UNISINOS

Programa de Pós-Graduação em

Computação Aplicada

Mestrado Acadêmico

Lucas Micol Policarpo

FEDERATED HOSPITAL:
A Multilevel Federated Learning Architecture for dealing with heterogeneous data distribution in the context of smart hospitals services

São Leopoldo, 2023

Lucas Micol Policarpo

FEDERATED HOSPITAL: **A Multilevel Federated Learning Architecture for dealing with heterogeneous data distribution in the context of smart hospitals services**

Dissertação apresentada como requisito parcial para a obtenção do título de Mestre pelo Programa de Pós-Graduação em Computação Aplicada da Universidade do Vale do Rio dos Sinos — UNISINOS

Advisor:
Prof. Dr. Rodrigo da Rosa Righi

São Leopoldo
2023

P766f Policarpo, Lucas Micol.

Federated hospital : a multilevel federated learning architecture for dealing with heterogeneous data distribution in the context of smart hospitals services / by Lucas Micol Policarpo. – 2023.

46 p. : il. ; 30 cm.

Dissertação (mestrado) — Universidade do Vale do Rio dos Sinos, Programa de Pós-Graduação em Computação Aplicada, São Leopoldo, RS, 2023.

“Advisor: Dr. Rodrigo da Rosa Righi”.

1. Federated learning. 2. Smart hospital. 3. Unbalance data distribution. 4. Healthcare. I. Title.

UDC: 004.85:64.024.8

Cataloging in Publication (CIP):
Librarian Alessandro Dietrich - CRB 10/2338

RESUMO

A integração de serviços de inteligência artificial e aprendizado de máquina na área da saúde revolucionou o atendimento ao paciente, abrangendo desde o monitoramento de saúde em tempo real até a análise complexa de imagens médicas. No entanto, a implementação desses serviços de aprendizado de máquina no contexto de hospitais inteligentes apresenta desafios significativos devido às diversas demandas de dados e preocupações com a privacidade. O Aprendizado Federado emerge como uma solução promissora, permitindo que os dados permaneçam com os usuários enquanto os modelos de aprendizado de máquina são treinados de forma colaborativa. O aprendizado federado garante a privacidade dos dados e oferece escalabilidade ao possibilitar o aprendizado distribuído entre vários usuários.

Nesta pesquisa, estendemos o paradigma do aprendizado federado para o domínio dos hospitais inteligentes e propomos o modelo "Hospital Federado"para enfrentar os desafios decorrentes da heterogeneidade entre diferentes departamentos hospitalares. Através da agregação em vários níveis, a arquitetura do Hospital Federado é projetada para acomodar as diversas demandas e situações de saúde dentro de cada departamento individual, fornecendo modelos de aprendizado de máquina personalizados e precisos para cada usuário.

Por meio de experimentação extensa e avaliação em cenários distintos, incluindo distribuições de dados homogêneas e heterogêneas, comparamos o desempenho do modelo do Hospital Federado em relação às abordagens padrão de aprendizado de máquina e aprendizado federado. Os resultados confirmam a eficácia de nossa proposta em termos de precisão, eficiência e velocidade de convergência. Além disso, o processo de agregação em vários níveis na arquitetura do hospital inteligente aprimora o desempenho do modelo, garantindo a geração de modelos de aprendizado de máquina personalizados específicos para as características únicas de cada departamento.

O modelo do Hospital Federado demonstra seu potencial para melhorar a execução de serviços orientados por aprendizado de máquina em hospitais inteligentes. Ao otimizar a precisão e o desempenho dos modelos de aprendizado de máquina para diversos departamentos de saúde, nossa proposta visa revolucionar a tomada de decisões baseada em dados, promovendo o atendimento personalizado ao paciente e serviços de saúde eficientes. O próximo passo desta pesquisa é implementar o Hospital Federado em hospitais reais na região metropolitana de Porto Alegre, Rio Grande do Sul.

Palavras-chave: Aprendizado federado. hospitais inteligentes. Distribuição de dados desbalanceada.

ABSTRACT

The integration of artificial intelligence (AI) and machine learning (ML) services in healthcare has revolutionized patient care, ranging from real-time health monitoring to complex medical image analysis. However, deploying these ML services in the context of smart hospitals poses significant challenges due to varying data demands and privacy concerns. Federated Learning (FL) emerges as a promising solution by allowing data to remain with users while training ML models collaboratively. FL ensures data privacy and offers scalability by enabling distributed learning across multiple users.

In this research, we extend the FL paradigm to the domain of smart hospitals and propose the "Federated Hospital" model to address the challenges posed by heterogeneity among different hospital departments. By leveraging multi-level aggregation, the Federated Hospital architecture is designed to accommodate the diverse demands and health situations within individual departments, providing personalized and accurate ML models for each user.

Through extensive experimentation and evaluation in distinct scenarios, including homogeneous and heterogeneous data distributions, we compare the performance of the Federated Hospital model against standard ML and FL approaches. The results confirm the effectiveness of our proposal in terms of accuracy, efficiency, and convergence speed. Moreover, the multi-level aggregation process in the smart hospital architecture enhances model performance, ensuring the generation of tailored ML models specific to each department's unique characteristics.

The Federated Hospital model demonstrates its potential to improve the execution of ML-oriented services in smart hospitals. By optimizing the accuracy and performance of ML models for diverse healthcare departments, our proposal aims to revolutionize data-driven decision-making, promoting personalized patient care and efficient healthcare services. The next step of this research is to execute Federated Hospital in real hospitals in the metropolitan area of Porto Alegre, Rio Grande do Sul.

Keywords: Federated Learning. Smart hospital. Unbalance data distribution.

LIST OF FIGURES

LIST OF TABLES

LIST OF ACRONYMS

AI	Artificial Intelligence
ANN	Artificial neural networks
DL	Deep Learning
DLG	Deep Leakage from Gradients
EHR	Electronic Health Record
FHE	Fully Homomorphic Encryption
HC	Health Care
HE	Homomorphic Encryption
IoMT	Internet of Medical Things
IoT	Internet of Things
FL	Federated Learning
FEC	Fog and edge computing
ML	Machine Learning

CONTENTS

1 INTRODUCTION

The emergence of machine learning (ML) as a way to program a computer to learn and improve automatically has dramatically innovated many areas. In some cases, the possible applications of ML were so helpful that new sub-areas were forming (BURKOV, 2019). One of the sub-areas that is constantly growing is ML applied to healthcare. Combining healthcare data with data analysis and ML techniques to identify patterns of interest is commonly known as health informatics (CALLAHAN; SHAH, 2017a). Even though the topic of health informatics is growing gradually, issues such as data security and privacy are topics that come into the debate because of the need for data from ML algorithms (ZHOU, 2023).

To address these challenges, federated learning (FL) has emerged as a promising solution. FL aims to decentralize training data by conducting collaborative learning sessions on users' devices (REHMAN MUHAMMAD HABIB UR, 2021). The users leverage pre-trained models and collaboratively improve a central model through the aggregation of distributed model updates. Figure 1 illustrates a standard FL diagram using the FedAvg protocol, where the global model is updated through weighted averaging of the trained models' weights until convergence.

Figure 1: Federated Learning diagram showing how communication proceeds between the aggregating server and individual clients.

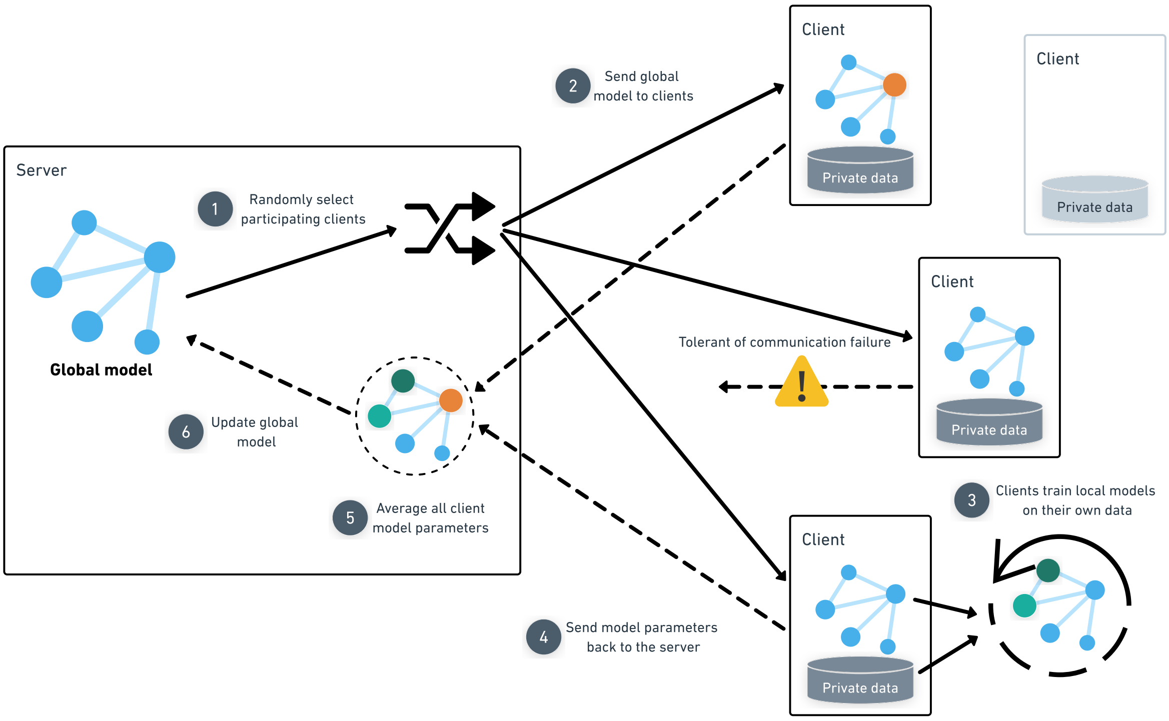
Server
Global model
1 Randomly select participating clients
2 Send global model to clients
Client
Private data
Client
Private data
Client
Private data
Tolerant of communication failure
3 Clients train local models on their own data
Client
Private data
4 Send model parameters back to the server
5 Average all client model parameters
6 Update global model

Source: Adapted from (REHMAN MUHAMMAD HABIB UR, 2021)

However, sharing gradients in collaborative learning poses significant security risks, as demonstrated by Zhu et al. (ZHU; HAN, 2020), who showed that pixel-wise accurate images and token-wise matching for texts can be recovered from shared gradients. This threat to privacy becomes a critical concern in healthcare informatics, where models often rely on sensitive patient information (ZHU et al., 2023). In such scenarios, collaborative learning methods lacking robust privacy measures may face resistance from users and healthcare institutions.

In this master's thesis, we propose the Federated Hospital model, a novel computational architecture tailored for executing federated learning services in intelligent hospitals. Our approach addresses the open issues in the literature, with a particular focus on model customization when dealing with data heterogeneity. The main contributions of this work include the proposed Federated Hospital architecture, a novel aggregation method centered around unbalanced data distribution, and the implementation of a prototype to evaluate different execution scenarios.

1.1 Motivation

In the context of smart cities, hospitals are critical in providing healthcare services. These services can range from answering an emergency call to monitoring a patient's vital signs. Within the scope of this work, we are looking for a method to improve FL services in smart hospitals. More precisely, services aimed at cooperation between hospitals for the improvement of ML models through the use of federated learning.

FL can be seen as a decentralized machine-learning technique that allows models to be trained in local devices without sending data to a central server. Smart hospitals have a large amount of sensitive data that needs to be protected, and FL can help ensure data privacy by keeping local information on each device. Additionally, FL can be applied to improve the accuracy of machine learning models in smart hospitals, allowing models to be trained locally with specific data from each hospital. While FL has several advantages, there are also disadvantages to consider. One of the main issues is that model performance can be affected by the heterogeneity of local data, which can lead to less accurate models (DASARADHARAMI REDDY; GADEKALLU et al., 2023). Additionally, FL requires a large amount of computational resources to run, which can be a challenge for hospitals with limited resources.

The study by Xu et al. (XU et al., 2021) reviews FL techniques oriented to healthcare informatics. The work shows some of the consequences of applying federated learning, for example, the FedAvg not being efficient when averaging models with high weight divergence, the communication overhead that may cause performance issues, and the privacy concerns about the sharing gradients. Moreover, tackling the open problems when combining FL with healthcare informatics, such as poor data quality and handling due to the absence of standardization.

