Promise Enyindah, 2026, 14:1
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Ransomware Detection Using Behavioral Analysis and Machine Learning: A Comprehensive Review

Promise Enyindah[1], Daniel Okon[2]
[1]Department of Computer Science, University of Port Harcourt, Rivers State, Nigeria.
[2]Department of Cybersecurity, University of Port Harcourt, Rivers State, Nigeria.

Abstract- Among notable threat within rising digital risks is ransomware, known for deep consequences - yearly monetary damages climb into multiple billions, alongside disruptions felt by critical operations such as healthcare systems, financial institutions, and government bodies. Since older methods based on fixed signatures prove ineffective against evolving code structures, modern approaches shift attention toward real-time conduct of harmful software after activation. Behavior seen during runtime provides clearer indicators compared to unchanging markers. With adaptive malware variants and widely available hacking resources increasing complexity, different defensive paths emerge as vital. A potential direction focuses on teaching models to identify faint signals of malicious behavior. Research indicates some sophisticated architectures spot ransomware accurately in over 99 out of every 100 trials. Among them are networks modeled after biological vision, those designed to follow temporal patterns, together with hybrid approaches merging several strategies. Despite lingering obstacles, advancement in automatic detection moves forward at a consistent pace. From observed behaviors - alterations in saved files, active system operations, internet-based data movements - to encryption patterns, distinctions emerge between malicious software and standard applications. These indicators form the basis for analysis within openly available datasets, while challenges such as false signals or evasion tactics adopted by adversaries are considered alongside them. Practical constraints influencing real-world deployment weave into each point made across the review. Future directions surface through interpretable artificial intelligence systems, adaptive defense frameworks, collaboration enriched by live threat intelligence feeds. Insight gained from this exploration aids progress toward more effective defenses targeting harmful file-locking behavior.

Keywords: Ransomware detection, behavioral analysis, machine learning, deep learning, cybersecurity, malware detection, LSTM, CNN, threat intelligence.

I. INTRODUCTION

Once dismissed as trivial, ransomware now risks collapsing institutions, urban centers, government systems. Colonial Pipeline's 2021 compromise triggered widespread fuel scarcity along America's East Coast. Change Healthcare's 2024 intrusion, carried out via BlackCat malware, halted medication fulfillment for countless individuals throughout the country. Such visible breakdowns reveal consequences deeper than monetary damage alone. Healthcare institutions faced repeated targeting, 181 verified ransomware incidents occurred in 2024 alone, affecting 25.6 million individuals' medical information. Although $5.7 million marked the typical extortion figure requested, actual payments settled around $900,000 where negotiations took place. Additional expenses emerged during restoration efforts; recovery averaged $2.57 million per organization, omitting long-term consequences such as damaged trust or regulatory penalties linked to exposed data.

Beyond Western economies which is often commonly reported, ransomware attacks have increasingly targeted institutions in Sub-Saharan Africa, including Nigeria. Government agencies, educational establishments and financial institutions, have faced escalating incidents, and yet these remain under-reported in global literature. The Nigeria Data Protection Commission (NDPC) and the NITDA have issued multiple advisories on ransomware threats which targets critical national infrastructure. The Central Bank of Nigeria (CBN) Cybersecurity Framework mandates incident reporting across deposit money banks (DMBs), yet

remediation capacity remains constrained by a limited skilled cybersecurity workforce. This geographic gap represents both a research opportunity and a practical risk that this review seeks to acknowledge.

In 2024, complex ransomware operations ran much like structured crime networks. With 89 verified incidents, RansomHub ranked highest in activity levels. LockBit came next, responsible for 83 breaches during the year. Medusa carried out 62 documented intrusions across various sectors. Play appeared behind them, linked to 57 separate events. Instead of encryption alone, these collectives often combined public exposure threats alongside system lockdowns. Their tools adapted constantly - malicious code altered form to bypass security checks. Access spread through shared frameworks allowing less skilled individuals to launch campaigns. Each attack reflected planning depth uncommon in earlier digital offenses.

In the Nigerian context, Ransomware-as-a-Service (RaaS) campaigns have consistently directly affected local enterprises and public institutions. There are various Documented incidents involving Nigerian financial infrastructure and state government ICT systems, which demonstrate that even partial encryption of transaction databases carries disproportionate consequences in economies where digital redundancy remains low. The proliferation of RaaS toolkits lowers the technical barrier for local threat actors, thereby compounding risk within regional networks that often lack dedicated Security Operations Centre (SOC) capabilities.

Limitations of Traditional Detection Methods

Older antivirus systems often rely on fixed signatures tied to known examples. Despite their accuracy, these approaches fall short now - mainly because malware changes too quickly. Looking only at history becomes ineffective as new versions emerge constantly throughout the day.

Without notice, hidden gaps emerge. When blind spots exist, new malicious code spreads unchecked until scanners adapt. Response starts solely post-examination.

A single trait defines some ransomware: constant change. Even when operation remains unchanged, the code transforms every time. Because of this shift, recognition through past examples proves ineffective. With each iteration, design evolves while keeping its damaging aim intact. Although behavior repeats, appearance differs always. What looks new still carries the old goal.

Once detection depends on signatures, encryption could start before an alert is triggered. Between infection and discovery, time passes - data remains exposed throughout. These systems aim to identify familiar dangers, yet lag when files are altered quickly. Rather than stopping damage ahead of time, reaction usually follows visible shifts. Seemingly defensive measures may turn into efforts chasing silent transformations masked by routine activity.

Fresh ransomware variants emerge rapidly due to service-based distribution, avoiding standard detection until security responses catch up. Because attack tools can be leased, updates propagate faster through malicious ecosystems than defensive measures often allow. Existing signature methods fall short under the constant influx of new iterations introduced each day.

