

CASA-IoT: Conformal Access Screening and Abstention for Managed IoT Systems Under Identity, Session, and Communication Uncertainty

Arnold Gaumond

Arnold@outlook.com

oklahoma

Method Article

Keywords: IoT security, adaptive authentication, conformal prediction, selective classification, biometric identity, access control, uncertainty quantification

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11394194/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

CASA-IoT: Conformal Access Screening and Abstention for Managed IoT Systems Under Identity, Session, and Communication Uncertainty^{*}

Arnold Gaumond

ARTICLE INFO

Keywords:

IoT security
adaptive authentication
conformal prediction
selective classification
biometric identity
access control
uncertainty quantification

ABSTRACT

Managed Internet of Things (IoT) interfaces increasingly combine biometric identity evidence, cryptographic session state, role permissions, device posture, and communication-security signals. A conventional classifier can still create an operational problem: every request is forced into an allow or deny decision even when the observed evidence is poorly supported by the calibration distribution. This paper presents CASA-IoT, a lightweight selective-access method that adds class-conditional conformal prediction to a probabilistic access-risk classifier. CASA-IoT maps a singleton benign prediction set to *allow*, a singleton attack set to *deny*, and ambiguous or empty sets to *step-up verification*. The method is intentionally modular: upstream biometric, QR/token, device, and channel-monitoring subsystems remain separate and provide normalized evidence to the access layer. A reproducible synthetic experiment was repeated across five seeds. Each development campaign contained 15,000 sessions spanning benign operation, four known attack families, and a blended family excluded from model fitting and calibration. CASA-IoT achieved AUROC 0.952 ± 0.004 , attack containment 0.926 ± 0.005 , and benign friction 0.176 ± 0.006 . On the held-out blended family, only $5.4\% \pm 0.5\%$ of sessions were allowed, while $87.2\% \pm 1.5\%$ were denied and the remainder were deferred to step-up verification. Known-class conformal coverage was 0.905 ± 0.007 at nominal 90% coverage. Under a shifted stress regime, fixed calibration became overly conservative; a small labeled calibration refresh reduced benign friction from 0.614 to 0.322 while retaining attack containment of 0.904. The results support conformal abstention as a practical policy layer, while also showing that calibration must be revisited under material distribution shift.

1. Introduction

Administrative and maintenance interfaces for managed IoT systems are rarely protected by a single piece of evidence. A request may arrive with a biometric match score, a cryptographically protected session token, a role assignment, a device-posture state, and a communication-health signal. Digital-identity guidance treats authentication as a lifecycle problem rather than a one-time password check [1], while zero-trust architecture rejects persistent implicit trust based only on network location or previous approval [2]. Risk-based authentication similarly uses contextual signals to decide when additional verification is appropriate [3]. These principles suggest that a practical IoT access layer should be able to combine heterogeneous evidence without assuming that any one upstream control is infallible.

The authorization boundary remains equally important. Role-based access control (RBAC) provides a mature basis for limiting actions according to organizational responsibilities [4], but authentication and authorization do not solve uncertainty by themselves. A classifier may still be forced to assign a definitive label to a request that differs substantially from its development distribution. This is especially problematic when the cost of an incorrect *allow* decision is asymmetric relative to the inconvenience of requesting an additional verification step. Authentication research has repeatedly emphasized the security–usability trade-off rather than treating maximum rejection as an acceptable objective [5].

This paper proposes **CASA-IoT** (Conformal Access Screening and Abstention for Managed IoT), a deliberately small method for converting probabilistic access evidence into three operational actions: *allow*, *step-up*, and *deny*. The core idea is selective classification: the model should make a direct decision only when the calibrated evidence supports one class strongly enough; otherwise it should abstain and request an additional control. This principle is related to

^{*} Preprint Version 2 – 15 October 2026. This version reports the full simulation protocol, five-seed results, a held-out blended-attack test, and a distribution-shift calibration study.

ORCID(s):

selective-classification theory [6], but the present study evaluates it specifically as an IoT access-control policy using identity, session, device, and channel evidence.

