

FRAMU: Attention-based Machine Unlearning using Federated Reinforcement Learning

Thanveer Shaik*, Xiaohui Tao*, Lin Li, Haoran Xie, Taotao Cai, Xiaofeng Zhu, and Qing Li

Abstract—Machine Unlearning, a pivotal field addressing data privacy in machine learning, necessitates efficient methods for the removal of private or irrelevant data. In this context, significant challenges arise, particularly in maintaining privacy and ensuring model efficiency when managing outdated, private, and irrelevant data. Such data not only compromises model accuracy but also burdens computational efficiency in both learning and unlearning processes. To mitigate these challenges, we introduce a novel framework, Attention-based Machine Unlearning using Federated Reinforcement Learning (FRAMU). This framework incorporates adaptive learning mechanisms, privacy preservation techniques, and optimization strategies, making it a well-rounded solution for handling various data sources, either single-modality or multi-modality, while maintaining accuracy and privacy. FRAMU’s strength lies in its adaptability to fluctuating data landscapes, its ability to unlearn outdated, private, or irrelevant data, and its support for continual model evolution without compromising privacy. Our experiments, conducted on both single-modality and multi-modality datasets, revealed that FRAMU significantly outperformed baseline models. Additional assessments of convergence behaviour and optimization strategies further validate the framework’s utility in federated learning applications. Overall, FRAMU advances Machine Unlearning by offering a robust, privacy-preserving solution that optimizes model performance while also addressing key challenges in dynamic data environments.

Index Terms—Machine Unlearning, Privacy, Reinforcement Learning, Federated Learning, Attention Mechanism.

I. INTRODUCTION

The widespread availability of decentralized and heterogeneous data sources has created a demand for Machine Learning models that can effectively leverage this data while preserving privacy and ensuring accuracy [1]. Traditional approaches struggle to handle the continual influx of new data streams, and the accumulation of outdated or irrelevant information hinders their adaptability in dynamic data environments [2], [3]. Moreover, the presence of sensitive or private data introduces concerns regarding data breaches and unauthorized

*Corresponding authors: Thanveer Shaik (email: Thanveer.Shaik@usq.edu.au) and Xiaohui Tao (email: Xiaohui.Tao@usq.edu.au) are with the School of Mathematics, Physics, and Computing, University of Southern Queensland, Queensland, Australia

Taotao Cai is with the School of Mathematics, Physics, and Computing, University of Southern Queensland, Queensland, Australia (e-mail: Taotao.Cai@usq.edu.au).

Lin Li is with the School of Computer and Artificial Intelligence, Wuhan University of Technology, China (e-mail: cathylilin@whut.edu.cn)

Haoran Xie is with the Department of Computing and Decision Sciences, Lingnan University, Tuen Mun, Hong Kong (e-mail: hrxie@ln.edu.hk)

Xiaofeng Zhu is with the University of Electronic Science and Technology of China (e-mail: seanzhuxf@gmail.com)

Qing Li is with the Department of Computing, Hong Kong Polytechnic University, Hong Kong Special Administrative Region of China (e-mail: qing-prof.li@polyu.edu.hk).

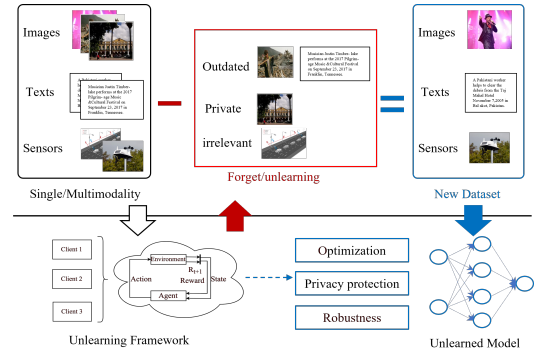


Fig. 1: Graphical abstract depicts the evolution of the FRAMU framework

access, necessitating the development of privacy-preserving techniques [4]. The concept of the “right to be forgotten” allows individuals to have their personal information removed from online platforms, although there’s no universal agreement on its definition or its status as a human right [5]. Despite this, countries like Argentina, the Philippines, and large parts of the EU are working on regulations¹. Therefore, there is a pressing need to advance the field of Machine Unlearning to ensure both adaptability and privacy in Machine Learning applications.

Example 1. *In a landmark 2014 decision that underscored the pressing need for Machine Unlearning, a Spanish court ruled in favor of an individual who sought the removal of specific, outdated Google search results related to a long-settled debt [6]. This verdict not only led to Google taking down the search results but also influenced broader European Union policies on the subject, emphasizing the urgent need for mechanisms that can efficiently erase outdated or private information from Machine Learning models without sacrificing accuracy. This critical requirement for Machine Unlearning is further highlighted by high-profile cases such as that of James Gunn, the famed writer and director, who was dismissed by Disney in 2018 when old, inappropriate tweets resurfaced [7]. Although social media platforms like Facebook offer features like “Off-Facebook Activity” to disconnect user data from third-party services, this does not guarantee the complete erasure of that data from the internet². Together, these instances accentuate the growing imperative for the development of robust Machine Unlearning technologies, especially in an era where data privacy regulations are continuously evolving*

¹<https://link.library.eui.eu/portal/The-Right-To-Be-Forgotten-A-Comparative-Study/tw0VHCyGcDc/>

²<https://www.facebook.com/help/2207256696182627>

and the "right to be forgotten" is increasingly recognized as essential.

Challenges. In today's digitally connected environment, data is distributed in various forms and from different sources, such as sensors, text documents, images, and time series data. For unlearning outdated or private data, Machine Unlearning presents unique challenges depending on whether it is a single type of data (known as single-modality) or multiple types of data (referred to as multimodality) [8]. With single-modality data, the issue primarily lies in the build-up of outdated or irrelevant information, which can negatively affect the model's effectiveness and precision [9], [10]. On the other hand, multimodality situations are even more complicated. Here, each type of data can have different characteristics and varying contributions to the overall model's performance [11], [12]. As we discussed in example 1, the need to unlearn outdated or private data is most important. This ensures individuals have the "right to be forgotten" about their information in publicly available platforms. However, the unlearning needs to happen in both single-modality and multimodality data to make it a holistic unlearning.

Distributed learning systems, particularly federated learning, have made significant strides forward in enabling Machine Learning models to train on decentralized data, offering the dual advantage of reduced communication costs and enhanced privacy [13], [14]. Notable efforts have been made to incorporate Differential Privacy (DP) into these systems [15], ensuring robust privacy safeguards through techniques like DP-SGD and DP-FedAvg [16], [17]. However, these existing frameworks face limitations when confronted with the dynamic nature of data distribution, an intrinsic challenge in distributed learning [18]. Although some efforts have been made in Machine Unlearning to address data irrelevancy over time, such as Sharded, Isolated, Sliced, and Aggregated(SISA) training methods, these solutions often operate in isolation from privacy-preserving mechanisms [19], [20]. This bifurcation leaves a crucial research gap: the absence of a unified approach that addresses both privacy concerns and the adaptability requirements in the face of ever-changing data landscapes. There is a need to bridge this gap by providing an integrated solution for robust privacy measures and efficient selective unlearning, thereby enabling Machine Learning models to be both secure and adaptable in dynamic, distributed environments.

The primary challenges in Machine Unlearning involve addressing the buildup of outdated or irrelevant information in single-modality data, which affects model precision, and handling the complexity of multimodality data where each type contributes differently to model performance. Additionally, current distributed learning systems, while advancing privacy and reducing communication costs, face limitations in adapting to dynamic data distributions and integrating robust privacy measures with efficient unlearning mechanisms, highlighting a need for a unified approach that ensures both security and adaptability in rapidly evolving data environments.

