CSEC-472 (Authentication) Final Paper

Walking the Tightrope: How Risk-Based Authentication Balances Trust and Friction

Michael Bhardwaj[1], Philip Colangelo[1], and Kevin McCarthy[1]

[1]Department of Computing Security, Rochester Institute of Technology

May 4, 2026

Abstract

Passwords remain the dominant authentication mechanism across the web, even though they are vulnerable to phishing, credential stuffing, password reuse, and session hijacking [1, 2]. Risk-Based Authentication (RBA) addresses this limitation by evaluating contextual signals during each login attempt, including IP address, geolocation, device fingerprint, browser characteristics, login history, time of day, and behavioral patterns [3, 4]. When risk is low, the user experiences a normal login. When risk is elevated, the system escalates to a step-up challenge, such as a one-time password, a biometric prompt, a push notification, or an account block.

This review synthesizes 35 sources from 2012–2025 across security, privacy, machine learning, human-computer interaction, mobile computing, cloud systems, and domain-specific deployments. Three findings emerge. First, machine-learning-based risk-scoring methods, including Random Forests [5], feed-forward neural networks [6], and deep reinforcement learning [7], often outperform classical probabilistic baselines when class imbalance and threshold selection are carefully handled. Second, real-world deployment remains limited: Gavazzi et al. [8] audited 208 popular websites and found that only 42.31% supported any form of multi-factor authentication and only 22.12% blocked an obvious account-hijacking attempt. Third, usability and privacy challenges remain unresolved. Users often rate RBA as more usable than always-on two-factor authentication [9, 10]. However, notification confusion, recovery deadlocks, administrator misconfiguration, and privacy risks from device fingerprinting continue to limit practical deployment [11, 12, 13, 14].

Overall, the literature shows that RBA is a technically mature and practical improvement over password-only authentication. However, its success depends on broader deployment, stronger evaluation standards, privacy-preserving risk signals, and user-centered challenge-and-recovery design.

Keywords

Risk-based authentication, adaptive authentication, machine learning, anomaly detection, usability, device fingerprinting, behavioral biometrics, privacy-preserving authentication, multi-factor authentication, credential stuffing, continuous authentication

1 Introduction

Authentication is the first line of defense for most digital systems, but the most common authentication mechanism remains the password. Passwords were designed as shared secrets, not as defenses against large-scale credential stuffing, phishing kits, automated account takeover, or stolen session tokens. Their persistence creates the central problem that Risk-Based Authentication (RBA) attempts to solve: how can systems improve account security without adding unnecessary friction to every login?

RBA matters because it aims to preserve what makes passwords deployable while addressing what makes them dangerous. Instead of replacing passwords outright, RBA adds contextual judgment around them. A login from a familiar device, network, and location can proceed with minimal interruption. A login from an unfamiliar device, an impossible travel pattern, an unusual browser fingerprint, or a suspicious network can trigger a stronger challenge. This makes authentication adaptive rather than static.

1.1 The Problem with Passwords

Bonneau et al. [1] showed why passwords remain difficult to replace. Although many authentication alternatives improve security, no proposed replacement satisfies the same combination of usability, deployability, and familiarity. Passwords survive not because they are secure, but because they are cheap, universal, easy to integrate, and already understood by users and services.

That deployability comes with serious security weaknesses. Credential stuffing reuses leaked username-password pairs across services. Phishing steals credentials directly from users. Session hijacking bypasses the login process once an attacker obtains a valid token [15, 2]. Password reset workflows also create additional attack surfaces when recovery procedures are weak or inconsistent [2]. These attacks exploit the same basic weakness: once a system accepts a password, it often treats the session as legitimate without enough contextual judgment.

Two-factor authentication (2FA) partially addresses this problem by requiring an additional proof of identity, but always-on 2FA introduces friction at every login. Users often accept this burden for high-value accounts, such as banking, but resist it for routine services [9, 10]. A better mechanism should add friction only when the risk justifies it.