The contributions are fourfold:

- a compact evidence model separating biometric identity, session integrity, authorization context, device posture, and communication health;
- a class-conditional conformal access rule that translates prediction-set structure into allow/step-up/deny actions rather than relying on a single probability threshold;
- a held-out blended-attack evaluation in which the attack family is excluded from fitting and calibration; and
- a distribution-shift experiment showing both the failure of stale calibration and the partial recovery obtained through a small labeled calibration refresh without retraining the underlying classifier.

The contribution is intentionally moderate. CASA-IoT does not replace a biometric recognition system, token protocol, RBAC engine, device-security monitor, or communication-security mechanism. It is a decision layer for combining their outputs and expressing uncertainty operationally.

2. Background and Related Work

2.1. Biometric identity and session assurance

Biometric authentication can contribute strong identity evidence, but its score should not be interpreted as a complete authorization decision. Zita and Khalil presented a fingerprint-recognition pipeline that combines engineered fingerprint features, LightGBM classification, and public/private key generation linked to recognized identity [7]. That work motivates the use of biometric confidence and key-binding evidence in CASA-IoT; the present study does not reproduce its fingerprint pipeline or derive cryptographic keys from synthetic images.

The broader biometric literature has long treated recognition accuracy, template protection, revocability, and privacy as distinct concerns. Maltoni et al. provide a comprehensive treatment of fingerprint-recognition design [8]; Ratha et al. discuss attacks and countermeasures across biometric authentication stages [9]; and Jain et al. survey biometric-template security requirements and protection approaches [10]. Accordingly, CASA-IoT treats the biometric subsystem as an upstream source of normalized evidence rather than a trusted root that overrides all other signals.

Session integrity provides a separate trust dimension. JSON Web Tokens define compact claim containers [11], while RFC 8725 documents implementation practices required to avoid common JWT failures such as algorithm confusion and insufficient validation [12]. Zita's QR-based login study combines short-lived cryptographic session material, replay controls, and RBAC for web applications [13]. Modern public-key authentication mechanisms such as WebAuthn similarly bind scoped credentials to authenticators while maintaining a distinct server-side authorization boundary [14]. CASA-IoT abstracts these implementation details into token freshness, nonce validity, key binding, and role-fit variables.

2.2. Uncertainty, selective decisions, and unsupported inputs

Probability estimates are not automatically reliable measures of uncertainty. Guo et al. showed that modern classifiers can be miscalibrated even when their accuracy is high [15], and Ovadia et al. demonstrated that predictive uncertainty can degrade substantially under dataset shift [16]. Conformal prediction provides a complementary framework in which a calibrated nonconformity score is converted into set-valued predictions with finite-sample coverage guarantees under exchangeability assumptions [17, 18].

The distinction between known and unsupported behavior also appears in open-set recognition and out-of-distribution research. Scheirer et al. formalized open-set recognition as a setting in which test inputs may lie outside the classes represented during training [19], while Hendrycks and Gimpel showed that confidence statistics can provide a useful baseline for identifying misclassified and out-of-distribution samples [20]. CASA-IoT does not claim formal zero-day recognition. Instead, it uses prediction-set ambiguity as an operational reason to abstain.

This uncertainty-aware decision philosophy has also appeared in quantum-security monitoring. Zita's QKD-XGuard preprint applies physics-informed residuals, conformal prediction, and explicit rejection to decoy-state BB84 monitoring [21]. CASA-IoT addresses a different problem—administrative access rather than QKD attack monitoring—but adopts the same high-level caution that unsupported observations should not automatically be forced into a confident known label.

2.3. IoT operating context and cyber-resilience

IoT security is complicated by heterogeneous devices, distributed communication, limited resources, and long operational lifecycles. Sicari et al. synthesize security, privacy, and trust challenges in IoT environments [22]; NISTIR 8259A defines a core baseline of device cybersecurity capabilities [23]; Roman et al. analyze security and privacy challenges in distributed IoT architectures [24]; and Mosenia and Jha survey attack surfaces and countermeasures across the IoT stack [25]. These sources motivate the device-posture and channel-health variables used in the present method.

Operational context is also relevant because an additional challenge or restriction has a cost. Autonomic-computing research established the principle that systems can adjust behavior according to changing state while preserving higher-level objectives [26]. Zita and Zita later proposed an explainable and energy-aware self-adaptive IoT architecture using Deep Q-Learning, SHAP, and LIME [27]. CASA-IoT does not perform reinforcement learning; battery and latency are used only as contextual variables in the synthetic operating regime, and security thresholds are never lowered solely because resources are constrained. This preserves the cyber-resilience principle that continued operation must remain bounded by explicit security objectives [28].