To address these challenges, we propose an Attention-based Machine Unlearning using Federated Reinforcement Learning (FRAMU) as shown in Fig. 1. By integrating federated learn-

ing, adaptive learning mechanisms, and privacy preservation techniques, FRAMU aims to leverage the diverse and dynamic nature of data in both single-modality and multimodality scenarios, while upholding privacy regulations and optimizing the learning process. An attention mechanism is incorporated into FRAMU to ensure responsible and secure handling of sensitive information across modalities. FRAMU leverages reinforcement learning and adaptive learning mechanisms to enable models to dynamically adapt to changing data distributions and individual participant characteristics in both single-modality and multimodality scenarios. This adaptability facilitates ongoing model evolution and improvement in a privacy-preserving manner, accommodating the dynamic nature of the data present in federated learning scenarios. In addition to addressing the challenges associated with unlearning outdated, private, and irrelevant data in both single-modality and multimodality scenarios, FRAMU offers valuable insights into the convergence behaviour and optimization of the federated learning process. The major contributions of our work are as follows:

- We propose an adaptive unlearning algorithm using an attention mechanism to adapt to changing data distributions and participant characteristics in single-modality and multimodality scenarios.
- We develop a novel design to personalize the unlearning process using the FedAvg mechanism [21] and unlearn the outdated, private, and irrelevant data.
- We propose an efficient unlearning algorithm that demonstrates fast convergence and achieves optimal solutions within a small number of communication rounds.
- We conduct extensive experiments to demonstrate the efficiency and effectiveness of the proposed approach using real-world datasets.

Organization. In Section II, we review related works. Section III outlines the problem addressed in this study. We present the proposed framework FRAMU in Section IV. The applications of FRAMU in single-modality and multimodality are discussed in Section V. In Section VI, we present the experimental setup and the evaluation results of the proposed framework, along with convergence and optimization analysis. Section VII delves into the implications of the proposed framework. Finally, in Section VIII, we conclude the paper.

II. RELATED WORKS

The importance of data privacy in distributed learning systems has garnered significant attention, especially when handling sensitive types of data like medical or behavioral information [22]. Differential Privacy (DP), a mathematically rigorous framework for ensuring individual privacy, has been widely adopted for this purpose [23], [24]. Efforts to integrate DP within distributed learning environments, particularly in federated learning, have been increasing [13], [14]. Abadi et al. [16] developed a seminal approach called Deep Learning with Differential Privacy (DP-SGD), which adapts the Stochastic Gradient Descent (SGD) algorithm to meet DP standards by clipping gradients and injecting noise, thereby offering stringent privacy safeguards during deep neural network (DNN) training. Building on this, McMahan et al. [17] further

tailored DP mechanisms for federated learning through an extension called DP-FedAvg. While these methods effectively address privacy concerns, they often fall short in dealing with dynamic data distributions, a prevalent issue in distributed learning [18]. Specifically, data sets can evolve over time, rendering some information outdated or irrelevant, and the persistence of such data in the learning process can compromise model efficacy. Although Machine Unlearning approaches like Sharded, Isolated, Sliced, and Aggregated (SISA) training [19] have emerged to tackle this issue by enabling efficient selective forgetting of data, these methods are not yet designed to work synergistically with privacy-preserving techniques like DP [20].

Federated learning has substantially revolutionized distributed learning, enabling the training of Machine Learning models on decentralized networks while preserving data privacy and minimizing communication costs [25]. Among the pioneering works in this area is the FedAvg algorithm by McMahan et al. [21], which relies on model parameter averaging across local models and a central server. However, FedAvg is not without its limitations, particularly when handling non-IID data distributions [26]. Solutions like FedProx by Li et al. [27] have sought to address this by introducing a proximal term for improved model convergence. While other researchers like Sahu et al. [28] and Konečný et al. [29] have made strides in adaptive learning rates and communication efficiency, the realm of federated learning still faces significant challenges in dynamic adaptability and efficient Machine Unlearning. While privacy has been partially addressed through Differential Privacy [30] and Secure Multiparty Computation [31], these techniques often compromise on model efficiency. Additionally, the applicability of federated learning in diverse sectors like healthcare and IoT emphasizes the unmet need for a model capable of dynamically adapting to varied data distributions, while preserving privacy and efficiency [32], [33].

Reinforcement Learning has garnered much attention for its ability to train agents to make optimal decisions through trial-and-error interactions with their environments [34], [35]. Several pivotal advancements have shaped the field, including the development of Deep Q-Networks (DQNs) [36]. DQNs combine traditional reinforcement learning techniques with DNNs, significantly enhancing the system’s ability to process high-dimensional inputs such as images. Furthermore, experience replay mechanisms have been integrated into them to improve learning stability by storing and reusing past experiences [37]. Mnih et al. [38] significantly accelerated the reinforcement learning field by implementing DQNs that achieved human-level performance on a variety of complex tasks. However, there are evident gaps in addressing challenges posed by non-stationary or dynamic environment situations where the statistical properties of the environment change over time. Under such conditions, a reinforcement learning agent’s ability to adapt quickly is paramount. Several approaches like meta-learning [39] and attention mechanisms [40], [41] have sought to remedy these issues to some extent. Meta-learning, for example, helps models quickly adapt to new tasks by training them on a diverse range of tasks. However, the technique does not offer a robust solution for unlearning or forgetting outdated

or irrelevant information, which is crucial for maintaining performance in dynamic environments. In a similar vein, attention mechanisms help agents focus on important regions of the input space, but they also fail to address the need for efficient unlearning of obsolete or irrelevant data. This leaves us with a significant research gap: the lack of mechanisms for efficient unlearning and adaptability in reinforcement learning agents designed for non-stationary, dynamic environments.

A key challenge for federated learning when faced with dynamic data distributions and the accumulation of outdated or irrelevant information is its adaptability in evolving environments. Reinforcement learning has been instrumental in training agents for optimal decision-making in dynamic environments, yet it too grapples with the need to efficiently unlearn outdated or irrelevant data. These challenges underscore the importance of integrating attention mechanisms into the Machine Unlearning process. Unlike selective data deletion, attention mechanisms assign reduced weights to outdated, private, or irrelevant information. The dynamic adjustment of attention scores allows these models to prioritize relevant data while disregarding obsolete or extraneous elements. By bridging the worlds of federated learning and reinforcement learning with attention mechanisms, our study addresses the pressing need for an integrated solution that optimizes decision-making in distributed networks with changing data landscapes [42]. In addition, this approach must preserve data privacy and adaptively forget outdated, private, or irrelevant information.

III. PRELIMINARIES & PROBLEM DEFINITION

This section establishes the foundational concepts and mathematical notations essential for the discussions and analyses presented in this paper. These concepts are summarized in Tab. I and form the basis for understanding the subsequent problem definitions and solution approaches. Our research is centered around the exploration of unlearning mechanisms in Machine Learning models, focusing on maintaining accuracy and computational efficiency while addressing the challenges posed by outdated or irrelevant data.

The problem is defined by two distinct settings: single-modality and multimodality. The single-modality setting is simpler and widely applicable in scenarios with uniform data types, such as sensor networks or content recommendation systems. However, it may lack the context provided by different types of data, potentially leading to less nuanced decisions. On the other hand, the multimodality setting is more complex but highly relevant in fields like healthcare, where a range of data types (e.g., medical imaging, patient history, etc.) can be used for more comprehensive understanding and decision-making. By exploring the problem in both these settings, we offer solutions that are both versatile and contextually rich.

A. Problem Definition - Single Modality

Problem Definition 1. Let $AG = \{ag_1, ag_2, \dots, ag_n\}$ be a set of agents, where each agent $ag \in AG$ represents an entity like an IoT device, traffic point, wearable device, edge computing node, or content recommendation system. Each agent ag observes states $S_i = \{s_1, s_2, \dots, s_m\}$ and takes actions $A = \{a_1, a_2, \dots, a_n\}$ based on a policy $\pi_i(s, a)$. Rewards $R_i(s, a)$ evaluate the quality of actions taken in