1.2 What Risk-Based Authentication Does Differently

RBA supplements traditional authentication with a dynamic risk score. Instead of treating every login attempt identically, the system compares each attempt against expected user behavior and environmental context [3, 4]. Common signals include IP address, autonomous system number, geolocation, user agent, browser and device characteristics, time of day, login history, and behavioral biometrics such as typing cadence, touch pressure, or application-use patterns [5, 16].

When the attempt resembles prior legitimate behavior, the login can proceed without interruption. When it deviates from the user's normal profile, the system escalates to an additional challenge, such as an email code, push notification, biometric prompt, or temporary block [17]. This model turns authentication into a context-aware decision process rather than a one-time gate.

The gap between the RBA's promise and deployment is substantial. Gavazzi et al. [8] audited 208 popular websites from the Tranco top 5,000 that supported account creation and found that only 42.31% supported any form of MFA and only 22.12% blocked an obvious account-hijacking attempt. Even when RBA exists, it is often inconsistent across services and account contexts [4, 18]. This suggests

that the RBA's main challenge is no longer basic feasibility, but reliable, privacy-conscious, and usable deployment.

1.3 Scope and Research Questions

RBA research spans security engineering, machine learning, human-computer interaction, privacy, mobile computing, healthcare, cloud infrastructure, and IoT. This review synthesizes 35 sources from 2012–2025 to identify what the literature currently shows and what problems remain unresolved.

Three research questions guide the review:

1. What machine learning methods and data sources are used for RBA risk scoring, and how do they compare in accuracy, cost, and generalizability?
2. How widely is RBA deployed in practice, and what barriers prevent broader adoption?
3. What security, usability, and privacy trade-offs shape whether RBA systems succeed in real environments?

2 Background and Significance

2.1 Password Persistence and Authentication Risk

Knowledge-based authentication has been the dominant paradigm in computer security since the early days of multi-user systems. Passwords are the most common implementation of this paradigm and remain embedded in websites, enterprise applications, cloud services, and mobile platforms. Their staying power is not accidental. Passwords require no specialized hardware, no biometric enrollment, no identity federation, and no external authenticator application [1]. For service providers, passwords are easy to store, reset, and integrate with existing login flows.

The same simplicity creates predictable weaknesses. Users reuse passwords, select weak secrets, store credentials insecurely, and fall for phishing attacks. Attackers exploit password reuse through credential stuffing and exploit recovery workflows through social engineering or poorly designed reset logic [2, 19, 20]. Once a credential is accepted, many systems do not meaningfully distinguish between the real user and an attacker who has obtained the correct secret.

RBA is significant because it does not require abandoning passwords completely. Instead, it adds contextual evaluation around an authentication process that users and services already understand. This makes RBA more deployable than many password-replacement schemes while still addressing attacks that password-only systems cannot distinguish from legitimate access.

2.2 RBA Architecture

Most RBA systems follow a common structure. First, the service collects contextual features during login. Second, it compares those features to user-specific and population-level baselines. Third, it computes a risk score. Fourth, it maps the score to an authentication response [3, 4, 16]. Low-risk logins proceed normally; medium-risk logins receive step-up authentication; high-risk logins may be blocked.

The features used in RBA vary by domain. Web services commonly use IP addresses, geolocation, browser fingerprints, operating systems, screen resolutions, cookies, login history, and time-of-day patterns. Mobile systems may include accelerometer readings, gyroscope patterns, app transition sequences, touch dynamics, and device-specific signals [21, 22, 23]. Enterprise and cloud systems may incorporate role, resource sensitivity, network segment, policy state, and organization-specific risk factors [24, 25, 26, 27].

Classical RBA systems often rely on prob-

abilistic models. These models estimate whether a login context is common for the user and uncommon for attackers. More recent systems use machine learning classifiers, including Random Forests, support vector machines, k-nearest neighbors, neural networks, and reinforcement learning [7, 6, 5]. These approaches aim to detect unusual access patterns more accurately than fixed rules. However, they also introduce new issues: model opacity, training-data dependence, class imbalance, computational cost, and inconsistent evaluation metrics.

