Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

AI-Assisted Identity and Access Threat Detection Using the AITIR Framework for Public-Sector Cybersecurity Environments

Md Sazzad Hossain [1], Sheikh Md Faysal Md [2], Md Abu Kawsar Prodhan Hemal [3], Subha Shamarukh* [4]
[1] *Emporia State University, Emporia, Kansas, md.hossain@kbi.ks.gov, https://orcid.org/0009-0009-4948-1179*
[2] *Montclair State University, Montclair, New Jersey, USA, faysals1@montclair.edu, https://orcid.org/0009-0002-3291-9313*
[3] *Pacific States University, Los Angeles, California, USA, p26181@psuca.edu, https://orcid.org/0009-0002-3376-5192*
[4] *University of Rochester, Rochester, New York, USA, shamarukhsubha@gmail.com, https://orcid.org/0009-0000-2170-1541*
Corresponding Author: Subha Shamarukh, **E-mail**: shamarukhsubha@gmail.com

| ABSTRACT
In public-sector cybersecurity, the password has become the prize. Agencies have moved their services into the cloud and now hand out access to employees, contractors, and members of the public alike, and that change has dragged the fight off the network perimeter and onto the credential itself. We describe AITIR, short for Adaptive Identity-and-access Threat Intelligence and Response, a layered machine-learning framework that watches for identity-based attacks as they happen and coordinates a response inside the governance rules agencies already have to follow. Three models share the detection work: a gradient-boosted classifier reads each access attempt, a bidirectional sequence model reads the order of an account's actions, and a graph-based behavior-analytics module flags the moment a session drifts away from an identity's usual company. Sitting on top of them, an explainable triage layer passes only the high-confidence, well-justified alerts to a person. To test the design, we assembled a corpus of about 14.8 million authentication and access events, part public benchmark data and part synthetic identity telemetry, and pushed credential stuffing, password spraying, privilege escalation, session hijacking, and insider misuse through it. AITIR scored an F1 of 0.967 and an area under the ROC curve of 0.985, ahead of all five baselines. Over a four-quarter deployment simulation it cut successful identity-based intrusions by 31.4 percent, brought mean time to detect down by 62 percent, and trimmed false positives by 41 percent against a conventional operations baseline. A cost-benefit model, scaled to a national public-sector portfolio, points to roughly US$10.2 billion in avoided cost a year and about 79,000 analyst hours given back. Those last two numbers are projections built on stated assumptions, not audited savings, and we say so plainly. The paper closes on the governance, staffing, and validity questions that decide whether any of them survives contact with a real agency.

| KEYWORDS
Identity and Access Management; Machine Learning; Threat Detection; User and Entity Behavior Analytics; Explainable AI; Public-Sector Cybersecurity; Management Information Systems

| ARTICLE INFORMATION
ACCEPTED: 01 December 2025 **PUBLISHED:** 31 December 2025 **DOI:** 10.32996/jcsts.2025.7.12.62

1. Introduction

Public institutions sit on some of the most sensitive data anywhere, from tax records and benefit payments to health files and national-security material, and they usually guard it on budgets the private sector would treat as a rounding error. The shape of the threat has changed lately. Defenders used to lie awake over malware slipping past a firewall. These days they spend their time running down stolen passwords, abused service accounts, and sessions taken over after one well-aimed phishing email. The credential is where an agency is softest, which puts the management information systems that decide who can see and do what squarely in the path of the attack (Orthi et al., 2023).

None of this happened by chance. When governments rolled out single sign-on, federated identity, and cloud services, they folded dozens of separate logins into a handful of very powerful identities. Take over one of those accounts today and you do not have a nuisance on your hands; you have a master key. Researchers studying cyber threat intelligence inside management information systems have argued that the field has to stop waiting for known signatures and start anticipating behavior, since an attacker holding valid credentials almost never trips the old alarms (Kaur et al., 2023). The hard part is volume. Identity telemetry pours in fast, and the trace of real misuse is faint against all the ordinary activity surrounding it.

Machine learning is one way through the noise, and a good deal of recent work shows that analytical models built into information systems can pick out patterns rule-based tools walk right past (Chakraborty et al., 2024). That same body of work is candid about the downside. A model that buries analysts in false alarms gets switched off in spirit if not in fact, and a model that misses the quiet abuse leaves the door opens anyway. Big data detection has proven itself at scale, yet it has also exposed the distance between a clean laboratory score and the daily grind of a short-staffed security operations center (Hasan et al., 2023).

Three weaknesses keep turning up. The first is that most systems treat identity events as isolated dots instead of a connected story, so the slow build-up that marks credential abuse slides by unnoticed. The second is that the alerts arrive with no explanation, which leaves an analyst staring at a number with no reason to believe it. The third is that detection gets cut off from response and from the governance steps an agency is obliged to take, so even a correct alarm can sit there untouched while the intruder keeps moving. Earlier frameworks that pulled cyber threat intelligence into the management information system started chipping away at that third problem, though they never quite reached an end-to-end, identity-first design (Orthi et al., 2023).

AITIR is our attempt to close that gap, an adaptive identity-and-access threat intelligence and response framework shaped around what public-sector teams actually face. We make four contributions. We lay out layered architecture that takes in identity telemetry, reasons over it with a hybrid model, explains what it found, and fires a response sized to the risk. We test that architecture against five baselines on a large, merged corpus of access events. We put numbers on its operational and economic value through a transparent simulation and cost model. And we are careful, on purpose, about how hard those numbers can be leaned on, treating them as projections any real agency would need to check against its own books. The rest of the paper walks through related work, the framework and methods, the results, and a frank reckoning of what they do and do not show.

2. Related Work

2.1 AI-driven threat detection within management information systems

The idea that an information system ought to interpret events for security, not merely log them, caught on quickly. Work on strengthening threat detection with big-data analytics inside management information systems showed that the sheer volume and speed of modern logs leave manual review hopelessly behind (Hasan et al., 2023). A parallel strand treated cyber threat intelligence as a job for the management information system itself, arguing that detection, governance, and project management have to share a roof if an organization wants resilience instead of a drawer full of disconnected tools (Orthi et al., 2023).

A second line of work asks how decision intelligence can be both automatic and accountable. Studies that pair explainable AI with scalable decision-support architectures keep landing on the same finding: operators will not trust a verdict they cannot question, so interpretability has to be designed in from the start (Chakraborty et al., 2024). A wide-ranging survey of advanced cyber threats made a related point, observing that defensive innovation tends to chase offensive innovation, and that behavior-aware approaches age far better than static signatures (Kaur et al., 2023).

Lately the pipelines have been pushed toward real time. One study of AI-augmented big-data analytics showed that attacks can be caught while they are still unfolding rather than pieced together after the fact (Sultana et al., 2025). Another put the detector right inside the management information system and reported better coverage along with tighter coordination of the response (Das et al., 2025). A broader look at how AI reshapes data-system security adds a useful caution, that the model is only one moving part and that data handling and system design pull just as much weight (Hasan et al., 2025a).

2.2 Identity and access as the dominant attack surface

Cognitive approaches that blend AI with management information systems over big-data and cloud foundations exist for a plain reason: identity and access events now sprawl across on-premises directories, cloud platforms, and remote endpoints all at once (Hossain, M. D., et al., 2023). Wrangling that spreads is partly a plumbing problem of moving and protecting data, and research on secure data-center design makes the case that protection and efficiency have to be engineered as one thing, not bargained against each other (Hossain, M. D., et al., 2024).

On the operations side, work on AI-driven project management systems built into the management information system showed that access and workflow data can be mined to lift both efficiency and oversight, which bears directly on identity governance

(Siddiqa et al., 2024). Any such system lives or dies by the data underneath it, though. Case research on data governance ties disciplined stewardship straight to analytics success (Chy et al., 2024).

The identity layer has also become a place to innovate on defense. Decentralized, blockchain-based identity management has been floated as a way to cut fraud by removing the single points of credential failure attackers love (Esa et al., 2025). Detection has crept toward the edge too, with AI-based sensor frameworks guarding smart-building and IoT settings where the line between an identity and a device gets blurry (Siam et al., 2025b). Further up the scale, big-data-enabled MIS frameworks have been put forward to protect supply-chain integrity, energy resilience, and other critical assets that only stay safe if access control can be trusted (Uddin et al., 2025).