In addition, services for the hospital may also be focused on the patients and their personal devices, such as health monitoring using smartwatches and smartphones. When running training sessions, those services tend to overload the user's device and need to be more efficient regarding the time to compute the training. As training occurs several times during the execution of an FL model, the dynamic execution of the model is still a problem that the user's device is in charge of. The survey from Nguyen et al. (NGUYEN et al., 2022) also presents future directions in federated learning for healthcare, such as the privacy problem. In the mentioned scenario, the systems deal with high-sensitivity data, and any attack or leakage can be a huge problem.

Summarizing, while FL presents a promising solution for improving machine learning services in smart hospitals, there are several open gaps that require further exploration and innovation, such as:

- Data heterogeneity distribution: The heterogeneity of local data within smart hospitals can significantly impact the performance of FL models. Addressing the challenge of varying data distributions and characteristics across different hospital sections is essential to ensure accurate and reliable ML models.
- Faster Aggregation Methods: As the Federated Hospitals architecture aims to improve the efficiency and convergence time of ML models, exploring and developing better aggregation methods will be crucial. Novel approaches to aggregating local gradients with reduced communication overhead can significantly enhance the convergence speed of the global model, leading to faster and more efficient training sessions.
- Privacy Concerns: Smart hospitals deal with sensitive and high-sensitivity data, making privacy a paramount concern. Despite employing FL to preserve data privacy, additional measures and advanced privacy-preserving techniques may be required to safeguard patient information effectively.

1.2 Research Question

The research question for this proposal is centered on exploring the effectiveness of a multilevel federated learning architecture for specialized training of service models in smart hospitals. Despite the potential benefits of federated learning for smart hospitals, several gaps and unresolved issues in the literature need to be addressed. One of the main gaps is the need for more research on the effectiveness of federated learning for training models specific to each hospital. Most studies in this area have focused on generalizing models across multiple hospitals, which may need to be more effective in capturing the unique characteristics of each hospital. Additionally, there is a need to explore the feasibility of a multilevel aggregation approach for training specialized service models in each hospital ward.

The research question to be worked on is defined as: *How can we effectively execute machine learning services in smart hospitals while addressing the of different hospital departments' demands?*

In the context of our research, the term "effectively" refers to the ability to efficiently and accurately execute machine learning services in smart hospitals, leading to improved patient care, optimized resource utilization, and enhanced healthcare outcomes. And, "Heterogeneity" pertains to the diversity and variation in data distribution, characteristics, and demands among different hospital departments. In brief, this research aims to contribute to developing effective and efficient machine-learning models for smart hospitals. By exploring the potential of federated learning for specialized training of service models, this project will provide valuable insights into the application of machine learning in healthcare and help address the gaps and limitations in the current literature.

1.3 Objectives

The main objective of this research is to develop and evaluate a multilevel federated learning model for specialized training of service models in smart hospitals. Inside this model, we propose an architecture that will address current approaches' limitations and improve the performance of machine learning models in healthcare. The multilevel federated learning model will consist of three main components: local models, global model, and aggregation algorithms. Local models will be trained with specific data from each hospital ward, while the global model will be trained by aggregating the local models. This multilevel federated learning model will also incorporate a multilevel aggregation approach to improve the accuracy and performance of the ML models. And an architecture that represents the execution flow and operability of the proposal.

The primary objective is to design and implement a novel model that leverages federated learning techniques to address the challenges of executing machine learning services in smart hospitals. The model aims to preserve data privacy, accommodate the heterogeneity of different hospital departments, and improve model accuracy and performance. By leveraging multilevel aggregation, the model architecture enables personalized and specialized ML models for each department, tailored to the unique characteristics of their data.

The specific objectives can be listed as follows:

(i) Raise the bibliographic concepts and evaluate state of the art to carry out this work.

(ii) Find the characteristics, similarities, and open gaps of the gathered works.

(iii) Design a multilevel federated learning model for specialized training of service models in smart hospitals.

(iv) Evaluate the performance of the proposed model and compare it with existing approaches.

(v) Demonstrate the potential of federated learning for improving the accuracy and efficiency of machine learning models in healthcare.

(vi) Contribute to the development of effective and efficient models for smart hospitals and provide insights for the application of machine learning in healthcare.

By achieving these goals, this research will contribute to advancing machine learning techniques in healthcare and provide valuable insights for developing smart hospitals. The proposed model can improve the accuracy and efficiency of machine learning models in healthcare while also handling data heterogeneity. Overall, this work aims to significantly contribute to the field of machine learning in healthcare and improve the quality of healthcare services in smart hospitals.

1.4 Hypothesis

We have the following hypotheses for this work:

- *Hypothesis 1:* Compared to the standard FL approaches, the proposed multilevel federated learning architecture model will improve the accuracy of the models in smart hospitals when dealing with unbalanced datasets;
- *Hypothesis 2:* The multilevel aggregation approach will effectively train specialized service models in each hospital ward, leading to improved model accuracy.

1.5 Text organization

The proposal is organized into six other Chapters. First, Chapter 2 presents the theoretical foundation related to this work, discussing the topics of artificial intelligence, Internet of things, and healthcare informatics. Chapter 3 discusses the state-of-the-art analysis, presenting the selection process and the academic proposals related to this work, as well as the open literature gaps. Then, in Chapter 4, the Federated Hospital model is proposed, presenting the project decisions, architecture and algorithms. Chapter 5 describes the evaluation methodology, while Chapter 6 presents the results, also bringing discussions and limitations. Finally, Chapter 7 brings the conclusion of the document, highlighting contributions and future work, in addition to analyzing the initial hypothesis.

2 THEORETICAL FOUNDATION

This chapter presents the main topics related to this work. Starting with Section 2.1 discusses the concepts of IoT and the definitions of Fog and Edge used in this work. The next part is Section 2.3, which describes AI and its evolution towards ML and, later, the creation and development of federated learning. Lastly, Section 2.4 approaches the smart cities context and discusses the possibilities regarding smart hospitals.

2.1 Internet of things

The Internet of Things (IoT) is a comprehensive environment that interconnects a large number of heterogeneous physical objects or things to the Internet. An IoT system mainly follows the architecture of the Cloud-centric Internet of Things (CIoT), in which the physical objects are represented in the Web resources managed by the servers in the global Internet. The standard IoT system involves three major technologies: embedded systems, middleware, and cloud services. Although the CIoT model is a common approach to implementing IoT systems, it faces growing challenges in IoT. Specifically, CIoT faces challenges in bandwidth, latency, uninterrupted, resource-constraint, and security (BUYYA; SRIRAMA, 2019).

In the case of bandwidth, the increasingly large and high-frequent rate data produced by objects in IoT will exceed the bandwidth availability, and entirely relying on the distant Cloud to manage things becomes impractical. Latency faces the challenges of achieving the requirement of controlling the end-to-end within tens of milliseconds. Applications that rely upon low-time responses cannot afford the consequences of latency in CIoT. The long distance between the Cloud and the front-end IoT devices can face unstable and intermittent network connectivity issues. Uninterrupted is crucial when a function cannot correctly execute due to the disconnection. Many front-end devices are typically resource-constrained and unable to conduct complex computational activities; thus, CIoT systems typically require front-end devices to continuously broadcast their data to the Cloud. However, such a design could be more feasible in many battery-powered devices since end-to-end data transfer via the Internet still requires significant energy. Many constrained front-end devices may need more resources to defend themselves against cyberattacks. Outdoor-based front-end devices, in particular, which rely on the distant Cloud to keep them updated with security software, can be targets for attackers, as the attackers are capable of performing a malicious activity at the edge network where the front-end devices are located and the Cloud does not have complete control over it (HERRERO, 2022).

In summary, CIoT has come a long way in the last decade. Several approaches have tried to extend centralized computing to a more geo-distributed manner. Industry-led fog computing architecture has gained the most attention. Academics and researchers have also explored mobile cloud computing models. These research vanguards resulted in multiple proposals for the meaning of Cloud, Fog, and Edge devices. For this reason, we present Figure 2, which eluci-

Figure 2: Network paradigm definition terminology

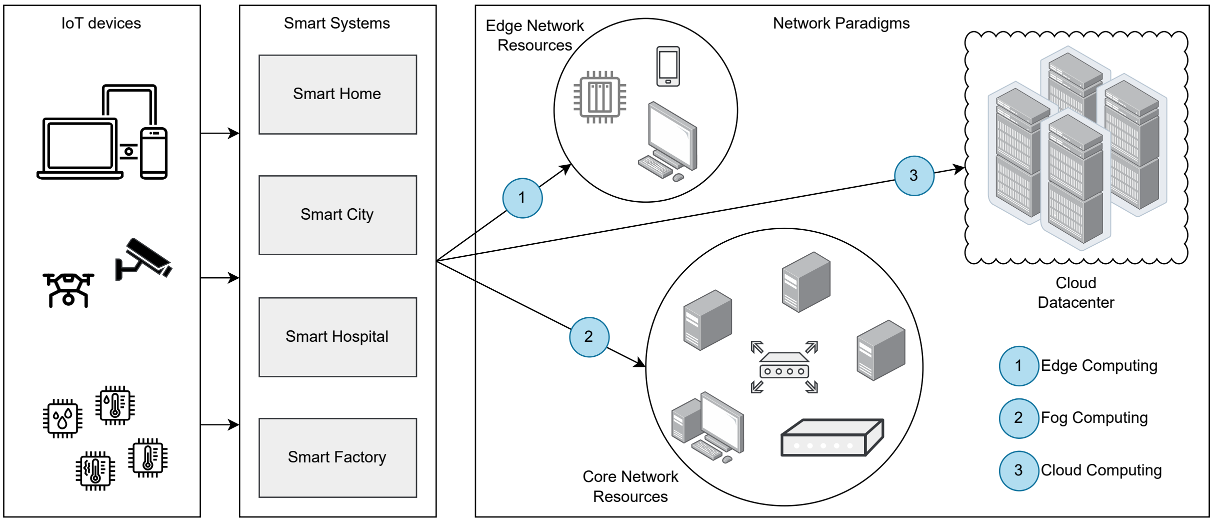

Source: Adapted from (BUYYA; SRIRAMA, 2019)

dates the regions and devices we are addressing as each specified terminology. In this work, we consider edge computing as everything executed near the user, such as its personal computer or smartphone. Fog computing happens when the processing is made in the local network or a set of devices connected to the same environment. Lastly, cloud computing occurs when the load is sent to execute in a data center. There is also the Fog-Cloud paradigm when the fog executes and loads part of its processing to the Cloud. However, we will not address this topic in this work.

2.2 Fog and Edge Computing

Fog computing is a conceptual model that addresses all the possibilities to extend the Cloud to the edge network of CIoT. The industry-led fog computing architecture has gained the most attention among the various approaches. For example, a fog-enabled IoT system can distribute the simple data-classification tasks to the IoT devices at the edge and assign the more complicated context reasoning tasks to the fog gateway devices. The decision of where the system should assign the tasks depends on efficiency and adaptability (BUYYA; SRIRAMA, 2019).

In the early phases of fog, mist computing was an alternative. However, recent publications (BUYYA; SRIRAMA, 2019) have defined mist as a subset of fog. As a result, mist elaborates on the necessity for deploying computational mechanisms to the IoT's extreme edge, where IoT devices are placed, to decrease communication delay between IoT devices in milliseconds. Mist computing is primarily motivated by the need to provide IoT devices with self-awareness through self-organizing, self-managing, and various self-mechanisms. As a result, IoT devices will be able to work constantly even when the Internet connection is inconsistent. In other words, fog can only deploy and manage itself by integrating edge computing technologies.

Fog and edge computing (FEC) supplements the Cloud in IoT by bridging the gap between

Figure 3: Agents interact with environments through sensors and actuators.

