PEOPLE'S DEMOCRATIC REPUBLIC OF ALGERIA
MINISTRY OF HIGHER EDUCATION AND SCIENTIFIC RESEARCH
UNIVERSITY OF MOHAMED BOUDIAF-MSILA
FACULTY OF MATHEMATICS AND COMPUTER SCIENCE
DEPARTMENT OF COMPUTER SCIENCE

N°: ……

THESIS

Submitted in Partial Fulfilment of the Requirements for the Degree of 3rd cycle Doctorate in Computer Science

Speciality: Computer Science

Option: Security of Computer Systems and Networks

By: Abdelhammid Bouazza

Entitled as:

Distributed, Fault Tolerant and Secure Time Allocation Algorithms for Scalable Communication in the Internet of Things

---

Publicly defended on: 29 /06 /2026

Board of Examiners:

M. Sayad Lamri	Prof.	University of M'sila	President
M. Debbi Hichem	Prof.	University of M'sila	Supervisor
M. Dabba Ali	MCA	University of M'sila	Examiner
M. Zouache Djaafar	Prof.	University of Bordj Bou Arréridj	Examiner
Mme. Aliouat Zibouda	Prof.	University of Sétif 1	Examiner
M. Lakhlef Hicham	Prof.	University of Bordeaux, France	Co-supervisor

Academic Year: 2025/2026

Dedication

"To those whose light illuminates my path…"

To my beloved parents,

whose prayers, sacrifices, and unwavering faith
have been the foundation of every achievement,
and the light guiding every step of this journey.

To my dear sisters and brother,
for their boundless love, encouragement,
and the quiet strength they have always offered.

To every teacher who lit the path of knowledge,
and to every researcher who dared to ask why…

— With gratitude —

Acknowledgements

First and foremost, I praise and thank **Almighty Allah**, the Most Gracious and the Most Merciful, for granting me the health, strength, patience, and perseverance necessary to complete this doctoral research. All praise is due to Him, without whose blessings and guidance this work would not have been possible.

The completion of this thesis would not have been possible without the guidance, support, and encouragement of many individuals, to whom I wish to express my heartfelt gratitude.

My deepest appreciation goes to my supervisor, **Prof. Debbi Hichem**, for his exceptional guidance, unwavering support, and scholarly vision throughout this research. His rigorous scientific approach, patience, and insightful observations have played a decisive role in shaping both this thesis and my growth as an independent researcher.

I am equally grateful to my co-supervisor, **Prof. Lakhlef Hicham**, for his invaluable mentorship and profound expertise in the field of the Internet of Things. His constructive feedback and constant involvement have greatly strengthened the scientific depth and rigor of this work.

I would also like to express my sincere appreciation to the **President of the Jury** for the honor of chairing the defense committee, as well as to the distinguished **members of the examining committee** for agreeing to evaluate this thesis. Their thoughtful remarks, relevant questions, and valuable suggestions have contributed significantly to improving the quality of this work.

I also wish to acknowledge the **University of Mohamed Boudiaf – M'sila**, the **Faculty of Mathematics and Computer Science**, and the **Department of Computer Science**. My thanks are equally extended to the teaching and administrative staff for their support and assistance throughout these years.

Finally, I offer my deepest gratitude to my **family** for their unconditional love, patience, and constant encouragement throughout my academic journey. Their unwavering faith in me has been a continuous source of strength and motivation. To all those who contributed, directly or indirectly, to the realization of this work, I extend my sincere thanks.

ABDELHAMMID Bouazza

Abstract

The Internet of Things (IoT) enables large-scale networks of resource-constrained devices that continuously generate and exchange data across critical domains such as healthcare, industry, and smart infrastructure. Ensuring scalable communication in such networks requires the joint optimization of resource efficiency, intrusion detection, privacy preservation, and routing reliability. This thesis addresses these challenges by proposing a unified intelligent framework for distributed, secure, and fault-tolerant communication in IoT, with the Internet of Medical Things (IoMT) adopted as the primary validation domain. The first contribution introduces a Time Allocation strategy based on dynamic Temporal Aggregation, in which the aggregation window T_{agg} is adaptively selected to transform raw data streams into compact feature vectors before transmission. This strategy significantly reduces computational cost, achieving training-time reductions of up to 97.2% and inference-time reductions of up to 97.5% on CICIoMT-2024, with comparable improvements on NF-UNSW-NB15-v2 and WUSTL-EHMS-2020, thereby enhancing scalability and contributing to lower energy consumption in resource-constrained IoT environments. The second contribution proposes FTL-HLSTM, a Federated Transfer Learning architecture based on Hierarchical Long Short-Term Memory networks for privacy-preserving intrusion detection. The framework addresses the Non-IID nature of distributed IoT data by distinguishing between globally shared and locally specific attack patterns, enabling collaborative learning without exchanging raw data. The model achieved 100.0% binary detection accuracy on the evaluated NF-UNSW-NB15-v2 split and 99.74% accuracy on CICIoMT-2024, while reducing training time compared with standard LSTM models. The third contribution develops a proactive fault-tolerant routing mechanism based on Multi-Criteria Decision Analysis using TOPSIS and AHP. The proposed approach computes a dynamic Trust Score for each node according to safety, energy, latency, and packet-loss criteria. Unlike reactive routing protocols, it excludes non-acceptable nodes from routing decisions and adaptively distributes traffic by assigning 70% to optimal nodes and 30% to acceptable nodes, thereby improving routing reliability under node failures and insider attacks. Finally, these components are integrated into a unified cyber-physical defense system in which intrusion detection results directly support fault-tolerant routing decisions. This closed-loop architecture enables autonomous self-protection and self-healing in IoT networks by linking detection intelligence with adaptive network control.

Keywords: Internet of Things, Internet of Medical Things, Scalable Communication, Temporal Aggregation, Federated Transfer Learning, Intrusion Detection, Fault Tolerance, MCDA-based Routing.

ملخص

تتيح إنترنت الأشياء (IoT) إنشاء شبكات واسعة النطاق من الأجهزة محدودة الموارد، والتي تولد وتتبادل البيانات باستمرار عبر مجالات حيوية مثل الرعاية الصحية، والصناعة، والبنى التحتية الذكية. ويتطلب ضمان اتصال قابل للتوسع في مثل هذه الشبكات تحسين كفاءة استخدام الموارد، والكشف عن الاختراقات، والحفاظ على الخصوصية، وموثوقية التوجيه. تعالج هذه الأطروحة هذه التحديات من خلال اقتراح إطار ذكي موحد للاتصال الموزع، الآمن، والمتسامح مع الأخطاء في إنترنت الأشياء، مع اعتماد إنترنت الأشياء الطبية (IoMT) كمجال رئيسي للتحقق والتقييم. تتمثل المساهمة الأولى في اقتراح استراتيجية لتخصيص الوقت قائمة على التجميع الزمني الديناميكي، حيث يتم اختيار نافذة التجميع T_{agg} بشكل تكيفي لتحويل تدفقات البيانات الخام إلى متجهات خصائص مضغوطة قبل إرسالها. تقلل هذه الاستراتيجية بشكل ملحوظ من التكلفة الحسابية، حيث حققت تخفيضا في زمن التدريب يصل إلى 97.2% وتخفيضا في زمن الاستدلال يصل إلى 97.5% على مجموعة بيانات CICIoMT-2024، مع تحسينات مماثلة على مجموعتي NF-UNSW-NB15-v2 و WUSTL-EHMS-2020. وتسهم هذه التخفيضات في تعزيز قابلية التوسع وتقليل استهلاك الطاقة في بيئات إنترنت الأشياء محدودة الموارد. تتمثل المساهمة الثانية في اقتراح FTL-HLSTM، وهي معمارية تعلم اتحادي بالنقل قائمة على شبكات الذاكرة طويلة قصيرة الأمد الهرمية، موجهة للكشف عن الاختراقات مع الحفاظ على الخصوصية. يعالج هذا الإطار طبيعة البيانات غير المستقلة وغير المتطابقة التوزيع (Non-IID) في بيئات إنترنت الأشياء الموزعة، من خلال التمييز بين أنماط الهجمات المشتركة والأنماط المحلية الخاصة، مما يتيح التعلم التعاوني دون تبادل البيانات الخام. حقق النموذج دقة كشف ثنائية بلغت 100.0% على مجموعة البيانات NF-UNSW-NB15-v2، ودقة بلغت 99.74% على مجموعة CICIoMT-2024، مع تقليل زمن التدريب مقارنة بنماذج LSTM القياسية. تتمثل المساهمة الثالثة في تطوير آلية توجيه استباقية متسامحة مع الأخطاء، قائمة على تحليل القرار متعدد المعايير باستخدام طريقتي TOPSIS و AHP. تحسب الطريقة المقترحة درجة ثقة ديناميكية لكل عقدة اعتمادا على معايير السلامة، والطاقة، وزمن التأخير، وفقدان الحزم. وخلافا لبروتوكولات التوجيه التفاعلية، تستبعد هذه الآلية العقد غير المقبولة من قرارات التوجيه، وتوزع حركة المرور تكيفيا بتخصيص 70% للعقد المثلى و 30% للعقد المقبولة، مما يحسن موثوقية التوجيه في ظل أعطال العقد والهجمات الداخلية. أخيرا، يتم دمج هذه المكونات ضمن نظام دفاع سيبراني-فيزيائي موحد، حيث تدعم نتائج الكشف عن الاختراقات قرارات التوجيه المتسامح مع الأخطاء بشكل مباشر. تتيح هذه البنية ذات الحلقة المغلقة تحقيق الحماية الذاتية والإصلاح الذاتي في شبكات إنترنت الأشياء، من خلال الربط بين ذكاء الكشف والتحكم التكيفي في الشبكة.

الكلمات المفتاحية: إنترنت الأشياء، إنترنت الأشياء الطبية، الاتصال القابل للتوسع، التجميع الزمني، التعلم الاتحادي بالنقل، كشف الاختراقات، التسامح مع الأخطاء، التوجيه القائم على تحليل القرار متعدد المعايير.

Contents

List of Figures

List of Tables

List of Abbreviations

AHP	Analytic Hierarchy Process
AODV	Ad hoc On-Demand Distance Vector
BFT	Byzantine Fault Tolerance
BLE	Bluetooth Low Energy
BPTT	Backpropagation Through Time
CI	Consistency Index
CICIoMT	Canadian Institute for Cybersecurity Internet of Medical Things (Dataset)
CNN	Convolutional Neural Network
CPS	Cyber-Physical Systems
CR	Consistency Ratio
DDoS	Distributed Denial of Service
DL	Deep Learning
DNN	Deep Neural Network
DODAG	Destination-Oriented Directed Acyclic Graph
DoS	Denial of Service
DP	Differential Privacy
ECG	Electrocardiogram
EHR	Electronic Health Record
FAR	False Alarm Rate
FedAvg	Federated Averaging
FL	Federated Learning
FTL	Federated Transfer Learning
FTL-HLSTM	Federated Transfer Learning with Hierarchical LSTM
FTL-TSLP	Federated Transfer Learning with Two-Stage LSTM Pipeline
GDPR	General Data Protection Regulation
GRU	Gated Recurrent Unit
HIPAA	Health Insurance Portability and Accountability Act
HLSTM	Hierarchical Long Short-Term Memory
IDS	Intrusion Detection System
IID	Independent and Identically Distributed
IoMT	Internet of Medical Things
IoT	Internet of Things

KNN	K-Nearest Neighbours
LSTM	Long Short-Term Memory
MAC	Medium Access Control
MCDA	Multi-Criteria Decision Analysis
MITM	Man-In-The-Middle
ML	Machine Learning
MLP	Multi-Layer Perceptron
MQTT	Message Queuing Telemetry Transport
MTTF	Mean Time To Failure
MTTR	Mean Time To Repair
NF-UNSW	NetFlow University of New South Wales (Dataset)
NIDS	Network Intrusion Detection System
NIS	Negative Ideal Solution
Non-IID	Non-Independent and Identically Distributed
PBFT	Practical Byzantine Fault Tolerance
PIS	Positive Ideal Solution
QoS	Quality of Service
RF	Random Forest
RI	Random Index
RNN	Recurrent Neural Network
RPL	Routing Protocol for Low-Power and Lossy Networks
SMOTE	Synthetic Minority Over-sampling Technique
SVM	Support Vector Machine
T_{agg}	Temporal Aggregation Window
TDMA	Time Division Multiple Access
TL	Transfer Learning
TOPSIS	Technique for Order of Preference by Similarity to Ideal Solution
WBAN	Wireless Body Area Network
WHO	World Health Organization
WSN	Wireless Sensor Network
WUSTL-EHMS	Washington University in St. Louis – Electronic Health Monitoring System (Dataset)

List of Publications

Journal Articles

[J1] **A. Bouazza**, H. Debbi, and H. Lakhlef, "FTL-TSLP: A Federated Transfer Learning Approach with a Two-Stage LSTM Pipeline for Fault-Tolerant and Privacy-Preserving Intrusion Detection in IoMT Networks," *Internet of Things* (Elsevier), vol. 29, p. 101832, 2025.
DOI: https://doi.org/10.1016/j.iot.2025.101832

International Conference Proceedings

[C1] **A. Bouazza**, H. Debbi, and H. Lakhlef, "Ensuring IoT System Fault Tolerance Using Deep Learning and Multi-Criteria Decision Analysis," in *Proceedings of the 39th International Conference on Advanced Information Networking and Applications* (AINA 2025), L. Barolli (Ed.), Lecture Notes on Data Engineering and Communications Technologies, vol. 247, pp. 14–25. Springer, Cham, 2025.
DOI: https://doi.org/10.1007/978-3-031-87769-8_2

[C2] **A. Bouazza**, H. Debbi, H. Lakhlef, and A. Smaili, "An Efficient Hierarchical LSTM-based Framework for Intrusion Detection in Internet of Things (IoT) Systems," in *Proceedings of the 31st International Conference on Software, Telecommunications and Computer Networks* (SoftCOM 2023), pp. 1–6. IEEE, Split, Croatia, 2023.
DOI: https://doi.org/10.23919/SoftCOM58365.2023.10271665

[C3] **A. Bouazza**, H. Debbi, and H. Lakhlef, "Machine Learning-based Intrusion Detection System Against Routing Attacks in the Internet of Things," in *Proceedings of the Tunisian-Algerian Joint Conference on Applied Computing* (TACC 2022), CEUR Workshop Proceedings, vol. 1613, p. 0073, 2022.

[C4] **A. Bouazza**, H. Debbi, and H. Lakhlef, "Two-Layer Adaptive Fault Tolerance Framework for Internet of Medical Things: A Weighted Autoencoder with Reinforcement Learning-Based Threshold Adaptation," submitted to the *45th International Conference on Computer Safety, Reliability and Security* (SAFECOMP 2026), Lecture Notes in Computer Science, Springer. Valencia, Spain, September 2026.

General Introduction

General Introduction

Ensuring optimal health is fundamental to human well-being, particularly as modern healthcare faces increasing demands for continuous, real-time patient monitoring. The World Health Organization (WHO) defines health as a state of complete physical, mental, and social well-being—not merely the absence of disease. Achieving this vision in practice requires healthcare systems that can monitor patients continuously, detect anomalies early, and enable timely interventions across distributed clinical environments.

The Internet of Things (IoT) has fundamentally transformed the healthcare industry by enabling interconnected networks of sensors and devices to collect, process, and transmit physiological data in real time. Within this broader ecosystem, the Internet of Medical Things (IoMT) has emerged as a critical domain, representing the convergence of biomedical engineering, wireless telecommunications, and advanced data analytics. Through the deployment of Wireless Body Area Networks (WBANs), vital physiological metrics—including cardiac rhythms, blood pressure, glucose levels, and respiratory rates—are continuously monitored, digitized, and transmitted to remote healthcare providers. According to recent industry projections, the number of connected IoMT devices is expected to reach tens of billions within the next decade, generating unprecedented volumes of medical telemetry data.

However, the realization of this vision is hindered by the inherent constraints of medical sensing devices and the critical nature of the data they generate. As the density of connected sensors increases, particularly in high-occupancy environments such as hospitals and elderly care facilities, the network confronts a growing challenge: managing the escalating volume of data transmissions within the limited computational, energetic, and bandwidth resources available at the network edge. Consequently, the challenge for modern computer science extends beyond merely establishing connectivity—it requires ensuring that this connectivity is scalable, secure, and fault-tolerant under strict real-time constraints.

Problem Statement

The integration of critical healthcare services into open wireless networks introduces a complex set of interrelated challenges that existing protocols fail to address simultaneously. These challenges can be articulated along three axes:

The Scalability Challenge (Time and Resource Allocation). Standard IoT

communication protocols transmit raw telemetry streams continuously, without considering the temporal dynamics of the data. In high-density IoMT networks, this approach leads to rapid bandwidth saturation and accelerated battery depletion of constrained sensor nodes. The fundamental problem lies in the absence of intelligent *Time Allocation* strategies—mechanisms capable of determining the optimal temporal granularity of data transmission to balance information fidelity with resource consumption. Concretely, in this thesis, Time Allocation refers to the strategic optimization of the temporal aggregation interval (T_{agg})—the window over which raw sensor readings are collected, summarized into feature vectors, and then transmitted—rather than a MAC-layer scheduling mechanism such as TDMA. Without such strategies, the network cannot scale to accommodate the growing number of medical devices without significant degradation in performance and reliability.

The Security and Privacy Paradox. IoMT data is inherently sensitive, as it encompasses protected health information subject to strict privacy regulations. Paradoxically, this data must also be inspected and analyzed to detect network intrusions and cyberattacks. Traditional Intrusion Detection Systems (IDS) are predominantly centralized, thereby creating single points of failure and exposing raw patient data to privacy risks during analysis. Furthermore, medical data collected across different hospitals and clinical environments is inherently Non-Independent and Identically Distributed (Non-IID): an attack pattern observed in one hospital may differ significantly from that in another, rendering standard Federated Learning (FL) approaches prone to significant degradation in model generalization and convergence speed when applied directly to heterogeneous IoMT environments.

The Reliability Imperative (Fault Tolerance). In a conventional network, a dropped packet constitutes an inconvenience; in IoMT, it can have life-threatening consequences. Existing routing protocols, such as the Routing Protocol for Low-Power and Lossy Networks (RPL) or Ad hoc On-Demand Distance Vector (AODV), are largely reactive—they initiate repair mechanisms only after a link failure has occurred. In a life-critical environment, such latency is unacceptable. The network requires a proactive mechanism capable of anticipating node failures and isolating compromised nodes before they disrupt the communication path.

Collectively, these three challenges form an interconnected trilemma: improving scalability without compromising security, enhancing security without sacrificing the reliability of data routing, and ensuring fault tolerance without introducing prohibitive resource overhead. Addressing these challenges in isolation is insufficient; what is required is a unified, intelligent framework that resolves them concurrently.

This leads to the central research question of this thesis:

How can we design a distributed, intelligent framework that optimizes Time Allocation for scalable communication, while concurrently ensuring robust security against sophisticated intrusions and proactive fault tolerance in the heterogeneous Internet of Medical Things?

Goals and Contributions

The primary aim of this thesis is to address the aforementioned trilemma by proposing a unified Intelligent IoMT Framework that jointly optimizes communication scalability, intrusion detection, and fault-tolerant routing. The proposed methodologies are designed to accommodate the resource-constrained nature of medical sensor nodes while overcoming the limitations of existing approaches. To validate the proposed framework, we conducted extensive experiments using specialized IoMT datasets (CICIoMT-2024, WUSTL-EHMS-2020), evaluating detection accuracy, routing resilience, and computational overhead under heterogeneous and adversarial conditions. In pursuit of this objective, we present four distinct contributions:

- **Contribution 1: Scalability via Strategic Time Allocation.** We introduce a novel data processing methodology based on Temporal Aggregation, where the aggregation window (T_{agg}) is dynamically optimized. By transforming raw telemetry streams into temporally aggregated feature vectors, we achieve a significant reduction in bandwidth usage and energy consumption, demonstrating that intelligent time allocation is a prerequisite for IoMT scalability.
- **Contribution 2: The FTL-HLSTM Framework for Non-IID Intrusion Detection.** We address the critical limitation of standard Federated Learning in heterogeneous medical environments. By integrating Transfer Learning with Hierarchical Long Short-Term Memory (HLSTM) networks, our model successfully distinguishes between *Common* (global) and *Isolated* (local) attack patterns, achieving superior detection accuracy while preserving strict data privacy across hospital sites. Furthermore, the proposed architecture is designed to be compatible with additional privacy-preserving mechanisms such as Secure Aggregation and Differential Privacy, ensuring extensibility to stricter regulatory requirements.
- **Contribution 3: Proactive Fault Tolerance via Multi-Criteria Decision Analysis.** Moving beyond conventional shortest-path routing logic, we propose a routing framework based on TOPSIS and AHP that calculates a dynamic *Trust Score* for every network node using multidimensional metrics encompassing safety, energy, and latency. Unlike reactive protocols, our system proactively

isolates nodes exhibiting degrading performance or suspicious behavior, ensuring high network availability even in hostile environments.

- **Contribution 4: A Unified Cyber-Physical Defense System.** A distinguishing feature of this work is the integration of deep learning-based intrusion detection with network-layer routing control. We demonstrate a closed-loop system where the *Network Intelligence* (IDS) and the *Network Control* (routing) operate in synergy, enabling the network to self-heal and self-protect autonomously by leveraging detection insights to trigger fault-tolerant routing decisions in real time.

Dissertation Outline

The thesis is structured into five chapters, organized across two primary parts: **Background** and **Contributions**. The Background part (Chapters 1–3) provides the reader with the necessary theoretical and contextual foundations, while the Contributions part (Chapters 4–5) presents our original research work.

The organizational structure of the dissertation unfolds as follows:

- **Chapter 1: Internet of Things and Context.** This chapter establishes the architectural landscape of IoT and IoMT. It defines the specific requirements of WBANs, the constraints of medical sensors, and the operational challenges encountered in healthcare-oriented IoT deployments.
- **Chapter 2: Theoretical Background.** This chapter consolidates the mathematical and theoretical foundations underpinning the research. It covers the principles of Time Allocation via Temporal Aggregation, the architecture of Deep Recurrent Neural Networks (LSTM), the paradigms of Distributed Learning (Federated and Transfer Learning), and the logic of Multi-Criteria Decision Analysis (MCDA) methods including TOPSIS and AHP.
- **Chapter 3: State of the Art.** This chapter provides a critical taxonomy and review of existing literature. It analyzes current Intrusion Detection Systems and routing protocols for IoMT, identifying specific gaps regarding Non-IID data handling and the absence of proactive fault tolerance mechanisms in current standards.
- **Chapter 4: The FTL-HLSTM Framework.** Focusing on security and scalability, this chapter details the design and implementation of our Federated Transfer Learning-based intrusion detection system. It presents the Intelligent Label

Classification algorithm, the temporal aggregation strategy, and evaluates the model's performance against benchmark datasets.

- **Chapter 5: Distributed Fault Tolerance and Routing.** Focusing on reliability, this chapter presents the MCDA-based routing mechanism. It details how the system leverages the outputs of the IDS to make intelligent routing decisions, and provides a comparative analysis demonstrating the framework's resilience against node failures and insider attacks.

The thesis concludes with a **General Conclusion** that summarizes the key findings, discusses the limitations of the proposed approaches, and outlines future research directions for the next generation of secure and scalable IoMT networks.

Part I

Background & Foundations

Chapters 1–3 establish the architectural, theoretical, and literature foundations upon which the contributions are built.

Chapter

1

Internet of Things

Chapter 1

Internet of Things

1.1 Introduction

The evolution of the Internet of Things (IoT) is anticipated to fundamentally transform the future internet landscape, introducing unprecedented opportunities for automation and the seamless integration of physical objects into digital ecosystems. This transformative trajectory pervades numerous domains; however, the medical field has distinguished itself as a particularly early and significant adopter of such technological innovation. At present, the healthcare sector is experiencing a rapid proliferation of IoT-driven applications, encompassing electronic health (e-Health) and mobile health (m-Health) paradigms, which are collectively designated as the Internet of Medical Things (IoMT). These advancements represent a paradigm shift in healthcare delivery, harnessing the capabilities of interconnected devices and sophisticated data analytics to enhance patient care, improve diagnostic precision, and optimize therapeutic outcomes. As IoMT continues to mature, it harbors immense potential to reshape established medical practices, empower patients through greater autonomy, and catalyze breakthroughs in healthcare research and innovation.

This chapter presents a comprehensive examination of the Internet of Medical Things (IoMT) and its principal enabling technology, namely Wireless Body Area Networks (WBANs). It provides an in-depth overview encompassing the definition, applications, architectures, as well as the associated benefits and challenges inherent to this domain. Through a systematic exploration of these multifaceted aspects, this chapter aims to establish a robust understanding of IoMT and its significance across various sectors, with particular emphasis on its transformative role within the healthcare industry.

1.2 Internet of Things

The Internet of Things (IoT) constitutes a contemporary computing paradigm that endeavors to transform conventional physical objects into intelligent, interconnected enti-

ties [94]. Widely acknowledged as one of the most disruptive technologies of the current era, IoT possesses the capacity to fundamentally alter the manner in which individuals perceive and interact with their surrounding environment [203]. This transformation is underpinned by substantial advancements in ubiquitous computing, embedded systems, communication technologies, sensor networks, Internet protocols, and web-based applications [70, 184]. These foundational technologies collectively endow everyday devices with computational intelligence, thereby actualizing the overarching vision of IoT [109].

The conceptual genesis of the Internet of Things can be traced to 1999, when researchers affiliated with the Auto-ID Center at the Massachusetts Institute of Technology originally articulated the foundational idea [204]. This seminal concept sought to imbue ordinary objects with computational intelligence and establish their connectivity to the Internet, thereby enabling pervasive machine-to-machine communication among real-world entities. The formal institutional recognition of IoT materialized in 2005 at the World Summit on the Information Society held in Tunisia, where the International Telecommunication Union (ITU) published two landmark reports delineating key enabling technologies, market opportunities, and emerging challenges. These reports characterized IoT as a paradigm destined to engender a dynamic and interconnected network of networks [156].

Since the inception of the ARPANET in the 1960s—the architectural precursor to the modern Internet—the number of network-connected devices has exhibited a trajectory of sustained growth, accelerating markedly following the commercialization and liberalization of the Internet in the late 1980s [182]. Contributing factors such as ubiquitous network connectivity and the introduction of IPv6, with its vastly expanded address space, have been instrumental in facilitating the evolution and proliferation of IoT [157]. Contemporary projections indicated that the global count of connected devices would surpass 25 billion by 2020, a substantial increase from approximately 10 billion in 2014, and would further exceed 100 billion by 2050 [87]. Figure 1.1 illustrates the historical and projected growth of Internet-connected devices spanning from the 1950s to 2050, as forecasted by IBM in 2015.

1.2.1 Architecture

Research pertaining to IoT architectures has been extensive, with investigators examining the associated challenges and design considerations from a multiplicity of perspectives. Consequently, a diverse array of architectural frameworks has been proposed to address domain-specific requirements. It is noteworthy, however, that no universally accepted architecture has emerged that satisfies the needs of every researcher or

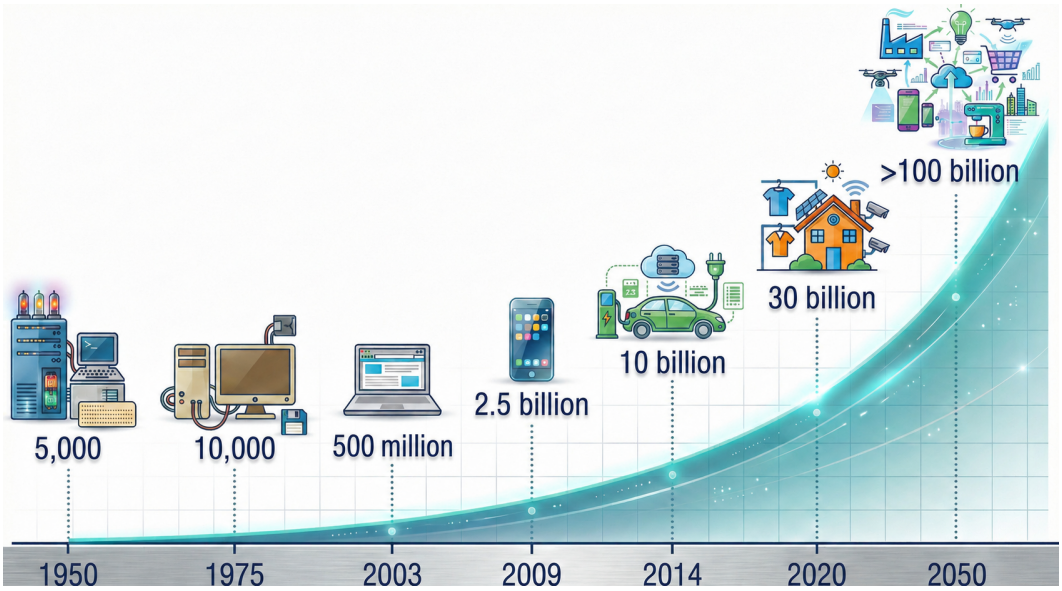

Figure 1.1: Growth of connected devices from 1950 to 2050 [157].

end-user, nor one that is universally applicable across all deployment scenarios [181]. During the nascent stages of IoT development, Wu et al. [205] proposed a foundational three-layer architecture comprising, from top to bottom, the Application Layer, the Network Layer, and the Perception Layer. This initial model was relatively rudimentary, as it aggregated numerous heterogeneous functions—now recognized as belonging to distinct operational domains—within a single architectural layer.

Subsequently, Khan et al. [119] expanded upon this framework by introducing a more granular five-layer architectural model. This refined architecture decomposes the monolithic Application Layer of the three-layer model into three specialized layers: the Middleware Layer, the Application Layer, and the Business Layer, while preserving the original Perception and Network Layers in their established form. Figure 1.2 presents a comparative illustration of both the three-layered and five-layered architectural models as reported in the literature.

1.2.2 Communication Technologies

A substantial number of communication technologies are presently available for deployment within IoT ecosystems, each characterized by its own distinctive set of advantages and inherent limitations. This heterogeneity in technical capabilities and trade-offs implies that no single communication technology is universally optimal for every deployment scenario, operational requirement, or application context. Accordingly, it is imperative to conduct a rigorous evaluation of each candidate technology against the specific demands of the target application to ascertain the most appropriate selection.

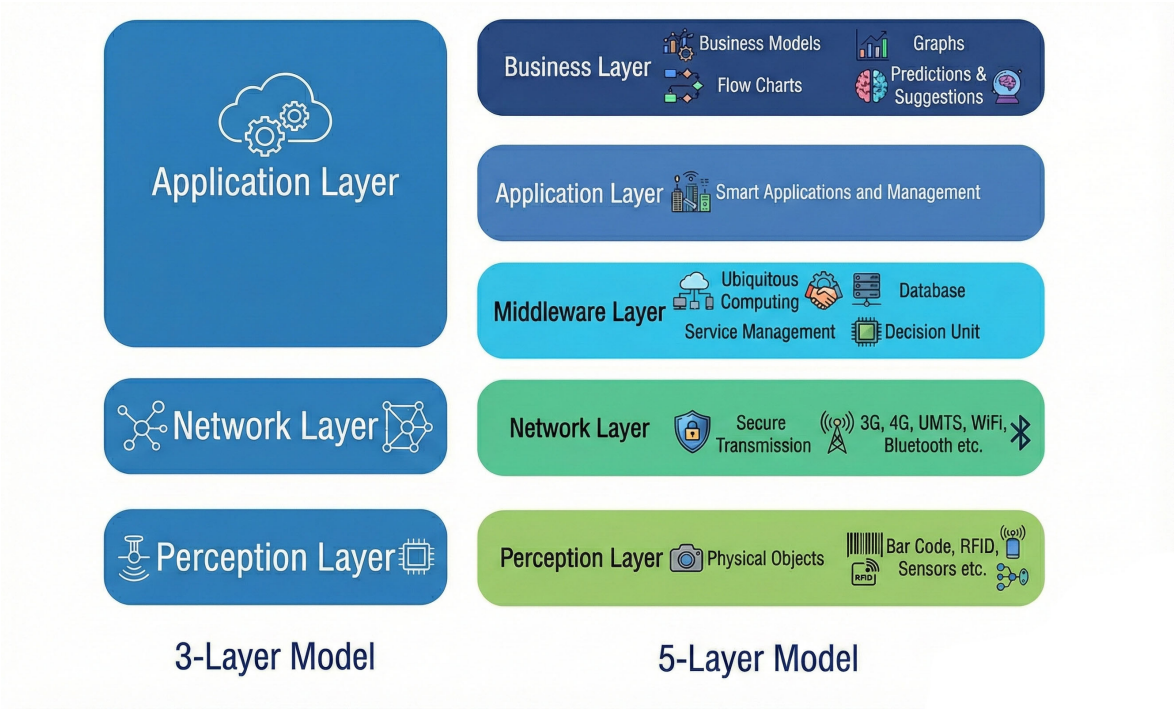

Figure 1.2: IoT layered architectures proposed in the literature [119].

The following subsections present a synthesized comparative analysis of the principal communication technologies documented in the extant literature.

Bluetooth

Bluetooth constitutes a widely adopted wireless communication standard recognized for its versatility in handling substantial volumes of data. It supports ad-hoc network formation and benefits from universal standardization, thereby ensuring broad device interoperability. Nevertheless, Bluetooth is characterized by relatively constrained data rates, elevated power consumption, and a documented susceptibility to external security threats [75, 197].

ZigBee

ZigBee is distinguished by its straightforward deployment characteristics and decentralized control architecture, which enables effective load distribution across multiple network nodes without necessitating a centralized coordinating authority. It offers advantageous properties including low power consumption, cost-effectiveness, and minimal latency. However, ZigBee exhibits notable deficiencies in robust security provisioning and may encounter interoperability challenges when interfacing with devices from heterogeneous manufacturers [52, 75].

WiMax

WiMax demonstrates particular proficiency in supporting high-throughput voice and data transmissions over extended geographical distances, with a single base station capable of simultaneously serving a considerable number of users. Notwithstanding these advantages, WiMax typically necessitates line-of-sight propagation paths and may suffer from bandwidth degradation under conditions of heavy user load [64, 75, 207].

Bluetooth Low Energy (BLE)

Bluetooth Low Energy (BLE) has been specifically engineered to prioritize energy efficiency, rendering it particularly well-suited for IoT applications operating under stringent power constraints. BLE achieves an effective balance between low power consumption and moderate data throughput; however, it remains limited by its constrained data handling capacity and susceptibility to certain classes of security attacks [145, 152, 177].

Wi-Fi

Wi-Fi is characterized by its high achievable data rates and pervasive global adoption, making it a viable candidate for a broad spectrum of IoT applications. However, Wi-Fi encounters several operational challenges, including escalating power consumption as user density increases, vulnerability to various attack vectors, and performance degradation in outdoor and physically obstructed propagation environments [137, 138, 183, 197].

LoRa and LoRaWAN

LoRa and LoRaWAN technologies provide extensive coverage areas and support for large-scale device populations, rendering them particularly suitable for IoT deployments spanning geographically vast regions. However, their predominantly point-to-point communication architecture and dependence on gateway infrastructure can introduce performance bottlenecks, and they may experience elevated packet loss rates under congested network conditions [64, 125, 207].

Wi-Fi HaLow

Wi-Fi HaLow, standardized as IEEE 802.11ah, extends the operational reach of conventional Wi-Fi through enhanced signal propagation characteristics and reduced power consumption, rendering it suitable for IoT deployments in challenging radio frequency environments. Nevertheless, the variability in achievable data rates and the absence of universally adopted frequency allocation standards represent potential limitations [65, 152].

MiWi and MiWi P2P

MiWi and MiWi Peer-to-Peer (P2P) protocols offer low power consumption and medium-range communication capabilities at zero licensing cost, positioning them as cost-effective solutions for certain IoT deployments. However, their proprietary nature and susceptibility to electromagnetic interference may constrain interoperability and operational reliability [3].

ISA100.11a

ISA100.11a prioritizes communication reliability and security provisioning, making it particularly well-suited for industrial IoT applications. Despite these merits, ISA100.11a confronts challenges including implementation complexity and limited interoperability with alternative communication technologies [52, 129].

WirelessHART

WirelessHART employs a self-organizing mesh network architecture complemented by robust security mechanisms, conferring resilience against interference and suitability for industrial IoT environments. However, its reliance on Time Division Multiple Access (TDMA) scheduling can introduce latency, and it may encounter difficulties with simultaneous multi-node communication in multi-drop configurations [52].

Z-Wave

Z-Wave is notable for its capacity to support a substantial number of simultaneous device connections, making it well-suited for dense IoT ecosystems. It incorporates mesh networking topology, ensuring robust communication even within topologically complex environments. With the elimination of single points of failure, Z-Wave provides enhanced reliability and fault resilience. Additionally, it is comparatively inexpensive

and exhibits low power consumption, making it appropriate for both residential and light commercial deployments. However, Z-Wave is subject to limitations including low data rates, line-of-sight operational dependencies, and documented instability in certain current implementations [197].

LTE, LTE-M, and LTE-A

Operating within licensed spectrum bands, LTE variants deliver increased throughput and reduced co-channel interference, ensuring reliable connectivity in spectrally congested environments. Their narrowband operational modes contribute to reduced power consumption, thereby extending the operational longevity of IoT devices. The cost-effectiveness of both terminal devices and base station infrastructure, coupled with the strategic reuse of existing cellular spectrum allocations, facilitates the densification of IoT deployments. Nevertheless, the associated cellular data plan charges may constitute a limiting factor for widespread adoption [64, 152, 207].

Ultra-Wideband (UWB)

Ultra-Wideband (UWB) technology operates within license-free spectral bands, enabling high data rates and efficient spectrum sharing with incumbent systems. Characterized by exceptionally low power consumption, UWB is well-suited for battery-constrained IoT devices. Its inherent immunity to multipath propagation effects and minimal interference generation further enhance operational reliability. However, the limited communication range and the engineering complexity of UWB antenna design pose challenges to practical deployment [112, 126, 152].

Wavenis

Wavenis supports ultra-low power operational modes and provides extensive communication ranges, rendering it appropriate for IoT deployments spanning large geographical areas. With reliable data transmission and support for heterogeneous network topologies, Wavenis ensures robust connectivity. However, its low achievable data rates, dependence on high link budgets for long-range operation, and requirement for line-of-sight communication paths may constrain its applicability in certain scenarios [112, 138, 197].

Insteon

Insteon integrates powerline and wireless communication modalities, facilitating extended operational ranges and multi-hop data transmission. Its decentralized architectural design enhances system reliability by eliminating single points of failure. However, Insteon is subject to limitations including low achievable data rates, elevated power consumption, and dependence on proprietary technology [112, 197].

Thread

Thread accommodates a large number of client devices and supports mesh networking topology, providing scalable and reliable communication for IoT ecosystems. With native IPv6 compatibility and robust security mechanisms, Thread ensures secure and efficient data transmission. However, its limited communication range, protocol complexity, and relatively demanding implementation requirements may present barriers to adoption [64, 152, 197].

EnOcean

EnOcean enables energy harvesting from ambient environmental sources, thereby supporting battery-free operation for IoT devices. Offering high data rates and compatibility with both indoor and outdoor deployment environments, EnOcean provides considerable versatility. However, its limited communication range, reliance on proprietary technology, and throughput constraints may affect its suitability for particular application domains [112, 177].

Li-Fi

Light Fidelity (Li-Fi) technology delivers high-speed optical wireless communication and elevated data transfer rates, providing efficient data transmission for IoT applications. Its low power consumption characteristics contribute to enhanced energy efficiency in connected environments, while its optical communication modality inherently provides superior security properties relative to radio frequency-based technologies. Furthermore, Li-Fi presents reduced health risks attributable to its reliance on visible light communication.

However, Li-Fi is confronted with several challenges, including comparatively high initial deployment costs, susceptibility to interference from ambient light sources, limited operational range, and signal attenuation by physical obstacles such as walls. These constraints render Li-Fi predominantly suitable for indoor application scenarios.

Moreover, the requisite infrastructure investment for Li-Fi deployment may constitute a significant impediment to its widespread adoption [9, 45, 134, 183].

In summary, each communication technology presents a distinctive combination of advantages and limitations. While certain technologies demonstrate superior performance in indoor environments, others exhibit enhanced efficacy in outdoor settings. Similarly, some technologies are optimized for scenarios involving a limited number of users, whereas others maintain satisfactory performance under high-density conditions. Consequently, identifying a single communication technology that universally satisfies the heterogeneous requirements of all IoT application domains remains impractical.

1.3 Emergence of IoT into the Medical Field

The penetration of Internet of Things (IoT) technologies into the medical domain has paved the way for the emergence of the Internet of Medical Things (IoMT). This evolution represents a transformative paradigm shift in healthcare delivery, enabled by the convergence of interconnected medical devices, sensor technologies, and advanced data analytics. IoMT encompasses a comprehensive spectrum of applications and technologies specifically tailored for medical utilization, spanning remote patient monitoring, telemedicine, intelligent healthcare facilities, and personalized medicine. Through the systematic integration of IoT technologies into healthcare delivery and management processes, IoMT substantially enhances the quality of patient care, improves clinical outcomes, and optimizes the operational efficiency of healthcare systems. Furthermore, IoMT facilitates real-time physiological monitoring, early detection of pathological conditions, and proactive clinical interventions, thereby fundamentally revolutionizing healthcare delivery and fostering a more interconnected, data-driven, and patient-centric healthcare ecosystem.

The integration of Wireless Sensor Networks (WSNs) into the broader IoT framework has played a pivotal role in facilitating the transition toward IoMT. WSNs provide a scalable and cost-effective infrastructure for the acquisition of data from diverse medical sensors and the subsequent transmission of this data to IoT platforms for computational analysis and clinical decision support. This integration has enabled seamless inter-device communication and data exchange among medical devices, supporting real-time monitoring and analysis of patient health data, and ultimately catalyzing the emergence and maturation of IoMT.

1.4 Internet of Medical Things (IoMT)

The Internet of Medical Things (IoMT) refers to an interconnected ecosystem of electronic devices equipped with specialized sensors designed for physiological sensing and continuous health monitoring purposes [11, 42]. The sensors employed within IoMT constitute a specialized subset of Wireless Sensor Networks (WSNs), specifically designated as Wireless Body Area Networks (WBANs). These sensors, commonly referred to as nodes, communicate and transmit acquired data via wireless communication links, typically through intermediary gateway devices. The collected information is subsequently relayed to centralized computational infrastructure, such as remote servers or cloud-based database systems, for storage, analysis, and further clinical utilization.

This section provides a systematic overview of IoMT, the sensor modalities utilized within this framework, and the underlying architectural design. Figure 1.3 depicts the layered architecture of IoMT based on the aforementioned architectural paradigm.

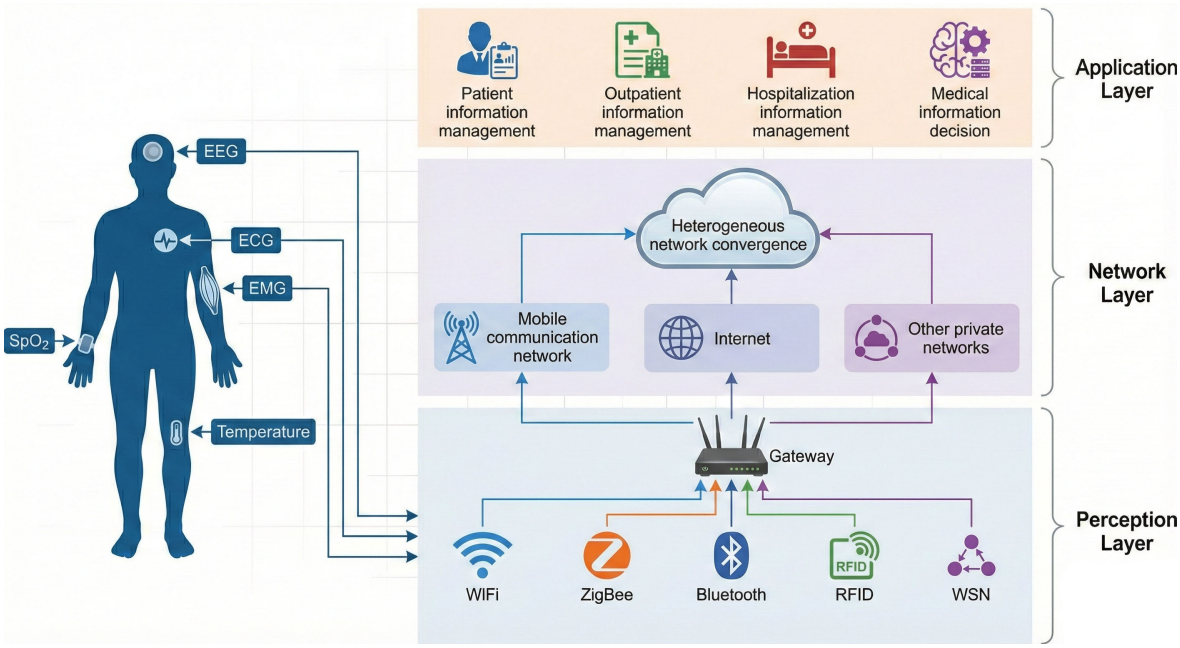

Figure 1.3: Layered IoMT architecture [189].

1.4.1 IoMT Architecture

Numerous researchers have investigated the architectural design of the Internet of Things, proposing various frameworks applicable to IoT-based systems. Among the most prominent architectural paradigms are the EPCglobal architecture, the Web of Things architecture, the sensor network-based architecture, the autonomous architecture, and the Machine-to-Machine (M2M) architecture. Of these, the M2M architecture

has emerged as the most widely adopted framework, incorporating essential architectural elements from both EPCglobal and WSN paradigms. Within the medical domain, the Internet of Medical Things (IoMT) represents a specialized and domain-specific instantiation of general-purpose IoT technology. In practice, IoMT applications predominantly adhere to the canonical three-tier architectural model of IoT, comprising the Application Layer, the Network Layer, and the Perception Layer [25, 189].

