Network Intrusion Detection Systems: A Comprehensive and Human-Centric Review of Machine Learning and Deep Learning Approaches

Dr. Rolly Gupta, Mohd Hamza, Lakshya Bhatt[⎚] and Avikal Bhatt[†]

Department of Computer Science and Engineering, SRM Institute of Science and Technology, India

Email: rollyg@srmist.edu.in, mh2330@srmist.edu.in, [⋆]lb8774@srmist.edu.in, [†]ab5923@srmist.edu.in

***Abstract*—As the world's dependence on digital communication deepens, the challenge of defending networks against sophisticated intrusions has become one of the defining technical problems of our era. Network Intrusion Detection Systems (NIDS) act as sentinels that analyse traffic to identify threats and anomalies. Yet the classical rule-based or signature-driven methods that once formed the backbone of intrusion detection are increasingly inadequate. Machine Learning (ML) and Deep Learning (DL) offer a paradigm shift, enabling systems to learn complex relationships directly from traffic data and to adapt as attack behaviours evolve.**

This review provides a human-centred synthesis of how ML and DL techniques have been applied to NIDS over roughly the last decade. It describes core algorithms, datasets, and evaluation practices, while also exploring practical questions that matter to real-world deployment—interpretability, resource constraints, and integration into security operations. By combining technical depth with an accessible narrative, we aim to make current research understandable to both practitioners and new researchers, outlining not only what works in the lab but what endures in production.

***Index Terms*—Network security, Intrusion detection, Machine learning, Deep learning, Cyber-threat analysis, Explainable AI**

I. Introduction

The internet has become an indispensable nervous system connecting individuals, businesses, governments, and critical infrastructure. Every online transaction, remote login, or sensor update travels through networks that must remain confidential, reliable, and available. Unfortunately, the same connectivity that enables innovation also provides fertile ground for malicious actors. Attacks now range from simple port scans and brute-force logins to coordinated botnets, ransomware, and stealthy zero-day exploits.

Early Intrusion Detection Systems were designed around manually written rules or known attack signatures. These worked reasonably well when threats changed slowly and traffic volumes were manageable. In the modern environment, however, static signatures cannot keep up with the creativity of attackers or the speed of emerging protocols. False positives overwhelm analysts, and false negatives allow sophisticated breaches to slip through unnoticed.

Machine learning introduces a learning component to this picture. Instead of depending entirely on explicit rules, ML algorithms learn statistical patterns that distinguish normal from abnormal behaviour. Deep learning extends this ability by automatically extracting hierarchical features from raw data, uncovering complex non-linear relationships that traditional models miss.

Yet accuracy alone does not guarantee a successful IDS. A model that performs well on a benchmark may still fail in production if it is too resource-hungry, too opaque for analysts to trust, or trained on unrealistic data. This paper therefore takes a holistic view: evaluating algorithmic progress while keeping sight of the operational realities that determine whether an ML-based IDS truly protects a network.

II. Background and Motivation

Intrusion detection research has a rich and evolving history, dating back to the early 1980s when Dorothy Denning proposed one of the first conceptual frameworks for real-time detection based on system audit logs. Her pioneering idea—that malicious activity could be identified by monitoring deviations from normal behavior—laid the foundation for decades of research that followed.

As computing environments expanded through the 1990s and early 2000s, the focus of intrusion detection shifted from isolated host-based systems toward large-scale **network traffic analysis**. This transition gave birth to landmark contributions such as the **KDD'99 dataset**, which became the cornerstone for early machine learning experiments in intrusion detection. Although simple by modern standards, those early efforts marked the beginning of using **data-driven intelligence** to detect cyber intrusions—an idea that now forms the core of nearly every modern security framework.

