–640, 2017.
[10] Nina Godbole, Sunit Belapure, "Cyber Security: Understanding Cyber Crimes, Computer Forensics and Legal Perspectives," Wiley India, pp. 1– 456, 2011.
[11] Veronika Švábenská, Jan Vykopal, Pavel Čeleda, "What Are We Protecting Against? A Systematic Review of Phishing Attacks," IEEE Access, Vol. 8, pp. 85732–85748, 2020.
[12] M. Aburrous, M. A. Hossain, K. Dahal, F. Thabtah, "Intelligent Phishing Website Detection System Using Fuzzy Techniques," Expert Systems with Applications, Vol. 38, No. 3, pp. 13403–13418, 2011.
[13] Guang Xiang, Jason Hong, Carolyn Rose, Lorrie Cranor, "CANTINA+: A Feature-Rich Machine Learning Framework for Detecting Phishing Web Sites," ACM Transactions on Information and System Security, Vol. 14, No. 2, pp. 1–28, 2011.
[14] Y. Sahingoz, E. Buber, O. Demir, B. Diri, "Machine Learning Based Phishing Detection from URLs," Expert Systems with Applications, Vol. 117, pp. 345–357, 2019.
[15] A. Basnet, S. Mukkamala, A. Sung, "Detection of Phishing Attacks: A Machine Learning Approach," International Conference on Soft Computing Applications, pp. 373–383, 2012.
[16] M. Khonji, Y. Iraqi, A. Jones, "Phishing Detection: A Literature Survey," IEEE Communications Surveys & Tutorials, Vol. 15, No. 4, pp. 2091–2121, 2013.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

[17] A. Le, A. Markopoulou, M. Faloutsos, "PhishDef: URL Names Say It All," IEEE INFOCOM Workshops, pp. 191–196, 2011.
[18] S. Garera, N. Provos, M. Chew, A. D. Rubin, "A Framework for Detection and Measurement of Phishing Attacks," ACM Workshop on Recurring Malcode, pp. 1–8, 2007.
[19] C. Whittaker, B. Ryner, M. Nazif, "Large-Scale Automatic Classification of Phishing Pages," Network and Distributed System Security Symposium (NDSS), pp. 1–14, 2010.
[20] R. Dhamija, J. D. Tygar, M. Hearst, "Why Phishing Works," Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, pp. 581–590, 2006.
[24] Wes McKinney, "Python for Data Analysis," O'Reilly Media, 3rd Edition, pp. 1–544, 2022.
[26] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research, Vol. 12, pp. 2825–2830, 2011.
[27] Kaggle Inc., "Phishing Website Detection Dataset," Kaggle Repository, 2024.
[28] University of California Irvine (UCI), "Phishing Websites Data Set," UCI Machine Learning Repository, Irvine, California.
[29] PhishTank Community, "PhishTank: Collaborative Clearing House for Data and Information about Phishing," OpenDNS, 2024.
[30] OpenPhish, "OpenPhish Phishing Intelligence Feed," OpenPhish Security Platform, 2024.