Perception Layer

The perception layer constitutes the most critical and architecturally complex stratum of the IoMT framework. It comprises two principal sublayers: the data access sublayer and the data acquisition sublayer. The data acquisition sublayer employs a diverse array of medical perception instruments and signal acquisition equipment to identify nodes within the IoMT network and to gather data pertaining to monitored entities. It leverages signal acquisition modalities including General Packet Radio Service (GPRS) technology [165], Radio Frequency Identification (RFID), image recognition, graphic coding, and heterogeneous sensor arrays encompassing physical signal sensors. Physiological signal sensors, chemical sensors, and DNA sensors collectively transform monitored entities within the network into readily identifiable Cyber-Physical Systems (CPS) nodes [132]. Within the IoMT context, nodes are taxonomically classified as passive CPS, active CPS, and Internet CPS, contingent upon the specific objects and operational requirements [132]. The data access sublayer establishes connectivity between data collected by the acquisition sublayer and the network layer through short-range wireless data transmission technologies, including Bluetooth, Wi-Fi, and ZigBee. The selection of appropriate access methods is governed by the prevailing IoMT deployment environment and the specific requirements of the various monitored objects [142].

Network Layer

The network layer is architecturally decomposed into two constituent sublayers: the service sublayer and the network transmission sublayer. The network transmission sublayer functions as the backbone communication infrastructure of the IoMT ecosystem, serving as the central nervous system for data routing and delivery. It leverages the Internet, mobile communication networks, and other specialized network infrastructure to transmit data acquired by the perception layer in a real-time, reliable, accurate, and uninterrupted manner [120]. Rather than supplanting existing network infrastructure, the objective of the IoMT network layer is to investigate and implement heterogeneous network integration technologies specifically tailored for healthcare environments [141].

The service sublayer integrates heterogeneous networks and diverse data types, encompassing data descriptions, data warehouses, and ancillary data repositories. It further develops a comprehensive support service system with standardized open interfaces for the various services operating at the application layer, thereby enabling third-party developers to construct domain-appropriate applications for use by medical professionals and other authorized personnel [144].

Application Layer

The application layer encompasses both health data decision-support applications and medical information management applications. Medical data management applications include, but are not limited to, material and equipment management, patient data management, inpatient treatment data management, and outpatient data management [178]. Medical data decision-support applications comprise patient data analysis, epidemiological data analysis, pharmaceutical data analysis, diagnostic support, and therapeutic data analysis [5].

1.4.2 IoMT and WBANs

IoMT and WBANs are fundamentally intertwined, fulfilling complementary and synergistic roles within modern healthcare systems. IoMT encompasses an extensive array of interconnected medical devices and sensors that collectively facilitate the acquisition, transmission, and analysis of health-related data. Within this architectural framework, WBANs serve as a critical enabling component by supporting seamless wireless communication between body-worn sensors and other IoMT devices. WBANs enable the continuous monitoring of individuals' vital signs and physiological parameters, providing real-time clinical insights into their health status. This synergistic relationship between IoMT and WBANs empowers healthcare providers to deliver personalized and proactive care, improve patient outcomes, and enhance overall healthcare system efficiency. Additionally, WBANs contribute substantively to the advancement of remote patient monitoring, telemedicine, and other innovative healthcare applications, thereby underscoring their indispensable role within the IoMT ecosystem.

1.4.3 Architecture of WBAN-based IoMT

A diversity of WBAN-based IoMT architectures has been documented in the literature, varying according to the adopted design methodology and the target application domain. Nevertheless, a fundamental three-level hierarchical architecture—analogous

to the layer-based architecture established for general IoMT—is consistently observed across existing works. This canonical architecture serves as the foundational blueprint for all WBAN-based IoMT applications. Figure 1.4 illustrates the representative architecture of WBAN-based IoMT networks, which is composed of three primary hierarchical levels.

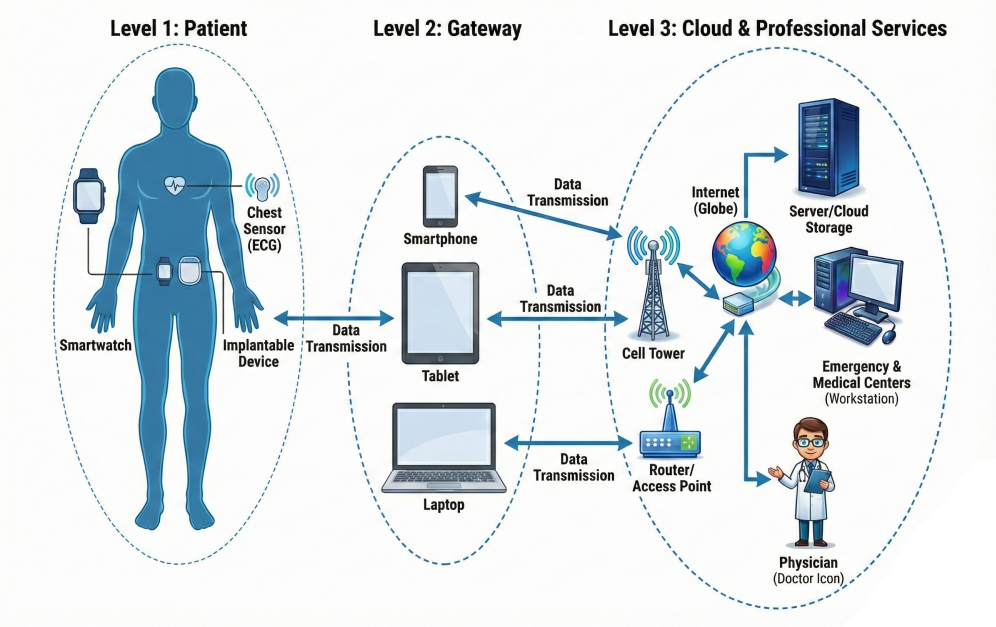

Figure 1.4: Architecture of WBAN-based IoMT.

Level 1 (Sensor Tier): This level comprises specialized transducers designated as medical sensors. These sensing nodes are engineered to continuously measure, monitor, and acquire specific biological signals from the human body. The gathered physiological data are subsequently transmitted to the gateway devices constituting Level 2.

Level 2 (Gateway Tier): Devices operating at this level primarily function as communication gateways, including personal digital assistants (PDAs), portable computers, and smartphones. They serve as intermediary nodes between Level 1 sensor devices and Level 3 back-end infrastructure, bearing responsibility for relaying the collected data from Level 1 sensing nodes to end-users at Level 3 via open communication channels.

Level 3 (Back-end Tier): At this hierarchical level, data and information received from Level 2 gateway devices are transmitted to end-users via the Internet. The nature of these end-users is contingent upon the specific IoMT network design and may encompass entities such as cloud computing platforms, emergency physicians, healthcare professionals, service providers, data analysts, family members, or the patients themselves.

1.4.4 Medical Sensors

Wireless medical sensors, as integral components of the IoMT infrastructure, fulfill a specialized function in the quantitative measurement of physiological metrics, including body temperature, arterial blood pressure, cardiac rate, electrocardiogram (ECG) waveforms, and respiratory parameters [63, 194]. These sensors transmit acquired biological data to a control device that is either worn on the body or positioned at an accessible proximal location. Medical sensors can be taxonomically classified into implantable nodes, garment-attached sensors, and body surface nodes (wearable devices), each serving distinct clinical purposes [40, 161]. Predominantly deployed as either implantable or wearable devices, their intimate association with the human body renders them indispensable across a wide range of medical and healthcare applications. Furthermore, medical IoT sensors encompass a diverse array of modalities tailored to specific physiological functions, including ECG, electroencephalogram (EEG), blood pressure, and body temperature sensors [108]. For example, ECG sensors are utilized to monitor cardiac rhythm and detect arrhythmic patterns, whereas EEG sensors are employed to assess and identify abnormalities in cerebral electrical activity. Comprehensive taxonomic information regarding medical sensor types is available in [108].

These sensors can be further categorized into three device classes based on their computational and storage capabilities:

Class 0: Devices that are severely constrained in memory and processing capabilities, with RAM capacity of less than 5 KB and flash memory of less than 10 KB. Class 0 devices cannot be secured through conventional cryptographic mechanisms.

Class 1: Devices that are moderately constrained in code space and processing capabilities, with RAM capacity of approximately 10 KB and flash memory capacity of approximately 100 KB. Class 1 devices provide adequate support for fundamental security functions and possess sufficient computational resources to operate a protocol stack.

Class 2: Devices that are comparatively less constrained, with RAM capacity of approximately 50 KB and flash memory capacity of up to 250 KB. Class 2 devices can effectively leverage lightweight, energy-efficient communication protocols and support the majority of standard protocol stacks.

Additionally, these sensors constitute a specialized category within Wireless Sensor Networks (WSNs) known as Wireless Body Area Networks (WBANs), which are extensively deployed in the IoMT domain. WBANs are specifically engineered to monitor physiological parameters and acquire data from the human body. They play a pivotal role in IoMT applications, facilitating continuous health monitoring, medical

diagnostics, and personalized healthcare delivery.

1.5 Applications of IoMT

A substantial number of IoMT applications depend on WBANs to achieve the requisite levels of operational efficiency and service quality. Given the considerable capabilities of WBANs, a diverse array of novel IoMT applications spanning domains such as clinical medicine, home-based healthcare, and continuous patient monitoring are being progressively adopted. The following subsections delineate and examine several of the most prevalent healthcare applications currently in deployment.

1.5.1 Health Records

This application domain encompasses various categories of electronic health records, which can be systematically classified into three principal forms:

- **Electronic Health Record (EHR):** A comprehensive digital representation of a patient's complete health history, providing a detailed longitudinal account of the patient's health status, with secure access restricted to authorized clinical users [43].
- **Electronic Medical Record (EMR):** A digital compilation of an individual patient's complete medical history within a specific clinical institution or healthcare facility [89].
- **Personal Health Record (PHR):** A patient-managed digital repository in which the individual securely maintains their own health-related data with appropriate confidentiality and privacy protections [124].

1.5.2 Remote Health Monitoring

Remote health monitoring constitutes an automated medical service that continuously tracks patients' vital signs through the deployment of WBANs. Various sensor modalities can be strategically positioned on or within the patient's body to monitor critical physiological indicators, including cardiac rate, arterial blood pressure, and body temperature. The acquired data is subsequently stored in a centralized control unit or transmitted to remote clinical facilities for analytical processing and further clinical evaluation [29].

1.5.3 Assisted Living

The integration of WBANs within the IoMT framework has introduced a paradigm in which patients can remain in their domiciliary environment while utilizing wearable medical sensors for continuous physiological monitoring. These sensors persistently track the patient's physiological parameters and can either store and transmit data at predefined intervals or, in specific clinical scenarios, autonomously administer prescribed medications—for example, insulin delivery upon detection of hyperglycemic episodes by blood glucose sensors. Moreover, the system is capable of triggering automated alerts to the nearest healthcare facility when clinically significant deviations are detected [29].

1.5.4 Telecare Medicine

An additional domain of IoMT that leverages WBAN technology is telecare medicine. This modality enables the remote delivery of healthcare services through the utilization of information and communication technologies, including WBANs [23]. By employing video conferencing and sensor-based data acquisition technologies, healthcare providers can remotely assess patients' clinical conditions and formulate medication prescriptions based on tele-sensed physiological data, thereby eliminating the necessity for physical co-location.

1.6 Benefits of IoMT

The integration of IoMT technology into the healthcare domain has precipitated a fundamental paradigm shift in healthcare delivery and management. This technological advancement has not only transformed Internet-mediated communication but has also exerted a significant impact across numerous sectors, with healthcare being a primary beneficiary. By establishing seamless connectivity among clinicians, patients, and healthcare services, IoMT offers unprecedented levels of convenience, diagnostic accuracy, and operational flexibility [102].

A principal benefit of IoMT lies in its capacity to enable healthcare professionals to discharge their clinical responsibilities with enhanced precision and efficiency. This integration has conferred substantial advantages upon patients, as IoMT devices are designed to be user-friendly and facilitate seamless access to healthcare services. The major benefits provided by IoMT are enumerated as follows:

- The integration of IoT-enabled devices into healthcare infrastructure enhances

operational convenience while simultaneously achieving significant reductions in healthcare expenditures. Through real-time disease management capabilities, patient outcomes are markedly improved, culminating in an overall enhancement of quality of life. Moreover, IoMT augments the user experience and elevates the standard of patient care while concurrently reducing costs through efficient resource utilization.

- The most consequential benefit of IoMT is the promotion of healthier and longer lives, achieved through comprehensive disease management and preventive care strategies. IoMT enables continuous monitoring of vulnerable populations, including pediatric patients and elderly individuals, thereby ensuring their sustained well-being.

- A critical advancement afforded by IoMT is its capacity for automated notification of relevant stakeholders upon detection of clinically significant alterations in a patient's health status, potentially preserving lives and reducing critical response times. Furthermore, the capabilities of IoMT extend beyond healthcare, facilitating connectivity and interoperability among heterogeneous IoT devices for enhanced operational efficiency and effectiveness [102].

- **Medication adherence and family notification:** IoMT enables the assurance of timely medication administration, guaranteeing that patients receive their prescribed pharmacological interventions at appropriate intervals. Additionally, IoMT systems can automatically notify family members regarding the status of patient care, thereby fostering improved communication and support networks [92].

- **Simplicity, affordability, and usability:** IoMT solutions are designed with an emphasis on simplicity, offering intuitive user interfaces and streamlined operational processes for enhanced accessibility. These technologies are engineered for affordability, ensuring that cost barriers are minimized and healthcare remains universally accessible [158].

- Clinicians can efficiently manage patient records through IoMT systems, facilitating systematic organization and rapid accessibility of critical medical information [122].

- IoMT systems promote energy efficiency by optimizing the utilization of various resources, including temporal and financial investments. Through process automation and streamlined workflows, IoMT solutions reduce overall energy consumption, yielding time and cost savings for both healthcare providers and patients [186].

- IoMT enables healthcare providers to deliver continuous medical services beyond conventional operating hours, ensuring round-the-clock access to healthcare through remote monitoring, telemedicine, and automated alerting mechanisms for timely clinical interventions [48].

In summary, IoMT has emerged as a transformative force within the medical field, offering manifold benefits to individuals, society, the environment, consumers, and healthcare organizations. From personalized healthcare solutions to improved operational efficiency, IoMT holds considerable promise for revolutionizing healthcare delivery and enhancing overall population wellness.

1.7 Challenges in IoMT

The integration of WBANs within the IoMT framework into healthcare systems has yielded numerous operational benefits; however, it simultaneously introduces a constellation of significant technical and practical challenges. This section systematically examines the principal challenges confronting IoMT.

1.7.1 Limited Resources

Sensor nodes are typically characterized by diminutive physical dimensions and severely constrained computational resources, encompassing processing capacity, storage capabilities, communication bandwidth, and battery energy reserves. The limited energy budget of these sensor nodes within the network poses substantial longevity challenges for WBANs deployed in IoMT applications. Addressing the issue of resource scarcity necessitates the adoption of efficient utilization strategies [76]. Consequently, the implementation of energy-efficient communication protocols becomes imperative to extend the operational lifespan of the network. Representative examples include energy-aware routing algorithms at the network layer and duty-cycling energy-saving modes at the Medium Access Control (MAC) layer. Additionally, optimizing the utilization of constrained memory resources in sensor nodes is of critical importance, particularly given the memory-intensive nature of tasks such as routing table maintenance, data replication, and cryptographic security operations.

1.7.2 Scalability

The scalability of WBANs within the IoMT framework constitutes a critical determinant of system effectiveness and the feasibility of widespread adoption. WBANs must

be capable of accommodating an increasing population of interconnected devices and sensors without incurring degradation in performance or operational efficiency. Scalability in WBANs enables the network to manage growing volumes of data generated by an expanding array of medical sensors and devices [117]. This scalability imperative extends beyond mere device count to encompass factors such as network coverage area, data transmission throughput, and interoperability with existing healthcare infrastructure [50]. A scalable WBAN infrastructure ensures that IoMT systems can dynamically adapt and evolve to accommodate the changing requirements of healthcare applications, supporting advancements in remote patient monitoring, telemedicine, and personalized healthcare delivery.

1.7.3 Cost of Sensor Platforms

The elevated cost of commercially available sensor platforms constitutes a significant economic barrier to widespread adoption and deployment. Additionally, the challenge of developing more affordable and disposable sensor platforms further compounds this economic constraint [149].

1.7.4 Environmental Factors

The environmental conditions within which WBANs operate exert a substantial influence on their performance and operational reliability. WBANs are designed to function across diverse environmental settings, ranging from controlled indoor environments to variable outdoor conditions. Factors including temperature fluctuations, humidity levels, electromagnetic interference, and physical obstructions can adversely impact the performance of WBANs. Furthermore, WBANs deployed in healthcare settings must comply with stringent regulatory standards and safety requirements to ensure patient safety and data integrity. Accordingly, robust design methodologies and resilient communication protocols are essential to mitigate the deleterious effects of environmental challenges on WBAN operation, thereby ensuring continuous and reliable functionality across diverse healthcare deployment scenarios.

1.7.5 Inconsistent Wireless Communication

Communication within Wireless Body Area Networks is frequently characterized by inherent unreliability, attributable to the error-prone nature of the wireless propagation medium, which exhibits elevated bit error rates and variable link capacity. Consequently, for a WBAN to operate effectively, it must demonstrate sufficient reliability

tailored to the specific quality-of-service requirements of its intended clinical applications [92]. It is imperative that acquired medical data be transmitted reliably to clinical specialists, ensuring accurate and timely delivery for informed decision-making in healthcare settings.

1.7.6 Susceptibility to Node Failures

In IoMT networks, sensor nodes are frequently susceptible to unforeseen operational failures arising from various causes, including energy depletion and physical damage. Furthermore, communication links between nodes may be permanently disrupted due to hardware malfunction or environmental factors. As a consequence, WBANs deployed in IoMT applications must demonstrate resilience in the presence of node failures. To enhance fault tolerance, IoMT WBANs may strategically deploy a surplus of sensor nodes beyond the minimum operational requirement, thereby introducing redundancy into the network topology.

1.7.7 Security

Given the deployment of sensor devices in typically unprotected and physically accessible environments, numerous IoMT applications mandate stringent security provisioning to satisfy fundamental security requirements and safeguard against diverse cyber-attack vectors. This imperative is critical for preventing adversaries from compromising network operations through unauthorized acquisition of control over sensor nodes [22, 97, 186, 196]. Additionally, communication security constitutes a significant challenge for WBANs in IoMT. Wireless communication channels inherently lack confidentiality guarantees, rendering transmissions vulnerable to eavesdropping and message tampering by malicious entities, with potentially severe clinical consequences. Furthermore, the resource-constrained nature of sensor devices renders the implementation of conventional cryptographic security schemes impractical in WBANs, as such schemes typically incur prohibitive computational, communication, and memory overheads. Consequently, ensuring comprehensive security in IoMT WBANs remains an open and challenging research problem.

1.8 Cloud Integration in IoMT

IoMT-based healthcare systems are projected to significantly enhance the quality of life, reduce aggregate healthcare costs, and expand patients' medical knowledge base. From the healthcare provider's perspective, IoMT possesses the potential to minimize device

downtime through proactive remote monitoring and predictive maintenance. Additionally, IoMT can accurately forecast optimal resupply schedules for various medical devices, ensuring their continuous and efficient operation. It further facilitates the effective scheduling of limited healthcare resources, optimizing their utilization and improving service delivery to patients.

Cloud computing provides a comprehensive suite of services—encompassing databases, computational servers, software applications, data analytics, and networking infrastructure—delivered over the Internet. This computational paradigm enables flexible resource allocation, accelerated application deployment, and cost efficiencies through economies of scale. IoMT devices can seamlessly integrate with cloud services, leveraging them for the storage and processing of voluminous medical datasets. This integration supports enhanced data management and analytical capabilities, contributing to more informed clinical decision-making and improved patient outcomes [36, 166, 195].

1.8.1 Cloud Computing in Healthcare Systems

Cloud computing represents a computational paradigm that delivers substantial computational resources and services over a network, typically the Internet. In essence, it involves the utilization of remote servers for data storage, management, and processing. This technology confers numerous operational benefits, including access to virtualized hardware, collaborative software platforms, scalable virtual storage, and on-demand virtual server instances.

Researchers have characterized cloud computing as a novel and rapidly evolving information and communication technology (ICT) service model [153]. Multiple formal definitions of cloud computing have been proposed in the literature, with one comprehensive study identifying 22 distinct characterizations [111]. Cloud computing services are canonically categorized into three principal service models: Platform as a Service (PaaS), Infrastructure as a Service (IaaS), and Software as a Service (SaaS) [49].

Cloud computing infrastructure can be deployed across four distinct models [36]:

- **Public Cloud:** Accessible to the general public and provisioned by cloud service providers such as Amazon Web Services (AWS).
- **Private Cloud:** Operated exclusively by and for a single organization.
- **Hybrid Cloud:** An integrated architecture combining elements of both public and private cloud deployments.
- **Community Cloud:** A shared infrastructure serving organizations with convergent interests and requirements.

In the healthcare domain, cloud computing enables on-demand access to a comprehensive spectrum of information from diverse sources, including insurance claims, electronic medical records (EMRs), laboratory data, and medication histories. It possesses the capability to alert clinicians regarding missed or conflicting prescriptions within pharmacological regimens, particularly for the management of chronic conditions.

The principal advantages of cloud computing in healthcare encompass [71]:

- **Customized Service:** Patient requests can be automatically fulfilled without necessitating human intervention.
- **Network Access:** Patient-facing applications can operate seamlessly across diverse device platforms, including laptops, tablets, and smartphones.
- **Remote Access:** Patients can access cloud-based services without requiring knowledge of the physical data location or underlying infrastructure topology.

These features underscore the transformative potential of cloud computing in enhancing the efficiency and quality of healthcare delivery and patient outcomes.

1.9 Conclusion

This chapter has established the general context of the thesis by presenting the foundational concepts of the Internet of Things (IoT) and its medical extension, the Internet of Medical Things (IoMT). It has examined the main architectural models, communication technologies, and application domains of IoMT, with particular attention to Wireless Body Area Networks (WBANs) as a key enabling component for continuous health monitoring and smart healthcare services.

The chapter has also highlighted the main challenges that arise in IoMT environments, including resource constraints, scalability limitations, unreliable wireless communication, node failures, and security risks. These challenges confirm that designing efficient, secure, and resilient IoMT systems requires more than simple connectivity; it requires a rigorous theoretical foundation capable of addressing communication efficiency, intelligent intrusion detection, and fault-tolerant decision-making in an integrated manner.

Against this background, the next chapter introduces the theoretical foundations that support the proposed framework of this thesis. In particular, it presents the principles of temporal aggregation for scalable communication, the deep learning models used

for sequential intrusion detection, the foundations of privacy-preserving distributed learning, and the multi-criteria decision-making methods used to support fault-tolerant routing.

Chapter

2

Theoretical Background

Chapter 2

Theoretical Foundations of Time Allocation, Scalability, Security, and Fault Tolerance in IoT and IoMT

2.1 Introduction

Designing scalable, fault-tolerant, and secure communication systems for the Internet of Medical Things (IoMT) requires a rigorous foundation spanning multiple theoretical disciplines. IoMT deployments combine resource-constrained sensing devices, heterogeneous clinical environments, stringent privacy regulations, and a critical dependence on data integrity and continuous availability. Addressing these constraints demands an integrated theoretical toolkit that simultaneously accounts for temporal data dynamics, efficient resource utilization, reliable decision-making under uncertainty, and privacy-preserving collaborative learning.

This chapter consolidates the core theoretical foundations underpinning the framework developed in this thesis. The presentation is organized around **six interconnected pillars**, each corresponding to a principal component of the proposed system:

1. **IoMT telemetry as time-series data.** We motivate treating IoMT network traffic as sequential telemetry, establishing the temporal perspective that informs both intrusion detection and communication strategies throughout the thesis (Section 2.2).

2. **Time allocation via temporal aggregation.** We formalize temporal aggregation as a controllable *time allocation* mechanism for scalable communication, defining the aggregation window T_{agg} and its fundamental trade-offs (Section 2.3).

3. **Fault tolerance and trust foundations.** We introduce fault models relevant to IoMT—including non-Byzantine and Byzantine failure modes—and motivate lightweight trust-based decisions over heavy consensus protocols (Section 2.4).

4. **Multi-Criteria Decision Analysis (MCDA).** We present AHP for criteria weighting and TOPSIS for alternative ranking, which together provide the formal decision logic for node selection and model prioritization under conflicting objectives (Section 2.5).

5. **Deep learning for sequential intrusion detection.** We review Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and the Hierarchical LSTM (HLSTM) architecture that constitutes the detection backbone of the proposed framework (Section 2.6).

6. **Privacy-preserving distributed learning.** We present Federated Learning (FL), Transfer Learning (TL), and Federated Transfer Learning (FTL) as scalable paradigms for collaborative model training across healthcare institutions without sharing raw patient data (Section 2.7).

Section 2.8 concludes the chapter with a synthesis of the interconnections among these pillars and transitions to the state-of-the-art review in Chapter 3.

2.2 IoMT Network Telemetry as Time-Series Data

2.2.1 Temporal Dependencies in IoMT Traffic

IoMT devices—from wearable biosensors and implantable monitors to edge gateways and hospital routers [25, 178]—generate continuous streams of telemetry that encode the real-time state of both monitored patients and the underlying communication infrastructure. At the network level, each flow record or packet observation is timestamped and naturally ordered. The traffic observed at a given time instant t is statistically dependent on observations at preceding instants $t-1, t-2, \ldots$, owing to the persistence of connection states, protocol handshake sequences, and session-level behaviors [59, 110, 139]. This inherent temporal dependency renders IoMT network data a natural *multivariate time-series*, in which windowing and sequence formation are not merely convenient preprocessing steps but essential mechanisms for capturing the full signature of both normal and anomalous behavior.

From the standpoint of intrusion detection, the sequential structure has two critical implications. First, attacks that unfold gradually over multiple time steps—such as slow-rate denial-of-service (DoS) campaigns, multi-stage reconnaissance operations, or data exfiltration attempts [20]—cannot be reliably detected by classifiers that treat each observation independently. A classifier operating on isolated feature vectors lacks the contextual information needed to distinguish a slowly escalating attack from normal

traffic fluctuations. Second, temporal models such as LSTM and Hierarchical LSTM are specifically designed to learn dependencies across time steps, making them naturally suited to IoMT traffic analysis. This time-series perspective therefore provides the foundational justification for the deep learning architectures presented in Section 2.6 and supports the distributed learning strategies discussed in Section 2.7.

2.2.2 Feature Stream Notation

To formalize the temporal structure of IoMT data, we adopt the following notation that is used consistently throughout this thesis. Let i index a network node (e.g., a hospital site, a sensor gateway, or a client in the federated learning framework) and let t denote a discrete time index corresponding to individual flow observations or packet records. The feature vector observed at node i and time t is denoted by:

$$\mathbf{x}_i(t) \in \mathbb{R}^d, \tag{2.1}$$

where d is the dimensionality of the feature space. Depending on the dataset and preprocessing pipeline, features may include flow duration, byte counts, inter-arrival times, protocol flags, and other network telemetry attributes. The ordered collection $\{\mathbf{x}_i(t)\}_{t=1}^{T_i}$ for a given node i constitutes a multivariate time-series of length T_i that serves as the input to both the local intrusion detection model and the distributed learning pipeline. The main notation used throughout this chapter is summarized in Table 2.1.

Table 2.1: Principal notation used in Chapter 2.

Symbol	Meaning
i	Node/client index (hospital, gateway, or site)
t	Discrete time index (flow or packet observation)
d	Feature dimensionality
$\mathbf{x}_i(t)$	Feature vector at node i, time t
T_{agg}	Temporal aggregation window (time allocation parameter)
$\phi(\cdot)$	Aggregation operator (mean, max, variance, etc.)
$\mathbf{z}_i(k)$	Aggregated feature vector for window k at node i
$\mathbf{A}$	AHP pairwise comparison matrix
w_j	Weight of criterion j (derived via AHP)
C_i	TOPSIS closeness coefficient (trust score) for alternative i
$\mathbf{w}$	Global model parameters (federated learning)
p_k	Aggregation weight of client k in FedAvg

2.3 Time Allocation via Temporal Aggregation

The raw feature streams defined in Section 2.2 can generate extremely high data volumes when transmitted at full temporal resolution. In IoMT environments, where bandwidth and energy are severely constrained, transmitting every individual observation creates a fundamental scalability bottleneck. *Temporal aggregation* provides a principled mechanism for managing this challenge by introducing a controllable **time allocation** parameter—the aggregation interval T_{agg}—that governs the granularity at which data is summarized prior to transmission and downstream processing.

2.3.1 Motivation: The Scalability–Latency–Accuracy Trade-off

Transmitting every individual flow record or packet observation from each IoMT node incurs costs along three interdependent dimensions:

1. **Communication bandwidth**, which is inherently limited in wireless hospital environments and shared among potentially hundreds of concurrent sensor streams [23].

2. **Energy consumption**, which directly determines the operational lifetime of battery-powered sensor nodes. Since active computation and wireless transmission are the primary energy consumers on constrained devices [10], reducing the volume and frequency of data transmissions translates directly into energy savings.

3. **Computational load**, which affects both training and inference times of downstream machine learning models. Processing fewer, more compact feature vectors reduces the per-sample cost of gradient computations, forward passes, and model updates.

Conversely, aggressive aggregation can suppress short-lived attack signatures, thereby potentially reducing detection accuracy and increasing detection latency. A brief anomalous flow that occurs entirely within one aggregation window may be obscured when averaged with many benign observations.

This three-way tension between scalability, detection latency, and classification accuracy constitutes a central design challenge in IoMT communication systems. The temporal aggregation framework formalized below provides a tunable mechanism for

navigating this trade-off. Its experimental validation across multiple benchmark datasets—including CICIoMT-2024, NF-UNSW-NB15-v2, and WUSTL-EHMS-2020—is presented in Chapter 4.

2.3.2 Formal Definition of the Aggregation Window

Given the per-node time-series $\{\mathbf{x}_i(t)\}$ defined in Section 2.2.2, let $\tau_i(t)$ denote the physical timestamp of the t-th observation at node i. Temporal aggregation over a window of duration T_{agg} produces a reduced sequence $\{\mathbf{z}_i(k)\}$ by applying an aggregation operator $\phi(\cdot)$ independently over each feature dimension:

$$\mathbf{z}_i(k) = \phi\left(\left\{\mathbf{x}_i(t) \;\middle|\; \tau_i(t) \in \left[k \cdot T_{\text{agg}},\, (k+1) \cdot T_{\text{agg}}\right)\right\}\right), \qquad k \in \mathbb{N}_0, \tag{2.2}$$

where $\phi(\cdot)$ denotes a summary-statistic operator. Common choices for ϕ include the arithmetic mean, maximum, variance, or a concatenation of multiple statistics (e.g., $[\text{mean}, \max, \text{std}]$), yielding an aggregated feature vector of dimensionality $d' \geq d$. A partially filled final window is aggregated over the samples it contains.

The aggregation interval T_{agg} is the primary *time allocation* parameter of the framework:

- **Small T_{agg}** (in the limit where T_{agg} approaches the native sampling interval of the stream, which corresponds operationally to "no aggregation") preserves fine-grained temporal information at the cost of high data volume and computational expense.
- **Large T_{agg}** reduces communication and computational overhead substantially but risks obscuring transient phenomena such as short-lived attack bursts.

In this thesis, T_{agg} is treated as an adaptive design parameter. Chapter 4 systematically evaluates multiple aggregation settings—including no aggregation (raw stream), 10 s, 30 s, 60 s, and 120 s intervals—demonstrating that moderate aggregation (e.g., $T_{\text{agg}} = 30\,\text{s}$) achieves substantial reductions in training time (up to 97.2%) and test-set inference time (up to 97.5%) on CICIoMT-2024, with comparable gains on NF-UNSW-NB15-v2 (84.8% and 93.1%) and WUSTL-EHMS-2020 (99.3% and 99.1%), while maintaining competitive classification accuracy under the evaluated conditions.

2.3.3 Impact on Distributed Learning Efficiency

Beyond reducing local computational load, temporal aggregation exerts a multiplicative effect on the efficiency of distributed learning. In a federated learning setting

(Section 2.7.1), each client must: (i) perform local training on its private dataset, (ii) transmit model updates to the central aggregator, and (iii) receive the updated global model. When temporal aggregation is applied prior to local training, the number of effective training samples at each client is reduced by a factor proportional to T_{agg}, which in turn reduces:

- the number of gradient computation steps per local epoch,
- the wall-clock training time and associated energy consumption per communication round, and
- potentially the total number of communication rounds required to reach a target performance level, as aggregated representations tend to exhibit lower short-term variance, yielding more stable optimization dynamics.

Furthermore, when learning dynamics are more stable, client drift—a phenomenon in which local models diverge excessively under non-IID conditions [105]—may be mitigated, leading to improved global model convergence. These effects are empirically evaluated in the federated learning experiments presented in Chapter 4.

2.4 Fault Tolerance and Trust Foundations in IoMT

The reliability of IoMT communication is contingent upon the ability of the network to sustain correct operation in the presence of component failures and adversarial behavior. This section introduces the foundational concepts of fault tolerance as they apply to the IoMT domain, defines key reliability metrics, and motivates the lightweight trust-based decision approach adopted in this thesis.

2.4.1 Fault Taxonomy: Non-Byzantine and Byzantine Faults

In distributed systems theory, faults are broadly classified according to the behavioral assumptions about faulty components [27]. The distinction between non-Byzantine and Byzantine failures has profound implications for the design of fault-tolerant mechanisms.

Non-Byzantine faults. Non-Byzantine faults encompass failure modes in which the faulty component either ceases functioning or behaves in a detectable, predictable manner. The principal types include:

- **Crash faults:** A node ceases operation entirely and does not resume without external intervention.
- **Omission faults:** A node fails to send or receive messages despite being otherwise operational, often due to transient communication errors or queue overflows.
- **Resource depletion faults:** A sensor node's battery or memory is exhausted, causing degraded service quality or complete operational cessation.

These fault types are particularly prevalent in IoMT environments, where sensor nodes operate under severe energy and memory constraints, and wireless links are inherently unreliable [217].

Byzantine faults. Byzantine faults represent the most general and adversarial failure model, first formalized by Lamport et al. [123]. A component exhibiting Byzantine behavior may act arbitrarily: sending contradictory information to different parts of the system, selectively dropping or modifying messages, or actively attempting to subvert the system's correctness guarantees. In healthcare environments, Byzantine behavior can arise from compromised nodes (e.g., through malware injection), firmware manipulation, or insider attacks conducted by adversaries with physical access to the network infrastructure [98, 211].

A system is said to be *f-Byzantine fault-tolerant* if it produces correct outputs despite the arbitrary misbehavior of up to f of its n components. In the classical unauthenticated oral-messages model of Byzantine agreement, Lamport, Shostak, and Pease [123] established that solvability requires at least $n \geq 3f + 1$ nodes; the signed-message setting relaxes this threshold but is not assumed here. This bound imposes both redundancy and communication costs that may be prohibitive in resource-constrained settings.

2.4.2 Reliability Metrics

The reliability of a distributed IoMT system is commonly characterized by the following metrics:

Inherent (steady-state) availability. For a repairable node, the inherent or steady-state availability is the long-run fraction of time that the node is operational, excluding external logistics delays:

$$A_{\mathrm{i}} = \frac{\mathrm{MTTF}}{\mathrm{MTTF} + \mathrm{MTTR}}, \tag{2.3}$$

where MTTF denotes the Mean Time To Failure and MTTR denotes the Mean Time To Repair. This formulation assumes steady-state operation of the repair process and is the quantity used throughout this thesis when a single-number operational target is required; point-in-time and interval availability are related but distinct notions that are not used here.

Reliability. The probability that a system performs its intended function without failure over a specified time interval $[0, T]$.

In IoMT, achieving high availability is essential because interruptions in telemetry can delay clinical interventions with potentially life-threatening consequences. The fault tolerance mechanisms developed in Chapter 5 are designed to maintain high availability by proactively identifying and isolating unreliable or compromised nodes *before* they disrupt communication paths, rather than reacting after failures have already occurred.

2.4.3 Lightweight Trust Decisions vs. Heavy Consensus

Classical Byzantine Fault Tolerant (BFT) consensus protocols—such as Practical Byzantine Fault Tolerance (PBFT) [44]—guarantee correctness under up to $f = \lfloor (n-1)/3 \rfloor$ Byzantine failures among n nodes. In PBFT-style replication, however, the normal-case message complexity grows as $O(n^2)$ per consensus round due to the all-to-all prepare and commit phases, which is prohibitive in resource-constrained IoMT networks for several reasons:

1. Sensor nodes have limited processing power and cannot sustain the cryptographic verification required by BFT protocols.
2. The wireless medium introduces variable latency and packet loss that exacerbate round complexity and convergence time.
3. The number of active nodes in a hospital network may fluctuate due to intermittent connectivity, mobility, and device duty-cycling, making a fixed-membership consensus group impractical.

Instead of relying on full consensus, this thesis adopts a **lightweight decision-based approach** grounded in Multi-Criteria Decision Analysis (Section 2.5). The MCDA framework computes a composite trust score for each node based on multiple observable criteria (safety classification, residual energy, communication latency, and packet loss rate), enabling the system to rank and select trustworthy nodes without

requiring agreement from all participants. This approach trades the strong consistency guarantees of BFT for substantially lower communication overhead and practical deployability on constrained devices, while still providing effective fault isolation as demonstrated experimentally in Chapter 5.

2.5 Multi-Criteria Decision Analysis (MCDA) for Node and Model Prioritization

Multi-Criteria Decision Analysis encompasses a family of formal methods for evaluating alternatives against multiple, often conflicting, criteria [46, 160]. In the context of IoMT, MCDA provides a rigorous framework for selecting trustworthy routing nodes, allocating network resources, and prioritizing candidate models. This section presents the two MCDA techniques employed in this thesis: the Analytic Hierarchy Process (AHP) for criteria weighting and TOPSIS for alternative ranking.

2.5.1 The Decision Problem in IoMT

IoMT systems must simultaneously optimize across objectives that are frequently in tension. A sensor node that offers low communication latency may consume excessive energy, shortening its operational lifetime. A highly reliable node may reside on a congested link with elevated packet loss. An intrusion detection model that achieves high detection accuracy may impose an unacceptable false alarm rate or computational overhead under certain traffic conditions. Making principled trade-offs among these criteria requires a structured methodology that can incorporate both domain-expert knowledge (through weighting) and objective performance measurements (through normalization and ranking).

2.5.2 AHP for Criteria Weighting

The Analytic Hierarchy Process (AHP), introduced by Saaty [169], decomposes a decision problem into a hierarchy consisting of three levels: a *goal* (e.g., selecting the most trustworthy routing node), a set of *criteria* (e.g., safety, energy, latency, packet loss), and a set of *alternatives* (e.g., available network nodes). The relative importance of criteria is determined through pairwise comparisons elicited from domain experts.

Pairwise comparison matrix. Given n criteria, a pairwise comparison matrix $\mathbf{A} = [a_{ij}]_{n \times n}$ is constructed, where the entry a_{ij} represents the relative importance of cri-

terion i over criterion j on Saaty's fundamental scale (1–9). The matrix satisfies the reciprocity property $a_{ij} \cdot a_{ji} = 1$ and $a_{ii} = 1$ for all i:

$$\mathbf{A} = \begin{bmatrix} 1 & a_{12} & \cdots & a_{1n} \\ 1/a_{12} & 1 & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 1/a_{1n} & 1/a_{2n} & \cdots & 1 \end{bmatrix}. \tag{2.4}$$

Weight derivation. The priority weight vector $\mathbf{w} = (w_1, w_2, \ldots, w_n)^\top$ is obtained as the principal eigenvector of $\mathbf{A}$:

$$\mathbf{A}\,\mathbf{w} = \lambda_{\max}\,\mathbf{w}, \tag{2.5}$$

where $\lambda_{\max}$ is the largest eigenvalue. The normalized eigenvector yields the relative importance weight for each criterion.

Consistency verification. To ensure that the expert judgments are coherent, AHP employs a Consistency Ratio. For $n \geq 3$:

$$\mathrm{CR} = \frac{\mathrm{CI}}{\mathrm{RI}}, \qquad \mathrm{CI} = \frac{\lambda_{\max} - n}{n-1}, \tag{2.6}$$

where RI is the Random Index—a tabulated constant that depends on matrix size n (e.g., RI$= 0.58$ for $n = 3$, RI$= 0.90$ for $n = 4$, RI$= 1.12$ for $n = 5$). For $n = 1$ or $n = 2$ the tabulated RI is zero and the ratio is not meaningful; in the regime $n \geq 3$, judgments are considered acceptably consistent when $\mathrm{CR} < 0.1$, otherwise the pairwise comparisons should be revised [169].

2.5.3 TOPSIS for Ranking and Trust Scoring

The Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS), developed by Hwang and Yoon [93], ranks alternatives based on their geometric proximity to a Positive Ideal Solution (PIS) and distance from a Negative Ideal Solution (NIS). The method proceeds through five steps.

Step 1: Decision matrix construction and normalization. Let $\mathbf{X} = [X_{ij}]$ be the $m \times n$ decision matrix, where m is the number of alternatives (nodes) and n is the number of criteria. Each entry is normalized by a min–max transform that simultaneously converts *cost criteria* (lower is better, e.g., latency or packet-loss rate) and *benefit criteria* (higher is better, e.g., node safety or residual energy) into a common benefit-

oriented score $r_{ij} \in [0,1]$. For cost criteria:

$$r_{ij} = \frac{X_{j,\max} - X_{ij}}{X_{j,\max} - X_{j,\min}}. \tag{2.7}$$

For benefit criteria:

$$r_{ij} = \frac{X_{ij} - X_{j,\min}}{X_{j,\max} - X_{j,\min}}. \tag{2.8}$$

Both transforms map the criterion to $[0,1]$ and orient it so that values approaching 1 denote superior performance. In the degenerate case where $X_{j,\max} = X_{j,\min}$ for some criterion j, the criterion provides no discriminative information and may be assigned a constant normalized value or excluded from the analysis.

Step 2: Weighted normalized matrix. Using the AHP-derived weights w_j with $\sum_{j=1}^{n} w_j = 1$, the weighted normalized matrix is computed as:

$$v_{ij} = w_j \cdot r_{ij}. \tag{2.9}$$

Step 3: Ideal and negative-ideal solutions. Because the normalization step of Equations (2.7)–(2.8) has already re-oriented every criterion so that larger weighted scores always correspond to better performance, the Positive Ideal Solution $\mathbf{A}^+$ and the Negative Ideal Solution $\mathbf{A}^-$ are identified component-wise as the maximum and minimum of the weighted normalized matrix:

$$A_j^+ = \max_{1 \le i \le m} v_{ij}, \qquad j = 1, \ldots, n, \tag{2.10}$$

$$A_j^- = \min_{1 \le i \le m} v_{ij}, \qquad j = 1, \ldots, n. \tag{2.11}$$

No cost/benefit split is required at this stage, since that split has been absorbed into the normalization step.

Step 4: Separation measures. The Euclidean separation of each alternative from the Positive and Negative Ideal Solutions is computed directly in the weighted normalized space:

$$D_i^+ = \sqrt{\sum_{j=1}^{n} \left(v_{ij} - A_j^+ \right)^2}, \tag{2.12}$$

$$D_i^- = \sqrt{\sum_{j=1}^{n} \left(v_{ij} - A_j^- \right)^2}. \tag{2.13}$$

Weights are not reintroduced here, because the criterion weights w_j have already been incorporated through the weighted normalized entries v_{ij} defined in Equation (2.9).

Step 5: Closeness coefficient and ranking. Provided that $D_i^+ + D_i^- > 0$, the relative closeness of each alternative to the ideal solution is:

$$C_i = \frac{D_i^-}{D_i^+ + D_i^-}, \qquad C_i \in [0,1]. \tag{2.14}$$

Alternatives are ranked in descending order of C_i; a higher closeness coefficient indicates superior overall performance against the evaluated criteria. In Chapter 5, this closeness score is interpreted as a dynamic *trust score*, and nodes are classified into operational categories—"Best," "Acceptable," or "Non-Acceptable"—using predefined thresholds applied to C_i.

2.5.4 Role of MCDA in the Proposed Framework

The MCDA methodology described above serves two distinct and complementary roles in this thesis:

- **Fault-tolerant routing (Chapter 5):** AHP-weighted TOPSIS computes dynamic trust scores for each routing node based on safety classification, residual energy, communication latency, and packet loss. These scores enable proactive isolation of unreliable or compromised nodes before they disrupt communication paths.
- **Client-side model selection (Chapter 4):** A related weighted multi-criteria scoring mechanism is employed to select among candidate intrusion detection models under conflicting performance metrics (e.g., accuracy, false alarm rate, detection latency).

The common thread is the use of structured multi-objective decision analysis to manage the inherent trade-offs in IoMT systems, whether the "alternatives" are network nodes or machine learning models.

2.6 Deep Learning for Sequential Intrusion Detection

Network traffic in IoMT environments exhibits temporal dependencies that classical machine learning algorithms—which treat each observation independently—fail to capture effectively. This section presents the deep learning architectures for sequential data that form the detection backbone of the proposed framework, progressing from basic

RNN foundations through LSTM gating to the hierarchical architecture employed in later chapters.