TABLE I: Summary of Notations and Descriptions

Symbol	Description
AG	Set of agents in the model
ag	An individual agent in the set AG
S_i	States observed by an agent ag
A	Set of possible actions
$\pi_i(s, a)$	Policy followed by the agent
$R_i(s, a)$	Rewards for actions in different states
θ_{ag}	Parameters of local models for agent ag
θ_g	Parameters of global model
$w_{i,ag}$	Attention score for a data point i in agent ag
M	Set of modalities in multimodality setting
X_m	Data vectors for modality m
θ_m	Parameters for modality m
$w_{i,m}$	Attention scores within a modality m
t	Time step
s_t	State at time step t
a_t	Action at time step t
r_t	Reward at time step t
R_t	Cumulative reward
$\pi(a_t s_t)$	Policy function
$Q(s_t, a_t)$	Q -function
γ	Discount factor
α_i	Attention score for feature i
$\Delta\theta_{ag}$	Update sent by agent ag
f	Function for calculating attention scores
$w_{g,ag}$	Global attention score for update from agent ag
AG	Number of local agents
α_{avg}	Average attention score
δ	Predetermined threshold for attention score
$ag \in AG$	A specific agent within the set of all agents AG
m	Number of modalities
x_1, x_2, \dots, x_m	Data vectors for each modality
v_i	Feature vector for modality i
\bar{w}_j	Averaged attention score across modalities for data point j
λ	Mixing factor
T	The total number of training rounds
α	Learning rate for Q -value function updates
η	Scaling factor for attention score updates
β	Mixing factor for combining global and local model parameters
ε	Convergence threshold for global model parameters
w_{ag}	Local model parameters for agent ag
W	Global model parameters
A_i	Attention score for data point i
$A_{i,ag}$	Attention score for data point i within agent ag
N	Total number of data points across all agents
n_{ag}	Number of data points in agent ag




Fig. 2: Single Modality Example

different states. Agents possess local models with parameters θ_i , while a central server maintains a global model with parameters θ_g .

Example 2. In the single-modality setting shown in Fig. 2, let $AG = \{ag_1, ag_2, \dots, ag_n\}$ be a set of agents. An agent ag can represent a real-world entity such as a traffic light in a city. These traffic lights observe various states $S_i = \{s_1, s_2, \dots, s_m\}$, such as the number and speed of passing cars, and the change of colors (actions $A = \{a_1, a_2, \dots, a_n\}$) according to an algorithmic policy $\pi_i(s, a)$. The system evaluates the effectiveness of the traffic light changes in reducing wait time or congestion (rewards $R_i(s, a)$). Each traffic light has its own local decision-making model characterized by parameters θ_i , and there is a global model for optimizing city-wide traffic flow with parameters θ_g .




Fig. 3: Multimodality Example

To address the challenge of preserving data privacy and adaptively forgetting private, outdated, or irrelevant information, attention scores w_{ij} are assigned to each data point j in the local dataset of agent $ag \in AG$. These attention scores, computed using a function f that considers the current model state or contextual information, guide the learning and unlearning process within each agent. By assigning higher attention scores to relevant data and potentially forgetting or down-weighting irrelevant data, the agents can effectively focus on the most informative and up-to-date information.

B. Problem Definition - Multimodality

Problem Definition 2. In the multimodality setting, let $M = \{1, 2, \dots, m\}$ represent the set of modalities, where m is the total number of modalities. Each modality $m \in M$ is associated with a set of data vectors $X_m = \{x_{m1}, x_{m2}, \dots, x_{mn}\}$, and has its local model with parameters θ_k . Attention scores w_{im} are assigned to individual data points x_{im} within each modality to guide the learning and unlearning process.

Example 3. In the multimodality setting shown in Fig. 3, consider a healthcare system as a collection of agents in set $M = \{1, 2, \dots, m\}$, where m represents different types of medical data (modalities) such as medical imaging and patient history. For instance, medical imaging (modality M_1) would have a set of MRI scans represented as data vectors $X_1 = \{x_{11}, x_{12}, \dots, x_{1n}\}$. Likewise, patient history (modality M_2) might involve a set of past diagnosis records that are represented as data vectors $X_2 = \{x_{21}, x_{22}, \dots, x_{2n}\}$. Each modality has a specialized model with parameters θ_1 for medical imaging and θ_2 for patient history. These models use attention mechanisms to weigh the importance of each data point, represented by attention scores w_{1m} for MRI scans and w_{2m} for patient history records. These scores guide the decision-making process in diagnosis and treatment.

In the multimodality setting, the complexity is elevated by the integration of heterogeneous data types and the application of specialized machine learning models for each modality. For example, in a healthcare system, combining data from disparate sources like medical imaging and patient history presents a unique challenge. Each data type, or modality, not only varies in format but also in the nature of the information it conveys, necessitating distinct processing and analysis methods. The key challenge here is to develop an integrated approach that effectively synthesizes these diverse data streams into a coherent understanding, enhancing decision-making in critical applications such as patient diagnosis and treatment. Attention mechanisms play a crucial role in this context, determining the relevance of each data point across different

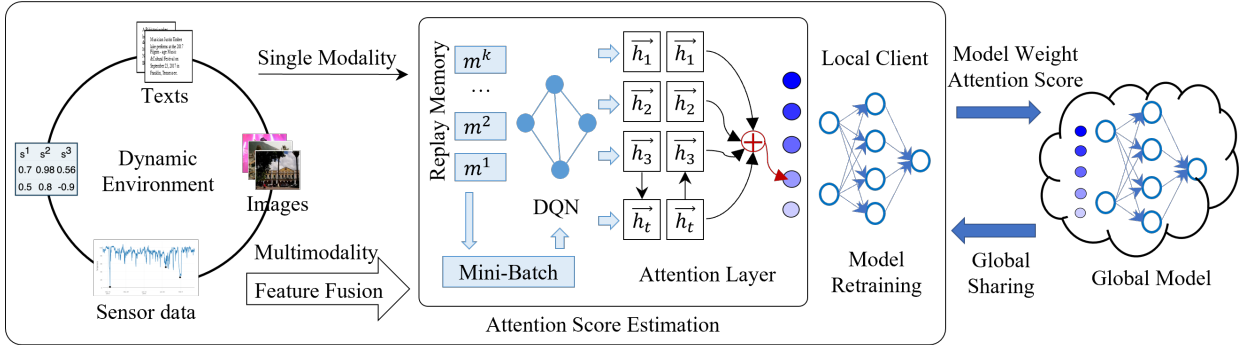


Fig. 4: An overview of the proposed FRAMU framework, illustrating its end-to-end adaptive algorithm that incorporates an attention mechanism. The figure is divided into multiple components, each corresponding to a specific phase in the federated learning process. Starting from the left, the diagram begins with data collection from diverse modalities. The framework applies an adaptive learning algorithm that not only updates the global model, but also incorporates an efficient unlearning mechanism for discarding outdated, private, or irrelevant data.

modalities. However, assigning and calibrating these attention scores is non-trivial and introduces an additional layer of complexity. The successful implementation of multimodal systems has profound implications, particularly in improving the accuracy and efficacy of decision-making processes.

IV. FRAMU FRAMEWORK

In an era marked by an ever-increasing influx of data, the need for adaptive Machine Learning models that can efficiently unlearn outdated, private, or irrelevant information is paramount. The methodology proposed in this paper addresses this necessity by introducing two key technical contributions. First, we propose an adaptive unlearning algorithm that utilizes attention mechanisms to tailor the learning and unlearning processes in a single-modality, and then extend the process to multimodality. This innovative approach allows the model to adapt to dynamic changes in data distributions, as well as variations in participant characteristics such as demographic information, behavioural patterns, and data contribution frequencies among others. Second, we put forth a novel design that employs the FedAvg mechanism [21] to personalize the unlearning process. This design ensures that the model is able to discard data that has become irrelevant, outdated, or potentially invasive from a privacy perspective, thus preserving the integrity of the learning model while adapting to new or changing data. The following sections will elaborate on these contributions, providing a detailed discussion of the proposed framework as depicted in Fig. 4.

The FRAMU framework adopts a federated learning architecture comprising Local Agents and a Central Server, each with distinct roles in model training, unlearning, and adaptation. It employs a reinforcement learning paradigm where each agent iteratively learns from its environment. This integration of federated learning and reinforcement learning is termed federated reinforcement learning. However, what sets FRAMU apart is the integration of attention mechanisms to weigh the relevance of each data point in learning and unlearning. The attention scores are then aggregated and processed at the Central Server to refine the global model.

- **Local Agents:** Responsible for collecting real-time data and performing local model updates. They observe states, take actions, and calculate rewards to update their Q-values and attention scores.
- **Central Server:** Aggregates local models and attention scores, filters out irrelevant data points, and updates the global model.
- **Attention Mechanism:** Dynamically calculates attention scores for each data point to inform the unlearning process.
- **FedAvg Mechanism:** Utilized for global model updates, ensuring that the global model represents a consensus across all agents.