3. CASA-IoT Method

3.1. Evidence vector

Each access session is represented by

$$\mathbf{x} = [b, k, t, n, r, d, c, e, \ell]^T, \quad (1)$$

where b is biometric confidence, k is key-binding confidence, t is token freshness, n is nonce/replay validity, r is role-action fit, d is device security posture, c is communication-channel health, e is normalized battery state, and ℓ is normalized latency quality. All variables are scaled to $[0, 1]$ with larger values representing more favorable evidence.

The first seven coordinates are security evidence. Battery and latency are included to expose the model to realistic operating variation and to test whether calibration remains stable under resource shift. They do not independently authorize a sensitive action.




Figure 1: CASA-IoT separates upstream evidence generation from calibrated access decisions. The method does not replace the biometric, session, RBAC, device, or channel subsystems.

3.2. Probabilistic attack model

A logistic model estimates

$$\hat{p}(y = 1 \mid \mathbf{x}) = \sigma(\mathbf{w}^T \tilde{\mathbf{x}} + \beta), \quad (2)$$

where $y = 1$ denotes malicious or policy-violating access, $\tilde{\mathbf{x}}$ is standardized using training data only, and $\sigma(\cdot)$ is the logistic function. Logistic regression was selected deliberately because the study focuses on the abstention policy rather than on maximizing predictive complexity.

For reference, ordinary probability-threshold access uses

$$A_p(\mathbf{x}) = \begin{cases} \text{allow,} & \hat{p}_1 \leq \tau_L, \\ \text{step-up,} & \tau_L < \hat{p}_1 < \tau_H, \\ \text{deny,} & \hat{p}_1 \geq \tau_H. \end{cases} \quad (3)$$

This baseline depends directly on the numerical probability scale, which may become unreliable under shift.

3.3. Class-conditional conformal screening

Let C_k be the held-out calibration set for class $k \in \{0, 1\}$. For calibration observation i with true label y_i , the nonconformity score is

$$a_i = 1 - \hat{p}_{y_i}(\mathbf{x}_i). \quad (4)$$

For a new request $\mathbf{x}$ and candidate class k ,

$$a_k(\mathbf{x}) = 1 - \hat{p}_k(\mathbf{x}), \quad (5)$$

and the class-conditional conformal p-value is

$$p_k(\mathbf{x}) = \frac{1 + \sum_{i \in C_k} \mathbb{I}[a_i \geq a_k(\mathbf{x})]}{|C_k| + 1}. \quad (6)$$

At significance level $\alpha = 0.10$,

$$\Gamma_{0.10}(\mathbf{x}) = \{k : p_k(\mathbf{x}) > 0.10\}. \quad (7)$$

The operational decision is

$$A_c(\mathbf{x}) = \begin{cases} \text{allow,} & \Gamma = \{0\}, \\ \text{deny,} & \Gamma = \{1\}, \\ \text{step-up,} & \Gamma \in \{\emptyset, \{0, 1\}\}. \end{cases} \quad (8)$$

Thus uncertainty is not converted into a falsely precise decision. The conformal coverage statement applies to exchangeable samples from the calibrated classes; it is not a theorem about unseen attack detection.

4. Experimental Methodology

4.1. Synthetic session model

The study uses a synthetic generator so that identity, session, authorization, device, communication, and resource evidence can be varied independently and reproducibly. No human-subject, biometric-image, or private operational data are used. Development-regime benign sessions draw security-evidence variables from high-confidence beta distributions, while battery and latency have wider but still favorable distributions. The stress regime lowers battery and latency substantially and moderately broadens several security variables to represent a deployment change rather than a simple repetition of the training environment.

Five attack scenarios are generated. Four are available during model development, while the fifth is held out entirely from fitting and calibration. Table 1 summarizes the intended perturbations.

Each seed produces 15,000 development sessions: 5,000 benign sessions and 2,000 sessions for each of the five attack families. The held-out blended family is reserved for evaluation. For the remaining classes, 55% of sessions are used for training, 20% for conformal calibration, and 25% for known-class testing. The final development test therefore contains 3,250 known-class sessions plus 2,000 unseen blended sessions.