2.6.1 Recurrent Neural Network Foundations

Recurrent Neural Networks (RNNs) extend feedforward architectures by introducing directed cycles in their computation graph, enabling the network to maintain a hidden state that serves as a form of "memory" over the input sequence [167]. At each time step t, the hidden state $\mathbf{h}_t$ is updated as:

$$\mathbf{h}_t = \tanh(\mathbf{W}_{xh}\,\mathbf{x}_t + \mathbf{W}_{hh}\,\mathbf{h}_{t-1} + \mathbf{b}_h)\,, \tag{2.15}$$

where $\mathbf{x}_t \in \mathbb{R}^d$ is the input vector at time t, $\mathbf{W}_{xh} \in \mathbb{R}^{h\times d}$ and $\mathbf{W}_{hh} \in \mathbb{R}^{h\times h}$ are learnable weight matrices, and $\mathbf{b}_h \in \mathbb{R}^h$ is a bias vector. The output logit at each step is computed as:

$$\mathbf{y}_t = \mathbf{W}_{hy}\,\mathbf{h}_t + \mathbf{b}_y. \tag{2.16}$$

For classification tasks, a final activation such as $\mathbf{softmax}(\cdot)$ is then applied to $\mathbf{y}_t$ to obtain class probabilities.

This recurrence enables the RNN to process sequences of variable length and, in principle, to capture dependencies across arbitrary time spans. In practice, however, the training of standard RNNs is severely limited by the vanishing gradient problem, as discussed in the following subsection.

2.6.2 The Vanishing Gradient Problem

Training RNNs via Backpropagation Through Time (BPTT) requires computing the gradient of the loss function with respect to parameters at early time steps. This computation involves an *ordered* product of Jacobian matrices across all intervening time steps:

$$\frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_k} = \frac{\partial \mathbf{h}_t}{\partial \mathbf{h}_{t-1}}\,\frac{\partial \mathbf{h}_{t-1}}{\partial \mathbf{h}_{t-2}} \cdots \frac{\partial \mathbf{h}_{k+1}}{\partial \mathbf{h}_k}, \qquad 0 \le k < t. \tag{2.17}$$

As the temporal distance $(t-k)$ increases, this product tends to either vanish (when the spectral radius of the Jacobians is consistently less than 1, causing exponential decay) or explode (when it is consistently greater than 1, causing exponential growth) [32, 154]. The vanishing gradient problem renders standard RNNs unable to learn dependencies spanning more than a few dozen time steps, which is insufficient for detecting sophisticated network attacks that evolve over extended periods or exhibit long-range temporal correlations.

2.6.3 LSTM Gating Mechanism

Long Short-Term Memory (LSTM) networks were introduced by Hochreiter and Schmidhuber [90] to address the vanishing gradient problem through an architecture that explicitly controls the flow of information via a memory cell and multiplicative gates. The gate equations adopted below correspond to the common later vanilla LSTM family: the original 1997 cell included input and output gates, whereas the forget gate was introduced in the subsequent extension and has since become standard. An LSTM cell thus comprises three multiplicative gates—the forget gate, the input gate, and the output gate—together with a cell state that enables information to persist across many time steps via an additive update rule. The internal operations and gating mechanism of the LSTM cell are illustrated in Figure 2.1. Given input $\mathbf{x}_t$ and the previous hidden state $\mathbf{h}_{t-1}$:

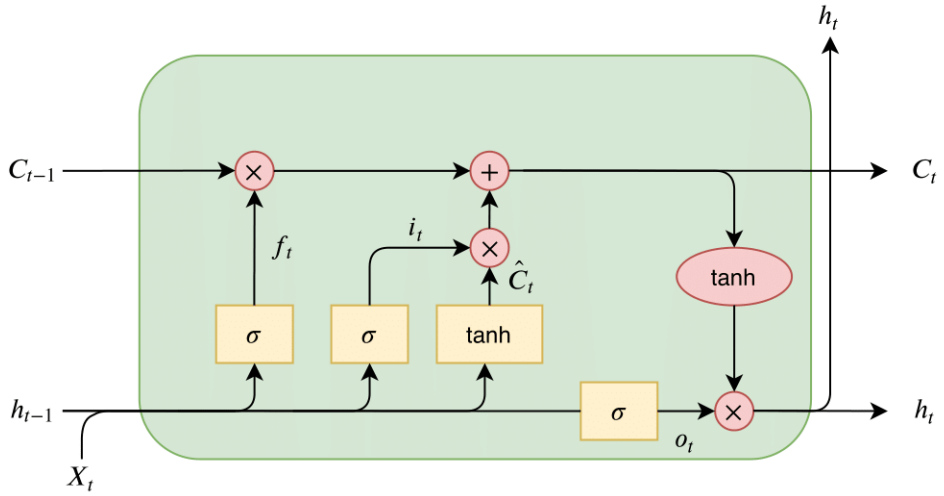

Figure 2.1: Internal architecture of a Long Short-Term Memory (LSTM) cell, illustrating the flow of data through the forget, input, and output gates.

Forget gate. The forget gate determines which components of the previous cell state should be retained:

$$\mathbf{f}_t = \sigma\big(\mathbf{W}_f[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_f\big), \tag{2.18}$$

where $\sigma(\cdot)$ denotes the sigmoid activation function and $[\cdot,\cdot]$ denotes vector concatenation.

Input gate and candidate state. The input gate controls what new information is added to the cell state:

$$\mathbf{i}_t = \sigma(\mathbf{W}_i[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_i), \tag{2.19}$$

$$\tilde{\mathbf{C}}_t = \tanh(\mathbf{W}_C[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_C). \tag{2.20}$$

Cell state update. The cell state is updated by selectively forgetting old information and incorporating new candidate values:

$$\mathbf{C}_t = \mathbf{f}_t \odot \mathbf{C}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{C}}_t, \tag{2.21}$$

where $\odot$ denotes the Hadamard (element-wise) product. The *additive* cell update (as opposed to a purely multiplicative one) is the key architectural innovation that enables gradients to flow across many time steps without vanishing, since the gradient of $\mathbf{C}_t$ with respect to $\mathbf{C}_{t-1}$ includes a direct additive path through the forget gate.

Output gate and hidden state. The output gate determines which parts of the cell state are exposed as the hidden state:

$$\mathbf{o}_t = \sigma(\mathbf{W}_o[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_o), \tag{2.22}$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{C}_t). \tag{2.23}$$

2.6.4 Hierarchical LSTM for Multi-Scale Temporal Modeling

While standard LSTM networks effectively capture temporal dependencies at a single granularity, IoMT traffic exhibits patterns across multiple temporal scales: fine-grained packet-level dynamics, medium-grained flow-level behaviors, and coarse-grained session-level trends. The Hierarchical LSTM (HLSTM) architecture [185] addresses this multi-scale requirement through a two-level processing pipeline.

In the HLSTM architecture employed in this thesis, the *first level* performs binary classification (normal vs. anomalous) on the raw or lightly aggregated input stream, providing a rapid initial filtering that separates benign traffic from potentially malicious flows. The *second level* receives the output of the first level—augmented with temporally aggregated features computed over window T_{agg}—and performs fine-grained multi-class attack classification to identify specific attack types.

This hierarchical decomposition offers two principal advantages:

1. **Computational efficiency:** By pre-filtering normal traffic at the first level, the multi-class classifier at the second level processes a substantially reduced volume of data, decreasing overall computational cost.

2. **Tunable detection granularity:** The aggregation interval T_{agg} serves as a tunable parameter that the system administrator can adjust to balance detection

granularity against communication and computational efficiency. This alignment between temporal aggregation (Section 2.3) and hierarchical processing is a central design principle of the proposed framework.

The complementarity between temporal aggregation and HLSTM processing is validated experimentally in Chapter 4, where the HLSTM architecture is shown to improve accuracy over standard LSTM by 6.79 percentage points while requiring 45.4% less training time.

2.7 Privacy-Preserving Distributed Learning: FL, TL, and FTL

Traditional centralized machine learning requires aggregating all training data at a single location, which is infeasible in IoMT environments where privacy regulations (e.g., HIPAA in the United States [201], GDPR in the European Union [66]) prohibit the sharing of protected health information across institutional boundaries. Distributed learning paradigms address this fundamental constraint by enabling collaborative model training without centralizing raw data.

2.7.1 Federated Learning and FedAvg

Federated Learning (FL), introduced by McMahan et al. [135], is a distributed machine learning paradigm in which multiple clients (e.g., hospitals, sensor gateways, or clinical sites) collaboratively train a shared model under the coordination of a central aggregator, without exchanging their raw data. Each client trains a local model on its private dataset and transmits only model parameters (e.g., neural network weights) to the aggregator.(Fig. 2.2).

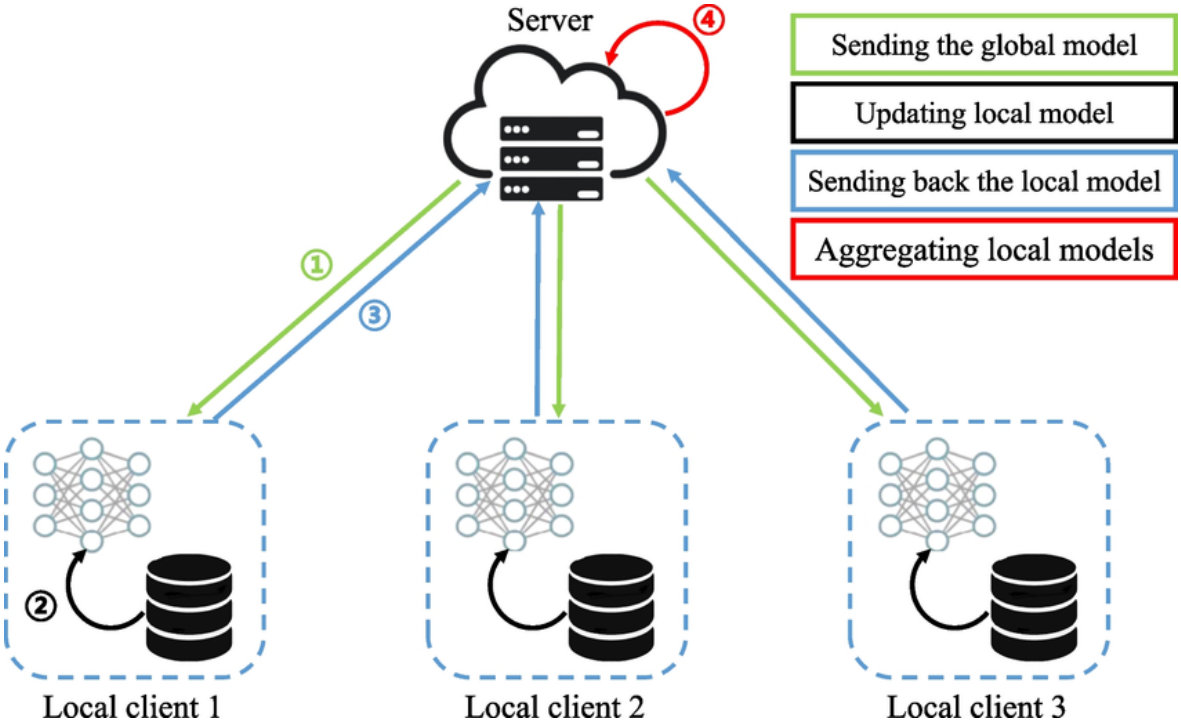

Figure 2.2: The Federated Averaging (FedAvg) framework for distributed machine learning.

The FedAvg algorithm. Let $S_r \subseteq \{1,2,\ldots,N\}$ denote the subset of clients that participate in communication round r. Federated Averaging (FedAvg) proceeds iteratively over rounds $r = 0,1,2,\ldots$:

1. The aggregator broadcasts the current global model parameters $\mathbf{w}^{(r)}$ to the clients in S_r.
2. Each selected client $k \in S_r$ performs several epochs of local stochastic gradient descent on its private data, producing updated local parameters $\mathbf{w}_k^{(r+1)}$.
3. Clients in S_r transmit their updated local parameters to the aggregator.
4. The aggregator forms the new global model as a weighted average over the participating clients:

$$\mathbf{w}^{(r+1)} = \sum_{k \in S_r} p_k^{(r)} \mathbf{w}_k^{(r+1)}, \qquad p_k^{(r)} = \frac{n_k}{\sum_{\ell \in S_r} n_\ell}, \tag{2.24}$$

where n_k is the number of training samples at client k. The corresponding global optimization objective, defined over the *full* federation of N clients, is:

$$\mathbf{w}^* = \arg\min_{\mathbf{w}} f(\mathbf{w}) = \arg\min_{\mathbf{w}} \sum_{k=1}^{N} \frac{n_k}{n} F_k(\mathbf{w}), \qquad n = \sum_{k=1}^{N} n_k, \tag{2.25}$$

where $F_k(\mathbf{w})$ is the local empirical loss function at client k. Equation (2.24) therefore defines the round-specific aggregation rule, whereas Equation (2.25) defines the global objective that the procedure is intended to minimize.

Privacy mechanisms. FL inherently preserves privacy by ensuring that raw data never leaves the client device. However, model updates may still leak information under certain inference attacks [215]. To mitigate this risk, additional privacy layers can be incorporated depending on the threat model:

- **Secure Aggregation** [39]: Cryptographic protocols ensure that the aggregator can compute the weighted average without observing any individual client's update.
- **Differential Privacy** [2, 60]: Calibrated noise is added to local updates before transmission, providing formal mathematical guarantees that the presence or absence of any single training sample cannot be inferred from the published model.

The framework proposed in Chapter 4 is designed to be compatible with both mechanisms, ensuring extensibility to deployment scenarios with varying privacy requirements.

2.7.2 Non-IID Data and System Heterogeneity in Healthcare IoMT

Despite its privacy advantages, FL in healthcare settings faces several practical challenges that motivate the hybrid approach adopted in this thesis.

Non-IID data distributions. In supervised learning, non-IID (non-Independent and Identically Distributed) conditions arise when each client k draws samples $(\mathbf{x}, y)$ from a unique joint distribution $P_k(\mathbf{x}, y)$ [100]. In IoMT, this is the norm rather than the exception: different hospitals encounter different patient populations, disease profiles, sensor configurations, and attack patterns. Non-IID effects are commonly decomposed into three canonical forms of statistical heterogeneity [208, 214], which manifest primarily as:

- **Attribute skew:** Significant differences in feature distributions $P_k(\mathbf{x})$ across clients, arising from heterogeneous sensor types, network configurations, or patient demographics.

- **Label distribution skew:** Variations in label distributions $P_k(y)$ across clients—for example, one hospital may experience predominantly DDoS attacks while another faces reconnaissance or man-in-the-middle attacks—despite consistent conditional feature distributions $P_k(\mathbf{x}|y)$.
- **Label preference skew:** Variations in conditional label distributions $P_k(y|\mathbf{x})$ despite identical feature distributions, common in subjective labeling scenarios.

Under severe non-IID conditions, FedAvg can suffer from *client drift* [105], where local models diverge significantly during multiple epochs of local training, and the resulting averaged global model performs poorly for all clients. This fundamental limitation motivates the integration of transfer learning with federated learning, as discussed in Sections 2.7.3 and 2.7.4.

Communication efficiency. Frequent exchange of model parameters between clients and the aggregator can be bandwidth-intensive, particularly for deep neural networks with millions of parameters. Temporal aggregation (Section 2.3) directly alleviates this burden by reducing the local dataset size and, consequently, the number of local training iterations and communication rounds required to reach a target performance level.

System heterogeneity. Clients in an IoMT federation may possess vastly different computational capabilities, ranging from resource-constrained edge gateways to powerful hospital servers. Stragglers—clients that are significantly slower than others due to hardware limitations, network congestion, or competing workloads—can delay the global training process, motivating hybrid strategies such as partial participation, client selection, and transfer-based adaptation.

2.7.3 Transfer Learning

Transfer Learning (TL) [150, 216] addresses the challenge of limited labeled data by leveraging knowledge from a source domain to improve learning in a related target domain. Formally, a *domain* $\mathscr{D} = \{\mathscr{X}, P(\mathscr{X})\}$ consists of a feature space $\mathscr{X}$ and a marginal distribution $P(\mathscr{X})$ over it, and a *task* $\mathscr{T} = \{\mathscr{Y}, f(\cdot)\}$ consists of a label space $\mathscr{Y}$ and a predictive function $f : \mathscr{X} \to \mathscr{Y}$ learned from training data. Given a source domain $\mathscr{D}_S$ with task $\mathscr{T}_S$ and a target domain $\mathscr{D}_T$ with task $\mathscr{T}_T$, transfer learning aims to improve the target predictive function $f_T(\cdot)$ using knowledge from $\mathscr{D}_S$ and $\mathscr{T}_S$ under the assumption that $\mathscr{D}_S \neq \mathscr{D}_T$ or $\mathscr{T}_S \neq \mathscr{T}_T$.

In the context of IoMT intrusion detection, TL enables a model pre-trained on a comprehensive general-purpose network traffic dataset to be fine-tuned on a smaller, domain-specific IoMT dataset. This is particularly valuable for three reasons:

1. Labeled IoMT attack data is scarce and costly to obtain due to privacy constraints and the difficulty of conducting controlled attack experiments in clinical settings.
2. Pre-training on diverse traffic captures generic feature representations (e.g., temporal flow patterns, statistical protocol features) that transfer effectively across network domains [133].
3. Fine-tuning requires substantially fewer computational resources and training samples than training from scratch, making it suitable for resource-constrained environments.

The principal TL strategies employed in deep learning are *feature extraction* (using the pre-trained model as a fixed feature extractor), *fine-tuning* (updating some or all layers on the target data), and *progressive unfreezing* (gradually unfreezing layers during adaptation). In this thesis, fine-tuning is the primary transfer mechanism, as it provides the most flexible adaptation to domain-specific IoMT traffic characteristics.The detailed mechanism presented in Fig. 2.3.

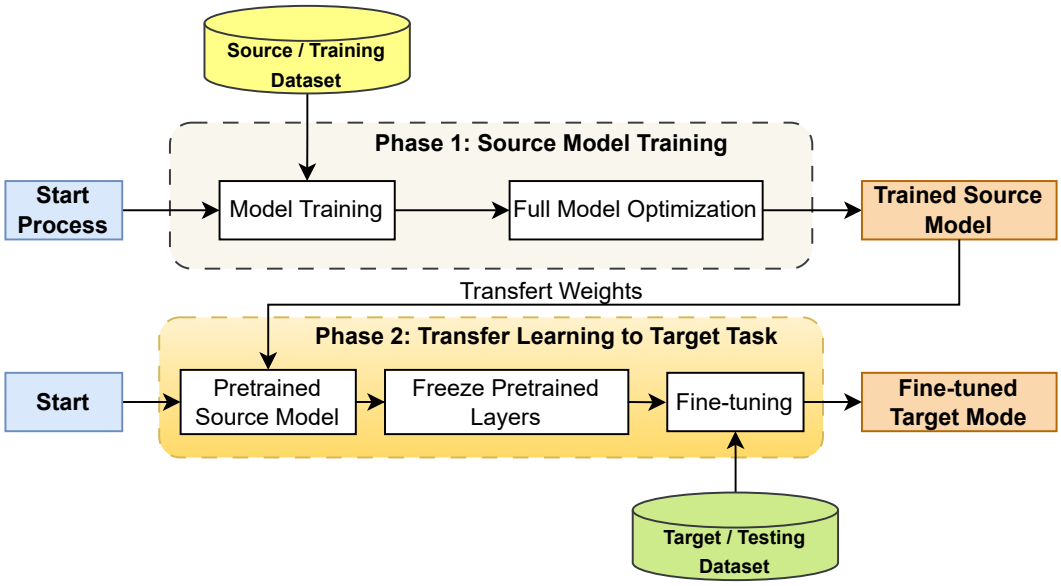

Figure 2.3: Flowchart of the transfer learning mechanism.

2.7.4 Federated Transfer Learning

Federated Transfer Learning (FTL) [99, 131] combines the privacy-preserving properties of FL with the domain adaptation capabilities of TL. FTL is designed for scenarios in which:

1. Different institutions hold data from related but distinct domains (domain shift due to differing patient populations, sensor types, or clinical workflows).
2. Direct data sharing between institutions is prohibited by privacy regulations.
3. Individual local datasets may be too small or insufficiently diverse for effective standalone model training.

In the FTL framework proposed in Chapter 4, an **Intelligent Label Classification** algorithm distinguishes between globally shared ("Common") attack patterns and locally specific ("Isolated") attack patterns. Common patterns—attack types observed across multiple hospital sites—are learned collaboratively through federated aggregation, benefiting from the statistical power of the combined data. Isolated patterns—attack types that appear only at specific sites due to local network configurations or targeted adversaries—are handled through local transfer learning, where knowledge from the globally trained model is adapted to the local distribution. This partitioning enables the framework to achieve high detection accuracy across heterogeneous hospital sites without requiring centralized data pooling, effectively addressing the non-IID challenge identified in Section 2.7.2.

2.8 Conclusion

This chapter has presented the theoretical foundations upon which the proposed framework is built. The exposition was organized around **six complementary pillars**:

1. **IoMT telemetry as time-series data:** IoMT network traffic was characterized as multivariate sequential telemetry, motivating the use of temporal models and establishing the feature stream notation used throughout the thesis.
2. **Temporal aggregation as time allocation:** The aggregation window T_{agg} was formalized as the primary time allocation parameter governing the scalability–latency–accuracy trade-off, with its impact on both local computation and distributed learning efficiency explicitly analyzed.

3. **Fault tolerance and trust foundations:** Non-Byzantine and Byzantine fault models were introduced, core reliability metrics (availability, MTTF, MTTR) were defined, and the motivation for lightweight MCDA-based trust decisions over heavy consensus protocols was established.

4. **MCDA decision logic:** AHP for criteria weighting and TOPSIS for alternative ranking were presented as the formal basis for trust-based node selection and multi-objective model prioritization.

5. **Sequential deep learning:** RNN foundations, the vanishing gradient problem, the LSTM gating mechanism, and the HLSTM multi-scale architecture were described, establishing the detection backbone of the proposed system.

6. **Privacy-preserving distributed learning:** FL, TL, and FTL were introduced as scalable learning paradigms that enable collaborative model training across heterogeneous, privacy-sensitive healthcare institutions.

These pillars are intentionally interconnected: temporal aggregation reduces the data volume that distributed learning must process and yields more stable optimization dynamics; MCDA provides the decision logic that both fault-tolerant routing and model selection require; and the HLSTM architecture leverages the temporal structure of IoMT traffic while benefiting from aggregation to achieve efficient, accurate intrusion detection.

Chapter 3 presents a comprehensive review of the state of the art, positioning the contributions of this thesis within the broader research landscape. Subsequently, Chapter 4 integrates these theoretical foundations into the unified FTL-HLSTM intrusion detection framework, and Chapter 5 develops the MCDA-based fault-tolerant routing mechanism.

Chapter

3

State of the Art

Chapter 3

State of the Art in Security and Resilience for IoT and IoMT

3.1 Introduction

The Internet of Things (IoT), and its medical extension the Internet of Medical Things (IoMT), now underpin infrastructures as varied as smart industrial plants and continuously monitored hospital wards. The cost of this ubiquity is an attack surface that has grown faster than the defensive tooling around it. Heterogeneous protocol stacks (MQTT, CoAP, Bluetooth LE, Zigbee, 6LoWPAN), severely resource-constrained end devices, intermittent edge connectivity, and—on the IoMT side specifically—some of the most demanding confidentiality and availability requirements found anywhere in IoT [55, 73, 91] combine to create a threat environment in which traditional network-security assumptions rarely hold unchallenged. Recent incidents reinforce this diagnosis: the 2024 Change Healthcare outage [202] and repeated advisories on medical infusion pumps [53, 199] have translated an abstract risk into a regulatory one, under the HIPAA Security Rule [201], the General Data Protection Regulation [66], and pre-market cybersecurity guidance from the U.S. Food and Drug Administration [200].

This chapter reviews the state of the art along three tightly coupled axes and uses that review to identify the gaps that the contribution chapters of this thesis set out to address. The first axis concerns intrusion-detection frameworks for IoT and IoMT, traced from classical cryptography and signature-based firewalls through centralized machine learning, deep learning, and federated learning. The second concerns hierarchical deep-learning architectures, which exploit the multi-scale structure that network traffic displays at the packet, flow, and session levels. The third concerns fault-tolerance and resilience mechanisms for distributed IoT systems, where routing, consensus, and trust management have to accommodate both benign failures and active adversaries. The objective is deliberately analytical rather than exhaustive: the review is intended to clarify where the literature is already strong, where it is thin, and where the present thesis can reasonably be expected to add value—without pre-validating Chapters 4 and 5, which are the empirical place for that validation.

3.2 Security Solutions in IoT and IoMT

This section reviews the IoT and IoMT security literature in three steps. Section 3.2.1 surveys the cryptography- and signature-based defenses that constitute the classical stack. Section 3.2.2 examines the structural conditions under which those defenses become insufficient in constrained and heterogeneous networks, and, consequently, why the literature has moved towards learning-based detection. Section 3.2.3 surveys the AI-driven intrusion-detection work that has emerged in response. A note on framing is useful at the outset: AI-driven intrusion detection should be read as a *complement* to cryptography, authentication, access control, and network segmentation, not as a substitute for them. None of the learning-based systems reviewed below removes the need for a correctly configured cryptographic baseline.

3.2.1 Traditional Security Mechanisms and Protocols

By *traditional security mechanisms* this review denotes the family of static protection schemes that predate the widespread adoption of machine learning in IoT and IoMT defenses: applied cryptography, entity authentication, integrity verification, and network-level controls such as firewalls. Their common property is that detection rules are fixed at design time rather than learned from observed traffic. Sadhu *et al.* [171] offer the most recent broad survey in this space, classifying secret-key cryptography, Physical Unclonable Functions (PUFs), and blockchain primitives across the three canonical IoT layers (perception, network, application). Abosata *et al.* [4] revisit the same primitives from the Industrial IoT side, and Attkan and Ranga [26] extend the discussion to cyber-physical systems, arguing—persuasively, though without quantitative backing—that AI-assisted key management is in the process of being absorbed into what used to be a purely classical stack.

Lightweight cryptography. Thakor *et al.* [192] compare more than fifty block- and stream-cipher variants tailored to resource-constrained devices, tabulating the trade-off between cycle count, memory footprint, and security margin. Their comparison makes it clear that no single cipher dominates across all three axes, which is itself informative: the apparent abundance of lightweight primitives masks a recurring need to pick the right one per deployment, and to re-pick when hardware or threat models change. Pandey *et al.* [151] update this landscape with recent ECC-based constructions and argue, on energy grounds, in favor of hybrid schemes that couple a lightweight block cipher with an elliptic-curve key-agreement primitive.

Elliptic-curve cryptography and authentication. El-Hajj *et al.* [61] propose a decentralized zone-based Public-Key Infrastructure for IoT that relies on Elliptic-Curve Cryptography (ECC) to shrink certificate sizes and distribute trust across administrative zones. Yang *et al.* [206] pair ECC with trusted-token authentication for Industrial IoT gateways, reporting a substantial reduction in handshake latency relative to classical TLS. Alanazi [15] extends the line further with a triple-layer authentication protocol and documents the baseline ECC handshake cost on Class-1 sensor nodes—a useful dose of realism given how often such costs are elided in IoT security papers.

PUFs, blockchain, and quantum key distribution. Al-Meer and Al-Kuwari [12] offer a decade-long retrospective on PUF-based primitives for IoT. The survey is genuinely useful, but it also underscores a recurring concern with PUFs that is rarely resolved: their susceptibility to machine-learning-based modeling attacks, which undermines the key-less-authentication selling point whenever an adversary has access to enough challenge–response pairs. Dhar *et al.* [56] investigate hybrid schemes that couple blockchain consensus with Quantum Key Distribution (QKD) as a long-horizon post-quantum pathway for IoT trust anchors—a plausible direction in principle, though the deployment-readiness for hospital-grade IoMT endpoints remains unproven.

Firewalling in constrained networks. Rajasoundaran *et al.* [163] review firewall architectures for resource-limited wireless networks and compare signature-based and policy-based designs. Their finding that both families degrade under multi-protocol IoMT traffic is a telling one: it is not just that signature databases lag, but that the firewall abstraction itself is under-specified for a regime where MQTT, CoAP, BLE, and Wi-Fi coexist on the same wristband-to-gateway path. This observation resurfaces in Section 3.2.2.

Taken together, these works form the non-AI baseline against which the AI-driven intrusion-detection literature surveyed in Section 3.2.3 is positioned. What they share—and what distinguishes them from the learning-based work discussed later—is that security is enforced *statically*, through cryptographic guarantees or manually curated signature databases, rather than *learned* from observed traffic. This distinction, rather than any of the individual contributions listed above, motivates the limitations analysis that follows.

3.2.2 Limitations of Conventional Approaches in Constrained Networks

Traditional mechanisms remain necessary in IoT and IoMT, but on their own they are increasingly insufficient in constrained and heterogeneous environments. The limitations reported most frequently in the reviewed literature, and consolidated in Table 3.1, can be grouped into five recurring patterns.

Computational and energy overhead. Classical cryptographic primitives such as RSA-2048 and AES-256 can exceed the compute budget of severely resource-constrained IoT and IoMT devices without sacrificing either sensing frequency or battery lifetime [151, 192]. Lightweight cryptography partially absorbs this burden, but the residual energy cost of per-packet cryptographic operations remains non-negligible for always-on medical wearables [15, 171]. The issue is not that the primitives are too expensive in absolute terms; it is that the power envelope of a wristband-class device is so tight that even a few percent of extra duty cycle per hour translates into a measurable reduction of deployment autonomy.

Zero-day and behavioral blindness. Signature-based firewalls are structurally limited against zero-day exploits or protocol-abuse attacks that ride on legitimate packet shapes, because their decision rules are anchored in a finite, manually curated knowledge base [163, 171]. Cryptography itself detects no anomaly: a compromised endpoint that still holds a valid key continues to transmit cryptographically well-formed traffic and so passes perimeter checks [20, 81]. This is the clearest structural reason why a cryptographic baseline, however correctly deployed, cannot be the only layer of defense.

Key-management brittleness. Centralized Public-Key Infrastructures remain a single point of failure whose compromise can invalidate large device populations at once [4, 61]. In heterogeneous IoMT fleets, key rotation and certificate revocation are rarely performed at the cadence implied by the underlying risk model, which widens the gap between paper guarantees and operational reality [15, 206].

Limited adaptivity. Static defenses cannot track either the slow drift of legitimate traffic or the faster mutation of attack campaigns. Gugueoth *et al.* [81] and Alwahedi *et al.* [20] both argue that learning-based mechanisms are needed to sustain detection accuracy over time in realistic IoT deployments; the evidence they provide is

compelling for IoT broadly, although less so for IoMT specifically, where comparable longitudinal evaluations are still scarce.

Unmanaged heterogeneity. Classical cryptography has difficulty enforcing a coherent policy across the multi-protocol, multi-vendor reality of IoMT, in which MQTT, CoAP, BLE, and Wi-Fi routinely coexist on the same wristband-to-gateway path [4, 171]. What is missing from the reviewed work is not a primitive but a cross-layer abstraction that could be instantiated uniformly across these stacks.

The principal limitations discussed above are summarized in Table 3.1.

Table 3.1: Limitations most frequently reported for traditional security mechanisms in IoT and IoMT, and the sources documenting them.

Limitation	**Description**	**Sources**
Computational complexity	The complexity of RSA/AES easily surpasses the limits of ultra-low-power devices.	[15, 151, 192]
Static signatures	Signature firewalls struggle with zero-day attacks	[163]
Key-management brittleness	Centralized PKI can create single points of failure	[61]
Behavioral blindness	Cryptography does not detect anomalous usage	[81, 171]
Energy overhead	Heavy crypto can shorten IoT battery life	[163, 192]
Limited adaptivity	Static schemes do not track evolving attacks	[20, 81]
Unmanaged heterogeneity	Single-policy cryptography struggles across multi-protocol IoMT	[4, 171]

3.2.3 Artificial Intelligence-Driven Intrusion Detection Frameworks

AI-driven intrusion detection systems (IDS) for IoT and IoMT fall into three paradigms that roughly follow the field's historical trajectory: *centralized* machine learning over hand-crafted features; *deep learning* with end-to-end representation learning; and *federated learning* under privacy-preserving distributed optimization. Framing them purely as technological generations, as most recent surveys do, risks understating how much each of them is in fact a *response* to a specific reported weakness of the previous one. The discussion that follows is organized accordingly: each paradigm is presented together with the problem it inherited and the new trade-off it introduced.

Centralized Machine Learning Models for Anomaly Detection

Centralized machine-learning models remain a surprisingly strong baseline in IoMT security, for three practical reasons: interpretability, low training cost, and tooling maturity. Binbusayyis and Vaiyapuri [37] couple an autoencoder with a one-class Support Vector Machine to detect anomalies without requiring labeled attack samples—a useful property in the IoMT regime where attack labels are expensive to collect. Gupta *et al.* [82] construct a two-stage classical pipeline specific to IoMT anomaly detection, and Alalhareth and Hong [14] introduce LRGU, a mutual-information feature-selection scheme whose downstream accuracy gains are consistent enough to treat as a reproducible contribution rather than a dataset-specific fluke. Aljuhani *et al.* [16] propose an interpretable stacked ensemble tailored to IoMT; Alsalman [19] compares adaptive machine-learning techniques for IoT anomaly detection under evolving threat patterns; and Saleh *et al.* [173] propose SG-IDS, a hybrid statistical-graph detector for generic IoT traffic. Dadkhah *et al.* [54] release the CICIoMT-2024 benchmark used later in this thesis and report a strong Random Forest baseline on it; Salehpour *et al.* [175] and Doménech *et al.* [58] close out this group with, respectively, a cloud-oriented ensemble and a comparative study of classical classifiers on IoMT data.

Khan *et al.* [118] evaluate a Random Forest classifier on the CIC-IoT taxonomy (seven categories, thirty-three attack types) and report high top-line accuracies, and El-Sofany *et al.* [62] benchmark seven classical algorithms on the same taxonomy, confirming that tree-based ensembles do well when feature engineering is carefully tuned—a caveat that deserves more weight than these papers give it, since careful feature engineering is precisely what does not transfer across IoMT sites with heterogeneous device fleets. Kantharaju *et al.* [104] introduce SAPGAN, a self-attention pruning GAN used as an IDS on IoT traffic. Two studies, Nguyen *et al.* [147] and Alkadi *et al.* [17], directly tackle the adversarial-robustness problem flagged in Section 3.2.2: the first hardens a decision-tree ensemble against PGD perturbations, while the second proposes RobEns, an ensemble whose accuracy stays within two points of its clean-data counterpart under adversarial traffic—a result that is modest in absolute terms but encouraging for a constrained-compute setting.

Two broader limitations characterize the centralized-ML literature as a whole. First, a clear majority of the reviewed studies pool training data centrally, an architectural choice that is difficult to reconcile with the privacy constraints of the IoMT domain and which no amount of post-hoc tuning can repair. Second, hand-crafted features struggle to capture the long-range temporal dependencies that characterize multi-stage intrusion campaigns; this is the structural shortcoming that the deep-learning paradigm is best understood as responding to.

Deep Learning Architectures for Threat Detection

Deep-learning architectures—convolutional, recurrent, and attention-based—progressively displace hand-crafted features in favor of representations learned end-to-end. In [146] propose Swarm-NN, a swarm-optimized neural classifier for IoT anomaly detection; Ghourabi [77] couples LightGBM with a lightweight Transformer for IoMT traffic; Faruqui *et al.* [68] present SafetyMed, a hybrid CNN–LSTM IDS for medical telemetry. Alalhareth and Hong [13] complement their feature-selection line of work with a fuzzy-LSTM classifier tolerant to uncertain inputs. Khan *et al.* [115] distribute an LSTM-based IDS across a fog–cloud fabric, offering a partial treatment of the distributed dimension that is still relatively rare at the time. Subsequent contributions include a CNN–LSTM detector [21], the two-LSTM stack L2D2 [8], an interpretability-oriented variant [198], and a multi-block IoMT pipeline [34]—collectively illustrating that deep-IDS design has entered a consolidation phase in which architectural variety is outpacing the rate at which new datasets are released.

A second, more recent cluster pushes two directions in parallel: architectural hybridization and adversarial robustness. Sajid *et al.* [172] evaluate an XGBoost–CNN–LSTM cascade across four benchmarks, while Sinha *et al.* [187] report that an LSTM–CNN attains 99.87 % accuracy on the BoT-IoT benchmark under clean traffic but drops to 90.2 % under adversarial perturbations—a ten-point gap that is more informative than either number in isolation, because it shows how fragile state-of-the-art accuracy becomes once the test-time distribution shifts even moderately. Bommana *et al.* [38] address that vulnerability with a combination of Quantum-inspired Coyote Optimization (Q-COA), a Restricted Boltzmann Machine, and a Recurrent CNN. Bao *et al.* [30] provide a systematic baseline for this line of research by quantifying CNN vulnerability to adversarial examples. Mengara *et al.* [136] propose IoTSecUT, an uncertainty-aware framework combining a conditional GAN, an autoencoder, and a Transformer; Ain *et al.* [7] specialize a CNN–LSTM–autoencoder to DDoS detection on the CICIoT-2023 dataset; and Chemmakha *et al.* [51] employ a GAN to balance the UNSW-NB15 distribution before a GRU classifier and report 99.36 % accuracy.

Notwithstanding these performance gains, deep IDS trained on a single, centrally aggregated dataset inherit the same privacy and scalability shortcomings that constrained their classical predecessors. The third paradigm, federated learning, is motivated precisely by the need to address those two shortcomings simultaneously—though, as the next subsection argues, it introduces a new class of problems of its own.

Federated Learning Frameworks for Privacy-Preserving Detection

Federated learning (FL) [99, 100] keeps raw traffic on each participating node and exchanges only model updates with a central aggregator, so it directly answers the privacy critique of centralized deep IDS. A first wave of FL-based IDS established feasibility in IoT and IoMT settings. Singh *et al.* [185] propose a dew-cloud hierarchical federated LSTM for IoT, which, as Section 3.2.4 discusses, is among the closest architectural neighbors to the FTL-HLSTM of Chapter 4. Sarhan *et al.* [180] federate an LSTM across multiple NetFlow feeds; Zukaib *et al.* [218] introduce Meta-Fed IDS for IoMT; Misbah *et al.* [140] replace vanilla FedAvg with a dynamic-averaging rule tolerant to mild client drift; Ioannou *et al.* [95] present GEMLIDS-MIOT, an energy-pruned Random Forest coupled with a One-Class SVM; Khan *et al.* [116] extend this line with a reinforcement-learning layer in Fed-Inforce; and Kazmi *et al.* [113] provide a comparatively rare head-to-head evaluation of federated versus centralized detection on a common task—rare because very few papers in this space actually commit to that comparison.

A second wave of IoMT-focused, Q1-indexed work has consolidated the approach. In *et al.* [67] propose FedIoMT, which substitutes Kolmogorov–Arnold Networks for the usual local model; Begum *et al.* [31] pair federated learning with blockchain in BFLIDS to obtain a tamper-evident audit trail; Ghourabi [78] evaluates federated XGBoost on both CICIoMT-2024 and WUSTL-EHMS-2020, giving it useful cross-regime coverage; Bensaid *et al.* [35] introduce SA-FLIDS, a self-adaptive federated IDS; Lian *et al.* [130] propose a two-stage FL protocol with blockchain-mediated aggregation; and Tawfik *et al.* [191] combine federated detection with explainability in FedMedSecure.

A third group attacks the harder problem that sits underneath all of the above: the non-IID setting, widely acknowledged as the principal difficulty of federated learning in healthcare applications [79, 107]. Before listing these studies it is worth spelling out why IoMT data is almost never IID. Every participating client in IoMT federations is typically a hospital, a clinical department, or a remote-monitoring provider, and the traffic each one observes reflects a particular combination of patient demographics (age distribution, chronic-condition prevalence, acuity), device fleet (manufacturer, firmware, wireless stack), physical environment (ICU versus home-monitoring gateway), and local usage patterns. The consequence is that $p(x)$, $p(y)$, and $p(y \mid x)$ commonly differ between clients at the same time, yielding several textbook forms of statistical heterogeneity—feature skew, label skew, concept drift, and quantity skew—simultaneously [107, 168, 212]. A related and often under-appreciated variant is the *isolated-label* regime: many attack classes are observed by only a subset of clients, and

some clients see certain classes exclusively, which is more pathological for learning than any of the skews listed above.

Under these conditions the default Federated Averaging (FedAvg) algorithm—which tacitly assumes that local empirical risks are unbiased estimators of the global risk—degrades in ways now well documented. Local SGD steps drift towards client-specific optima rather than towards the global minimizer, a phenomenon formalized by Karimireddy *et al.* [107] as *client drift*; the global model under-performs on rare or client-local attack classes because averaging dilutes the gradient signal of the clients that hold them [193, 212]; and the privacy–accuracy trade-off is sharpened rather than softened by heterogeneity, since differentially-private aggregation rules have been reported to degrade more severely on heterogeneous clients than on IID ones, because additive noise is applied on top of already-biased updates [168]. The studies below attempt to mitigate these effects through hierarchical aggregation, knowledge distillation, personalization, or transfer learning.

Islam *et al.* [96] propose PP-HFFL, a privacy-preserving hierarchical federated-learning framework; Zhao *et al.* [213] apply knowledge distillation to realign heterogeneous client models; Peng *et al.* [155] introduce FD-IDS, a distillation-based intrusion detector; Soomro *et al.* [188] release SecureDyn-FL for dynamic client populations; Thein *et al.* [193] propose pFL-IDS, a personalized variant targeted at heterogeneous IoT; Han *et al.* [86] introduce CFMT, a cross-silo federated multi-task formulation; Zhang *et al.* [212] report one of the closest architectural neighbors to the FTL-HLSTM of Chapter 4, combining federated learning with transfer learning in an Industrial IoT setting; Ruzafa-Alcázar *et al.* [168] quantify the accuracy cost of the privacy budget under differential privacy; and Rahmati *et al.* [162] combine federated learning, a GRU classifier, and homomorphic encryption for additional confidentiality at substantial compute cost—the trade-off being exactly the sort that operational IoMT deployments will have to weigh explicitly.

The federated-learning literature taken as a whole suggests that privacy-preserving IoMT intrusion detection is now empirically feasible. But it also exposes two tensions that remain only partially resolved in the reviewed studies: the non-IID bottleneck, which can degrade aggregated model quality by amounts that matter clinically; and the limited coupling between federated detection and fault-tolerance at the network layer, which means that a correctly flagged compromise rarely translates into a routing response. Both tensions resurface, sharpened, in the gap analysis of Section 3.4.

3.2.4 Hierarchical Deep Learning Architectures for IoT Security

Hierarchical deep learning (HDL) exploits the multi-scale structure of network traffic—from individual packets to flows and complete sessions—by stacking models whose receptive fields grow progressively in both temporal span and semantic abstraction. This section first makes the conceptual case for the hierarchy, then reviews the reviewed empirical evidence, and finally identifies what is missing in the parts of that evidence most relevant to IoMT.

Why hierarchy is a natural fit for IoT/IoMT traffic. Network traffic exhibits structure at three scales at once: the packet level (header fields, payload patterns), the flow level (inter-arrival times, burstiness, session length), and the session level (multi-step behaviors typical of multi-stage attacks). A flat classifier collapses all three scales into a single feature vector, which makes it difficult for any single optimizer to balance them. A hierarchical architecture, in contrast, assigns distinct stages to distinct scales, so the statistics that live at each scale are processed at the appropriate granularity. For intrusion detection specifically, a two-stage decomposition—binary normal-versus-anomaly separation followed by fine-grained multi-class categorization—does double duty: it respects the scale hierarchy *and* it provides a computational filter, since the expensive multi-class stage only runs on traffic the first stage has already flagged. That computational argument is easy to overlook in the methodological literature but matters a great deal at IoMT deployment scale.

Reviewed empirical evidence. Diro and Chilamkurti [57] were among the first to show, in an IoT context, that distributed sub-models trained at fog nodes and aggregated at a master node can outperform a single flat classifier on the NSL-KDD dataset—a result whose main limitation today is the age of its benchmark rather than the architectural idea itself. Saba *et al.* [170] use a Convolutional Neural Network as a first-stage spatial extractor on an NID-style benchmark, while Alkahtani and Aldhyani [18] feed CNN-extracted spatial features into an LSTM (CNN–LSTM) for botnet detection on the N-BaIoT benchmark, which covers nine commercial IoT devices. Kim *et al.* [121] compare LSTM and GRU networks as single-stage predictors on a UNSW-NB15 subset; their results are consistent with the case for complementing a first stage with a second, finer-grained classifier, even though the paper itself does not frame the finding that way.

Subsequent work extends the hierarchical idea in related configurations. Kasongo [110] proposes a recurrent-neural-network framework (LSTM, GRU, and Simple-RNN vari-

ants) coupled with XGBoost-based feature selection and reports a best binary-classification test accuracy of 88.13 % on NSL-KDD (XGBoost–LSTM) and 87.07 % on UNSW-NB15 (XGBoost–Simple-RNN). Dong *et al.* [59] introduce MCA-LSTM, which couples information-gain feature selection with a Multi-correlation-Analysis Triangle-Area Map and a final LSTM stage, reaching 82.15 % on NSL-KDD in the five-class setting and 77.74 % on UNSW-NB15 in the ten-class setting. Kabir *et al.* [103] propose a stacking ensemble that combines an Extra Trees classifier with Mutual-Information-Gain feature selection and reports 96.24 % accuracy on UNSW-NB15. Finally, Singh *et al.* [185] transpose the hierarchical idea into a federated setting through a dew-cloud HLSTM. The evaluation in [185] is performed on a pre-IoMT dataset, which is what prevents it from closing the hierarchical-plus-federated-plus-IoMT loop that this thesis pursues.

What is missing. Two observations follow from the evidence reviewed. Hierarchical designs tend to outperform flat counterparts when the task is genuinely multi-class and the traffic is genuinely multi-scale—an alignment that is especially natural in IoMT. But among the reviewed studies, the only one that combines the hierarchical structure with federated aggregation [185] is evaluated on legacy IoT traffic rather than on modern IoMT-specific benchmarks such as CICIoMT-2024 or WUSTL-EHMS-2020. The gap is therefore narrow but specific: a hierarchical-plus-federated architecture, validated on current IoMT benchmarks, under the non-IID conditions that are the default in IoMT. This is the gap that motivates the FTL-HLSTM contribution of Chapter 4.