The FRAMU framework, as outlined in Algorithm 1, has been carefully designed to facilitate adaptive decision-making in distributed networks through federated reinforcement learning. Each step within the algorithm is crafted with specific intentions: The initialization stage (Lines 1-3) sets the groundwork by initializing local and global model parameters, as well as attention scores. These initializations are crucial for ensuring that both local and global perspectives are considered right from the start of the learning process. The iterative learning process (Lines 4-24) involves several key components. Local Agent Decision-Making (Lines 5-11) enables each local agent to observe states, take actions, and update its Q-values and attention scores, ensuring that local knowledge is continuously updated to reflect the dynamic nature of the agents' environments. Central Server Aggregation (Lines 12-17) plays a pivotal role in integrating local updates and refining the global model. By assessing the attention scores, the server can identify and diminish the influence of less relevant data points, thereby enhancing the model's focus on significant information. Model Synchronization (Lines 18-24) involves the dissemination of global model parameters back to local agents for fine-tuning, ensuring a bi-directional flow of information that keeps local models informed by their immediate environment and aligned with the broader objectives of the global model.

Algorithm 1: FRAMU Framework

Input: a set of Local Agents, a Central Server, $T, \theta, \alpha, \eta, \gamma, \beta, \varepsilon$

Output: \bar{W} : Trained global model parameters for federated reinforcement learning

- 1 Initialize local model parameters w_{ag} for each agent ag ;
- 2 Initialize global model parameters W at the central server;
- 3 Initialize attention scores $A_{i,ag,m}$ for each data point i in agent ag and modality m ;
- 4 **while** $t \leq T$ **do**
- 5 **foreach** local agent ag **do**
- 6 Observe current states $s_{i,m}$ for each modality m ;
- 7 Take action a_t based on policy derived from $Q(s, a; w_{ag})$;
- 8 Observe reward r_t and next states $s'_{i,m}$ for each modality m ;
- 9 Compute TD error $\delta = r_t + \gamma \max_a Q(s'_{i,m}, a; w_{ag}) - Q(s_{i,m}, a_t; w_{ag})$;
- 10 Update $Q(s_{i,m}, a_t; w_{ag}) \leftarrow Q(s_{i,m}, a_t; w_{ag}) + \alpha \delta$;
- 11 Update attention scores $A_{i,ag,m} \leftarrow A_{i,ag,m} + \eta |\delta|$;
- 12 Send local model parameters w_{ag} and attention scores $A_{i,ag,m}$ to Central Server;
- 13 **foreach** data point i in modality m **do**
- 14 **if** $\sum_{ag} \frac{1}{m} \sum_m A_{i,ag,m} / N_{ag} < \theta$ **then**
- 15 Reduce influence of data point i in the global model;
- 16 Aggregate local model parameters to update global parameters: $W \leftarrow \sum_{ag} \left(\frac{n_{ag}}{N} \right) w_{ag}$;
- 17 Send updated global model parameters W to local agents;
- 18 **foreach** local agent ag **do**
- 19 Fine-tune local model with global model:
- 20 $w'_{ag} \leftarrow \beta W + (1 - \beta) w_{ag}$;
- 21 **if** $|P(W_{t+1}) - P(W_t)| < \varepsilon$ **then**
- 22 Break;
- 23 Increment t ;
- 24 **return** W

V. APPLICATIONS OF FRAMU

This section explores the practical applications of the FRAMU framework across different settings, single-modality and multimodality, and its continuous adaptation and learning.

A. FRAMU with Single Modality

Central to FRAMU is an attention layer that functions as a specialized approximator, augmenting the learning capability of individual agents. This attention layer distinguishes itself by assigning attention scores to individual data points during the function approximation process. These scores serve as indicators of each data point's relevance in the agent's local learning. The agent updates these scores as it interacts with its environment and receives either rewards or penalties, thereby continually refining its model. Specifically, an agent operates in discrete time steps, current state s_t , taking action a_t , and receiving reward r_t , at each time step t . The ultimate goal is to determine an optimal policy $\pi(a_t|s_t)$ that maximizes the accumulated reward R_t . The Q -function, which quantifies expected accumulated rewards with a discount factor γ , is given by Equation 1.

$$Q(s_t, a_t) = \mathbb{E}[R_t | s_t, a_t] = r_t + \gamma \mathbb{E}[Q(s_{t+1}, a_{t+1}) | s_t, a_t] \quad (1)$$

The attention layer further characterizes each state s_t by its features $[x_1, x_2, \dots, x_n]$, and assigns attention scores α_i as per:

$$\alpha_i = \text{Attention}(x_i, \text{context}) \quad (2)$$

Here, the context may include additional data such as previous states or actions. The Q -function is then approximated using a weighted sum of these features:

$$Q(s_t, a_t) \approx \sum (\alpha_i \cdot x_i) \quad (3)$$

After completing their respective learning cycles, agents forward their model updates θ and attention scores α to the Central Server as a tuple (θ, α) .

1) Local and Global Attention Score Estimation

FRAMU estimates attention scores both locally and globally. On the local front, each agent employs its attention mechanism to compute scores for individual data points based on their relevance to the task at hand. For an agent ag with local model parameters θ_{ag} , the attention score w_{ij} for data point j is given by:

$$w_{ij} = f(s_j, \theta_{ag}) \quad (4)$$

At the global level, these scores assist the Central Server in prioritizing updates or pinpointing data points for global unlearning. For global parameters θ_g , the global attention score derived from the updates of agent ag is:

$$w_{g,ag} = f(\Delta\theta_{ag}, \theta_g) \quad (5)$$

In this equation, $\Delta\theta_{ag}$ is the model update from agent ag , and the function f calculates attention scores while taking into account the aggregated local scores and other global contextual cues.

2) Global Model Refinement and Unlearning

Model updates from local agents are aggregated at the Central Server using FedAvg [43]. The attention scores are instrumental in the global unlearning process, with the average attention score calculated as:

$$\alpha_{\text{avg}} = \frac{1}{AG} \sum \alpha_{ag} \quad (6)$$

When α_{avg} falls below a predetermined threshold δ , the server adjusts the contribution of the respective feature in the global model as given by Equation 7:

$$\theta_{\text{global}'} = g(\theta_{\text{global}}, \alpha_{\text{avg}}) \quad (7)$$

Once refined, this global model is sent back to the local agents. The enhanced model shows improved adaptability and robustness to changes in data distributions due to the integration of aggregation and unlearning mechanisms. Consequently, the local agents are better positioned to excel within their particular operational environments. These revised global model parameters, denoted as $\theta_{\text{global}'}$, are then dispatched from the Central Server to the local agents, where $\theta_k = \theta_{\text{global}'}$.

B. FRAMU with Multimodality

The multimodal FRAMU Framework extends its capabilities to seamlessly incorporate various data types, including images, text, audio, and sensor readings. This integration not only enriches decision-making but also optimizes the performance of local agents. By fine-tuning their models to multiple data types, agents are better equipped to operate in complex environments.

1) Modality-Specific Attention Mechanisms

To effectively manage data from diverse sources, the framework employs specialized attention mechanisms for each modality. These mechanisms generate unique attention scores for data points within a given modality, aiding in both learning and unlearning processes. By doing so, the framework allows local agents to focus on the most relevant and informative aspects of each modality.

The attention scores for a specific modality j for an agent $ag \in AG$ can be mathematically represented as:

$$w_{ij} = f_j(s_{ij}, \theta_i), \quad (8)$$

Here, s_{ij} signifies a data point from modality j related to agent $ag \in AG$, while θ_i represents that agent's local model parameters. The function f_j considers modality-specific attributes and context to compute these attention scores.

For a feature vector v_i derived from modality j within agent $ag \in AG$, feature-level fusion can be represented as:

$$v_i = [x_{i1}, x_{i2}, \dots, x_{im}] \quad (9)$$

2) Unlearning and Adaptation across Modalities

In a multimodal setup, attention scores from all modalities collectively inform the unlearning process. If a data point consistently receives low attention scores across different modalities, it indicates that the point is either irrelevant or outdated. The Central Server uses this multimodal insight to refine the global model.