Under such limitations, focus moves to analyzing ransomware actions mid-attack rather than depending only on file signatures. Thanks to progress in identifying patterns via machine learning, tools today foresee risks by tracking steps common in harmful operations. Instead of reacting to familiar clues, safeguards trigger upon spotting odd encryption trends among data. By watching ongoing exchanges within a system, shifts away from usual habits signal possible intrusion sooner. When conduct develops moment by moment, algorithms shaped by varied cases detect irregularities even without earlier exposure to the danger. Although once dependent on fixed code indicators, modern techniques focus on behaviors emerging from real-time system activity. As every move is measured relative to standard patterns, conclusions develop gradually amid shifting contexts. Unfamiliar intrusions still expose their nature via actions typical of information theft sequences. Rather than comparing known templates, spotting threats relies

upon interpreting intent within recurring entry efforts. Gradually, persistent characteristics across harmful operations enable recognition even when masked.

Research Objectives

This study aims to describe ways of detecting ransomware by observing system activities while applying machine learning models. Examination centers on how these strategies operate, the threats they uncover, yet also considers their limitations and typical usage patterns

1. Examine the behavioral indicators that distinguish ransomware from legitimate software.
2. Examining techniques begins by observing behavior in machine learning models tasked with spotting ransomware. Structural choices shape the analytical framework applied throughout investigations. Detection depends heavily on which attributes are chosen, influencing experimental findings. What emerges from testing is often tied to metrics revealing performance within defined settings. What also matters is how different deep learning frameworks were used in earlier studies. Findings tend to rely on whether results across published works show similar patterns.
3. Analyze publicly available datasets used in ransomware detection research.
4. Unexpected complications arise due to incorrect alarms combined with reduced speed. Even with protections in place, hostile strategies slip past unnoticed. Adjusting frameworks to fit updated guidelines introduces delays. Defensive updates lag behind evolving bypass techniques. Real dangers get overlooked too often because of misleading signals. With more checks appearing without warning, speed begins to drop. Despite installing updates, some threats still go unnoticed.
5. Factors when setting up actual systems.
6. Propose future research directions to advance the field
7. Evaluate the feasibility of deploying behavioral machine learning detection frameworks in resource-constrained environments which is characteristic of African higher education institutions and public-sector organizations. These environments are where legacy hardware, intermittent power supply, and limited bandwidth impose deployment constraints not addressed in prevailing studies.

This analysis draws on recent research, patterns seen during actual intrusions, followed by industry responses across domains. Out of these elements forms a structure - neither theoretical nor detached, instead rooted in practical demands guiding systems designed to detect ransomware sooner. Every part fit within scholarly standards together with real-world constraints, influencing approaches before new dangers fully appear.

II. METHODOLOGY

For this study the use of a structured review method focused on spotting ransomware through actions taken during attacks alongside smart systems trained to detect patterns.

Literature Search Strategies

A close look emerged through reviewing academic papers, field records, one real-world case after another - each revealing how behavioral traces paired with machine-driven analysis detect ransomware. Different linked approaches shaped the exploration, allowing wider angles to surface naturally at every stage.

Within leading scholarly databases, investigations unfolded across IEEE Xplore together with the ACM Digital Library, arXiv, Google Scholar, and specialized periodicals focused on cyber defense. Keyword groupings tied concepts such as ransomware detection to observable changes in program execution, architectures using layered neural networks - like convolutional or recurrent types - as well as techniques monitoring fluctuations in data randomness or API usage. Despite emphasis falling on works published from 2020 through 2025 to reflect current risks and advances in machine-based reasoning, earlier foundational contributions stayed included if they introduced key structural approaches.

Inclusion and Exclusion Criteria

Inclusion depended on meeting set terms. When studies aligned with established criteria, they qualified for consideration. Those fulfilling all stipulations moved forward without exception. Passage through every checkpoint allowed entry into evaluation. Meeting every defined criterion was required to qualify. Only when each specified condition existed did inclusion follow

Focused on ransomware detection specifically (not general malware detection)
Emerging from repeated exposure, algorithmic adjustments followed behavioral cues. Patterns within information guided the formation of structures over time. Without clear instruction, understanding deepened across levels of complexity
Numbers described how well things worked, using exactness along with consistency. Alongside these, completeness joined F1-score when outcomes were shared. Mistaken identifications counted within the assessment method. From each measurement came a separate understanding of what occurred. What was written down depended on figures instead of opinions.

A single methodological outline appears, offering clarity through measured detail. Reproducibility becomes possible due to specific procedural notes included. Transparency emerges not from volume but from precision of content shared. Explanations stop before reaching surplus, maintaining restraint throughout. Structure allows comparison across different applications naturally appearing within journals assessed by expert review, alongside established trade-focused analysis platforms.

Only cryptography drew attention - other studies moved elsewhere. Without behavioral components, network surveillance failed entry. Theories untouched by practical trials held no place in this space. Absence of method details erased entire contributions from view.

Data Extraction and Analysis

Close observation of file handling can reveal behaviors that raise concern. After a program runs, changes in stored records frequently remain visible. Hidden processes occasionally show themselves through odd patterns among active tasks. How systems exchange information may display uncommon protocol choices. The presence of scrambled transmissions could indicate ongoing concealment techniques.

Several machine learning frameworks, differing in architecture, found application. Structure-based variations among models influenced implementation choices. Diverse configurations emerged through iterative design processes. Architectural distinctions shaped how each model was utilized. Application methods depended on underlying framework layouts.

Information on data collections, covering volume, makeup, and origin
Performance measurement combines accuracy with precision. Alongside evaluation metrics, recall joins the F1-score. Because it highlights missed detections, the F2-score finds use. At regular intervals, false positives appear - monitored via their distinct rate. A different measure tracks missed detections, recorded on its own. Depending on the situation, each gauge applies when errors of identification take separate forms.