2.3 The Security-Usability Trade-off

RBA's value depends on the threshold design. If the threshold is too strict, legitimate users are challenged too often, which can make them frustrated. If the threshold is too permissive, attackers can pass without step-up authentication. Wiefling et al. [3] showed that small configuration changes can significantly affect both attacker detection and legitimate-user re-authentication rates.

User studies show the same tension. RBA is often perceived as more usable than always-on 2FA because it reduces unnecessary interruptions [9, 10]. However, if users are never challenged, they may not perceive that the system is providing any additional protection [9]. Effective RBA, therefore, needs not only accurate scoring but also visible, understandable, and trustworthy security communication.

The trade-off is not only an end-user issue. Administrators must choose thresholds, escalation actions, and recovery policies. Markert et al. [13] found that administrators can misunderstand threshold semantics and rely heavily on insecure defaults. This means that an RBA system can be technically strong and still fail in practice if its configuration interface does not help administrators understand the security and usability consequences of their choices.

2.4 Privacy and Contextual Data Collection

RBA depends on contextual data that can also expose sensitive user information. IP addresses, device fingerprints, geolocation data, and behavioral traces can reveal where users are, what devices they use, when they access services, and how their behavior changes over time. These signals improve account protection, but they also create privacy and compliance concerns [11, 14].

Browser fingerprinting illustrates this tension. A detailed fingerprint can help detect suspicious access, but it can also be used to track users across sessions or services. Lin et al. [15] showed that fingerprint-related assumptions can become attack surfaces when phishing pages collect and replay signals that services treat as trusted. Wiefling et al. [14] further showed that changes in browser identification mechanisms, such as the shift from user-agent strings to Client Hints, can affect both privacy and the quality of the RBA signal.

For this reason, privacy cannot be treated as a secondary concern. A useful RBA system should collect only the signals it needs, retain them only as long as necessary, explain collection practices clearly, and prefer less invasive signals when possible. Otherwise, RBA can become a security mechanism that protects accounts by creating a separate surveillance risk.

3 Related Work

Prior authentication research generally follows two paths: replacing passwords or strengthening password-based systems with additional controls. Bonneau et al. [1] provide a foundational framework for the replacement path. Their comparison of authentication schemes shows why passwords persist despite clear weaknesses. Many alternatives improve security, but sacrifice deployability, cost, recoverability, or user familiarity.

RBA belongs to the second path. It does not try to remove passwords from the ecosystem. Instead, it recognizes that passwords will remain common and attempts to make password-based authentication less brittle. This distinguishes RBA from passwordless approaches, hardware-token schemes, and biometric-only systems. RBA is best understood as an adaptive layer that can coexist with passwords, 2FA, biometrics, and account recovery workflows.

Real-world measurement studies show that this adaptive layer is not yet common enough. Gavazzi et al. [8] audited 208 popular websites and found weak adoption of both MFA and suspicious-login detection. Wiefling et al. [4] found that deployed RBA behavior changes across services and time. Makowski and Pöhn [18] found that major platforms differ significantly in how they respond to suspicious logins. Together, these studies suggest that RBA deployment is fragmented and difficult to reason about from the outside.

Banking research reflects similar inconsistency. Karim et al. [19] reviewed authentication practices in online banking and found that only a minority of the surveyed banks used behavioral risk scoring. This is notable because banking is a high-value domain where stronger authentication has obvious benefits. If RBA adoption remains limited even there, the barrier is likely not only technical performance but also operational complexity, costs, regulations, and user support burden.

A separate line of related work examines continuous and context-aware authentication. Mobile studies use touch, motion, application usage, and sensor patterns to identify users throughout a session rather than only at login [21, 22, 23]. Cloud and enterprise studies extend risk evaluation to resource access and policy state [24, 25, 26, 27]. These systems blur the boundary between RBA and broader Zero Trust architectures, in which identity is continuously evaluated, and access is adapted to risk [28].