The motivation for revisiting this field through a modern lens arises from several converging technological and social trends that have transformed both the scale and the nature of cybersecurity threats:

- **Exploding traffic volume.** The Internet has grown into a dynamic, hyper-connected ecosystem of billions of devices—from smart watches and industrial IoT sensors to autonomous vehicles and cloud data centers. These inter-connected systems generate **terabytes of network data every minute**, rendering manual monitoring and rule maintenance virtually impossible. Traditional signature-based intrusion detection systems simply cannot keep up with the pace, diversity, and volume of this data deluge.
- **Evolving attack tactics.** Cyber adversaries today are far more sophisticated than in the past. They employ **poly-

morphic malware, encrypted communication channels, and multi-stage infiltration tactics** that easily bypass static rule-based detectors. Attacks are often adaptive, stealthy, and distributed across multiple systems. Detecting such threats requires models capable of **learning subtle, temporal, and contextual patterns** that static systems fail to recognize.

- **Rise of AI infrastructure.** The emergence of open-source frameworks such as **TensorFlow, PyTorch, and Scikit-learn** has democratized access to artificial intelligence. Complex architectures once limited to computer vision and natural language processing can now be **seamlessly applied to network data**—analyzing packet flows, connection logs, and time-series traffic with unprecedented precision. This technological accessibility has empowered researchers to explore deep learning-based intrusion detection at a scale never before possible.
- **Need for automation.** The modern Security Operations Centre (SOC) faces an overwhelming flood of alerts and an acute shortage of skilled cybersecurity analysts. Manual triage and log inspection are no longer sustainable. Automated machine learning and deep learning models offer **intelligent prioritization**—filtering false positives, correlating alerts, and helping human analysts focus on truly anomalous or high-risk activities.

These converging forces justify the **renewed global interest in applying machine learning (ML) and deep learning (DL)** to Network Intrusion Detection Systems (NIDS). Beyond their raw detection capability, ML and DL models promise something far more valuable—**adaptability**. Unlike static rule-based systems, these models can learn continuously from evolving network behavior, discovering patterns and correlations that were previously invisible. They bring **scalability**, capable of processing high-speed, high-volume data streams in near real-time, and **resilience**, adapting to emerging threats without the need for exhaustive rule-writing or human intervention.

Yet, amid this enthusiasm, it is essential to approach the field with **critical reflection** rather than unguarded optimism. Despite impressive accuracy metrics reported in recent studies, many existing benchmarks still rely on **synthetic or outdated datasets** such as NSL-KDD, KDD'99, or UNSW-NB15—datasets that no longer mirror the complexity of modern cyber traffic. Moreover, **pre-labeled attacks** in these datasets represent a closed-world assumption that rarely holds true in the real world, where new and unknown attacks surface daily. Models trained under such constraints may perform well in controlled experiments but struggle to generalize in operational settings.

Another recurring issue is the **imbalance in class distribution**—most datasets contain a disproportionately high number of normal traffic samples compared to actual attack instances. Many studies report high overall accuracy while overlooking poor recall for minority attack classes, creating a misleading sense of effectiveness. Equally underexplored is the **computational cost** of these models, including energy consumption, latency, and scalability challenges in real-time network environments.

This review, therefore, seeks not merely to summarize existing methods, but to provide **a human-readable, context-aware analysis** of what each model type truly contributes to the field. By tracing how different algorithmic families—decision trees, support vector machines, convolutional and recurrent networks, hybrid models—approach the same problem, we aim to illuminate their **strengths, weaknesses, and practical trade-offs**. The goal is not to identify a single "best" model but to understand how the discipline has evolved, where it stands today, and what gaps still remain between **academic innovation and deployable, real-world intrusion detection systems**.

Ultimately, the motivation behind this study is rooted in a simple belief: **the future of cybersecurity will depend not just on smarter algorithms, but on systems that think, learn, and adapt as intelligently as the threats they face.** practical limits lie.

III. Methodology of Review

1) A. Review Methodology: A **structured and reproducible review process** was adopted to ensure both the **breadth** and **depth** of coverage across the selected literature. This systematic approach helped in maintaining objectivity, transparency, and consistency throughout the review. The overall workflow, as conceptually represented in *Figure˜**??**{fig:method}*, can be visualized as a sequential yet iterative process encompassing five major stages: **literature search**, **selection**, **data extraction**, **synthesis**, and **qualitative analysis**. Each stage was designed to progressively refine the pool of studies, ensuring that only the most relevant and methodologically sound works were included for deeper investigation.