Limitations and challenges identified by authors

A detailed examination at actual ransomware incidents depends on recorded data - sources involve official alerts, sector analyses, vendor summaries, alongside bodies like CISA. Shifts in attack methods often emerge alongside evolving financial pressure tactics, access approaches, while effects differ across industries.

Quality Assessment

The dataset's scale prompted attention, although careful method design shaped conclusions more. Not only size but variety within data held equal weight. How samples divided between training and testing shifted interpretation meaningfully. With verification techniques applied, insights grew clearer step by step. Over months, availability of executable code-built confidence slowly. Reporting on configurations appeared clear throughout. Depth in evaluation arose via multiple indicators of function.

Differences from standard techniques brought context forward. Views on ease of use changed because processing demands grew. Constraints tied to live operation held similar importance. Should discussion shift toward incorrect alerts, consequences face closer examination. As fragments align, confidence in outcomes increases accordingly. Though initial doubts remain, trust builds where patterns emerge clearly.

III. RESULTS

Our review identified four primary categories of behavioral indicators that distinguish ransomware from legitimate software: file system activity, process behavior, network communication patterns, and cryptographic operations. Understanding these indicators is fundamental to developing effective detection systems.

File System Activity Patterns

Ransomware exhibits distinctive file system behaviors during the encryption phase. Research has identified several key patterns:

1. **Mass file modifications:** Ransomware typically encrypts large numbers of files in rapid succession, creating an unusually high rate of file write operations that differs significantly from typical user or application behavior.
2. **Sequential access patterns:** Unlike normal applications that may access files randomly based on user interaction, ransomware often processes files in systematic patterns (directory traversal, alphabetical order, or file type prioritization).
3. **File extension changes:** Many ransomware families append specific extensions to encrypted files (e.g., locked, encrypted, or custom extensions like. wannacry), creating a detectable pattern of simultaneous extension modifications.
4. **Shadow copy deletion:** To prevent recovery, ransomware commonly deletes Windows Volume Shadow Copies using vssadmin.exe or wmic.exe commands, a behavior rarely performed by legitimate software.
5. **Ransom note creation:** Placement of text files containing ransom demands (commonly named README.txt, DECRYPT_INSTRUCTIONS.txt, or similar) across multiple directories provides a strong indicator.
6. From a practitioner point of view, these file system behaviors map directly to MITRE ATTACK Techniques T1490 (Inhibit System Recovery) and T1486 (Data Encrypted for Impact). Detection engineers can author composite rules within SIEM platforms such as Wazuh or Elastic SIEM by correlating high-frequency file write events with concurrent VSS deletion calls. This indicator substantially reduces false positive rates compared to monitoring either signal in isolation. Group Policy Objects (GPOs) which restricts non-administrative access to vssadmin.exe and wmic.exe provide a complementary preventive control

Process Behavior

Ransomware frequently aims for elevated access to extend influence over stored data. Not limited to basic boundaries, it may bypass security prompts designed to block changes. Authority levels shift when identity markers get manipulated improperly. Instead of those methods, flaws left unfixed in essential services could serve the same purpose. When safeguards meant to protect core functions weaken, broader network presence becomes possible.

Hidden within ordinary operations, harmful scripts take shape using tactics such as process hollowing. Though indistinguishable from regular activity, these strategies sidestep traditional monitoring meant for obvious dangers. Execution unfolds quietly beneath approved applications instead of standing apart. Access to higher-level rights emerges not through force but by mimicking trusted workflows. System features, built for routine tasks, are repurposed - intentionally - not due to flaws. Behind seamless appearances lie a shift in how control is silently assumed.

Survival after restart turns feasible if ransomware adjusts registry keys, ties actions to scheduled tasks, or sets up silent services. Continued function past system start happens by concealed setup changes, self-triggering mechanisms, or enduring sequences.

Resistance to shutdown emerges from altered access points, postponed initiators, or always-on modules. Steady performance results from changed boot patterns, built-in waiting periods, together with sustained run methods. After a reset, activity persists because of saved configuration updates, timed startups, or enlisted support tools.

Upon analysis, sophisticated ransomware can identify managed environments through signs of sandboxing or virtualization. Rather than proceeding normally, such malware detects debugging tools to hinder disassembly efforts. Delays triggered by time-related evaluations often reveal non-physical setups. Combined awareness of surroundings enhances avoidance capabilities significantly.

Close inspection of how processes start shows subtle hints regarding what a system is doing. Parent and offspring links catch notice due to unusual shapes they form at times. Paths beginning inside file interpreters sometimes appear clearly apart. Rather than merge with routine tasks, certain patterns break away in noticeable ways. Unexpected patterns emerged during observation of reproductive irregularities tied to message-driven access channels. In various instances, particular behaviors proved more significant than alternative indicators for recognizing risks. Shifts in behavior frequently followed moments of expanded user permissions. Context improved when hidden document appearances were properly separated. Notable changes often preceded documented system intrusions. Outcomes shifted noticeably once pattern precision took priority over sheer volume. Under particular circumstances, encryption signatures began to appear. Advanced phases frequently aligned with methods designed to resist scrutiny. A small fraction of behavioral indicators - approximately 100 - showed actual relevance. Distinct signal categories boosted classification reliability when emphasized.