The privacy literature adds an important counterbalance. Wiefling et al. [11] argue that RBA's data collection practices are underexamined relative to their security benefits. Lin et al. [15] show that attackers can exploit device-fingerprint assumptions. These findings prevent the literature from treating more data as automatically better. RBA improves authentication by adding context, but context is valuable precisely because it reveals information about users.

4 Research Method

This paper uses a systematic literature review as its research method. RBA spans multiple research communities, including machine learning, usable security, web measurement, privacy, mobile computing, healthcare, cloud infrastructure, and IoT. A review is appropriate because no single empirical study captures the full range of RBA's technical methods, deployment realities, and human factors.

The guiding research question was: *What are the current methods, data sources, and usability-privacy trade-offs in RBA research, and what open problems remain?*

4.1 Search Strategy

We conducted keyword searches across Google Scholar, IEEE Xplore, the ACM Digital Library, and the USENIX open-access archive. Search terms included combinations of "risk-based authentication," "adaptive authentication," "behavioral biometrics," "security-usability trade-off," "anomaly detection authentication," "continuous authentication," "device fingerprinting," and venue names such as USENIX Security, ACM CCS, NDSS, IEEE S&P, and ACM Transactions on Privacy and Security.

The review prioritizes work from 2020–2025 to capture the recent shift toward machine-learning-based RBA, while also including foundational authentication work from 2012 on-

ward. We considered venue quality using CORE and Scimago when applicable, but did not restrict the corpus only to top-ranked venues. Some relevant work appears in workshops, domain-specific journals, and system papers, and is valuable because it addresses implementation, privacy, or deployment barriers rather than only model performance.

4.2 Inclusion Criteria

Sources were included when they met at least one of the following criteria:

1. They proposed, evaluated, or implemented an RBA or adaptive authentication system.
2. They measured real-world adoption of MFA, RBA, suspicious-login detection, or account recovery controls.
3. They studied user perception, usability, notification design, or administrator configuration for RBA-like systems.
4. They analyzed privacy, fingerprinting, or contextual-data risks directly relevant to RBA.
5. They provided the foundational authentication context needed to explain why RBA is necessary.

We extracted each source's research method, study design, dataset, population, authentication signals, model type, evaluation metrics, security findings, usability findings, and privacy implications. We then grouped the results into three recurring themes: risk scoring and machine learning; data sources and deployment; and security, usability, and privacy trade-offs.

4.3 Limitations

The main limitation is the scarcity of production-scale RBA datasets. Many technical studies rely on small public behavioral datasets, such as the CRAWDAD MPU dataset with 42 users [7] or the LiveLab dataset with 35 users [22]. These datasets are useful for controlled evaluation but do not fully represent noisy, adversarial, long-term production login streams. Another limitation is publication bias: successful models are more likely to be published than negative or inconclusive results. To reduce this bias, the review includes deployment measurement studies, privacy analyses, administrator studies, and user-experience studies in addition to model-performance papers [11, 8, 4, 12, 9, 29, 13].

A final limitation is that many deployed RBA systems are proprietary. Major platforms rarely disclose their exact features, thresholds, models, or escalation logic. As a result, black-box measurement studies are necessary but imperfect. They can observe how a service responds to controlled login scenarios, but they cannot fully explain why a challenge was or was not triggered.

5 Findings

Analyzing the 35 sources in the corpus, three themes stand out: the move toward machine-learning-based risk scoring, the gap between research datasets and real-world deployment, and the persistent security-usability-privacy trade-off.

5.1 Theme One: Risk Scoring and Machine Learning

The clearest technical trend is the move from rule-based or probabilistic scoring toward machine-learning classifiers. Earlier RBA systems used environmental features such as IP addresses, login times, and device fingerprints in probabilistic formulas [3]. Newer work attempts to model more complex patterns using supervised, unsupervised, and reinforcement-learning methods.