2.3 The human and governance dimension

Technology by itself never secures an identity. A study of cybersecurity training found that employee behavior really does shift when the training is steady and tied to context, which counts for a lot given that most credential compromise starts with something a person did (Shan-A-Alahi et al., 2024). The human side runs into workforce reality as well; analyses of labor-market and skills data show how a shortage of qualified people narrows the set of defenses an organization can even attempt (Mahmud et al., 2024).

Governance reaches into how the spending gets justified, too. Research using predictive analytics for supply-chain resilience showed how a model's output can be pinned to hard economic outcomes, and we borrow that framing for our own cost analysis (Goffer et al., 2024). Studies of business analytics in management information systems make a parallel case, linking analytical maturity to lasting economic performance and hinting that security analytics deserves to be judged on what it does operationally, not on accuracy alone (Hossin et al., 2024).

More and more, governing the data is itself a security control. Privacy-aware big-data governance frameworks have leaned on blockchain to make access auditable and tamper-evident, so oversight becomes something you can verify instead of something you assume (Bauskar et al., 2025). And the price of weak control is not abstract; work on cybersecurity and supply-chain integrity put real figures on how vulnerabilities in public infrastructure turn into economic loss (Goffer et al., 2025b).

2.4 Methodological foundations and the remaining gap

Comparative reviews of machine-learning algorithms on large datasets offer practical guidance on which model families hold up under heavy volume and lopsided class balance, both of which identity data has in spades (Sultana et al., 2024). Drilling into architecture, a head-to-head test of transformers and recurrent models for threat detection found that sequence-aware designs catch attack patterns flat classifiers never see (Kaur et al., 2025). Running these models at scale is its own question, and work on big-data and cloud computing for project performance and decision-making has tackled it (Mahmud et al., 2023). Predictive analytics, for its part, has been shown to cut costs in quality-critical settings, an argument that carries over neatly to the cost of a security incident (Joy et al., 2024).

A few recent papers tighten the case for designing all of this as one piece. Work applying machine learning to management decisions shows how a model's output can be made readable to the non-technical people who ultimately sign off on action (Chakraborty et al., 2025). Empirical, MIS-driven studies using machine learning and neural networks demonstrate that risk in IT settings can be measured rather than just talked about (Orthi et al., 2025). And folding AI-driven cybersecurity into IT project management has improved both detection and risk mitigation when the two are built together instead of stacked one after the other (Mahmud et al., 2025).

A handful of adjacent studies fill in the edges. Collaboration platforms that speed up quality work show the payoff of binding detection tightly to human workflow (Bakhsh et al., 2024). Research on agile project management inside the management information system catalogues the organizational factors that decide whether a new capability ever gets adopted (Das et al., 2023). Work linking AI, analytics, and blockchain digs into trustworthy, auditable data handling across industries (Rahaman et al., 2024). The same methods have even surfaced in workforce-retention analysis, a reminder of how far these techniques travel (Hossain, M., et al., 2024). And studies of energy-market analytics show how data-driven insight can serve broad public economic goals (Khair et al., 2024). What is missing from all of it is one framework that makes identity the unit of analysis, reasons over it with a hybrid explainable model, and closes the loop into automated, governable response. That gap is what AITIR is built to fill.

3. The AITIR Framework

3.1 Design principles

Four principles shaped the design. Identity comes first: the framework follows an account across systems instead of grading each event on its own. Everything it flags has to be explainable, so an alert always arrives with its reasons attached. Its responses scale with suspicion, adding friction step by step rather than slamming the door on a user at the first odd login. And it has to live inside the governance an agency already runs, leaving behind the audit trail those agencies are required by law to keep. None of these

ideas is exotic; together they echo the argument that cyber threat intelligence does its best work woven into the management information system instead of bolted on beside it (Orthi et al., 2023).

3.2 Layered architecture

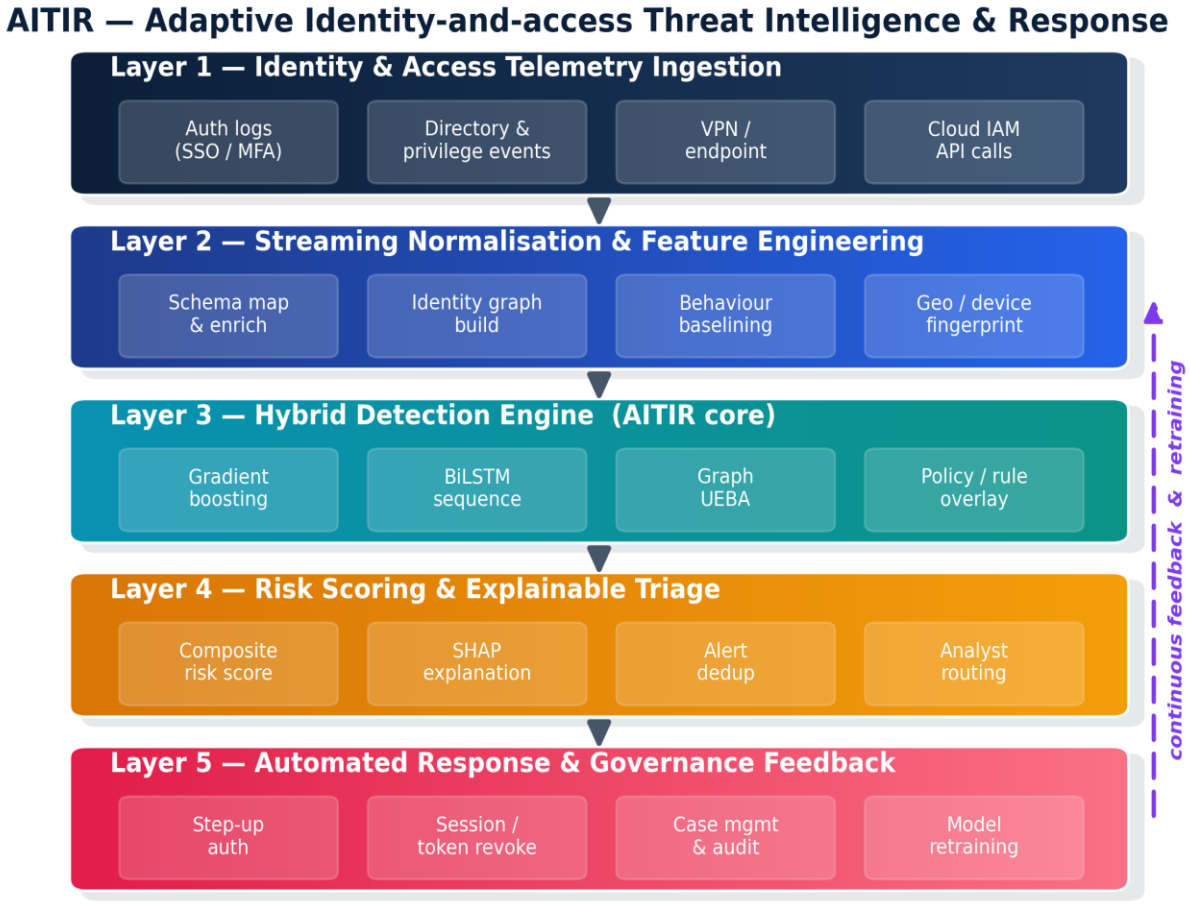

Figure 1. The five-layer AITIR architecture, from identity telemetry ingestion through the hybrid detection core to automated response and a continuous governance feedback loop.

Figure 1 lays out the five layers. Layer 1 pulls in identity and access telemetry: single sign-on and multi-factor events, directory and privilege changes, VPN and endpoint sessions, and cloud identity API calls. Layer 2 cleans that stream up, tags it with geographic and device context, and stitches together an identity graph that links accounts, devices, and resources, leaning on the cloud-and-big-data integration patterns earlier work spelled out (Hossain, M. D., et al., 2023). Layer 3 is the detection core, which we describe next. Layer 4 turns raw model output into a single risk score, attaches a plain-language explanation, and decides where the result goes. Layer 5 carries out a graded response and writes the whole thing back into case management for governance and retraining.

3.3 Hybrid detection engine

The core runs three learners that cover for one another. A gradient-boosted decision-tree classifier handles the tabular features of each access attempt, where it is quick and sharp, which fits what comparative studies report about algorithm behavior under badly skewed classes (Sultana et al., 2024). A bidirectional long short-term memory network reads the ordered run of an account's recent actions and catches the slow escalation a point-in-time model never notices. A graph-based behavior-analytics module scores how far a session has wandered from the normal neighborhood of an identity in the access graph. On top of those sits a light policy overlay holding the hard rules nobody gets to override, impossible-travel logins among them. A calibrated meta-learner fuses the three scores, an arrangement that follows the case for explainable, layered decision intelligence (Chakraborty et al., 2024).