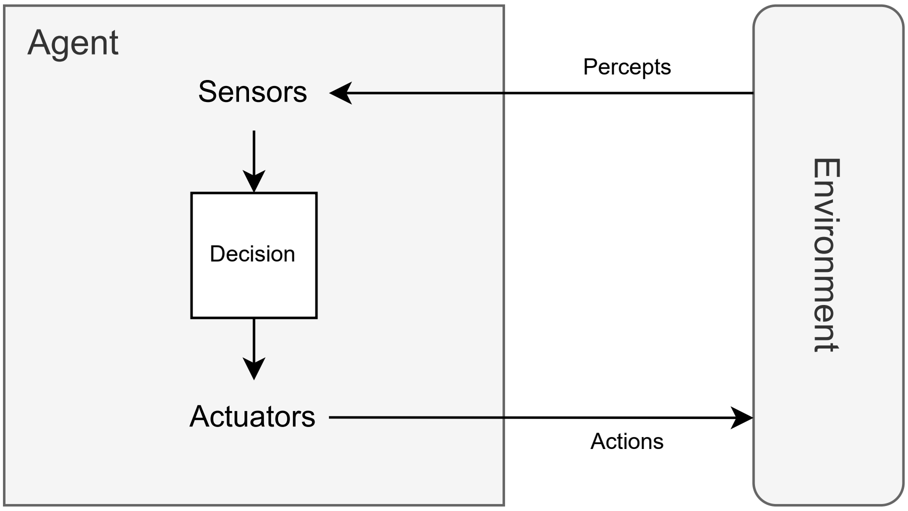

Source: Adapted from (RUSSELL; NORVIG, 2005).

the Cloud and the devices to offer service continuity. In FEC architectures, the edge is the closest hardware to the user. Furthermore, there are five main advantages when using edge computing. They are security, cognition, agility, short latency, and efficiency. Edge devices are less exposed to attacks in the local network. In traditional CIoT architectures, even sensitive data will go through the network to be processed, and this can be a problem for sensitive smart systems, such as smart hospitals.

FEC architectures also allow decision-making from the user, the cognition guaranteeing not only adaption on the edge device but also agility when compared to cloud services that depend on business holders to establish, deploy, and manage the infrastructure. Lastly, if the time to respond is crucial, the edge has the lowest latency of all of the players in CIoT architectures. Moreover, in smart systems, large data flows are expected, and the execution in the local device can increase the efficiency in execution time when communication bandwidth is limited.

2.3 Artificial intelligence

AI is one of the most recent scientific and engineering fields. Work began in earnest shortly after World War II, and the name was created in 1956. AI now includes a wide range of subfields, from the general (learning and perception) to the specific (playing chess, proving mathematical theorems, writing poetry, driving a car on a crowded street, and diagnosing diseases). AI applies to any intellectual task (RUSSELL; NORVIG, 2005). In the field of AI, an agent is defined as anything that perceives its environment through sensors and acts on that consciousness through actuators. The term percept refers to the agent's perceptual inputs at any given time. The percept sequence of an agent is the complete history of everything the agent has ever perceived. Figure 3 illustrates this concept.

According to (RUSSELL; NORVIG, 2005), there are five main types of agent programs.

The first is the Simple reflex agents. The simple reflex agent is the most basic type of agent. These agents choose actions based solely on the current percept, disregarding the rest of the percept history. Even in complex environments, simple reflex behaviors occur. Consider yourself the driver of an automated taxi. If the car in front of you brakes and its brake lights illuminate, you should take notice and begin braking. To put it another way, some processing is done on the visual input to determine the condition "The car in front is braking." This activates some previously established connection in the agent program to the action "initiate braking." This is known as the if-then rule.

The second type of agent is the Model-based reflex. The most effective way to deal with partial observability is for the agent to keep track of what it cannot see. That is, the agent should maintain some sort of internal state dependent on the percept history and thus reflects at least some of the current state's unobserved aspects. For the braking example, the internal state is limited to the previous frame from the camera, allowing the agent to detect when two red lights at the vehicle's edge turn on or off simultaneously. Updating this internal state information over time necessitates the inclusion of two types of knowledge in the agent program. First, we need to know how the world evolves independently of the agent. Second, we need to know how the agent's actions affect the rest of the world.

The third approach is the goal-based agents. Understanding the current state of the environment is only sometimes sufficient to decide what to do. For example, the taxi can turn left, right, or continue straight at a road junction. The destination of the taxi determines the correct decision. In other words, in addition to a current state description, the agent requires goal information that describes desirable situations, such as arriving at the passenger's destination. The agent program can use this information in conjunction with the model to select actions that will achieve the goal. A goal-based agent could reason that if the car in front of it has its brake lights on, it will slow down. Given how the world typically unfolds, the only action that will achieve the goal of not colliding with other cars is to brake. Although the goal-based agent appears less efficient, it is more adaptable because the knowledge supporting its decisions is explicitly represented and mutable.

The fourth agent is utility-based. In most situations, more than goals is required to generate high-quality behavior. Many action sequences, for example, will get the taxi to its destination, but some are faster, safer, more reliable, or less expensive than others. Goals are simply a binary distinction. The utility function of an agent is an internalization of the performance measure. If the internal utility function and the external performance measure agree, then an agent choosing actions to maximize its utility will be rational regarding the external performance measure. Furthermore, goals are insufficient in two cases, but a utility-based agent can still make rational decisions. First, when competing goals can only be met in part, the utility function specifies the appropriate trade-off. Second, when the agent has several goals to pursue, none of which can be achieved with certainty, the utility allows the agent to weigh the likelihood of success against the importance of the goals.

Figure 4: Learning agent overview

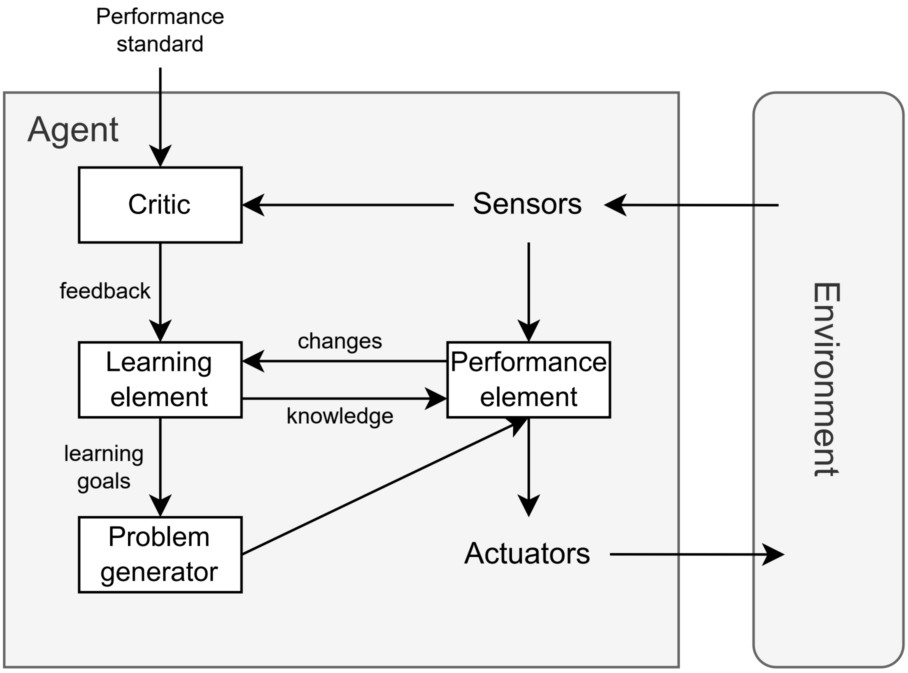

Source: Adapted from (RUSSELL; NORVIG, 2005).

Lastly, we have the Learning agents. This is now the preferred method for developing cutting-edge AI systems in many fields. As previously stated, learning has another advantage: it allows the agent to operate in initially unknown environments and become more competent than its initial knowledge alone might allow. A learning agent is made up of four conceptual components. The most crucial distinction is between the learning element, which is in charge of improving, and the performance element, which is in charge of choosing external actions.

Four conceptual parts construct a learning agent. The performance element is what we previously thought of as the entire agent: it perceives and decides on actions. The critic informs the learning element of the agent's performance in relation to a predefined performance standard. The critic is required because the percepts alone do not indicate the agent's success. The learning element takes the critic's feedback on how the agent is performing and determines how the performance element should be modified to perform better in the future. The problem generator is the final component of the learning agent. It is in charge of suggesting actions that will result in new and informative experiences. Figure 4 summarizes the ideas of the learning agents.

In this proposal, we will be mainly addressing the definition of learning agents. An agent learns if its performance on upcoming tasks increases due to its observations of the outside world. Learning can range from trivial to complex tasks. We will focus on one type of learning problem that appears limited but has broad applicability: learning a function that predicts the output for new inputs from a collection of input-output pairs.

2.3.1 Machine learning

An agent learns if it improves its performance on future tasks after making observations about the world (RUSSELL; NORVIG, 2005). Learning can range from the mundane, such as writing down a phone number, to the profound and complex, such as image recognition. There are three main reasons why an agent should learn. For starters, the designers cannot anticipate every possible situation in which the agent may find itself. A maze-navigating robot, for example, must learn the layout of each new maze it encounters. Second, the designers can only anticipate some changes over time; a program designed to forecast stock market prices for tomorrow must learn to adapt when conditions shift from boom to bust. Third, human programmers only sometimes know how to create a solution. Three types of feedback determine the three main types of learning:

Unsupervised learning occurs when the agent learns patterns in the input without explicit feedback. Clustering is the most common unsupervised learning task: detecting potentially useful clusters of input examples. A taxi driver, for example, may gradually develop an understanding of "good traffic days" and "bad traffic days" without ever being given labeled examples of each. The agent learns from a series of reinforcements—rewards or punishments—in **reinforcement learning**. For example, the lack of a tip at the end of the journey indicates to the taxi driver that something went wrong. The two points for a win at the end of a chess game indicate to the agent that something went well. It is up to the agent to determine which of the preceding actions was most responsible for the reinforcement.

The agent observes some example input-output pairs and learns a function that maps from input to output in **supervised learning**. The inputs in component 1 above are percepts, and the output is provided by a teacher who says, "Brake!" or "Turn left." Component 2's inputs are camera images, and the outputs are again from a teacher saying, "That's a bus." The braking theory in 3 is a function of states and braking actions to stopping distance in feet. The output value is obtained directly from the agent's percepts (after the fact). In this case, the environment is the teacher.

In this proposal, we will be tackling the supervised learning method. Even though there are approaches to using reinforcement and unsupervised learning methods with FL, those topics are out of the scope of this research. That said, the current state of the art in supervised learning is around implementing different types of artificial neural networks.

2.3.1.1 Artificial neural networks

Neural network learning methods provide a robust approach to approximating real, discrete, and vector-valued target functions. Artificial neural networks (ANNs) are among the most effective learning methods for certain types of problems, such as learning to interpret complex real-world sensor data. For example, the back-propagation algorithm has proven successful in many practical problems, such as learning to recognize handwritten characters, recognizing spoken words, and recognizing faces, all of this back in the 1990s (MITCHELL, 1997).

The study of ANN has been inspired partly by the observation that biological learning systems are built on very complex webs of interconnected neurons. In a rough analogy, ANN is built out of a densely interconnected set of simple units, where each unit takes several real-valued inputs (possibly the outputs of other units) and produces a single real-valued output (which may become the input to many other units).

One type of ANN system is based on a unit called a perceptron, illustrated in Figure 5. A perceptron takes a vector of real-valued inputs, calculates a linear combination of these inputs, then outputs a one if the result is more significant than some threshold and -1 otherwise. Learning a perceptron involves choosing values for the weights (wo, . . . , wn). Therefore, the space H of candidate hypotheses considered in perceptron learning is the set of all possible real-valued weight vectors.

Figure 5: A perceptron

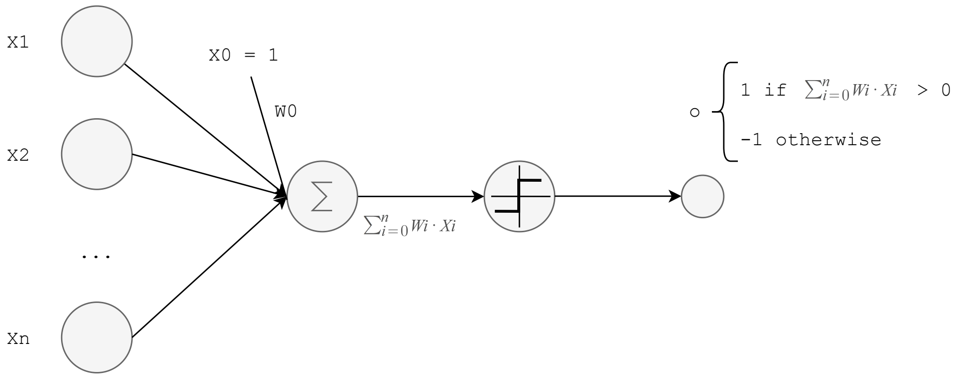
X1
X2
...
Xn
X0 = 1
W0
$\sum$
$\sum_{i=0}^{n} Wi \cdot Xi$
o
1 if $\sum_{i=0}^{n} Wi \cdot Xi$ > 0
-1 otherwise

Source: Adapted from (MITCHELL, 1997)

2.3.2 Federated learning

Federated learning, a paradigm at the intersection of machine learning, data privacy, and distributed systems, has emerged as a powerful approach to address the challenges of training machine learning models on decentralized data sources. Traditional machine learning methods often require centralizing data, which can raise concerns about data privacy, security, and legal compliance. Federated learning, however, offers a compelling alternative by enabling model training across multiple locations while keeping the data itself distributed and secure.