Picard and Pierre [7] built RLAuth, a deep reinforcement learning system that treats au-

thentication as a sequential decision problem. Tested on the CRAWDAD MPU dataset, RLAuth achieved a G-Mean of 92.62%, outperforming SVM, DNN, Gaussian Naïve Bayes, k-NN, Isolation Forest, and one-class SVM baselines. The study is especially relevant because it addresses class imbalance, a core problem in authentication data where legitimate logins vastly outnumber attacks.

Rotter et al. [6] showed that a feed-forward neural network trained on likelihood-style features derived from IP addresses and user-agent strings could outperform a classical statistical baseline while reducing unnecessary re-authentication for legitimate users. Their evaluation across multiple thresholds is valuable because RBA performance depends on the operating point chosen by the service. A model that looks strong at one threshold may be unacceptable at another if it challenges too many legitimate users or allows too many attackers through.

Other studies expand the feature space. Ashibani and Mahmoud [22] used temporal app-access patterns for smartphone authentication. Papaioannou et al. [21, 23] incorporated accelerometer, gyroscope, and touch dynamics into mobile novelty-detection pipelines. Banerjee and Singh [30] added hardware-level signals such as MAC address binding alongside AI-driven behavioral scoring. These studies show that RBA is not limited to web-login metadata; it can incorporate behavioral and device-level features when the deployment context supports them.

The overall finding is that machine learning improves RBA's detection capability, but comparisons remain difficult. Studies use different datasets, threat models, thresholds, and metrics. Accuracy alone is often misleading because of class imbalance. A classifier can appear accurate by labeling most attempts legitimate, even if it fails to catch the rare malicious attempts that matter most. Better reporting should include false accept rates, false reject rates, re-authentication frequency, attacker success rates, computational cost, and usability impact.

5.2 Theme Two: Data Sources and Deployment

RBA research often depends on datasets that are too small or controlled to support broad generalization. Mobile and behavioral studies commonly use a few dozen users, including CRAWDAD MPU with 42 users [7] and LiveLab with 35 users [22]. These datasets help evaluate algorithms, but they do not fully capture account age, device churn, travel, VPN use, accessibility needs, shared devices, or adversarial adaptation.

Production-scale evidence is rare. Wiefling et al. [3] analyzed 31.3 million login attempts from 3.3 million Norwegian users, making it one of the strongest empirical studies in the corpus. Their work shows that RBA can reduce risk at a large scale, but also that performance depends heavily on parameter tuning and operational constraints. The difference between a controlled dataset and a production login stream is therefore not just scale. It is also messy: users change phones, use corporate VPNs, travel, clear cookies, share devices, and log in from networks with unstable geolocation.

Deployment measurement studies show that adoption remains limited. Gavazzi et al. [8] audited 208 popular websites and found that only 42.31% supported any form of MFA and only 22.12% blocked an obvious account-hijacking attempt. Wiefling et al. [4] found that deployed RBA systems behave inconsistently, with challenge frequency varying across services and over time. Makowski and Pöhn [18] similarly found inconsistent behavior across major platforms.

Implementation cost is another barrier. RBA systems must collect signals, store history, compute risk scores, select challenges, manage recovery, and avoid becoming denial-of-service targets. Wiefling et al. [3] reported

that unoptimized RBA can take over five seconds per login, making performance a real operational concern. Unsel et al. [26] showed that even implementing a known RBA algorithm in OpenStack required substantial engineering work. This suggests that broader adoption will require practical reference implementations, safer defaults, and clearer administrative tooling.

Deployment also depends on organizational maturity. Omotade et al. [31] frame authentication as one component of a broader information security management system. From that perspective, RBA is not just a model or a login feature. It depends on governance, incident response, logging, account recovery, user support, privacy policy, and security monitoring. Services that lack those supporting processes may struggle to deploy RBA responsibly even if the underlying model is available.

5.3 Theme Three: Security, Usability, and Privacy

RBA succeeds only if users and administrators can understand and tolerate it. Wiefling et al. [9] found that RBA can be more usable than always-on 2FA, but also that users who are never challenged may not perceive any added security. Jan and Khan [29] similarly found that users prefer stronger authentication for sensitive services, even when it adds friction. This supports an important design principle: users do not evaluate authentication only by convenience. They also consider whether the friction level matches the account's sensitivity.