2) B. Search Strategy: To ensure a comprehensive collection of relevant studies, multiple **digital libraries** were systematically explored, including **IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink,** and **Scopus**. These repositories were chosen because they host a wide range of peer-reviewed research articles spanning computer science, cybersecurity, and data analytics.

A combination of **keywords and logical operators** (AND, OR) was used to form complex queries that could capture the diverse terminology used in the domain. Typical search terms included *"network intrusion detection," "machine learning," "deep learning," "cyber-attack classification,"* and specific **dataset identifiers** such as *"CICIDS2017," "NSL-KDD,"* or *"UNSW-NB15."* This ensured that both generic and dataset-specific studies were retrieved.

The **time window** for the primary search was **2017 to 2020**, a period that witnessed a significant acceleration in the use of deep learning (DL) for intrusion detection applications. However, to maintain historical and conceptual continuity, a small number of **pre-2017** and **post-2020** papers were also reviewed. These outliers provided valuable context about how techniques evolved before and after the DL surge, helping to connect emerging approaches with foundational research.

3) C. Inclusion and Exclusion Criteria: To maintain the quality and relevance of the review, a set of **inclusion and exclusion criteria** was applied after the initial search. Only studies that satisfied all inclusion requirements were retained for detailed analysis.

Included studies were those that:

1) **Proposed an original ML or DL-based model** specifically designed for **network intrusion detection**.
2) **Provided empirical evaluation** using at least **one publicly available dataset**, ensuring reproducibility.
3) **Reported standard performance metrics** such as **accuracy, precision, recall, F1-score,** or **ROC-AUC**, allowing meaningful comparison across studies.

On the other hand, the following studies were **excluded**:

- Works that were **purely theoretical** or conceptual, offering no implementation or empirical testing.
- Studies that **reused public datasets** without introducing **any methodological novelty**.
- Research focusing on **host-based**, **physical**, or **sensor-level intrusion detection** systems, as these fall outside the **network intrusion** domain considered in this review.

This filtration ensured that the final pool of literature consisted only of papers that contributed tangible technical advancements in **Network Intrusion Detection Systems (NIDS)** using ML and DL methods.

4) D. Data Extraction: For every study meeting the inclusion criteria, a **structured data extraction process** was followed. This stage involved creating a detailed summary table capturing multiple aspects of each paper to facilitate systematic comparison. The following attributes were extracted:

1) **Dataset(s) used** – including dataset names, size, type of attacks, and any preprocessing or balancing methods applied.
2) **Algorithmic approach** – specifying whether the model belonged to the **tree-based**, **support vector machine (SVM)**, **convolutional neural network (CNN)**, **recurrent neural network (RNN)**, or **hybrid/ensemble** family.
3) **Feature engineering techniques** – describing how raw network data (e.g., packet headers, flow statistics) were transformed into meaningful input features.
4) **Evaluation metrics** – noting the reported performance results such as **accuracy**, **precision**, **recall**, **F1-score**, **ROC-AUC**, or **detection rate**.
5) **Additional remarks** – highlighting any discussion about **computational efficiency**, **training time**, **model interpretability**, or **scalability** for real-world deployment.

This standardized data extraction framework enabled the review to maintain **objectivity and comparability** across studies that otherwise varied in experimental design and reporting standards.

5) E. Synthesis Approach: Once all relevant data were collected, a **two-level synthesis** approach was performed — combining both **quantitative** and **qualitative** methods.

On the **quantitative** side, studies that used the same datasets (such as NSL-KDD or CICIDS2017) were compared directly in terms of their reported metrics. This allowed an assessment of which algorithmic families or preprocessing techniques yielded the most consistent improvements.

The **qualitative synthesis** focused on identifying **patterns, design strategies, and emerging research themes**. Studies were grouped according to their core algorithmic family — for instance, **tree-based models**, **neural networks**, **hybrid ensembles**, or **unsupervised feature learners**. Within these groups, recurring approaches such as **feature fusion**, **autoencoder-based dimensionality reduction**, or **hybridization** of traditional ML and DL methods were discussed in depth.

