

Development and Testing of AI-Driven Intrusion Detection Systems, Anomaly Detection, and Predictive Analytics for Preventing Data Breaches in Enugu

Nebo Anne Nkeledilim

Department of Computer Education, Enugu State College of Education (Technical), Enugu State, Nigeria

Corresponding author: Nebo Anne Nkeledilim

ABSTRACT

Data breaches against banks, tertiary institutions, hospitals, telecommunication operators, and government agencies in Enugu have grown in frequency and cost, while most organisations in the metropolis still depend on signature-based intrusion detection that cannot recognise novel attacks. This study developed and tested a hybrid artificial intelligence (AI) intrusion detection prototype that combines gradient-boosted classification, autoencoder-based anomaly detection, and a long short-term memory (LSTM) predictive analytics module, and it examined how AI-driven security capabilities relate to breach prevention effectiveness among information technology (IT) personnel in Enugu. A two-phase quantitative design was adopted. In Phase 1, a corpus of 200,000 labelled network flows (120,000 from the CIC-IDS2017 benchmark and 80,000 captured from three pilot organisations in Enugu) was used to train and test six detectors on a stratified hold-out set of 40,000 flows. In Phase 2, a proportionate stratified random sample of 206 IT and cybersecurity personnel drawn from 38 organisations completed a validated 22-item questionnaire. The hybrid prototype achieved 98.54% accuracy, an F1 score of 97.07%, and a false positive rate of 0.89%, outperforming the signature-based baseline (accuracy = 90.66%, F1 = 79.88%); a McNemar test confirmed the difference, $\chi^2(1, N = 40,000) = 2,774.54$, $p < .001$, Cohen's $g = .44$. Multiple regression showed that AI intrusion detection capability ($\beta = .34$, $p < .001$, $f^2 = .14$), anomaly detection capability ($\beta = .28$, $p < .001$, $f^2 = .09$), and predictive analytics capability ($\beta = .21$, $p = .001$, $f^2 = .06$) jointly explained 48.5% of the variance in breach prevention effectiveness, $F(3, 202) = 63.29$, $p < .001$. The study concludes that layered AI detection substantially strengthens breach prevention in Enugu and recommends phased adoption supported by local traffic data, analyst training, and compliance with the Nigeria Data Protection Act 2023.

Keywords: intrusion detection; anomaly detection; predictive analytics; machine learning; data breach prevention; Enugu

1. Introduction

The digital transformation of public and private organisations in Nigeria has expanded the volume of sensitive data that institutions collect, store, and transmit. Banks process customer identity and transaction records, tertiary institutions hold student results and staff payroll files, hospitals maintain patient histories, and government ministries, departments, and agencies (MDAs) manage citizen records. Each of these repositories is an attractive target for cybercriminals. Globally, the average cost of a data breach reached about US\$4.45 million in 2023, and organisations that deployed security AI and automation extensively recorded markedly lower breach costs and shorter breach life cycles than those that did not (IBM Security, 2023). For Nigerian organisations, the enactment of the Nigeria Data Protection Act 2023 has raised the legal and reputational stakes of a breach, since data controllers are now expected to implement appropriate technical and organisational measures to protect personal data (Federal Republic of Nigeria, 2023).

Enugu, the capital of Enugu State and a long-standing administrative and educational hub of South-East Nigeria, illustrates the tension between rapid digitalisation and limited cyber-defence capacity. The metropolis hosts federal and state tertiary institutions, commercial bank regional offices, microfinance banks, fintech agents, telecommunications switching facilities, teaching and specialist hospitals, and a growing number of state MDAs that have migrated records to online portals. Many of these organisations protect their networks with firewalls and signature-based intrusion detection systems (IDS) that compare traffic against libraries of known attack patterns. Signature-based systems are effective against documented threats but are blind to zero-day exploits, polymorphic malware, credential misuse, and slow data exfiltration that do not match stored rules (Garcia-Teodoro et al., 2009; Khraisat et al., 2019).

Artificial intelligence offers three complementary responses to this weakness. First, supervised machine learning and deep learning classifiers can learn complex, non-linear boundaries between benign and malicious traffic from labelled data, thereby improving detection of attack variants (Buczak & Guven, 2016; Vinayakumar et al., 2019). Second, anomaly detection methods model normal behaviour and flag deviations, which allows the detection of previously unseen attacks without prior labels (Chandola et al., 2009; Ahmed et al., 2016). Third, predictive analytics uses historical incident and alert data to forecast periods or assets at elevated risk, enabling proactive allocation of scarce security resources. Although these approaches are well

developed in international literature, evidence on their development, testing, and perceived value in Nigerian operational settings, and in Enugu in particular, remains scarce.

2. Statement of the Problem

Organisations in Enugu continue to report phishing-led account compromises, unauthorised access to student and payroll databases, ransomware incidents, and insider data leakage, yet most rely on rule-based tools that generate large numbers of alerts while missing novel attacks. Security teams are small, and analysts in many institutions combine network administration with security monitoring. International studies demonstrate that AI-driven detectors outperform signature systems on benchmark datasets, but benchmark performance does not automatically transfer to local networks whose traffic composition, bandwidth constraints, and user behaviour differ from those represented in public datasets (Sommer & Paxson, 2010; Ring et al., 2019). In addition, little is known about whether IT personnel in Enugu perceive AI intrusion detection, anomaly detection, and predictive analytics as capabilities that actually improve breach prevention. Without locally generated technical and organisational evidence, decision-makers lack the basis for investing in AI-driven security. This study addresses that gap by developing and testing a hybrid AI detection prototype on local traffic and by examining the relationship between AI-driven capabilities and breach prevention effectiveness among IT personnel in Enugu.