A separate stress set contains 8,000 sessions: 3,000 benign and 1,000 from each attack family. The original classifier and calibration data are applied without modification in the first stress experiment. A second experiment performs a *calibration refresh*: the trained classifier is kept fixed, while conformal score distributions are rebuilt from a small labeled sample of 400 benign sessions and 150 sessions from each known attack family in the shifted regime. The blended family remains excluded.

The complete experiment is repeated for seeds 17, 43, 101, 307, and 911. NIST risk-assessment guidance motivates reporting both undesirable acceptance and excessive operational restriction rather than optimizing one metric in isolation [29].

Table 1
Synthetic access scenarios.

Scenario	Primary evidence change
Benign	High biometric, key-binding, token, nonce, role, posture, and channel evidence.
Biometric spoof	Depressed biometric and key-binding confidence with otherwise plausible session state.
Token replay	Reduced token freshness and nonce validity, with mild channel inconsistency.
Privilege misuse	Low role-action fit with moderate posture degradation despite plausible identity evidence.
Channel anomaly	Low communication-health evidence and modest device/session inconsistency.
Blended unseen	Moderate simultaneous degradation across identity, session, role, posture, and channel coordinates; excluded from fitting and calibration.

4.2. Baselines

CASA-IoT is compared with two deliberately simple alternatives:

1. **Probability thresholds:** the same logistic model with fixed thresholds $\tau_L = 0.35$ and $\tau_H = 0.65$.
2. **Fixed weighted rule:** a transparent weighted sum of unfavorable evidence, with step-up and deny thresholds fixed before the final evaluation.

Because CASA-IoT and the probability-threshold baseline share the same underlying classifier, differences in action behavior arise from the decision layer rather than model capacity.

4.3. Metrics

AUROC measures binary ranking quality. For access policy, the primary metrics are

$$\text{Attack containment} = P(A \neq \text{allow} \mid y = 1), \quad (9)$$

$$\text{Dangerous false accept} = P(A = \text{allow} \mid y = 1), \quad (10)$$

$$\text{Benign friction} = P(A \neq \text{allow} \mid y = 0), \quad (11)$$

$$\text{Benign deny} = P(A = \text{deny} \mid y = 0). \quad (12)$$

The step-up rate is reported separately because a challenge is operationally less severe than a hard denial. For known calibrated classes, empirical conformal coverage is also reported. Calibration quality is conceptually distinct from ranking quality; probabilistic scoring rules such as the Brier score illustrate why confidence quality deserves separate attention from accuracy [30].

5. Results

5.1. Development-regime performance

Table 2 reports mean $\pm$ standard deviation across five seeds. The three methods achieve similar ranking quality because the two learned policies use the same logistic classifier and the weighted rule uses the same evidence. Their action profiles differ substantially.

CASA-IoT denies a larger share of attacks than the probability-threshold policy (0.854 ± 0.014 versus 0.802 ± 0.005), while using step-up verification less frequently overall. The weighted rule contains a similar fraction of attacks but does so primarily by routing a large portion of sessions to step-up verification rather than by confidently distinguishing them. These results do not establish that conformal screening is universally superior; they show that prediction-set semantics produce a different and interpretable intervention profile with the same base classifier.

Known-class empirical conformal coverage is 0.905 ± 0.007 at nominal 90% coverage, and 0.927 ± 0.012 of known-class requests produce singleton prediction sets. This is close to the intended operating point and provides a calibration check that a fixed probability threshold does not directly offer.

Table 2

Development-regime results across five seeds. Lower dangerous false accept and benign friction are better.

Method	AUROC	Attack containment	Benign friction	Benign deny
CASA-IoT	0.952 ± 0.004	0.926 ± 0.005	0.176 ± 0.006	0.098 ± 0.009
Probability thresholds	0.952 ± 0.004	0.932 ± 0.004	0.184 ± 0.006	0.069 ± 0.007
Fixed weighted rule	0.943 ± 0.003	0.935 ± 0.004	0.225 ± 0.013	0.004 ± 0.001

Table 3

Stress-regime results. The calibration refresh keeps the trained classifier fixed and updates only conformal score distributions using a small labeled shifted-regime sample.