In the realm of federated learning, two primary processes come into play: model training and inference. During model training, information is exchanged between participating parties, but crucially, the raw data remains localized and protected. This safeguard ensures that sensitive or private data is never exposed during the collaborative learning process. Once the model is trained, it can be retained by a single party or shared among several, depending on the specific use case and agreements in place (YANG et al., 2019).

The advantages of federated learning are manifold. Firstly, it is inherently designed to protect user privacy and data security. By allowing model updates to be the only information shared, it sidesteps the need for raw data transfer, a practice that can introduce significant privacy and security risks. Secondly, federated learning promotes collaboration among multiple parties, enabling them to collectively train machine learning models that are superior to what any individual entity could achieve independently. This collaborative approach can lead to enhanced model accuracy and performance.

However, like any technology, federated learning has its challenges. One key consideration is the communication infrastructure between the owners of local data and the central aggregation server. It must be fast and reliable to support the timely exchange of model updates. Additionally, the scalability of federated learning can be complex, especially in scenarios where a large number of participants, such as mobile users, are involved. The potential for many participants can introduce instability and unpredictability into the system.

Data distribution disparities among federated learning participants can also pose issues. Different parties may have data with non-identical distributions, unbalanced numbers of data samples, or unique data characteristics. These disparities can lead to biased models or even hinder the training process altogether. Consequently, when designing a federated learning system, it is imperative to account for these variations and implement strategies to mitigate their impact on model quality.

The establishment of a federated learning ecosystem extends beyond technical considerations; it is an economic challenge as well. Creating mechanisms that ensure fair and transparent profit distribution among participating parties is crucial for fostering long-term engagement and incentivizing involvement. These mechanisms must also serve as deterrents against malicious participants who may seek to undermine the integrity of the federated learning process.

As federated learning continues to evolve, it draws from a diverse array of disciplines. Its foundation in machine learning and statistics is complemented by insights from information security, encryption techniques, model compression, and game theory. Economic principles and mechanism design play a pivotal role in structuring federated learning ecosystems that encourage collaboration while safeguarding against potential pitfalls. This interdisciplinary approach reflects the multifaceted nature of federated learning and its expanding applications in diverse domains.

2.4 Healthcare Informatics and ML in Healthcare

Health Informatics started as Medical and Nursing Informatics in the 1970s, a period described as undergoing exponential development due to the increasing availability of steadily less expensive hardware, more robust software, and the introduction of microcomputers (HOVENGA, 2010). During the 1980s, there was much interest in using computers to help with medical decisions, including artificial intelligence. System linking emerged in 1989 when multiple disciplines collaborated to develop integrated systems utilizing new database technology and network power.

While Health Informatics strives to define its role in healthcare, other healthcare professionals continue incorporating the technologies into their respective fields. Computing systems, for example, are widely used in radiological imaging. Among the lessons to be drawn from the history of health informatics is that the discipline of health informatics must be aware of and involved in the goals and activities of health care itself. Technologies are becoming more widely available, with ever more powerful tools enabling healthcare workers to create systems for their own benefit easily.

With the expansion of multi-modality data in the last decade, the role of data analytics in health informatics has expanded fast. This has also increased interest in developing analytical, data-driven models in health informatics based on ML (RAVÌ et al., 2016). ML has facilitated the development of more data-driven solutions in health informatics by allowing for the automatic generation of features, reducing the amount of human intervention in this process. This benefits many health informatics problems and has eventually supported a significant leap forward for unstructured data such as those generated by medical imaging, medical informatics, and bioinformatics.

In a growing industry of smartwatches, smart wristbands, and devices that constantly collect a plethora of health data, the use of ML to analyze this data is gaining traction. ML may be the answer to both lowering healthcare costs and improving patient-doctor relationships. ML and big data solutions can be used for various health-related purposes, including assisting doctors in developing more personalized prescriptions and treatments for patients and assisting patients in determining when and if they should schedule follow-up appointments (BHARDWAJ; NAMBIAR; DUTTA, 2017).

ML techniques applied to EHR data can yield actionable insights ranging from improving patient risk scoring systems to predicting disease onset and streamlining hospital operations. Statistical models that take advantage of the variety and richness of EHR-derived data (rather than a small set of expert-selected and/or traditionally used features) are still uncommon. However, they represent an exciting avenue for future research. New data sources, such as wearables, bring with them new opportunities and challenges (CALLAHAN; SHAH, 2017b).

3 RELATED WORK

Based on the research topic, works that encompass the same scope or that solve similar problems were analyzed. These serve as a basis for possible solution attempts. That said, a search for related works was conducted to survey the available scientific works. The remainder of this chapter presents in Section 3.1 the process of selecting and choosing related works. Section 3.2 presents the chosen papers and their proposals. Finally, Section 3.3 analyzes the works and open gaps.

3.1 Selection process

The selection process was based on three steps. 1) Definition of the search. 2) Job filtering. 3) Analysis and correlation of topics. The search strings were created based on the three main topics covered by this work: FL, FOG/EDGE, and HC. The searches were based on the last five years, using combinations of the mentioned terms such as "federated learning on healthcare," "fog/edge federated learning on smart hospitals," and "federated learning on edge and fog for healthcare informatics."

We chose five databases to conduct the searches: ACM, IEEE Xplore, Google Scholar, Science Direct, and Springer. The search was based on articles published from 2017 to 2022. In total, we found more than 300 articles in the databases. The next step was to carry out the filtering of the works. In the first stage of selecting pages, duplicate works were removed, and articles containing more than five pages were selected. Then, the titles and abstracts of the works were evaluated to validate the cohesion with the theme sought. The search focused on the architecture of federated learning projects for healthcare, focusing mainly on edge/fog computing. The surveys were removed from the related works. However, they were separated for further analysis as open issues were found that could be tackled. By objective, all the works that presented the architecture proposal related to the search theme were the last ones, thus defining a scope. Finally, nine works were left at the end of the filtering processes.

3.2 Analyzing the State-of-the-Art

Sanyal et al. (SANYAL et al., 2019) present a framework for executing federated learning in IoMT devices constrained in power and computational capabilities. They propose an alternative solution to the issues of energy efficiency, latency, and privacy for resource-constrained that are present on those. The framework predicts a data matrix using aggregated model average, computes and delivers filter parameters for the IoMT devices, and performs decision-making using the aggregated data matrix. They also present a workaround for the generated eigenvalue perturbation of the data matrix, using Matrix Perturbation Theory to help to fix this issue.

Guo et al. (GUO et al., 2020) propose a federated edge learning system for efficient privacy-preserving mobile healthcare. Specifically, an edge-based training task offloading strategy to improve training efficiency. During model training, the model uses Gaussian perturbation to prevent gradient leaking. This is necessary since they use an offloading mechanism without using a homomorphically encrypted weight matrix when aggregating the models. They implemented a system prototype to evaluate the training efficiency, inference performance, and noise sensitivity. Overall, the architecture reduces the resource requirements of the mobile device and improves the efficiency of training models, compromising a small quantity of the model performance due to the privacy scheme.

Hakak et al. (HAKAK et al., 2020) propose an Edge-assisted data analytics framework that uses Federated Learning to re-train local ML models using user-generated data. This framework could leverage pre-trained models to extract user-customized insights while preserving privacy and Cloud resources. They discuss applications such as disease management/prevention, mental health tracking, and real-time health monitoring, as well as the challenges of this technology, for example, cyberattacks and acceptance of the technology. The paper is only a conceptual framework, so the authors do not present an evaluation methodology or further results.

Rahman et al. (RAHMAN et al., 2020) present a lightweight hybrid FL framework in which blockchain smart contracts manage the edge training plan, trust management, and authentication of participating federated nodes, the distribution of global or locally trained models, the reputation of edge nodes, and their uploaded datasets or models. The framework also supports the complete encryption of a dataset, the model training, and the inferencing process. Each federated edge node performs additive encryption, while the blockchain uses multiplicative encryption to aggregate the updated model parameters. The framework supports lightweight differential privacy to support the total privacy and anonymization of the IoMT data. Moreover, this work mainly focuses on the security and privacy of FL in IoT rather than the FL execution method itself.

Wu et al. (WU et al., 2020) develop a cloud-edge-based federated learning framework for in-home health monitoring, which learns a shared global model in the cloud from multiple homes at the network edges and achieves data privacy protection by keeping user data locally and using homomorphic encryption aggregation. To cope with the imbalanced and non-independent and identically distributed distribution inherent in the user's monitoring data, they designed a generative convolutional autoencoder to achieve accurate and personalized health monitoring by refining the model with a generated class-balanced dataset from the user's personal data.

Zhao et al. (ZHAO et al., 2020) propose a system that uses edge devices to implement activity and health monitoring locally and applies federated learning to facilitate the training process. The devices use the Databox platform to manage sensor data collected in people's homes, conduct activity recognition locally, and collaboratively train a deep neural network model without transferring the collected data into the cloud. The paper results show that the processing time of local inference on an edge device is acceptable. Meanwhile, the inference accuracy of the system can converge to a sufficient and stable level after a few rounds of communication between the clients and the server. Still, this architecture remains susceptible to DLG.

Połap et al. (POłAP et al., 2021) propose an architecture of a system that ensures the security of private data and allows the addition and modification of the used classification methods. The main advantages of the proposed system are based on the implementation of blockchain technology elements and threaded federated learning. The individual elements are located on the agents who exchange information. Additionally, they propose building an agent with a consortium mechanism for classification results from many ML solutions. This offers a new model of agents that can be implemented as a system for processing medical data in real-time. They compared the approach with other methods and showed that the proposition could improve the IoMT solutions by presenting a new idea of a multi-agent system that can separate different tasks like security or classification and, as a result, minimize operation time and increase accuracy.

Xue et al. (XUE et al., 2021) explicitly consider the problem of decentralized clinical decision problem for sequential clinical treatment. Using a double deep Q-Network based on a fully decentralized federated framework enabled by an integrated system named SMEC. Provides a way to infer real-time treatment policy from large amounts of distributed observational electronic medical records. The system's performance shows good results for real-time sequential clinical treatment policy for patients when implementing clinical decision support systems.

Wang et al. (WANG et al., 2022) propose a privacy protection scheme for federated learning under edge computing. They first propose a lightweight privacy protection protocol based on a shared secret and a weight mask based on a random mask scheme of secret sharing. It can protect gradient privacy without losing model accuracy and resist equipment dropping and collusion attacks between devices. Secondly, they design an algorithm based on a digital signature and hash function, which achieves the integrity and consistency of the message, as well as resisting replay attacks. Lastly, they propose a periodic average training strategy to prove that their model is faster than deferential privacy ones. Meanwhile, compared with classical federated learning, their system efficiency is slightly lower but ensures data safety.

3.3 Analysis and Opportunities

From the selected related works, a grouping was performed. This grouping is based on the most common topics covered in the works. We aim to highlight the less tackled areas that arouse the ideas of open works. The topics highlighted were: proposal of aggregation method, use of fog-computing, use of edge-computing, use of encryption methods for privacy, architecture focused on smart hospitals, architecture focused on home healthcare, and analysis in heterogeneous data. Table 1 summarizes the ideas presented by the analyzed works. Each column represents one of the mentioned topics, respectively. Each marking with "X" represents whether the referring work addresses one of the mentioned topics.

Table 1: Related work analysis.