3. Purpose, Questions, and Hypotheses

The general purpose of the study was to develop and test AI-driven intrusion detection, anomaly detection, and predictive analytics for preventing data breaches in Enugu. Specifically, the study sought to (a) develop a hybrid AI intrusion detection prototype integrating supervised classification, anomaly detection, and predictive analytics; (b) compare the detection performance of the prototype with a signature-based baseline and with single-model AI detectors; (c) determine the relationship between AI intrusion detection capability and breach prevention effectiveness; (d) determine the relationship between anomaly detection capability and breach prevention effectiveness; and (e) determine the relationship between predictive analytics capability and breach prevention effectiveness.

Two research questions guided the work. First, how does the hybrid AI intrusion detection prototype perform, in terms of

accuracy, precision, recall, F1 score, false positive rate, and area under the receiver operating characteristic curve (AUC), relative to a signature-based baseline and single-model AI detectors? Second, to what extent do AI intrusion detection capability, anomaly detection capability, and predictive analytics capability explain breach prevention effectiveness in organisations in Enugu?

Three null hypotheses were tested at the .05 level of significance. Anticipated effect sizes were specified a priori using Cohen's (1988) conventions for f^2 , where .02, .15, and .35 denote small, medium, and large effects.

H_{01} : AI intrusion detection capability does not significantly predict breach prevention effectiveness in organisations in Enugu (anticipated effect: small-to-medium, $f^2 \approx .10$).

H_{02} : Anomaly detection capability does not significantly predict breach prevention effectiveness in organisations in Enugu (anticipated effect: small-to-medium, $f^2 \approx .08$).

H_{03} : Predictive analytics capability does not significantly predict breach prevention effectiveness in organisations in Enugu (anticipated effect: small, $f^2 \approx .05$).

4. Significance of the Study

The study is significant at technical, organisational, and policy levels. Technically, it provides a tested reference architecture and performance evidence derived partly from traffic captured in Enugu networks, which local developers and university laboratories can replicate. Organisationally, it offers IT managers empirical grounds for prioritising investment among intrusion detection, anomaly detection, and predictive analytics. For regulators such as the Nigeria Data Protection Commission, the findings indicate the kind of technical measures that can support compliance with statutory breach-prevention obligations. Academically, the study contributes to a thin body of Nigerian quantitative evidence at the intersection of machine learning and information security management.

5. Theoretical Framework

5.1 Dynamic Capabilities Theory

The study was anchored on Dynamic Capabilities Theory and supported by Protection Motivation Theory, while the technical development drew on Denning's (1987) intrusion detection model. Dynamic Capabilities Theory, advanced by Teece et al. (1997), holds that organisations sustain performance in rapidly changing environments not merely by possessing resources but by developing the capacity to integrate, build, and reconfigure internal and external competencies. Dynamic capabilities are commonly described as the abilities to sense changes in the environment, seize opportunities or respond to threats, and transform organisational routines accordingly. The cyber-threat landscape is a quintessential turbulent environment: attack techniques mutate continuously, and static defences decay in effectiveness over time.

Within this study, AI-driven security functions are conceptualised as technology-enabled dynamic capabilities. AI intrusion detection represents a seizing capability because it classifies and blocks recognised malicious traffic in real time. Anomaly detection represents a sensing capability because it identifies deviations from normal behaviour that may signal new threats for which no signature exists. Predictive analytics represents a transforming capability because forecasts of elevated risk allow organisations to reallocate monitoring effort, patch priorities, and access controls before incidents occur. The theory therefore predicts that organisations with stronger AI-enabled sensing, seizing, and transforming capabilities will achieve higher breach prevention effectiveness, which is the logic underpinning the three hypotheses.

5.2 Protection Motivation Theory

Protection Motivation Theory (PMT), proposed by Rogers (1975), explains protective behaviour as the product of threat

appraisal (perceived severity and vulnerability) and coping appraisal (perceived efficacy of the recommended response and self-efficacy to carry it out). Applied to organisational security, PMT suggests that IT personnel will engage actively with protective technologies when they perceive breaches as severe and likely and when they believe that the tools available to them are effective. The theory is relevant because the second phase of this study measured personnel perceptions: capability ratings reflect coping appraisal, whereas breach prevention effectiveness reflects the outcome of protective action. PMT therefore complements Dynamic Capabilities Theory by explaining the human pathway through which technical capability translates into prevention outcomes.

5.3 Intrusion Detection Model

Denning (1987) proposed that intrusions could be detected by monitoring audit records for abnormal patterns of system use, relying on statistical profiles of subjects and objects. This formulation established the distinction between misuse (signature) detection and anomaly detection that continues to structure the field (Khraisat et al., 2019). The prototype developed in this study operationalises both branches: a supervised classifier addresses misuse detection with learned rather than hand-written signatures, and an autoencoder and isolation forest provide profile-based anomaly detection consistent with Denning's original model.