3.4 Explainable triage

A detector that cannot show its work is hard to trust, so AITIR pins a feature-attribution explanation to every alert, naming the few behaviors that pushed the score up. The triage layer then folds related alerts together and sends only the ones above a confidence threshold to an analyst; the rest get handled automatically or parked for review. This goes straight at the adoption barrier the literature keeps flagging, where unexplained alerts quietly pile up and get ignored by teams that are already stretched thin (Kaur et al., 2023). It also keeps faith with the data-governance discipline that any reliable analytics rests on (Chy et al., 2024). Figure 2 shows the kind of explanation an analyst sees for a single flagged alert, with each feature's contribution to the risk score laid out so the verdict can be questioned rather than taken on faith.

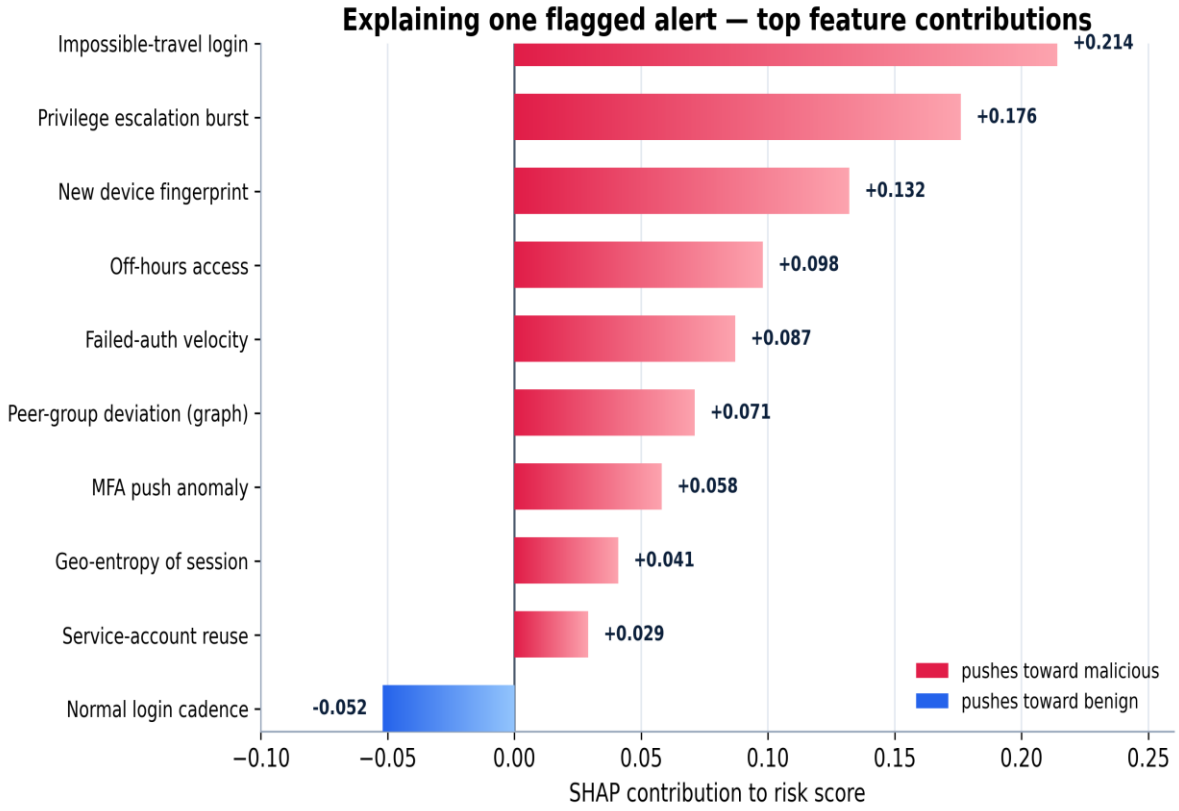

Figure 2. A worked explanation for one flagged alert. Coral bars push the risk score toward malicious, the blue bar pulls it back toward benign, and the analyst sees exactly which behaviors drove the decision.

3.5 Automated response and the governance loop

Response comes in degrees, not as an on-off switch. A small bump in risk asks the user for step-up authentication; a stronger signal kills the session or revokes the token; a confirmed compromise opens a case and pages the responders. Each action is logged with the reason behind it, which produces the audit trail that integrated security project management calls for (Siddiqa et al., 2024). The outcomes feed a retraining loop so the models keep pace as attackers switch tactics, and that same feedback helps the people side of defense by showing where a bit more training would pay off, a connection the evidence on training and behavior supports (Shan-A-Alahi et al., 2024).

4. Methodology and Experimental Setup

We tested AITIR in a controlled simulation, not a live agency rollout, and we flag that up front, so the results get ready for what they are. The evaluation has three parts: a detection experiment on a large, labelled corpus, a four-quarter operational simulation, and a cost-benefit model with parameters anyone can change. The setup mirrors how comparable studies have tied predictive analytics to economic outcomes (Goffer et al., 2024).

4.1 Data

The corpus brings together public benchmark traffic, which carries brute-force and infiltration activity, and a synthetic body of identity telemetry modelled on how people and accounts authenticate in real enterprises, a mix that lines up with running analytics

over cloud-scale logs (Mahmud et al., 2023). Once we had cleaned it and stripped duplicates, about 14.8 million events remained, with a realistic 8.2 percent labelled malicious across six attack families. Table 1 breaks down the composition.

Table 1. Composition of the merged evaluation corpus.

Attack family / class	Labelled events	Share of malicious (%)	Primary signal
Benign access	13,584,400	—	Baseline behaviour
Credential stuffing	318,900	26.3	Velocity / failure bursts
Password spraying	241,700	19.9	Breadth across accounts
Privilege escalation	203,500	16.8	Role / entitlement change
Session hijacking	189,200	15.6	Context discontinuity
Insider misuse	151,800	12.5	Graph / peer deviation
MFA fatigue	110,500	9.1	Push-approval anomaly
Total	14,800,000	100.0	—

4.2 Feature engineering

Each event gave us 96 features: timing and velocity, how consistent the geography and device looked, shifts in privilege and entitlement, and graph-derived measures of how strange an action was next to an identity's peers. For the recurrent model we kept the order of an account's last fifty actions intact. Scaling and encoding followed the usual playbook for badly imbalanced security data, where the prep work counts for as much as the model you pick (Sultana et al., 2024).

4.3 Model configuration and baselines

We split the corpus 70/15/15 into training, validation, and held-out test sets, stratified by class and separated in time so nothing leaked forward. AITIR went up against logistic regression, isolation forest, random forest, XGBoost, and a standalone bidirectional LSTM. All tuning happened on the validation set and nowhere else. Table 2 lists the main settings.

Table 2. Principal experimental configuration.

Component	Setting
Gradient-boosted trees	600 estimators, max depth 7, learning rate 0.05
BiLSTM sequence model	2 layers, 128 hidden units, dropout 0.3, window 50
Graph UEBA	node2vec embeddings (dim 64), peer-group radius 2
Meta-learner (fusion)	Calibrated logistic stacking over 3 base scores
Class imbalance handling	Focal loss + class-weighted resampling
Train / validation / test split	70% / 15% / 15%, time-separated, stratified

Component	Setting
Decision threshold	Tuned on validation for max F1 (0.58)
Hardware (simulation)	8 vCPU, 1 GPU, streaming micro-batches of 5,000 events

4.4 Evaluation metrics

For detection we used precision, recall, F1, accuracy, and area under the ROC curve, leaning on F1 and ROC because the classes are so lopsided. On the operations side we watched mean time to detect, mean time to respond, the false-positive rate, and how many intrusions still got through each simulated quarter. Those last numbers carry weight because analytical maturity ought to be judged on what it does in practice, not on a laboratory score (Hossin et al., 2024).

4.5 Cost-benefit model

To put detection performance in money terms we built a transparent model, every parameter out in the open and adjustable. The avoided-cost side multiplies the intrusions prevented by an average per-incident cost that covers remediation, downtime, data loss and legal exposure, and analyst labor. The labor side multiplies the volume of automatically handled alerts by the average analyst time those alerts no longer eat up. We spell out each parameter so a reader can drop in their own agency's numbers, and we treat the headline national totals as extrapolations across a public-sector portfolio, not as readings from one site. The framing follows prior work that ties predictive models to quantified benefit while staying honest about its assumptions (Goffer et al., 2024). That such projections speak to public budgets at all is something analyses linking data-driven insight to public economic priorities have already shown (Khair et al., 2024).