Method	AUROC	Attack containment	Benign friction
CASA-IoT, stale calibration	0.897 ± 0.005	0.976 ± 0.005	0.614 ± 0.033
CASA-IoT + calibration refresh	0.897 ± 0.005	0.904 ± 0.008	0.322 ± 0.021
Probability thresholds	0.897 ± 0.005	0.978 ± 0.005	0.626 ± 0.030
Fixed weighted rule	0.884 ± 0.005	0.997 ± 0.001	0.908 ± 0.007

5.2. Held-out blended attack

The blended family is intentionally more difficult than the single-mechanism attacks because no one coordinate is driven to an extreme value. It is never used to fit the logistic model or construct the development calibration sets. Across seeds, CASA-IoT allows 0.054 ± 0.005 of these requests, sends 0.074 ± 0.012 to step-up verification, and denies 0.872 ± 0.015 . Mean classifier attack probability is 0.819 ± 0.005 .

This result should be interpreted as a held-out stress test, not a proof of generic zero-day detection. The blended generator is still defined by the same nine-dimensional evidence space, so an attack that preserves all monitored signals could remain undetected.

5.3. Distribution shift and calibration refresh

Table 3 shows the main limitation of the method. When a shifted operating regime is presented without updating calibration, CASA-IoT becomes overly conservative: attack containment increases, but benign friction reaches 0.614 ± 0.033 . Probability thresholds exhibit a similar effect, and the fixed weighted rule is even more restrictive.

The refresh trades some containment for a large reduction in benign friction. This is consistent with the general observation that uncertainty estimates degrade under dataset shift [16]. Importantly, the refresh does not retrain the classifier and does not use the held-out blended family. It therefore illustrates a maintenance path for the decision layer, not a claim that calibration drift can be solved without trustworthy labeled data.

6. Discussion

The first useful property of CASA-IoT is architectural separation. A biometric recognizer, QR/token workflow, RBAC engine, device monitor, and communication-security monitor retain their own security assumptions. The access layer consumes normalized evidence instead of pretending that one mechanism establishes complete trust. This is particularly important for the four application domains represented by the related Zita studies: biometric-linked cryptographic identity [7], secure QR/RBAC session handling [13], uncertainty-aware QKD monitoring [21], and explainable resource-aware IoT adaptation [27]. CASA-IoT does not combine those systems into a single deployment; it uses them as examples of distinct evidence-producing subsystems that may exist around an administrative decision.

Second, abstention has a clear operational interpretation. Selective classification has long treated rejection as a legitimate response when the model is insufficiently confident [6]. In access control, that rejected region maps naturally to step-up verification rather than to an unspecified “unknown” label. This can make uncertainty actionable without granting access based on a weak score or automatically denying every atypical request.

Third, the stress experiment shows why calibration maintenance is a security requirement rather than a cosmetic machine-learning adjustment. Under shifted conditions, stale calibration became excessively restrictive. Such behavior may be safe in the narrow sense of preventing access, but it can undermine availability and operator trust. Cyber-resilient

engineering explicitly treats continued essential function, recovery, and adaptation as security objectives alongside prevention [28].

CASA-IoT should therefore be understood as a policy mechanism, not a universal authentication architecture. In production, the evidence sources would require independent assurance. A forged biometric score, compromised token-validation service, inaccurate device posture, or manipulated channel monitor could corrupt the input vector. Zero-trust architecture remains relevant precisely because evidence must be re-evaluated rather than permanently trusted [2].

7. Limitations

This study has five important limitations. First, all observations are synthetic. The distributions were constructed to create interpretable overlap among benign, known-attack, and held-out blended scenarios; they do not represent measured production frequencies. Second, the base classifier is logistic regression. This makes the role of the conformal layer easier to isolate, but nonlinear real-world interactions may require richer models. Third, the conformal coverage statement relies on exchangeability within the calibrated distribution; the stress experiment demonstrates the practical consequence when that assumption is weakened. Fourth, the held-out blended attack is not equivalent to an unconstrained adversary. Fifth, a calibration refresh requires trustworthy labeled samples from the changed environment, which may be difficult during an active incident.