Network Communication Patterns

Network-level behavioral indicators include:

1. Communication with remote servers occurs when ransomware sends data about infected systems. Such traffic might travel through unusual ports instead of common ones. External connections can involve rare protocols rather than typical network methods. Encryption keys sometimes arrive via secure links after initial contact. Instructions from attackers occasionally come after the malware activates. Data transfer happens even if standard pathways are unavailable. Contact with command centers continues until blocked or completed.
2. Outbound data flows grow larger than normal when attackers copy information prior to locking systems. Uncommon external endpoints begin receiving steady streams of internal files. Movement toward unknown servers appear during early attack phases.
3. Information leaves the network in bulk shortly after access is gained. Transfers occur at odd intervals instead of consistent schedules. Sensitive material exits through channels not typically used for communication. Volume spikes stand out compared to regular traffic baselines.
4. Occasionally, DNS lookups display irregularities - such as connections to newly created addresses, strings resembling automated naming, or signs tied to anonymized networks. Patterns emerge when domain requests mirror machine-generated structures rather than typical human-assigned names. Domains appearing freshly registered often appear within suspicious lookup sequences. Instead of standard naming logic, some queries point toward labels built through formulaic methods. Anonymity infrastructure markers sometimes surface during these atypical resolution attempts. Uncommon name resolutions may align with network behaviors seen in privacy-focused routing systems. Freshly observed domains occasionally correlate with nonstandard request frequencies.

Cryptographic Indicators

Clear signals sometimes come from cryptographic operations. High file entropy - often seen during encryption - is uncommon in ordinary data. A sudden change in entropy across disk regions can suggest ongoing encoding activity. Rather than measuring quantity alone, scrutiny falls upon system

calls tied to recognized crypto APIs. Patterns that deviate from typical request flows tend to draw notice. When analyzing usage of functions such as CryptEncrypt or BCryptDecrypt, patterns begin to appear - not solely in frequency, yet also through sequence and supplied data. From sandbox testing, contrasts surface: destructive encryption tools act unlike common utilities. Observed closely, legitimate backup software acts unlike its dangerous counterparts. Central to identification lies setting - moment, recurrence, arrangement; these unite into a unique signature. Accuracy improves once actions drift from established baselines.

Machine Learning Approaches and Architectures

Our review identified diverse machine learning and deep learning approaches applied to ransomware detection, ranging from traditional supervised learning to sophisticated hybrid neural network architectures

Traditional Machine Learning Methods

A few studies applied standard machine learning together with manually designed behavioral markers. Notable performance emerged from the Ransom Wall configuration, reaching 98.25% precision and nearly zero false positives through Gradient Tree Boosting. Although reliant on careful feature engineering, it offered transparent reasoning behind decisions while maintaining minimal computational load. Owing to such qualities, deployment within environments constrained by resources turned possible.

Older methods include Random Forests, handling multiple behavior patterns without difficulty. Where classification splits data into just two groups, SVM emerges - particularly with thoughtful kernel selection. Because explanations matter, Decision Trees appear; their logic unfolds step by step. Accuracy tends to fall short of deep learning, usually ranging from 85% to 95%, yet training completes more quickly. Despite lower precision, speed offers a different kind of advantage. Demand for computing resources remains low. Outcomes often prove easier to interpret than those from intricate systems. Gains in efficiency arise through simple operations rather than peak output. A transparent design can matter more than maximum prediction power. Simpler methods might be chosen where resources are constrained, even with lower precision. When resources are limited, simple solutions gain value. Despite shortcomings in speed or function, real-world usefulness may favor them instead

Deep Learning Architectures

While hybrid CNN-LSTM architectures deliver impressive benchmark accuracy, practitioners must always evaluate GPU memory requirements against available hardware. In resource-constrained deployments, a two-tier architecture is recommended. Lightweight gradient boosting models (XGBoost or Random Forest) execute locally on the endpoint as a first-pass filter, escalating only anomalous events to a centralized cloud inference layer running the full deep learning pipeline. This design reduces per-endpoint computational overhead by an estimated 60–75% while preserving sensitivity for ambiguous cases.

Deep learning methods handled complex patterns and timing details better than others. Notably, some frameworks delivered steady outcomes over repeated tests. What set them apart emerged during transitions between computation phases. Performance relied heavily on layer coordination evolving freely. A key feature began with managing ordered data using built-in retention systems. Unexpected improvements appeared primarily through direct learning of unprocessed data streams. Despite difficulties faced by conventional approaches, sustained precision characterized these structures across extended periods.

Features presented as images or grid-like structures often work well with Convolutional Neural Networks. Although rare classes pose challenges, merging CNNs with artificially created data brought accuracy close to 98.90 percent. Because spatial positioning matters in organized formats, these models excel - particularly when analyzing visual representations of API behavior. Rather than building anew, certain approaches used earlier systems like ResNet50, applying knowledge from large-scale image studies. Thanks to inherited understanding, robust outcomes

appeared despite small quantities of malicious software examples.

From timing structures onward, Long Short-Term Memory models manage ordered data effectively. In examining API requests, peak results emerged with sequences of seven actions grouped together. Both speed and precision stayed steady in this configuration. Rather than omitting phases, the approach traces intrusion steps carefully - initial access first, then reconnaissance, concluding with data encryption.

One configuration achieved 99.87% accuracy, with precision reaching 99.89%, recall standing at 99.85%, yet producing just 0.13% false positives. Following similar lines, an alternative structure iCNN-LSTM - attained 99.57% on the F2 metric, despite overlooking 4.89% of real incidents and signaling incorrectly in 0.16% of instances. Visual features were extracted via convolutional layers; meanwhile, temporal dynamics were captured using memory-equipped recurrent segments. Because spatial patterns and sequential behavior in malware activity were processed jointly, detection improved subtly. When run in parallel, attention-guided mechanisms supported speed, maintaining consistency throughout evaluations.

From simple activity records, some systems detect trends without human guidance. Through step-by step analysis, one model forms insights where traditional methods require hand-designed features. Ninety-nine percent precision has emerged when spotting malicious programs via this technique. By reviewing timelines of behavior, evolving frameworks adjust instead of depending on fixed signals. Previously vital expertise becomes less central due to such autonomous adaptation.