The average attention score across all modalities for a specific data point is:

$$\bar{w}_j = \frac{1}{m} \sum_{i=1}^m w_{ij} \quad (10)$$

If \bar{w}_j falls below a predefined threshold, the Central Server de-emphasizes or removes that data point from the global model, ensuring that only current and relevant data contribute to decision-making.

During the adaptation phase, local agents utilize the updated global model to enhance their local models. The interplay between global and local parameters is regulated by a mixing factor, which allows local agents to leverage shared insights while preserving modality-specific skills. This relationship can be denoted by:

$$\theta_i^{\text{new}} = \lambda \theta_{\text{global}} + (1 - \lambda) \theta_i^{\text{old}} \quad (11)$$

Here, θ_i^{new} represents the updated local model parameters, θ_{global} signifies the global model parameters, θ_i^{old} is the previous local parameters, and λ serves as the mixing factor.

Through this, the multimodal FRAMU framework maintains an up-to-date and relevant global model, while enabling local agents to make better decisions across a range of data types.

C. Continuous Adaptation and Learning in the FRAMU Framework

Continuous adaptation and learning are critical in the FRAMU framework, enabling it to thrive in dynamic and changing environments. These processes create an iterative exchange of knowledge between local agents and a Central Server, which leads to consistent model refinement on both local and global scales.

1) Local-Level Adaptation

Local agents need the ability to adapt in real time to changes in their operational environments. Within reinforcement learning paradigms, agents continually update their policies in response to actions taken and rewards observed. Furthermore, attention scores allocated to data points or features can vary dynamically based on new data or shifts in relevance. This adaptability ensures that the models of individual local agents remain current. Let s_t denote the state of the environment at time t , and a_t represent the action taken by the agent. After receiving a reward r_t and transitioning to a new state s_{t+1} , the agent aims to maximize the expected cumulative reward. The Q-value function $Q(s, a)$ serves as a proxy for this cumulative reward, and it is updated using temporal-difference learning algorithms as follows:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left[r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \right] \quad (12)$$

Here, α is the learning rate, and γ is the discount factor.

Attention scores, denoted by A_i for data point i , are updated based on the temporal-difference error δ :

$$A_i \leftarrow A_i + \eta |\delta|, \quad (13)$$

where η is a scaling factor, and $\delta = r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)$.

2) Global Model Aggregation and Adaptation

As local agents continuously update their models, these adaptations are communicated to the Central Server. It aggregates this information to refine the global model while also tracking the attention scores from local agents. If these scores reveal diminishing importance for certain data points, the server may initiate global unlearning. This ensures the global model remains current and avoids obsolescence. Local agents send their updated model parameters, w_{ag} for agent ag , and attention scores $A_{i,ag}$ to the Central Server. The server aggregates these to update the global model parameters W as follows:

$$W \leftarrow \frac{1}{AG} \sum_{ag} w_{ag}, \quad (14)$$

where AG represents the total number of local agents.

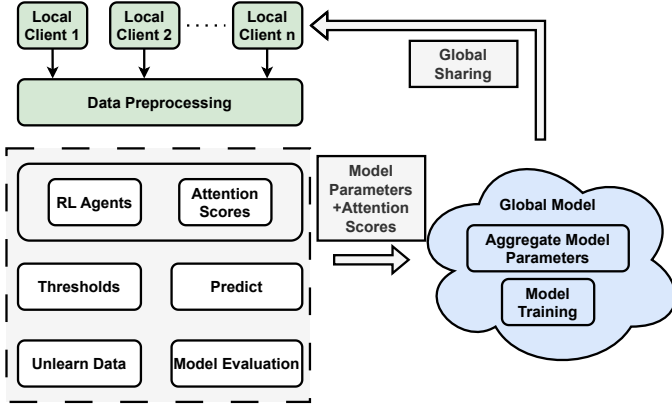


Fig. 5: Experimental Setup: This diagram showcases the architecture of the FRAMU framework, detailing the interaction between local and global models within a federated learning environment.

3) Feedback Mechanisms

After the global model is updated, it is disseminated back to local agents through a feedback loop. This cyclic interaction allows local agents to either initialize or further refine their models based on the global one. This is particularly beneficial when local agents confront new or unfamiliar data points that other agents have encountered. Through this mechanism, the global model acts as a repository of shared knowledge, enhancing the decision-making capabilities of all local agents. The global model parameters W are sent to local agents, who then adjust their local models using a mixing factor β as follows:

$$w'_{ag} \leftarrow \beta W + (1 - \beta)w_{ag}, \quad (15)$$

where β ranges from 0 to 1 and regulates the influence of the global model on local models.

VI. EXPERIMENTAL SETUP AND RESULTS ANALYSIS

To effectively evaluate the performance of the FRAMU framework, we undertook comprehensive experiments using real-world datasets. These experiments were designed to validate not only the efficiency and effectiveness of our approach but also to establish the practical utility of FRAMU in real-world applications. Our experimental setup encompassed several components, including datasets, baseline models, evaluation metrics, and specific FRAMU configurations, as depicted in Fig. 5. A critical aspect of our experimentation involved fine-tuning key thresholds to guide the unlearning process, particularly the `outdated_threshold` and `irrelevant_threshold`. These parameters were adjusted based on domain expertise and sensitivity analysis, with the `outdated_threshold` defining the time frame for data obsolescence and the `irrelevant_threshold` setting criteria for data’s statistical insignificance. Additionally, we introduced a `privacy_epsilon` parameter to balance data utility with privacy preservation, aligning with GDPR regulations.

Deep learning methods are known for their ability to learn features autonomously and automate model-building processes. Despite criticisms of neural network family algorithms for their ‘black box’ nature, deep learning models are

renowned for their robust and efficient performance. These models are widely adopted by the research community [44], [45]. In our work, we utilized a Convolutional Neural Network (CNN) for image and sensor data, and a Long Short-Term Memory (LSTM) network for time series and text data, specifically tailored for federated learning scenarios. This model choice was made to efficiently handle both single and multimodal data types, integrating attention mechanisms and unlearning processes to enhance overall functionality.

We found the tuning of parameters such as the outdated threshold, irrelevant threshold, and the β value for local model fine-tuning to be crucial. It was essential to strike the right balance in the outdated threshold to prevent premature data discarding or retention of outdated information, which could affect model accuracy and relevancy. Similarly, careful calibration of the irrelevant threshold was necessary to maintain a balance between data comprehensiveness and quality, ensuring useful data was not excluded nor excessive noise retained. The β value, crucial in determining the extent of global model influence on local models, required fine-tuning to ensure an optimal balance between local and global learning. This was key for local models to benefit from global insights while preserving their unique learning characteristics. The interplay of these hyperparameters significantly influenced FRAMU’s performance, particularly in its ability to adapt to new data and retain relevant historical information. Through sensitivity analyses, we determined their optimal ranges, aiming to maximize FRAMU’s efficiency and adaptability in various real-world scenarios.

TABLE II: Datasets for evaluation

Modality	Dataset	OD*	PD*	ID*	Description
Single Modality	AMPds2 [46]	✓	✓	✓	Electricity, water, and natural gas consumption data from a Canadian household.
	METR-LA [47]	✓	✗	✓	Traffic speed data from over 200 sensors in Los Angeles Metropolitan area.
	MIMIC-III [48]	✓	✓	✓	Health-related data from critical care units, including demographics, vital signs, laboratory results, and medications.
Multi Modality	NYPD [49]	✓	✓	✓	Records of complaints filed with the New York City Police Department.
	MIMIC-CXR [50]	✓	✓	✓	Chest radiographs with associated radiology reports for medical image analysis tasks.
	Smart Home EnergyDataset (SHED) [51]	✓	✓	✓	Energy consumption data from smart home devices and appliances.

*OD - Outdated Data, PD - Privacy Data, ID - Irrelevant Data.

A. Datasets

In this study, publicly available datasets that encompass various modalities and address specific challenges related to outdated, private, and irrelevant data are adopted. Tab. II provides detailed information about each dataset, including the data modality, number of instances, attributes, target variables, and specific characteristics pertinent to our study. In order to evaluate FRAMU, we conducted a comprehensive comparison of its performance against several contemporary baseline models.