6. Review of Empirical Studies

6.1 Datasets and the validity of benchmark evaluation

Tavallae et al. (2009) conducted a statistical analysis of the KDD Cup 99 dataset and found a very large proportion of redundant and duplicated records in both the training and test sets. They showed that this redundancy biased learning algorithms towards frequent records and inflated reported accuracy, and they proposed the refined NSL-KDD dataset. Moustafa and Slay (2015) generated the UNSW-NB15 dataset in a controlled cyber range, combining realistic modern normal activity with synthetic contemporary attack behaviour across nine attack families. Sharafaldin et al. (2018) produced CIC-IDS2017, capturing benign background traffic modelled on user profiles alongside brute force, denial of service, web attacks, infiltration, and botnet activity. Ring et al. (2019), in a systematic survey of network-based intrusion detection datasets, concluded that no single dataset is universally representative and recommended evaluation on more than one dataset and, where possible, on data from the researcher's own operating environment. That recommendation directly informs the present study's combination of benchmark and locally captured traffic.

6.2 Machine learning and deep learning detectors

Buczak and Guven (2016) surveyed data mining and machine learning methods for cyber intrusion detection and found that ensemble and tree-based methods frequently achieved strong results, but that comparison across studies was hampered by inconsistent datasets and metrics. Breiman's (2001) random forest and Chen and Guestrin's (2016) XGBoost, the two ensemble learners used in this study, were both shown in their original evaluations to deliver high accuracy and robustness. Kim et al. (2016) applied a long short-term memory recurrent neural network to KDD Cup 99 and reported a high detection rate. Yin et al. (2017) developed an RNN-based model on NSL-KDD and found that it outperformed J48, artificial neural networks, random forest, and support vector machines. Javaid et al. (2016) used self-taught learning with a sparse autoencoder and softmax classifier on NSL-KDD.

Shone et al. (2018) proposed a non-symmetric deep autoencoder stacked with a random forest classifier and reported higher accuracy and lower training time than a deep belief network on KDD Cup 99 and NSL-KDD. Vinayakumar et al. (2019) evaluated deep neural networks across KDD Cup 99, NSL-KDD, UNSW-NB15, and CIC-IDS2017 and found that

deep networks generally outperformed classical classifiers. Ferrag et al. (2020) compared seven deep learning models on CSE-CIC-IDS2018 and Bot-IoT and showed that performance varied by model and attack type. Ahmad et al. (2021) concluded that hybrid approaches were promising but that most studies relied on outdated datasets and seldom reported deployment conditions.

6.3 Anomaly detection, imbalance, evaluation, and acceptance

Chandola et al. (2009) emphasised that the definition of normal behaviour, the availability of labels, and the type of anomaly determine which technique is appropriate. Liu et al. (2008) introduced the isolation forest, isolating anomalies through random partitioning with linear time complexity. Mirsky et al. (2018) developed Kitsune, an ensemble of lightweight autoencoders for online unsupervised network intrusion detection, and demonstrated that it could run on inexpensive hardware, a finding particularly relevant to resource-constrained organisations in Enugu. Ahmed et al. (2016) noted that high false positive rates remain the principal barrier to operational adoption.

Intrusion data are typically imbalanced. Chawla et al. (2002) proposed SMOTE and showed that combining it with undersampling improved minority-class performance in ROC space. Saito and Rehmsmeier (2015) showed that precision-recall plots are more informative than ROC plots on imbalanced data, which justifies reporting precision, recall, and F1 alongside AUC. Sommer and Paxson (2010) argued that machine learning is far less used in real intrusion detection than in research because of the high cost of errors, traffic diversity, the semantic gap between alerts and actionable meaning, and flawed evaluation practices. Lundberg and Lee (2017) proposed SHAP values as a unified approach to interpreting model predictions. Davis (1989) and Venkatesh et al. (2003) established that perceived usefulness shapes technology acceptance, justifying the examination of how personnel perceive AI-driven capabilities.

6.4 Gap in the literature

The review reveals three gaps. First, nearly all reviewed technical studies evaluated detectors exclusively on public benchmark datasets generated in North American, European, or Australian laboratories; none tested detectors on traffic from Nigerian networks, whose composition is shaped by shared bandwidth, mobile-first access, heavy use of third-party payment gateways, and intermittent power-related outages. Second, studies typically evaluated intrusion detection, anomaly detection, or forecasting in isolation rather than as an integrated system. Third, technical performance studies seldom connected detector performance to organisational breach prevention outcomes as perceived by the personnel who operate the tools. The present study addresses all three gaps by developing an integrated hybrid prototype, testing it on a combination of benchmark and Enugu traffic, and examining the organisational relationship between AI-driven capabilities and breach prevention effectiveness.

7. Method

7.1 Research design and study area

The study adopted a two-phase quantitative design. Phase 1 used a developmental and experimental (design-and-test) approach consistent with design science research in information systems, in which an artefact is constructed and rigorously evaluated against defined criteria (Hevner et al., 2004). Phase 2 used a cross-sectional correlational survey design to test the hypothesised relationships between AI-driven security capabilities and breach prevention effectiveness. Combining the two phases allowed the study to establish both what the technology can achieve under controlled testing and how AI-enabled capabilities relate to prevention outcomes in operating organisations.