The method also does not claim that biometric confidence proves identity, that a valid JWT proves authorization, or that communication-health evidence proves end-to-end security. Each upstream subsystem remains subject to its own attack surface and assurance requirements.

8. Reproducibility and Version Note

The accompanying `casa_iot_simulation.py` script generates the synthetic sessions, performs all five seeded experiments, fits the classifier, constructs class-conditional conformal sets, applies the baseline policies, performs the stress calibration refresh, and writes the CSV results used in Tables 2 and 3. The package includes the complete results across seeds, the held-out blended-attack results, the calibration summary for seed 17, and representative development and stress datasets.

This manuscript is explicitly labeled **Preprint Version 2**. The version label is bibliographic/versioning information; it does not imply that the present study has been peer reviewed or accepted by a journal.

9. Conclusion

CASA-IoT introduces a compact conformal decision layer for managed IoT access. Rather than treating a probabilistic classifier as an automatic allow/deny oracle, it converts calibrated prediction sets into three actions: allow when only benign behavior is supported, deny when only attack behavior is supported, and request step-up verification when the prediction set is ambiguous or empty. Across five synthetic development campaigns, the method achieved AUROC 0.952 ± 0.004 and contained 0.926 ± 0.005 of attack sessions with benign friction of 0.176 ± 0.006 . The held-out blended attack was allowed in only about 5.4% of sessions. The stress experiment also exposed the method's main weakness: stale calibration can become operationally over-conservative under distribution shift. A small labeled calibration refresh reduced that friction substantially without retraining the classifier. These results justify further evaluation on real authentication traces and instrumented IoT testbeds, but not a claim of production readiness.

Declarations

Ethics approval and consent to participate: Not applicable; the study uses only synthetic data.

Consent for publication: Not applicable.

Data and code availability: The simulation code and generated result files accompany this manuscript.

References

- [1] P. A. Grassi, M. E. Garcia, J. L. Fenton, Digital identity guidelines: Authentication and lifecycle management, Tech. Rep. NIST SP 800-63B, National Institute of Standards and Technology (2017). doi:10.6028/NIST.SP.800-63b.