From detection to resilience. The literature reviewed so far treats IoT/IoMT security as a detection problem: each cited work ends once an intrusion is flagged. Operationally, that is not where the problem ends. A detected compromise is the beginning of a potential failure cascade at the network layer. A hospital gateway whose telemetry has been poisoned, a wearable whose firmware has been hijacked, or a federated client that has become a Byzantine participant is not simply a classifier event; from the network's point of view, it is a *faulty node*—one that can corrupt routing decisions, consume energy on its neighbors, delay time-critical medical traffic, and propagate its malfunction along the paths that still use it [33, 88]. The boundary between "security" and "fault tolerance" is therefore more porous than the reviewed IDS literature typically acknowledges: a sufficiently accurate intrusion detector is a necessary but not sufficient condition for network-level resilience, because the detection signal has to be translated into a routing response. That translation is the detection-to-response loop that most reviewed IDS papers leave unclosed, and that Chapter 5 aims to close through a proactive, trust-aware MCDA routing layer. The rest of the

present chapter accordingly shifts focus from detection to its downstream operational consequence—network resilience—and surveys the mechanisms by which distributed IoT and IoMT systems can keep operating correctly when a subset of their nodes is known, or suspected, to be compromised.

3.3 Fault Tolerance Mechanisms and Network Resilience in Distributed IoT Systems

A short taxonomic note helps read the literature that follows. *Fault tolerance* is the ability of a system to keep operating correctly in the presence of faults—benign or malicious—typically through redundancy or replication. *Resilience* is broader: it adds the ability to recover a degraded operating point after a disruption. *Byzantine tolerance* is the subclass of fault tolerance designed to withstand arbitrary (including adversarial) behavior by a bounded subset of participants, traditionally via consensus protocols. *Trust-aware routing*, finally, uses runtime trust scores to bias path selection away from nodes whose behavior appears suspicious—this is the category into which the routing layer of Chapter 5 falls.

Against that taxonomy, the reviewed fault-tolerance literature for distributed IoT splits into four complementary axes: replication and consensus protocols, trust-aware and QoS-aware routing, hardware-level resilience, and, more recently, Multi-Criteria Decision Analysis (MCDA) for proactive node selection. Each axis is reviewed in turn below, with an emphasis on whether the cited work couples its fault-tolerance mechanism with the security literature of Section 3.2—a coupling that turns out to be rarer than one might expect.

Consensus-based resilience. Zafar *et al.* [210] propose a practical Byzantine Fault Tolerance (pBFT) overlay on a lightweight blockchain for IoT, and Qi *et al.* [159] refine it into B-RBFT with reduced message complexity for resource-constrained devices. More recently, Beniwal *et al.* [33] introduce RB-BFT X for healthcare IoT, which is explicitly multi-dimensional—it simultaneously addresses fault tolerance, security, IoT, IoMT, and distributed deployment—and is, on those grounds, one of the most demanding comparators on the first five dimensions of the coverage analysis in Section 3.4.1. What it does not quantify is behavior under non-IID data, which leaves the sixth axis of the comparison open. More generally, heavy consensus mechanisms such as pBFT buy strong guarantees with a message complexity and latency that are hard to reconcile with the energy and bandwidth budgets of IoMT endpoints—this is the practical reason the thesis pursues a lighter-weight MCDA-based alternative in Chapter 5 rather

than a consensus-based one.

Quality-of-Service and trust-aware routing. Reyana *et al.* [164] propose QoS-EO, a quality-of-service-driven evolutionary optimizer for IoT routing under node failures. Chanak and Samanta [47] propose an intelligent fault-tolerant routing scheme for IoT-enabled wireless sensor networks that reuses partially faulty nodes to tolerate failures without additional hardware overhead—a pragmatic design choice that prefigures the thesis's own reluctance to add dedicated fault-tolerance hardware to already-constrained IoMT endpoints. Khaleel [114] proposes Bi-OWSP, a bidirectional overlay augmented with waypoint protection. Haseeb *et al.* [88] present LSDAR, a lightweight structure-based data-aware routing scheme for trust-based IoT networks, which is the closest architectural neighbor to the routing layer developed in Chapter 5. Agarwal *et al.* [6] extend the line further by applying deep reinforcement learning to fault-tolerant routing, with experimental validation on topologies of several hundred nodes. A common thread across all of these schemes is that they remain predominantly *reactive*: a node's trust score or routing weight is adjusted after a failure or suspicious behavior has already been observed. For time-critical IoMT settings, where detection after a clinical consequence is itself a failure mode, that is an uncomfortable default.

Hardware-level resilience. For completeness, Ferreira *et al.* [69] present ReViTA, a resilience layer for the Internet-of-Vehicles overlay, whereas Joardar *et al.* [101] target ReRAM-based fault tolerance at the accelerator level. These contributions establish the lower bound of the resilience stack against which higher-level routing decisions are taken; they are included here for context rather than as direct comparators for the contribution chapters.

MCDA-based proactive routing. More recently, Yu *et al.* [209] report a TOPSIS-based fault-tolerant Network-on-Chip routing scheme. The importance of this reference for the thesis is not the NoC application itself but the validation it provides for the MCDA primitive adopted in Chapter 5—TOPSIS with AHP-weighted criteria, augmented by a dynamic trust score—in an adjacent domain. Within the reviewed sample, we did not identify a prior work that couples this primitive in a closed loop with a federated hierarchical-LSTM intrusion detector in the IoMT setting.

A broader pattern is visible across this body of literature: the fault-tolerance and security communities remain largely disjoint within the reviewed sample. Few works address both concerns simultaneously, and among those that do, none in the reviewed sample jointly addresses the non-IID dimension characteristic of federated IoMT deployments. Section 3.4.1 quantifies this observation and it is one of the principal

motivations of the thesis.

3.4 Research Gaps and Motivation

Before identifying the gaps, it is worth briefly naming what the reviewed literature does well. The AI-IDS community has produced mature, accurate detectors for both IoT and IoMT; the fault-tolerance community has produced principled consensus and routing mechanisms that hold up under adversarial assumptions; and the federated-learning community has demonstrated that privacy-preserving intrusion detection is now empirically feasible in clinically realistic settings. Each of these contributions is substantial, and the present thesis takes them as a starting point rather than as targets for dismissal.

Where the reviewed literature is collectively thinner is at the intersection of those three strands. The core trade-offs—between detection accuracy, privacy preservation, fault tolerance, scalability, and the management of heterogeneous client data—are tightly coupled, but the reviewed work tends to address them in isolation. Centralized ML approaches typically reach high detection accuracy but compromise privacy; federated learning preserves privacy but pays a well-documented performance penalty under statistical heterogeneity; consensus-based fault tolerance guarantees correctness but at a message-complexity and energy cost that IoMT endpoints cannot comfortably absorb. Six research gaps emerge from this review, of which one is both broader in scope and structurally harder than the others. For that reason, this thesis treats it as the *principal* gap that drives its design choices: the robust handling of strongly non-IID client data, and specifically *extreme label skew*.

Why non-IID data with strong label skew is the principal open problem. Across recent Q1-journal surveys and empirical assessments, statistical heterogeneity—and label-distribution skew in particular—is consistently identified as the single most corrosive obstacle to federated and distributed learning. Ye *et al.* [208], in one of the most cited recent surveys of heterogeneous FL, place statistical heterogeneity first among five heterogeneity categories and report that the bulk of open research challenges in the field trace back to it. Zhu *et al.* [214] likewise show that models trained under non-IID clients systematically underperform their centralized counterparts, with the performance gap widening as the divergence between client label distributions grows. Karimireddy *et al.* [107] give this degradation its now-standard formal name, *client drift*: FedAvg converges to biased stationary points whenever local empirical risks are not unbiased estimators of the global risk, a condition that is systematically violated under label skew. Jimenez-Gutierrez *et al.* [83] quantify the effect across four types of

non-IIDness (label, feature, quantity, spatiotemporal) and conclude that *label* and *spatiotemporal* skew are the two regimes in which FL performance collapses most sharply, with identifiable Hellinger-distance thresholds beyond which accuracy drops precipitously. Li *et al.* [128], in the IIoT context closest to this thesis, likewise find that local class-distribution skew is the dominant driver of federated accuracy degradation, and Babar *et al.* [28] confirm the same pattern in federated medical imaging. Five independent lines of evidence reaching the same diagnosis is unusual in this literature; it is why the present thesis treats the problem as the principal one rather than as one concern among many.

Why the IoMT setting makes this problem harder still. In IoMT the label-skew problem takes an especially severe form that the generic FL literature only partially addresses. Every participating client is typically a hospital, a clinical department, or a home-monitoring gateway, and the attack traffic each one observes reflects its own device fleet, patient demographics, and local threat exposure. As a direct consequence, many attack classes are observed by only a subset of clients, and some clients see certain classes *exclusively*—what Thein *et al.* [193] and Zhang *et al.* [212] highlight as the *isolated-label* regime. Under that regime, naive FedAvg aggregation dilutes the gradient signal of the clients that hold rare classes, so rare-attack recall collapses even when global accuracy looks acceptable. Privacy sharpens the trade-off rather than smoothing it: Ruzafa-Alcázar *et al.* [168] show that differential-privacy noise degrades more severely on heterogeneous clients than on IID ones, since additive noise is applied on top of already-biased updates. Li *et al.* [127] report a parallel effect for Byzantine-robust aggregation rules: several of them, effective under IID assumptions, fail outright in the non-IID regime. Fotohi *et al.* [72] reach the same conclusion for blockchain-mediated FL, where performance on label-skewed clients degrades sharply even when security guarantees are preserved. The convergent diagnosis is that non-IID robustness under strong label skew is not one engineering concern among many; it appears to be the defining bottleneck of privacy-preserving IoMT intrusion detection.

Why this gap persists in the reviewed IDS sample. Despite how prominent this problem is in the surveys above, the coverage analysis of the reviewed AI-IDS studies in Table 3.2 shows that only a small minority of IDS-focused papers address non-IID data with an explicit experimental protocol, and almost none address it *jointly* with fault tolerance, scalability, and IoMT-specific validation. Three structural reasons explain the asymmetry. First, the dominant IDS benchmarks (NSL-KDD, UNSW-NB15, CICIDS 2017) were released as centralized datasets and have no standard federated split that actually exposes label skew. Second, the prevailing evaluation protocol in the

federated-IDS literature uses 3–8 synthetically-partitioned clients with mild Dirichlet skew ($\alpha \approx 0.5$ or higher), which does not reproduce the extreme skew observed in real IoMT fleets. Third, the algorithmic responses that do address label skew—hierarchical aggregation, personalization, knowledge distillation, clustered FL, transfer learning—have been studied largely in isolation from fault-tolerance and routing concerns. The combined effect is that the hardest version of the problem, the one that a hospital IoMT deployment would actually face, is systematically under-tested in the reviewed IDS sample. That is not a criticism of any individual paper; it is a coordination problem of the subfield.

Building on this diagnosis, the six research gaps that emerge from the reviewed sample are now listed below, with the non-IID / label-skew gap first and treated as the principal one that the contributions of this thesis are designed to address:

1. **Non-IID data with strong label skew (principal gap).** The reviewed federated-IDS studies predominantly evaluate under mild Dirichlet skew or a small number of clients, and rarely report dedicated metrics under the extreme label-skew and isolated-label regimes that characterize real IoMT deployments [83, 107, 128, 208, 214].

2. **Scalability to realistic client populations.** Many federated frameworks in the reviewed sample are evaluated with three to eight clients. That does not reflect the hospital-fleet scale at which an IoMT IDS would operate, and it tends to mask the non-IID bottleneck identified above—which is a specific reason to treat the two gaps as intertwined rather than independent.

3. **Privacy–performance trade-off under heterogeneity.** FL preserves privacy but is reported to incur a detection-accuracy cost relative to centralized baselines, and this cost is amplified rather than absorbed under heterogeneous clients [168, 214].

4. **Integrated security–resilience solutions.** Few of the reviewed approaches integrate fault tolerance, security, and scalable communication within a single framework that closes the detection-to-response loop.

5. **Hierarchical detection.** A majority of the reviewed IDS still use single-stage classification, which tends to underperform on fine-grained multi-class attack taxonomies and does not exploit the multi-scale structure of IoMT traffic.

6. **Proactive fault management.** Most of the reviewed fault-tolerance mechanisms are reactive: they respond to failures rather than anticipate them, so a compromised node continues to receive traffic until a downstream symptom is observed.

Positioning of this thesis. In response to these gaps—and, first and foremost, to the non-IID label-skew bottleneck—this thesis proposes an integrated framework that combines federated learning, transfer learning, and a hierarchical LSTM topology (FTL-HLSTM, Chapter 4), coupled in a closed loop with an MCDA-based proactive routing layer (Chapter 5). The central design hypothesis is that the combination of *(i)* a hierarchical two-stage topology, which decouples localized feature extraction from global semantic classification, and *(ii)* a transfer-learning head, which re-projects the aggregated representation onto each client's local label space, is specifically suited to absorbing strong label skew at the feature level, before it can bias the global decision boundary. Framed this way, the contribution is positioned where the reviewed literature is thinnest: at the intersection of strong-label-skew non-IID robustness, IoMT-specific benchmarking, scalable federated evaluation, and proactive fault-tolerant routing. Whether the resulting framework actually delivers on this hypothesis is precisely the empirical question that Chapters 4 and 5 are written to answer—not to pre-empt.

The literature is evaluated along seven dimensions (FT, Sec, IoT, IoMT, Dist, Scal, and N-IID), from which six principal research gaps are derived. The three subsections that follow substantiate this overview: Section 3.4.1 reports the seven-dimension coverage table that evidences the gaps; Section 3.4.2 examines the benchmark-dataset landscape; and Section 3.4.3 describes, gap by gap, the mechanisms by which the contribution chapters are *intended* to respond.

3.4.1 Unified Coverage Analysis of the State of the Art

In order to aggregate the material surveyed in Sections 3.2–3.3 into an actionable gap analysis, the sixty-three retained AI-centered studies, together with the present thesis (line [T]), are evaluated along seven dimensions that jointly characterize the scalability–security–fault-tolerance trilemma at the core of this work: **FT** (Fault Tolerance), **Sec** (Security, understood as intrusion detection or authentication), **IoT** (generic IoT applicability), **IoMT** (IoMT-specific validation), **Dist** (distributed training or deployment), **Scal** (scalability to realistic client populations rather than three-to-eight-client prototypes), and **N-IID** (explicit handling of heterogeneous client data distributions). Table 3.2 reports, for each study, whether a given dimension is fully addressed (✓), partially addressed (∼), or not addressed (−).

Classification criteria. For transparency and reproducibility, the rules adopted to assign the three symbols are made explicit below. A dimension is scored ✓ (*fully addressed*) when it is explicitly designed into the method *and* evaluated empirically; ∼ (*partially addressed*) when the paper either discusses the dimension without evaluating

it, evaluates it only on a proxy benchmark, or reports results that cover the dimension only partially; and – (*not addressed*) when the paper neither designs nor evaluates the dimension. Concretely: (a) **Dist** is scored ✓ when the training or inference pipeline is operated across multiple physical or logical nodes, ∼ when a fog–cloud or edge overlay is described but not empirically evaluated in a distributed regime, and – for purely centralized pipelines. (b) **Scal** is scored ✓ when the reported evaluation exceeds eight clients or equivalent computational units, ∼ when evaluation remains within the three-to-eight-client range frequently used in the federated-IDS literature, and – when scalability is not discussed. (c) **N-IID** is scored ✓ when the paper explicitly simulates non-IID label or feature skew and reports dedicated metrics, ∼ when mild heterogeneity is mentioned without targeted evaluation, and – when the IID assumption is retained. Although these rules are applied systematically, some categorizations inevitably involve interpretive judgment, especially when studies report partial or indirect treatment of a given dimension; the table should therefore be read as an analytical framework rather than as an absolute classification.

Table 3.2: Unified coverage analysis of the sixty-three reviewed AI-centered studies, together with the positioning of this thesis ([T]), across the seven dimensions of the scalability–security–fault-tolerance trilemma.

Reference	**Year**	**FT**	**Sec**	**IoT**	**IoMT**	**Dist**	**Scal**	**N-IID**
[37]	2022	–	✓	✓	–	–	–	–
[82]	2022	–	✓	✓	✓	–	–	–
[14]	2023	–	✓	✓	✓	–	–	–
[16]	2023	–	✓	✓	✓	–	–	–
[19]	2024	–	✓	✓	✓	–	–	–
[173]	2023	–	✓	✓	–	–	–	–
[54]	2024	–	✓	✓	✓	–	–	–
[175]	2024	–	✓	✓	✓	∼	∼	–
[58]	2024	–	✓	✓	✓	–	–	–
[146]	2022	–	✓	✓	–	–	–	–
[77]	2023	–	✓	✓	✓	–	–	–
[68]	2023	–	✓	✓	✓	–	–	–
[13]	2023	–	✓	✓	✓	–	–	–
[115]	2023	∼	✓	✓	✓	✓	∼	–

continued on next page

Table 3.2 — continued from previous page

Reference	Year	FT	Sec	IoT	IoMT	Dist	Scal	N-IID
[21]	2023	–	✓	✓	–	–	–	–
[8]	2024	–	✓	✓	✓	–	–	–
[198]	2023	–	✓	✓	✓	–	–	–
[34]	2024	–	✓	✓	✓	–	–	–
[185]	2023	∼	✓	✓	✓	✓	✓	–
[218]	2024	–	✓	✓	✓	✓	–	∼
[180]	2023	–	✓	✓	–	✓	∼	∼
[140]	2025	–	✓	✓	✓	✓	∼	∼
[95]	2024	∼	✓	✓	✓	✓	∼	–
[116]	2024	–	✓	✓	✓	✓	–	–
[113]	2024	–	✓	✓	✓	✓	–	✓
[164]	2023	✓	∼	✓	–	✓	✓	–
[210]	2023	✓	✓	✓	–	✓	✓	–
[159]	2023	✓	✓	✓	–	✓	✓	–
[47]	2021	✓	–	✓	–	✓	✓	–
[114]	2023	✓	–	✓	–	✓	✓	–
[69]	2023	✓	–	∼	–	∼	∼	–
[101]	2023	✓	–	∼	–	–	–	–
[67]	2025	–	✓	✓	✓	✓	∼	∼
[31]	2024	–	✓	✓	✓	✓	∼	–
[78]	2025	–	✓	✓	✓	✓	∼	–
[35]	2024	∼	✓	✓	✓	✓	∼	–
[130]	2023	–	✓	✓	✓	✓	∼	✓
[96]	2025	–	✓	✓	–	✓	✓	✓
[213]	2023	–	✓	✓	–	✓	∼	✓
[155]	2025	–	✓	✓	–	✓	∼	✓
[191]	2025	–	✓	✓	✓	✓	∼	–
[188]	2026	–	✓	✓	–	✓	✓	✓

continued on next page

Table 3.2 — continued from previous page

Reference	Year	FT	Sec	IoT	IoMT	Dist	Scal	N-IID
[193]	2023	–	✓	✓	–	✓	∼	✓
[86]	2024	–	✓	✓	–	✓	✓	✓
[212]	2023	–	✓	✓	–	✓	∼	✓
[168]	2023	–	✓	✓	–	✓	∼	✓
[88]	2022	✓	✓	✓	–	✓	✓	–
[6]	2022	✓	∼	✓	–	✓	✓	–
[209]	2025	✓	–	∼	–	✓	✓	–
[33]	2025	✓	✓	✓	✓	✓	✓	–
[118]	2024	–	✓	✓	∼	–	–	–
[62]	2024	–	✓	✓	–	–	–	–
[104]	2024	–	✓	✓	–	–	–	–
[147]	2023	–	✓	✓	–	–	–	–
[17]	2024	–	✓	✓	–	–	–	–
[172]	2024	–	✓	✓	–	–	–	–
[187]	2025	–	✓	✓	–	–	–	–
[38]	2025	–	✓	✓	–	–	–	–
[30]	2022	–	✓	✓	–	–	–	–
[136]	2024	–	✓	✓	–	–	–	–
[7]	2025	–	✓	✓	–	–	–	–
[51]	2024	–	✓	✓	–	–	–	–
[162]	2025	–	✓	✓	–	✓	∼	–
[T] (FTL-HLSTM + MCDA)	**2026**	✓	✓	✓	✓	✓	✓	✓

Aggregating Table 3.2 column by column yields a synthetic picture of the reviewed sample. The *Security* and *IoT* dimensions are, unsurprisingly given the scope of the review, nearly saturated: 56 of 63 studies (89 %) fully address the security dimension, and 60 of 63 (95 %) fully address the IoT dimension. The *IoMT validation* and *distributed* dimensions reach moderate coverage, at 27/63 (43 %) and 32/63 (51 %) respectively—IoMT-specific benchmarking and distributed deployment, though increasingly common, are not yet systematic practices in the reviewed sample. *Scalability* follows a

related pattern: only 13 of 63 studies (21 %) fully scale their evaluation to realistic client populations, and a further 18 (29 %) report only partial scalability evidence, typically bounded by the three-to-eight-client protocol that dominates the reviewed federated-IDS literature. *Fault Tolerance* and *Non-IID* are, within the reviewed sample, the two most conspicuous blind spots, fully addressed by only 11 studies (17 %) and 10 studies (16 %) respectively.

Among the reviewed studies, we did not identify a paper that simultaneously addresses all seven dimensions. The closest comparator is Beniwal *et al.* [33], which fully addresses six of the seven dimensions but leaves the non-IID setting unaddressed; other strong candidates—Lian *et al.* [130], Haseeb *et al.* [88]—fully address five of the seven. The picture is therefore not that the literature is weak everywhere, but rather that the specific combination of fault tolerance *plus* non-IID robustness *plus* IoMT-specific validation has not, to the best of this review, been covered together in a single reviewed study. Within that analytical framework, the thesis contribution [T] is positioned as *targeting* that combination; whether the targeted coverage is empirically achieved is the subject of Chapters 4 and 5.

3.4.2 Dataset Landscape and Rationale for Dataset Selection

Reproducibility depends, at first approximation, on the benchmarks on which the reviewed studies are evaluated. Table 3.3 contrasts the benchmarks most commonly used in the recent literature along four practical criteria: IoT applicability, IoMT specificity, recency (post-2020), and the richness of the attack taxonomy.

Dataset-selection criteria of this thesis. The three benchmarks retained for the empirical validation of Chapter 4 were selected according to five explicit criteria: (i) *recency*—publication year at or after 2020, so that captured traffic reflects current IoT and IoMT protocol stacks; (ii) *public accessibility*, so that the evaluation can be reproduced independently; (iii) *reproducibility of features*, i.e. documented feature sets rather than raw undocumented traces; (iv) *traffic diversity*, i.e. non-trivial attack taxonomies covering several categories of threat; and (v) *complementarity of regimes*, so that each retained dataset exercises a distinct traffic abstraction (multi-protocol IoMT, NetFlow v2, WBAN telemetry). The criteria are applied uniformly; the retained benchmarks are highlighted in bold in Table 3.3.

Three observations from Table 3.3 motivate the three-dataset strategy of Chapter 4. First, legacy datasets (KDD-99, NSL-KDD, UNSW-NB15) remain dominant in the reviewed literature despite their well-documented inability to capture modern IoMT protocols such as MQTT, CoAP, and Bluetooth LE—a persistence that is eas-

Table 3.3: Landscape of benchmark datasets used across the reviewed IoT and IoMT security literature

Dataset	Year	Features	Rows	Labels	IoT	IoMT	Thesis
KDD Cup 99	1999	41	4,898,431	5	–	–	–
NSL-KDD [190]	2009	41	148,517	5	–	–	–
UNSW-NB15 [143]	2015	49	2,540,044	10	–	–	–
Bot-IoT	2018	42	>72M	5	✓	–	–
N-BaIoT	2018	115	7,062,606	11	✓	–	–
TON-IoT	2020	43	21,978,630	10	✓	–	–
WUSTL-EHMS-2020 [84]	**2020**	**43**	**16,318**	**3**	✓	✓	✓
NF-UNSW-NB15-v2 [179]	**2022**	**43**	**2,390,275**	**10**	–	–	✓
NF-BoT-IoT-v2	2022	43	37,763,497	5	✓	–	–
Edge-IIoTSet	2022	61	2,219,201	15	✓	–	–
RT-IoT 2022	2022	83	123,117	12	✓	–	–
CICIoT-2023	2023	46	46,686,579	34	✓	–	–
CICIoMT-2024 [54]	**2024**	**39**	**8,775,013**	**19**	✓	✓	✓

ier to explain by path dependence than by scientific preference. Second, among the reviewed benchmarks, only two recent public IoMT-specific datasets clearly satisfy the recency and reproducibility criteria adopted in this thesis: CICIoMT-2024 [54] and WUSTL-EHMS-2020 [84]. This scarcity is itself one of the contributors to the limited diversity of reproducible IoMT benchmarking. Third, NF-UNSW-NB15-v2 [179] is retained in this thesis not as an IoMT dataset—which it is not—but as a complementary NetFlow-oriented regime that exercises the detector on a distinct traffic abstraction. On those three grounds, the empirical evaluation in Chapter 4 adopts a coordinated three-dataset protocol consisting of CICIoMT-2024 (multi-protocol IoMT), NF-UNSW-NB15-v2 (NetFlow v2), and WUSTL-EHMS-2020 (WBAN telemetry), with the goal of a coverage breadth that, within the reviewed sample, appears to remain uncommon.

3.4.3 Alignment with Thesis Contributions

The following paragraphs describe, gap by gap, the architectural mechanism by which the contribution chapters are *intended* to respond. The emphasis is on *how*, not on *what*; the language is deliberately future-facing, since the empirical evidence that the mechanisms achieve their intended effect is developed in Chapters 4 and 5, not here. Reviewers should therefore read this subsection as a design rationale, not as an outcome report.

Gap G1 — Non-IID data with strong label skew. Chapter 4 proposes FTL-HLSTM, which couples transfer learning with a hierarchical LSTM topology. The hierarchical topology is intended to decouple localized feature extraction—performed

by the lower stage on each client's own traffic distribution—from global semantic classification, performed by the upper stage on aggregated representations. The design intent is that client drift is absorbed at the feature level before it can bias the global decision boundary. The transfer-learning head is intended to re-project the aggregated representation onto each client's local label space, which is the mechanism by which the architecture is hypothesized to handle the isolated-label regime identified in Sections 3.2.3 and 3.4. Whether the mechanism actually absorbs strong label skew in practice is an empirical question answered in Chapter 4.

Gap G2 — Scalability to realistic client populations. Chapter 4 evaluates the framework under federation sizes larger than those commonly reported in the reviewed federated-IDS literature, and it uses an asynchronous hierarchical aggregation rule that is designed to decouple the cost of global aggregation from the number of local workers. The scalability claim is *evaluated* rather than *asserted*: the measurements—training time, convergence rounds, per-client cost—are the content of Chapter 4, not of this review.

Gap G3 — Privacy–performance trade-off. The transfer-learning component of FTL-HLSTM is designed to re-inject task-specific knowledge into the locally personalized model without exposing raw client data to the aggregator. The intent is to mitigate the accuracy penalty commonly reported for federated learning on heterogeneous clients [113, 168] by paying that penalty at personalization time rather than at inference time. Whether the trade-off shifts favorably—and by how much—is quantified in Chapter 4.

Gap G4 — Integrated security and resilience. Chapters 4 and 5 are *jointly* intended to respond along all seven dimensions of Table 3.2. The integration mechanism is a closed-loop coupling between the two chapters: detection insights produced by FTL-HLSTM (Chapter 4) are streamed to the MCDA routing layer (Chapter 5) as real-time trust-score updates, while the routing layer in turn supplies FTL-HLSTM with topologically filtered traffic that excludes nodes already flagged as untrustworthy. The argument is that neither chapter on its own would close the detection-to-response loop identified in Section 3.2.4; the coupling itself is the contribution. Whether the loop delivers measurable improvements is evaluated in the contribution chapters.

Gap G5 — Hierarchical versus single-stage detection. The FTL-HLSTM architecture is assessed as a two-stage alternative to single-stage classification, with the first LSTM stage performing binary normal-versus-anomaly separation and the sec-

ond stage performing fine-grained multi-class attack categorization. The two-stage decomposition is evaluated empirically on CICIoMT-2024 and NF-UNSW-NB15-v2 in Chapter 4. A negative result on this gap—i.e., no meaningful improvement over a flat counterpart—remains a possibility and would itself be informative.

Gap G6 — Proactive fault management. Chapter 5 proposes a proactive routing layer in which trust-aware scoring, continuously fed by the Chapter 4 detector and combined with TOPSIS and AHP, is used to anticipate and avoid degrading or suspicious nodes. The intended behavior is that malicious or failing nodes are de-ranked and progressively isolated before a service-level disruption occurs, which would contrast with the predominantly reactive schemes reviewed in Section 3.3. Whether the observed behavior is in fact proactive (as opposed to fast-reactive, which is a different claim) is the subject of Chapter 5's evaluation.

3.5 Conclusion

The review developed in this chapter has covered three tightly connected axes: traditional security mechanisms and their limits in constrained networks (Section 3.2); hierarchical deep-learning architectures for intrusion detection (Section 3.2.4); and fault-tolerance mechanisms for distributed IoT systems (Section 3.3). Sixty-three AI-centered studies and seventeen traditional-security studies were consolidated into a seven-dimension coverage table (Table 3.2), a benchmark-dataset landscape (Table 3.3), and a gap-to-mechanism mapping (Section 3.4.3).

What this synthesis suggests, rather than proves, can be stated compactly. Three principal observations emerge from the review, each of which shapes the remainder of the thesis. *First*, within the reviewed sample, AI-based intrusion detection has matured along the Security and IoT axes but remains thin along Fault Tolerance and non-IID robustness—each fully covered by only 16% of the reviewed studies (Table 3.2). *Second*, IoMT-specific experimental validation rests on a narrower benchmark base than the field would ideally have: CICIoMT-2024 and WUSTL-EHMS-2020 are the two principal IoMT-specific options retained by this thesis, complemented by NF-UNSW-NB15-v2 as a NetFlow-oriented regime to broaden coverage. *Third*, and most consequentially for the positioning of this work, the joint treatment of scalability, privacy-preserving detection, non-IID robustness, and proactive fault tolerance remains, to the best of this review, unaddressed in any single prior work in the reviewed sample.

That diagnosis is the point of departure for the rest of the thesis, not its conclusion. Chapter 4 develops and evaluates the FTL-HLSTM framework; Chapter 5 develops

the proactive, trust-aware MCDA routing layer that closes the detection-to-response loop. Whether the two chapters succeed in transforming the gaps identified here into working mechanisms—and by what quantitative margin—is the question Chapters 4 and 5 are written to answer. Readers should judge the value of the framework against what those chapters actually measure, not against the expectations this review chapter might otherwise set.

Part II

Original Contributions

Chapters 4–5 present the novel algorithms and frameworks
developed in this thesis for scalable, secure, and fault-tolerant IoT.

Chapter

4

Federated Intrusion Detection

Chapter 4

FTL-HLSTM: Federated Transfer Learning for Privacy-Preserving Intrusion Detection in IoMT Networks

4.1 Introduction

The Internet of Medical Things (IoMT) has emerged as a cornerstone of contemporary healthcare infrastructures, enabling continuous physiological monitoring, real-time clinical decision support, and the coordination of therapeutic actions across interconnected edge–cloud tiers [91]. As established in Chapter 1.1, the operational profile of IoMT networks is distinctive in two respects: the underlying devices are severely constrained in computation, memory, and energy, and the data they generate are among the most privacy-sensitive encountered in any Internet of Things (IoT) vertical. The critical review conducted in Chapter 3 further demonstrated that traditional centralised perimeter defences, originally designed for enterprise networks, scale poorly to such heterogeneous and decentralised deployments and leave a substantial residual attack surface [74].

Against this backdrop, Machine Learning (ML) and Deep Learning (DL) have become the dominant paradigms for intrusion detection in IoT and IoMT. Their effectiveness, however, is tightly coupled to the availability of large, representative training corpora, which in conventional settings presupposes the centralised aggregation of raw traffic. In the IoMT context, such aggregation conflicts directly with the confidentiality obligations imposed by the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR), both of which mandate the protection of Personal Health Information (PHI) [66, 201]. Federated Learning (FL) has consequently been advanced in the recent literature as a principled response to this tension, since it allows multiple sites to contribute to a common model without ever exposing their local data.

The present chapter argues, however, that standard FL constitutes a necessary but not sufficient answer to the requirements identified in the preceding chapters of this

thesis. Three structural limitations, each documented in the state-of-the-art analysis of Chapter 3.1, continue to restrict its applicability to realistic IoMT deployments. First, clinical sites inherit divergent case mixes, device populations, and attack exposures, so that their local data distributions are markedly non-independent and non-identically distributed (non-IID); under label skew and feature shift, standard Federated Averaging (FedAvg) converges slowly, if at all, and generalises poorly [106]. Second, a substantial proportion of the published literature continues to rely on benchmark datasets that predate current IoMT traffic patterns or on idealised experimental assumptions that understate the true resource envelope of edge nodes [80]. Third, and central to the overall argument of this thesis, existing frameworks seldom treat *time as a strategic resource*; yet, as emphasised in the temporal-aggregation analysis of Chapter 2, the way in which raw streams are windowed and summarised prior to transmission is a decisive determinant of communication overhead, energy consumption, and, ultimately, scalability.

To address these three limitations in a unified manner, this chapter introduces the **Federated Transfer Learning framework with Hierarchical Long Short-Term Memory (FTL-HLSTM)**, which constitutes the first major original contribution of the thesis. FTL-HLSTM is designed as a privacy-preserving intrusion detection architecture specifically tailored to resource-constrained IoMT environments, and combines three complementary ingredients: (i) a temporal-aggregation front end that converts raw traffic into compact feature vectors, (ii) a hierarchical LSTM backbone that exploits the multi-scale structure of IoMT sessions, and (iii) an intelligent label-classification mechanism that routes globally shared attack patterns through federated aggregation while delegating locally specific patterns to a dedicated transfer-learning pathway.

The remainder of the chapter is organised as follows. Section 4.2 presents the system architecture and server- and client-side methodologies of FTL-HLSTM, including the Intelligent Label Classification algorithm and the weighted multi-criteria model-selection procedure. Section 4.3 reports the experimental protocol and the empirical evaluation on the NF-UNSW-NB15-v2, CICIoMT-2024, and WUSTL-EHMS-2020 benchmarks, covering centralised, federated, and severely non-IID regimes, as well as a scalability analysis over federation sizes ranging from four to one hundred clients. Section 4.4 discusses the findings in relation to the research objectives defined in the General Introduction and contrasts FTL-HLSTM with contemporary state-of-the-art intrusion detection systems. Section 4.5 summarises the contributions of the chapter and motivates the trust-aware routing layer developed in Chapter 5, which closes the cyber-physical feedback loop between detection intelligence and network control.

4.2 Methodology

This chapter develops the *Federated Transfer-Learning Hierarchical Long Short-Term Memory* (FTL-HLSTM) framework, designed for privacy-preserving intrusion detection in *Internet of Medical Things* (IoMT) environments. The primary objective is to address the statistical heterogeneity challenges that arise from diverse network configurations, heterogeneous IoMT device specifications (non-independent and identically distributed (non-IID) data), and distinct intrusion patterns. To mitigate these heterogeneities, FTL-HLSTM integrates *Federated Learning* (FL) for the collaborative detection of common intrusion patterns across multiple clients, together with client-specific *Transfer Learning* (TL) for handling isolated intrusion labels unique to individual clients.

FTL-HLSTM adopts a decentralised training strategy in which edge clients independently analyse local IoMT traffic data. Clients communicate exclusively model parameters and non-sensitive label metadata to a central aggregator, ensuring strict adherence to healthcare data privacy regulations. Sensitive patient and institutional data remain securely stored within local environments. An architectural overview of the proposed framework is illustrated in Fig. 4.1.

4.2.1 System Architecture

The FTL-HLSTM framework consists of three primary layers:

- **IoMT Device Layer:** This foundational layer comprises diverse healthcare monitoring devices, such as wearable sensors and bedside medical equipment. These devices continuously generate multivariate network and telemetry data streams, forming the primary input for the intrusion detection system (IDS).
- **Edge Client Layer:** This intermediate layer consists of clinical institutions, including hospitals and specialised medical departments. Each institution operates dedicated edge computing nodes responsible for local data preprocessing, temporal aggregation of data, and localised HLSTM model training.
- **Central Cloud Aggregator:** At the top tier, the central cloud aggregator coordinates FL for common intrusion labels and distributes both the global federated model and specialised TL models for isolated intrusion labels.

The HLSTM architecture is selected for its hierarchical structure, explicitly designed to leverage client-level LSTM models. This design maintains the predictive

performance of traditional LSTM architectures while significantly reducing training and inference times. HLSTM is particularly effective in modelling the sequential and temporal characteristics of IoMT-based cyberattacks, which often involve multi-stage processes such as reconnaissance, exploitation, and exfiltration. Furthermore, it effectively supports continuous clinical monitoring scenarios, where anomalies gradually develop over extended periods.

The operational workflow of the FTL-HLSTM framework comprises three distinct phases, as shown in Fig. 4.2.

- **Phase 1: Intelligent Label Classification:** Initially, each client transmits a metadata vector indicating the presence of labels to the central aggregator. The server employs Algorithm 1 to categorise these labels into *common labels*—present across multiple clients—and *isolated labels*—unique to individual clients. Direct integration of isolated labels into FL can negatively affect aggregation efficiency (e.g., in Federated Averaging (FedAvg)) due to their exclusivity. This classification enables targeted learning strategies without compromising data privacy, as raw IoMT data are never transmitted.
- **Phase 2: Hybrid Learning Execution:** For common labels, clients independently train local models using relevant data subsets. The central aggregator consolidates these updates via FedAvg to produce a unified global detection model. For isolated labels, clients clone their local HLSTM models, modify the final classification layer for binary classification specific to the isolated label, and fine-tune the model locally. The specialised TL models are then sent to the central aggregator for distribution alongside the global federated model.
- **Phase 3: Optimised Deployment:** Clients use Algorithm 3 to select the optimal detection model—either the global federated model or a client-specific TL model—for each intrusion label. This selection process applies a multi-criteria evaluation considering detection accuracy, inference latency, and computational efficiency, ensuring efficient and context-specific deployment.

Throughout the operational workflow, the FTL-HLSTM framework maintains strict privacy preservation by restricting communication to secure transmissions of model parameters and non-sensitive metadata. Sensitive IoMT and patient data remain entirely local, ensuring compliance with established healthcare data privacy regulations and standards.

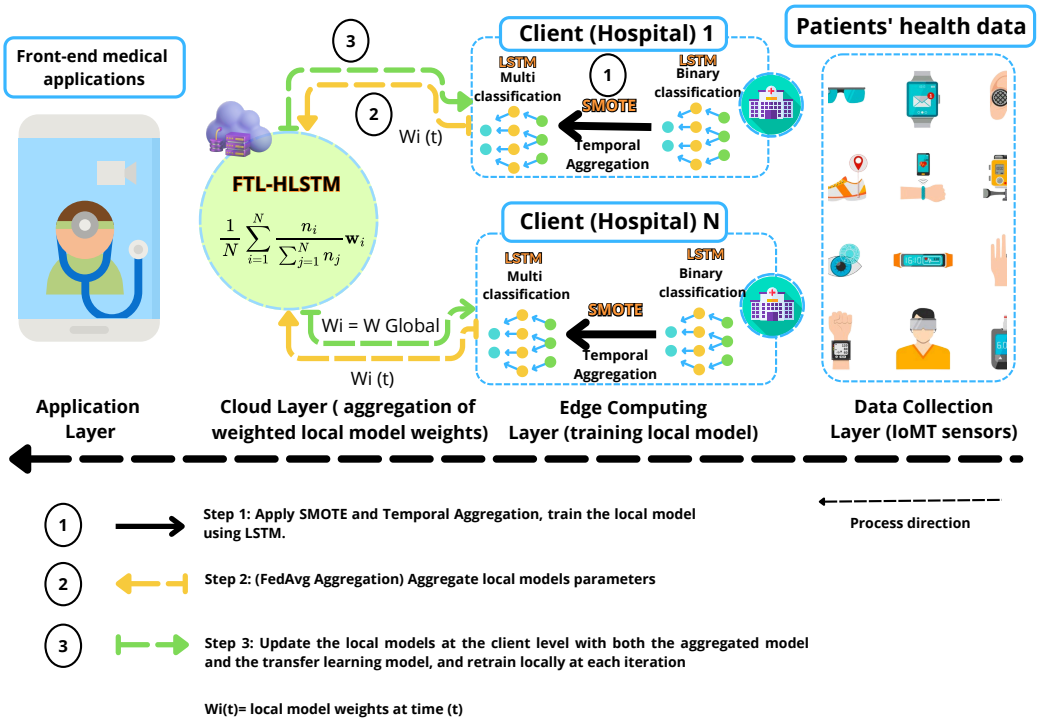

Figure 4.1: Overview of the FTL-HLSTM architecture showing the federated learning approach for intrusion detection in IoMT networks.

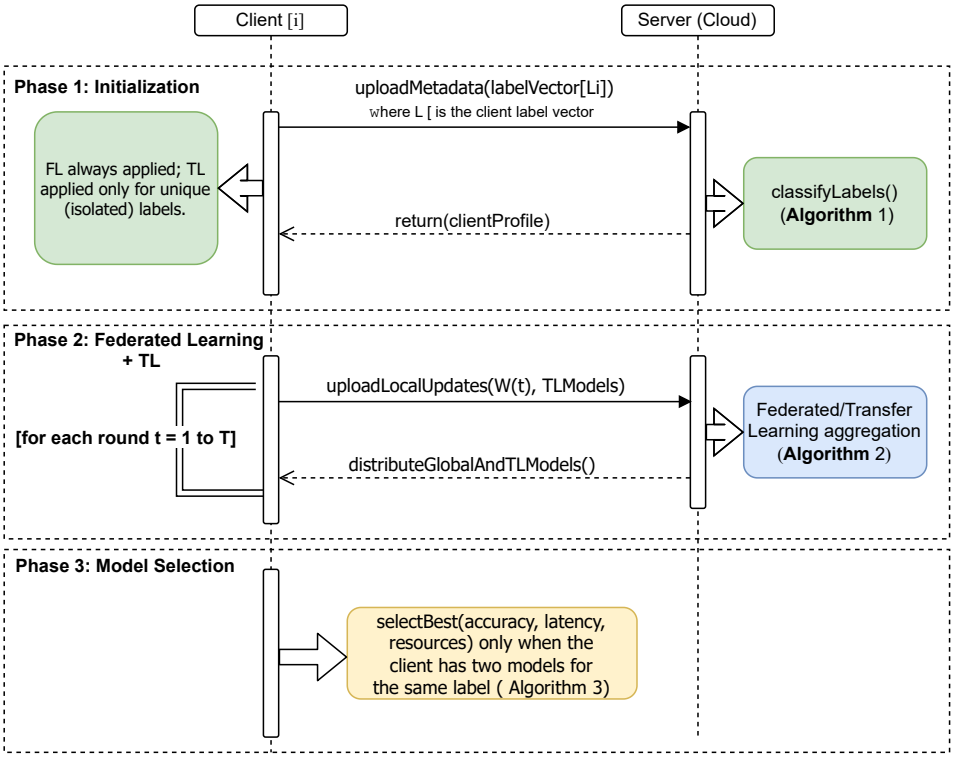

Figure 4.2: Sequence diagram of the proposed Federated Transfer Learning (FTL-HLSTM) workflow.

4.2.2 Server-Side Methodology for Federated Transfer Learning in IoMT

The server functions as a supervisory control entity, orchestrating collaborative training processes among participating clients. At the start of each training cycle, the server executes *Label Categorisation* (Algorithm 1) to partition the label set into *common labels* ($\boldsymbol{L}_{\text{common}}$) and *isolated labels* ($\boldsymbol{L}_{\text{isolated}}$). Isolated labels, which occur exclusively within individual client datasets, introduce significant non-independent and identically distributed (non-IID) challenges, thereby adversely impacting federated learning performance. To mitigate these effects, the server assigns isolated labels to a dedicated *Transfer Learning* (TL) pathway.

Following label categorisation, the server identifies the specific client $c^*(\ell)$ owning each isolated label $\ell \in \boldsymbol{L}_{\text{isolated}}$ and sends a TL notification to that client. Upon receiving the notification, the client constructs a TL model by:

1. Cloning its existing multi-class *Hierarchical Long Short-Term Memory* (HLSTM) model,
2. Converting the output layer into a binary classifier tailored to the isolated label ℓ, and
3. Fine-tuning the classifier using its local dataset $\mathscr{D}_{c^*(\ell)}$.

Simultaneously, clients with non-empty intersections $\mathscr{L}_c \cap \boldsymbol{L}_{\text{common}}$ perform training on subsets of their data corresponding to the common labels. The server aggregates these updates using *Federated Averaging* (FedAvg; Algorithm 2), yielding an updated global model $w^{(t)}$. Subsequently, the server redistributes both the global model and the specialised TL models $\{\text{TL\_MODEL}(\ell)\}$ to all clients for continued training and deployment.

Throughout this entire process, only model parameters and non-sensitive label metadata are exchanged. No raw IoMT data or personally identifiable information leaves the local client environment, ensuring that all privacy and regulatory constraints are maintained and that sensitive data remain securely localised.

Intelligent Label Classification

The Intelligent Label Classification procedure (Algorithm 1) runs at the start of each federated cycle or whenever the client set or label taxonomy changes. Each client $c \in \mathscr{C}$ transmits a binary label-presence vector $v_c \in \{0,1\}^{|\mathscr{Y}|}$, where $v_c[\ell] = 1$ iff label ℓ

is observed locally. The server computes the cross-client support

$$s(\ell) = \sum_{c \in \mathscr{C}} v_c[\ell], \tag{4.1}$$

and partitions the taxonomy $\mathscr{Y}$ using a minimum-support threshold $k_{\min}$:

$$L_{\text{common}} = \{\ell : s(\ell) \geq k_{\min}\}, \qquad L_{\text{isolated}} = \mathscr{Y} \setminus L_{\text{common}}. \tag{4.2}$$

We set $k_{\min} = 2$ by default so that only labels with multi-site evidence contribute to the shared objective, thereby reducing negative transfer from idiosyncratic labels and improving privacy against label-uniqueness inference. (Optional) Secure aggregation and DP noise on $s(\ell)$ may be applied prior to thresholding.