Hardware-Based Detection

Recent studies have looked into identifying threats through physical device metrics, observed via artificial intelligence methods like XGBoost, LightGBM, neural networks with multiple layers, along with convolutional structures. Instead of relying on surface software behavior, these systems examine underlying processor traits - examples include failed cache access, incorrect branching predictions, or how many operations complete per cycle all shaped by ransomware activity. Because malicious actions leave traces in chip-level operations, even programs designed to hide from conventional scans may still be noticed. Detection rooted in equipment signals might catch advanced attacks that slip past application-focused tools, since their footprints appear directly within hardware operation, independent of hidden or scrambled code.

Incremental Learning Systems

For operational deployment, incremental learning systems should always follow a structured model update cadence. A recommended practice is to trigger batch retraining upon accumulation of at least 50 novel ransomware samples, or following any confirmed zero-day incident affecting peer organizations. Integration with threat intelligence platforms such as MISP (Malware Information Sharing Platform) enables automated ingestion of Indicators of Compromise (IOCs), reducing manual dataset curation overhead in organizations with limited cybersecurity staffing.

Facing shifting ransomware behaviors, scientists turned to batch-driven update methods allowing models to evolve gradually instead of restarting training entirely. Such a method reduces heavy computing demands tied to refreshing deep learning frameworks, yet keeps accuracy stable when confronting fresh malware forms. As novel ransomware instances appear, these adaptive systems absorb them into existing knowledge, adjusting responses over time minus the burden of rebuilding everything from scratch.

Datasets and Evaluation Benchmarks

A notable gap is the near-total absence of ransomware behavioral datasets sourced from African network environments. Malware behavior varies significantly based on network topology, OS patch levels, and application stack. all this differ systematically between high income and lower middle income country deployments. Establishing a Nigerian or pan-African ransomware behavioral dataset, potentially coordinated through NITDA,

NCC, and partner universities, represents a high-priority research gap. Such a dataset would enable regionally-calibrated detection model training and contribute to global threat intelligence sharing under the ECOWAS Cybersecurity Policy framework. When building systems to detect ransomware behavior, diverse high-quality information proves essential. Our review identified several openly available datasets, thoroughly described and frequently cited, as particularly useful instances.

One hundred thirty-eight thousand forty-seven items made up a broad dataset, seventy per cent labeled as ransomware, others marked benign, used to train neural models thoroughly. Despite fewer samples, the RISS group added value: five hundred eighty-two hostile files paired with nine hundred forty-two clean counterparts, all analyzed via Cuckoo sandbox for exact behavioral traces. When attention turned to handheld devices, information from CICAndMal2017 arrived - over ten thousand eight hundred fifty-four cases harvested straight from working phones, divided clearly between risky and harmless. Progress followed with CCCS-CIC-AndMal-2020, delivering four hundred thousand Android apps, covering fourteen categories of threats and one hundred ninety-one separate origins.

Twenty-five live strains, gathered across five years, formed the foundation. From 2019 to 2024, this span allowed models to absorb trends mirroring current threats. Static images evolved into timed data markers during analysis. Where sample sizes were small, pertained frameworks filled gaps. Even as outdated variants disappeared, value held firm - rooted in emerging actions.

Even with earlier results, obtaining real-world attack samples in open databases continues to be rare; most datasets still come from experimental setups. Rather than depend on artificial environments, some researchers have used synthetic generation techniques like SMOTE to enlarge existing logs. Because ransomware behaves differently across families, broad collections are needed to capture diverse behaviors. The way harmful code runs can change depending on hardware or software settings. A small number of groups have looked outside their field, borrowing insights from related disciplines instead.

Performance Metrics and Comparative Analysis

Table -1: Name of the Table

Approach	Study	Accuracy	Precision	Recall	FPR
Hybrid LSTM-CNN	Kara & Aydos (2024)	99.87%	99.89%	99.85%	0.13%
iCNN-LSTM	Homayoun et al. (2024)	F2: 99.57%	—	—	0.16%
CNN + SMOTE	Wei et al. (2023)	98.90%	—	—	—
Gradient Tree Boosting	Ransom Wall (2023)	98.25%	—	—	~0%
ResNet50 Transfer Learning	Poudyal et al. (2024)	Superior	—	—	—

Appearing in evaluations, recurring traits emerge. While combinations such as CNN-LSTM achieve high marks, their strength reveals itself mainly when precision exceeds 99 percent. Since repeated errors reduce trust, keeping incorrect warnings rare becomes essential after implementation. Top performers in trials maintained fewer than two mistakes per thousand cases, showing ability to distinguish actual risks from routine system activity. When overlooking dangers carries heavier consequences than triggering unnecessary alarms, focus shifts toward the F2 metric this method emphasizes detecting almost every event rather than minimizing each error.

Sometimes outcomes shifted based on timing of identification. When choices needed speed, monitoring actions live faced greater difficulty

compared to evaluations made afterward, particularly if activities remained incomplete. Instead of using just a single indicator, systems that merged insights from networks, devices, and files performed more effectively. Strength in defense emerged most clearly when multiple approaches combined distinct yet connected signals through cooperation.

IV. DISCUSSION

Despite impressive detection accuracies reported in research, several significant challenges complicate real-world deployment of behavioral machine learning systems for ransomware detection.

False Positive Management

Certain standard applications behave similarly to harmful software; detection systems may respond incorrectly. Rather than operating quietly, backup utilities that duplicate numerous files can mimic unauthorized access patterns. When compression tools reorganize information, the resulting irregularity occasionally prompts alerts. Updates applied widely across directories draw attention because of extensive system alterations. Software creators working with encoded routines encounter false signals since their processes include concealed transformations.