Notification design is a major weakness. Wei et al. [12] found that 46% of users suspected a legitimate RBA notification email was phishing, and 89% preferred receiving security alerts on their phones. They also found that 56% preferred enabling 2FA after an alert, while 34.4% considered password changes too burdensome. Poorly designed alerts can therefore reduce the value of accurate detection by confusing users at the moment they need to respond.

RBA can also create recovery deadlocks. If a user must access an email account to retrieve a code, but that email account also triggers RBA, the user may be locked out of both services [9]. This shows that step-up authentication and recovery workflows must be designed together rather than treated as separate features. A secure challenge is not useful if legitimate users cannot complete it during realistic failure scenarios.

Administrator behavior is equally important. Markert et al. [13] showed that administrators can misunderstand RBA threshold controls and rely on defaults without fully understanding their consequences. This matters because thresholds determine the practical balance between security and friction. A poor administrative interface can produce insecure deployments even when the underlying detection model is reasonable.

Privacy remains a central concern. Effective RBA often depends on device fingerprints, geolocation, IP history, behavioral traces, and other sensitive signals [11]. These data can protect accounts, but they also create tracking and retention risks. Wiefling et al. [3] showed that validated round-trip time may substitute for some location data, reducing reliance on precise geolocation while making VPN-based spoofing harder. More work is needed on privacy-preserving signals that protect accounts without creating unnecessary surveillance infrastructure.

6 Implications and Future Work

The literature shows that RBA is technically promising, but future work should focus less on isolated model improvement and more on deployable, measurable, and privacy-conscious systems.

First, RBA needs stronger production-scale validation. Future studies should evaluate sys-

tems under realistic class imbalance, long-term behavior change, account recovery scenarios, travel, VPN use, device changes, accessibility needs, shared-device environments, and adversarial adaptation. A model that performs well on a small behavioral dataset may fail when deployed across millions of heterogeneous users.

Second, the field needs shared evaluation standards. Current papers report accuracy, FAR, FRR, G-Mean, re-authentication rate, or usability outcomes inconsistently. A stronger standard would report both security and friction: attacker-blocking rate, legitimate-user challenge rate, false accept rate, false reject rate, computational cost, recovery failure rate, and user comprehension of alerts. This would make it easier to determine whether one system is genuinely better than another, rather than optimized for a different metric.

Third, privacy must be treated as a design constraint. RBA systems should minimize data collection, explain what signals are used, limit retention, and explore less invasive alternatives such as coarse-grained network features, validated latency, anonymized behavioral profiles, or privacy-preserving computation. This is especially important in regulated environments such as healthcare, finance, and education.

Fourth, RBA needs a better human-facing design. Users must recognize legitimate alerts, understand why a challenge occurred, and recover safely when step-up authentication fails. Administrators need clearer threshold controls, safer defaults, and interfaces that communicate the effect of configuration choices. Without better tooling, even accurate RBA systems can be weakened by misconfiguration or user mistrust.

Finally, future work should treat RBA as part of a broader authentication ecosystem. It interacts with passwords, MFA, password reset, account recovery, fraud detection, session management, and Zero Trust access control. Studying RBA only at the login prompt misses many of the places where authentication succeeds or fails in practice.

7 Conclusion

Passwords remain dominant because they are familiar and easy to deploy, but they are not sufficient against modern credential-based attacks. RBA improves password-based authentication by adding contextual judgment: routine logins can remain low-friction, while suspicious attempts receive stronger scrutiny.

This review of 35 sources supports three conclusions. First, RBA is a technically credible approach, and machine-learning methods can improve risk scoring compared with simpler probabilistic baselines. Second, adoption remains far behind capability; many popular websites still lack meaningful suspicious-login detection, and deployed systems behave inconsistently. Third, the RBA's future depends on more than classifier accuracy. Privacy exposure, confusing notifications, recovery deadlocks, administrator misconfiguration, and inconsistent evaluation standards all limit real-world effectiveness.