Algorithm 1 Label Categorization at Server Level

Require: Client set $\mathscr{C}$; local label sets $\{L_c\}_{c \in \mathscr{C}}$

Ensure: L_{isolated} (labels appearing in exactly one client); L_{common} (labels appearing in at least two clients)

```
LabelCount ← empty map
for each c ∈ 𝒞 do
    for each ℓ ∈ L_c do
        LabelCount[ℓ] ← LabelCount[ℓ] +1 (initialize to 1 if absent)
    end for
end for
L_isolated ← ∅,   L_common ← ∅
for each ℓ in keys(LabelCount) do
    if LabelCount[ℓ] = 1 then
        L_isolated ← L_isolated ∪ {ℓ}
    else
        L_common ← L_common ∪ {ℓ}
    end if
end for
return L_isolated, L_common
```

Transfer Learning for Isolated Labels

For each $\ell \in L_{\text{isolated}}$, the server identifies its unique owner $c^*(\ell)$ and initiates client-side transfer learning (TL). The client clones the shared encoder θ_s from the global HLSTM, removes the multi-class head, and instantiates a label-specific binary classifier

$$g_\ell(x; \theta_s, \theta_\ell) = \sigma\big(W_\ell^\top h(x; \theta_s) + b_\ell\big), \tag{4.3}$$

where $h(\cdot;\theta_s)$ is the (frozen) encoder output and $\theta_\ell = (W_\ell, b_\ell)$. The head is trained with weighted binary cross-entropy to address class imbalance:

$$\mathscr{L}_\ell(\theta_\ell) = -w_+ \, y_\ell \log g_\ell(x) \; - \; w_- \left(1 - y_\ell\right) \log\bigl(1 - g_\ell(x)\bigr), \tag{4.4}$$

with $y_\ell = \mathbf{1}\{y = \ell\}$ and (w_+, w_-) derived from local class frequencies. A calibrated threshold τ_ℓ may be selected on a validation split to satisfy a site policy, e.g., FPR $\leq 1\%$.

Federated Learning for Common Labels

For client c, let $L_c^{\text{com}} = L_c \cap L_{\text{common}}$. Local training minimizes a masked objective:

$$\min_{w} \; \mathbb{E}_{(x,y)\sim\mathscr{D}_c}\bigl[\,\mathbf{1}\{y \in L_c^{\text{com}}\} \cdot \ell(f(x;w),y)\,\bigr], \tag{4.5}$$

ensuring samples outside L_c^{com} contribute zero gradient. Server-side aggregation employs sample-weighted FedAvg over the set $S_t = \{c \in \mathscr{C} : \; |L_c^{\text{com}}| > 0\}$:

$$w^{(t)} = \frac{\sum_{c\in S_t} n_c^{\text{com}} \, w_c^{(t)}}{\sum_{c\in S_t} n_c^{\text{com}}}, \qquad n_c^{\text{com}} = \bigl|\{(x,y) \in \mathscr{D}_c : y \in L_c^{\text{com}}\}\bigr|. \tag{4.6}$$

Federated and Transfer Learning Workflow

Algorithm 2 Federated and Transfer Learning Workflow

Require: Clients $\mathscr{C}$; initial global model $w^{(0)}$; local epochs E; rounds T; $L_{\text{isolated}}, L_{\text{common}}$ from Algorithm 1

Ensure: $w^{(T)}$ (global model for common labels); $\{w_c^{\text{TL}}\}$ (TL heads for isolated labels)

for all $c \in \mathscr{C}$ **in parallel do**
  $w_c^{\text{local}} \leftarrow w^{(0)}$; train on $\mathscr{D}_c$ for E epochs
end for
for $t = 1$ to T **do**
  Broadcast $w^{(t-1)}$
  for all $c \in \mathscr{C}$ **in parallel do**
    if $L_c \cap L_{\text{common}} \neq \emptyset$ **then**
      $w_c^{(t)} \leftarrow w^{(t-1)}$; train on $\mathscr{D}_c$ masked to L_c^{com} for E epochs
      Upload $w_c^{(t)}$ (or $\Delta w_c^{(t)}$) via secure aggregation
    end if
    if $L_c \cap L_{\text{isolated}} \neq \emptyset$ **then**
      Create w_c^{TL} from w_c^{local}; replace head with binary head(s) for L_c^{iso}
      Freeze encoder; fine-tune head(s) on $\mathscr{D}_c|_{L_c^{\text{iso}}}$; upload w_c^{TL} (params + metrics)
    end if
  end for
  Aggregate (common labels): $w^{(t)} \leftarrow \frac{\sum_{c \in S_t} n_c^{\text{com}} w_c^{(t)}}{\sum_{c \in S_t} n_c^{\text{com}}}$
  collected $\{w_c^{\text{TL}}\}$ **and Redistribute** $w^{(TL)}$ to all clients
end for
return $w^{(T)}$, $\{w_c^{\text{TL}}\}$

Only parameters and metadata are exchanged:

$$\left\{ w^{(t)},\ \{w_c^{(t)}\}_{c \in S_t},\ \{\text{TL\_MODEL}(\ell)\} \right\},$$

and never raw IoMT traffic. Secure aggregation reveals only weighted sums (not individual updates). Optional differential privacy can be applied by adding calibrated noise to clipped gradients or to label-support counts prior to thresholding.

4.2.3 Client-Side Methodology: Hierarchical LSTM with Federated and Side Transfer Learning

This section outlines the client-side implementation of a Hierarchical Long Short-Term Memory (HLSTM) architecture, specifically designed for efficient and effective intrusion detection in Internet of Medical Things (IoMT) deployments (Figure 4.3). The HLSTM architecture seeks to achieve high detection accuracy while minimizing computational overhead and latency, essential for resource-constrained IoMT environments. The methodology employs a structured hierarchical approach consisting of data preprocessing, binary anomaly detection, temporal aggregation, and multi-class classification. Additionally, an optional Side Transfer Learning (STL) module enhances detection of rare, client-specific attack types without disrupting the shared global model.

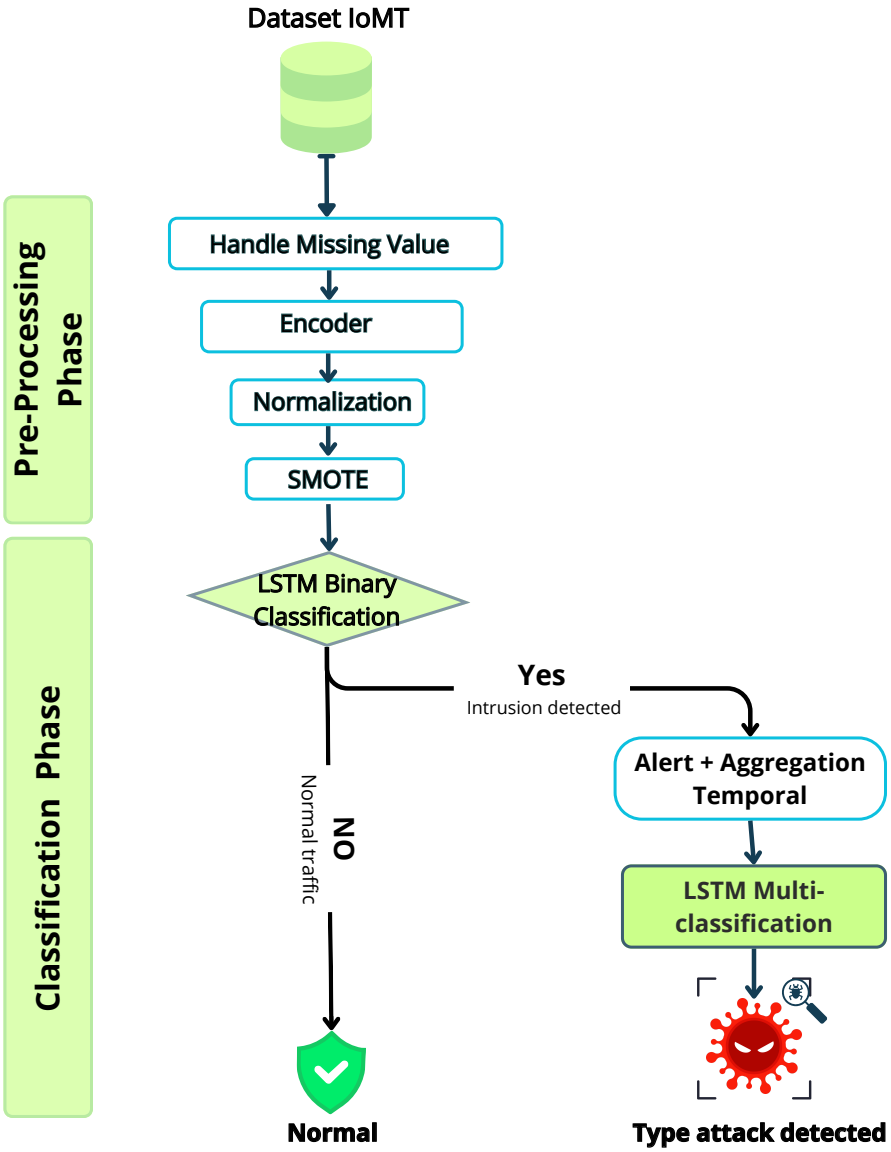

Yes: Intrusion detected, proceed to alert and classify type of attack
No: Traffic is normal, no further action needed.

Figure 4.3: HLSTM architecture for client-side intrusion detection in IoMT networks.

Notation

Let $x_t \in \mathbb{R}^d$ represent the feature vector at time t. Streaming data are segmented into overlapping windows:

$$W_k = \{t_k - L + 1, \ldots, t_k\}, \quad X_k = \{x_t : t \in W_k\}, \tag{4.7}$$

where L denotes window length, s is the stride, k indexes the windows, and t indexes positions within each window.

Client-Side Pre-processing

- **Feature Encoding:** Categorical features are encoded numerically using one-hot encoding, while numerical features remain continuous.
- **Missing Data Handling:** To preserve the integrity of the empirical data, records with missing values are removed during the preprocessing stage. No imputation is applied, as the objective is to base the analysis solely on real observed values and to avoid introducing synthetic estimates that may affect the reliability of the experimental results.
- **Normalization:** Features undergo min-max scaling based on training-set statistics:

$$x'_{i,j} = \frac{x_{i,j} - a_j}{b_j - a_j}, \quad a_j = \min_{\text{train}} x_{\cdot,j}, \quad b_j = \max_{\text{train}} x_{\cdot,j} \tag{4.8}$$

- **Class Imbalance Mitigation:** The Synthetic Minority Oversampling Technique (SMOTE) is applied solely to the training dataset post-windowing to address class imbalance and prevent temporal leakage.

Shared Encoder

An LSTM-based encoder processes each segmented window to generate hidden state representations:

$$(h_{k,1}, \ldots, h_{k,L}) = \text{LSTM}_{\theta_s}(X_k), \quad h_k^\star = \text{pool}(h_{k,1:L}), \tag{4.9}$$

where $h_k^\star$ is a compact representation derived via pooling (e.g., mean-pooling).

- **Stage 1 – Binary Anomaly Detection:** A logistic regression classifier serves as an initial anomaly detection gate:

$$p_k = \sigma(w^\top h_k^\star + b), \quad g_k = \mathbf{1}\{p_k \geq \tau\},$$

  with threshold τ optimized on validation data to balance recall and false-positive rates. Normal traffic ($g_k = 0$) exits early, reducing computational load.

- **Stage 2 – Temporal Aggregation:** Anomalous traffic from Stage 1 is sum-

marized through temporal aggregation:

$$r_k = \frac{1}{L} \sum_{t \in W_k} h_{k,t}, \quad R_k = [r_{k-M+1}, \ldots, r_k] \in \mathbb{R}^{M \times d_h},$$

where R_k captures short-range dynamics using a buffer of the latest M aggregated summaries.

- **Stage 3 – Multi-class Classification:** Aggregated anomalous traffic is classified into specific attack categories via a secondary LSTM:

$$\tilde{h}_k = \text{LSTM}_{\theta_{mc}}(R_k), \quad \hat{y}_k = \text{softmax}(U\tilde{h}_k + c).$$

Side Transfer Learning (STL) for Isolated Labels

For rare or client-specific attack labels ℓ, STL binary classifiers are instantiated using the frozen shared encoder:

$$s_\ell(k) = \sigma(W_\ell^\top h_k^\star + b_\ell), \quad d_\ell(k) = \mathbf{1}\{s_\ell(k) \geq \tau_\ell\}. \tag{4.10}$$

STL classifiers undergo local fine-tuning using class-weighted binary cross-entropy or focal loss, with thresholds τ_ℓ optimized on client-specific validation data. STL parameters are securely transmitted back to the server, ensuring confidentiality of the raw IoMT data.

4.2.4 Client-Side Optimal Model Selection via Weighted Multi-Criteria Analysis

This section presents a weighted multi-criteria decision-making framework designed for optimal client-side model selection in intrusion detection within *Internet of Medical Things* (IoMT) devices. Clients typically possess two models: a global *Federated Learning* (FL) model and a specialized *Transfer Learning* (TL) model tailored for isolated, client-specific attack labels. The overall process is illustrated in Figure 4.4, while Algorithm 3 details the procedure for selecting the best-performing model based on multiple evaluation criteria, such as accuracy, false alarm rate (FAR), and inference latency (TT).

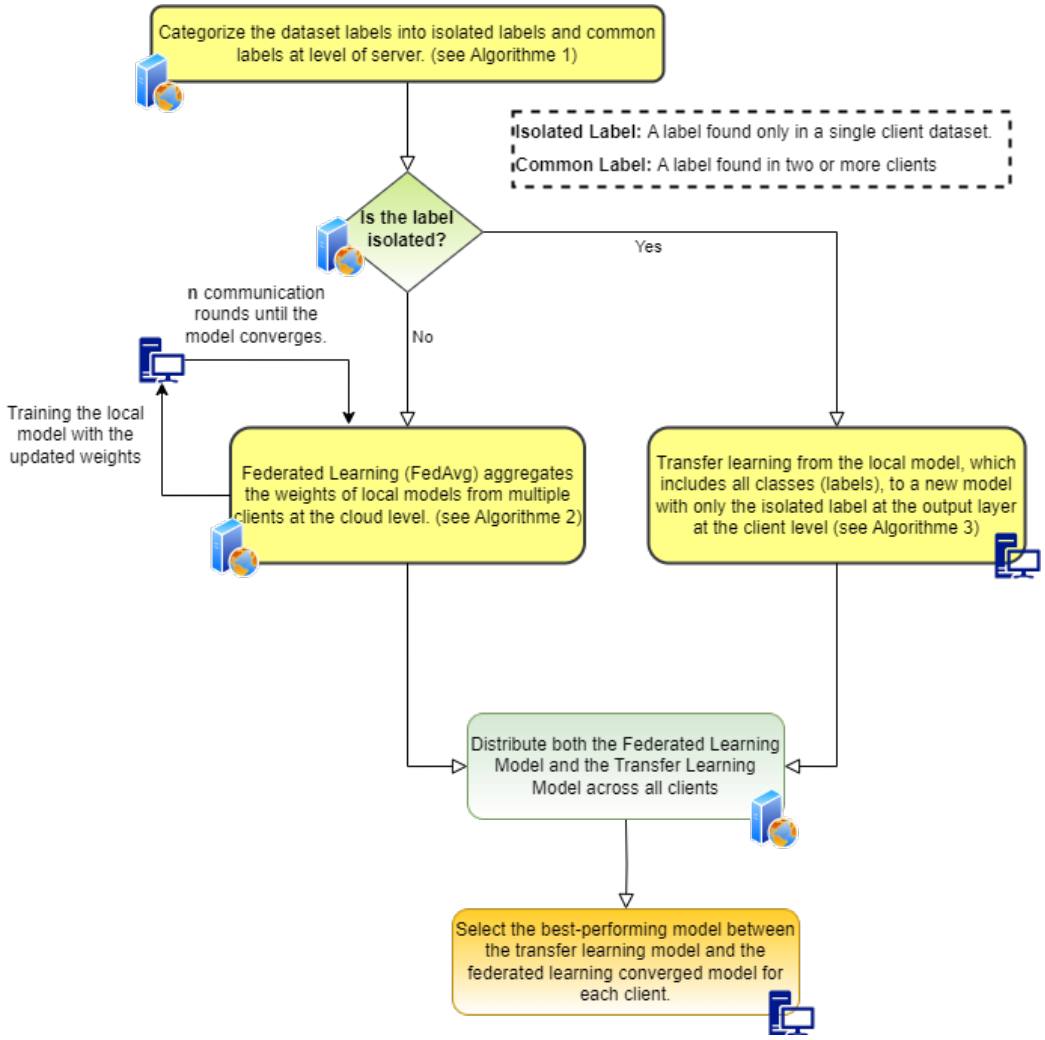

Figure 4.4: Overview of Federated Transfer Learning Approach for IDS in IoMT, highlighting the interaction between client nodes and the central server, showing both the training and inference phases of the system.

Notation Let $\mathscr{M} = \{m_{\text{FL}}, m_{\text{TL}}\}$ denote the set of available candidate models, with m_{FL} representing the global federated learning model and m_{TL} representing the specialized side transfer learning model. The set of evaluation criteria is given by $\mathscr{C} = \{\text{Accuracy}, \text{FAR}, \text{TT}\}$. For each criterion $c \in \mathscr{C}$:

- $T_c > 0$ is the normalization target (e.g., $T_{\text{Accuracy}} = 1$, policy-defined T_{FAR}, or latency target T_{TT}).
- $W_c \geq 0$ is the criterion weight satisfying $\sum_{c \in \mathscr{C}} W_c = 1$.
- $\delta_c \in \{+1, -1\}$ indicates maximization $(+1)$ or minimization (-1).

Let $M_{i,c}$ represent the performance of model m_i on criterion c. A small positive constant $\varepsilon > 0$ is used to prevent numerical instability. Optionally, a subset $\mathscr{H} \subseteq \mathscr{C}$ defines mandatory constraints models must meet to remain feasible.

Each criterion is normalized onto a unified higher-is-better scale through direction-aware normalization:

$$S_{i,c} = W_c \left(\frac{M_{i,c} + \varepsilon}{T_c + \varepsilon} \right)^{\delta_c} \tag{4.11}$$

Metrics for maximization (e.g., accuracy) use $\delta_c = +1$, while metrics for minimization (e.g., FAR, latency) use $\delta_c = -1$. The total utility score for each model is calculated as:

$$S_i = \sum_{c \in \mathscr{C}} S_{i,c}, \quad \text{BestModel} = \arg\max_{m_i \in \mathscr{M}} S_i. \tag{4.12}$$

This approach identifies the model offering the optimal overall trade-off across all evaluation criteria.

4.2.5 Algorithm 3 – Weighted Multi-Criteria Model Selection

Algorithm 3 Client-Side Weighted Multi-Criteria Model Selection

Require: Candidate models $\mathscr{M} = \{m_{\text{FL}}, m_{\text{TL}}\}$; criteria $\mathscr{C}$; targets $\{T_c\}$; weights $\{W_c\}$ with $\sum_c W_c = 1$; directions $\{\delta_c\}$; performance metrics $\{M_{i,c}\}$; tolerance $\varepsilon > 0$; optional hard constraints $\mathscr{H} \subseteq \mathscr{C}$.

Ensure: Optimal model selection $\text{BestModel} \in \mathscr{M}$.

Verify $\sum_c W_c = 1$; if not, normalize W_c.
$\text{BestModel} \leftarrow \text{None}$; $\text{BestScore} \leftarrow -\infty$
for each model $m_i \in \mathscr{M}$ **do**
  if $\exists c \in \mathscr{H}$ violated by m_i (maximized metric below target or minimized metric above target) **then**
    continue ▷ Exclude infeasible models
  end if
  $\text{ModelScore} \leftarrow 0$
  for each criterion $c \in \mathscr{C}$ **do**
    $R \leftarrow \frac{M_{i,c}+\varepsilon}{T_c+\varepsilon}$
    $S_{i,c} \leftarrow W_c \cdot R^{\delta_c}$
    $\text{ModelScore} \leftarrow \text{ModelScore} + S_{i,c}$
  end for
  if $\text{ModelScore} > \text{BestScore}$ **then**
    $\text{BestScore} \leftarrow \text{ModelScore}$; $\text{BestModel} \leftarrow m_i$
  else if $\text{ModelScore} = \text{BestScore}$ **then**
    Apply tie-break criteria (e.g., inference latency, memory usage)
    if tie-break criteria favor m_i **then**
      $\text{BestModel} \leftarrow m_i$
    end if
  end if
end for
return BestModel

4.3 Experimental results

This section presents a comprehensive evaluation of the proposed FTL-HLSTM framework for intrusion detection in IoMT environments. The evaluation progresses systematically from centralised learning through federated learning under IID conditions to challenging non-IID scenarios, demonstrating the framework's robustness across diverse deployment contexts.

4.3.1 Datasets

The practical evaluation employs three benchmark datasets, each selected to represent distinct characteristics of the IoMT security landscape. These datasets capture the inherent heterogeneity and complexity of IoMT traffic patterns and are widely used in intrusion detection research. Collectively, they encompass a broad spectrum of real-world attack scenarios, ranging from general network intrusions to sophisticated threats targeting healthcare-specific devices and infrastructure. To mitigate the adverse effects of class imbalance on model performance, the Synthetic Minority Oversampling Technique (SMOTE) was applied. To preserve the integrity of the evaluation process and avoid data leakage, SMOTE was applied exclusively to the training set, ensuring that the validation and test sets remained unbiased and reflective of real-world intrusion distributions. This targeted oversampling strategy enhances the representation of minority intrusion classes, thereby improving model robustness and significantly strengthening the IDS's capability to detect both frequent and rare attack types.

NF-UNSW-NB15-v2 Dataset

The NF-UNSW-NB15-v2 dataset, [179], is a NetFlow-based cybersecurity dataset encompassing nine distinct attack categories: Exploits, Fuzzers, Generic, Reconnaissance, Denial-of-Service (DoS), Analysis, Backdoor, Shellcode, and Worms. The dataset was constructed by converting publicly accessible packet capture (pcap) files from the original UNSW-NB15 dataset [143] into a structured format comprising 43 features derived via the NetFlow v9 protocol, using the nprobe tool. The NF-UNSW-NB15-v2 dataset contains a total of 2,390,275 network flow records, among which 95,053 (3.98%) represent attack instances, while the remaining 2,295,222 flows (96.02%) constitute benign traffic.

The foundational dataset, UNSW-NB15, is a well-established resource within the network intrusion detection research community, developed and released in 2015 by the Cyber Lab of the Australian Centre for Cyber Security (ACCS). The original UNSW-NB15 dataset employed the IXIA PerfectStorm tool to simulate a combination of normal network traffic and diverse synthetic attack scenarios, providing researchers a comprehensive environment to evaluate network intrusion detection systems (NIDS). The selection of the NF-UNSW-NB15-v2 dataset is particularly relevant for evaluating IoMT intrusion detection systems, as it includes attack categories that closely mirror those faced by medical networks. Notably, the dataset contains backdoor attacks, which threaten patient data integrity, as well as reconnaissance activities—often precursors to targeted attacks on medical devices. These attack patterns are highly representative of the security challenges faced in modern healthcare environments, making this dataset

an ideal candidate for testing the effectiveness of the proposed FTL-HLSTM framework.

This dataset therefore allows the ability of the FTL-HLSTM model to detect both high-frequency and rare attacks to be assessed while maintaining accuracy, particularly in the context of IoMT applications in which security is paramount.

Table 4.1: NF-UNSW-NB15-v2 dataset

Label	**Code**	**Count**
Analysis	0	2299
Backdoor	1	2169
Benign	2	2295222
DoS	3	5794
Exploits	4	31551
Fuzzers	5	22310
Generic	6	16560
Reconnaissance	7	12779
Shellcode	8	1427
Worms	9	164
Total		**2390275**

WUSTL EHMS 2020 Dataset

The WUSTL-EHMS-2020 dataset is specifically designed as a cybersecurity resource to identify and mitigate vulnerabilities within Internet of Medical Things (IoMT) systems through realistic simulation of cyber-attacks within healthcare contexts [85, 174]. Leveraging an integrated testbed that incorporates real-time patient biometric data alongside network traffic, this dataset primarily emphasizes two critical attack scenarios: Man-in-the-Middle (MITM) spoofing and data injection attacks. These scenarios pose significant threats by directly undermining the integrity and confidentiality of sensitive medical data, thus highlighting the dataset's importance for developing robust intrusion detection systems (IDS) tailored to healthcare environments [174].

The dataset comprises over 16,000 labeled samples categorized into normal operations (87.5%) and attack scenarios (12.5%), effectively balancing benign and malicious data traffic. Each record contains 44 unique features, encompassing network flow metrics and real-time biometric data from patients. The integration of real-time biometric information sets this dataset apart from other IoT or IoMT datasets, providing deeper insights into how cybersecurity threats can significantly impact healthcare system performance and patient safety.

Crucial network performance metrics documented in the dataset include packet count, average packet size, and inter-arrival times. These metrics are supplemented with patient-specific biometric parameters, such as heart rate and blood oxygen lev-

els, enabling researchers to authentically simulate and analyze the effects of network disruptions or malicious data manipulations on real-time patient monitoring.

The MITM spoofing attack scenario simulates conditions in which an adversary gains unauthorized access, intercepts, and modifies data transmitted between IoMT devices and healthcare servers or applications. Such attacks can mislead clinical decision-making, causing delayed or incorrect medical responses. Conversely, the data injection attack scenario involves introducing malicious or falsified data into legitimate data streams, compromising patient information accuracy and potentially misleading healthcare providers into making inappropriate clinical decisions.

The WUSTL-EHMS-2020 dataset is broadly applicable in developing advanced IDS through machine learning techniques aimed at detecting and counteracting cyber threats within IoMT infrastructures. Its distinct combination of biometric and network traffic data allows for sophisticated anomaly detection systems capable of identifying both network-level and physiological abnormalities associated with cyber-attacks or system malfunctions. This makes the dataset particularly valuable for anomaly detection research, facilitating the identification of unusual patterns indicative of cyber threats.

By incorporating authentic patient data with detailed simulated cyber-attack scenarios, the WUSTL-EHMS-2020 dataset allows researchers to comprehensively assess the potential impacts of cybersecurity threats on patient safety, medical device reliability, data integrity, and system latency. Furthermore, it supports investigations into medical personnel responses to potential false alerts generated by spoofing or data injection attacks, contributing to enhanced preparedness and resilience within healthcare cybersecurity frameworks.

Table 4.2: wustl-ehms-2020 Dataset

Label	**Code**	**Count**
Data Alteration	0	922
Spoofing	1	1124
Normal	2	14272
Total		**16318**

CICIoMT Dataset 2024

The CICIoMT-2024 dataset, developed by the Canadian Institute for Cybersecurity at the University of New Brunswick, constitutes a comprehensive resource designed explicitly to address cybersecurity vulnerabilities within the rapidly evolving Internet of Medical Things (IoMT) domain [54]. Given the increasing integration of IoMT technologies into healthcare infrastructures, securing medical devices and safeguarding

communication privacy have emerged as critical research priorities. Consequently, the CICIoMT-2024 dataset is intended to facilitate focused research aimed at detecting and mitigating various cyber threats against IoMT systems.

IoMT systems encompass diverse interconnected medical devices, software applications, and related services that exchange data via the Internet to enhance healthcare delivery outcomes. Prominent examples include wearable medical devices, patient-monitoring systems, intelligent hospital beds, and automated medication dispensers. However, such devices frequently exhibit vulnerabilities due to limited computational resources and inadequate security architectures, making them particularly susceptible to cyberattacks. Potential security breaches could lead to severe consequences, including manipulation of sensitive medical data, disruption of essential healthcare services, or unauthorized disclosure of personal health information [54].

The CICIoMT-2024 dataset comprises network traffic data collected from 40 IoMT devices—25 real physical devices and 15 simulated devices representative of typical healthcare settings. These devices utilize multiple communication protocols, including Wi-Fi, Message Queuing Telemetry Transport (MQTT), and Bluetooth, accurately reflecting the heterogeneous nature of real-world IoMT environments. The dataset is structured into two primary directories: the first, titled the Bluetooth Traffic Directory, contains data from Bluetooth-enabled devices. This segment is particularly valuable given the widespread use of Bluetooth in wearable medical devices such as fitness trackers and heart rate monitors, thereby capturing both legitimate and malicious communications specific to the Bluetooth protocol.

A distinctive feature of CICIoMT-2024 is the inclusion of simulated scenarios covering 18 distinct cyber-attacks, effectively highlighting IoMT vulnerabilities. Major attack categories include Distributed Denial-of-Service (DDoS) attacks, characterized by overwhelming targeted devices or networks with excessive traffic to disrupt legitimate access—posing significant threats to healthcare services reliant on device availability. Additionally, the dataset captures various Denial-of-Service (DoS) attack scenarios, typically originating from a single source to disrupt individual nodes or services [54].

The dataset is primarily used for research involving machine learning and anomaly detection systems. Its comprehensive nature, integrating both actual and simulated devices, enables robust testing and validation across multiple cybersecurity scenarios. Beyond cybersecurity applications, the CICIoMT-2024 dataset offers significant value for research involving network analysis, traffic classification, and protocol optimization. The use of actual IoMT devices ensures the dataset authentically represents real operational behaviors in clinical healthcare environments.

CICIoMT-2024 represents a significant advancement in IoMT security research by

providing an extensive and diverse collection of network traffic data encompassing both benign and malicious activities. This resource provides essential empirical evidence to address critical cybersecurity challenges in healthcare. As IoMT adoption continues to expand, datasets such as CICIoMT-2024 will serve as fundamental tools for enhancing the security and reliability of healthcare technology infrastructures [54].

Table 4.3: CICIoMT-2024 Dataset

Label1	**Label2**	**Code**	**Count**
ARP_Spoofing	ARP_Spoofing	0	16047
Benign	Benign	1	192732
MQTT-DDos-Connect_Flood MQTT-DDos-Publish_Flood	MQTT-DDos	2	200659
MQTT-DoS-Publish_Flood MQTT-DoS-Connect_Flood	MQTT-DoS	3	57149
MQTT-Malformed_Data	MQTT-Malformed_Data	4	5130
Recon-Port_Scan Recon-VulScan Recon-Ping_Sweep Recon-os_Scan	Recon	5	103726
TCP_IP-DDos-TCP TCP_IP-DDos-UDP TCP_IP-DDos-ICMP TCP_IP-DDos-SYN	TCP_IP-DDos	6	4779859
TCP_IP-DoS-TCP TCP_IP-DoS-UDP TCP_IP-DoS-ICMP TCP_IP-DoS-SYN	TCP_IP-DoS	7	1805529
Total			**7042831**

4.3.2 Performance Metrics and Evaluation

To evaluate the performance of the proposed framework, this chapter employs widely recognised statistical metrics, namely *accuracy*, *precision*, *recall*, and *F1-score*, computed from the confusion matrix as follows.

Accuracy quantifies the overall correctness of the model:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{4.13}$$

Precision measures the proportion of correctly identified positive instances among all instances predicted as positive:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{4.14}$$

Recall (or sensitivity) evaluates the proportion of actual positive instances that are correctly identified:

$$\text{Recall} = \frac{TP}{TP + FN} \tag{4.15}$$

F1-Score represents the harmonic mean of precision and recall, providing a balanced measure between the two:

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{4.16}$$

Where:

- TP (*True Positive*): Correctly identified anomalies.
- FP (*False Positive*): Normal instances incorrectly identified as anomalies.
- TN (*True Negative*): Correctly identified normal instances.
- FN (*False Negative*): Anomalies incorrectly identified as normal.

4.3.3 Model Parameters and Hyperparameters

Table 4.4 presents a detailed comparison of the model architectures and hyperparameter configurations employed in this chapter. The temporal architectures, specifically LSTM and GRU models, are selected for their capability to capture effectively the temporal dynamics inherent in IoMT attack patterns. The MLP model serves as a non-temporal baseline due to its simpler feedforward structure and comparatively lower computational complexity. XGBoost is additionally included as a tree-based ensemble baseline in order to evaluate the effectiveness of gradient boosting methods in intrusion detection.

Table 4.4: Hyperparameter configurations and architectural specifications for baseline models.

Category	Parameter	MLP	LSTM	GRU	XGB
Architecture	Layer 1	Dense(256)	LSTM(256)	GRU(256)	—
	Layer 2	Dense(128)	LSTM(128)	GRU(128)	—
	Layer 3	Dense(64)	LSTM(64)	GRU(64)	—
	Output	Softmax(n_c)	Softmax(n_c)	Softmax(n_c)	—
Regularization	Dropout	0.40	0.40	0.40	—
	Batch Norm.	Yes	Yes	Yes	—
	L2 Penalty	—	0.01	0.01	—
Optimization	Optimizer	Adam	Adam	Adam	—
	Learning Rate	0.001	0.001	0.001	0.1
	Batch Size	512	64	64	—
Training	Epochs	100	100	100	—
	Loss Function	Sparse Categorical CE			softmax
Tree-Specific	n_estimators	—	—	—	100
	max_depth	—	—	—	6
	subsample	—	—	—	0.8
	colsample_bytree	—	—	—	0.8
	tree_method	—	—	—	hist
	eval_metric	—	—	—	mlogloss

Note: n_c = number of classes; CE = Cross-Entropy; XGB = XGBoost; Batch Norm. = Batch Normalization.

4.3.4 Performance Analysis of Centralized Learning

This subsection rigorously evaluates the performance of the proposed Two-Stage LSTM Pipeline (HLSTM) architecture under centralized training conditions incorporating temporal aggregation. The evaluation establishes essential baseline metrics and systematically investigates the effects of varying temporal aggregation intervals on intrusion detection accuracy and computational efficiency, aiming to optimize predictive performance and training efficiency.

Temporal Aggregation Impact on Accuracy and Efficiency

Temporal aggregation significantly improved the performance metrics of the LSTM-based IDS across all evaluated datasets, as illustrated comprehensively in Table 4.5. Notably, the NF-UNSW-NB15-V2 dataset exhibited the most pronounced improvement, with accuracy increasing from 87.18% at no aggregation to 99.28% at 30-second

intervals. Precision also followed a parallel enhancement, rising from 89.80% to 99.29%. Comparable improvements were observed on the CICIoMT-2024 dataset, where accuracy and precision improved from baseline values of approximately 90.50% to 98.97% at 30-second aggregation intervals. The WUSTL-EHMS-2020 dataset consistently maintained exceptionally high performance, achieving 100% accuracy and precision at both 5-second and 30-second intervals.

Figure 4.5 illustrates the monotonic relationship between aggregation intervals and classification performance across the evaluated metrics. The improvement curves demonstrated rapid gains between 0 and 15 seconds, after which performance metrics plateaued. This trend remained consistent across accuracy, precision, recall, and F1-score metrics for all three datasets, suggesting an optimal trade-off existed between temporal granularity and classification accuracy.

Furthermore, temporal aggregation significantly reduced computational demands in both training and testing phases, as highlighted in Figure 4.5. The NF-UNSW-NB15-V2 dataset experienced the most significant efficiency gains, with training time decreasing by 84.8%, from 32,319.75 seconds to 4,902.53 seconds, as aggregation intervals increased from 0 to 30 seconds. Testing duration demonstrated an even more substantial reduction of 93.1%, decreasing from 260.89 seconds to 18.11 seconds. Similar efficiency enhancements were observed on the CICIoMT-2024 dataset, where training time decreased by 97.2% (from 12,552.78 to 352.92 seconds), and testing time by 97.5% (from 142.42 to 3.55 seconds). The WUSTL-EHMS-2020 dataset, despite being smaller in scale, achieved notable efficiency improvements with reductions of 99.3% and 99.1% in training and testing durations, respectively.

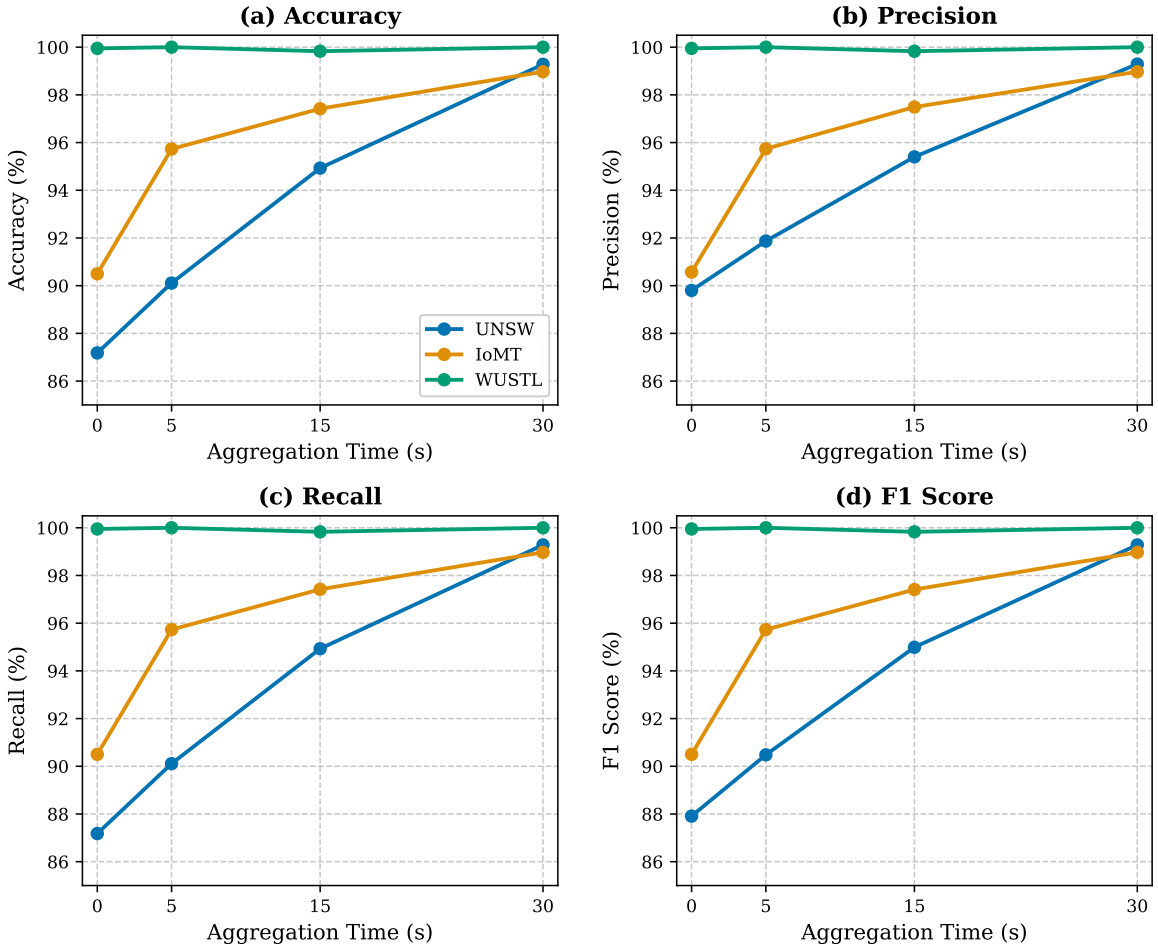

Figure 4.5: Performance Metrics and Computational Efficiency as Functions of Temporal Aggregation

Table 4.5: Temporal Aggregation Effects on LSTM Classification Performance: Cross-Dataset Comparative Analysis

T (s)	Acc (%)	Prec (%)	Rec (%)	F1 (%)	Train Time (s)	Test Time (s)
NF-UNSW-NB15-V2						
0	87.18	89.80	87.18	87.91	32,319.75	260.89
5	90.11	91.87	90.11	90.48	29,427.07	107.91
15	94.93	95.40	94.93	94.99	9,705.09	36.14
30	99.28	99.29	99.28	99.28	4,902.53	18.11
CICIoMT-2024						
0	90.50	90.57	90.50	90.50	12,552.78	142.42
5	95.73	95.74	95.73	95.73	2,216.72	17.63
15	97.42	97.49	97.42	97.41	743.87	6.34
30	98.97	98.97	98.97	98.97	352.92	3.55
WUSTL-EHMS-2020						
0	99.95	99.95	99.95	99.95	4,174.80	11.00
5	99.99	99.99	99.99	99.99	856.16	2.02
15	99.83	99.83	99.83	99.83	48.24	0.14
30	99.99	99.99	99.99	99.99	30.99	0.10

The computational efficiency gains presented in Figure 4.6, utilising logarithmic scaling, underscore the exponential nature of the relationship between temporal aggregation intervals and processing times. The consistent downward trends observed across all datasets confirmed that temporal aggregation not only enhanced IDS performance but also significantly reduced computational resource requirements, making this approach particularly advantageous for deployment in resource-constrained IoMT environments.

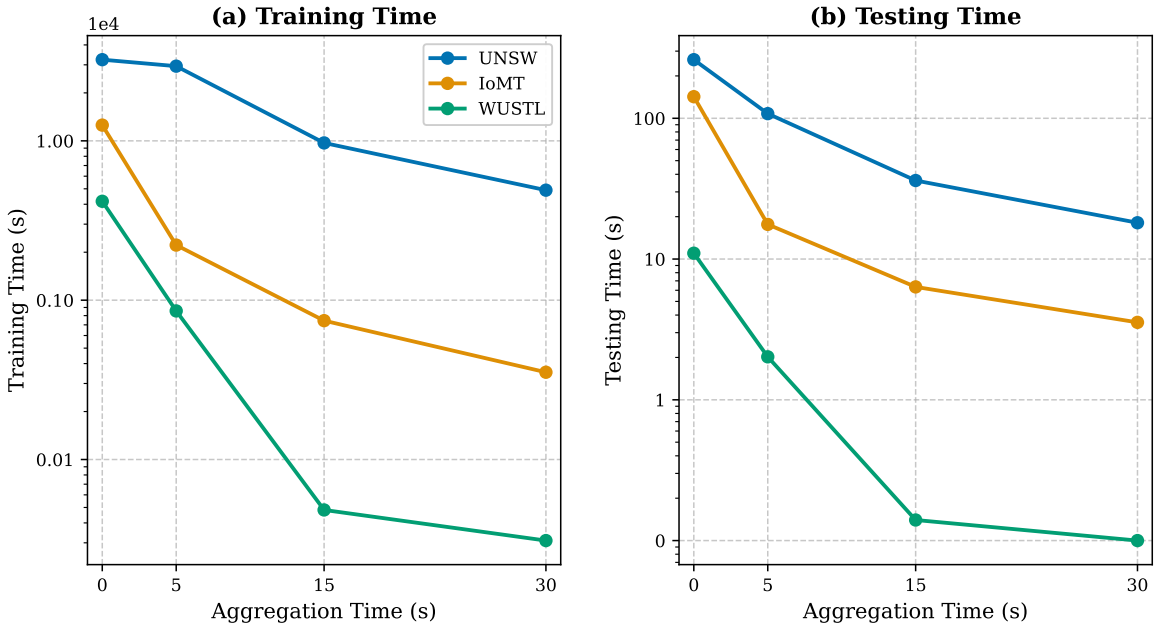

Figure 4.6: Computational Efficiency Gains Through Temporal Aggregation: Training and Testing Time Analysis

Binary versus Multi-Class Detection Comparison

Table 4.6 presents a detailed comparative analysis of binary and multi-class classification approaches employing LSTM models for intrusion detection. Binary classification consistently outperformed multi-class classification across all evaluated datasets, particularly demonstrating significant advantages in scenarios characterised by heterogeneous network traffic.

The NF-UNSW-NB15-V2 dataset exhibited the most pronounced difference in performance. Binary classification achieved perfect detection rates of 99.99% across accuracy, precision, recall, and F1-score metrics. Conversely, multi-class classification attained only 92.84% accuracy and a 93.02% F1-score, indicating a substantial accuracy gap of 7.16 percentage points. This notable discrepancy highlighted the comparative ease of distinguishing regular traffic from malicious activity compared to accurately categorising multiple distinct attack types within a complex network environment.

The CICIoMT-2024 dataset followed a similar trend, albeit with a slightly narrower performance differential. Binary classification achieved an accuracy of 99.74%, accompanied by balanced precision (99.54%) and recall (99.94%), reinforcing its robust detection capabilities. Multi-class classification exhibited considerably lower accuracy (91.62%), representing an 8.12 percentage point gap. Consistency in precision and recall metrics (91.64% and 91.62%, respectively) indicated the absence of systematic bias toward either false positives or false negatives, underscoring the inherent challenges associated with precise categorisation of IoMT-specific attack types.

In contrast, the WUSTL-EHMS-2020 dataset showed minimal differences between binary and multi-class classification performances. Binary classification recorded an accuracy of 99.98%, closely matched by multi-class classification at 99.95%, demonstrating a negligible differential of 0.03 percentage points. Identical precision and recall values (99.95%) in multi-class classification suggested that attack patterns within this specialised medical monitoring dataset possessed sufficiently distinct characteristics, facilitating accurate classification irrespective of the detection approach.

Table 4.6: Comparative Analysis of Binary and Multi-Class LSTM

Dataset	**Method**	**Acc. (%)**	**Prec. (%)**	**Rec. (%)**	**F1 (%)**
NF-UNSW-NB15-V2	Binary	99.99	99.99	99.99	99.99
	Multi-class	92.84	93.51	92.84	93.02
CICIoMT-2024	Binary	99.74	99.54	99.94	99.74
	Multi-class	91.62	91.64	91.62	91.63
WUSTL-EHMS-2020	Binary	99.98	99.99	99.96	99.98
	Multi-class	99.95	99.95	99.95	99.95

Moreover, the integration of temporal aggregation alongside multi-class classification in the second stage, following an initial binary classification in the first stage, established a two-stage detection pipeline that optimised both detection accuracy and computational efficiency. This HLSTM approach, beginning with rapid binary anomaly detection followed by refined temporal aggregation for detailed attack categorisation, enabled real-time intrusion detection essential for resource-constrained IoMT deployments. The substantial improvements detailed in Section 4.3.4, coupled with the proven responsiveness and superior performance of binary classification, effectively validated this two-stage pipeline methodology as both accurate and computationally efficient for real-time IoMT intrusion detection without delays.