The study was conducted in Enugu metropolis, comprising Enugu North, Enugu South, and Enugu East Local Government Areas of Enugu State, Nigeria. The metropolis concentrates the state's major data-intensive organisations, including federal and state tertiary institutions, regional offices of commercial banks, microfinance banks, telecommunications and internet service providers, public and private hospitals, and state MDAs whose services have been digitised. These organisations maintain networks that process sensitive personal and financial data and are therefore subject to the breach-prevention obligations of the Nigeria Data Protection Act 2023.

7.2 Phase 1: Development of the hybrid AI detection prototype

A labelled corpus of 200,000 bidirectional network flows was assembled. Of these, 120,000 flows were sampled from CIC-IDS2017 (Sharafaldin et al., 2018), covering benign traffic and attack classes including brute force, denial of service, web attacks, port scanning, infiltration, and botnet activity. The remaining 80,000 flows were captured over 30 days at the network perimeters of three pilot organisations in Enugu (a tertiary institution, a microfinance bank, and a hospital) using port mirroring and CICFlowMeter-compatible flow extraction. Local flows were labelled by combining signature alerts, analyst verification of each alert, and a controlled red-team exercise conducted with written authorisation from each organisation, during which scanning, credential brute-forcing, and simulated exfiltration were injected. All payloads were discarded, and internal IP addresses were pseudonymised before analysis. Across the combined corpus, 150,000 flows (75%) were benign and 50,000 (25%) were malicious.

Constant and duplicate features were removed, missing and infinite values were replaced with feature medians, and continuous features were standardised. A stratified 80/20 split produced a training set of 160,000 flows and a hold-out test set of 40,000 flows (30,000 benign; 10,000 attack). To address class imbalance, SMOTE (Chawla et al., 2002) was applied only within training folds, so that no synthetic records leaked into validation or test data.

Six detectors were built. The baseline was a signature-based IDS (Snort) running a current community rule set. Four single-model AI detectors were trained: an isolation forest (Liu et al., 2008) fitted on benign training flows; a random forest (Breiman, 2001); an XGBoost classifier (Chen & Guestrin, 2016); and a two-layer LSTM network (Hochreiter & Schmidhuber, 1997) operating on sequences of ten consecutive flows per host. The hybrid prototype integrated three layers. The anomaly layer comprised an autoencoder trained on benign flows, whose reconstruction error served as an anomaly score, together with the isolation forest score. The detection layer was an XGBoost classifier that received the original flow features plus the two anomaly scores. The predictive analytics layer was an LSTM forecaster that used daily aggregates of alerts, failed logins, and vulnerability-scan findings to predict the probability that the next 24 hours would be a high-risk day. SHAP values (Lundberg & Lee, 2017) were generated for each alert.

Hyperparameters were tuned by five-fold stratified cross-validation on the training set, with F1 as the optimisation criterion. Final models were retrained on the full training set and evaluated once on the untouched test set. Performance was measured with accuracy, precision, recall (detection rate), F1 score, false positive rate (FPR), and AUC. Following Saito and Rehmsmeier (2015), precision, recall, and F1 were given primary weight because of class imbalance. The difference between the hybrid prototype and the signature baseline was tested with McNemar's test on paired predictions, with Cohen's κ and the discordant-pair odds ratio as effect sizes. The predictive analytics layer was trained on 18 months of daily incident and alert records from the three pilot organisations and tested on the final 90 days.

7.3 Phase 2: Survey of IT and cybersecurity personnel

The population comprised 412 IT and cybersecurity personnel (network and system administrators, security analysts, IT managers, and database or software engineers) employed in 38 organisations in Enugu metropolis. The Krejcie and Morgan (1970) formula with $\chi^2 = 3.841$, $N = 412$, $P = .50$, and $d = .05$ yielded a minimum sample of 199. An a priori power analysis in G*Power 3 (Faul et al., 2007) for multiple regression with three predictors, a medium effect ($f^2 = .15$), $\alpha = .05$, and power = .95 indicated a minimum of 119 respondents. The larger value (199) was adopted, and 240 questionnaires were distributed. A proportionate stratified random sampling technique was used, with five organisational sectors as strata. Of 240 questionnaires, 214 were returned and eight were discarded for incomplete responses, leaving 206 usable questionnaires (usable response rate = 85.8%).

Table 1. Population, proportionate sample, and usable responses by sector.

Sector	Orgs	Pop.	Sample	Dist.	Usable
Banks and fintech	10	118	57	69	60
Tertiary institutions	7	96	46	55	48
Telecoms and ISPs	5	64	31	37	32
Hospitals	6	52	25	30	26
Government MDAs	10	82	40	49	40
Total	38	412	199	240	206

Note. Proportionate sample = (stratum population / 412) $\times$ 199, rounded. ISPs = internet service providers; MDAs = ministries, departments, and agencies.

Data were collected with the AI-Driven Security Capability and Breach Prevention Questionnaire (AISCBPQ). Section A captured demographic data. Section B contained 22 items rated from 1 (strongly disagree) to 5 (strongly agree), measuring four constructs: AI intrusion detection capability (AIDC; six items), anomaly detection capability (ADC; five items), predictive analytics capability (PAC; five items), and breach prevention effectiveness (BPE; six items). Items were adapted from the intrusion detection literature and refined for the Nigerian context.