RBA should therefore be understood not as a password replacement, but as a bridge between static authentication and continuous, context-aware security. Its value lies in balancing protection and friction. That balance is fragile, and the next phase of research should focus on proving that RBA can be deployed at scale in ways that are measurable, privacy-preserving, and trusted by users.

References

[1] Bonneau J, Herley C, van Oorschot PC, Stajano F. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In: 2012 IEEE Symposium on Security and Privacy. IEEE; 2012. p. 553-67.

[2] Innocenti T, Mirheidari SA, Kharraz A, Crispo B, Kirda E. You've Got (a Reset) Mail: A Security Analysis of Email-Based Password Reset Procedures. In: Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA 2021). Springer, Cham; 2021. p. 1-20.

[3] Wiefling S, Jørgensen PR, Thunem S, Iacono LL. Pump Up Password Security! Evaluating and Enhancing Risk-Based Authentication on a Real-World Large-Scale Online Service. ACM Transactions on Privacy and Security. 2022;26(1).

[4] Wiefling S, Dürmuth M, Iacono LL. What's in Score for Website Users: A Data-Driven Long-Term Study on Risk-Based Authentication Characteristics. In: Financial Cryptography and Data Security. Springer Berlin Heidelberg; 2021. p. 361-81.

[5] Al-Rumaim A, Pawar JD. Enhancing User Authentication: An Approach Utilizing Context-Based Fingerprinting With Random Forest Algorithm. IEEE Access. 2024;12:110850-61.

[6] Rotter D, Schwabe T, Dürmuth M. That's not you! Applying Neural Networks to Risk-Based Authentication to Detect Suspicious Logins. In: Proceedings of the 18th ACM Workshop on Artificial Intelligence and Security (AISec). Association for Computing Machinery; 2025. p. 100-10.

[7] Picard C, Pierre S. RLAuth: A Risk-Based Authentication System Using Reinforcement Learning. IEEE Access. 2023;11:61129-43.

[8] Gavazzi A, Williams R, Kirda E, Lu L, King A, Davis A, et al. A Study of Multi-Factor and Risk-Based Authentication Availability. In: 32nd USENIX Security Symposium (USENIX Security 23). USENIX Association; 2023. p. 2043-60.

[9] Wiefling S, Dürmuth M, Iacono LL. Verify It's You: How Users Perceive Risk-Based Authentication. IEEE Security & Privacy. 2021;19(6):47-57.

[10] Wiefling S, Dürmuth M, Iacono LL. More Than Just Good Passwords? A Study on Usability and Security Perceptions of Risk-based Authentication. In: Proceedings of the 36th Annual Computer Security Applications Conference (ACSAC). Association for Computing Machinery; 2020. p. 203-18.

[11] Wiefling S, Tolsdorf J, Iacono LL. Privacy Considerations for Risk-Based Authentication Systems. In: 2021 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW); 2021. p. 320-7.

[12] Wei T, Wang D, Li Y, Wang Y. Who is Trying to Access My Account? Exploring User Perceptions and Reactions to Risk-based Authentication Notifications. In: Network and Distributed System Security Symposium (NDSS); 2025. .

[13] Markert P, Schnitzler T, Golla M, Dürmuth M. How Administrators Configure Risk-based Authentication. In: Proceedings of the 18th Symposium on Usable Privacy and Security (SOUPS). USENIX Association; 2022. .

[14] Wiefling S, Hönscheid M, Iacono LL. A Privacy Measure Turned Upside Down? Investigating the Use of HTTP Client Hints on the Web. In: Proceedings of the 19th International Conference on Availability, Reliability and Security (ARES). Association for Computing Machinery; 2024. .

[15] Lin X, Ilia P, Solanki S, Polakis J. Phish in Sheep's Clothing: Exploring the Authentication Pitfalls of Browser Fingerprinting. In: 31st USENIX Security Symposium (USENIX Security 22). USENIX Association; 2022. p. 1651-68.