Comparative Analysis of HLSTM and Baseline Models

Table 4.7 presents a comprehensive evaluation of the proposed HLSTM against established baseline models: MLP, GRU, LSTM, and XGBoost (XGB). The comparative analysis was conducted across three heterogeneous IoMT datasets. The results consistently demonstrate that HLSTM achieves a superior balance between detection performance and computational efficiency, thereby validating its suitability for real-time intrusion detection in resource-constrained IoMT environments.

On the NF-UNSW-NB15-V2 dataset, HLSTM demonstrated superior balanced performance, achieving an F1-score of 99.63% with corresponding precision and recall values of 99.63%. This result is particularly significant when contrasted with XGBoost's

performance. Although XGBoost achieved high accuracy (99.12%), its F1-score deteriorated substantially to 66.58%, accompanied by a precision of only 64.87%. This pronounced disparity between accuracy and F1-score indicates a critical deficiency in addressing class imbalance, likely manifesting as an elevated false-positive rate that would undermine practical deployment. The recurrent baseline models, GRU (F1-score: 93.32%) and LSTM (F1-score: 93.02%), were also substantially outperformed by HLSTM while incurring considerably higher training times of 14,538.94 seconds and 16,465.25 seconds, respectively. The non-temporal MLP baseline exhibited the poorest performance with an F1-score of 85.06%, confirming the necessity of temporal modeling for effective intrusion detection in this domain.

The CICIoMT-2024 dataset further underscored HLSTM's advantages in handling complex IoMT-specific attack scenarios. HLSTM achieved a near-perfect F1-score of 99.82%, whereas alternative deep learning baselines exhibited substantial performance degradation. The standard LSTM attained an F1-score of 91.63%, while GRU's performance deteriorated markedly to 64.34%, suggesting that its simplified gating mechanism is insufficient for capturing the intricate temporal patterns characteristic of this dataset. XGBoost emerged as the strongest baseline; however, its F1-score of 93.16% remained 6.66 percentage points below HLSTM, with notably lower precision (90.35%) indicative of suboptimal class discrimination. Beyond superior detection metrics, HLSTM's pipeline architecture demonstrated exceptional computational efficiency, achieving a 64.6% reduction in training time and an 83.8% reduction in inference time relative to standard LSTM. These improvements are particularly critical for real-time deployment scenarios where both accuracy and response latency are paramount.

On the WUSTL-EHMS-2020 dataset, all evaluated models achieved near-perfect detection performance (F1-score $>$ 99.84%). This convergence in detection accuracy is likely attributable to the dataset's reduced temporal complexity and well-separated class distributions, which enable even non-temporal models such as MLP and XGBoost to achieve effective classification. Under these conditions, computational efficiency emerged as the primary differentiating factor. XGBoost demonstrated the fastest training time (1.11 seconds), benefiting from its tree-based ensemble architecture. However, HLSTM proved to be the most efficient among temporal models, requiring only 264.03 seconds for training—representing reductions of 35.8% and 31.4% compared to GRU and LSTM, respectively. Moreover, HLSTM achieved an inference time of 0.04 seconds, corresponding to a 96% reduction relative to standard LSTM (1.00 seconds), while remaining competitive with XGBoost (0.01 seconds). This exceptional inference efficiency is particularly valuable for edge-deployed IoMT security systems operating under strict latency constraints.

The comparative analysis presented in Table 4.7 reveals fundamental trade-offs be-

tween computational efficiency and robust, balanced detection performance. XGBoost, while demonstrating exceptional training efficiency across all datasets, exhibited a critical vulnerability: significantly degraded F1-scores and precision on the more complex NF-UNSW-NB15-V2 and CICIoMT-2024 datasets. This deficiency underscores XGBoost's inherent limitation in modeling sequential attack patterns—a capability essential for reliable intrusion detection systems that must identify temporally correlated malicious behaviors.

The standard LSTM and GRU architectures, despite their temporal modeling capabilities, were consistently outperformed by HLSTM in both detection accuracy and computational efficiency. This performance gap can be attributed to HLSTM's architectural innovations. Specifically, the two-stage pipeline—comprising binary classification for initial traffic filtering (Stage 1) followed by temporal aggregation and multi-class classification (Stage 2)—enables efficient extraction and exploitation of long-range temporal dependencies while maintaining computational tractability. The temporal aggregation mechanism in particular allows HLSTM to capture complex attack patterns that unfold over extended time windows, a capability that proves decisive on datasets featuring sophisticated, multi-step attack sequences.

Table 4.7: Performance Evaluation and Comparison of HLSTM with Baseline Models

Arch.	**T (s)**	**Acc. (%)**	**Prec. (%)**	**Rec. (%)**	**F1 (%)**	**Train Time (s)**	**Test Time (s)**
NF-UNSW-NB15-V2							
MLP	0	84.33	87.29	84.33	85.06	816.11	30.99
GRU	0	93.20	93.70	93.20	93.32	14,538.94	67.76
LSTM	0	92.84	93.51	92.84	93.02	16,465.25	65.64
XGBoost	0	99.12	64.87	77.01	66.58	748.56	3.06
HLSTM	**30**	**99.63**	**99.63**	**99.63**	**99.63**	**8,995.80**	**18.94**
CICIoMT-2024							
MLP	0	88.48	88.63	88.48	88.50	1,084.29	35.17
GRU	0	74.34	79.35	74.34	64.34	17,835.57	45.64
LSTM	0	91.62	91.64	91.62	91.63	22,146.85	90.33
XGBoost	0	99.77	90.35	97.73	93.16	1,138.25	7.18
HLSTM	**30**	**99.82**	**99.82**	**99.82**	**99.82**	**7,838.79**	**14.62**
WUSTL-EHMS-2020							
MLP	0	99.99	99.99	99.99	99.99	14.39	0.31
GRU	0	99.96	99.97	99.96	99.96	411.63	1.16
LSTM	0	99.95	99.95	99.95	99.95	384.99	1.00
XGBoost	0	99.97	99.82	99.85	99.84	1.11	0.01
HLSTM	**30**	**99.99**	**99.99**	**99.99**	**99.99**	**264.03**	**0.04**

Furthermore, HLSTM demonstrated significant reductions in both training and testing times across all datasets, confirming its computational efficiency. These improvements are attributed to HLSTM's two-stage pipeline architecture and temporal aggregation, which efficiently capture long-term dependencies while minimizing computational costs. The incorporation of binary classification in the first stage further enhances real-time intrusion detection without compromising response time, making HLSTM a highly effective solution for IoMT security.

4.3.5 Evaluation of Federated Learning

This section presents a comprehensive analysis of federated learning performance, starting with idealized Independent and Identically Distributed (IID) conditions and progressing to more realistic non-IID scenarios. These scenarios reflect the heterogeneous nature of healthcare networks, where data may be distributed unevenly across clients.

Federated Learning Evaluation Under IID Conditions

The FL evaluation commenced under IID conditions using the FedAvg algorithm to establish baseline performance metrics for the HLSTM model. The experimental setup involved 10 federated rounds, each comprising 20 local epochs, with uniformly distributed data shards across participating clients.

Table 4.8 summarises the significant performance gains achieved through temporal aggregation. For the NF-UNSW-NB15-V2 dataset, HLSTM accuracy improved notably from 89.87% without aggregation to 95.46% with 30-second intervals, marking a 5.59 percentage point enhancement. Precision correspondingly rose from 91.14% to 95.64%, and F1-score increased from 90.22% to 95.50%, affirming the efficacy of temporal aggregation in mitigating distributed learning challenges.

The CICIoMT-2024 dataset demonstrated even greater improvement, with HLSTM accuracy rising sharply to 99.82%, an increase of 8.37 percentage points from the baseline LSTM performance of 91.45%. Consistent precision, recall, and F1-score metrics (each at 99.82%) indicated balanced and unbiased detection performance, emphasising temporal aggregation's capability to manage complex IoMT traffic patterns effectively.

For the WUSTL-EHMS-2020 dataset, accuracy slightly increased from 99.99% to 99.99% due to temporal aggregation. While accuracy improvements were minimal, computational efficiency significantly benefited from reduced communication overhead and accelerated convergence.

Table 4.8: Impact of Temporal Aggregation on Performance Metrics across Datasets: A Federated Learning Approach using FedAvg and LSTM

Method	T (s)	Acc (%)	Prec (%)	Rec (%)	F1 (%)
NF-UNSW-NB15-V2 Dataset					
Federated Learning (FedAvg) using LSTM	0	89.87	91.14	89.87	90.22
Federated Learning (FedAvg) using LSTM	30	95.46	95.64	95.46	95.50
CICIoMT-2024 Dataset					
Federated Learning (FedAvg) using LSTM	0	91.45	91.47	91.45	91.45
Federated Learning (FedAvg) using LSTM	30	99.82	99.82	99.82	99.82
WUSTL-EHMS-2020 Dataset					
Federated Learning (FedAvg) using LSTM	0	99.99	99.99	99.99	99.99
Federated Learning (FedAvg) using LSTM	30	99.99	99.99	99.99	99.99

Figure 4.7 illustrates convergence dynamics across datasets. Specifically, for NF-UNSW-NB15-V2 (Figure 4.7a), temporal aggregation enabled rapid convergence within 2–3 rounds to about 95% accuracy, outperforming the non-aggregated baseline, which plateaued around 85% after 5–6 rounds. The loss curves confirmed more stable optimisation with aggregation.

For CICIoMT-2024 (Figure 4.7b), the aggregated model rapidly reached 99% accuracy within three rounds, considerably faster than the baseline method, which stabilised around 90% after five rounds. Loss reduction was notably substantial, further validating aggregation's advantages.

In the WUSTL-EHMS-2020 dataset (Figure 4.7c), temporal aggregation consistently maintained stable performance, contrasting with noticeable instability and fluctuations observed in non-aggregated loss curves during later rounds. This highlighted the crucial role of temporal aggregation in ensuring optimisation stability.

Overall, temporal aggregation substantially enhanced FL dynamics, improving accuracy, accelerating convergence, reducing communication overhead, and ensuring model stability.

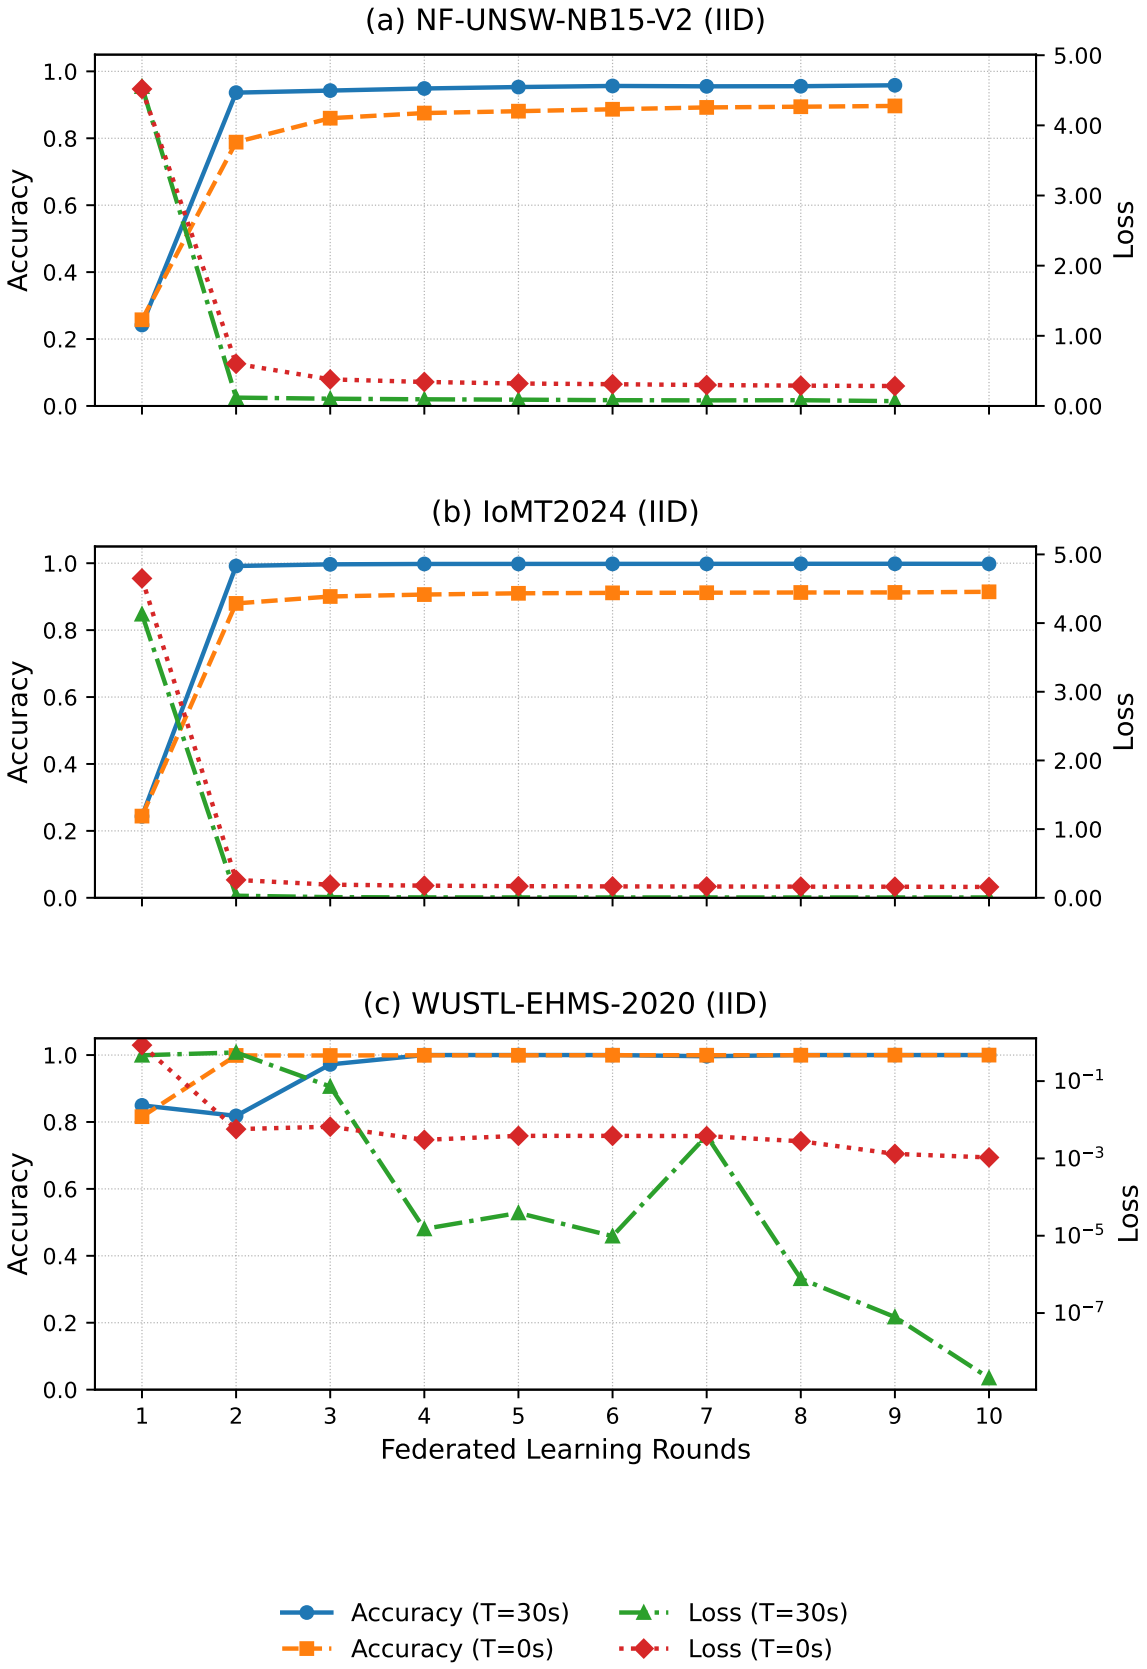

Figure 4.7: Comparison of accuracy and loss across different temporal aggregation intervals in Federated Learning using IID data. The figure illustrates the impact of 0s and 30s aggregation on training performance across multiple datasets.

Evaluation of FTL-HLSTM Under Non-IID Conditions

Real-world IoMT deployments inherently exhibit non-IID data characteristics, posing significant challenges to traditional FL methodologies. Healthcare institutions typically encounter distinct threat profiles influenced by factors such as specialisation, geographic location, and patient demographics, resulting in highly heterogeneous data distributions among federated participants. This section presents a comprehensive evaluation of the FTL-HLSTM framework under substantial data heterogeneity conditions.

Performance on Isolated Labels. This subsection assessed the performance of the FTL-HLSTM model in non-IID scenarios, comparing four FL strategies: standard FedAvg, Transfer Learning (TL) for binary classification, and two hybrid modes. Hybrid Mode 1 employed TL exclusively for isolated labels, whereas Hybrid Mode 2 used TL for isolated labels and adaptively selected between TL and FedAvg based on optimal performance across all labels using Algorithm 3. The experimental setup involved three healthcare institutions, each characterised by unique attack patterns:

- **Client 1 (General Hospital):** Common labels 0, 1, and 2.
- **Client 2 (Cardiac Facility):** Common labels 0 and 2, isolated label 3.
- **Client 3 (Pediatric Institution):** Common labels 1 and 2, isolated label 4.

The experiment comprised 10 federated rounds, each with 20 local epochs, utilising the LSTM(30) architecture.

Table 4.9 demonstrated substantial performance variations among the evaluated strategies, especially concerning isolated labels. Standard FedAvg failed to classify isolated labels, achieving 0.00% accuracy for Labels 3 and 4 across both NF-UNSW-NB15-V2 and CICIoMT-2024 datasets. However, FedAvg delivered satisfactory performance on common labels, with Label 2 achieving 99% accuracy due to its consistent presence across multiple clients.

In contrast, TL consistently provided superior accuracy, exceeding 99% for both isolated and common labels across all datasets. Specifically, TL attained accuracies of 99.87% and 99.85% for isolated Labels 3 and 4 on the NF-UNSW-NB15-V2 dataset, confirming its efficacy in managing isolated label scenarios through binary classification.

Hybrid Mode 1 effectively combined FedAvg for common labels and TL for isolated labels, preserving FedAvg's efficiency while significantly improving accuracy for isolated labels. Hybrid Mode 2 further enhanced adaptability by selectively applying TL to both isolated and poorly performing common labels. Notably, this approach improved

Label 1 accuracy from 40.37% to 99.97%, providing a balanced, performance-oriented solution.

The results obtained from the CICIoMT-2024 dataset mirrored these outcomes, reinforcing the inadequacy of FedAvg for isolated labels and underscoring the effectiveness of TL. Hybrid strategies demonstrated clear advantages in effectively handling heterogeneous IoMT data.

Table 4.9: Per-Label Performance Metrics on Non-IID Data

Method	**Metric**	**Label**				
		0	**1**	**2**	**3**	**4**
NF-UNSW-NB15-V2						
FedAvg	Acc/F1	99.9/77.0	40.4/57.5	100/100	0/0	0/0
TL(Binary)	Acc/F1	99.1/95.9	100/99.0	100/100	99.9/99.4	99.9/99.3
Hybrid-Mode-1	Acc/F1	99.9/77.0	40.4/57.5	100/100	99.9/99.4	99.9/99.3
Hybrid-Mode-2	Acc/F1	99.9/77.0	100/99.0	100/100	99.9/99.4	99.9/99.3
CICIoMT-2024						
FedAvg	Acc/F1	99.1/95.5	91.5/95.1	100/100	0/0	0/0
TL(Binary)	Acc/F1	99.8/99.1	99.7/94.3	100/99.3	100/97.3	99.7/99.1
Hybrid-Mode-1	Acc/F1	99.1/95.5	91.5/95.1	100/100	100/97.3	99.7/99.1
Hybrid-Mode-2	Acc/F1	99.1/95.5	97.8/94.3	100/100	100/97.3	99.7/99.1

B. Computational Complexity Analysis of FL Strategies The computational complexity of each strategy was evaluated in terms of storage, inference, communication, temporal, and spatial complexities, as illustrated in the following table. The parameters used for complexity calculations are defined as follows:

- d: Model dimension
- n: Total number of labels
- k: Number of isolated labels
- p: Number of poorly performing labels
- r: Number of federated rounds
- m: Number of clients

Table 4.10: Computational Complexity Analysis of Federated Learning Strategies for HLSTM -Based Intrusion Detection

Strategy	**Number of Models per Client**	**Temporal Complexity**	**Spatial Complexity**	**Complexity Level**
FedAvg	1 (global model)	$\mathcal{O}(r)$	$\mathcal{O}(d)$	Low
TL	n (binary classifiers)	$\mathcal{O}(r \times n)$	$\mathcal{O}(n \times d)$	High
Hybrid Mode 1	$1+k$ (global + isolated)	$\mathcal{O}(r)$	$\mathcal{O}((1+k) \times d)$	Medium
Hybrid Mode 2	n (adaptive selection)	$\mathcal{O}(r \times n)$	$\mathcal{O}(n \times d)$	High

Table 4.10 summarises the computational complexity of various FL strategies, highlighting the critical trade-offs between performance and resource usage:

- **FedAvg:** Exhibited the lowest complexity ($O(r)$ temporal, $O(d)$ spatial), but delivered insufficient accuracy for isolated labels.
- **TL (Binary):** Demonstrated the highest complexity ($O(r \times n)$ temporal, $O(n \times d)$ spatial), limiting its practical feasibility for resource-constrained devices.
- **Hybrid Mode 1:** Offered moderate complexity ($O(r)$ temporal, $O((1+k) \times d)$ spatial), efficiently addressing isolated labels.
- **Hybrid Mode 2:** Presented adaptive complexity depending on label performance, effectively balancing resource consumption and detection accuracy.

The proposed FTL-HLSTM framework successfully balanced performance and complexity, achieving overall accuracies of 95.46% on the NF-UNSW-NB15-V2 dataset and 99.82% on the CICIoMT-2024 dataset. Its two-stage pipeline architecture facilitated optimised feature sharing across related attack categories, significantly reducing computational redundancy while maintaining high accuracy for isolated labels. This balanced methodology underscored the practical suitability of FTL-HLSTM for real-world IoMT deployments, effectively managing accuracy requirements and computational constraints.

Performance Evaluation Under Dirichlet-Based Non-IID Conditions

In the previous subsection (Section 4.3.5, "Addressing Isolated Labels with FTL-HLSTM"), labels were manually allocated to simulate non-IID conditions. In contrast, this subsection employs a systematic and automated approach using Dirichlet-based partitioning, which provides a more rigorous evaluation of the FTL-HLSTM framework's robustness and generalizability under realistic non-IID scenarios.

The Dirichlet distribution serves as a robust framework for simulating heterogeneous data partitions, with a concentration parameter (α) controlling the degree of statistical heterogeneity. The parameter α was varied across four levels: $\alpha \in \{20.0, 1.0, 0.5, 0.1\}$, transitioning from near-homogeneous distributions ($\alpha = 20.0$) to extreme heterogeneity ($\alpha = 0.1$), thereby encompassing a broad spectrum of federated learning scenarios.

Class allocation patterns across federated clients under different Dirichlet parameters were analyzed using the CICIoMT-2024 dataset (Figure 4.8). Under weak heterogeneity ($\alpha = 20.0$), class proportions were relatively uniform with minimal variance. As heterogeneity increased ($\alpha = 1.0$), variability in class proportions grew significantly, resulting in broader ranges (with a maximum of approximately 0.32 and a minimum approaching zero). At stronger heterogeneity ($\alpha = 0.5$), pronounced client-specific data imbalances emerged, with class proportions fluctuating between 0.00 and 0.43. In scenarios of extreme heterogeneity ($\alpha = 0.1$), the class distributions became highly skewed, with certain classes becoming overwhelmingly dominant within individual client datasets, reaching proportions as high as 0.90.

These findings underscore the effectiveness of Dirichlet-based partitioning in simulating realistic, diverse non-IID conditions, providing a robust foundation for the evaluation of federated learning methodologies.

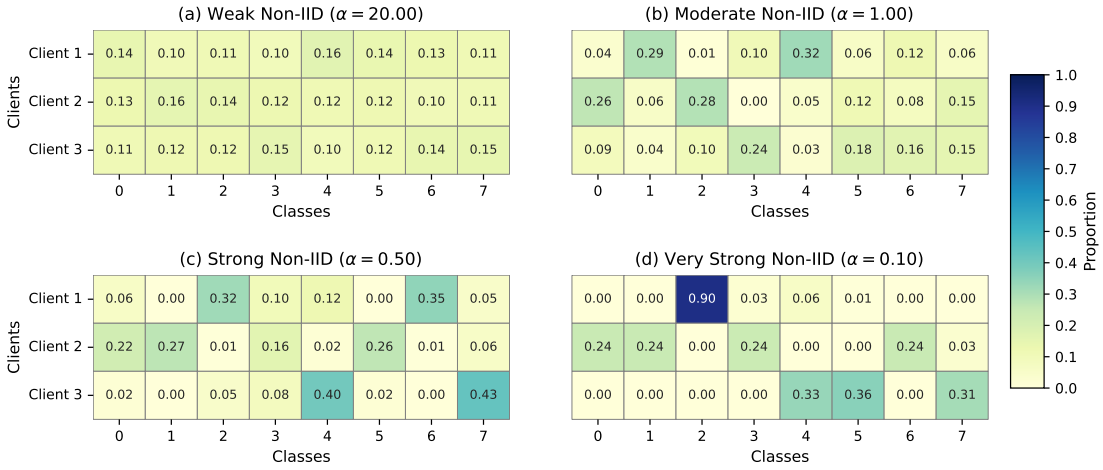

Figure 4.8: Class distribution across clients for the CICIoMT-2024 dataset using Dirichlet-based label allocation.

C. Baseline FedAvg-HLSTM Performance on CICIoMT-2024 Dataset Table 4.11 highlights the significant impact of increasing data heterogeneity on the baseline FedAvg-HLSTM model's performance with the CICIoMT-2024 dataset. Under low heterogeneity conditions ($\alpha = 20$), the model achieved an average accuracy of 96.01%. Despite robust overall performance, notable disparities emerged at the class level, par-

ticularly for Class 6, which exhibited significantly lower accuracy (73.06%). Precision (98.89%) and recall (73.06%) discrepancies indicated difficulties in effectively capturing minority-class characteristics even under near-IID conditions.

Moderate heterogeneity ($\alpha = 1.0$) yielded an improvement in overall accuracy to 99.51%, accompanied by enhancements across all classes. These results indicate beneficial regularisation effects associated with moderate distributional diversity, aligning with ensemble learning principles. However, these performance gains were sensitive to further increases in heterogeneity.

At higher heterogeneity levels ($\alpha = 0.5$), the model maintained high accuracy (99.36%), yet exhibited localised performance reductions, such as lower recall for Class 7 (97.70%). Despite overall stability, these subtle declines highlighted FedAvg's limitations in managing pronounced class imbalance.

Under extreme heterogeneity ($\alpha = 0.1$), the baseline model experienced a substantial performance degradation, achieving an accuracy of only 49.50%. Complete detection failures occurred in several classes (C3, C5, C6), each scoring 0.00% across all metrics. Class 1, despite a high recall (99.73%), showed notably low precision (36.68%), resulting in a high false-positive rate. Consequently, the F1-score declined sharply to 36.86%, clearly reflecting the model's severe inability to manage extreme non-IID scenarios effectively.

D. Performance of Proposed FTL-HLSTM Architecture on CICIoMT-2024 Dataset Table 4.12 demonstrates the proposed FTL-HLSTM architecture's effectiveness in mitigating performance degradation due to increasing data heterogeneity. The adaptive transfer learning mechanism progressively activated as heterogeneity intensified, consistently ensuring robust performance.

Under low and moderate heterogeneity ($\alpha \in \{20, 1\}$), FTL-HLSTM matched the baseline's performance exactly, achieving accuracies of 96.01% and 99.51%, respectively. This equivalence confirmed that transfer learning mechanisms did not introduce unnecessary computational complexity under relatively homogeneous conditions.

The superiority of FTL-HLSTM became evident at more substantial heterogeneity ($\alpha = 0.5$), maintaining high accuracy at 99.36%, identical to baseline performance but significantly enhancing specific class metrics. Class 6, previously vulnerable under baseline conditions, notably improved in recall (99.87%) and balanced precision (97.82%), leading to a significantly enhanced F1-score of 98.84%.

Under conditions of extreme heterogeneity ($\alpha = 0.1$), FTL-HLSTM demonstrated exceptional resilience, achieving an accuracy of 99.72%, substantially surpassing the baseline's 49.50%. Precision improved markedly from 43.59% to 98.16%, recall in-

creased significantly from 49.50% to 98.00%, and the F1-score rose substantially from 36.86% to 98.07%, resolving critical baseline limitations.

Table 4.13 and Figure 4.9 further illustrate the robustness and superiority of FTL-HLSTM compared to FedAvg under severe data heterogeneity. Notably, the FTL-HLSTM demonstrated remarkable consistency. Under extreme heterogeneity ($\alpha = 0.1$), the model achieved accuracies of 99.72% on CICIoMT-2024.

Table 4.11: Performance Evaluation of the FedAvg-HLSTM (30s) Model on the CICIoMT-2024 Dataset under Varying Dirichlet Alpha Settings.

Alpha	**Method**	**Metric**	**C0**	**C1**	**C2**	**C3**	**C4**	**C5**	**C6**	**C7**	**Avg**
20	FedAvg-HLSTM (30s)	Accuracy	99.75	98.46	99.99	99.47	99.99	98.06	73.06	99.26	96.01
		Precision	96.78	99.86	99.47	99.99	99.98	99.88	98.89	78.61	96.69
		Recall	99.75	98.46	99.99	99.47	99.99	98.06	73.06	99.26	96.01
		F1 Score	98.24	99.15	99.74	99.73	99.99	98.96	84.04	87.74	95.95
1	FedAvg-HLSTM (30s)	Accuracy	99.99	99.32	99.99	99.40	99.99	98.89	99.11	99.35	99.51
		Precision	98.35	99.87	99.41	99.99	99.96	99.98	99.39	99.16	99.51
		Recall	99.99	99.32	99.99	99.40	99.99	98.89	99.11	99.35	99.51
		F1 Score	99.17	99.59	99.70	99.70	99.98	99.43	99.25	99.26	99.51
0.5	FedAvg-HLSTM (30s)	Accuracy	99.92	99.29	99.99	99.95	99.99	98.53	99.87	97.70	99.36
		Precision	97.91	99.89	99.55	99.99	99.96	99.89	97.82	99.93	99.37
		Recall	99.92	99.29	99.99	99.55	99.99	98.53	99.87	97.70	99.36
		F1 Score	98.91	99.59	99.77	99.77	99.98	99.20	98.84	98.80	99.36
0.1	FedAvg-HLSTM (30s)	Accuracy	3.42	99.73	99.99	0.00	93.11	0.00	0.00	99.77	49.50
		Precision	99.09	36.68	80.92	0.00	99.56	0.00	0.00	32.44	43.59
		Recall	3.42	99.73	99.99	0.00	93.11	0.00	0.00	99.77	49.50
		F1 Score	6.60	53.63	89.45	0.00	96.23	0.00	0.00	48.96	36.86

Table 4.12: Performance Evaluation of the **Proposed Model: FTL-HLSTM (30s)** on the CICIoMT-2024 Dataset under Varying Dirichlet Alpha Settings.

Alpha	**Method**	**Metric**	**C0**	**C1**	**C2**	**C3**	**C4**	**C5**	**C6**	**C7**	**Avg**
20	Proposed Model: FTL-HLSTM (30s)	Accuracy	99.75	98.46	99.99	99.47	99.99	98.06	73.06	99.26	96.01
		Precision	96.78	99.86	99.47	99.99	99.98	99.88	98.89	78.61	96.69
		Recall	99.75	98.46	99.99	99.47	99.99	98.06	73.06	99.26	96.01
		F1 Score	98.24	99.15	99.74	99.73	99.99	98.96	84.04	87.74	95.95
1	Proposed Model: FTL-HLSTM (30s)	Accuracy	99.99	99.32	99.99	99.40	99.99	98.89	99.11	99.35	99.51
		Precision	98.35	99.87	99.41	99.99	99.96	99.98	99.39	99.16	99.51
		Recall	99.99	99.32	99.99	99.40	99.99	98.89	99.11	99.35	99.51
		F1 Score	99.17	99.59	99.70	99.70	99.98	99.43	99.25	99.26	99.51
0.5	Proposed Model: FTL-HLSTM (30s)	Accuracy	99.92	99.29	99.99	99.95	99.99	98.53	99.87	97.70	99.36
		Precision	97.91	99.89	99.55	99.99	99.96	99.89	97.82	99.93	99.37
		Recall	99.92	99.29	99.99	99.55	99.99	98.53	99.87	97.70	99.36
		F1 Score	98.91	99.59	99.77	99.77	99.98	99.20	98.84	98.80	99.36
0.1	Proposed Model: FTL-HLSTM (30s)	Accuracy	99.79	99.67	99.96	99.96	99.66	99.85	99.88	99.80	99.72
		Precision	98.12	90.98	99.01	98.13	98.21	98.55	99.86	99.85	98.16
		Recall	99.99	97.80	99.57	96.55	99.99	91.23	99.95	99.92	98.00
		F1 Score	99.05	94.27	99.29	97.33	99.10	94.75	99.90	99.88	98.07

Table 4.13: Global Performance Comparison of FedAvg-HLSTM and Proposed FTL-HLSTM Models (30s Interval) on CICIoMT-2024 Dataset Across Varying Dirichlet Alpha Settings

Alpha	**Metric**	**FedAvg-HLSTM**	**FTL-HLSTM (Proposed)**
20	Accuracy	96.01	96.01
	Precision	96.69	96.69
	Recall	96.01	96.01
	F1 Score	95.95	95.95
1	Accuracy	99.51	99.51
	Precision	99.51	99.51
	Recall	99.51	99.51
	F1 Score	99.51	99.51
0.5	Accuracy	99.36	99.36
	Precision	99.37	99.37
	Recall	99.36	99.36
	F1 Score	99.36	99.36
0.1	Accuracy	49.50	99.72
	Precision	43.59	98.16
	Recall	49.50	98.00
	F1 Score	36.86	98.07

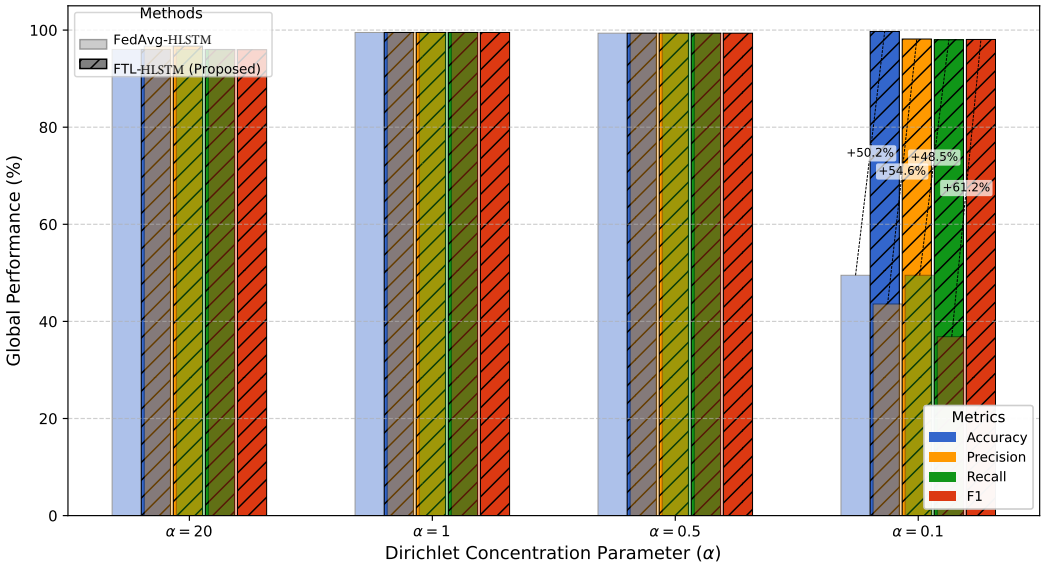

Figure 4.9: **Comparative global performance of FedAvg-HLSTM and the proposed FTL-HLSTM (30s)** across varying Dirichlet alpha values on the **CICIoMT-2024** dataset.

4.3.6 Comparative Performance Analysis Against Non-IID Robust Baselines

Following the preliminary heterogeneity analysis presented in Section 4.3.5, this subsection benchmarks FTL-HLSTM against federated learning algorithms specifically designed to address non-IID data distributions. The comparison includes two established baselines: FedProx, which constrains local model drift using proximal regularisation, and SCAFFOLD, which employs control variates to correct client-drift-induced gradient variance.

The evaluation employs the CICIoMT-2024 dataset under severe statistical heterogeneity. Client datasets were generated using a Dirichlet distribution with concentration parameter $\alpha = 0.1$, producing extreme label-distribution skew. Under this configuration, each client observes only a small, highly imbalanced subset of the global label space, with some classes absent—a scenario representative of realistic IoMT deployments where devices monitor distinct patient groups or pathological conditions.

Federated training was performed with three clients over 20 communication rounds, with each client executing five local epochs per round. All algorithms utilised the same HLSTM backbone architecture, ensuring that performance differences stem solely from the aggregation strategies rather than architectural variations.

All methods were implemented using their published optimal configurations. FedProx was evaluated with proximal terms $\mu \in \{0.01, 0.1\}$ using the Adam optimizer

Table 4.14: Performance under extreme non-IID (Dirichlet $\alpha = 0.1$) on CICIoMT-2024.

Algorithm	**Acc.**	**Prec.**	**Rec.**	**F1**
FedProx ($\mu = 0.1$)	69.36	33.15	56.61	36.31
FedProx ($\mu = 0.01$)	68.44	33.66	55.44	36.11
SCAFFOLD	66.49	33.24	43.14	28.33
FTL-HLSTM	**99.72**	**98.16**	**98.00**	**98.07**

(learning rate $\eta = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$), batch size $B = 32$, L2 regularization weight $\lambda = 0.01$, and dropout rate $p = 0.4$. SCAFFOLD employed SGD with learning rate $\eta = 0.0001$ and momentum $\gamma = 0.9$, batch size $B = 32$, and a two-tier gradient clipping mechanism consisting of a local clip norm $\tau_1 = 1.0$ and a global threshold $\tau_2 = 5.0$. FTL-HLSTM adopted the same optimiser configuration as FedProx but incorporated the dual federated transfer learning mechanism, including pre-trained feature extractors and adaptive knowledge distillation.

Table 4.14 reports the performance after 20 rounds of federated training. Under extreme non-IID conditions, the baseline algorithms exhibit substantial degradation. FedProx achieves an F1-score of 36.31% for $\mu = 0.1$, demonstrating that proximal regularisation alone cannot reconcile gradient conflicts derived from divergent local objectives. Precision remains low at 33.15%, reflecting significant misclassification of minority classes absent from local datasets. SCAFFOLD exhibits even greater performance collapse, achieving an F1-score of 28.33% and a recall of only 43.14%. Despite its variance-reduction mechanism, the method struggles when client objectives diverge significantly across the network.

In contrast, FTL-HLSTM demonstrates highly stable convergence and near-optimal performance, achieving 99.72% accuracy, 98.16% precision, 98.00% recall, and a 98.07% F1-score. These outcomes correspond to improvements of 61.76 percentage points over FedProx and 69.74 points over SCAFFOLD in F1-score alone, highlighting the substantial benefits of integrating targeted transfer learning with federated aggregation.

These findings establish FTL-HLSTM as a robust and computationally efficient solution for federated IoMT intrusion detection under extreme non-IID conditions. The framework consistently outperforms specialised non-IID baselines while maintaining feasibility for deployment on resource-constrained edge devices.

Scalability Evaluation of FTL-HLSTM

To evaluate the scalability of the FTL-HLSTM architecture, experiments were conducted across federation sizes ranging from 4 to 100 clients under IID conditions. Three benchmark datasets representing distinct application domains were employed for this

evaluation, namely NF-UNSW-NB15-V2, CICIoMT-2024, and WUSTL-EHMS-2020.

Across all datasets and federation sizes, the FTL-HLSTM consistently demonstrated balanced precision and recall, with closely aligned F1-scores and accuracy metrics. This consistency highlights the architecture's robust feature extraction capabilities and its ability to maintain predictive accuracy as the federation size increases.

Results presented in Table 4.15 demonstrate that FTL-HLSTM consistently maintained high accuracy levels exceeding 92%, even at the largest evaluated scale of 100 clients. The observed stability and predictable performance trends across increasing federation sizes underscore the architecture's reliability and practical suitability for large-scale federated learning deployments. These findings collectively confirm the scalability, stability, and adaptability of the FTL-HLSTM architecture, emphasising its effectiveness for federated learning scenarios requiring significant client scalability.

Table 4.15: Scalability performance of FTL-HLSTM (30s) under IID conditions

Clients	**Dataset**	**Acc. (%)**	**Prec. (%)**	**Rec. (%)**	**F1 (%)**
4	UNSW	95.53	96.52	95.53	95.47
	IoMT	99.65	99.65	99.65	99.65
	WUSTL	99.99	99.99	99.99	99.99
10	UNSW	94.23	94.78	94.23	94.26
	IoMT	99.47	99.47	99.47	99.47
	WUSTL	99.99	99.99	99.99	99.99
20	UNSW	93.45	95.44	93.45	93.22
	IoMT	98.74	98.76	98.74	98.74
	WUSTL	99.99	99.99	99.99	99.99
50	UNSW	92.42	94.30	92.42	92.18
	IoMT	98.67	98.72	98.67	98.68
	WUSTL	99.99	99.99	99.99	99.99
100	UNSW	95.25	96.34	95.25	95.17
	IoMT	97.48	97.58	97.48	97.49
	WUSTL	99.99	99.99	99.99	99.99

4.4 Discussion

The experimental findings presented robustly validate the effectiveness of the proposed FTL-HLSTM architecture for IoMT environments. These outcomes comprehensively fulfil both theoretical propositions and practical objectives defined in this research.

A notable enhancement within FTL-HLSTM is the implementation of temporal aggregation, which considerably enhances intrusion detection capabilities. Empirical

evaluations using the NF-UNSW-NB15-V2 dataset demonstrated substantial accuracy improvements from 87.18% without temporal aggregation to 99.28% with a 30-second aggregation interval—representing a relative enhancement of approximately 13.88%. Additionally, temporal aggregation notably optimised computational efficiency by significantly reducing training duration by 84.84% (from 32,319.75 s to 4,902.53 s) and inference latency by 93.06% (from 260.89 s to 18.11 s) (Table 4.5; Figure 4.5). Similar enhancements were observed in FL contexts, where accuracy improved from 89.87% to 95.46% under independent and IID conditions (Table 4.8). These findings underscore temporal aggregation's clear superiority over conventional LSTM and GRU-based methods, particularly in terms of detection accuracy and computational resource utilisation (Table 4.7).

Moreover, integrating TL into FL frameworks effectively mitigated statistical heterogeneity common in IoMT scenarios. Under severely non-IID conditions characterised by a Dirichlet distribution parameter ($\alpha = 0.1$), traditional Federated Averaging (FedAvg) methods suffered significant performance deterioration, achieving near-random accuracy levels of 49.5%. Conversely, the proposed FTL-HLSTM architecture maintained exceptional accuracy, reaching 99.72% (Table 4.13). This robust performance results from the targeted application of TL, specifically addressing isolated or underrepresented attack labels. Such targeted interventions effectively resolve gradient conflicts and preserve essential client-specific information, as confirmed by detailed per-label analyses (Table 4.12).

Additionally, the proposed hybrid FL methodology presents a significant theoretical advancement by optimally balancing accuracy and computational complexity. This methodology uses FedAvg for common labels and applies TL selectively to isolated or underperforming labels. Consequently, spatial and model-count complexity scales linearly with the number of isolated labels ($\mathcal{O}(1+k)$), as opposed to the $\mathcal{O}(n)$ complexity of traditional binary-classification-based TL methods. This characteristic makes the hybrid approach particularly suitable for resource-constrained IoMT contexts (Table 4.10).

Scalability analyses further confirm the practical applicability and robustness of FTL-HLSTM. Evaluations across federation sizes ranging from 4 to 100 clients consistently demonstrated accuracy exceeding 92%, alongside balanced precision and recall metrics across multiple benchmark datasets (NF-UNSW-NB15-V2, CICIoMT-2024, WUSTL-EHMS-2020) (Table 4.15). These findings highlight the model's consistent performance and scalability, reinforcing its suitability for real-world deployment in clinical environments where accurate and real-time cybersecurity is crucial for patient safety and care continuity.

This research contributes significantly across three critical dimensions: (i) propos-

ing and validating temporal aggregation techniques to enhance intrusion detection performance in IoMT; (ii) developing an innovative hybrid FL architecture capable of addressing statistical heterogeneity and isolated attack categories effectively; and (iii) empirically verifying scalability, robustness, and practical relevance of the proposed model across diverse IoMT scenarios. Collectively, these contributions significantly advance both theoretical understanding and practical implementation of cybersecurity measures in healthcare infrastructures.

4.4.1 Comparison with Related Work

Table 4.16 presents an extensive comparative evaluation of the proposed FTL-HLSTM framework against contemporary state-of-the-art IDS across multiple benchmark datasets specifically tailored for IoMT environments. The comparative analysis distinguishes between federated and centralised learning paradigms, elucidating the performance advantages and architectural innovations of the proposed method.