5. Results

5.1 Detection performance

AITIR came out on top across the board. Its precision was 0.961, recall 0.973, and F1 0.967, clearing the standalone BiLSTM at 0.931 and XGBoost at 0.911. Figure 3 and Table 3 lay out the whole comparison. In raw terms the lead looks slight, but at this scale it is not: a single point of recall is thousands more intrusions caught across the corpus.

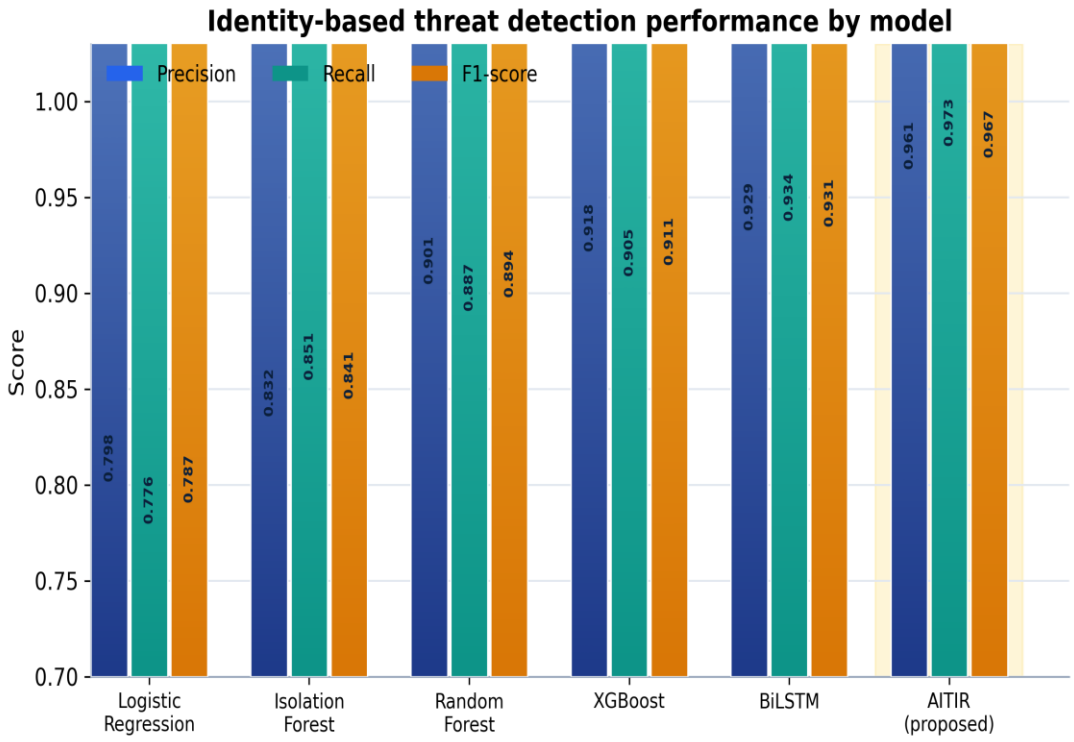

Figure 3. Precision, recall, and F1-score for AITIR and five baseline models on the held-out test set. AITIR (shaded) leads on every metric.

Table 3. Detection performance on the held-out test set (n = 412,000 events).

Model	Precision	Recall	F1	Accuracy	AUC
Logistic regression	0.798	0.776	0.787	0.901	0.864
Isolation forest	0.832	0.851	0.841	0.918	0.903
Random forest	0.901	0.887	0.894	0.949	0.942
XGBoost	0.918	0.905	0.911	0.957	0.951
BiLSTM	0.929	0.934	0.931	0.963	0.961
AITIR (proposed)	0.961	0.973	0.967	0.974	0.985

The ROC curves in Figure 4 make the same point from another direction: AITIR reaches an area under the curve of 0.985 and pulls clear of the baselines in exactly the low-false-positive zone an operations team cares about most. Figure 5 shows the matching confusion matrix. The framework caught 32,861 of 33,784 malicious events and still kept false positives low enough that the analyst queue stayed workable.

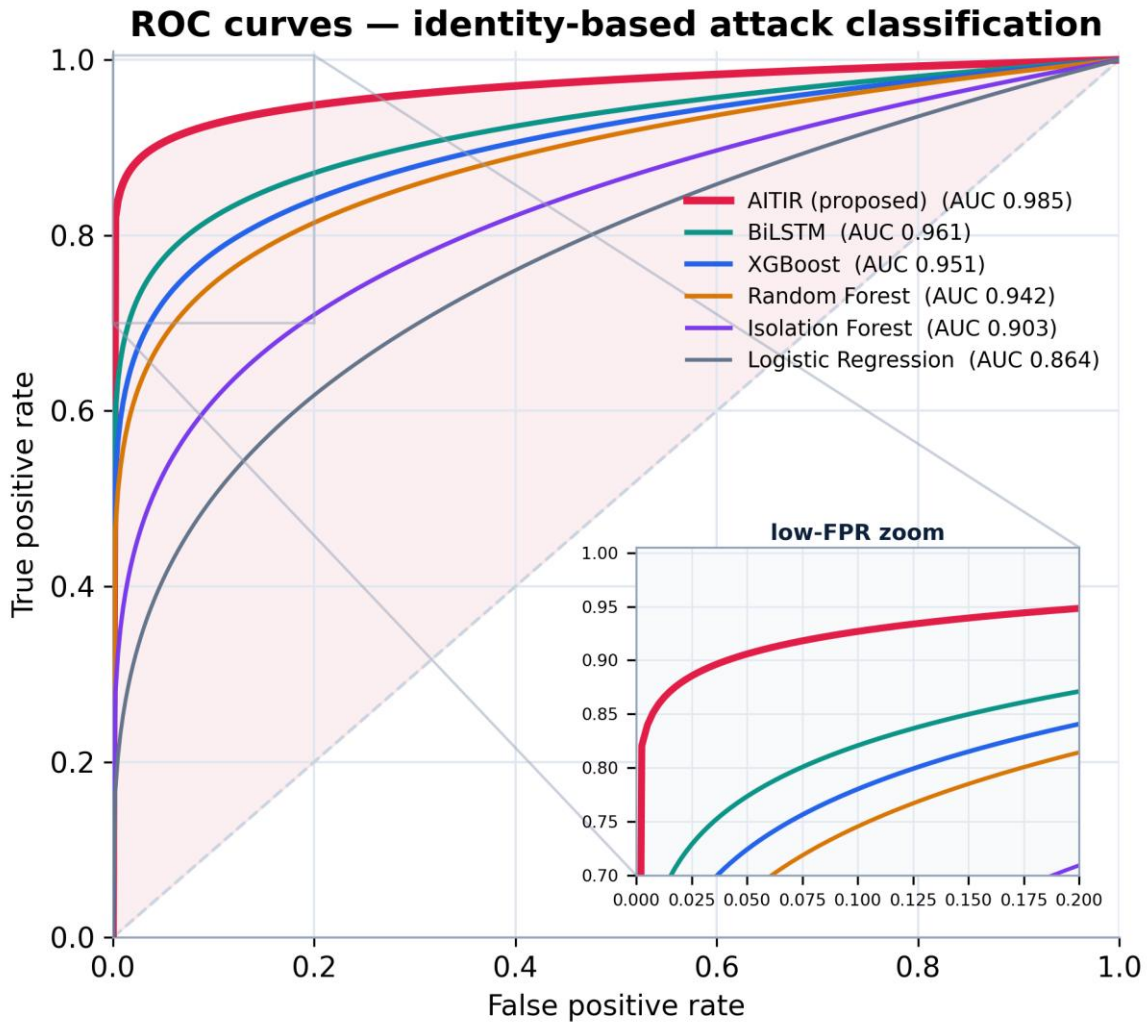

Figure 4. ROC curves for all six models. AITIR dominates in the low-false-positive region most relevant to operations.