It is known that trustworthy systems sometimes repeat damaging behaviors seen elsewhere. Within vast infrastructures housing many machines, tiny flaws still generate numerous false alerts daily. Gradually, professionals grow less responsive to frequent notifications, possibly missing actual threats. Without control, minor faults multiply into overwhelming volumes of notices. Identical operations challenge recognition logic when judging underlying purpose. Even a tiny fraction of errors becomes loud chaos at scale. When unimportant data floods in, attention splits apart.

Though separate origins exist, alerts merge into confusion. As confidence slips bit by bit, reaction times stretch longer. False alarms consume effort while yielding nothing. Across disconnected instances, repetition blurs source clarity. Accuracy in assessing intent falters when relying only on actions. Existing models leave contextual voids untouched. Identical movements appear even when goals differ entirely. Over time, unnoticed errors take deeper root without challenge.

Security measures start by permitting execution solely for authenticated programs. Instead of one-time verification, multiple inspection layers increase confidence step by step. Previous behavior and established reliability influence evaluation of present conduct. Upon triggering notifications, interpretable machine-based logic reveals the basis for judgments. Such visibility enables quicker reassessment of incidents wrongly marked as risky.

Performance Overhead and Scalability

Heavy computational loads arise when analyzing real-time behavioral information. Although advanced models such as integrated CNN-LSTM architectures perform well, their runtime requirements strain available processing power. Continuous surveillance - monitoring active files, executing programs, along with communication patterns - generates sustained pressure on system capacity because of unceasing input volume. Uninterrupted scrutiny, therefore, introduces measurable latency across oversight mechanisms.

At deployment to endpoints, monitoring tools require thoughtful structure to avoid reducing device performance. In enterprise environments, collecting activity records from numerous systems introduces heavier demands on infrastructure below. A possible method places initial threat spotting within individual devices themselves - unusual behaviors alone move forward, flagged for extended analysis through cloud resources. Compact models form via removal of redundant parameters, flattening architectures, or drawing knowledge from bulkier predecessors. Faster processing appears where visual computation units or specialized hardware manage artificial intelligence workloads locally.

Adversarial Machine Learning

As machine learning becomes more common in spotting threats, ransomware developers modify their tactics to escape detection. Adversarial

techniques now challenge automated systems by manipulating how predictions are made. Rather than lock files quickly, some malware proceeds slowly - mirroring normal user behavior step by step. In much the same way, particular types mimic operations of legitimate software to appear ordinary. Meanwhile, alternate forms insert distorted inputs designed to trick analysis tools into overlooking them.

Should virtual machines or debugging software appear, ransomware may shift behavior, executing benign routines instead. Harmful samples entering training data risk eroding system precision gradually. Though deception tactics evolve, exposure to synthetic attack behaviors at initialization can strengthen resistance. Combining dissimilar model structures often yields more resilient outcomes across conditions. Frequent updates help maintain relevance when adversaries adopt novel methods. Behavioral changes in programs triggered by environment cues sometimes reveal threats in unanticipated ways.

Cold Start and Zero-Day Detections

When systems face unknown ransomware types unlike prior examples, recognition becomes difficult. This issue arises due to lack of representative data during development stages. Although monitoring actions instead of fixed codes offers broader reach, some malicious software avoids identification by using unfamiliar tactics or ciphers. Recognition often waits until enough instances appear for system updates. One path involves spotting irregular activities distinct from standard operations, bypassing reliance on established templates. Another applies methods allowing adjustment based on scarce fresh evidence. A further option draws insight from similar threats, supplying initial response ability even without direct exposure.

Integration with Existing Security Infrastructure

Deployment works only when linked properly to current security setups. Where tools already operate, new components fit beside them without disruption. Connections form through shared protocols, allowing one system to respond as another detects change. Real-time performance depends on speed, not just compatibility across platforms. Alerts arrive quickly where data flows freely between systems.

Privacy and Compliance

Tracking actions such as file usage, running programs, or even what files contain may conflict with personal privacy expectations. When applying these methods, institutions face legal frameworks including GDPR, CCPA, and HIPAA that govern data handling. To balance safety needs against individual rights, some approaches keep data on devices instead of sending it elsewhere, strip identities from behavior records without losing analysis value, yet inform those affected about surveillance practices in clear terms.

Response and Remediation

Detection falls short unless paired with swift reaction. When a threat is clearly identified, processes stop without delay through automation. Files that might be harmful are moved aside quickly so they cannot encrypt more data. Devices showing signs of infection get separated from the network to block spread. Evidence gets preserved automatically by capturing system states before changes occur. How much autonomy a system uses ties directly to how certain the identification is. Clear threats lead to instant countermeasures applied by software. Cases lacking clarity pass upward for examination by experts instead. Only when signals strongly point one way does full automation engage.

Future Research Directions

The following are future research directions to research on in the behavioral analysis system:

Explainable AI for Ransomware Detection

One reason security teams struggle is that deep learning works like a hidden mechanism, leaving them unsure about how alerts get triggered. Because of this opacity, spotting the cause behind a flagged event becomes guesswork. To help, tools from explainable AI reveal what behaviors mattered most during analysis - so decisions are no longer invisible. These insights let reviewers check if an alert makes sense, recognize tactics used, or decide whether automation deserves reliance. Moving ahead, progress will depend on shaping those

explanation methods around ransomware needs without weakening accuracy.

Federated Learning for Privacy-Preserving Detection

Federated learning lets groups train models together without handing over personal behavior records, easing worries about privacy yet still tapping into shared attack knowledge. Instead of passing around confidential information, companies might boost their detection systems by exchanging only the tweaks made to those systems - a move that can speed up spotting new ransomware tricks while keeping data under local control and within legal boundaries.