Table 4.16: Performance comparison of recent IDS for IoMT

Method	**Learning Paradigm**	**Acc. (%)**	**Prec. (%)**	**Rec. (%)**	**F1 (%)**
CICIoMT-2024 Dataset					
RF [54](2024)	Centralized	73.3	69.1	57.7	55.1
Feature selection+RF [176](2025)	Centralized	93.5	94.0	93.0	93.0
Two-stacked LSTM [8](2025)	Centralized	98.0	98.0	98.0	98.0
FL+RF [140](2025)	Federated [Non-IID]	99.2	99.4	99.2	99.1
BiGRU-BiLSTM [198](2025)	Centralized	99.9	99.7	99.9	99.9
FTL-HLSTM	**Federated [Non-IID†]**	**99.7**	**98.2**	**98.0**	**98.1**
NF-UNSW-NB15-v2 Dataset					
FL+LSTM [180](2023)	Federated [Non-IID]	NA	NA	85.3	NA
FTL-HLSTM	**Federated [Non-IID†]**	**99.9**	**98.7**	**99.9**	**99.5**
WUSTL-EHMS-2020 Dataset					
FL+PCA+DNN [113](2024)	Federated [Non-IID]	88.0	66.0	50.0	57.0
FTL-HLSTM	**Federated [IID]**	**99.99**	**99.99**	**99.99**	**99.99**

†Extreme Non-IID: Dirichlet($\alpha = 0.1$)

Federated approaches such as FL combined with LSTM [180] and FL with PCA + DNN [113] have demonstrated varying performance under non-IID data conditions. Specifically, the FL+LSTM method achieved a recall of only 85.3% on the NF-UNSW-NB15-v2 dataset under less severe non-IID settings. In contrast, the proposed FTL-HLSTM demonstrates substantial improvement, achieving a recall of 99.9% under extreme non-IID conditions ($\alpha = 0.1$), corresponding to a relative increase of approximately 14.6%. Similarly, FL+PCA+DNN reported a critically limited recall of 50% on the WUSTL-EHMS-2020 dataset under non-IID conditions, whereas the proposed framework achieved near-perfect accuracy, precision, recall, and F1-score metrics

(99.99%), underscoring its exceptional resilience to data-distribution challenges. On the CICIoMT-2024 dataset, the federated Random Forest (FL+RF) approach [140] attained a competitive accuracy of 99.2%; however, it inherently lacks mechanisms to model effectively the temporal dependencies that are crucial for detecting sophisticated attack patterns in dynamic IoMT systems. The proposed FTL-HLSTM framework addresses these limitations, providing a superior accuracy of 99.7% together with a well-balanced precision, recall, and F1-score, which together reflect its robust capability to manage severe non-IID conditions while preserving data privacy. Centralised learning methodologies display varying degrees of performance when evaluated on the CICIoMT-2024 dataset. Specifically, the baseline Random Forest model [54] demonstrates limited efficacy, achieving an accuracy of only 73.3% along with a notably low F1-score of 55.1%. In contrast, centralised approaches such as optimised feature selection combined with Random Forest [176] and stacked Long Short-Term Memory (LSTM) architectures [8] have demonstrated commendable performance. However, these methods rely on the centralised aggregation of data, and therefore pose potential risks to patient privacy—an essential consideration within healthcare contexts. The FTL-HLSTM framework, while matching or surpassing these centralised systems in performance, notably preserves data locality and privacy, which are critical for regulatory compliance. The BiGRU-BiLSTM approach [198] displays marginally superior accuracy (99.9%) on CICIoMT-2024; however, it requires centralised data aggregation and incurs significantly greater computational complexity due to its bidirectional processing architecture. In contrast, the proposed framework achieves comparable accuracy (99.7%) while substantially reducing computational costs, with training time reduced by 84.8% and inference latency reduced by 93.1% through effective temporal aggregation. The proposed FTL-HLSTM framework represents a significant methodological advancement in federated IDS for IoMT, effectively managing statistical heterogeneity and stringent privacy requirements, while consistently demonstrating superior or comparable performance relative to existing federated and centralised methodologies.

4.5 Conclusion

This chapter has presented FTL-HLSTM, a Federated Transfer Learning framework built upon a hierarchical Long Short-Term Memory backbone, developed in response to the research gaps identified in Chapter 3.1 concerning privacy-preserving intrusion detection in resource-constrained and statistically heterogeneous IoMT environments. The framework was designed around an explicit reconciliation of three otherwise competing objectives—detection accuracy, computational efficiency, and data privacy—and

has been shown, through a systematic empirical campaign, to deliver each of them simultaneously rather than at the expense of the others.

The temporal aggregation mechanism integrated into the FTL-HLSTM architecture proved to be a decisive architectural choice. On the NF-UNSW-NB15-v2 dataset, it yielded an absolute accuracy improvement of up to 13.87 %, together with an 84.84 % reduction in training time and a 93.06 % reduction in inference latency. These results provide concrete evidence that time, when treated as a first-class design variable, carries detection information that is systematically under-exploited by packet-level pipelines. They also substantiate the central claim of the thesis, articulated in the General Introduction, that intelligent time allocation is not merely an efficiency concern but a genuine enabler of scalable and secure IoT communication.

The selective application of transfer learning to isolated attack categories, combined with federated averaging for commonly observed threats, successfully mitigated the statistical heterogeneity that characterises real-world IoMT deployments. Under a severe non-IID regime parameterised by a Dirichlet distribution with $\alpha = 0.1$, FTL-HLSTM attained an accuracy of 99.86 %, corresponding to a relative gain of approximately 190 % over standard Federated Averaging. This result directly addresses the second research objective of the thesis and confirms that hybridising federated and transfer learning provides a principled mechanism for handling label skew without any exchange of raw clinical data.

The scalability analysis further supports the practical viability of the framework. Accuracy consistently exceeded 92 % across federation sizes ranging from four to one hundred clients, and these levels were preserved across three independent benchmark datasets covering complementary traffic regimes. The framework therefore satisfies the combined requirements of accuracy, privacy preservation, and deployability that were set out as prerequisites for a clinically viable intrusion detection system.

At the same time, the results expose limitations that the thesis acknowledges transparently. The spatial complexity of the hybrid federated–transfer pathway scales linearly in the number of isolated labels, $\mathcal{O}(1+k)$, and may become constraining in deployments that exhibit a large number of institution-specific threat categories. Moreover, FTL-HLSTM has not been instrumented against adversarial Byzantine behaviour at the parameter level, so that mechanisms such as gradient poisoning, model replacement, and membership-inference attacks remain outside its current protective scope. These limitations delineate a well-defined research agenda, which is taken up both in the General Conclusion of the thesis and in the complementary defensive layer introduced next.

The contributions reported in this chapter should not, however, be regarded as a

self-contained solution. Detection intelligence is only operationally meaningful when the underlying network is able to act upon it, which presupposes a routing substrate capable of excluding nodes that have been identified as compromised or unreliable. This observation motivates the second original contribution of the thesis. Chapter 5 therefore develops a proactive, trust-aware routing framework based on multi-criteria decision analysis, in which the attack labels produced by FTL-HLSTM are consumed as a safety signal by the routing layer. Together, the two chapters realise the closed-loop, self-protecting cyber-physical architecture announced as the overarching contribution of this dissertation.

Chapter

5

Proactive Fault Tolerance

Chapter 5

Proactive Fault-Tolerant Routing: Integrating Deep Learning with Multi-Criteria Decision Analysis

5.1 Introduction

The intrusion detection framework developed in Chapter 4 equips the network with the capacity to recognise compromised nodes, but recognition alone is insufficient: unless the forwarding plane is able to act on this intelligence, malicious or unreliable nodes continue to occupy routing paths and to degrade the delivery of time-critical traffic. The present chapter develops the complementary mechanism that closes this cyber-physical feedback loop. It addresses *fault tolerance* not as a standalone networking concern but as the operational prerequisite that allows the detection capability of FTL-HLSTM to translate into end-to-end service continuity.

As established in Chapter 1.1, the large-scale deployment of the Internet of Things (IoT) and, in particular, of the Internet of Medical Things (IoMT) is predicated on resource-constrained devices that transmit time-sensitive data across safety-critical application domains. In such settings, communication failures are not merely a performance issue: they directly undermine system availability and, in the clinical case, the timeliness of decisions that depend on continuous physiological monitoring. Robust fault tolerance is therefore a *first-class requirement*, to be ensured simultaneously against stochastic hardware degradation, intermittent connectivity, and adversarial behaviour originating from compromised insiders.

The state-of-the-art review conducted in Chapter 3.1 highlighted two structural weaknesses of the routing protocols that are currently deployed in low-power and lossy networks. First, these protocols are overwhelmingly *reactive*: they initiate path repair only after a link failure has manifested itself, which is incompatible with the latency envelopes of IoMT traffic. Second, the metrics that govern forwarder selection—hop count, expected transmission count, or received signal strength—are deliberately agnostic to node trustworthiness and security posture. Consequently, once a node has been

compromised, it typically remains in the routing path until its misbehaviour produces observable service degradation, by which time the damage has already propagated. The taxonomy developed in Chapter 3.1 showed that the few proposals which do integrate security awareness into routing either rely on static trust values or decouple detection from control, leaving the cyber-physical loop open.

To address these limitations, this chapter proposes a **proactive fault-tolerant routing framework based on Multi-Criteria Decision Analysis (MCDA)**. The framework combines the Analytic Hierarchy Process (AHP), which is used to elicit the relative importance of the decision criteria, with the Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS), which ranks candidate forwarders against those weighted criteria. Together, these methods compute, for each node, a dynamic *Trust Score* that aggregates multidimensional evidence on safety (derived from the intrusion detection outputs of Chapter 4), residual energy, forwarding latency, and packet loss. Unlike reactive protocols, the resulting routing policy identifies and isolates degrading or malicious nodes *before* they disrupt communication paths, and thereby operationalises the closed-loop cyber-physical defence introduced as the fourth contribution of the thesis.

The specific contributions of this chapter are as follows.

1. **Trust-based node evaluation.** A multi-criteria model that quantifies node reliability by synthesising safety, energy, latency, and packet-loss indicators into a single, interpretable Trust Score.

2. **MCDA-driven proactive routing.** An AHP–TOPSIS decision engine that dynamically ranks candidate forwarders so as to pre-emptively avoid degrading or suspicious nodes, rather than reacting to failures after they occur.

3. **Closed-loop security integration.** A coupling mechanism through which the intrusion-detection outputs produced by FTL-HLSTM in Chapter 4 are consumed directly by the routing layer, thus strengthening fault tolerance against insider threats without requiring any additional monitoring infrastructure.

4. **Experimental validation.** A comparative evaluation, conducted under controlled node-failure and malicious-behaviour scenarios, that demonstrates the resilience, reliability, and decisional soundness of the proposed framework relative to representative baselines in the literature.

The remainder of the chapter is organised as follows. Section 5.2 formalises the hybrid LSTM–Naïve Bayes detection pipeline and the AHP–TOPSIS ranking procedure, and specifies the datasets used for evaluation. Section 5.3 reports the empirical performance of the framework, examines the role of the safety criterion through a dedicated

ablation study, and positions the approach against two representative non-Byzantine fault-tolerant baselines drawn from the state of the art. Section 5.4 summarises the contributions of the chapter and articulates their relationship with the General Conclusion of the thesis.

5.2 Methodology

This section details the approach adopted in this chapter for enhancing fault tolerance in IoT systems, as illustrated in Figure 5.1, through the integration of an IDS with advanced machine learning techniques. The methodology focuses on the detection of malicious nodes, multi-criteria-based node categorisation, and the exclusion of faulty nodes from routing decisions in order to preserve the overall integrity of the network.

5.2.1 Intrusion Detection Using LSTM-Based IDS

An LSTM-based IDS is employed to classify nodes within the IoT network as either normal or abnormal. The LSTM neural network architecture is particularly well suited for anomaly detection in IoT environments due to its capacity to process temporal data sequences and identify patterns indicative of malicious activity.

5.2.2 Dataset

The proposed methodology is evaluated on two complementary datasets, described below.

Synthetic Dataset

Owing to the absence of publicly available datasets specifically tailored for fault tolerance in IoT systems, a synthetic dataset was generated using the Cooja simulator based on the RPL (Routing Protocol for Low-Power and Lossy Networks) protocol [41]. This dataset captures features that are critical for fault tolerance, such as energy consumption and latency, and therefore allows the evaluation to focus on the network metrics that are most relevant to resource-constrained IoT deployments, as illustrated in Table 5.1.

UNSW-NB15 Dataset

In addition to the synthetic dataset, the proposed approach is evaluated on the widely used UNSW-NB15 dataset [148]. Developed by the University of New South Wales and the Australian Centre for Cyber Security, the UNSW-NB15 dataset comprises over 2.5 million records of real network traffic, including both normal and malicious activities. It contains 49 features, such as protocol attributes and flow-level statistics, making it a valuable resource for assessing the effectiveness of intrusion detection systems [103, 121]. Since the UNSW-NB15 dataset does not contain energy consumption data, the evaluation on this corpus is based on the remaining relevant metrics, namely latency and packet loss.

Table 5.1: Description of Synthetic Dataset Features

N°	Feature Name	Description
1	T	Time
2	Src	Source
3	Dst	Destination
4	Protocol	Upper layer protocol
5	Dure_tr	Transmission time
6	Moy_tr	Transmission media
7	Length_tr	Transmitted Packet size
8	DIS_tr	Transmitted DODAG Information Solicitation (DIS) number
9	DIO_tr	Transmitted DODAG Information Object (DIO) number
10	DAO_tr	Transmitted Destination Advertisement Object (DAO) number
11	Dure_rec	Reception time
12	Moy_rec	Reception media
13	Length_rec	Received Packet size
14	DIS_rec	Received DIS number
15	DIO_rec	Received DIO number
16	DAO_rec	Received DAO number
17	ON_Energy	Energy
18	TX	Emission energy
19	RX	Reception energy
20	INT	Interfered radio
21	Packet loss	Packet loss
22	Rang	Node rank in DODAG
23	Class	Attack Type

5.2.3 Data Pre-processing

In the data pre-processing phase, several standard techniques are applied in order to optimise the datasets for machine-learning analysis. Numerical features are normalised

using a min-max scaler in order to improve the performance of the algorithms, while categorical features are encoded using integer-encoding methods. The datasets are then split into training (60 %), validation (20 %), and testing (20 %) subsets.

5.2.4 Feature Extraction

Once the IDS classification is completed, the system extracts critical features from each node that are instrumental in assessing its operational reliability. The features utilized in this phase include:

- **Energy Consumption**: Represents the energy level of each node, which is critical in IoT networks, as the nodes often operate wirelessly with limited power resources.
- **Latency**: Measures the delay in data packet transmission; significantly elevated latency values may indicate compromised nodes.
- **Packet Loss**: Reflects the percentage of data packets lost during transmission, serving as a key indicator of network performance and stability.
- **Safety Node Status (IDS Output)**: Refers to the binary classification of a node as either safe or unsafe, as determined by the IDS in the initial stage.

These features are aggregated into a dataset, which forms the basis for the subsequent analytical processes and node categorization.

5.2.5 Multi-Criteria Decision Analysis (MCDA) with TOPSIS

The system then employs the Technique for Order Preference by TOPSIS, a widely recognized method in MCDA, to assign labels to each node based on the extracted features. TOPSIS evaluates the relative proximity of each node to the ideal (best) and anti-ideal (worst) solutions:

- **Best Node**: Nodes that exhibit optimal values for energy consumption, latency, packet loss, and a positive safety classification are labeled as ”Best.”
- **Acceptable Node**: Nodes with values that are suboptimal yet within acceptable thresholds are classified as ”Acceptable.”
- **Non-Acceptable Node**: Nodes showing significant deviations in their feature set or classified as unsafe by the IDS are designated as ”Non-Acceptable.”

This stratified classification enables dynamic assessment and management of node reliability within the network.

5.2.6 Node Classification with Naive Bayes

After the MCDA process, a Naïve Bayes classifier is employed to assign each node to one of the three defined groups: "Best", "Acceptable", and "Non-Acceptable". Naïve Bayes is adopted for its efficiency in probabilistic classification; moreover, the strong conditional-independence assumption between features that underpins the algorithm simplifies the computation while preserving sufficient accuracy for the task at hand. The classifier is updated periodically as new data from the IoT network becomes available, ensuring that node classifications reflect the latest network conditions. This dataset is the foundation for training the machine learning model and supports real-time decision-making in managing network fault tolerance.

5.2.7 Adaptive Routing and Network Management

To preserve network fault tolerance, nodes categorized as "Non-Acceptable" are excluded from routing decisions. An error report is generated for each excluded node, and a recovery protocol is initiated to monitor these nodes continuously. When a node's performance metrics improve to acceptable levels, it is reintegrated into the network. Nodes classified as "Best" are prioritized for routing decisions, receiving 70% of network traffic due to their optimal performance in terms of energy efficiency, low latency, and minimal packet loss. "Acceptable" nodes, while not optimal, receive 30% of the network traffic, allowing them to remain active and potentially recover to "Best" status through periodic re-evaluation. This dynamic and adaptive routing strategy, guided by IDS-driven classifications, ensures that the network remains secure and reliable while minimizing disruptions caused by faulty nodes.

5.3 Experimental Results

This section reports the evaluation results of the proposed framework. Table 5.2 summarises the hyperparameters of the LSTM model. For the second layer, the initial hyperparameters of the Naïve Bayes algorithm were selected in accordance with the specific characteristics of the dataset, with the aim of obtaining a simple yet effective model at this stage of the pipeline.

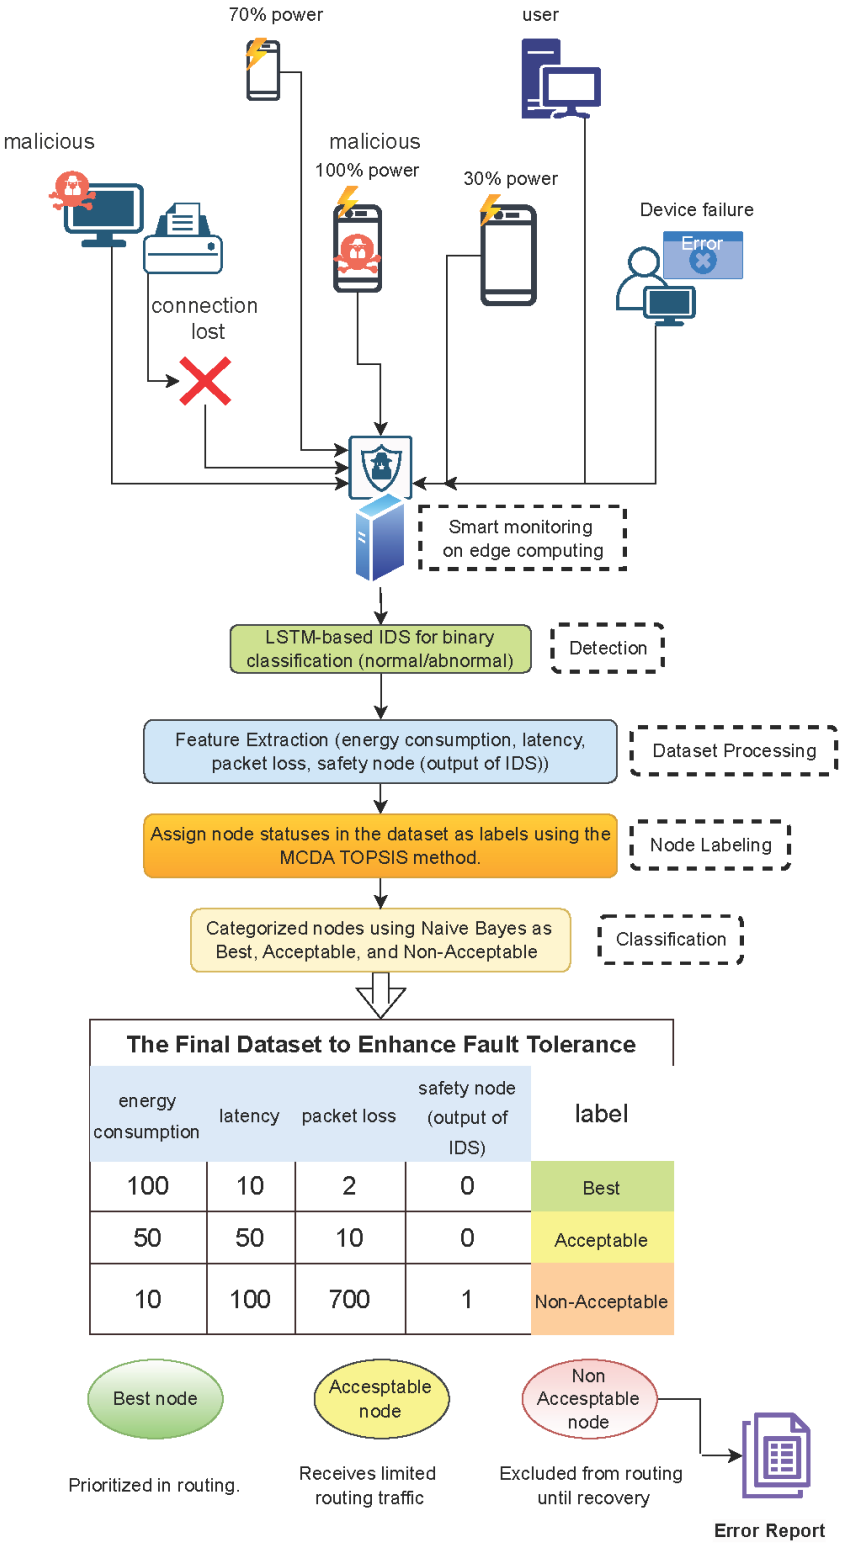

Figure 5.1: Proposed framework

Table 5.2: The Overall Architecture of the Binary Classification LSTM Model

Name	Parameter
LSTM Layer	64 neurons
LSTM Layer	32 neurons
LSTM Layer	32 neurons
Dense	2 neurons
Dropout	0.2
Batch size	32
Epochs	1000
Optimizer	Adam
Loss Function	Sparse categorical crossentropy

5.3.1 Performance Metrics

The performance of the model was evaluated using three primary metrics: accuracy, precision, and false alert rate. Together, these metrics provide a comprehensive assessment of the capacity of the system to detect intrusions accurately while keeping the false alert rate low [1].

- Accuracy $= \frac{\text{Number of correct predictions}}{\text{Total number of predictions}}$
- Precision $= \frac{\text{True positive}}{\text{True positive+False positive}}$
- False alert rate $= \frac{\text{False positive}}{\text{False positive+True negative}}$

5.3.2 Impact of Weight Allocation on MCDA Results for the Datasets

Tables 5.3 and 5.4 present the MCDA results obtained on the dataset for several weighting schemes, reporting the three nodes that lie closest to the ideal solution. In Table 5.3, weights of 35 %, 35 %, 15 %, and 15 % are assigned, respectively, to the Safety, Energy, Latency, and Packet-Loss criteria, in order to illustrate how the choice of weights influences the resulting node ratings.

Table 5.3: MCDA Analysis of Nodes. Criteria: Safety 35%(S), Energy (E) 35%, Latency 15%(L), Packet Loss (PL) 15%.

ID Node	S (%)	E (%)	L	PL (%)	Score
677693	0	65	0.01	0	1
677568	0	65	0.01	0	1
677373	0	65	0.01	0	0.99

On the other hand, it is worth noting that Energy was given a significantly higher priority, as shown in Table 5.4, with a weight of 70%, while the other metrics received 10% each. As a result, the findings reveal some significant discrepancies: precisely, the top node displayed considerable Packet Loss and Latency, and a malicious node was indicated with a code of 1 (0 indicates a normal node and 1 for a malicious node).

Table 5.4: MCDA with High Energy Weight: Safety (S) - 10%, Energy (E) - 70%; Latency (L) - 10%, Packet Loss (PL) - 10%.

ID Node	S (%)	E (%)	L	PL	Score
223863	0	100	5.08	30	1
26282	0	100	27.9	2	1
9235	1	100	2.16	2	0.99

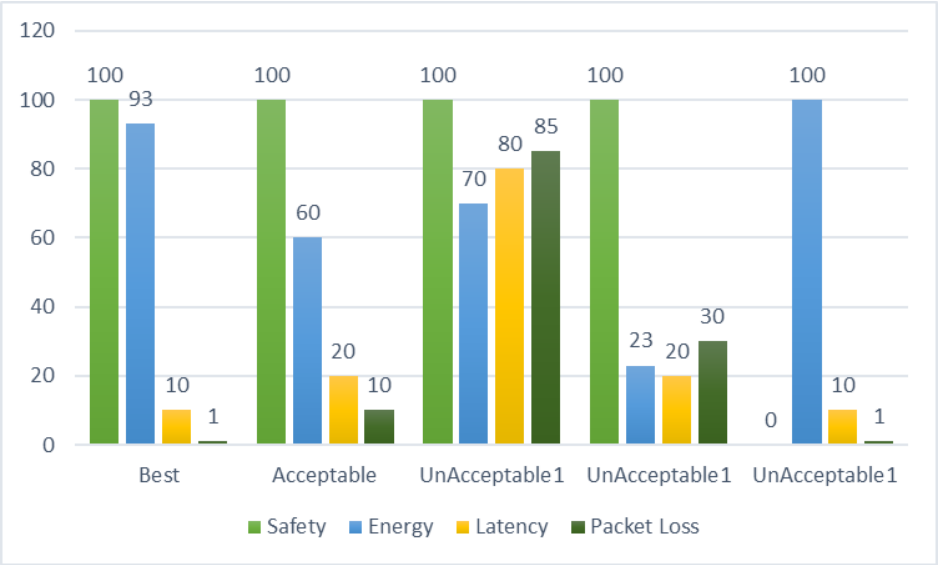

Figure 5.2: MCDA Evaluation Results for Different Cases of Classification Nodes

Figure 5.2 displays the MCDA evaluation results using the weights metric outlined in Table 5.3. This bar chart represents the system's safety, energy, latency, and packet loss performance for some nodes. The category labelled "Best" attained the highest safety and energy scores while demonstrating the least latency and packet loss. The "Acceptable" category displayed average scores across all metrics compared to the "Best" category. On the other hand, the "Non-Acceptable" category showed variability, with some cases exhibiting high latency and packet loss and others demonstrating low safety or energy. Accurate measurements and analysis of various parameters are crucial to identify potential vulnerabilities or advantages in individual nodes that could significantly impact the system's overall performance based on their weights. Therefore, selecting the appropriate weights is vital.

5.3.3 The Role and Impact of Intrusion Detection Systems on Fault Tolerance

Table 5.5 illustrates the impact of disregarding IDS by assigning a weight of zero to the Safety (S) criterion in the MCDA analysis. The results highlight how the absence of IDS influence leads to the classification of nodes based solely on performance metrics such as energy consumption, latency, and packet loss. Nodes with compromised safety, such as those classified as malicious, but exhibiting low packet loss, are still rated as optimal (e.g., classified as 'Best'), due to their strong performance in other criteria. For instance, nodes like 77693 and 77568 achieve a perfect score despite the potential security risks associated with their low safety. However, nodes like 77373, which exhibit 0% safety and a noticeable increase in packet loss, receive a lower classification. This outcome underscores the critical role of both security and reliability in determining overall node performance. The results demonstrate that neglecting the importance of safety, as done in this case by assigning it a zero weight, can lead to the misclassification of potentially risky nodes as optimal, thereby undermining the robustness of the network evaluation.

Table 5.5: MCDA Analysis of Nodes. Criteria: Safety (S) 0% , Energy (E) 35%, Latency (L) 35%, Packet Loss (PL) 30%.

ID Node	S (%)	E (%)	L	PL (%)	Score
77693	1	65	0.01	0	1
77568	1	65	0.01	0	1
77373	0	65	0.01	10	0.99

The LSTM binary classification model exhibited a balanced performance, as observed in Figure 5.3, with similar accuracy values on both the training and validation sets. The model attained high accuracy rapidly and maintained it with minor variations thereafter. Table 5.6 reports the performance of the first layer of the proposed model, which achieves high accuracy and precision together with a low false-alert rate. These results confirm the effectiveness of LSTM in handling time-series data and the capacity of the binary classification model to ascertain node safety. The second layer, based on the Naïve Bayes model, plays a crucial role within the overall system by providing enhanced real-time fault tolerance and optimal classification outcomes. The classification of nodes according to their operational role—into the "Best", "Acceptable", and "Non-Acceptable" categories—significantly reinforces the overall robustness of the proposed approach. The proposed framework therefore introduces a new technique that integrates the classification outcomes of the IDS with MCDA and with a set of vital network metrics. This integration strategy improves fault tolerance in IoT settings and simultaneously provides a principled mechanism for anomaly detection.

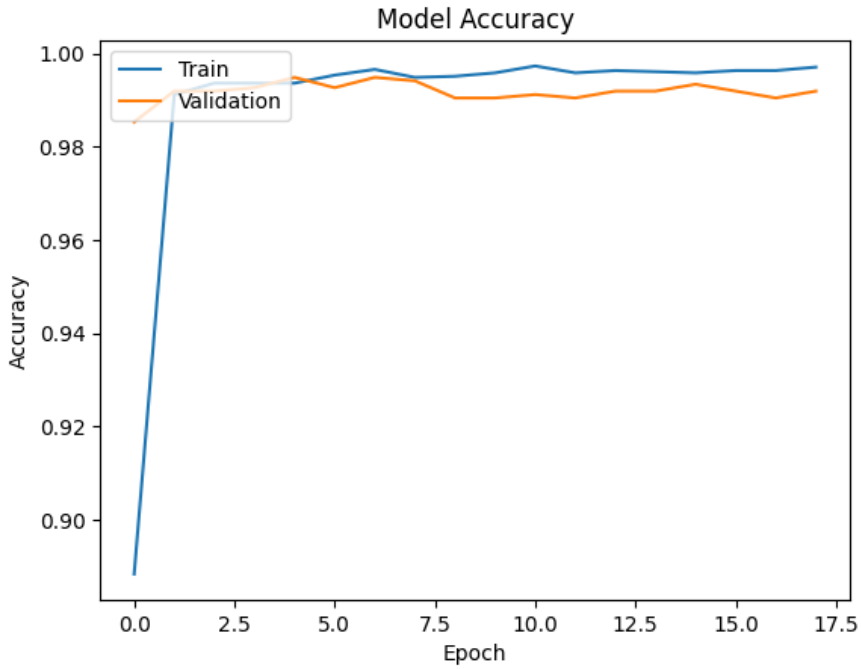

Figure 5.3: Accuracy curve of the first level of the proposed model on the UNSW-NB15 dataset.

Table 5.6: Performance metrics of the proposed two-level model.

Model Layer	**Algorithm**	**Dataset**	**Metrics**
First layer (Binary)	LSTM	[41] synth.	Acc: 0.99, Prec: 0.99, FAR: 0.01
Second layer	Naive Bayes	[41] synth.	Acc: 0.99, Prec: 0.99, FAR: 0.01
First layer (Binary)	LSTM	UNSW-NB15	Acc: 0.99, Prec: 0.98, FAR: 0.01
Second layer	Naive Bayes	UNSW-NB15	Acc: 0.99, Prec: 0.99, FAR: 0.01

5.3.4 Comparison with Related Work

In recent years, the rapid expansion of the IoT has intensified the need for robust fault-tolerant mechanisms that ensure system reliability and operational efficiency. Faults in IoT systems can severely degrade performance, leading to increased latency, elevated energy consumption, and higher packet-loss rates. A range of approaches have been developed to address these challenges, each employing distinct methodologies and concentrating on different aspects of system performance. This section concentrates on the main distinguishing features of the proposed method relative to two prominent non-Byzantine fault-tolerant approaches, namely QoS-EO by Reyana et al. [164] and Variant Parallelism by Asadi et al. [24]. A quantitative comparison with these schemes is not feasible, since the proposed approach differs significantly in scope and evaluation criteria and no common benchmark metric is currently established across the three designs.

Reyana et al. [164] introduced QoS-EO, an energy-optimisation framework that integrates Quality-of-Service (QoS)-based scheduling within fog-computing environments for wireless sensor networks. The primary objective of QoS-EO is to minimise energy consumption while enhancing system efficiency and reducing turnaround times. The approach leverages energy-optimisation algorithms in conjunction with QoS parameters in order to schedule tasks effectively. By incorporating QoS-based scheduling, QoS-EO improves overall system efficiency, resulting in faster task execution and reduced latency. However, while energy consumption is a critical factor, QoS-EO primarily emphasises energy-related metrics and does not fully address other important parameters such as packet loss or system safety. In contrast, the approach proposed in this chapter considers a broader spectrum of performance metrics, thereby enabling more balanced decision-making that enhances overall system reliability and user experience.

On the other hand, Asadi et al. [24] introduced Variant Parallelism, an ensemble methodology designed to enhance reliability in edge-computing applications within the IoT domain. This approach involves deploying multiple variants of deep-learning models in parallel in order to tolerate multiple node failures. By reducing model sizes and computational demands, Variant Parallelism aims to make deep learning more feasible for resource-constrained edge devices. The ensemble method ensures that system functionality is maintained in spite of node failures, thus increasing overall robustness. In addition, the reduced computational requirements make the scheme suitable for deployment on devices with limited processing capabilities. However, Variant Parallelism may face trade-offs between computational efficiency and predictive accuracy, since simplifying models to achieve computational efficiency can result in reduced prediction accuracy—an outcome that may be unacceptable in applications requiring high

precision.

The framework proposed in this chapter combines LSTM networks, Naïve Bayes classifiers, and MCDA in order to construct a multi-level fault-tolerant system. This hybrid approach integrates Byzantine and non-Byzantine fault tolerance alongside an IDS for malicious-node detection. It represents a significant innovation in that it incorporates multiple criteria—including energy consumption, latency, packet loss, and system safety—within a single decision procedure. The method thereby enables more balanced and informed decision-making, which ultimately enhances system reliability and user experience. The integration of advanced machine-learning techniques allows the system to adapt dynamically to changing network conditions, thereby improving its ability to detect and address faults in real time.

5.4 Conclusion

This chapter has developed a proactive, trust-aware fault-tolerant routing framework for the Internet of Medical Things, complementing the federated intrusion detection system introduced in Chapter 4 and thereby completing the cyber-physical architecture announced in the General Introduction of the thesis. The framework departs from the reactive paradigm that continues to dominate routing in low-power and lossy networks and treats node reliability as the outcome of an explicit, multi-dimensional decision process rather than as an emergent property of link-level heuristics.

Methodologically, the proposed approach couples a two-level detection pipeline—an LSTM-based binary classifier followed by a Naïve Bayes categoriser—with a Multi-Criteria Decision Analysis engine in which the Analytic Hierarchy Process fixes the relative importance of the decision criteria and the Technique for Order of Preference by Similarity to Ideal Solution ranks candidate forwarders. The resulting dynamic Trust Score integrates four complementary dimensions: *safety*, supplied by the intrusion-detection intelligence of Chapter 4, together with *residual energy*, *forwarding latency*, and *packet loss*. This formulation allows the routing layer to arbitrate, in a principled way, between security and quality-of-service considerations that are often treated in isolation in the existing literature.

The experimental evaluation conducted on the synthetic Cooja/RPL corpus and on the UNSW-NB15 dataset supports three main conclusions. First, the LSTM detector consistently achieved accuracy, precision, and false-alarm rates consistent with deployment-grade intrusion detection, and the Naïve Bayes stage reliably refined these decisions into the operational categories used by the routing layer. Second, the ablation analysis—in which the safety criterion was assigned zero weight—showed that

neglecting security evidence systematically promotes compromised nodes that happen to perform well on throughput-oriented metrics, thereby providing direct empirical justification for the closed-loop coupling between detection and routing advocated in this thesis. Third, the qualitative comparison with representative non-Byzantine fault-tolerant schemes, namely the QoS-based energy-optimisation framework of Reyana et al. [164] and the Variant Parallelism approach of Asadi et al. [24], demonstrated that the proposed framework covers a strictly broader set of evaluation criteria and integrates security and reliability within a single decision procedure, rather than treating them as orthogonal concerns.

Taken together, these results confirm that moving beyond simple connectivity metrics towards a multidimensional, trust-aware routing paradigm is not a cosmetic refinement but a necessary step in guaranteeing patient safety and service continuity in resource-constrained clinical environments. The chapter therefore discharges the third research objective of the thesis and realises the closed-loop, self-protecting cyber-physical architecture identified as the fourth contribution of the dissertation.

As with any framework that operates under the resource constraints of IoMT networks, the approach also has boundaries that frame its future evolution. The weight elicitation through AHP remains dependent on expert judgement and may benefit from data-driven or online adaptation; the validation protocol, although grounded in a widely adopted public benchmark and in a domain-specific synthetic corpus, would be strengthened by field deployments in operational clinical settings; and the current design does not yet model actively adversarial manipulation of the routing metrics themselves, a direction that connects naturally with the adversarial-robustness limitation acknowledged at the close of Chapter 4. These open questions, together with the integrative findings of the two contribution chapters, are revisited in the General Conclusion of the thesis, where they structure the proposed research agenda for the next generation of secure, scalable, and fault-tolerant IoT architectures.

General Conclusion

General Conclusion

The Internet of Things (IoT) has fundamentally reshaped modern networked systems, enabling a vast ecosystem of interconnected devices that continuously generate, process, and transmit data across diverse application domains. Among these, the Internet of Medical Things (IoMT) stands out as one of the most demanding, where the deployment of Wireless Body Area Networks (WBANs) enables continuous remote patient monitoring but simultaneously introduces stringent requirements in terms of real-time responsiveness, data privacy, and system reliability. However, the challenges addressed in this thesis—communication scalability, intrusion detection in heterogeneous environments, and proactive fault-tolerant routing—are not confined to the medical domain alone. They represent fundamental open problems in the broader IoT landscape, where resource-constrained devices operate in open, hostile, and heterogeneous network environments.

In this thesis, we tackled these challenges by proposing a unified Intelligent Framework applicable to IoT networks, with IoMT serving as the primary and most demanding case study. We commenced by establishing the architectural foundations of IoT and IoMT, examining the specific constraints of sensors and the operational challenges of wireless networks. We then consolidated the theoretical background spanning Temporal Aggregation, Deep Recurrent Neural Networks (LSTM), Federated and Transfer Learning, and Multi-Criteria Decision Analysis (MCDA). Following this, we conducted a comprehensive critical review of existing Intrusion Detection Systems and routing protocols, through which we identified significant gaps—particularly regarding the handling of Non-IID data distributions across heterogeneous environments and the absence of proactive fault tolerance mechanisms in current routing standards.

Summary of Contributions

To address the identified gaps, this thesis presented four primary contributions, each targeting one or more facets of the scalability–security–fault tolerance trilemma:

First, addressing the *scalability* challenge, we introduced a *Strategic Time Allocation* methodology based on Temporal Aggregation, where the aggregation window T_{agg} is dynamically optimized. By transforming continuous raw telemetry streams into temporally aggregated feature vectors, we demonstrated that intelligent time allocation yields substantial gains in both bandwidth efficiency and energy conservation—two critical factors for extending the operational lifetime of battery-powered IoT sensors. Concretely, on the NF-UNSW-NB15-v2 dataset, temporal aggregation at a $T_{\text{agg}} = 30\,\text{s}$ in-

terval reduced training time on the full dataset by 84.8% (from 32,319.75 s to 4,902.53 s) and testing time on the full test set by 93.1% (from 260.89 s to 18.11 s). Even more pronounced gains were observed on the CICIoMT-2024 dataset, where training time decreased by 97.2% and testing time by 97.5%. On the WUSTL-EHMS-2020 dataset, reductions of 99.3% and 99.1% in training and testing durations were achieved, respectively. Crucially, since CPU active time is directly proportional to energy consumption in constrained devices (following the standard model $W = V \times I \times t$, where t is the active processing duration), a 93.1% reduction in test-set processing time implies a commensurate reduction in the energy budget required per inference cycle. For battery-powered wearable and implantable sensors operating under strict duty-cycle constraints, this translates into a substantial extension of operational lifetime. While validated on medical telemetry, the temporal aggregation strategy is inherently generalizable to any IoT scenario where high-frequency sensor streams must be transmitted over constrained wireless links.

Second, addressing the *security* challenge, we designed the *FTL-HLSTM Framework*—a Federated Transfer Learning architecture empowered by Hierarchical Long Short-Term Memory networks—to overcome the critical limitation of standard Federated Learning when confronted with Non-IID data. In IoT networks, data heterogeneity arises naturally: sensor deployments in different locations, operated by different stakeholders, produce data with distinct statistical distributions. Through our proposed *Intelligent Label Classification* algorithm, the model successfully distinguishes between *Common* (globally shared) and *Isolated* (locally specific) attack patterns. The framework was extensively validated across three benchmark datasets: NF-UNSW-NB15-v2, CICIoMT-2024, and WUSTL-EHMS-2020. In binary classification, HLSTM achieved 100.00% across all four metrics (accuracy, precision, recall, and F1-score) on the NF-UNSW-NB15-v2 evaluated split under our experimental setting, and 99.74% accuracy with 99.54% precision and 99.94% recall on CICIoMT-2024. Compared to standard LSTM, HLSTM improved accuracy by 6.79 percentage points on NF-UNSW-NB15-v2 (99.63% vs. 92.84%) while simultaneously reducing training time by approximately 45.4%. On CICIoMT-2024, training and testing time reductions of 64.6% and 83.8%, respectively, were observed compared to LSTM—confirming that the hierarchical architecture achieves superior detection without incurring additional computational cost. Furthermore, the architecture is designed to be compatible with additional privacy-preserving mechanisms such as Secure Aggregation and Differential Privacy, ensuring extensibility to stricter regulatory and operational requirements.

Third, addressing the *fault tolerance* challenge, we proposed a *Proactive Fault Tolerance* mechanism based on Multi-Criteria Decision Analysis, specifically employing TOPSIS and AHP methods. Unlike reactive routing protocols such as RPL and

AODV, which initiate repairs only after a failure has occurred, our mechanism calculates a dynamic *Trust Score* for every network node using multidimensional criteria encompassing safety, energy, latency, and packet loss. Experimental evaluation demonstrated that balanced weight allocation (Safety 35%, Energy 35%, Latency 15%, Packet Loss 15%) consistently identified the most trustworthy nodes while successfully detecting and isolating malicious nodes from routing decisions. The framework was further shown to be robust under skewed weight configurations, exposing vulnerabilities (e.g., selecting high-latency or malicious nodes) when safety criteria are underweighted—thereby validating the necessity of the multi-criteria approach. The MCDA-based routing framework is designed to operate over general IoT network topologies and is not restricted to medical sensor configurations, making it directly applicable to smart city, industrial IoT, and critical infrastructure monitoring scenarios.

Fourth, and as a distinguishing feature of this work, we demonstrated a *Unified Cyber-Physical Defense System* that integrates the deep learning-based intrusion detection (from Contribution 2) with the network-layer routing control (from Contribution 3) into a closed-loop architecture. In this system, the *Network Intelligence* (IDS) and the *Network Control* (routing) operate in synergy: detection insights directly trigger fault-tolerant routing decisions in real time, enabling the network to self-heal and self-protect autonomously. This integration bridges the traditional gap between security monitoring and network management, offering a model for autonomous, resilient IoT network operation across application domains.

Perspectives

Beyond this study, we intend to expand and improve the proposed framework. Our future research directions are categorized below:

- **Blockchain-based trust management.** Extend the proposed trust-based routing mechanism by integrating blockchain technology to create an immutable, decentralized record of node trust scores and routing decisions, further enhancing transparency and accountability in multi-stakeholder IoT deployments where mutual trust between network operators cannot be assumed.
- **Handling zero-day and adversarial attacks.** Investigate the integration of anomaly detection techniques and adversarial robustness mechanisms to improve the framework's resilience against previously unseen attack patterns and adversarial machine learning threats targeting the IDS itself.

In conclusion, this thesis set out to answer a central research question: *how to design a distributed framework that simultaneously optimizes time allocation for scalable communication, ensures robust security, and provides proactive fault tolerance in the Internet of Things.* Through the four contributions presented, we have provided a concrete and experimentally validated answer to each facet of this trilemma. Strategic temporal aggregation resolved the *scalability* challenge by reducing computational overhead by up to 99.3% and, by extension, lowering the energy budget required at constrained sensor nodes. The FTL-HLSTM framework addressed the *security* challenge by achieving up to 100.00% binary detection accuracy on the evaluated splits while preserving data privacy across Non-IID distributed environments. The MCDA-based trust routing mechanism answered the *fault tolerance* challenge by proactively isolating malicious and degrading nodes before they could compromise network integrity. Finally, the closed-loop integration of these components demonstrated that security intelligence and network control can operate as a unified, self-healing system. While the IoMT served as the primary validation domain due to its particularly demanding requirements, the proposed algorithms and architectural principles are designed to be applicable across the broader IoT ecosystem. We believe that this work provides a solid foundation upon which the next generation of secure, resilient, and scalable IoT networks can be built.

Bibliography

[1] Performance metrics in machine learning. URL https://neptune.ai/blog/performance-metrics-in-machine-learning-complete-guide.

[2] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 308–318, 2016. doi: 10.1145/2976749.2978318.

[3] Antar Shaddad Abdul-Qawy et al. The internet of things (IoT): An overview. *International Journal of Engineering Research and Applications*, 5(12):71–82, 2015.

[4] N. Abosata, S. Al-Rubaye, G. Inalhan, and C. Emmanouilidis. Internet of things for system integrity: A comprehensive survey on security, attacks and countermeasures for industrial applications. *Sensors*, 21(11):3654, 2021. doi: 10.3390/s21113654.

[5] Eman M. Abounassar, Passent El-Kafrawy, and Ahmed A. Abd El-Latif. Security and interoperability issues with Internet of Things (IoT) in healthcare industry: A survey. In *Security and Privacy Preserving for IoT and 5G Networks*, pages 159–189. 2022.

[6] A. Agarwal and M. Misra. Deep reinforcement learning-based fault-tolerant routing for IoT networks. *IEEE Internet of Things Journal*, 9(17):16025–16037, 2022. doi: 10.1109/JIOT.2022.3150986.