B. Baseline Models

In the evaluation of the FRAMU framework’s performance and robustness, we have carefully selected several baseline models for comparison. The models in baseline models were adopted from the original work. The rationale behind choosing each model and its relevance to our study is elaborated below:

• Single-modality

- **FedLU [52]**: FedLU represents a significant advance in federated learning, integrating knowledge graph embedding with mutual knowledge distillation. Its selection as a baseline is due to its innovative approach to collaborative learning, which is closely aligned with FRAMU’s objectives in single-modality settings. FedLU’s methodology provides a comparative framework for assessing FRAMU’s efficiency in knowledge synthesis and distribution.
- **Zero-shot MU [53]**: Zero-shot MU specializes in Machine Unlearning, employing error-minimizing-maximizing noise and gated knowledge transfer. This model was chosen for its novel approach to unlearning, providing a benchmark to evaluate FRAMU’s capability in effectively removing learned information without extensive retraining, a crucial aspect in dynamic environments.
- **SISA Training [19]**: The SISA Training framework is a strategic model that limits data points for optimized unlearning. Its inclusion as a baseline allows us to compare FRAMU’s efficiency in data management and unlearning processes, especially in scenarios where data minimization is key to performance and privacy.

• Multimodality

- **MMoE [54]**: The MMoE model, optimized for handling multimodal data via ensemble learning, serves as a benchmark for evaluating FRAMU’s performance in multimodality settings. Its approach, employing expert networks for different data modalities, provides a comparative perspective for FRAMU’s adaptability and efficiency in handling diverse data types.
- **CleanCLIP [55]**: CleanCLIP, a fine-tuning framework that mitigates spurious associations from backdoor attacks, is pivotal for comparing FRAMU’s robustness against data security threats. Its focus on weakening spurious correlations offers insights into FRAMU’s capabilities in maintaining data integrity and security.
- **Privacy-Enhanced Emotion Recognition (PEER) [56]**: The PEER model, utilizing adversarial learning for privacy-preserving emotion recognition, aligns well with FRAMU’s privacy objectives. Its comparison with FRAMU highlights the effectiveness of FRAMU in safeguarding privacy while performing complex analytical tasks.

C. Evaluation Metrics

The FRAMU framework is evaluated using several important metrics: Mean Squared Error (MSE) [57], Mean Absolute Error (MAE) [58], Reconstruction Error (RE) [59], and Activation Distance (AD) [60]. A lower MSE or MAE score




Fig. 6: Comparative Analysis of MSE and MAE Differences between Original and Unlearned Single Modality Data

shows that the unlearning process is closely aligned with what was expected, indicating a high quality of unlearning. The RE measures how well the model can rebuild data that it has unlearned, with a lower score being better. AD measures the average distance between the predictions of the model before and after unlearning, using what is known as L2-distance, on a specific set of forgotten data. These metrics together give a well-rounded evaluation of how well the unlearning process is working.

All the experiments were run using Python programming language (version 3.7.6) and related TensorFlow, Keras, Open Gym AI, and stable_baselines3 packages.

D. FRAMU Unlearning Results in Single Modality Context

To assess the effectiveness of FRAMU in unlearning outdated, private, and irrelevant data, we analyzed the results from various experiments. FRAMU’s performance was benchmarked against that of established baseline models: FedLU, Zero-shot MU, and SISA Training. It’s important to note that the METR-LA dataset [47] was excluded from the private data unlearning evaluation due to its lack of privacy-sensitive data. For a thorough comparison, we present the performance metrics of FRAMU in unlearning outdated, private, and irrelevant data alongside the results from baseline models in Tab. III. The p-values in these comparisons are indicative of the statistical significance of FRAMU’s performance improvements.

1) Outdated Data

The unlearning of outdated data is vital for maintaining model accuracy and relevance. Outdated data might introduce noise, biases, or outdated patterns. By selectively unlearning such data, FRAMU aims to align the model with the latest data distribution. FRAMU consistently achieved lower MSE and MAE than the baseline models in unlearning outdated data across various datasets. This improvement, evident from the low p-values in Tab. III, demonstrates FRAMU’s statistically significant superiority in adapting models to current data distributions.

2) Private Data

The retention of private data in models can pose significant privacy and legal risks. To mitigate this, FRAMU incorporates techniques for unlearning private data while preserving privacy. Excluding the METR-LA dataset from this analysis, FRAMU consistently outperformed the baseline models in

TABLE III: FRAMU - Evaluation Results in Single Modality Context

Unlearning	Dataset		FedLU [52]			Zero-shot [53]			SISA [19]			FRAMU (Ours)	
			MSE	MAE	p-value	MSE	MAE	p-value	MSE	MAE	p-value	MSE	MAE
Outdated Data	Original	AMPds2	0.063	6.740	0.024	0.061	6.890	0.031	0.059	6.760	0.041	0.046	5.570
		METR-LA	0.079	7.140	0.016	0.082	7.210	0.038	0.078	7.090	0.029	0.065	5.930
		MIMIC-III	0.099	12.800	0.031	0.102	12.930	0.045	0.097	12.680	0.032	0.083	10.650
	Unlearned	AMPds2	0.060	6.630	0.015	0.055	6.860	0.029	0.056	6.690	0.036	0.038	4.670
		METR-LA	0.075	7.020	0.029	0.077	7.100	0.025	0.072	6.960	0.032	0.052	4.910
		MIMIC-III	0.095	12.650	0.023	0.098	12.820	0.041	0.094	12.580	0.017	0.069	8.900
Private Data	Original	AMPds2	0.052	6.780	0.014	0.054	6.930	0.037	0.053	6.810	0.041	0.041	5.540
		MIMIC-III	0.078	12.870	0.035	0.080	13.010	0.043	0.079	12.760	0.045	0.064	10.600
		MIMIC-III	0.078	12.870	0.035	0.080	13.010	0.043	0.079	12.760	0.045	0.064	10.600
	Unlearned	AMPds2	0.049	6.670	0.011	0.052	6.910	0.035	0.051	6.740	0.015	0.033	4.590
		MIMIC-III	0.075	12.720	0.031	0.077	12.900	0.038	0.076	12.650	0.016	0.053	8.860
		MIMIC-III	0.075	12.720	0.031	0.077	12.900	0.038	0.076	12.650	0.016	0.053	8.860
Irrelevant Data	Original	AMPds2	0.047	6.700	0.035	0.050	6.850	0.044	0.048	6.730	0.031	0.037	5.440
		METR-LA	0.054	7.100	0.027	0.056	7.170	0.041	0.055	7.050	0.025	0.043	5.830
		MIMIC-III	0.070	12.730	0.038	0.072	12.870	0.031	0.071	12.620	0.039	0.057	10.410
	Unlearned	AMPds2	0.045	6.590	0.011	0.047	6.830	0.036	0.046	6.660	0.029	0.030	4.510
		METR-LA	0.052	6.980	0.014	0.054	7.070	0.019	0.053	6.930	0.022	0.035	4.750
		MIMIC-III	0.068	12.580	0.029	0.070	12.760	0.024	0.069	12.510	0.027	0.047	8.690

TABLE IV: Comparative analysis of FRAMU’s performance in single modality against baseline models in RE and AD metrics.

Unlearning	Dataset	FedLU [52]		Zero-shot MU [53]		SISA training [19]		FRAMU (Ours)	
		RE	AD	RE	AD	RE	AD	RE	AD
Outdated Data	AMPds2	0.03	0.66	0.029	0.68	0.028	0.67	0.024	0.57
	METR-LA	0.038	0.7	0.039	0.71	0.037	0.69	0.033	0.59
	MIMIC-III	0.048	1.26	0.049	1.28	0.047	1.25	0.043	1.15
Private Data	AMPds2	0.031	0.67	0.032	0.69	0.03	0.67	0.026	0.57
	MIMIC-III	0.049	1.27	0.051	1.29	0.048	1.27	0.044	1.17
Irrelevant Data	AMPds2	0.028	0.66	0.029	0.68	0.027	0.66	0.023	0.56
	METR-LA	0.034	0.7	0.035	0.71	0.033	0.69	0.029	0.59
	MIMIC-III	0.05	1.26	0.052	1.28	0.049	1.25	0.045	1.15

both MSE and MAE metrics in scenarios involving private data. For example, in the AMPds2 dataset, FRAMU’s superior performance in MSE (0.038) and MAE (4.670) is a testament to its effective federated reinforcement learning approach that respects privacy concerns. The significance of these performance gains is reinforced by the associated p-values.