Placeholder figure for review methodology diagram.

Fig. 1. Conceptual stages of the systematic review: search, selection, extraction, and synthesis.

Several **trends and observations** emerged from this synthesis:

- A **steady shift toward deep learning architectures** (CNNs, RNNs, and autoencoders) for feature extraction and anomaly detection.
- The **increased adoption of IoT-specific datasets**, reflecting a growing emphasis on securing smart and distributed environments.
- A rising interest in **explainable AI (XAI)**, driven by the need for transparency and trust in automated security systems.

Rather than identifying a single "best-performing" algorithm, the review aimed to **map the evolution of the field**—from classical ML-based systems to sophisticated, data-driven DL architectures. This perspective helped in identifying **existing research gaps**, such as the lack of real-time adaptability, interpretability challenges, and limited cross-dataset generalization. Ultimately, the goal of this systematic review was to **trace the maturation of network intrusion detection research** and to highlight the **barriers that still separate academic prototypes from practical, deployable NIDS solutions** capable of operating efficiently in modern network environments. .

IV. Datasets and Their Characteristics

Machine Learning and Deep Learning models rely entirely on the quality of the data used to train them. In the context of network intrusion detection, datasets act as the foundation upon which every model is built, validated, and benchmarked. Unfortunately, one of the major challenges in IDS research is the absence of a universally accepted, up-to-date dataset that accurately reflects today's network environments.

A dataset for NIDS typically contains captured network traffic represented through packet-level or flow-level records, along with corresponding labels such as *normal*, *DoS*, *probe*, or *malware*. The completeness, accuracy, and realism of these records determine whether an ML or DL model can generalize beyond the training data.

A. Historical Evolution of IDS Datasets

The journey of publicly available IDS datasets began in the late 1990s with DARPA98 and the KDD Cup'99 datasets. They represented a breakthrough: a standardized corpus of attack and normal traffic traces. However, these early datasets quickly became outdated as network technologies evolved. They used simulated environments with limited protocols, unrealistic traffic distributions, and repeated records that bias learning algorithms.

To address these flaws, researchers introduced NSL-KDD, a refined version of KDD'99 that removes duplicates and balances class proportions. Although it remains a popular benchmark due to its accessibility, NSL-KDD still represents network behaviour from an era when most traffic was unencrypted and attack patterns were relatively simple.

The next major step came with the UNSW-NB15 dataset, released in 2015 by the Australian Centre for Cyber Security. It introduced nine modern attack families and 49 features derived from real network traffic captured at a university network. Around the same period, the Canadian Institute for Cybersecurity produced the CICIDS2017 dataset, followed by CSE-CIC-IDS2018. These datasets attempted to mimic real-world enterprise traffic with both benign and malicious flows, covering a wide range of attacks such as DDoS, brute-force, infiltration, and botnet activity.

TABLE I
COMPARATIVE SUMMARY OF COMMONLY USED NIDS DATASETS

Dataset	Period/Source	Attack Types	Features
NSL-KDD	2009, simulated LAN	DoS, Probe, R2L, U2R	41
UNSW-NB15	2015, ACCS	9 types incl. exploits	49
CICIDS2017	2017, CIC	DDoS, brute-force, infiltration, botnet	80+
BoT-IoT	2019, IoT testbed	DoS, theft, exfiltration	46
TON_IoT	2020, IoT logs	Multi-source, cross-layer	40+

B. Emergence of IoT and Edge-Oriented Datasets

With the proliferation of Internet of Things (IoT) devices, new datasets have been curated to capture threats unique to constrained networks. The BoT-IoT dataset (2019) focuses on IoT-oriented attacks such as data exfiltration, DoS, and information theft using lightweight protocols like MQTT. More recently, the TON_IoT dataset extended this work by collecting telemetry data not only from network flows but also from sensors and logs across multiple IoT layers.