Reference	Aggregation	Fog	Edge	Encryption	Hospital	Home	Heterog. data
(SANYAL et al., 2019)	X	X	X		X		
(GUO et al., 2020)	X	X			X		
(HAKAK et al., 2020)			X		X	X	
(RAHMAN et al., 2020)	X	X	X	X	X		
(WU et al., 2020)	X	X	X	X		X	X
(ZHAO et al., 2020)			X			X	
(POłAP et al., 2021)			X		X		
(XUE et al., 2021)	X	X	X	X	X		
(WANG et al., 2022)		X	X	X	X		

Source: Made by the author

Looking at the current state-of-art, it is noticeable that there are open issues to be tackled. We mainly point out the lack of model specification based in data heterogeneity. Only one of the analyzed papers brings user models customization as the main topic and focuses on having a model that adapts better for each system user. Moreover, in IoMT, privacy is crucial, and since the DLG attacks a system without the use of homomorphic encryption or any other differential privacy approach is susceptible to data reconstruction. Some of the open questions pointed out coincide with those found in the surveys that served as a motivational basis for the beginning of this dissertation. Thus, this means the searches for related works are under the chosen scope and available themes in the literature. Reiterating the main gaps found were:

- Lack of model adaptation based on data heterogeneity in federated learning systems;
- Only a couple of works used some encryption for gradient leakage prevention;
- Multi-aggregation methods were not found in the literature.

4 PROPOSAL: FEDERATED HOSPITAL MODEL

Based on the related works analysis and the surveys in the literature, here we are presenting the Federated Hospital model in order to fill out the aforementioned gaps. In particular, the proposed model executes federated learning services in the context of smart hospitals taking into account heterogeneity data distribution from the users and addressing this problem with a multi layer architecture proposal. The rest of this chapter is organized in other four sections. First, Section 4.1 presents some motivational situations, highlighting where we can employ the developed model. Secondly, Section 4.2 presents the project decisions, detailing our premises and project scope. Moreover, Section 4.3 describes the architecture, giving examples of problems and how the solution addresses data heterogeneity in different scenarios. Lastly, Section 4.4 presents the algorithms and the system operability.

4.1 Use cases

In this section, we highlight some motivation scenarios where the current proposal could bring health benefits for the society. The proposed multilevel architecture for federated learning in smart hospitals offers a plethora of use case scenarios, showcasing its versatility and potential impact. Firstly, in the context of patient monitoring, the architecture can be utilized to train personalized models for each hospital section, capturing unique patient behaviors and characteristics. This enables accurate real-time monitoring of vital signs, disease progression, and treatment responses, leading to timely interventions and improved patient outcomes.

Secondly, in the domain of predictive analytics, the architecture allows for localized model training on historical patient data within specific hospital sections. This empowers healthcare professionals to develop predictive models tailored to their section's patient population, enabling early detection of disease patterns, optimal resource allocation, and effective preventive measures.

Lastly, in the area of medical image analysis, the architecture can facilitate the training of specialized models for different imaging modalities within each hospital section. This enhances diagnostic accuracy and efficiency, supporting radiologists in detecting abnormalities, prioritizing critical cases, and facilitating timely treatment decisions. The proposed architecture addresses critical challenges in these scenarios, such as the need for personalized models, localized expertise, and efficient utilization of data resources, making it a valuable solution for improving healthcare delivery in smart hospitals.

4.2 Design Decisions

In developing the proposed Federated Hospitals, several design decisions were made to guide the implementation and evaluation process. The key design decisions are as follows:

1. Prototype Development: To validate the effectiveness and feasibility of the proposed architecture, a prototype will be developed. This prototype will simulate the expected execution flow and aggregation steps of the architecture within a controlled environment. The prototype will serve as a proof of concept and enable us to assess the architecture's performance and potential benefits in a practical setting.

2. Evaluation Metrics: The evaluation of the proposed architecture will focus on comparing the performance of the initial root model and the converged ward-specific models. The evaluation metrics will include classification accuracy and convergence speed. By measuring these metrics, we can assess the improvement achieved by the multilevel federated learning approach and the effectiveness of the specialized ward-specific models.

3. Scope Limitations: Due to the focused nature of this research project, certain aspects such as network security, gradient leakage, and privacy-preserving techniques will not be explicitly addressed in this stage. While these aspects are important considerations for real-world implementation, their exploration is beyond the immediate scope of this proposal. However, the impact of these factors can be considered in future works to enhance the architecture's robustness and privacy preservation capabilities.

By making these design decisions, we aim to develop a prototype that demonstrates the expected flow of execution and aggregation steps of the Federated Hospital. Through the evaluation of the initial root model and the converged ward-specific models, we can gain insights into the architecture's effectiveness and potential for improving model performance in healthcare settings. While the scope limitations focus the project's efforts, future works can explore network security, gradient leakage, and privacy-preserving techniques to further enhance the architecture's practical implementation and ensure the privacy and security of sensitive healthcare data. Also, it is worth mentioning that the architecture was designed and developed based on the needs and open gaps in the topic of FL in smart hospitals. However, this does not preclude its use in another scenario.

4.3 Architecture

To address the challenges posed by the heterogeneity of data and the need for specialized services in smart hospitals, we propose a novel multilevel architecture for federated learning. Our architecture is designed to leverage the similarities of behaviors within hospital sections while capitalizing on the specialized services provided in each area. This multilevel approach consists of three levels: local training, global aggregation, and local-global refinement. The overview of the architecture is presented in Figure 6.

At the local level, models are trained within specific hospital sections using localized data. This enables the capture of section-specific patterns and nuances, enhancing the model's performance for local users. By training models on data specific to each section, the architecture accounts for variations in patient demographics, disease prevalence, and treatment protocols. This localized training facilitates the customization of services to meet the specific needs of each hospital section, ultimately promoting better patient outcomes.

Moving to the global level, the local models are aggregated into a centralized global model using the federated averaging technique. This aggregation process combines the expertise from multiple hospital sections, enabling the global model to achieve better generalization and capture broader trends across the entire hospital. By incorporating diverse perspectives and knowledge from various sections, the global model ensures a balanced representation and enhances its ability to recognize global patterns and trends in healthcare data.

Once the global model is aggregated, it is sent back to the section level for local-global refinement. At this stage, the global model is aggregated with the local models again, but using a weighted averaging approach. This refinement step aims to strike a balance between the global knowledge captured by the centralized model and the specialized expertise gained at the local level. By employing weighted averaging, the refinement process ensures that the final models maintain both the global perspective and the section-specific specialization, thus optimizing the accuracy and relevance of the models for each hospital section.

Figure 6: Architecture overview of the system

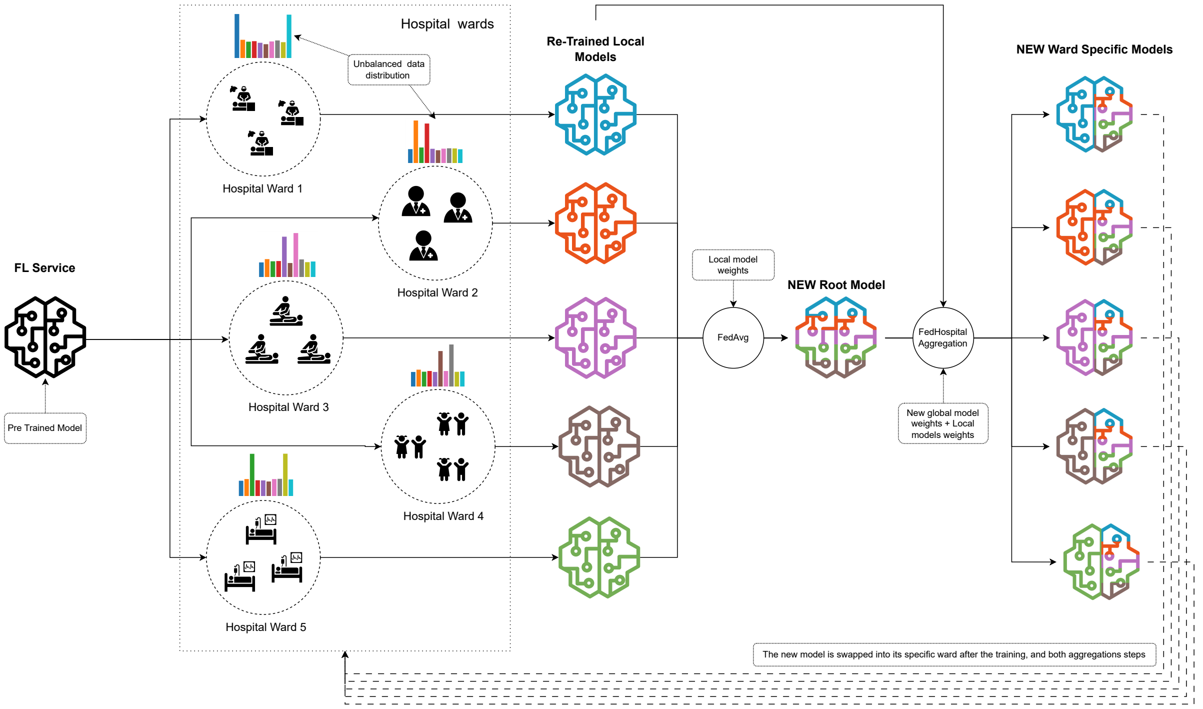

Source: Made by the author

The proposed architecture offers several advantages. Firstly, it leverages the inherent similarities within hospital sections, allowing for targeted model training and service customization. This promotes personalized healthcare delivery and improved patient outcomes. Secondly, the global aggregation step enables knowledge sharing and enhances the global model's generalization capabilities by incorporating insights from various sections. Lastly, the local-global refinement process ensures a balance between local specialization and global representation, resulting in accurate and context-aware models.

The proposed architecture holds immense potential for smart hospitals. It addresses the challenges of heterogeneity and specialized services by capitalizing on local expertise while capturing global trends. Through this multilevel approach, our architecture facilitates enhanced model performance, optimized resource utilization, and improved decision-making in intelligent healthcare systems.

4.4 Algorithms

In this section we describe the algorithms that will be in use by the architecture. The main algorithm that will be in use is the FedAvg (MCMAHAN et al., 2016). The algorithm utilizes FedAvg for local aggregation and global model update, and weighted averaging for the aggregation of the local model with the new global model.

Algorithm 1 Multilevel Federated Learning Aggregation for Smart Hospitals

```
Input: Local training datasets $D_1, D_2, ..., D_N$ for $N$ hospital sections
Output: Global model $M_{\text{global}}$
Initialize global model $M_{\text{global}}$ with random parameters;
for each round $t = 1$ to $T$ do
    for each hospital section $i = 1$ to $N$ do
        Local model update:    $M_i \leftarrow \text{FedAvg}(M_{\text{global}}, D_i)$;
    end
    Global model aggregation:    $M_{\text{global}} \leftarrow \text{FedAvg}(M_1, M_2, ..., M_N)$;
end
for each hospital section $i = 1$ to $N$ do
    Weighted averaging for local model:    $M_i \leftarrow \alpha \cdot M_i + (1 - \alpha) \cdot M_{\text{global}}$;
end
return $M_{\text{global}}$
```

In Algorithm 1, T represents the number of communication rounds, M_i denotes the local model for hospital section i, and α is a weighting factor for the weighted averaging at the local level. The FedAvg function performs the federated averaging algorithm, which aggregates the models or updates by taking their average.

This demonstrates the iterative process of training the local models within each hospital section, aggregating them at the global level, and then performing weighted averaging to combine the global model with the local models. The algorithm captures the multilevel aggregation approach of the proposed architecture, enabling specialization and collaboration among hospital sections while maintaining a cohesive global model.

5 EVALUATION METHODOLOGY

In this research, we employed a comparative evaluation methodology to assess the performance of the proposed Federated Hospitals architecture in smart hospitals. Due to the absence of well-established benchmarks for FL in healthcare, we compared our approach against two standard models: traditional ML approach and FL using FedAvg. The remainder of this chapter describes how we assembly the evaluation of the presented model. First, we detail the employed technologies to develop a prototype. Second, we detail the considered input workload. Third, we present the evaluation scenarios and their input parameters. Lastly, for each scenario we present a set of metrics for evaluation purposes.

5.1 Prototype

For the prototype, we used Python 3.11 for the modules code, Keras for the machine learning and federated learning implementation. Also, Pandas and Numpy for the data process and analysis. As depicted in Figure 7, the flowchart illustrates the step-by-step implementation of our proposed architecture, realized through a programming library. This library encompasses a well-structured set of functions that execute the various stages of the architecture's flow. Such an approach enhances manageability and flexibility, as it allows for easy parameterization and adaptation of the modules to accommodate different FL algorithms within the architecture.