3) Irrelevant Data

Unlearning irrelevant data helps reduce noise and interference from non-contributory data points, enhancing model accuracy and prediction. FRAMU showed exceptional performance in unlearning irrelevant data, recording the lowest MSE and MAE values across all datasets in comparison to the baseline models. For instance, in the AMPds2 dataset, FRAMU’s MSE of 0.033 and MAE of 5.600 surpassed other models. The low p-values validate FRAMU’s significant advantage in discarding irrelevant data.

Fig. 6 visually compares the differences in MSE and MAE between original and unlearned data across various datasets and models. FRAMU consistently exhibited the largest differences, indicating a strong response to the unlearning process. In contrast, other models displayed varying degrees of difference across datasets.

Moreover, in the comparison of RE and AD metrics as illustrated in Tab. IV, FRAMU consistently outperformed its counterparts. Specifically, in the AMPds2 dataset, FRAMU’s RE and AD values (0.024 and 0.57, respectively) were superior to those of FedLU (0.03 and 0.66). Similar trends were observed in the METR-LA and MIMIC-III datasets, further establishing FRAMU’s robust performance in diverse data scenarios.




Fig. 7: Comparative Analysis of MSE and MAE Differences between Original and Unlearned multimodality Data

E. FRAMU Unlearning Results in Multimodality Context

In the multimodality experiment, the FRAMU framework demonstrated its capability to handle diverse data types, including images, text, and sensor data. The aim was to assess FRAMU’s effectiveness in unlearning outdated, private, and irrelevant data in a multimodal context. For this, we utilized benchmark datasets like MIMIC-CXR [50], NYPD Complaint Data [49], and SHED [51]. The key focus was on evaluating error reduction and performance improvements in comparison to baseline models, with p-values highlighting the statistical significance of FRAMU’s advancements.

1) Outdated Data

FRAMU consistently outperformed baseline models across all datasets in handling outdated data. In the NYPD Complaint Data [49], for instance, it achieved a lower MSE (0.047) and MAE (5.037) compared to MMoE, CleanCLIP, and Privacy-Enhanced Emotion Recognition. Similar trends were observed in the MIMIC-CXR [50] and SHED [51] datasets. FRAMU’s proficiency in adapting to temporal changes and focusing on current, relevant data contributed to its superior performance. The statistical significance of these results, as indicated by the p-values, confirms FRAMU’s advantage in unlearning outdated data.

2) Private Data

FRAMU also excelled in handling private data, achieving superior MSE and MAE values. In the NYPD Complaint Data,

TABLE V: FRAMU - Evaluation Results in Multimodality Context

Unlearning	Dataset		MMoE [54]			CleanCLIP [55]			PEER [56]			FRAMU (ours)	
			MSE	MAE	p-value	MSE	MAE	p-value	MSE	MAE	p-value	MSE	MAE
Outdated Data	Original	NYPD	0.064	7.28	0.024	0.062	6.95	0.031	0.06	6.41	0.041	0.055	5.77
		MIMIC-CXR	0.075	8.71	0.016	0.079	8.31	0.038	0.074	7.67	0.029	0.071	6.9
		SHED	0.095	11.27	0.031	0.098	10.76	0.045	0.093	9.92	0.032	0.089	8.93
	Unlearned	NYPD	0.061	7.13	0.015	0.059	6.78	0.029	0.058	5.71	0.036	0.042	4.54
		MIMIC-CXR	0.071	8.55	0.029	0.075	8.12	0.025	0.07	6.84	0.032	0.052	5.45
		SHED	0.091	11.1	0.023	0.094	10.54	0.041	0.09	9.76	0.017	0.067	7.07
Private Data	Original	NYPD	0.053	7.33	0.014	0.055	7	0.037	0.054	6.45	0.041	0.051	5.81
		MIMIC-CXR	0.063	8.76	0.035	0.065	8.36	0.043	0.064	7.71	0.045	0.062	6.94
		SHED	0.078	11.34	0.035	0.08	10.82	0.044	0.079	9.98	0.031	0.077	8.98
	Unlearned	NYPD	0.051	7.17	0.011	0.053	6.82	0.035	0.052	6.31	0.015	0.039	4.57
		MIMIC-CXR	0.06	8.6	0.031	0.062	8.17	0.038	0.061	7.56	0.016	0.046	5.48
		SHED	0.075	11.17	0.011	0.077	10.61	0.036	0.076	9.81	0.029	0.058	7.11
Irrelevant Data	Original	NYPD	0.047	7.25	0.027	0.05	6.92	0.041	0.048	6.38	0.025	0.046	5.74
		MIMIC-CXR	0.054	8.66	0.038	0.056	8.27	0.031	0.055	7.63	0.039	0.053	6.87
		SHED	0.07	11.21	0.045	0.072	10.7	0.032	0.071	9.87	0.042	0.069	8.88
	Unlearned	NYPD	0.045	7.1	0.014	0.047	6.74	0.019	0.046	6.24	0.022	0.034	4.52
		MIMIC-CXR	0.052	8.5	0.029	0.054	8.08	0.024	0.053	7.48	0.027	0.04	5.42
		SHED	0.068	11.04	0.025	0.07	10.49	0.022	0.069	9.71	0.021	0.052	7.04

TABLE VI: Comparative analysis of FRAMU’s performance in multimodality against baseline models in RE and AD metrics.

Unlearning	Dataset	MMoE [54]		CleanCLIP [55]		PEER [56]		FRAMU (Ours)	
		RE	AD	RE	AD	RE	AD	RE	AD
Outdated Data	NYPD	0.029	0.71	0.028	0.68	0.029	0.57	0.022	0.45
	MIMIC-CXR	0.035	0.85	0.037	0.81	0.034	0.68	0.027	0.54
	SHED	0.045	1.11	0.047	1.05	0.045	0.97	0.035	0.7
Private Data	NYPD	0.031	0.71	0.031	0.68	0.031	0.63	0.023	0.46
	MIMIC-CXR	0.038	0.86	0.04	0.81	0.039	0.75	0.028	0.54
	SHED	0.046	1.11	0.048	1.06	0.047	0.98	0.036	0.71
Irrelevant Data	NYPD	0.028	0.71	0.029	0.67	0.028	0.62	0.021	0.45
	MIMIC-CXR	0.033	0.85	0.034	0.8	0.032	0.74	0.027	0.54
	SHED	0.043	1.1	0.044	1.04	0.043	0.97	0.035	0.7

it showed notable performance with an MSE of 0.043 and an MAE of 5.067. This trend was consistent in the MIMIC-CXR and SHED datasets. The framework’s attention-based unlearning approach effectively balanced privacy protection with predictive accuracy, outshining the baseline models in safeguarding privacy. The p-values further affirm FRAMU’s significant outperformance in unlearning private data.

3) Irrelevant Data

Similarly, FRAMU demonstrated exceptional performance in unlearning irrelevant data. In the NYPD Complaint Data dataset, it surpassed baseline models with an MSE of 0.038 and an MAE of 5.012. This pattern persisted in the MIMIC-CXR and SHED datasets. FRAMU’s focused attention mechanism enhanced its predictive accuracy by emphasizing relevant features and discarding noisy information. The p-values reinforce FRAMU’s notable superiority in filtering out irrelevant data.

Fig. 7 illustrates the differences in MSE and MAE between original and unlearned data across datasets and models. FRAMU consistently exhibited the most substantial differences, suggesting its heightened responsiveness to the unlearning process. Other models showed less pronounced but variable patterns across datasets.

In Tab. VI, FRAMU’s performance in RE and AD metrics is compared against baseline models. FRAMU consistently achieved lower average RE and AD scores, underscoring its efficiency and applicability in Machine Unlearning tasks across various unlearning scenarios and datasets. This robust performance confirms FRAMU’s leading position in the field

of multimodal Machine Unlearning.