Face and content validity were established by three experts: two in computer science (cybersecurity) and one in measurement and evaluation. A pilot test with 30 IT personnel in Awka, Anambra State, who were not part of the study sample, confirmed clarity and yielded preliminary Cronbach's alpha values above .80 for all constructs. For the main study, construct reliability and validity were assessed through confirmatory factor analysis. Internal consistency was evaluated with Cronbach's alpha and composite reliability ($CR \geq .70$). Convergent validity was assessed through standardised factor loadings ($\geq .70$) and average variance extracted ($AVE \geq .50$), and discriminant validity through the Fornell and Larcker (1981) criterion (Nunnally & Bernstein, 1994; Hair et al., 2019).

Questionnaires were administered in person and electronically by the researcher and four trained research assistants over six weeks. Participation was voluntary, informed consent was obtained, and no names or organisational identifiers were recorded. For Phase 1, written permission was obtained from the management of each pilot organisation, red-team activities were restricted to agreed windows and targets, packet payloads were not retained, and all personal data were handled in line with the Nigeria Data Protection Act 2023. Phase 1 data were analysed with Python (scikit-learn, XGBoost, and TensorFlow). Phase 2 data were analysed with IBM SPSS Statistics and AMOS. Means and standard deviations described the constructs, confirmatory factor analysis established the measurement model, and multiple linear regression tested the three hypotheses. Regression assumptions were checked through normality of residuals, the Durbin-Watson statistic, and variance inflation factors. Common method bias was assessed

with Harman's single-factor test (Podsakoff et al., 2003). Hypotheses were rejected where $p < .05$.

8. Results

8.1 Demographic profile of respondents

Table 2 shows the characteristics of the 206 respondents. Most were male (69.4%), held a first degree or Higher National Diploma (54.4%), and had 5–10 years of experience (39.3%). Network and system administrators formed the largest role category (34.5%), and fewer than half (45.6%) held a professional security certification.

Table 2. Demographic characteristics of respondents ($N = 206$).

Characteristic	Category	n	%
Gender	Male	143	69.4
	Female	63	30.6
Role	Network/system administrator	71	34.5
	Security analyst/SOC staff	52	25.2
	Database/software engineer	45	21.8
Experience	IT manager	38	18.4
	Less than 5 years	58	28.2
	5–10 years	81	39.3
Highest qualification	More than 10 years	67	32.5
	HND/bachelor's degree	112	54.4
	Master's degree	79	38.3
Security certification	Doctorate	15	7.3
	Yes	94	45.6
	No	112	54.4

Note. SOC = security operations centre; HND = Higher National Diploma.

8.2 Measurement model

Table 3 presents factor loadings and reliability statistics. All 22 standardised loadings ranged from .74 to .85. Cronbach's alpha ranged from .894 (PAC) to .916 (BPE), and composite reliability ranged from .894 to .916. AVE values ranged from .628 to .646, confirming convergent validity. Harman's single-factor test showed that one factor accounted for 31.4% of the total variance, well below 50%, suggesting that common method bias was not a serious concern.

Table 3. Factor loadings, reliability, and convergent validity per construct.

Construct	Items	Loadings	α	CR	AVE
AIDC	6	.74–.85	.912	.913	.636
ADC	5	.76–.84	.900	.900	.644
PAC	5	.75–.83	.894	.894	.628
BPE	6	.77–.84	.916	.916	.646

Note. Item loadings: AIDC = .82, .79, .85, .77, .81, .74; ADC = .84, .80, .78, .83, .76; PAC = .79, .83, .81, .75, .78; BPE = .80, .84, .78, .82, .77, .81. CR = composite reliability; AVE = average variance extracted.

Table 4 reports construct means, standard deviations, and inter-construct correlations, with the square root of AVE on the diagonal. Respondents rated AI intrusion detection capability highest ($M = 3.87$, $SD = 0.68$) and predictive analytics capability lowest ($M = 3.54$, $SD = 0.79$), suggesting that forecasting is the least mature of the three capabilities in Enugu organisations. All correlations were positive and significant ($r = .52$ to $.61$, $p < .001$), and each diagonal value (.792 to .804) exceeded the largest correlation in its row and column (.61), satisfying the Fornell-Larcker criterion.

Table 4. Means, standard deviations, correlations, and discriminant validity.

Construct	M	SD	1	2	3	4
1. AIDC	3.87	0.68	(.797)			
2. ADC	3.71	0.72	.58***	(.803)		
3. PAC	3.54	0.79	.52***	.55***	(.792)	

4. BPE	3.79	0.70	.61***	.59***	.54***	(.804)
--------	------	------	--------	--------	--------	--------

Note. $N = 206$. Values in parentheses are square roots of AVE. *** $p < .001$.

8.3 Research Question 1: Performance of the hybrid prototype

Table 5 presents the performance of the six detectors on the 40,000-flow hold-out test set. The hybrid prototype recorded the best results on every metric: accuracy of 98.54%, precision of 97.31%, recall of 96.84%, F1 of 97.07%, FPR of 0.89%, and AUC of .994. It correctly detected 9,684 of the 10,000 attack flows while raising 268 false alarms among 30,000 benign flows (584 misclassifications). The signature baseline detected only 7,412 attacks (recall = 74.12%) and produced 1,146 false alarms (3,734 misclassifications). XGBoost was the strongest single model (F1 = 96.04%), followed by random forest and LSTM. The isolation forest achieved higher recall than the signature baseline (82.36% vs. 74.12%) but the highest FPR (7.70%). Adding anomaly scores to XGBoost reduced missed attacks from 438 to 316 (27.9%) and false alarms from 351 to 268 (23.6%).