Figure 7: Prototype implementation flowchart

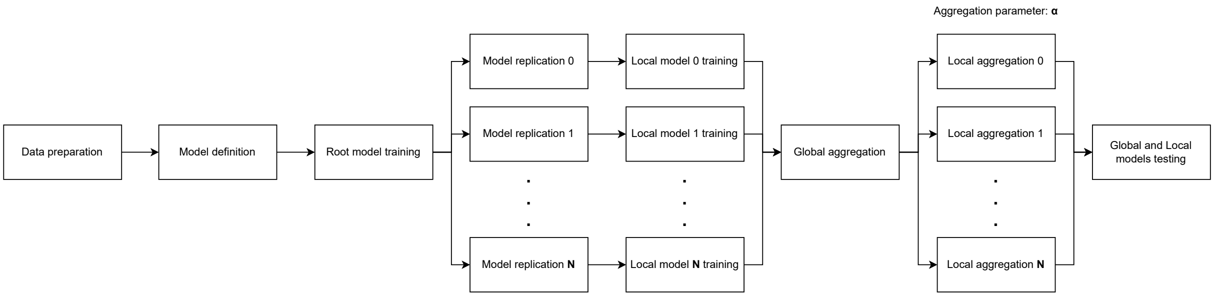

Source: Made by author

By encapsulating the flowchart into a programming library, the implementation becomes highly modular and extensible. This modularity empowers researchers and practitioners to customize and fine-tune the modules according to the specific requirements and complexities of different federated learning scenarios. As a result, the prototype offers a robust and adaptable platform for testing and experimenting with various FL algorithms within the proposed architecture.

Figure 8: Example data from MNIST dataset

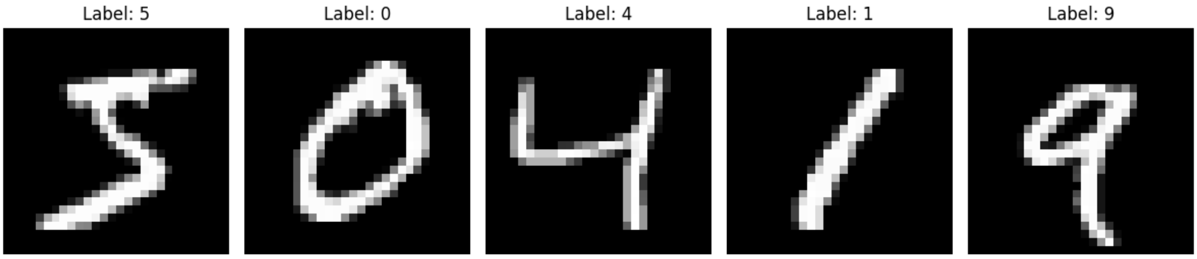

Source: Made by author

5.2 Input workload

For the benchmark of the proposal, we opt to use the MNIST dataset (DENG, 2012). The MNIST is a widely used benchmark dataset in the field of machine learning and computer vision. It consists of a collection of handwritten digit images, where each image is a grayscale 28x28 pixel representation of a single digit from 0 to 9. The dataset is divided into a training set with 60,000 examples and a test set with 10,000 examples. MNIST serves as a standard dataset for evaluating and comparing the performance of various image classification algorithms and models. It provides a diverse range of digit samples with varying styles and handwriting, making it a suitable choice for assessing the effectiveness and generalization capabilities of models trained through federated learning in the context of smart hospitals. In our case we will be using it to compare the different scenarios and how the architecture performs in comparison to standard machine learning approach, standard federated learning approach using FedAvg, and the proposed model. Figure 8 presents a sample of the dataset.

To evaluate the performance of our proposed Multilevel Federated Learning Architecture for Smart Hospitals, we partitioned the MNIST dataset into distinct subsets for training, validation, and testing. Initially, a portion of the dataset was dedicated to training the first global model. This allowed the global model to gain initial knowledge and capture broad patterns in digit classification, we created a subset containing 10,000 examples for the global model. Figure 9a shows the data distribution for the first training set. The remaining data was then divided into five groups, representing each hospital ward within the smart hospital infrastructure. Each group was carefully curated to exhibit a high level of class heterogeneity, ensuring that the models trained on these subsets would specialize in recognizing the specific patterns. This division into distinct groups enabled localized training and expertise within each section. Figure 9 presents the label distribution in multiple shards. Each shard can be view as data produced by a hospital ward while running a service. The subdivision of the test group into five divisions was done without any specific reason, as the same test could be conducted with any number of subdivisions, allowing for flexibility in evaluating the proposed architecture.

Additionally, we have a separate set exclusively for testing the models and assessing the classifier's accuracy. This evaluation set served as an objective benchmark, allowing us to gauge

Figure 9: Train data distribution for the global model and each hospital section

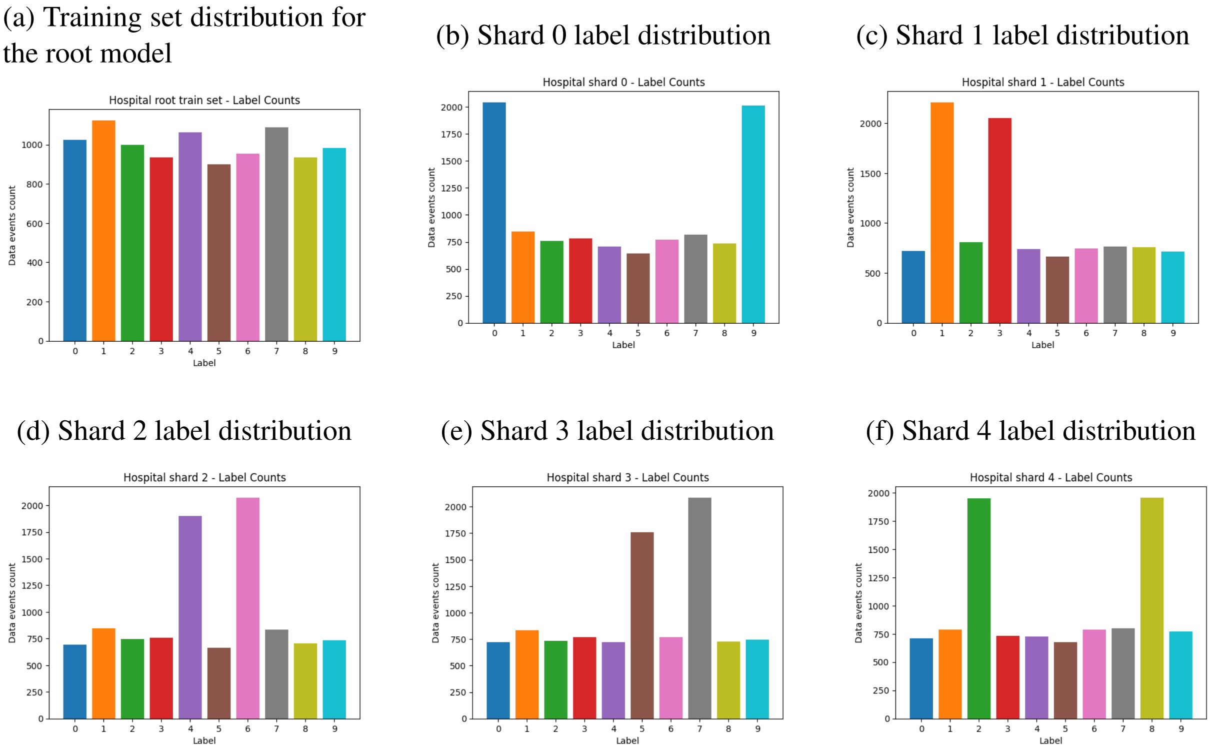

(a) Training set distribution for the root model

(b) Shard 0 label distribution

(c) Shard 1 label distribution

(d) Shard 2 label distribution

(e) Shard 3 label distribution

(f) Shard 4 label distribution

Source: Made by the author

the generalization capabilities and overall performance of the trained models across all hospital sections. By employing this data division strategy, we aimed to simulate real-world scenarios within smart hospitals while providing a robust framework for evaluating the effectiveness of our proposed architecture.

5.3 Scenarios and parameters

To comprehensively evaluate the performance of the proposed Federated Hospitals, we conducted experiments in two distinct evaluation scenarios. Their main difference is how the data is split for training and evaluation.

- Scenario 1 - Homogeneous data distribution: The data from the MNIST is split into equally distributed portions. The training procedure occurs for 200 epochs (100 epochs for global model and more 100 epochs with FL) and the evaluation is done by using the 10,000 test dataset fraction. It is assessed the time to converge and the accuracy of the final models;

- Scenario 2 - Heterogeneous data distribution: The data from the MNIST is split into unequally distributed portions, as presented by Figure 9. The training procedure occurs for 200 epochs (100 epochs for global model and more 100 epochs with FL) and the evaluation is done by using the 10,000 test dataset fraction. It is assessed the time to converge and the accuracy of the final models.

The reasoning behind these evaluation scenarios lies in simulating real-world data distribution that could be encountered in smart hospitals and healthcare settings. By designing these distinct scenarios, we aimed to comprehensively evaluate the performance of the proposed Federated Hospitals architecture under varying data distribution conditions. As well as understanding its performance under a controlled environment.

Each scenario encompassed a standard machine learning approach, a standard federated learning approach utilizing FedAvg, and our novel Federated Hospital approach. We defined specific parameters for testing and evaluating each scenario, allowing for a comparative analysis of their respective outcomes.

Every model is using the same feed-forward neural network designed for classification tasks. It takes input data with 784 features from MNIST, which are flattened into a vector. The network consists of two fully connected layers and a non-linear activation function applied in between. The network applies a linear transformation with a ReLU activation function, and then applies another linear transformation to produce the final output of size 10, which represents the scores for each class in a classification task. The batch size is set to 64, and the learning rate 0.001.

- Evaluation Model 1 - Standard Machine Learning Approach: This evaluation model is a traditional machine learning approach using the MNIST dataset. We trained a single model on the entire dataset, disregarding any hospital-specific information. The model was trained solely on the training set and evaluated on the separate testing set;

- Evaluation Model 2 - Standard Federated Learning Approach (FedAvg): For the second model, we implemented a standard federated learning approach, utilizing the popular FedAvg algorithm. We divided the MNIST dataset into multiple groups, simulating the hospital wards within the smart hospital. Each group represented a separate hospital section, and the models were trained independently on their respective local datasets. The parameters for this scenario included the fraction of data sampled from each hospital for local training, and the aggregation algorithm. The global model was updated iteratively by aggregating the local models' weights. The final model's performance was evaluated using the separate testing set;

- Evaluation Model 3 - Federated Hospital Approach: In the novel Federated Hospital scenario, we implemented our proposed Multilevel Federated Learning Architecture for Smart Hospitals. Similar to the standard federated learning approach, we divided the MNIST dataset into groups representing hospital sections. The parameters for this scenario encompassed the same parameters as the standard federated learning approach, along with additional parameters specific to our architecture. These additional parameters included the allocation of data and models across hospital sections, and the aggregation mechanism used at each level. The final model's performance was evaluated using the testing set, and compared with the results from the other scenarios.

To compare the performance of the two scenarios, we assessed metrics such as classification accuracy and convergence speed. Given the problem statement, using accuracy as a metric is a reasonable choice because it provides a straightforward measure of the model's performance in correctly classifying images. By comparing the results across the scenarios, we aimed to identify the advantages and limitations of each approach, highlighting the effectiveness of the Federated Hospital architecture in improving local model performance while maintaining a strong global representation.

6 RESULTS

The results chapter presents a comprehensive analysis of the proposed Federated Hospitals architecture in smart hospitals. Through two distinct evaluation scenarios as outlined in Section 5.3, we compare the performance of our approach against standard ML and FL models using FedAvg. The evaluation encompasses accuracy and convergence time, providing valuable insights into the effectiveness of our architecture. The remainder of this Chapter is organized as follows. Section 6.1 details our models conversion speed time and our prototype setup. We present the first evaluation scenario in Section 6.2. Moreover, we detail the results from the second scenario in Section 6.3. Lastly, we discuss the results in Section 6.4.