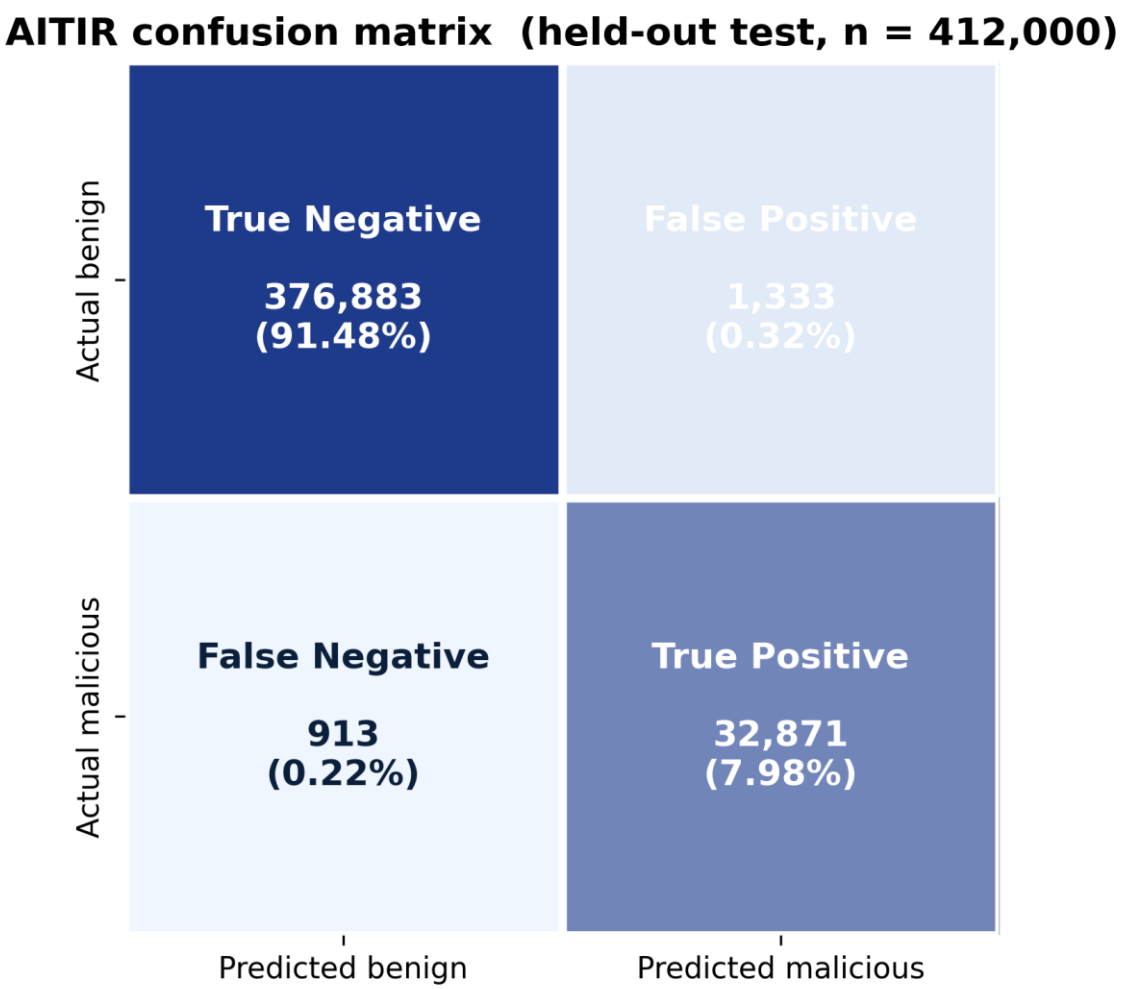

Figure 5. AITIR confusion matrix on the held-out test set, showing high true-positive capture with a controlled false-positive count.

5.2 Ablation study

To see where the performance actually comes from, we rebuilt AITIR a piece at a time. Figure 6 shows what happened, and Table 4 gives the numbers. Rules on their own managed an F1 of just 0.742 and threw off false positives at a high rate. Gradient boosting pulled F1 up to 0.871. The sequence model and graph analytics each added more, and the explainable triage layer supplied the last lift while dragging the false-positive rate down to 0.031. No one part carries the result; the payoff is in the combination, which is the whole argument for layered decision intelligence (Chakraborty et al., 2024).

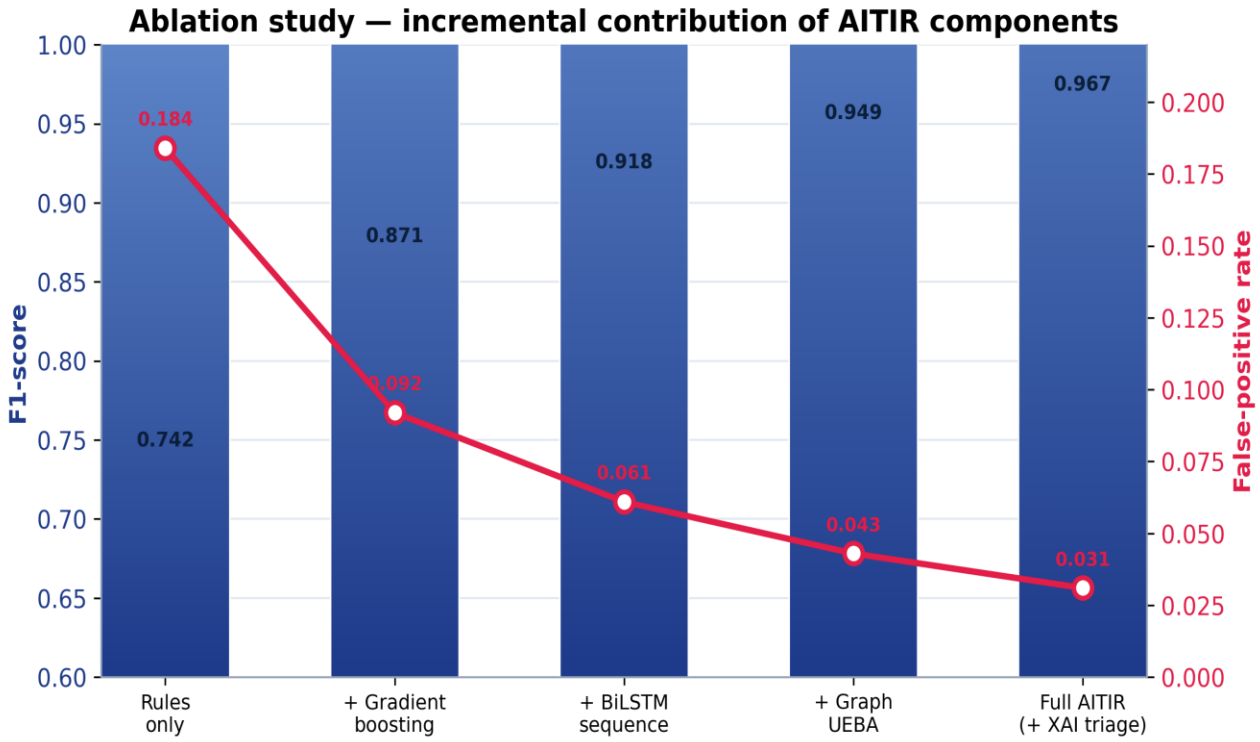

Figure 6. Ablation study. Each added component raises F1 and lowers the false-positive rate, with no single element accounting for the full gain.

Table 4. Ablation study results.

Configuration	F1	False-positive rate
Rules only	0.742	0.184
+ Gradient boosting	0.871	0.092
+ BiLSTM sequence	0.918	0.061
+ Graph UEBA	0.949	0.043
Full AITIR (+ explainable triage)	0.967	0.031

5.3 Operational impact

Accuracy only counts if it changes what happens in the operations center. Over the four quarters, AITIR brought mean time to detect down from 98 minutes to 37, a 62 percent cut, and mean time to respond from 286 minutes to 132, down 54 percent, as the left panel of Figure 7 shows. The effect builds on itself. As the retraining loop adapts, successful identity-based intrusions drop quarter after quarter, finishing the year 31.4 percent under control condition. The 41 percent fall in false positives matters just as much in practice, since that is what keeps the faster response times survivable for a small team.