Table 5. Detection performance on the 40,000-flow hold-out test set.

Detector	Acc. %	Prec. %	Rec. %	F1 %	FPR %	AUC
Signature (Snort)	90.66	86.61	74.12	79.88	3.82	—
Isolation forest	89.81	78.10	82.36	80.17	7.70	.903
Random forest	97.71	95.93	94.87	95.40	1.34	.987
LSTM	97.69	95.44	95.31	95.38	1.52	.985
XGBoost	98.03	96.46	95.62	96.04	1.17	.991
Hybrid prototype	98.54	97.31	96.84	97.07	0.89	.994

Note. Test set = 40,000 flows (10,000 attack; 30,000 benign). Hybrid: TP = 9,684, FN = 316, FP = 268, TN = 29,732. Signature: TP = 7,412, FN = 2,588, FP = 1,146, TN = 28,854. AUC is not reported for the signature baseline because it produces binary alerts rather than scores.

McNemar's test on paired predictions (Table 6) confirmed that the hybrid prototype's advantage over the signature baseline was statistically significant. Of 3,574 discordant flows, the hybrid was correct and the baseline wrong in 3,362 cases, whereas the baseline was correct and the hybrid wrong in only 212 cases, $\chi^2(1, N = 40,000) = 2,774.54$, $p < .001$. The discordant-pair odds ratio was 15.86, and Cohen's g was .44, a large effect (Cohen, 1988).

Table 6. Paired classification outcomes of the hybrid prototype and signature baseline.

Hybrid prototype	Signature correct	Signature incorrect	Total
Correct	36,054	3,362	39,416
Incorrect	212	372	584
Total	36,266	3,734	40,000

Note. McNemar $\chi^2(1)$ with continuity correction = $(|3,362 - 212| - 1)^2 / (3,362 + 212) = 2,774.54$, $p < .001$; odds ratio = $3,362/212 = 15.86$; Cohen's $g = .44$.

The predictive analytics layer was evaluated on the final 90 days of incident records, of which 21 days were designated high-risk on the basis of confirmed incidents or intrusion attempts that required escalation. The forecaster flagged 23 days as high-risk, of which 17 were correct, yielding a recall of .81 (17/21), a precision of .74 (17/23), an overall accuracy of .89 (80/90), and an AUC of .89. In a supplementary 30-day live pilot, mean inference latency for the hybrid detector was 1.8 milliseconds per flow on a standard CPU server, and SHAP explanations were available for every alert.

8.4 Research Question 2 and hypothesis testing

Regression assumptions were satisfied. Residuals were approximately normal, skewness and kurtosis for all constructs fell within ± 1 , the Durbin-Watson statistic was 1.94, and VIF values ranged from 1.57 to 1.73. As shown in Table 7, the three AI-driven capabilities jointly explained 48.5% of the variance in breach prevention effectiveness, $R^2 = .485$, adjusted $R^2 = .477$, $F(3, 202) = 63.29$, $p < .001$, a large overall effect ($f^2 = 0.94$).

Table 7. Multiple regression of breach prevention effectiveness on AI-driven security capabilities.

Predictor	B	SE	β	t	p	VIF	sr ²	f ²
Constant	0.78	0.22		3.49	<.001			
AIDC	0.35	0.07	.34	5.23	<.001	1.65	.070	.14
ADC	0.27	0.07	.28	4.17	<.001	1.73	.044	.09
PAC	0.19	0.06	.21	3.33	.001	1.57	.028	.06

Note. $N = 206$. $R^2 = .485$; adjusted $R^2 = .477$; $F(3, 202) = 63.29$, $p < .001$. sr^2 = squared semipartial correlation.

AI intrusion detection capability significantly and positively predicted breach prevention effectiveness, $\beta = .34$, $t(202) = 5.23$, $p < .001$, uniquely accounting for 7.0% of the variance ($f^2 = .14$), a small-to-medium effect slightly larger than the anticipated .10. H_{01} was rejected. Holding the other capabilities constant, a one-point increase in AIDC was associated with a 0.35-point increase in BPE.

Anomaly detection capability significantly and positively predicted breach prevention effectiveness, $\beta = .28$, $t(202) = 4.17$, $p < .001$, uniquely explaining 4.4% of the variance, $f^2 = .09$, a small-to-medium effect close to the anticipated .08. H_{02} was rejected.

Predictive analytics capability significantly and positively predicted breach prevention effectiveness, $\beta = .21$, $t(202) = 3.33$, $p = .001$, uniquely explaining 2.8% of the variance, $f^2 = .06$, a small effect consistent with the anticipated .05. H_{03} was rejected. The three unique contributions sum to 14.2% of variance; the remaining 34.3% of explained variance was shared among the predictors. Table 8 summarises the decisions.

Table 8. Summary of hypothesis tests.

H	Path	β	p	f ²	Effect	Decision
H_{01}	AIDC → BPE	.34	<.001	.14	Small to medium	Rejected
H_{02}	ADC → BPE	.28	<.001	.09	Small to medium	Rejected
H_{03}	PAC → BPE	.21	.001	.06	Small	Rejected

Note. Cohen's (1988) f^2 benchmarks: .02 = small, .15 = medium, .35 = large.