6.1 Models training performance

Each model was trained for 200 epochs in each scenario. The standard ML approach underwent training for 200 epochs, with a learning rate of 0.001 and a batch size of 64. For the standard FL and the FedHospital scenarios, the global model was trained for 100 epochs, followed by an additional 100 epochs for the federated learning part. The batch size and learning rate for both federated approaches were set to the values used in the standard ML training.

Figure 10 illustrates the loss incurred by each of the models under evaluation during the 200 epochs. Notably, both federated approaches consisted of five sections, and Figure 10c displays the individual loss values for each hospital section in the FedHospital scenario.

Figure 10: Loss across training sessions. For each chart, X-axis represents each epoch, and Y-axis is the loss value for that epoch.

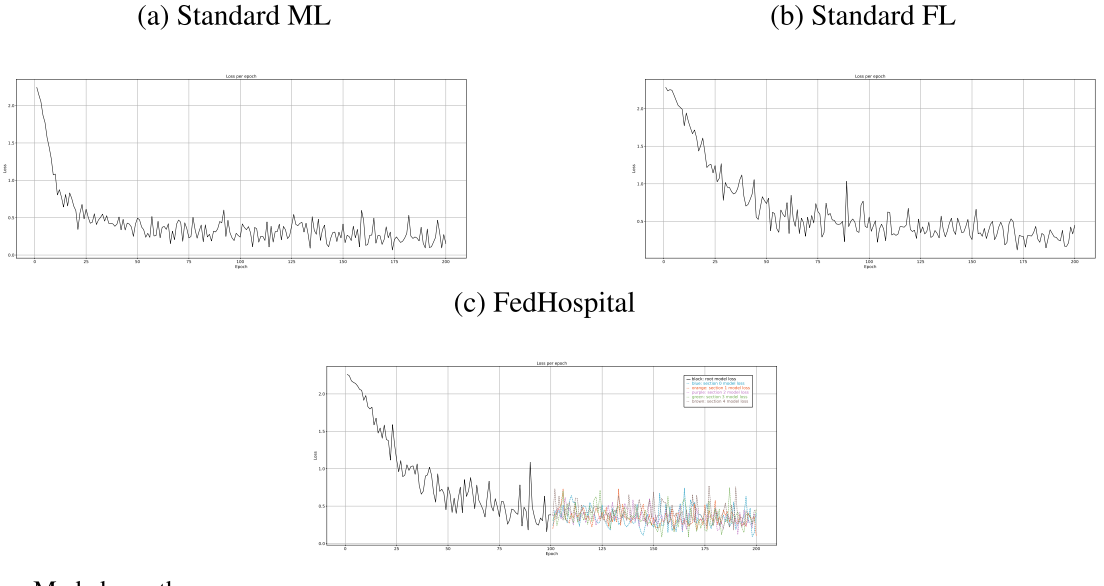

(a) Standard ML

(b) Standard FL

(c) FedHospital

Source: Made by author

The time taken to train each of the models for 200 epochs varied across the different scenarios. In the standard ML approach, the training process lasted around 800 seconds. However,

Figure 11: Time in seconds to execute 200 epochs for each model

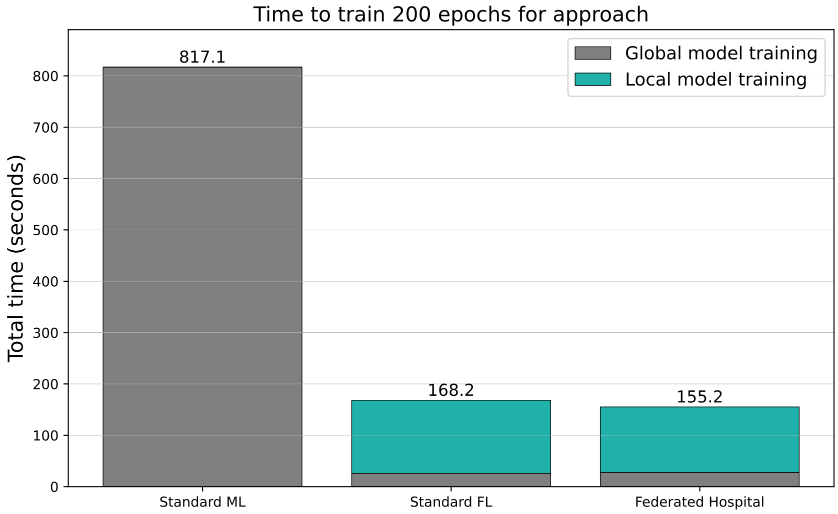

Source: Made by author

as we transitioned to FL scenarios, counter-intuitively, the training time decreased significantly, we believe this is due to the smaller batch size, the parallelism of FL, and the smaller subsets of data used for training the local models. In the standard FL scenario, the global model underwent training for 100 epochs before initiating the federated learning process, where each local model in the five sections was further trained for another 100 epochs. The decentralized nature of federated learning ended up reducing the overall training time compared to the standard ML approach. For the FedHospital scenario, the training process involved similar steps as the standard FL scenario, but with the added local-global refinement stage. The global model was aggregated with the local models at each hospital section using weighted averaging. Although this refinement introduced an additional aggregation step, the training time remained comparable to that of the standard FL scenario.

It is essential to note that the time to train each model may vary depending on the hardware and computational resources available. In our experiments, we used an Intel® Core™ i7-10750H as the processor, and utilized the NVIDIA's 1660Ti GPU to accelerate the training process. While FL methods can sometimes demonstrate faster training times, we should consider that we are using a simulated environment and the FL users may not have this type of hardware in a real world scenario. However, this does not invalidate the proposed multilevel architecture solution. In Figure 11 we disclose the time to execute the mentioned epochs.

It is important to highlight that the results presented for execution time and loss were computed using the second scenario. The same tests and analysis were conducted for the first scenario, but the outcomes were found to be nearly identical. The training time and loss per epoch exhibited negligible differences, which did not warrant further discussion. Therefore, we have chosen to present the results from only one of the scenarios to avoid redundant information. In both scenarios, the execution time and loss metrics demonstrated consistency and stability, validating the robustness of the proposal. The close alignment of results between the two scenarios provides further confidence in the architecture's performance and effectiveness across various healthcare settings and datasets. By focusing on the more informative scenario and presenting results that capture the architecture's essence, we ensure a clear and concise presentation of the research findings. The chosen approach allows us to maintain the text's clarity and readability while providing valuable insights into the architecture's potential and benefits for smart hospital applications.

6.2 Scenario 1 - Homogeneous data distribution

In this section, we present the results obtained from Scenario 1, where we distributed the data using its label in a homogeneous manner. This distribution ensured that each hospital section had the same, or nearly the same, amount of labeled data, facilitating a fair and balanced comparison between the models. We evaluated three models in this scenario: the standard machine learning (ML) model, the standard federated learning (FL) model, and our proposed Federated Hospital (α is set to 0.7). Table 2 showcases the accuracy achieved by each model based on the data label.

Table 2: Accuracy achieved by each model based on homogeneously distributed data label. The Federated Hospital (FH) accuracy is computed by averaging the accuracy of each section. We used five sections with the α set to 0.7.

	Label Accuracy (%)									
	0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
ML	98.06	97.53	**90.50**	**91.29**	**94.30**	**87.44**	**94.99**	**92.12**	**90.35**	**89.99**
FL	**98.16**	**97.62**	89.05	91.09	93.38	86.43	94.36	91.83	88.71	89.10
FH	97.88	97.49	88.39	90.85	92.64	85.47	94.17	91.45	88.19	88.9

Source: Made by the author

The standard ML model demonstrated the best overall results with 92.75% accuracy. The standard FL model, operating in a decentralized manner, showcased comparable accuracy across labels, illustrating its ability to leverage localized data and achieve balanced results. Even outperforming the standard ML for labels '0' and '1', with the final 91,97% accuracy. Lastly, our proposal enabled the architecture to leverage the specialized knowledge of each hospital section while capturing global trends. However, it did not outperform any of the other models within any specific label. The final accuracy for the Federated Hospital was 91.54%.

6.3 Scenario 2 - Heterogeneous data distribution

In this section, we present the results obtained from Scenario 2, where we distributed the data using its label in a heterogeneous manner. This distribution ensured that each hospital section had varying quantities of labeled data, simulating real-world scenarios where data distribution is not uniform across sections, as described in Section 5.2. We compared the performance of three models in this scenario: the standard ML model, the standard FL model, and our proposed Federated Hospital (α is set to 0.7). Table 3 displays the accuracy achieved by each model based on the data label.

Table 3: Accuracy achieved by each model based on heterogeneous distributed data label. For the Federated Hospital α set to 0.7.

	Label Accuracy (%)									
	0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
ML	98.06	97.53	**90.50**	91.29	94.30	87.44	94.99	92.12	90.35	89.99
FL	97.76	97.36	86.92	89.41	92.26	82.62	93.32	89.88	86.04	87.91
FH Sec. 0	**98.67**	97.18	87.89	90.10	90.22	85.09	94.05	89.69	87.37	**91.87**
FH Sec. 1	97.55	**97.89**	87.60	**93.56**	92.97	82.96	94.47	90.86	86.45	88.80
FH Sec. 2	97.65	97.53	88.37	90.10	**95.21**	84.30	**95.62**	91.25	86.86	86.22
FH Sec. 3	97.65	97.44	87.60	89.31	93.38	**89.35**	93.32	**93.58**	86.86	86.52
FH Sec. 4	97.86	97.00	90.41	89.50	92.16	83.52	94.05	90.76	**91.79**	88.21

Source: Made by the author

It is important to disclose that the standard ML model was trained on the entire dataset, so it is performance will remain consistent when compared to the previous data distribution. That said, the standard ML exhibited a consistent accuracy across labels, keeping the same 92.75% overall accuracy. However, this time our proposal, outperformed both the standard ML and standard FL models significantly when evaluating the local levels. By capitalizing on the multilevel aggregation approach, the architecture effectively addressed the heterogeneity challenge. The specialized models in each hospital section demonstrated high accuracy for their respective labels, while the global aggregation process ensured a cohesive and comprehensive model representation. The final accuracy for the Federated Hospital was 91.17%.

In contrast the standard FL model, did not outperformed any of the other models this time. It is accuracy varied across labels based on data availability, but, the Federated Hospital approach consistently achieved superior accuracy for it is specialized models. Also, the architecture have the ability to combine localized expertise with global knowledge sharing proved invaluable in recognizing patterns across the hospital while preserving the uniqueness of each section. The final accuracy for the standard FL was 90.34%

6.4 Discussion

The comparison of resource usage and convergence time revealed our proposed model to be significantly advantageous. In a controlled environment, using the same hardware, it achieved 200 epochs within an impressive 155 seconds. This marginal improvement over the standard FL approach can be attributed to the distinct aggregation methodologies employed.

While the standard FL approach aggregates all models and the global model at each round, our proposal performs aggregation only twice: once at the end of the 200th epoch to combine local models with the global model, and another to propagate the updated global model to each local model. This streamlined aggregation process contributes to the efficiency of our Federated Hospital architecture, resulting in faster convergence and more efficient resource utilization.

Regarding inference performance, Scenario 1 demonstrated our proposal's capability to handle homogeneous data distribution effectively. Although it did not surpass the accuracy of the other approaches, it provided a comparable performance, particularly when compared to the standard FL model. Notably, the standard FL model even outperformed the standard ML model for some labels under homogeneous data distribution, showcasing its ability to leverage localized data for enhanced accuracy.

However, the true potential of our proposal emerged in Scenario 2, where we introduced heterogeneous data distribution. The results from this scenario demonstrated the substantial benefits of the proposed Federated Hospital architecture in handling varying data distributions within smart hospitals. The architecture's adaptability to localized data patterns and robustness in capturing global trends underscore its potential for revolutionizing intelligent healthcare systems.

With this, we list our main achievements as follow:

- Efficient Resource Utilization: Our proposed Federated Hospital architecture demonstrated significant advantages in resource usage and convergence time. With a 155-second completion time for 200 epochs, it outperformed the standard FL approach due to streamlined aggregation methodologies.

- Handling Heterogeneous Data Distribution: The true potential of our proposal emerged in Scenario 2, where we introduced heterogeneous data distribution. The results highlighted the substantial benefits of the Federated Hospital architecture in adapting to varying data distributions within smart hospitals.

We also point the main limitation of our proposal:

- Privacy and Security Concerns: While FL offers data privacy advantages, ensuring robust security measures is crucial to safeguard patient information against potential threats.