Adaptive and Self-Learning Systems

Reinforcement learning (RL)-based adaptive detection shows particular promise where labelled ransomware samples are scarce. By framing detection as a sequential decision problem, where an agent observes system call sequences and receives rewards for correct classification, Reinforcement learning agents can generalize from simulated environments to real-world deployments. Organizations should schedule adversarial simulation exercises (red team engagements) at minimum bi-annually to generate novel evasion scenarios that feed back into the learning pipeline, preventing model plateau on known attack patterns. Systems of tomorrow must adjust themselves as ransomware changes, requiring no human involvement. One path explores reinforcement learning, letting detection models discover effective reactions by engaging with fake attack scenarios. Instead of fixed rules, some methods rely on constant updates drawn from live streams of behavior patterns. Another route uses meta-learning, which allows quick shifts toward unknown variants using only a few examples. Adaptability emerges not from bulk but from how fast knowledge transfers across rare instances.

Multi-Modal Detection

Combining various detection methods - like monitoring device activity, examining data flows across networks, using current threat reports, and studying how users interact with systems - within a shared structure may lead to more reliable results and stronger resistance to bypass attempts. Starting from endpoint signals, graph-based neural models might map links among devices, running programs, and communication channels, exposing threats that isolated checks would miss. Rather than relying on singular inputs, fusion strategies have potential to include newer types of information: processor usage metrics, sequences of operating system requests, and how software accesses memory regions.

Quantum-Resistant Ransomware Detection

With progress in quantum computing, ransomware could begin using quantum-driven encryption methods. Such a shift might make today's decryption-focused defenses obsolete. Instead of relying on code-breaking approaches, future systems may need to track activity patterns unrelated to cryptography. One possible path forward lies in monitoring how threats behave over time. Unexpected insights may emerge when detection focuses less on data structure, more on sequence anomalies. Quantum-enhanced machine learning might play a role by identifying subtle irregularities across vast datasets. Importance grows steadily for early exploration into non-traditional defense models. Detection strategies that ignore cipher strength entirely may gain relevance. Over time, reliance on classical analysis alone appears riskier. Shifts in computational power demand parallel shifts in defensive logic.

Recommendations for Practitioners

Based on our review, we offer the following recommendations for organizations implementing ransomware detection systems:

1. Adopt layered defense strategies combining behavioral machine learning detection with signature-based antivirus, network security controls, regular backups, and user training. No single approach provides complete protection.
2. Prioritize solutions with low false positive rates (below 0.5%) to maintain analyst effectiveness and user trust. Pilot deployments should carefully measure false positive rates in production environments before full rollout.
3. Implement automated response capabilities for high-confidence detections while maintaining

human oversight for borderline cases. Define clear escalation procedures and response playbooks.

4. Maintain comprehensive, tested backup and recovery procedures as the ultimate safeguard. Even the best detection systems cannot guarantee 100% protection.
5. Regularly update detection models to address emerging ransomware families and evasion techniques. Establish processes for rapid model retraining when new threats emerge.
6. Invest in threat intelligence integration to leverage collective knowledge about ransomware campaigns, indicators of compromise, and attack techniques.
7. Conduct regular tabletop exercises and red team assessments to validate detection capabilities against realistic attack scenarios.
8. For organizations in developing economies with constrained budgets, it's advisable to adopt open-source security stacks as cost-effective alternatives to commercial tooling. Platforms such as Wazuh (EDR), MISP (threat intelligence), and TheHive (incident case management) provide enterprise-grade ransomware detection and response capabilities at minimal licensing cost. This makes layered defense accessible to public-sector institutions, universities, and Small and Medium Scale Enterprises.
9. Invest in cybersecurity workforce development alongside technical controls. Detection systems are only as effective as the analysts interpreting their outputs. Structured capacity-building programmes aligned with internationally recognized certifications (CEH, CompTIA Security+) will ensure that personnel can distinguish genuine incidents from false positives and respond decisively to confirmed ransomware activity. National cybersecurity training initiatives and academic institution partnerships can address the workforce gap at scale.

IV. CONCLUSION

One wrong click can still unravel months of security work; especially as digital hostage attacks keep evolving faster than defenses. Last year, medical providers faced nearly two hundred breaches - each hitting hundreds of thousands, sometimes millions. These aren't just random break-ins anymore; they twist their code constantly, dodge old-style scanners, then spread through ready-made attack kits sold online. Spotting them by known patterns fails now because nothing stays the same long enough to be recognized.

Right off the bat, watching how programs behave - instead of just scanning files - shows real potential when paired with smart algorithms that learn over time. Notably, studies from the past few years point to patterns in behavior that stand out once systems start encrypting data unusually fast. One thing becomes clear after sifting through multiple experiments: odd file access sequences often come up before locks appear. It turns out timing matters too - the speed at which folders get hit can signal trouble ahead. What researchers keep seeing is that processes talking to strange web addresses right before mass changes occur tend to be red flags. Another clue? When apps suddenly stop responding normally while background tasks spike. Some models even catch shifts in memory use that most users would never notice. Taken together, these signs form a trail that newer tools are getting better at following

File changes, how programs run, internet traffic, and encryption actions together help spot ransomware. Mass edits to files often stand out. One after another, files get accessed fast. Gaining higher access levels can be a red flag. Inserting code into running processes raises alarms. Talking to command-and-control servers is suspicious. High randomness inside file data may point to encoding. These signs, seen together, separate real tools from harmful ones. Models that mix convolutional and recurrent layers spot threats extremely well, hitting above 99% correct identifications while staying under 0.2% wrong alarms. What makes them strong is how one part pulls out visual-like patterns, another tracks changes over time - seeing not just single actions but chains of behavior. Learning what matters straight from system activity, without hand-crafted

rules, marks a shift from older techniques built on predefined traits.