9. Discussion

The hybrid prototype detected almost 97% of attack flows with a false positive rate below 1%, whereas the signature baseline missed roughly one in four attacks. This gap is consistent with Buczak and Guven (2016), Vinayakumar et al. (2019), and Ahmad et al. (2021): learned models outperform rule-based detection. The large McNemar effect ($g = .44$) confirms that the advantage is not an artefact of chance. The superiority of XGBoost among the single models agrees with Chen and Guestrin (2016) and with the strong showing of tree ensembles in the surveys reviewed. That the LSTM performed comparably to, but not better than, the tree ensembles mirrors Ferrag et al. (2020), who found that deep architectures do not uniformly dominate across datasets and attack types.

Adding anomaly scores to the supervised classifier reduced both missed attacks and false alarms. This supports the logic of hybrid architectures advanced by Shone et al. (2018) and extends it: rather than using the autoencoder only for feature compression, the prototype used reconstruction error as explicit evidence of novelty. The isolation forest on its own produced the highest false positive rate, echoing Ahmed et al. (2016) and Sommer and Paxson (2010). The hybrid design converts the sensitivity of anomaly detection into a feature that the supervised layer can weigh. The low inference latency observed in the pilot is consistent with Mirsky et al. (2018) and is an important consideration for Enugu organisations operating with constrained budgets.

AI intrusion detection capability was the strongest organisational predictor. In terms of Dynamic Capabilities Theory (Teece et al., 1997), this reflects the seizing capability: organisations able to classify and act on malicious traffic in real time are better placed to contain intrusions before data is exfiltrated. The significant contribution of anomaly detection capability supports the sensing dimension of the theory and Denning's (1987) argument that profiling normal behaviour enables detection of threats for which no signature exists. Predictive analytics capability made the smallest but still significant unique contribution and was also the lowest-rated capability. Forecasting affects prevention indirectly, by shaping staffing, patching, and monitoring priorities, so its effect may be partly carried through the other two capabilities. From the standpoint of Protection Motivation Theory (Rogers, 1975), personnel who perceive their tools as efficacious are more likely to act on their outputs (Davis, 1989; Venkatesh et al., 2003).

Taken together, the two phases converge. The technical tests show that layered AI detection is substantially more accurate than signature-based detection on a mixture of benchmark and local traffic, and the survey shows that organisations with stronger AI-driven capabilities report more effective breach prevention. The relatively low rate of professional security certification among respondents suggests, however, that human capacity may limit how fully these technical gains are realised.

10. Implications

For IT managers in Enugu, the results indicate that replacing or augmenting signature-based IDS with AI-driven intrusion detection should be the first investment priority, followed by anomaly detection integrated into, rather than deployed separately from, the classification layer. Predictive analytics should be introduced once reliable incident logging is in place. For developers and university computer science departments, the study demonstrates that effective detectors can be built with open-source tools and trained partly on local traffic. For regulators, the findings suggest that guidance issued under the Nigeria Data Protection Act 2023 could recognise AI-assisted monitoring, supported by explainable alerts, as good practice for data controllers handling large volumes of personal data.

11. Limitations and Further Research

Local traffic was captured from only three organisations over 30 days, and labels relied partly on signature alerts and analyst judgement, which may have missed stealthy attacks. Attacks injected during the red-team exercise, though realistic, may not represent the full diversity of threats faced by Enugu organisations. Phase 2 relied on self-reported perceptions measured at a single point in time, which limits causal inference. The survey was restricted to Enugu metropolis. Further research should extend data capture to more organisations and longer periods, evaluate concept drift and periodic retraining, test adversarial robustness of the hybrid model, use longitudinal designs linking capability to recorded incidents, and replicate the study in other South-East and national contexts.

12. Conclusion

This study developed and tested a hybrid AI prototype that integrates supervised intrusion detection, anomaly detection, and predictive analytics, and it examined the relationship between these capabilities and breach prevention effectiveness in Enugu. The prototype outperformed a signature-based baseline and all single-model detectors, detecting 96.84% of attacks with a false positive rate of 0.89%. AI intrusion detection, anomaly detection, and predictive analytics capabilities each significantly predicted breach prevention effectiveness, with small-to-medium, small-to-medium, and small effects, respectively, and together explained nearly half of its variance. The study concludes that AI-driven, layered

detection offers a practical and substantial improvement in data breach prevention for organisations in Enugu.

13. Recommendations

Organisations in Enugu should progressively integrate machine learning classifiers with existing signature-based systems rather than rely on signatures alone. Anomaly detection should be embedded as an input to supervised detection to reduce false alarms. Organisations should institute consistent incident logging to enable predictive analytics, and they should sponsor IT staff for professional security certification and training in interpreting AI alerts. Tertiary institutions in Enugu should establish shared, anonymised traffic repositories to support local model training. The Nigeria Data Protection Commission should issue technical guidance that encourages explainable AI-assisted monitoring as part of statutory breach-prevention measures.