The Federated Hospital architecture empowers smart hospitals to train specialized machine learning models for each department, leading to tailored and more effective patient care. The proposed architecture optimizes model accuracy by leveraging localized data patterns and global trends, contributing to better healthcare decision-making and treatment outcomes. In summary, this proposal could be beneficial for patient care and the improvement of personalized models. With the model being used in a hospital, we could give better classification results for possible diseases, control stress levels from hospital staff more precisely, and overall, increase life quality from its users.

7 CONCLUSION

In this proposal, we introduced the Federated Hospital: A Multilevel Federated Learning Architecture for Smart Hospitals, aimed at leveraging the benefits of federated learning in the context of intelligent healthcare systems through specialized models. Our architecture addressed the challenges of model adaptation based on data heterogeneity. By dividing the MNIST dataset into hospital sections and training specialized models at each level, we achieved promising results that showcased the advantages of our approach. Through extensive evaluation in two distinct scenarios using three models, including a standard machine learning approach, a standard federated learning approach using FedAvg, and our novel Federated Hospital approach, we were able to compare their performance and highlight the contributions of our proposed architecture.

The results demonstrated that the Federated Hospital approach outperformed the standard machine learning approach in terms of accuracy, especially when considering the hospital section specific characteristics of data. Compared to the standard federated learning approach, our architecture exhibited improved performance for local users in terms of patient-level accuracy and convergence speed, while maintaining a strong global representation.

In Section 1.4 we presented our hypotheses for this proposal. With the development of this work, we can take the following conclusion: Hypotheses 1 - Confirmed. Our experimental results confirmed Hypothesis 1, demonstrating that the Federated Hospitals architecture outperformed the standard federated learning approach in terms of accuracy when dealing with unbalanced datasets; Hypothesis 2 - Confirmed. We can also confirm the second hypothesis, the multilevel aggregation approach was able to create specialized service models for each hospital ward, having performed better even when compared to standard ML method.

The evaluation metrics revealed significant improvements in classification accuracy, convergence speed, and resource utilization when compared to both the standard machine learning approach and the standard federated learning approach. The results demonstrated that the Federated Hospital architecture effectively captured the local characteristics and behaviors within each hospital ward, while maintaining a strong global representation.

Our contribution lies in bridging the gap between localized expertise and global knowledge sharing. The Federated Hospital architecture effectively combines the strengths of both approaches, resulting in accurate models that specialize in local data while maintaining a cohesive global representation.

Overall, our proposed Federated Hospital architecture offers a promising solution for addressing the challenges of heterogeneous data distribution in smart hospitals. By combining the power of federated learning with the localized expertise within hospital sections, we demonstrated improved performance, patient-level accuracy, and convergence speed. The architecture's flexibility, efficiency, and superior inference performance highlight its potential to revolutionize data-driven decision-making in healthcare, promoting improved patient care and tailored services across diverse hospital environments.

7.1 Contributions

The contributions of our proposal are twofold. First, we introduced a multilevel architecture that enables specialized training at local levels while ensuring effective aggregation at a global level. This architecture provides a practical solution for utilizing localized expertise within smart hospitals, enhancing the accuracy and efficiency of healthcare services.

Secondly, our evaluation highlighted the effectiveness and advantages of the Federated Hospital approach. By leveraging the unique characteristics of each hospital section, our approach achieved improved performance for local users, ensuring personalized and accurate predictions while maintaining the global knowledge representation.

7.2 Publications

During the development of this mater thesis, we had a set of publications. First we detail a list of articles that were important to develop this work, but does not comprehend the core of the current proposal.

- PriBB: A Benchmark Proposal to Analyze Blockchain Applications Performance - Conference: 2020 XLVI Latin American Computing Conference (CLEI);
- Blockchain in the reverse agrochemical supply chain: a systematic mapping study - Journal: International Journal of Business Information Systems, 2021;
- Fuzzy time series for predicting phenological stages of apple trees - Conference: SAC '21: The 36th ACM/SIGAPP Symposium on Applied Computing;
- Machine learning through the lens of e-commerce initiatives: An up-to-date systematic literature review - Journal: Computer Science Review, 2021;
- Otimizando o diagnóstico automatizado de glaucoma a partir de imagens de fundo de olho - Conference: XXII Escola Regional de Alto Desempenho da Região Sul (ERAD-RS), 2022;
- Development and testing of methods for detecting off-wrist in actimetry recordings - Journal: Sleep, Volume 45 Issue 8, 2022;
- Fraud detection and prevention in e-commerce: A systematic literature review - Electronic Commerce Research and Applications, 2022;
- Aiding Glaucoma Diagnosis from the Automated Classification and Segmentation of Fundus Images - Book: Intelligent Systems, 2022;

- A Blockchain-Based End-to-End Data Protection Model for Personal Health Records Sharing: A Fully Homomorphic Encryption Approach - Journal: Sensors, 2023.

In addition, to the afford mentioned set of publication we also have a group of closed related research papers.

- Unindo Aplicações Críticas e Sensores IoT com QoS Individual e Adaptativo em Hospitais Inteligentes - Conference: Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS) 2021;
- Looking at Smart Cities Through the Lens of a Pandemic Era: A Systematic Literature Review - Journal: International Journal of Technology Management, 2022;
- Uma arquitetura escalável e segura para a execução de aprendizado federado no contexto de hospitais inteligentes - Conference: XXII Escola Regional de Alto Desempenho da Região Sul (ERAD-RS), 2022;
- Tracking machine learning models for pandemic scenarios: a systematic review of machine learning models that predict local and global evolution of pandemics - Network Modeling Analysis in Health Informatics and Bioinformatics, 2022.

7.3 Future works

While our research on the Multilevel Federated Learning Architecture for Smart Hospitals has yielded promising results, there are several limitations that need to be acknowledged. These limitations open up avenues for future work and research to further enhance the proposed architecture:

1. Privacy and Security: Although federated learning preserves data privacy by keeping the data within each hospital section, potential privacy and security concerns should be thoroughly addressed. Developing robust mechanisms for secure model aggregation, preventing model poisoning attacks, and ensuring compliance with privacy regulations are important aspects that require further investigation.
2. Generalizability: While our research demonstrates the efficacy of the proposed architecture within the specific context of smart hospitals, further investigation is needed to assess its generalizability to other domains and application scenarios. Exploring its applicability in diverse healthcare settings, such as outpatient clinics or remote healthcare systems, could provide valuable insights into its versatility and adaptability.

By addressing these limitations and advancing research in these areas, we can unlock the full potential of the Multilevel Federated Learning Architecture for Smart Hospitals and pave the way for intelligent, privacy-preserving healthcare systems that optimize patient care and enable data-driven decision-making.

REFERENCES

BHARDWAJ, R.; NAMBIAR, A. R.; DUTTA, D. A study of machine learning in healthcare. In: IEEE 41ST ANNUAL COMPUTER SOFTWARE AND APPLICATIONS CONFERENCE (COMPSAC), 2017., 2017. **Anais. . .** [S.l.: s.n.], 2017. v. 2, p. 236–241.

BURKOV, A. **The hundred-page machine learning book**. [S.l.: s.n.], 2019. 978–1999579500 p.

BUYYA, R.; SRIRAMA, S. N. **Fog and edge computing**: principles and paradigms. [S.l.]: John Wiley & Sons, 2019.

CALLAHAN, A.; SHAH, N. H. Chapter 19 - machine learning in healthcare. In: SHEIKH, A. et al. (Ed.). **Key advances in clinical informatics**. [S.l.]: Academic Press, 2017. p. 279–291.

CALLAHAN, A.; SHAH, N. H. Machine learning in healthcare. In: **Key advances in clinical informatics**. [S.l.]: Elsevier, 2017. p. 279–291.

DASARADHARAMI REDDY, K.; GADEKALLU, T. R. et al. A comprehensive survey on federated learning techniques for healthcare informatics. **Computational Intelligence and Neuroscience**, [S.l.], v. 2023, 2023.

DENG, L. The mnist database of handwritten digit images for machine learning research [best of the web]. **IEEE signal processing magazine**, [S.l.], v. 29, n. 6, p. 141–142, 2012.

GUO, Y. et al. Feel: a federated edge learning system for efficient and privacy-preserving mobile healthcare. In: INTERNATIONAL CONFERENCE ON PARALLEL PROCESSING - ICPP, 49., 2020, New York, NY, USA. **Anais. . .** Association for Computing Machinery, 2020. (ICPP '20).

HAKAK, S. et al. A framework for edge-assisted healthcare data analytics using federated learning. In: IEEE INTERNATIONAL CONFERENCE ON BIG DATA (BIG DATA), 2020., 2020. **Anais. . .** [S.l.: s.n.], 2020. p. 3423–3427.

HERRERO, R. **Fundamentals of iot communication technologies**. [S.l.]: Springer, 2022.

HOVENGA, E. **Health informatics**: an overview. [S.l.]: IOS Press, 2010. (Studies in health technology and informatics).

MCMAHAN, H. B. et al. Communication-efficient learning of deep networks from decentralized data. , [S.l.], 2016.

MITCHELL, T. **Machine learning**. [S.l.]: McGraw hill Burr Ridge, 1997.

NGUYEN, D. C. et al. Federated learning for smart healthcare: a survey. **ACM Computing Surveys (CSUR)**, [S.l.], v. 55, n. 3, p. 1–37, 2022.

POłAP, D. et al. Agent architecture of an intelligent medical system based on federated learning and blockchain technology. **Journal of Information Security and Applications**, [S.l.], v. 58, p. 102748, 2021.

RAHMAN, M. A. et al. Secure and provenance enhanced internet of health things framework: a blockchain managed federated learning approach. **IEEE Access**, [S.l.], v. 8, p. 205071–205087, 2020.

RAVÌ, D. et al. Deep learning for health informatics. **IEEE journal of biomedical and health informatics**, [S.l.], v. 21, n. 1, p. 4–21, 2016.

REHMAN MUHAMMAD HABIB UR, G. M. M. **Federated learning systems**. [S.l.]: Springer, 2021.

RUSSELL, S.; NORVIG, P. Ai a modern approach. **Learning**, [S.l.], v. 2, n. 3, p. 4, 2005.

SANYAL, S. et al. A federated filtering framework for internet of medical things. In: ICC 2019 - 2019 IEEE INTERNATIONAL CONFERENCE ON COMMUNICATIONS (ICC), 2019. **Anais. . .** [S.l.: s.n.], 2019. p. 1–6.

WANG, R. et al. Privacy-preserving federated learning for internet of medical things under edge computing. **IEEE Journal of Biomedical and Health Informatics**, [S.l.], 2022.

WU, Q. et al. Fedhome: cloud-edge based personalized federated learning for in-home health monitoring. **IEEE Transactions on Mobile Computing**, [S.l.], v. PP, p. 1–1, 12 2020.

XU, J. et al. Federated learning for healthcare informatics. **Journal of Healthcare Informatics Research**, [S.l.], v. 5, n. 1, p. 1–19, 2021.

XUE, Z. et al. A resource-constrained and privacy-preserving edge-computing-enabled clinical decision system: a federated reinforcement learning approach. **IEEE Internet of Things Journal**, [S.l.], v. 8, n. 11, p. 9122–9138, 2021.

YANG, Q. et al. Federated learning. **Synthesis Lectures on Artificial Intelligence and Machine Learning**, [S.l.], v. 13, n. 3, p. 1–207, 2019.

ZHAO, Y. et al. Privacy-preserving activity and health monitoring on databox. In: THIRD ACM INTERNATIONAL WORKSHOP ON EDGE SYSTEMS, ANALYTICS AND NETWORKING, 2020, New York, NY, USA. **Proceedings. . .** Association for Computing Machinery, 2020. p. 49–54. (EdgeSys ’20).

ZHOU, T. Hierarchical federated learning with gaussian differential privacy. In: INTERNATIONAL CONFERENCE ON ADVANCED INFORMATION SCIENCE AND SYSTEM, 4., 2023, New York, NY, USA. **Proceedings. . .** Association for Computing Machinery, 2023. (AISS ’22).

ZHU, J. et al. Blockchain-empowered federated learning: challenges, solutions, and future directions. **ACM Comput. Surv.**, New York, NY, USA, v. 55, n. 11, feb 2023.

ZHU, L.; HAN, S. Deep leakage from gradients. In: **Federated learning**. [S.l.]: Springer, 2020. p. 17–31.