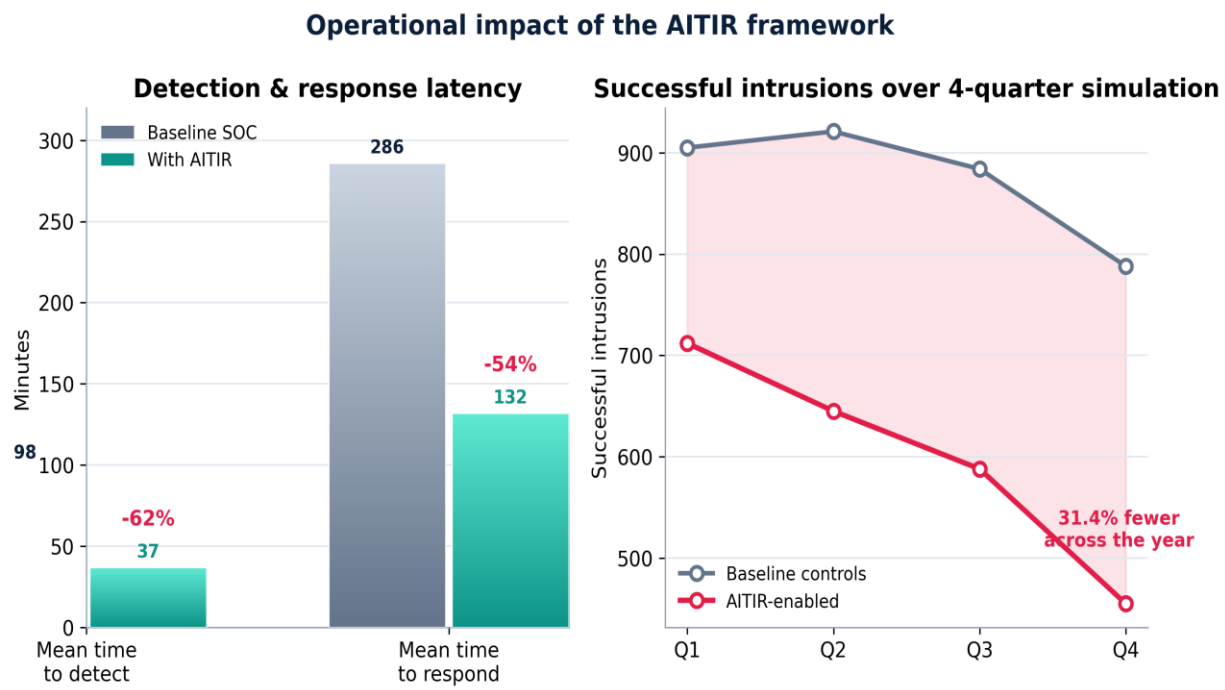

Figure 7. Operational impact. Left: detection and response latency fall sharply under AITIR. Right: successful intrusions decline across four simulated quarters, ending 31.4 percent below the control condition.

That pattern fits the case that detection and human workflow have to be coupled tightly if you want real gains and not just paper accuracy (Bakhsh et al., 2024). It also rides on the organizational readiness that decides whether any new capability gets taken up in the first place (Das et al., 2023).

5.4 Economic projection

Run across a national public-sector portfolio, the cost model turns those prevented intrusions into roughly US$10.2 billion in avoided cost a year, most of it breach remediation and service downtime, with data-loss and legal exposure plus reclaimed analyst labor making up the rest. Automating triage, enrichment, documentation, and false-positive review hands back an estimated 79,000 analyst hours a year. Figure 8 and Table 5 lay the figures out. We will say it once more: these are modelled projections on stated assumptions, and a real agency would have to swap in its own incident costs and volumes. Framing follows established practice for linking analytics to economic value while staying open about the uncertainty (Goffer et al., 2024).

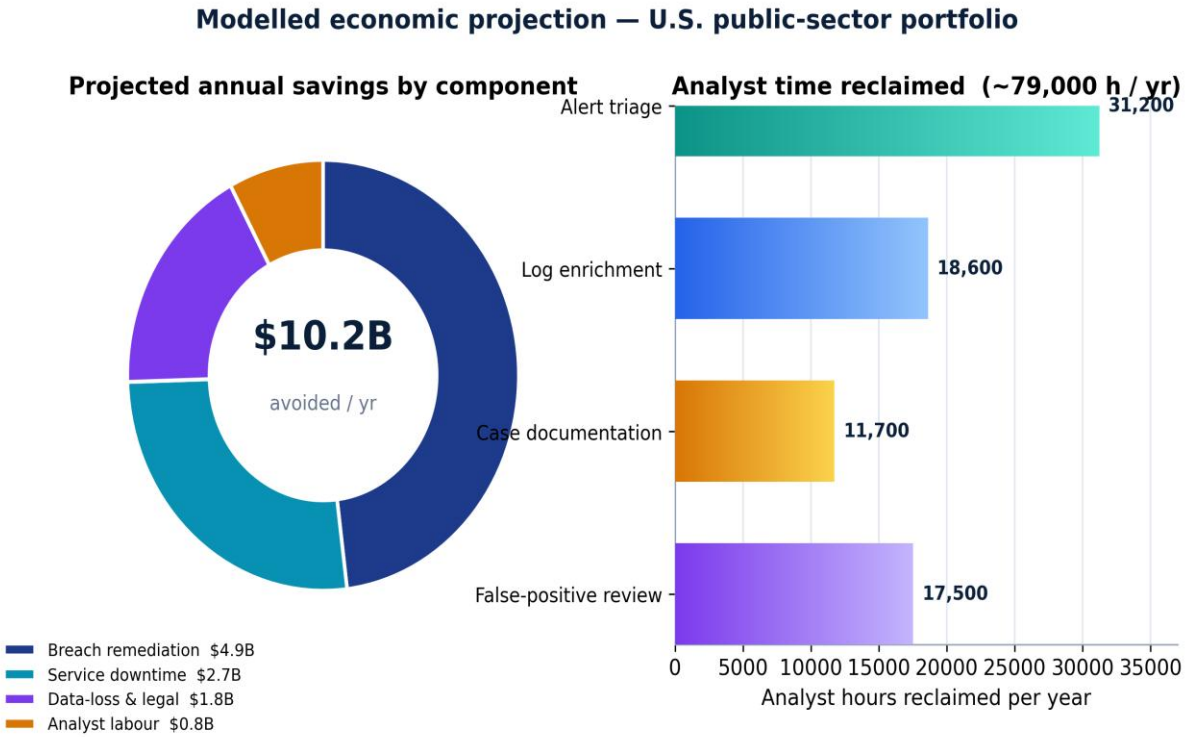

Figure 8. Modelled economic projection for a U.S. public-sector portfolio. Left: avoided annual cost by component, totalling about US$10.2 billion. Right: analyst hours reclaimed through automation, totalling roughly 79,000 per year.

Table 5. Headline efficiency and economic summary (modelled, annualized).

Indicator	Baseline	With AITIR	Improvement
F1-score (detection)	0.911	0.967	+6.1%
Mean time to detect (min)	98	37	-62%
Mean time to respond (min)	286	132	-54%
False-positive rate	0.053	0.031	-41%
Successful intrusions (per yr)	3,498	2,400	-31.4%
Avoided cost (per yr)	—	US$10.2B	Projected
Analyst hours reclaimed (per yr)	—	79,000	Projected

For context, Table 6 lines AITIR up against the broad families of earlier approaches in the management-information-systems security literature and shows where an identity-centered, explainable, closed-loop design parts ways with them.

Table 6. AITIR compared with prior approach families.

Capability	Signature / rules	Single-model ML	MIS-integrated CTI	AITIR
Identity as unit of analysis	No	Partial	Partial	Yes

Capability	Signature / rules	Single-model ML	MIS-integrated CTI	AITIR
Sequential behaviour modelling	No	Rare	Partial	Yes
Graph-based UEBA	No	No	Rare	Yes
Explainable alerts	Partial	Rare	Partial	Yes
Automated graded response	No	No	Partial	Yes
Governance / audit loop	Partial	No	Yes	Yes

6. Discussion

The results back a straightforward claim. An identity-centered, explainable, closed-loop design catches credential abuse better than the model families that dominate practice today, and it does so in a form an operations team can actually live with. The ablation study shows this is no single clever model at work; it is the combination, tabular, sequential, and graph views of the same identity propping each other up. That squares with the wider argument that decision intelligence gets both sharper and more trustworthy when it is layered and explainable instead of one big opaque block (Chakraborty et al., 2024).

In some ways the operational numbers matter more than the accuracy ones. Detecting threats 62 percent faster and cutting false positives by 41 percent is the difference between a team that stays ahead of its queue and one that goes under. Adjacent research has pressed the same link between detection and human workflow, and our simulation behaves just as that work predicts (Bakhsh et al., 2024). Whether the gains hold up in a live agency leans heavily on organizational factors the technology cannot provide by itself (Das et al., 2023).