- [2] S. Rose, O. Borchert, S. Mitchell, S. Connelly, Zero trust architecture, Tech. Rep. NIST SP 800-207, National Institute of Standards and Technology (2020). doi:10.6028/NIST.SP.800-207.
- [3] S. Wiefeling, L. Lo Iacono, M. Dürmuth, Is this really you? an empirical study on risk-based authentication applied in the wild, in: 34th IFIP TC-11 International Conference on Information Security and Privacy Protection, Vol. 562 of IFIP Advances in Information and Communication Technology, 2019, pp. 134–148. doi:10.1007/978-3-030-22312-0_10.
- [4] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, C. E. Youman, Role-based access control models, Computer 29 (2) (1996) 38–47. doi:10.1109/2.485845.
- [5] J. Bonneau, C. Herley, P. C. van Oorschot, F. Stajano, The quest to replace passwords: A framework for comparative evaluation of web authentication schemes, in: 2012 IEEE Symposium on Security and Privacy, 2012, pp. 553–567. doi:10.1109/SP.2012.44.
- [6] R. El-Yaniv, Y. Wiener, On the foundations of noise-free selective classification, Journal of Machine Learning Research 11 (2010) 1605–1641.
- [7] W. Zita, M. S. Khalil, Fingerprint-based cryptographic identity: A custom recognition pipeline with key pair generation, East Journal of Computer Science 1 (3) (2025) 17–34. doi:10.63496/ejcs.Vol1.Iss3.112.
- [8] D. Maltoni, D. Maio, A. K. Jain, S. Prabhakar, Handbook of Fingerprint Recognition, 2nd Edition, Springer, 2009. doi:10.1007/978-1-84882-254-2.
- [9] N. K. Ratha, J. H. Connell, R. M. Bolle, Enhancing security and privacy in biometrics-based authentication systems, IBM Systems Journal 40 (3) (2001) 614–634. doi:10.1147/sj.403.0614.
- [10] A. K. Jain, K. Nandakumar, A. Nagar, Biometric template security, EURASIP Journal on Advances in Signal Processing 2008 (2008) 579416. doi:10.1155/2008/579416.
- [11] M. Jones, J. Bradley, N. Sakimura, Json web token (jwt), Tech. Rep. RFC 7519, Internet Engineering Task Force (2015). doi:10.17487/RFC7519.
- [12] Y. Sheffer, D. Hardt, M. Jones, Json web token best current practices, Tech. Rep. RFC 8725, Internet Engineering Task Force (2020). doi:10.17487/RFC8725.
- [13] W. Zita, Secure qr code login system with role-based access for web applications, TechRxiv preprint (2025). doi:10.36227/techrxiv.175616741.16408346/v1.
- [14] World Wide Web Consortium, Web authentication: An api for accessing public key credentials level 2, W3C Recommendation, 8 April 2021 (2021).
- [15] C. Guo, G. Pleiss, Y. Sun, K. Q. Weinberger, On calibration of modern neural networks, in: Proceedings of the 34th International Conference on Machine Learning, Vol. 70, 2017, pp. 1321–1330.
- [16] Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. V. Dillon, B. Lakshminarayanan, J. Snoek, Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift, in: Advances in Neural Information Processing Systems, Vol. 32, 2019.
- [17] V. Vovk, A. Gammerman, G. Shafer, Algorithmic Learning in a Random World, 2nd Edition, Springer, 2022. doi:10.1007/978-3-031-06649-8.
- [18] A. N. Angelopoulos, S. Bates, A gentle introduction to conformal prediction and distribution-free uncertainty quantification, Foundations and Trends in Machine Learning 16 (4) (2023) 494–591. doi:10.1561/2200000101.
- [19] W. J. Scheirer, A. d. R. Rocha, A. Sapkota, T. E. Boult, Toward open set recognition, IEEE Transactions on Pattern Analysis and Machine Intelligence 35 (7) (2013) 1757–1772. doi:10.1109/TPAMI.2012.256.
- [20] D. Hendrycks, K. Gimpel, A baseline for detecting misclassified and out-of-distribution examples in neural networks, in: International Conference on Learning Representations, 2017.
- [21] W. Zita, Qkd-xguard: Physics-informed conformal open-set monitoring for unseen and adaptive attacks in decoy-state bb84, Preprint (2026).
- [22] S. Sicari, A. Rizzardi, L. A. Grieco, A. Coen-Porisini, Security, privacy and trust in internet of things: The road ahead, Computer Networks 76 (2015) 146–164. doi:10.1016/j.comnet.2014.11.008.
- [23] M. Fagan, K. N. Megawati, K. Scarfone, M. Smith, Iot device cybersecurity capability core baseline, Tech. Rep. NISTIR 8259A, National Institute of Standards and Technology (2020). doi:10.6028/NIST.IR.8259A.
- [24] R. Roman, J. Zhou, J. Lopez, On the features and challenges of security and privacy in distributed internet of things, Computer Networks 57 (10) (2013) 2266–2279. doi:10.1016/j.comnet.2012.12.018.
- [25] A. Mosenia, N. K. Jha, A comprehensive study of security of internet-of-things, IEEE Transactions on Emerging Topics in Computing 5 (4) (2017) 586–602. doi:10.1109/TETC.2016.2606384.
- [26] J. O. Kephart, D. M. Chess, The vision of autonomic computing, Computer 36 (1) (2003) 41–50. doi:10.1109/MC.2003.1160055.
- [27] W. Zita, N. Zita, Explainable and energy-aware self-adaptive iot architecture using deep q-learning, Sakarya University Journal of Computer and Information Sciences 9 (2) (2026) 364–376. doi:10.35377/SAUCIS...1736606.
- [28] R. Ross, V. Pillitteri, R. Graubart, D. Bodeau, R. McQuaid, Developing cyber-resilient systems: A systems security engineering approach, Tech. Rep. NIST SP 800-160 Vol. 2 Rev. 1, National Institute of Standards and Technology (2021). doi:10.6028/NIST.SP.800-160v2r1.
- [29] R. M. Blank, P. D. Gallagher, Guide for conducting risk assessments, Tech. Rep. NIST SP 800-30 Rev. 1, National Institute of Standards and Technology (2012). doi:10.6028/NIST.SP.800-30r1.
- [30] G. W. Brier, Verification of forecasts expressed in terms of probability, Monthly Weather Review 78 (1) (1950) 1–3. doi:10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2.