C. Key Characteristics and Attributes

Commonly extracted features include:

- **Basic TCP/IP features:** source/destination IP, ports, protocol type, and service.
- **Time-based features:** duration, packets per second, average packet size.
- **Statistical features:** byte count, flag counts, connection ratios.
- **Content features:** payload strings when available.

D. Dataset Limitations

Even with advances, existing datasets share several weaknesses:

1) Lack of realism from lab-generated traffic.
2) Severe class imbalance between benign and attack flows.
3) Reproducibility issues caused by feature-drift.
4) Privacy constraints limiting public release of modern captures.
5) Increasing encryption that hides payload content.

E. Comparison of Common Datasets

F. Dataset Design Considerations

Future **datasets** play a crucial role in advancing the reliability and adaptability of **Network Intrusion Detection Systems (NIDS)**. For these systems to remain relevant and effective, datasets should not be static snapshots of outdated network behavior but rather **living, evolving resources** that grow alongside technological progress and emerging threat landscapes.

To achieve this, future datasets must be **continuously updated** to capture the **latest attack vectors**, network protocols, and traffic patterns that characterize today's rapidly changing digital ecosystem. Cyber threats evolve almost daily, and models trained on stale or obsolete data struggle to recognize new or modified attack behaviors. Regular updates ensure that detection models remain capable of identifying both known and previously unseen (zero-day) attacks with greater accuracy.

Equally important is the need for **privacy preservation** during data collection and dissemination. Real-world network traffic often contains sensitive personal or organizational information. Therefore, datasets must employ **robust anonymization and encryption techniques** to protect user identities, IP addresses, and confidential communications without compromising the data's analytical value. This balance between **data utility and data privacy** is essential to maintain ethical standards and encourage collaboration among researchers who might otherwise hesitate to share real-world data due to privacy concerns.

Another critical requirement is that datasets should be **verified and validated by domain experts** before public release. Expert validation ensures that the dataset accurately represents true attack behaviors rather than synthetic or mislabeled events. Such verification minimizes inconsistencies, labeling errors, and redundancies that can mislead model training and evaluation. Expert-reviewed datasets contribute significantly to the **trustworthiness and scientific credibility** of intrusion detection research.

In addition, future datasets should aim for **contextual diversity**, reflecting the broad spectrum of modern computing environments—**cloud infrastructures, mobile devices, Internet of Things (IoT) networks, edge computing nodes,** and **industrial control systems (ICS)**. Each of these domains exhibits distinct traffic patterns, security requirements, and potential vulnerabilities. Including such diversity helps build models that are **generalizable and adaptive**, capable of functioning effectively across multiple real-world scenarios rather than being limited to narrow or laboratory-specific settings.

To further promote **reproducibility and fairness**, datasets should be **distributed with standardized train/test splits**. Inconsistent splitting practices across studies often lead to

incomparable results and inflated performance claims. By providing official benchmark partitions, researchers can evaluate their algorithms under identical conditions, allowing for fair, transparent, and meaningful comparison of results. Standardization in this area fosters collaboration and accelerates collective progress by ensuring that improvements in accuracy or efficiency reflect true methodological advancements rather than variations in experimental setup.

In summary, the future of NIDS research depends heavily on the **quality, relevance, and ethical design of datasets**. Datasets that are **continuously updated, privacy-preserving, expert-validated, contextually diverse,** and **distributed with standard splits** will serve as the backbone of credible, reproducible, and deployable intrusion detection solutions. Such efforts will bridge the gap between experimental research and real-world application—paving the way for intelligent, trustworthy, and adaptive cybersecurity systems that evolve in tandem with the threats they are designed to detect.

G. Conclusion of Dataset Discussion

Datasets form the **invisible backbone** of all Intrusion Detection System (IDS) research. They are the **silent enablers** behind every breakthrough in cybersecurity intelligence, quietly shaping the trajectory of algorithmic progress. In truth, many of the apparent leaps in detection accuracy or robustness over the years have not stemmed solely from revolutionary models, but rather from **better-curated, more representative datasets** that allow these models to truly learn the complex behavior of network traffic. The **quality, diversity, and freshness** of data often define the ceiling of what an IDS can achieve.