[16] Papaioannou M, Pelekoudas-Oikonomou F, Mantas G, Serrelis E, Rodriguez J, Fengou MA. A Survey on Quantitative Risk Estimation Approaches for Secure and Usable User Authentication on Smartphones. Sensors. 2023;23(6).

[17] Wiefling S, Patil T, Dürmuth M, Iacono LL. Evaluation of Risk-Based Re-Authentication Methods. In: ICT Systems Security and Privacy Protection (IFIP SEC). Springer International Publishing; 2020. p. 280-94.

[18] Makowski JP, Pöhn D. Evaluation of Real-World Risk-Based Authentication at Online Services Revisited: Complexity Wins. In: Proceedings of the 18th International Conference on Availability, Reliability and Security (ARES 2023); 2023. p. 73:1-73:9.

[19] Karim NA, Khashan OA, Kanaker H, Abdulraheem WK, Alshinwan M, Al-Banna AK. Online Banking User Authentication Methods: A Systematic Literature Review. IEEE Access. 2024;12:741-57.

[20] Qbea'h M, Alneyadi A, Alkaabi J, Alderei M, Almansouri S. Classification of Authentication Approaches to Stop the Next Breaking: Challenges, Benefits, Drawbacks, Awareness, and Recommendations. In: 2023 International Conference on Computer and Applications (ICCA); 2023. p. 1-5.

[21] Papaioannou M, Zachos G, Mantas G, Rodriguez J. Novelty Detection for Risk-based User Authentication on Mobile Devices. In: 2022 IEEE Global Communications Conference (GLOBECOM); 2022. p. 837-42.

[22] Ashibani Y, Mahmoud QH. An Intelligent Risk-Based Authentication Approach for Smartphone Applications. In: 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC); 2020. p. 3807-12.

[23] Papaioannou M, Sucasas V, Mantas G, Aaraj N, Essop A, Rodriguez J. Risk Estimation for a Secure and Usable User Authentication Mechanism for Mobile Passenger ID Devices. In: 2022 IEEE Global Communications Conference (GLOBECOM); 2022. .

[24] Liu J, Liu R, Lai Y. Risk-Based Dynamic Identity Authentication Method Based on the UCON Model. Security and Communication Networks. 2022;2022.

[25] Vajpayee P, Hossain G. Multi-Tenant Cloud Security-Risk Prediction through Cyber-Value-at-Risk (CVaR). In: 2024 12th International Symposium on Digital Forensics and Security (ISDFS); 2024. p. 1-7.

[26] Unsel V, Wiefling S, Gruschka N, Iacono LL. Risk-Based Authentication for OpenStack: A Fully Functional Implementation and Guiding Example. In: Proceedings of the Thirteenth ACM Conference on Data and Application Security and Privacy (CODASPY). Association for Computing Machinery; 2023. p. 237-43.

[27] Büttner A, Pedersen AT, Wiefling S, Gruschka N, Iacono LL. Is It Really You Who Forgot the Password? When Account Recovery Meets Risk-Based Authentication. In: Ubiquitous Security. Springer Nature Singapore; 2024. p. 401-19.

[28] Kumar P. Next-Generation Secure Authentication and Access Control Architectures: Advanced Techniques for Securing Distributed Systems in Modern Enterprises. International Journal of Computational and Experimental Science and Engineering. 2025;11(3):4966-95.

[29] Jan SU, Khan HU. Identity and Aggregate Signature-Based Authentication Protocol for IoD Deployment Military Drone. IEEE Access. 2021;9:130247-63.

[30] Banerjee J, Singh D. Adaptive Two-Factor Authentication using AI-based Risk Scoring and MAC Address Verification. In: 2025 6th International Conference on Electronics and Sustainable Communication Systems (ICESC); 2025. p. 1140-5.

[31] Omotade AL, Ghimire A, Sultan M. Cybersecurity, Data Privacy, and Information Security Management. In: 2025 3rd International Conference on Artificial Intelligence and Automation Control (AIAC); 2025. p. 27-33.