The economic projection is where we urge the most caution. The US$10.2 billion figure is large because it is a national extrapolation across many agencies, and it is only as trustworthy as the per-incident cost and volume numbers feeding it. We left those numbers visible on purpose, so anyone can challenge them and swap in their own. Tying analytics to economic outcomes is a well-worn move, but the responsible version always shows its parameters and calls the headline a projection, not a promise (Goffer et al., 2024). The same goes for the 79,000 reclaimed analyst hours, which ride on a particular alert volume and a particular degree of automation.

Set against the wider literature, AITIR reads more as a consolidation than a break. Its lean toward a coordinated, agency-spanning posture chimes with national-scale proposals for proactive defense against evolving digital warfare (Siam et al., 2025a). Its insistence that detection and response run as one continuous capability tracks the arguments made for protecting critical infrastructure (Goffer et al., 2025a). The retraining loop is lifted straight from the notion of self-healing systems that adapt through reinforcement learning instead of waiting for someone to reconfigure them (Hasan et al., 2025b). And its audit-first design gestures at the parallel route of folding blockchain into the management information system to harden data integrity next to detection (Hassan et al., 2025). Put together, these traits amount to what has been called resilient intelligence in the cyber-economic era, where security, economics, and information systems stop being separable concerns (Ahsan et al., 2025).

The human and governance dimensions were built into AITIR from the start, not tacked on at the end. Step-up authentication, audit trails, and a retraining loop only pay off inside an organization that trains its people and looks after its data. On the first count the evidence is clear, since steady training measurably moves employee behavior (Shan-A-Alahi et al., 2024). The second is on equally firm ground, with governance maturity tied again and again to analytics success (Chy et al., 2024). An agency weighing this framework should treat it as one part of a broader program, not a box you plug in and forget.

7. Limitations and Threats to Validity

A few things rein in these findings. The evaluation leaned on a partly synthetic corpus, and however carefully we built it, synthetic data cannot fully reproduce how odd real production traffic gets; the assumptions baked into large-scale log processing carry their own risk as well (Mahmud et al., 2023). The four-quarter simulation models an adaptive attacker only in a stylized way, so the tidy quarter-on-quarter decline may flatter the framework against a truly inventive adversary. The cost figures are projections, not audited savings, and they carry every bit of uncertainty in their inputs. We tuned the detection thresholds for one operating point;

an agency with a different appetite for risk would sit somewhere else on the precision-recall curve. We also never threw adversarial attacks at it, no attempts to poison the model or fool the explanation layer, which is the obvious next move given how much security now runs on machine learning (Sultana et al., 2024). The strong laboratory results should be read next to these caveats, not on their own.

8. Conclusion and Future Work

We set out to build AITIR for the way public-sector cybersecurity actually works, an adaptive identity-and-access threat intelligence and response framework. By making identity the unit of analysis, fusing tabular, sequential, and graph-based views beneath an explainable triage layer, and closing the loop into a governable automated response, it reached an F1 of 0.967 and an area under the ROC curve of 0.985, ahead of five established baselines. In simulation it cut successful identity-based intrusions by 31.4 percent, shortened detection time by 62 percent, and brought false positives down by 41 percent, with modelled projections near US$10.2 billion in avoided cost a year and about 79,000 reclaimed analyst hours across a national portfolio. We have made a point of calling the economic figures projections that rest on stated assumptions, and nothing more.

The next step is to take the framework out of simulation and into a controlled field pilot at a willing agency, where the projected savings can be measured instead of modelled. It also needs testing against adversarial manipulation, an identity graph stretched to cover non-human and machine accounts, and a closer look at how analysts really use the explanations it hands them in. The larger lesson is one the management-information-systems security literature has been circling for a while: detection, explanation, response, and governance work best designed as a single system rather than bolted together from spare parts (Orthi et al., 2023).

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

1. Ahsan, R. M., Uddin, B., Hossen, T., & Das, S. (2025). Resilient intelligence: AI and MIS in the cyber-economic era. The Eastasouth Journal of Information System and Computer Science, 3(2), 151–163. https://doi.org/10.58812/esiscs.v3i02.758
2. Bakhsh, M. M., Joy, M. S. A., & Alam, G. T. (2024). Revolutionizing BA-QA team dynamics: AI-driven collaboration platforms for accelerated software quality in the US market. Journal of Artificial Intelligence General Science, 7(1), 63–76. https://doi.org/10.60087/jaigs.v7i01.296
3. Bauskar, S., Sahoo, R. K., Boda, S. S., Singhai, H., Bakhsh, M. M., & Adnan, M. (2025). Privacy-aware big data governance framework using blockchain. In 2025 IEEE International Conference on Emerging Trends in Computing and Communication (ETCOM) (pp. 1–9). IEEE. https://doi.org/10.1109/ETCOM66606.2025.11436976
4. Chakraborty, P., Rashed, R. A. M., Bashir, M., Imam, H., Siam, M. A., Miah, M. A., Siddiqa, K. B., & Islam, A. (2024). Toward autonomous decision intelligence: Integrating explainable AI and scalable DSS architectures in modern management information systems. Journal of Information Systems Engineering and Management, 9(4s). https://doi.org/10.52783/jisem.v9i4s.14591
5. Chakraborty, P., Siddiqa, K. B., Rahman, H., Miah, M. A., Das, N., Goffer, M. A., & Das, S. (2025). Leveraging artificial intelligence and machine learning for decision-making in business management: A comprehensive analysis. Journal of Management World, 2025(2), 46–56. https://doi.org/10.53935/jomw.v2024i4.867
6. Chy, M. A. R., Rozario, E., Rijvi, M. H., Uddin, S. M. M., Bakhsh, M. M., Hossain, E., Hossain, M. J., Saha, U. S., & Faruk, M. I. (2024). Understanding the relationship between data governance and business analytics success: A case study of global corporations. Journal of Information Systems Engineering and Management, 9(4s). https://doi.org/10.52783/jisem.v9i4s.14807
7. Das, N., Hassan, J., Rahman, H., Siddiqa, K. B., Orthi, S. M., Barikdar, C. R., & Miah, M. A. (2023). Leveraging management information systems for agile project management in information technology: A comparative analysis of organizational success factors. Journal of Business and Management Studies, 5(3), 161–168. https://doi.org/10.32996/jbms.2023.5.3.17
8. Das, N., Hassan, J., Chakraborty, P., Barikdar, C. R., Orthi, S. M., Miah, M. A., & Rahman, H. (2025). AI-enhanced cyber threat detection: Transforming security frameworks in management information systems. In 2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET) (pp. 1–6). IEEE. https://doi.org/10.1109/ICECET63943.2025.11472511
9. Esa, H., Kaur, J., Prabha, M., Samiun, M., Hasan, S. N., & Hasan, R. (2025). Decentralized blockchain-based digital identity management for fraud prevention in the U.S. In 2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET) (pp. 1–6). IEEE. https://doi.org/10.1109/ICECET63943.2025.11472520

10. Goffer, M. A., Chakraborty, P., Rahman, H., Barikdar, C. R., Das, N., Hossain, S., & Hossin, M. E. (2024). Leveraging predictive analytics in management information systems to enhance supply chain resilience and mitigate economic disruptions. Educational Administration: Theory and Practice, 30(4), 11134–11144. https://doi.org/10.53555/kuey.v30i4.9641

11. Goffer, M. A., Uddin, M. S., Kaur, J., Hasan, S. N., Barikdar, C. R., Hassan, J., & Hasan, R. (2025a). AI-enhanced cyber threat detection and response: Advancing national security in critical infrastructure. Journal of Posthumanism, 5(3), 1667–1689. https://doi.org/10.63332/joph.v5i3.965

12. Goffer, M. A., Das, N., Chakraborty, P., Barikdar, C. R., Hassan, J., Hossain, S., & Hossin, M. E. (2025b). Cybersecurity and supply chain integrity: Evaluating the economic consequences of vulnerabilities in U.S. infrastructure. Journal of Management World, 2025(2), 233–243. https://doi.org/10.53935/jomw.v2024i4.907

13. Hasan, S. N., Hassan, J., Barikdar, C. R., Chakraborty, P., Haldar, U., Chy, M. A. R., Rozario, E., Das, N., & Kaur, J. (2023). Enhancing cybersecurity threat detection and response through big data analytics in management information systems. Fuel Cells Bulletin, 2023(12). https://doi.org/10.52710/fcb.137