Even though tests in labs look good, using this in everyday settings brings big hurdles - like dealing with too many false alarms, slowing systems down, dodging clever attackers, plus spotting new types of ransomwares that haven't been seen before. Making it work means weaving it carefully into current security tools, respecting user privacy and rules, while also ensuring responses can happen without constant human oversight.

Fast changes mark today's ransomware threats, as hackers adopt sharper methods while those fighting back scramble to catch up. Tools driven by artificial intelligence are appearing, adding pressure. Ahead lies uncertainty brought by quantum computing's effect on encryption, pushing experts to rethink how attacks are spotted. New ideas in defense stay essential, shaped by these unfolding shifts.

One step ahead, transparent AI methods need more attention so machines can show how they spot risks. Collaboration without sharing raw data might work better through decentralized learning models. As dangers change, defenses must shift on their own, quietly adjusting behind the scenes. Pulling clues from different kinds of inputs could strengthen overall accuracy. When computers grow powerful enough to crack today's codes, new ways to resist those attacks will matter more than ever.

Here's what matters to those doing the work: spotting threats through behavior-based machine learning works best when it's one part of many, not something standing alone. Protection that actually holds up blends smart detection methods alongside solid backups and ways to restore data quickly; it includes teaching people how to recognize risks, pulling in current threat details, then acting on them clearly when trouble hits - teams need to run through their steps often enough so they know exactly what to do.

Out here, new tricks from ransomware makers push defenders to come up with fresh countermeasures. Because of this back-and-forth, progress doesn't stop - each side adapts fast. Staying ahead means pouring resources into studies that dig deep. People who guard networks must swap findings quickly, no delays. While high-end tools scan for threats, basic habits like updates and access checks still matter just as much. Through it all, one thing stays true - the fight evolves constantly.

A cross-cutting limitation of the reviewed literature is its implicit assumption of high resource deployment environments- modern GPU infrastructure, dedicated SOC teams, and reliable high-bandwidth connectivity. This assumption excludes a substantial portion of the global attack surface, particularly institutions in Africa, Southeast Asia, and Latin America, where ransomware impact is growing fastest yet detection research remains least represented. Future research agendas must deliberately incorporate resource constrained deployment contexts and validate detection architectures under conditions representative of these environments. Collaborative frameworks between African academic institutions, national cybersecurity agencies, and international research bodies can accelerate this inclusive research agenda. Though total security stays out of reach, recent progress in spotting odd behavior through smart algorithms gives teams a real edge against ransomware attacks. Because every method has its weak points and certain needs on the ground, knowing what works - and where it falls short - helps groups pick stronger defenses that fit how they actually operate.

REFERENCES

1. Lashkari, A.H., et al. (2020). CCCS-CIC-AndMal-2020: Android malware dataset. Canadian Institute for Cybersecurity.
2. Economic Community of West African States. (2021). ECOWAS regional cybersecurity and cybercrime strategy. ECOWAS Commission.
3. Chen, Q., & Bridges, R.A. (2023). Automated behavioral malware analysis using machine learning. IEEE Symposium on Security and Privacy, 45–62.
4. Kim, A., et al. (2023). RansomWall: Gradient tree boosting for ransomware detection. ACM

Conference on Computer and Communications Security, 892–908.
5. Moussaileb, R., et al. (2023). Ranker: Kernel-level behavioral analysis for stealthy ransomware detection. Network and Distributed System Security Symposium.
6. National Assembly of the Federal Republic of Nigeria. (2023). Nigeria Data Protection Act 2023. Federal Republic of Nigeria.
7. Wei, J., et al. (2023). Ransomware detection with CNN and SMOTE: Addressing class imbalance. Future Generation Computer Systems, 138, 158–172.
8. Zolfaghari, B., & Koshiba, T. (2023). RISS dataset: Ransomware analysis using Cuckoo sandbox. Data in Brief, 45, 108634.
9. Al-Rimy, B.A.S., et al. (2024). Ransomware detection using behavioral analysis and machine learning: A systematic review. Computers & Security, 138, 103647.
10. Central Bank of Nigeria. (2024). Risk-based cybersecurity framework and guidelines for deposit money banks and payment service banks. CBN.
11. Deloitte Nigeria. (2024). Nigeria cybersecurity outlook 2025. Deloitte Nigeria.
12. Homayoun, S., et al. (2024). iCNN-LSTM: Deep learning approach for ransomware detection using behavioral features. Journal of Cybersecurity and Privacy, 4(1), 89–107.
13. Kara, I., & Aydos, M. (2024). Hybrid LSTM-CNN model for real-time ransomware detection. Computers & Security, 142, 103825.
14. Maniath, S., et al. (2024). Ransomware detection using hardware performance counters. IEEE Transactions on Dependable and Secure Computing, 21(3), 1567–1582.
15. Poudyal, S., et al. (2024). Transfer learning with ResNet50 for image-based ransomware detection. IEEE Access, 12, 23445–23461.
16. Shaukat, K., et al. (2024). Deep learning-based API call analysis for ransomware detection. Expert Systems with Applications, 215, 119342.
17. Singh, J., & Singh, J. (2024). Batch-based incremental learning for ransomware detection. Journal of Information Security and Applications, 72, 103412.
18. U.S. Cybersecurity and Infrastructure Security Agency (CISA). (2024). StopRansomware.gov: Ransomware statistics and trends. Retrieved from
19. Verizon. (2024). 2024 Data Breach Investigations Report. Verizon Enterprise.
20. Zimba, A., et al. (2024). Healthcare sector ransomware: 2024 analysis and impact assessment. Health Information Management Journal.