A well-designed dataset serves as the **mirror of the digital world**, capturing its complexity, unpredictability, and evolving threat landscape. It provides the contextual foundation upon which algorithms are trained to recognize what is normal and what is malicious. When datasets fail to reflect real-world network conditions—whether by being outdated, biased, or incomplete—the resulting models become **blind to modern attack patterns**, detecting yesterday's threats while missing today's. Thus, even the most sophisticated deep neural networks, stacked with millions of parameters, are rendered ineffective if they are trained on **stagnant or unrepresentative data**.

Equally vital is the **ethical stewardship of data**. In an era where privacy breaches are as damaging as the attacks themselves, it is imperative that datasets are collected and shared in a manner that **respects privacy, preserves anonymity, and ensures compliance** with ethical standards. True progress in IDS research cannot come at the cost of user trust. The datasets of the future must therefore integrate **privacy-preserving mechanisms** such as anonymization, encryption, and controlled access — balancing the dual responsibility of fostering innovation and safeguarding individual rights.

Moreover, datasets should evolve to embody the **diversity of today's digital ecosystem**. Modern networks are no longer confined to traditional desktops and servers—they extend into **cloud infrastructures, IoT devices, mobile networks, smart homes, and industrial systems**. Each of these environments generates distinct traffic behaviors and unique security challenges. By incorporating this diversity, future datasets can support the development of **more resilient and adaptive IDS models**, capable of handling the heterogeneity of real-world networks rather than thriving only in isolated laboratory conditions.

Finally, the community must move toward **standard- ized dataset protocols**—including predefined train-test splits, clearly documented labeling processes, and benchmark scenarios—to ensure that comparisons between different IDS models are **fair, transparent, and reproducible**. Without these shared standards, progress becomes fragmented, and the scientific merit of new methods remains uncertain.

In essence, datasets are not merely supporting tools; they are the **lifeblood of innovation** in intrusion detection. The future of IDS will depend not only on more powerful algorithms, but on **how truthfully and ethically we capture the world they are meant to defend**. Only when datasets evolve in parallel with technology and threats can IDS systems become genuinely intelligent, trustworthy, and future-ready—capable of protecting digital spaces with the same agility with which those spaces evolve.

V. Future Scope

The evolution of Network Intrusion Detection Systems (NIDS) is far from complete. While the past decade has demonstrated remarkable progress through the integration of machine learning (ML) and deep learning (DL), the road ahead presents both opportunities and challenges that demand a more holistic and human-centered approach. The next generation of intrusion detection must move beyond accuracy metrics and focus on **trust, adaptability, and ethical intelligence**—qualities that will define the true maturity of cybersecurity research.

A. Realistic and Evolving Datasets

One of the most urgent needs for the future is the creation of **dynamic, continuously updated datasets** that mirror the complexity of real-world network environments. Current benchmarks such as NSL-KDD or CICIDS2017, though valuable, represent only static snapshots of a rapidly evolving threat landscape. To ensure genuine progress, future datasets must integrate **live, real-time traffic**, **emerging attack types**, and **cross-domain scenarios** such as IoT, mobile, and cloud ecosystems. These datasets should also prioritize **privacy preservation** through techniques like **differential privacy** and **data anonymization** to encourage open collaboration without compromising user confidentiality. Moreover, expert verification and community-driven validation should become standard practice, ensuring data authenticity and consistency across research works [1], [2].

B. Explainability and Trustworthy AI

As ML and DL models grow in complexity, they often become **black boxes**—highly accurate but poorly understood. Future NIDS research must focus on **explainable AI (XAI)** frameworks that make model decisions transparent to human operators. An explainable system not only improves trust but also aids analysts in understanding why a certain connection or

packet was flagged as suspicious. Techniques such as **Layer-wise Relevance Propagation (LRP)**, **SHAP values**, or **LIME-based interpretations** can be integrated into IDS pipelines to reveal how features contribute to classifi- cation decisions. This shift from opaque prediction to human-comprehensible reasoning will play a crucial role in bridging the gap between automation and accountability [3], [4].