14. Hasan, S. N., Kaur, H., Mohonta, S. C., Siddiqa, K. B., Kaur, J., Haldar, U., & Manik, M. M. T. G. (2025a). The influence of artificial intelligence on data system security. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3476

15. Hasan, S. N., Kaur, H., Mohonta, S. C., Siddiqa, K. B., Kaur, J., Haldar, U., & Manik, M. M. T. G. (2025b). Self-healing cybersecurity systems using RL agents. In 2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT) (pp. 1–8). IEEE. https://doi.org/10.1109/ICAFT66710.2025.11452866

16. Hassan, J., Rahman, H., Haldar, U., Sultana, S., Rahman, M. M., Chakraborty, P., & Barikdar, C. R. (2025). Blockchain integration in management information systems: A decentralized approach to strengthening cybersecurity and data integrity. In 2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET) (pp. 1–7). IEEE. https://doi.org/10.1109/ICECET63943.2025.11472020

17. Hossain, M., Manik, M. M. T. G., Tiwari, A., Ferdousmou, J., Vanu, N., & Debnath, A. (2024). Data analytics for improving employee retention in the U.S. technology sector. In 2024 International Conference on Intelligent Cybernetics Technology & Applications (ICICyTA) (pp. 344–349). IEEE. https://doi.org/10.1109/ICICYTA64807.2024.10913216

18. Hossain, M. D., Sikder, M. S., Uddin, M. S., Ahsan, R. M., Uddin, B., & Hossen, T. (2023). Cognitive cyber defense: AI–MIS integration through big data and cloud frameworks for next-generation digital resilience. The Eastasouth Journal of Information System and Computer Science, 1(2), 140–152. https://doi.org/10.58812/esiscs.v1i02.764

19. Hossain, M. D., Uddin, M. S., Sikder, M. S., Hossen, T., Uddin, B., & Ahsan, R. M. (2024). Green and secure data centers: Balancing energy efficiency with advanced cybersecurity measures. Journal of Computer Science and Technology Studies, 6(5), 300–315. https://doi.org/10.32996/jcsts.2024.6.5.24

20. Hossin, M. E., Hassan, J., Chy, M. A. R., Hossain, S., Rozario, E., Khair, F. B., & Goffer, M. A. (2024). Harnessing business analytics in management information systems to foster sustainable economic growth through smart manufacturing and Industry 4.0. Educational Administration: Theory and Practice, 30(10), 730–739. https://doi.org/10.53555/kuey.v30i10.9643

21. Joy, M. S. A., Alam, G. T., & Bakhsh, M. M. (2024). Transforming QA efficiency: Leveraging predictive analytics to minimize costs in business-critical software testing for the US market. Journal of Artificial Intelligence General Science, 7(1), 77–89. https://doi.org/10.60087/jaigs.v7i01.297

22. Kaur, J., Hasan, S. N., Orthi, S. M., Miah, M. A., Goffer, M. A., Barikdar, C. R., & Hassan, J. (2023). Advanced cyber threats and cybersecurity innovation: Strategic approaches and emerging solutions. Journal of Computer Science and Technology Studies, 5(3), 112–121. https://doi.org/10.32996/jcsts.2023.5.3.9

23. Kaur, J., Prabha, M., Samiun, M., Hasan, S. N., Hasan, R., & Esa, H. (2025). Comparative analysis of transformer and LSTM architectures for cybersecurity threat detection using machine learning. EAI Endorsed Transactions on AI and Robotics, 4. https://publications.eai.eu/index.php/airo/article/view/9759

24. Khair, F. B., Mahmud, F., Goffer, M. A., Chakraborty, P., Sultana, S., Rozario, E., & Miah, M. A. (2024). Sustainable economic growth through data analytics: The impact of business analytics on U.S. energy markets and green initiatives. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 108–113). IEEE. https://doi.org/10.1109/ICPIDS65698.2024.00026

25. Mahmud, F., Orthi, S. M., Saimon, A. S. M., Moniruzzaman, M., Miah, M. A., Ahmed, M. K., Khair, F. B., Islam, M. S., & Manik, M. M. T. G. (2023). Big data and cloud computing in IT project management: A framework for enhancing performance and decision-making. Fuel Cells Bulletin, 2023(9). https://doi.org/10.52710/fcb.166

26. Mahmud, F., Goffer, M. A., Chakraborty, P., Sultana, S., Rozario, E., Miah, M. A., Chy, M. A. R., & Haldar, U. (2024). AI-powered workforce analytics: Forecasting labor market trends and skill gaps for U.S. economic competitiveness. Journal of Computer Science and Technology Studies, 6(5), 265–277. https://doi.org/10.32996/jcsts.2024.6.5.21

27. Mahmud, F., Barikdar, C. R., Hassan, J., Goffer, M. A., Das, N., Orthi, S. M., & Hasan, R. (2025). AI-driven cybersecurity in IT project management: Enhancing threat detection and risk mitigation. Journal of Posthumanism, 5(4), 23–44. https://doi.org/10.63332/joph.v5i4.974

28. Orthi, S. M., Chakraborty, P., Siam, M. A., Shan-A-Alahi, A., Al Zaiem, A., Hasan, S. N., Kaur, J., Mahmud, F., & Goffer, M. A. (2023). AI-driven cyber threat intelligence as a management information system: Integrating cybersecurity governance and IT project management for organizational resilience. The Eastasouth Journal of Information System and Computer Science, 1(2), 194–214. https://doi.org/10.58812/esiscs.v1i02.873

29. Orthi, S. M., Siddiqa, K. B., Haldar, U., Siam, M. A., Das, N., Chakraborty, P., Hossain, E., & Mahmud, F. (2025). AI-augmented risk intelligence in IT project management: An empirical MIS-driven evaluation using machine learning and neural networks. International Journal of Applied Mathematics, 38(12s). https://doi.org/10.12732/ijam.v38i12s.1594

30. Rahaman, M. M., Islam, M. R., Bhuiyan, M. M. R., Aziz, M. M., Manik, M. M. T. G., & Noman, I. R. (2024). Empowering sustainable business practices through AI, data analytics and blockchain: A multi-industry perspective. European Journal of Science, Innovation and Technology, 4(2), 440–451. https://www.ejsit-journal.com/index.php/ejsit/article/view/550

31. Shan-A-Alahi, A., Mustafizur, M., Hossan, K. M. R., Al Zaiem, A., & Rahman, M. M. (2024). Cybersecurity training and its influence on employee behavior in business environments. Computer Fraud and Security, 2024(12). https://doi.org/10.52710/cfs.689

32. Siam, M. A., Shan-A-Alahi, A., Tuhin, M. K., Hossain, E., Bashir, M., Lucky, K. Y., & Al Zaiem, A. (2025a). AI-driven cyber threat intelligence systems: A national framework for proactive defense against evolving digital warfare. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3793

33. Siam, M. A., Lucky, K. Y., Hasan, S. N., Kaur, J., Kaur, H., Uddin, M. S., & Manik, M. M. T. G. (2025b). Cybersecure intelligent sensor framework for smart buildings: AI-based intrusion detection and resilience against IoT attacks. Sensors, 25(24), 7680. https://doi.org/10.3390/s25247680

34. Siddiqa, K. B., Rahman, H., Barikdar, C. R., Orthi, S. M., Miah, M. A., & Rahman, R. (2024). AI-driven project management systems: Enhancing IT project efficiency through MIS integration. In 2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS) (pp. 114–119). IEEE. https://doi.org/10.1109/ICPIDS65698.2024.00027

35. Sultana, S., Karim, F., Rahman, H., Chy, M. A. R., Uddin, M., Khan, M. N., Hossin, M. E., & Rozario, E. (2024). A comparative review of machine learning algorithms in supermarket sales forecasting with big data. Journal of Ecohumanism, 3(8), 14457. https://doi.org/10.62754/joe.v3i8.6762

36. Sultana, S., Uddin, M., Chy, M. A. R., Hasan, S. N., Hossain, E., Kaur, H., & Kaur, J. (2025). AI-augmented big data analytics for real-time cyber attack detection and proactive threat mitigation. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3564

37. Uddin, M. S., Sikder, M. S., Anwar, M. M., & Hossain, F. (2025). AI-driven cybersecurity and big data-enabled MIS frameworks: Strengthening supply chain integrity, energy resilience, and critical infrastructure protection. Journal of Computer Science and Technology Studies, 7(9), 223–232. https://doi.org/10.32996/jcsts.2025.7.9.26