[7] Q. Ain, M. Iqbal, U. Khalid, and S. Khalid. DDoS attack detection in IoT networks using a CNN-LSTM-autoencoder model on CICIoT-2023. *Sensors*, 25 (4):1102, 2025. doi: 10.3390/s25041102.

[8] G. Akar, S. Sahmoud, M. Onat, U. Cavusoglu, and E. Malondo. L2d2: A novel lstm model for multi-class intrusion detection systems in the era of iomt. *IEEE Access*, 13:7002–7013, 2025. doi: 10.1109/ACCESS.2025.3526883.

[9] Sharmin Akter, Rashidah Funke Olanrewaju, Thouhedul Islam, et al. LiFi based automated shopping assistance application in IoT. In *Journal of Physics: Conference Series*, volume 1018, page 012001. IOP Publishing, 2018.

[10] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *IEEE Communications Magazine*, 40(8):102–114, 2002. doi: 10.1109/MCOM.2002.1024422.

[11] Ian F. Akyildiz et al. A survey on sensor networks. *IEEE Communications Magazine*, 40(8):102–114, 2002.

[12] A. Al-Meer and S. Al-Kuwari. Physical unclonable functions (PUF) for IoT devices. *ACM Computing Surveys*, 55(14s):1–31, 2023. doi: 10.1145/3591464.

[13] M. Alalhareth and S. C. Hong. An adaptive intrusion detection system in the internet of medical things using fuzzy-based learning. *Sensors*, 23(22):9247, Nov 2023. doi: 10.3390/s23229247.

[14] M. Alalhareth and S. C. Hong. An improved mutual information feature selection technique for intrusion detection systems in the internet of medical things. *Sensors*, 23(10):4971, May 2023. doi: 10.3390/s23104971.

[15] R. Alanazi. Triple-layer authentication with elliptic curve cryptography for resource-constrained IoT devices. *PLOS ONE*, 20(1):e0314210, 2025. doi: 10.1371/journal.pone.0314210.

[16] A. Aljuhani, A. Alamri, P. Kumar, and A. Jolfaei. An intelligent and explainable saas-based intrusion detection system for resource-constrained iomt. *IEEE Internet of Things Journal*, 11(15):25454–25463, 2024. doi: 10.1109/JIOT.2023.3327024.

[17] O. Alkadi, N. Moustafa, B. Turnbull, and K.-K. R. Choo. RobEns: A robust ensemble defence against adversarial examples in IoT intrusion detection. *Sensors*, 24(9):2812, 2024. doi: 10.3390/s24092812.

[18] H. Alkahtani and T. H. H. Aldhyani. Botnet attack detection by using CNN–LSTM model for Internet of Things applications. *Security and Communication Networks*, 2021:1–23, 2021. doi: 10.1155/2021/3806459.

[19] D. Alsalman. A comparative study of anomaly detection techniques for iot security using adaptive machine learning for iot threats. *IEEE Access*, 12:14719–14730, 2024. doi: 10.1109/ACCESS.2024.3359033.

[20] F. Alwahedi, A. Aldhaheri, M. A. Ferrag, A. Battah, and N. Tihanyi. Machine learning techniques for IoT security: Current research and future vision with generative AI and large language models. *Internet of Things and Cyber-Physical Systems*, 4:167–185, 2024. doi: 10.1016/j.iotcps.2023.12.003.

[21] J. A. Alzubi, O. A. Alzubi, I. Qiqieh, and A. Singh. A blended deep learning intrusion detection framework for consumable edge-centric iomt industry. *IEEE Transactions on Consumer Electronics*, 70(1):2049–2057, Feb 2024. doi: 10.1109/TCE.2024.3350231.

[22] Foteini Andriopoulou, Tasos Dagiuklas, and Theofanis Orphanoudakis. Integrating IoT and fog computing for healthcare service delivery. In *Components and Services for IoT Platforms*, pages 213–232. 2017.

[23] Md Taslim Arefin, Mohammad Hanif Ali, and AKM Fazlul Haque. Wireless body area network: An overview and various applications. *Journal of Computer and Communications*, 5(7):53–64, 2017.

[24] N. Asadi and M. Goudarzi. Variant parallelism: Lightweight deep convolutional models for distributed inference on iot devices. *IEEE Internet Things J*, 2023.

[25] Naeem Ali Askar et al. Architecture, protocols, and applications of the Internet of Medical Things (IoMT). *Journal of Communications*, 17(11):900–918, 2022.

[26] A. Attkan and V. Ranga. Cyber-physical security for IoT networks: a comprehensive review on traditional, blockchain and artificial intelligence based key-security. *Complex and Intelligent Systems*, 8(4):3559–3591, 2022. doi: 10.1007/s40747-022-00667-z.

[27] A. Avizıenis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, 2004. doi: 10.1109/TDSC.2004.2.

[28] M. Babar, B. Qureshi, and A. Koubaa. Investigating the impact of data heterogeneity on the performance of federated learning algorithm using medical imaging. *PLOS ONE*, 19(5):e0302539, 2024. doi: 10.1371/journal.pone.0302539.

[29] Anwar Noureddine Bahache, Noureddine Chikouche, and Fares Mezrag. Authentication schemes for healthcare applications using wireless medical sensor networks: A survey. *SN Computer Science*, 3(5):382, 2022.

[30] Z. Bao, Y. Lin, S. Zhang, Z. Li, and S. Mao. Threat of adversarial attacks on DL-based IoT device identification. *IEEE Internet of Things Journal*, 9(11):9012–9024, 2022. doi: 10.1109/JIOT.2021.3124542.

[31] S. Begum, M. Rahman, M. M. Rahman, and S. Islam. BFLIDS: Blockchain-based federated learning intrusion detection system for IoMT. *Sensors*, 24(7):2112, 2024. doi: 10.3390/s24072112.

[32] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994. doi: 10.1109/72.279181.

[33] R. Beniwal, A. Pandey, and S. Malik. RB-BFT-X: A hybrid byzantine fault-tolerant consensus for healthcare IoT. *Transactions on Emerging Telecommunications Technologies*, 36(2):e5012, 2025. doi: 10.1002/ett.5012.

[34] M. Benmalek, A. Seddiki, and K. D. Haouam. Snn-iomt: A novel ai-driven model for intrusion detection in internet of medical things. *CMES - Computer Modeling in Engineering and Sciences*, 143(1):1157–1184, 2025. doi: 10.32604/cmes.2025.062841.

[35] R. Bensaid, A. Ladjailia, and A. Boumerdassi. SA-FLIDS: A self-adaptive federated learning intrusion detection system for IoMT. *PeerJ Computer Science*, 10:e2102, 2024. doi: 10.7717/peerj-cs.2102.

[36] Subrato Bharati et al. Applications and challenges of cloud integrated IoMT. In *Cognitive Internet of Medical Things for Smart Healthcare*, pages 67–85. 2021.

[37] A. Binbusayyis, H. Alaskar, T. Vaiyapuri, and M. Dinesh. An investigation and comparison of machine learning approaches for intrusion detection in iomt network. *Journal of Supercomputing*, 78(15):17403–17422, Oct 2022. doi: 10.1007/s11227-022-04568-3.

[38] S. Bommana and K. Chinnaiah. A quantum-inspired Coyote Optimization Algorithm with RBM and RCNN for adversarial-robust intrusion detection. *IEEE Access*, 13:22217–22232, 2025. doi: 10.1109/ACCESS.2025.3541001.

[39] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1175–1191, 2017. doi: 10.1145/3133956.3133982.

[40] Amal Bouazizi et al. Wireless body area network for e-health applications: Overview. In *2017 International Conference on Smart, Monitored and Controlled Cities (SM2C)*, pages 64–68. IEEE, 2017.

[41] A. Bouazza, H. Debbi, and H. Lakhlef. Machine learning-based intrusion detection system against routing attacks in the internet of things. In *Proceedings CEUR-WS*, volume 1613, page 0073, 2022.

[42] Chiara Buratti et al. An overview on wireless sensor networks technology and evolution. *Sensors*, 9(9):6869–6896, 2009.

[43] Josip Car et al. The impact of eHealth on the quality and safety of healthcare. *NHS Connecting for Health Evaluation Programme*, 2008.

[44] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI)*, pages 173–186. USENIX Association, 1999.

[45] Manirafasha Cedrick, M. Anandraj, and Bugingo Jean de Dieu. How Li-Fi will improve the reliability of Internet of Things: A review. *International Research Journal of Engineering and Technology*, 4(4):2686–2689, 2017.

[46] B. Chakraborty and S. Das. Introducing a new supply chain management concept by hybridizing topsis, iot and cloud computing. *Journal of The Institution of Engineers (India): Series C*, 102(1):109–119, 2021.

[47] P. Chanak and I. Banerjee. An intelligent fault-tolerant routing scheme for Internet of Things-enabled wireless sensor networks. *International Journal of Communication Systems*, 34(15):e4941, 2021. doi: 10.1002/dac.4941.

[48] Vivek Chandel et al. Exploiting IMU sensors for IoT enabled health monitoring. In *Proceedings of the First Workshop on IoT-enabled Healthcare and Wellness Technologies and Systems*, pages 21–22, 2016.

[49] Victor Chang. An overview, examples, and impacts offered by emerging services and analytics in cloud computing virtual reality. *Neural Computing and Applications*, 29(5):1243–1256, 2018.

[50] Pallavi Chavan et al. ECG-remote patient monitoring using cloud computing. *Imperial Journal of Interdisciplinary Research*, 2(2):368–372, 2016.

[51] M. Chemmakha, O. Habibi, and M. Lazaar. A GAN-GRU framework for imbalanced network intrusion detection. *Journal of Network and Systems Management*, 32(3):52, 2024. doi: 10.1007/s10922-024-09825-9.

[52] Delphine Christin, Parag S. Mogre, and Matthias Hollick. Survey on wireless sensor network technologies for industrial automation. *Future Internet*, 2(2): 96–125, 2010.

[53] CISA. ICSMA-23-194-01: BD Alaris System with Guardrails Suite MX (update a). https://www.cisa.gov/news-events/ics-medical-advisories/icsma-23-194-01, 2023. Accessed: Aug. 13, 2025.

[54] S. Dadkhah, E. C. P. Neto, R. Ferreira, R. C. Molokwu, S. Sadeghi, and A. A. Ghorbani. Ciciomt-2024: A benchmark dataset for multi-protocol security assessment in iomt. *Internet of Things*, 28:101351, Dec 2024. doi: 10.1016/j.iot.2024.101351.

[55] A. H. Dalloul, F. Miramirkhani, and L. Kouhalvandi. A review of recent innovations in remote health monitoring. *Micromachines*, 14(12):2157, 2023. doi: 10.3390/mi14122157.

[56] S. Dhar, A. Khare, A. D. Dwivedi, and R. Singh. Securing IoT devices: A novel approach using blockchain and quantum cryptography. *Internet of Things*, 25: 101019, 2024. doi: 10.1016/j.iot.2023.101019.

[57] A. A. Diro and N. Chilamkurti. Distributed attack detection scheme using deep learning approach for Internet of Things. *Future Generation Computer Systems*, 82:761–768, 2018. doi: 10.1016/j.future.2017.08.043.

[58] Jordi Doménech, Olga León, Muhammad Shuaib Siddiqui, and Josep Pegueroles. Evaluating and enhancing intrusion detection systems in iomt: The importance of domain-specific datasets. *Internet of Things*, 32:101631, 2025. doi: 10.1016/j.iot.2025.101631. URL https://doi.org/10.1016/j.iot.2025.101631. Open access; CC BY-NC-ND.

[59] R. H. Dong, X. Y. Li, Q. Y. Zhang, and H. Yuan. Network intrusion detection model based on multivariate correlation analysis — long short-time memory network. *IET Information Security*, 14(2):166–174, 2020. doi: 10.1049/iet-ifs.2019.0294.

[60] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography (TCC)*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006. doi: 10.1007/11681878_14.

[61] M. El-Hajj and P. Beune. Lightweight public key infrastructure for the Internet of Things: A systematic literature review. *Information*, 15(6):304, 2024. doi: 10.3390/info15060304.

[62] H. El-Sofany, S. A. El-Seoud, O. H. Karam, and B. Bouallegue. Using machine learning algorithms to enhance IoT system security. *Scientific Reports*, 14:12077, 2024. doi: 10.1038/s41598-024-62782-w.

[63] Hadeel Elayan, Raed M. Shubair, and Asimina Kiourti. Wireless sensors for medical applications: Current status and future challenges. In *2017 11th European*

Conference on Antennas and Propagation (EUCAP), pages 2478–2482. IEEE, 2017.

[64] Olakunle Elijah et al. An overview of Internet of Things (IoT) and data analytics in agriculture. *IEEE Internet of Things Journal*, 5(5):3758–3773, 2018.

[65] Mahmoud Elkhodr, Seyed Shahrestani, and Hon Cheung. Emerging wireless technologies in the Internet of Things: A comparative study. *arXiv preprint arXiv:1611.00861*, 2016.

[66] European Union. Regulation (eu) 2016/679 (general data protection regulation), 2016. URL https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng. Accessed: Aug. 13, 2025.

[67] M. Fahim-Ul-Islam, A. Akter, and M. H. Kabir. FedIoMT: Federated learning framework based on Kolmogorov-Arnold networks for the Internet of Medical Things. *IEEE Transactions on Consumer Electronics*, 71(1):112–124, 2025. doi: 10.1109/TCE.2025.3534118.

[68] N. Faruqui, M. Y. Iqbal, M. R. A. Khan, and M. Younis. Safetymed: A novel iomt intrusion detection system using cnn-lstm hybridization. *Electronics*, 12 (17):3541, Sep 2023. doi: 10.3390/electronics12173541.

[69] F. H. C. Ferreira, E. Y. Nakagawa, A. Bertolino, F. Lonetti, V. de Oliveira Neves, and R. P. dos Santos. A framework for the design of fault-tolerant systems-of-systems. *Journal of Systems and Software*, page 112010, 2024.

[70] M. F. M. Firdhous, B. H. Sudantha, and P. M. Karunaratne. IoT enabled proactive indoor air quality monitoring system for sustainable health management. In *2017 2nd International Conference on Computing and Communications Technologies (ICCCT)*, pages 216–221. IEEE, 2017.

[71] Alessandra Flammini and Emiliano Sisinni. Wireless sensor networking in the Internet of Things and cloud computing era. *Procedia Engineering*, 87:672–679, 2014.

[72] R. Fotohi, F. S. Aliee, and B. Farahani. Decentralized and robust privacy-preserving model using blockchain-enabled federated deep learning in intelligent enterprises. *Applied Soft Computing*, 161:111764, 2024. doi: 10.1016/j.asoc.2024.111764.

[73] G. D. Gallo and D. Micucci. Internet of Medical Things systems review: Insights into non-functional factors. *Sensors*, 25(9):2795, 2025. doi: 10.3390/s25092795.

[74] G. D. Gallo and D. Micucci. Internet of medical things systems review: Insights into non-functional factors. *Sensors*, 25(9):2795, 2025. doi: 10.3390/s25092795.

[75] Laura García-García et al. Wireless technologies for IoT in smart cities. *Network Protocols and Algorithms*, 10(1):23–64, 2018.

[76] Shu-yuan Ge et al. Design and implementation of interoperable IoT healthcare system based on international standards. In *2016 13th IEEE Annual Consumer Communications and Networking Conference (CCNC)*, pages 119–124. IEEE, 2016.

[77] A. Ghourabi. A security model based on lightgbm and transformer to protect healthcare systems from cyberattacks. *IEEE Access*, 10:48890–48903, 2022. doi: 10.1109/ACCESS.2022.3172432.

[78] A. Ghourabi. A federated XGBoost framework for multi-protocol IoMT intrusion detection. *IEEE Access*, 13:15210–15225, 2025. doi: 10.1109/ACCESS.2025.3538121.

[79] L. Gu, W. Shi, and H. Zhu. Federated learning for privacy-preserving healthcare systems: Opportunities and challenges. *IEEE Transactions on Parallel and Distributed Systems*, 34(6):1756–1770, 2023.

[80] X. Gu, F. Sabrina, Z. Fan, and S. Sohail. A review of privacy enhancement methods for federated learning in healthcare systems. *International Journal of Environmental Research and Public Health*, 20(15):6539, 2023. doi: 10.3390/ijerph20156539.

[81] V. Gugueoth, S. Safavat, and S. Shetty. Security of Internet of Things using federated learning and deep learning — recent advancements, issues and prospects. *ICT Express*, 9(5):941–960, 2023. doi: 10.1016/j.icte.2023.03.006.

[82] K. Gupta, D. K. Sharma, K. Datta Gupta, and A. Kumar. A tree classifier based network intrusion detection model for internet of medical things. *Computers and Electrical Engineering*, 102:108158, Sep 2022. doi: 10.1016/j.compeleceng.2022.108158.

[83] D. M. Jimenez Gutierrez, H. Hassan, L. Landi, A. Vitaletti, and I. Chatzigiannakis. A thorough assessment of the non-IID data impact in federated learning. *arXiv preprint arXiv:2503.17070*, 2025. doi: 10.48550/arXiv.2503.17070.

[84] A. A. Hady, A. Ghubaish, T. Salman, D. Unal, and R. Jain. Intrusion detection system for healthcare systems using medical and network data: A compar-

ison study. *IEEE Access*, 8:106576–106584, 2020. doi: 10.1109/ACCESS.2020.3000421.

[85] A. A. Hady, A. Ghubaish, T. Salman, D. Unal, and R. Jain. Intrusion detection system for healthcare systems using medical and network data: A comparison study. *IEEE Access*, 8:106576–106584, 2020. doi: 10.1109/ACCESS.2020.3000421.

[86] J. Han, S. Zhang, and X. Wang. CFMT: Cross-silo federated multi-task learning for IoT intrusion detection. *IEEE Internet of Things Journal*, 11(6):10112–10125, 2024. doi: 10.1109/JIOT.2023.3336221.

[87] Asma Haroon et al. Constraints in the IoT: The world in 2020 and beyond. *International Journal of Advanced Computer Science and Applications*, 7(11), 2016.

[88] K. Haseeb, N. Islam, T. Saba, A. Rehman, and Z. Mehmood. LSDAR: A lightweight structure based data-aware routing protocol for trust-based IoT networks. *IEEE Internet of Things Journal*, 9(15):13211–13220, 2022. doi: 10.1109/JIOT.2022.3141876.

[89] Richard Hillestad et al. Can electronic medical record systems transform health care? *Health Affairs*, 24(5):1103–1117, 2005.

[90] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735.

[91] C. Huang, J. Wang, S. Wang, and Y. Zhang. Internet of medical things: A systematic review. *Neurocomputing*, 557:126719, 2023. doi: 10.1016/j.neucom.2023.126719.

[92] Chao-Hsi Huang and Kung-Wei Cheng. RFID technology combined with IoT application in medical nursing system. *Bulletin of Networking, Computing, Systems, and Software*, 3(1):20–24, 2014.

[93] C.-L. Hwang and K. Yoon. *Multiple Attribute Decision Making: Methods and Applications*, volume 186 of *Lecture Notes in Economics and Mathematical Systems*. Springer, Berlin, Heidelberg, 1981. doi: 10.1007/978-3-642-48318-9.

[94] Jorge E. Ibarra-Esquer et al. Tracking the evolution of the Internet of Things concept across different application domains. *Sensors*, 17(6):1379, 2017.

[95] I. Ioannou et al. Gemlids-miot: A green effective machine learning intrusion detection system based on federated learning for medical iot network security

hardening. *Computer Communications*, 218:209–239, Mar 2024. doi: 10.1016/j.comcom.2024.02.023.

[96] S. Islam, M. Ahmed, and R. Haque. PP-HFFL: Privacy-preserving hierarchical federated learning for non-IID IoT traffic. *Sensors*, 25(5):1421, 2025. doi: 10.3390/s25051421.

[97] Robert S. H. Istepanian et al. Internet of m-health things "m-IoT". In *IET Seminar on Assisted Living 2011*, pages 1–3. IET, 2011.

[98] Z. Jiang, Z. Cao, B. Krishnamachari, S. Zhou, and Z. Niu. Senate: A permissionless byzantine consensus protocol in wireless networks for real-time internet-of-things applications. *IEEE Internet of Things Journal*, 7(7):6576–6588, 2020.

[99] Y. Jin, H. Zhu, J. Xu, and Y. Chen. *Federated Learning Fundamentals and Advances*. Machine Learning: Foundations, Methodologies, and Applications. Springer, Singapore, 2022. doi: 10.1007/978-981-19-7083-2.

[100] Y. Jin, H. Zhu, J. Xu, and Y. Chen. *Federated learning fundamentals and advances*. Springer, 2022.

[101] B. K. Joardar, A. I. Arka, J. R. Doppa, and P. P. Pande. Fault-tolerant deep learning using regularization. In *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, page 1–6, 2022.

[102] Gulraiz J. Joyia et al. Internet of Medical Things (IoMT): Applications, benefits and future challenges in healthcare domain. *Journal of Communications*, 12(4):240–247, 2017.

[103] M. H. Kabir, M. S. Rajib, A. S. M. T. Rahman, M. M. Rahman, and S. K. Dey. Network intrusion detection using unsw-nb15 dataset: Stacking machine learning based approach. In *2022 International Conference on Advancement in Electrical and Electronic Engineering, ICAEEE 2022*, 2022. doi: 10.1109/ICAEEE54957.2022.9836404.

[104] V. Kantharaju, H. Suresh, and B. S. Jayasri. SAPGAN: Self-attention pruning GAN for intrusion detection in the IoT. *Scientific Reports*, 14:23712, 2024. doi: 10.1038/s41598-024-74921-4.

[105] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *Proceedings of the International Conference on Machine Learning*, pages 5132–5143, 2020.

[106] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh. SCAFFOLD: Stochastic controlled averaging for federated learning. In *Proceedings of the International Conference on Machine Learning*, pages 5132–5143, 2020.

[107] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh. SCAFFOLD: Stochastic controlled averaging for federated learning. In *Proc. Int. Conf. Machine Learning (ICML)*, pages 5132–5143, 2020.

[108] G. S. Karthick and P. B. Pankajavalli. A review on human healthcare Internet of Things: A technical perspective. *SN Computer Science*, 1(4):1–19, 2020.

[109] G. G. K. W. M. S. I. R. Karunarathne, K. A. D. T. Kulawansa, and M. F. M. Firdhous. Wireless communication technologies in Internet of Things: A critical evaluation. In *2018 International Conference on Intelligent and Innovative Computing Applications (ICONIC)*, pages 1–5. IEEE, 2018.

[110] S. M. Kasongo. A deep learning technique for intrusion detection system using a Recurrent Neural Networks based framework. *Computer Communications*, 199: 113–125, 2023. doi: 10.1016/j.comcom.2022.12.010.

[111] R. Katz, P. Goldstein, and R. Yanosky. Cloud computing in higher education, 2010.

[112] H. Kaur and S. Sharma. A comparative study of wireless technologies: ZigBee, Bluetooth LE, EnOcean, Wavenis, Insteon and UWB. In *Proceedings of the International Conference on Recent Trends in Computing and Communication Engineering (RTCCE'13)*, 2013.

[113] Syed Hussain Ali Kazmi, Rosilah Hassan, Faizan Qamar, Kashif Nisar, and Dahlila Putri Dahnil. Threat intelligence in iomts with federated learning using non-iid data: An experimental analysis. In *2024 IEEE 7th International Symposium on Telecommunication Technologies (ISTT)*, pages 120–125. IEEE, 2024. doi: 10.1109/ISTT59650.2024.10738501.

[114] M. I. Khaleel. A fault tolerance aware green iot workflow scheduling algorithm for multi-dimensional resource utilization in sustainable cloud computing. *Internet of Things*, 23:100909, 2023.

[115] F. Khan, M. A. Jan, R. Alturki, M. D. Alshehri, S. T. Shah, and A. U. Rehman. A secure ensemble learning-based fog-cloud approach for cyberattack detection in iomt. *IEEE Transactions on Industrial Informatics*, 19(10):10125–10132, Oct 2023. doi: 10.1109/TII.2022.3231424.

[116] I. A. Khan et al. Fed-inforce-fusion: A federated reinforcement-based fusion model for security and privacy protection of iomt networks against cyber-attacks. *Information Fusion*, 101:102002, Jan 2024. doi: 10.1016/j.inffus.2023.102002.

[117] Jamil Y. Khan and Mehmet R. Yuce. *Wireless Body Area Network (WBAN) for Medical Applications*. InTechOpen, 2010.

[118] M. Khan, M. Iqbal, and M. A. Khan. Enhanced IoT intrusion detection using Random Forest on the CIC-IoT taxonomy. *Scientific Reports*, 14:11842, 2024. doi: 10.1038/s41598-024-62814-5.

[119] Rafiullah Khan et al. Future internet: The Internet of Things architecture, possible applications and key challenges. In *2012 10th International Conference on Frontiers of Information Technology*, pages 257–260. IEEE, 2012.

[120] Samiya Khan and Mansaf Alam. Wearable Internet of Things for personalized healthcare. In *Health Informatics: A Computational Perspective in Healthcare*, pages 43–60. 2021.

[121] S. Kim, L. Chen, and J. Kim. Intrusion prediction using lstm and gru with unsw-nb15. In *2021 Computing, Communications and IoT Applications (ComComAp)*, page 101–106. IEEE, 2021.

[122] Duddela Dileep Kumar and Pratti Venkateswarlu. Secured smart healthcare monitoring system based on IoT. *Imperial Journal of Interdisciplinary Research*, 2(10), 2016.

[123] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982. doi: 10.1145/357172.357176.

[124] Kumar Laxman, Sharanie Banu Krishnan, and Jaspaljeet Singh Dhillon. Barriers to adoption of consumer health informatics applications for health self management. *Health Science Journal*, 9(5):1, 2015.

[125] Huang-Chen Lee and Kai-Hsiang Ke. Monitoring of large-area IoT sensors using a LoRa wireless mesh network system. *IEEE Transactions on Instrumentation and Measurement*, 67(9):2177–2187, 2018.

[126] Jin-Shyan Lee, Yu-Wei Su, and Chung-Chou Shen. A comparative study of wireless protocols: Bluetooth, UWB, ZigBee, and Wi-Fi. In *IECON 2007–33rd Annual Conference of the IEEE Industrial Electronics Society*, pages 46–51. IEEE, 2007.

[127] S. Li, E. C.-H. Ngai, and T. Voigt. An experimental study of Byzantine-robust aggregation schemes in federated learning. *IEEE Transactions on Big Data*, 10 (6):975–988, 2023. doi: 10.1109/TBDATA.2023.3237397.

[128] Z. Li, Y. He, H. Yu, J. Kang, X. Li, Z. Xu, and D. Niyato. Data heterogeneity-robust federated learning via group client selection in industrial IoT. *IEEE Internet of Things Journal*, 9(18):17844–17857, 2022. doi: 10.1109/JIOT.2022.3161943.

[129] Zhuguo Li et al. The evolution of IoT wireless networks for low-rate and real-time applications. 18(1):175–188, 2017.

[130] Z. Lian, Q. Zeng, W. Wang, T. R. Gadekallu, and C. Su. Blockchain-based two-stage federated learning with non-IID data in IoMT system. *IEEE Transactions on Computational Social Systems*, 10(4):1701–1710, 2023. doi: 10.1109/TCSS.2022.3216802.

[131] Y. Liu, Y. Kang, C. Xing, T. Chen, and Q. Yang. A secure federated transfer learning framework. *IEEE Intelligent Systems*, 35(4):70–82, 2020. doi: 10.1109/MIS.2020.2988525.

[132] R. Madhumathi, T. Arumuganathan, and R. Shruthi. Internet of Things in precision agriculture: A survey. In *Intelligent Sustainable Systems: Proceedings of ICISS 2021*, pages 539–553. 2022.

[133] E. Mahdavi, A. Fanian, A. Mirzaei, and Z. Taghiyarrenani. Itl-ids: Incremental transfer learning for intrusion detection systems. *Knowledge-Based Systems*, 253: 109542, Oct 2022. doi: 10.1016/j.knosys.2022.109542.

[134] Manas Ranjan Mallick. A comparative study of wireless protocols with Li-Fi technology: A survey. In *Proceedings of 43rd IRF International Conference*, pages 8–12, 2016.

[135] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282, 2017.

[136] O. Mengara, Y. Cherief, and N. Hassan. IoTSecUT: An uncertainty-aware Transformer-cGAN-autoencoder framework for IoT intrusion detection. *IEEE Internet of Things Journal*, 11(19):30977–30992, 2024. doi: 10.1109/JIOT.2024.3412334.

[137] Wu Mengdi. Wireless communication technologies in Internet of Things (IoT). Master's thesis, Faculty of Technology, Communication and Systems Engineering, 2017.

[138] Lionel Metongnon and Ramin Sadre. Fast and efficient probing of heterogeneous IoT networks. *International Journal of Network Management*, 28(1):e1997, 2018.

[139] Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai. Kitsune: An ensemble of autoencoders for online network intrusion detection. In *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS)*, 2018. doi: 10.14722/ndss.2018.23211.

[140] A. Misbah, A. Sebbar, and I. Hafidi. Securing internet of medical things: An advanced federated learning approach. *International Journal of Advanced Computer Science & Applications*, 16(2), 2025.

[141] Alaa Hamid Mohammed, Raad M. Khaleefah, Ihsan Amjad Abdulateef, et al. A review software defined networking for Internet of Things. In *2020 International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, pages 1–8. IEEE, 2020.

[142] Satarupa Mohanty et al. Smart healthcare analytics using Internet of Things: An overview. In *Smart Healthcare Analytics: State of the Art*, pages 1–11. 2022.

[143] N. Moustafa and J. Slay. Unsw-nb15: a comprehensive data set for network intrusion detection systems (unsw-nb15 network data set). In *2015 Military Communications and Information Systems Conference (MilCIS)*, pages 1–6, Canberra, Australia, 2015. IEEE. doi: 10.1109/MilCIS.2015.7348942.

[144] Ghulam Muhammad, Mohammed F. Alhamid, and Xiaomi Long. Computing and processing on the edge: Smart pathology detection for connected healthcare. *IEEE Network*, 33(6):44–49, 2019.

[145] Sankar Mukherjee and G. P. Biswas. Networking for IoT and applications using existing communication technology. *Egyptian Informatics Journal*, 19(2): 107–127, 2018.

[146] S. Nandy, M. Adhikari, M. A. Khan, V. G. Menon, and S. Verma. An intrusion detection mechanism for secured iomt framework based on swarm-neural network. *IEEE Journal of Biomedical and Health Informatics*, 26(5):1969–1976, May 2022. doi: 10.1109/JBHI.2021.3101686.

[147] T. A. Nguyen, M. Conti, N. V. Abhishta, and M. H. Au. Towards adversarially robust decision tree ensembles for IoT intrusion detection. *Internet of Things*, 22:100756, 2023. doi: 10.1016/j.iot.2023.100756.

[148] Lab of UNSW Canberra. The unsw-nb15 dataset. URL https://research.unsw.edu.au/projects/unsw-nb15-dataset.

[149] Luís M. L. Oliveira and Joel J. P. C. Rodrigues. Wireless sensor networks: A survey on environmental monitoring. *Journal of Communications*, 6(2):143–151, 2011.

[150] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010. doi: 10.1109/TKDE.2009.191.

[151] A. K. Pandey, M. Raj, and A. Kumar. A survey of lightweight cryptographic algorithms for IoT-based applications. *Wireless Networks*, 30(7):6655–6677, 2024. doi: 10.1007/s11276-023-03409-2.

[152] Mani Pareek and Sushil Buriya. A study of link layer protocols in IoT. *International Journal on Future Revolution in Computer Science and Communication Engineering*, 4(2):355–359, 2018.

[153] A. H. M. Shahariar Parvez et al. Effect of fault tolerance in the field of cloud computing. In *Inventive Computation Technologies*, volume 4, pages 297–305. Springer, 2020.

[154] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, volume 28, pages 1310–1318, 2013.

[155] H. Peng, X. Liu, and K. Wang. FD-IDS: Federated and distilled intrusion detection for resource-constrained IoT. *Sensors*, 25(2):512, 2025. doi: 10.3390/s25020512.

[156] Charith Perera et al. Context aware computing for the Internet of Things: A survey. *IEEE Communications Surveys and Tutorials*, 16(1):414–454, 2013.

[157] Veena Pureswaran and Paul Brody. Device democracy: Saving the future of the Internet of Things. *IBM Corporation*, 23, 2015.

[158] Chetanya Puri et al. iCarMa: Inexpensive cardiac arrhythmia management—an IoT healthcare analytics solution. In *Proceedings of the First Workshop on IoT-enabled Healthcare and Wellness Technologies and Systems*, pages 3–8, 2016.

[159] Y. Qi, S. Shao, S. Wu, X. Qiu, S. Guo, and S. Guo. A distributed intelligent service trusted provision approach for iot. *IEEE Internet Things J*, 2023.

[160] C. Z. Radulescu, M. Radulescu, and R. Boncea. A linear trade-off group topsis method with application for internet of things devices ranking. *Procedia Computer Science*, 242:528–535, 2024.

[161] G. K. Ragesh and K. Baskaran. An overview of applications, standards and challenges in futuristic wireless body area networks. *International Journal of Computer Science Issues (IJCSI)*, 9(1):180, 2012.

[162] A. Rahmati, M. Salehi, and M. Dehghan. Federated GRU with homomorphic encryption for privacy-preserving intrusion detection in IoT. *Informatics*, 12(1): 18, 2025. doi: 10.3390/informatics12010018.

[163] S. Rajasoundaran, S. V. N. Santhosh Kumar, M. Selvi, and A. Kannan. A comprehensive survey of firewall mechanisms for securing resource-limited wireless networks. *Security and Privacy*, 7(4):e392, 2024. doi: 10.1002/spy2.392.

[164] A. Reyana, S. Kautish, K. A. Alnowibet, H. M. Zawbaa, and A. Wagdy Mohamed. Opportunities of iot in fog computing for high fault tolerance and sustainable energy optimization. *Sustainability*, 15(11):8702, 2023.

[165] Ognjen Riđ ić et al. The smart city, smart contract, smart health care, Internet of Things (IoT), opportunities, and challenges. In *Blockchain Technologies for Sustainability*, pages 135–149. 2022.

[166] Md Robiul Alam Robel et al. Fault tolerance in cloud computing—an algorithmic approach. In *Innovations in Bio-Inspired Computing and Applications (IBICA 2019)*, pages 307–316. Springer, 2021.

[167] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986. doi: 10.1038/323533a0.

[168] P. Ruzafa-Alcázar, P. Fernández-Saura, E. Cánovas, M. Gil-Pérez, A. Huertas Celdrán, and F. J. García Clemente. Intrusion detection based on privacy-preserving federated learning for the industrial IoT. *IEEE Transactions on Industrial Informatics*, 19(2):1145–1154, 2023. doi: 10.1109/TII.2021.3126728.

[169] T. L. Saaty. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. McGraw-Hill, New York, NY, USA, 1980.

[170] T. Saba, A. Rehman, T. Sadad, H. Kolivand, and S. A. Bahaj. Anomaly-based intrusion detection system for IoT networks through deep learning model. *Computers and Electrical Engineering*, 99:107810, 2022. doi: 10.1016/j.compeleceng.2022.107810.

[171] P. K. Sadhu, V. P. Yanambaka, and A. Abdelgawad. Internet of things: Security and solutions survey. *Sensors*, 22(19):7433, 2022. doi: 10.3390/s22197433.

[172] M. Sajid, K. Malik, and F. A. Khan. Hybrid XGBoost-CNN-LSTM framework for network intrusion detection. *Journal of Cloud Computing*, 13:72, 2024. doi: 10.1186/s13677-024-00627-7.

[173] H. M. Saleh, H. Marouane, and A. Fakhfakh. Stochastic gradient descent intrusions detection for wireless sensor network attack detection system using machine learning. *IEEE Access*, 12:3825–3836, 2024. doi: 10.1109/ACCESS.2023.3349248.

[174] A. Salehpour and K. Samadzamini. A bibliometric analysis on the application of deep learning in economics, econometrics, and finance. *International Journal of Computational Science and Engineering*, 27(2):167–181, 2024. doi: 10.1504/IJCSE.2024.137074.

[175] A. Salehpour, M. Norouzi, M. A. Balafar, and K. SamadZamini. A cloud-based hybrid intrusion detection framework using xgboost and adasyn-augmented random forest for iomt. *IET Communications*, 18(19):1371–1390, Dec 2024. doi: 10.1049/cmu2.12833.

[176] A. Salehpour, M. A. Balafar, and A. Souri. An optimized intrusion detection system for resource-constrained iomt environments: enhancing security through efficient feature selection and classification. *Journal of Supercomputing*, 81(6):72–53, Apr 2025. doi: 10.1007/s11227-025-07253-3.

[177] Tara Salman and Raj Jain. A survey of protocols and standards for Internet of Things. *arXiv preprint arXiv:1903.11549*, 2019.

[178] Omaji Samuel et al. IoMT: A COVID-19 healthcare system driven by federated learning and blockchain. *IEEE Journal of Biomedical and Health Informatics*, 27(2):823–834, 2022.

[179] M. Sarhan, S. Layeghy, and M. Portmann. Towards a standard feature set for network intrusion detection system datasets. *Mobile Networks and Applications*, 27(1):357–370, 2022. doi: 10.1007/s11036-021-01843-0.

[180] M. Sarhan, S. Layeghy, N. Moustafa, and M. Portmann. Cyber threat intelligence sharing scheme based on federated learning for network intrusion detection. *Journal of Network and Systems Management*, 31(1):3, Mar 2023. doi: 10.1007/s10922-022-09691-3.

[181] Pallavi Sethi and Smruti R. Sarangi. Internet of Things: Architectures, protocols, and applications. *Journal of Electrical and Computer Engineering*, 2017.

[182] Arbia Riahi Sfar et al. A roadmap for security challenges in the Internet of Things. *Digital Communications and Networks*, 4(2):118–137, 2018.

[183] Pradip Kumar Sharma, Young-Sik Jeong, and Jong Hyuk Park. EH-HL: Effective communication model by integrated EH-WSN and hybrid LiFi/WiFi for IoT. *IEEE Internet of Things Journal*, 5(3):1719–1726, 2018.

[184] Bhagya Nathali Silva, Murad Khan, and Kijun Han. Internet of Things: A comprehensive review of enabling technologies, architecture, and challenges. *IETE Technical Review*, 35(2):205–220, 2018.

[185] P. Singh, G. S. Gaba, A. Kaur, M. Hedabou, and A. Gurtov. Dew-cloud-based hierarchical federated learning for intrusion detection in iomt. *IEEE Journal of Biomedical and Health Informatics*, 27(2):722–731, Feb 2023. doi: 10.1109/JBHI.2022.3186250.

[186] Rashmi Singh. A proposal for mobile e-care health service system using IoT for Indian scenario. *Journal of Network Communications and Emerging Technologies (JNCET)*, 6(1), 2016.

[187] A. Sinha, S. Chakraborty, and R. Chatterjee. Adversarially robust LSTM-CNN intrusion detection for IoT networks. *Scientific Reports*, 15:9821, 2025. doi: 10.1038/s41598-025-94512-3.

[188] A. Soomro, A. Rafique, and M. Waqas. SecureDyn-FL: Dynamic federated intrusion detection with adaptive client selection. *IEEE Transactions on Network and Service Management*, 23(1):112–127, 2026. doi: 10.1109/TNSM.2026.3510112.

[189] Lanfang Sun et al. Edge-cloud computing and artificial intelligence in Internet of Medical Things: Architecture, technology and application. *IEEE Access*, 8:101079–101092, 2020.

[190] M. Tavallaee, E. Bagheri, W. Lu, and A. A. Ghorbani. A detailed analysis of the KDD CUP 99 data set. In *Proc. IEEE Symp. Computational Intelligence for Security and Defense Applications (CISDA)*, pages 1–6, 2009. doi: 10.1109/CISDA.2009.5356528.

[191] M. Tawfik, R. Hassan, and S. Elsayed. FedMedSecure: An explainable federated intrusion detection framework for the Internet of Medical Things. *Scientific Reports*, 15:8154, 2025. doi: 10.1038/s41598-025-90514-w.

[192] V. A. Thakor, M. A. Razzaque, and M. R. A. Khandaker. Lightweight cryptography algorithms for resource-constrained IoT devices: A review, comparison and research opportunities. *IEEE Access*, 9:28177–28193, 2021. doi: 10.1109/ACCESS.2021.3052867.

[193] T. T. Thein, Y. Shiraishi, and M. Morii. pFL-IDS: Personalized federated learning for intrusion detection in heterogeneous IoT. *Future Generation Computer Systems*, 146:146–160, 2023. doi: 10.1016/j.future.2023.04.014.

[194] Eric J. Topol, Steven R. Steinhubl, and Ali Torkamani. Digital medical tools and sensors. *JAMA*, 313(4):353–354, 2015.

[195] Hong-Linh Truong and Schahram Dustdar. Principles for engineering IoT cloud systems. *IEEE Cloud Computing*, 2(2):68–76, 2015.

[196] Vasileios Tsoutsouras et al. Software design and optimization of ECG signal analysis and diagnosis for embedded IoT devices. In *Components and Services for IoT Platforms*, pages 299–322. 2017.

[197] Nidal M. Turab. IoT wireless home automation technologies and their relation to specific absorption rate. *Journal of Theoretical and Applied Information Technology*, 96(14):4597–4609, 2018.

[198] Z. Turgut and M. S. Başarslan. Xbideep: A novel explainable artificial intelligence based intrusion detection system for internet of medical things environment. *Internet of Things*, 33:101675, Sep 2025. doi: 10.1016/j.iot.2025.101675.

[199] U.S. Food and Drug Administration. Early alert: Infusion pump software issue from Baxter. https://www.fda.gov/medical-devices/medical-device-recalls/, 2025. Accessed: Aug. 13, 2025.

[200] U.S. Food and Drug Administration. Cybersecurity in medical devices: Quality system considerations and content of premarket submissions (final guidance), 2025. URL https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations\and-content-premarket-submissions. Accessed: Aug. 13, 2025.

[201] U.S. HHS. Summary of the hipaa security rule, 2024. URL https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html. Accessed: Aug. 13, 2025.

[202] U.S. HHS Office for Civil Rights. Change healthcare cybersecurity incident: Frequently asked questions, 2025. URL https://www.hhs.gov/hipaa/for-professionals/special-topics/change-healthcare-cybersecurity-incident-frequently-asked-questions/index.html. Accessed: Aug. 13, 2025.

[203] Alem Čolaković and Mesud Hadžialić. Internet of Things (IoT): A review of enabling technologies, challenges, and open research issues. *Computer Networks*, 144:17–39, 2018.

[204] Feng Wang et al. A survey from the perspective of evolutionary process in the Internet of Things. *International Journal of Distributed Sensor Networks*, 11(3): 462752, 2015.

[205] Miao Wu et al. Research on the architecture of Internet of Things. In *2010 3rd International Conference on Advanced Computer Theory and Engineering (ICACTE)*, volume 5, pages V5–484. IEEE, 2010.

[206] Y. Yang, H. Cao, X. Chen, and L. Ge. A lightweight authentication protocol for industrial IoT based on ECC and trusted token. *Sensors*, 23(20):8558, 2023. doi: 10.3390/s23208558.

[207] Ibrar Yaqoob et al. Enabling communication technologies for smart cities. *IEEE Communications Magazine*, 55(1):112–120, 2017.

[208] M. Ye, X. Fang, B. Du, P. C. Yuen, and D. Tao. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Computing Surveys*, 56(3):1–44, 2023. doi: 10.1145/3625558.

[209] L. Yu, J. Zhang, and Z. Li. TOPSIS-based fault-tolerant routing algorithm for Network-on-Chip. *Scientific Reports*, 15:6581, 2025. doi: 10.1038/s41598-025-88914-4.

[210] N. Zafar, A. Khanna, S. Jain, Z. Ali, and J. Ahamed. Safeguarding iot: Harnessing practical byzantine fault tolerance for robust security. In *International Conference on Data Analytics & Management*, page 287–301. Springer, 2023.

[211] G. Zhang, F. Pan, Y. Mao, S. Tijanic, M. Dang'Ana, S. Motepalli, S. Zhang, and H.-A. Jacobsen. Reaching consensus in the byzantine empire: A comprehensive review of bft consensus algorithms. *ACM Computing Surveys*, 56(5):1–41, 2024.

[212] Z. Zhang, X. Zhang, Q. Niu, and P. Zhang. Federated transfer learning for intrusion detection in industrial IoT. *IEEE Transactions on Industrial Informatics*, 19(8):8859–8868, 2023. doi: 10.1109/TII.2022.3225727.

[213] Y. Zhao, J. Chen, J. Zhang, D. Wu, and R. Ranjan. Semi-supervised federated learning with knowledge distillation for IoT intrusion detection. *IEEE Internet of Things Journal*, 10(11):9935–9947, 2023. doi: 10.1109/JIOT.2022.3231412.

[214] H. Zhu, J. Xu, S. Liu, and Y. Jin. Federated learning on non-IID data: A survey. *Neurocomputing*, 465:371–390, 2021. doi: 10.1016/j.neucom.2021.07.098.

[215] L. Zhu, Z. Liu, and S. Han. Deep leakage from gradients. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, pages 14747–14756, 2019.

[216] F. Zhuang et al. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, Jan 2021. doi: 10.1109/JPROC.2020.3004555.

[217] Y. Zou, L. Yang, G. Jing, R. Zhang, Z. Xie, H. Li, and D. Yu. A survey of fault tolerant consensus in wireless networks. *High-Confidence Computing*, page 100202, 2024.

[218] U. Zukaib, X. Cui, C. Zheng, D. Liang, and S. U. Din. Meta-fed ids: Meta-learning and federated learning based fog-cloud approach to detect known and zero-day cyber attacks in iomt networks. *Journal of Parallel and Distributed Computing*, 192:104934, Oct 2024. doi: 10.1016/j.jpdc.2024.104934.