C. Lightweight and Energy-Efficient Models

With the expansion of IoT and edge computing, intrusion detection cannot remain confined to powerful cloud servers. The future will demand **lightweight, energy-efficient NIDS models** capable of running on low-power devices without sacrificing accuracy. Research in **model compression**, **quantization**, and **federated learning** offers promising directions for achieving scalable and distributed detection frameworks. By decentralizing learning, networks can detect anomalies closer to their source, minimizing latency and improving response time while maintaining data privacy. Such innovations are critical for sustaining cybersecurity in **resource-constrained, ubiquitous computing environments** [5], [6].

D. Integration of Multimodal and Behavioral Data

Traditional intrusion detection has primarily relied on network packet and flow data. However, the future lies in integrating **multimodal information**—including system logs, user behavior, application telemetry, and even threat intelligence feeds from external sources. Fusing these diverse data streams can provide a **360-degree situational awareness** of network health, enabling early detection of coordinated or insider attacks that single-source models often miss. Behavioral analytics, powered by unsupervised learning and temporal modeling, could allow NIDS to **anticipate** threats before they manifest, marking a transition from reactive defense to **proactive security intelligence**.

E. Human-in-the-Loop Learning

Despite automation's advantages, the human element remains irreplaceable in cybersecurity. Future systems should incorporate **human-in-the-loop mechanisms**, where analysts can guide or correct machine learning models interactively. This synergy between human intuition and machine precision can lead to more robust detection pipelines that learn from real operational feedback rather than static datasets. Such collaboration will also foster **continual learning**, where systems evolve over time, adapting to new threats as they appear in the wild—a form of "living AI" that grows with the network it protects [7].

F. Ethical and Global Collaboration

Finally, the future of intrusion detection must embrace a sense of **global collaboration and ethical responsibility**. Cyber threats are borderless, and so must be our defense strategies. Collaborative data-sharing frameworks, open-source benchmarking initiatives, and transparent evaluation standards can transform IDS research from fragmented individual efforts into a **cohesive global movement**. Moreover, ethical considerations—such as fairness, non-discrimination, and responsible AI deployment—must be woven into the design of future models. A truly intelligent IDS is not only one that detects threats efficiently, but one that does so **responsibly, transparently, and in harmony with human values**.

G. Looking Ahead

In conclusion, the future of NIDS research is about more than technical innovation—it is about **rethinking security as an evolving, adaptive, and ethical discipline**. By uniting cutting-edge algorithms with human understanding, fostering collaboration across academia and industry, and ensuring data diversity and transparency, the next generation of intrusion detection systems can move closer to achieving what has long been the dream of cybersecurity: a defense mechanism that learns, reasons, and protects—just as intelligently as the adversaries it seeks to outsmart.

References (sample placeholders):

REFERENCES

[1] M. Ring, S. Wunderlich, D. Gru¨ dl, D. Landes, and A. Hotho, "Flow-based network traffic generation using Generative Adversarial Networks," *Computers Security*, vol. 82, pp. 156–172, 2019.

[2] N. Moustafa, J. Slay, and G. Creech, "Novel geometric area analysis technique for anomaly detection using trapezoidal area estimation on large-scale networks," *IEEE Transactions on Big Data*, 2019.

[3] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv preprint arXiv:1702.08608*, 2017.

[4] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Mu¨ller, "Causability and explainability of artificial intelligence in medicine," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 2020.

[5] T. Li, A. S. Sahu, M. Zaheer, et al., "Federated Learning: Challenges, Methods, and Future Directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.

[6] K. Bonawitz et al., "Towards Federated Learning at Scale: System Design," *Proceedings of Machine Learning and Systems*, 2019.

[7] S. Amershi, A. Begel, C. Bird, et al., "Software Engineering for Machine Learning: A Case Study," *Proceedings of the 41st International Conference on Software Engineering*, 2019.