References

- Ahmad, Z., Shahid Khan, A., Wai Shiang, C., Abdullah, J., & Ahmad, F. (2021). Network intrusion detection system: A systematic study of machine learning and deep learning approaches. *Transactions on Emerging Telecommunications Technologies*, 32(1), Article e4150. <https://doi.org/10.1002/ett.4150>
- Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. <https://doi.org/10.1016/j.jnca.2015.11.016>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176. <https://doi.org/10.1109/COMST.2015.2494502>
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15. <https://doi.org/10.1145/1541880.1541882>
- Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357. <https://doi.org/10.1613/jair.953>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). ACM. <https://doi.org/10.1145/2939672.2939785>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Denning, D. E. (1987). An intrusion-detection model. *IEEE Transactions on Software Engineering*, SE-13(2), 222–232. <https://doi.org/10.1109/TSE.1987.232894>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/BF03193146>
- Federal Republic of Nigeria. (2023). *Nigeria Data Protection Act, 2023*. Federal Government Printer.
- Ferrag, M. A., Maglaras, L., Moschogiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. *Journal of Information Security and Applications*, 50, Article 102419. <https://doi.org/10.1016/j.jisa.2019.102419>
- Fornell, C., & Larcker, D. F. (1981). Evaluating structural equation models with unobservable variables and measurement error. *Journal of Marketing Research*, 18(1), 39–50. <https://doi.org/10.1177/002224378101800104>
- Garcia-Teodoro, P., Diaz-Verdejo, J., Maciá-Fernández, G., & Vázquez, E. (2009). Anomaly-based network intrusion detection: Techniques, systems and challenges. *Computers & Security*, 28(1–2), 18–28. <https://doi.org/10.1016/j.cose.2008.08.003>
- Hair, J. F., Risher, J. J., Sarstedt, M., & Ringle, C. M. (2019). When to use and how to report the results of PLS-SEM. *European Business Review*, 31(1), 2–24. <https://doi.org/10.1108/EBR-11-2018-0203>
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.2307/25148625>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- IBM Security. (2023). *Cost of a data breach report 2023*. IBM Corporation.
- Javadi, A., Niyaz, Q., Sun, W., & Alam, M. (2016). A deep learning approach for network intrusion detection system. In *Proceedings of the 9th EAI International Conference on Bio-inspired Information and Communications Technologies* (pp. 21–26). <https://doi.org/10.4108/eai.3-12-2015.2262516>

- Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. *Cybersecurity*, 2(1), Article 20. <https://doi.org/10.1186/s42400-019-0038-7>
- Kim, J., Kim, J., Thu, H. L. T., & Kim, H. (2016). Long short term memory recurrent neural network classifier for intrusion detection. In 2016 International Conference on Platform Technology and Service (PlatCon) (pp. 1–5). IEEE. <https://doi.org/10.1109/PlatCon.2016.7456805>
- Krejcie, R. V., & Morgan, D. W. (1970). Determining sample size for research activities. *Educational and Psychological Measurement*, 30(3), 607–610. <https://doi.org/10.1177/001316447003000308>
- Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In 2008 Eighth IEEE International Conference on Data Mining (pp. 413–422). IEEE. <https://doi.org/10.1109/ICDM.2008.17>
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765–4774). Curran Associates.
- Mirsky, Y., Doitsman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An ensemble of autoencoders for online network intrusion detection. In *Proceedings of the Network and Distributed System Security Symposium (NDSS 2018)*. Internet Society. <https://doi.org/10.14722/ndss.2018.23204>
- Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems. In 2015 Military Communications and Information Systems Conference (MilCIS) (pp. 1–6). IEEE. <https://doi.org/10.1109/MilCIS.2015.7348942>
- Nunnally, J. C., & Bernstein, I. H. (1994). *Psychometric theory* (3rd ed.). McGraw-Hill.
- Podsakoff, P. M., MacKenzie, S. B., Lee, J.-Y., & Podsakoff, N. P. (2003). Common method biases in behavioral research: A critical review of the literature and recommended remedies. *Journal of Applied Psychology*, 88(5), 879–903. <https://doi.org/10.1037/0021-9010.88.5.879>
- Ring, M., Wunderlich, S., Scheuring, D., Landes, D., & Hotho, A. (2019). A survey of network-based intrusion detection data sets. *Computers & Security*, 86, 147–167. <https://doi.org/10.1016/j.cose.2019.06.005>
- Rogers, R. W. (1975). A protection motivation theory of fear appeals and attitude change. *The Journal of Psychology*, 91(1), 93–114. <https://doi.org/10.1080/00223980.1975.9915803>
- Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), Article e0118432. <https://doi.org/10.1371/journal.pone.0118432>
- Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP 2018)* (pp. 108–116). SciTePress. <https://doi.org/10.5220/0006639801080116>
- Shone, N., Ngoc, T. N., Phai, V. D., & Shi, Q. (2018). A deep learning approach to network intrusion detection. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2(1), 41–50. <https://doi.org/10.1109/TETCI.2017.2772792>
- Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. In 2010 IEEE Symposium on Security and Privacy (pp. 305–316). IEEE. <https://doi.org/10.1109/SP.2010.25>
- Tavallae, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. In 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications (pp. 1–6). IEEE. <https://doi.org/10.1109/CISDA.2009.5356528>
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533.
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- Vinayakumar, R., Alazab, M., Soman, K. P., Poornachandran, P., Al-Nemrat, A., & Venkatraman, S. (2019). Deep learning approach for intelligent intrusion detection system. *IEEE Access*, 7, 41525–41550. <https://doi.org/10.1109/ACCESS.2019.2895334>
- Yin, C., Zhu, Y., Fei, J., & He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. *IEEE Access*, 5, 21954–21961. <https://doi.org/10.1109/ACCESS.2017.2762418>