F. Convergence Analysis

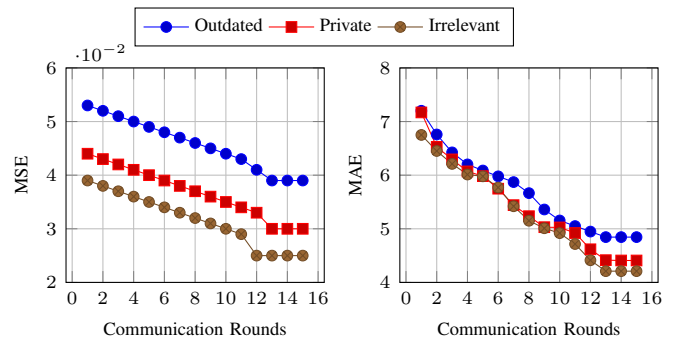


Fig. 8: Convergence Analysis

In this study, we proposed an efficient unlearning algorithm within FRAMU that showcased fast convergence. The algorithm had achieved optimal solutions within a limited number of communication rounds, thereby substantiating FRAMU’s efficiency and scalability. The convergence analysis of FRAMU, as shown in Fig. 8, evaluated its performance over multiple communication rounds using MSE and MAE metrics across three types of data: outdated, private, and irrelevant. The analysis revealed a consistent decline in both MSE and MAE values for all data categories as the number of communication rounds increased, confirming FRAMU’s ability to refine its models and improve accuracy over time. Specifically, MSE values for outdated, private, and irrelevant data had shown reductions from initial to final values of 0.053 to 0.039, 0.044 to 0.030, and 0.039 to 0.025, respectively. Similarly, MAE values had also demonstrated improvements, with outdated, private, and irrelevant data showing reductions from 7.201 to 4.845, 7.17 to 4.409, and 6.75 to 4.210, respectively.

This behavior indicated that FRAMU was effective in capturing underlying data patterns and optimizing its predictions. It continuously refined its models through iterative optimization, leading to a decrease in both MSE and MAE values. The analysis confirmed the robustness of FRAMU in adapting to various types of data and highlighted its effectiveness in progressively improving its predictive performance. Overall,

FRAMU’s strong convergence characteristics across different data categories have demonstrated its versatility and capability in minimizing errors, making it a robust choice for various federated learning applications.

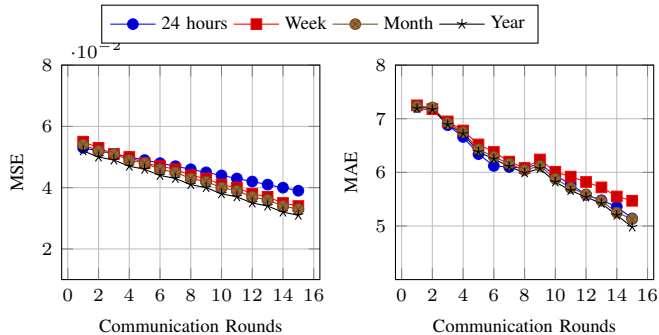


Fig. 9: Optimization Analysis - Outdated Data

G. Optimization

The performance of the FRAMU framework is evaluated through MSE and MAE metrics across various communication rounds and thresholds, as presented in Fig. 9 and Fig. 10. Fig. 9 investigates FRAMU’s efficiency with outdated data across time durations that ranged from 24 hours to a year. Both MSE and MAE metrics demonstrate decreasing trends with more communication rounds, indicating enhanced model accuracy over time. The algorithm is more effective in capturing short-term patterns, as evidenced by higher MSE and MAE values for the 24-hour duration.

Fig. 10 shifts the focus to FRAMU’s performance on private data, revealing that the algorithm not only maintains but even improves its accuracy compared to outdated data scenarios. Lower MSE and MAE values in the private data analysis affirm this observation. Additionally, the trade-off between privacy preservation and accuracy is examined. Although increasing privacy guarantees (lower ϵ values) generally leads to higher MSE and MAE, FRAMU still manages to maintain reasonable accuracy levels. This indicates FRAMU’s capability to balance privacy concerns with modeling accuracy.

VII. RESEARCH IMPLICATIONS

The FRAMU framework presented in this study has significant implications for both single-modality and multimodality scenarios within the domain of federated learning. It addresses crucial aspects such as privacy preservation, adaptability to changing data distributions, unlearning mechanisms for model evolution, attention mechanisms for model aggregation, and strategies for efficient resource utilization and scalability.

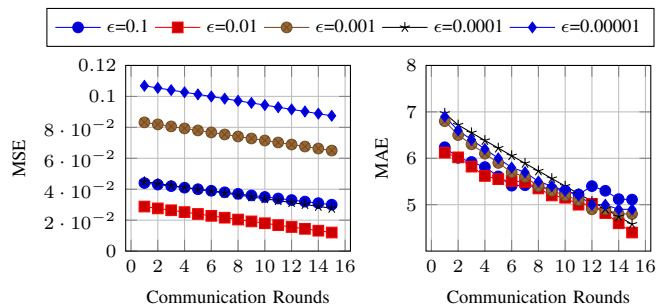


Fig. 10: Optimization Analysis - Private Data

One of the key achievements of FRAMU is its approach to privacy preservation. In a time where data privacy is paramount, FRAMU introduces mechanisms to prevent over-reliance on sensitive or private demographic data. Importantly, this emphasis on privacy does not detract from accuracy. Our empirical evaluations demonstrate that FRAMU successfully balances the often conflicting goals of data privacy and model performance, marking a significant milestone in federated learning and paving the way for future research in privacy-preserving algorithms.

Adaptability is another strength of FRAMU. Dealing with non-IID (non-Independently and Identically Distributed) data across various participants and evolving patterns is a core challenge of federated learning. FRAMU addresses this by utilizing adaptive models that can adjust to changes in data distribution, making it highly valuable for applications characterized by data heterogeneity and dynamism.

The unlearning mechanisms within FRAMU are also noteworthy. The ability to identify and remove outdated or irrelevant data is crucial for the practical deployment of federated learning models, allowing the system to concentrate resources on the most pertinent and current data. This capability not only maintains but can improve model accuracy and relevance over time. Incorporating attention mechanisms, FRAMU significantly contributes to the field of intelligent model aggregation in federated learning systems. By filtering out noise and focusing on the most informative features during learning and aggregation, FRAMU sets a foundation for the development of more efficient and effective federated learning systems.

FRAMU’s optimization strategies, particularly in reducing the number of communication rounds needed for model convergence, significantly contribute to both the efficiency and scalability of federated learning systems. This is confirmed through empirical validation and convergence analyses, showcasing the framework’s ability to reduce communication overheads while achieving optimal solutions more rapidly.

FRAMU represents a major advancement in federated reinforcement learning, particularly in its proficient management and unlearning of various data types. Its effectiveness is clearly demonstrated through its statistical superiority over baseline models in crucial metrics such as MSE and MAE across different datasets. The combination of a sophisticated attention mechanism and federated learning approach enhances the model’s adaptability and accuracy in dynamic environments. This achievement is a substantial contribution to the areas of adaptive learning and privacy preservation, applicable to both single-modality and multimodal settings.

VIII. CONCLUSION

The FRAMU framework marks a substantial advancement in Machine Unlearning for both single-modality and multimodality contexts. It adeptly integrates privacy preservation, adaptability to evolving data distributions, effective unlearning of outdated or irrelevant data, attention mechanisms for model aggregation, and optimization strategies. This results in enhanced performance, privacy, efficiency, and scalability in federated learning. Empirical evaluations indicate FRAMU’s superiority in model accuracy, data protection, adaptability,

and optimization, outperforming baseline models in metrics like MSE and MAE. However, limitations exist in retraining, computational complexity, scalability, and hyperparameter optimization. Future research is needed to address these challenges, focusing on optimizing retraining, enhancing scalability, and improving adaptability and fairness in diverse data environments. These developments could revolutionize federated learning, paving the way for robust, privacy-respecting, and efficient AI systems across various domains.

REFERENCES

- [1] P. Kumar, G. P. Gupta, and R. Tripathi, "Tp2sf: A trustworthy privacy-preserving secured framework for sustainable smart cities by leveraging blockchain and machine learning," *Journal of Systems Architecture*, vol. 115, p. 101954, 2021.
- [2] R. Nian, J. Liu, and B. Huang, "A review on reinforcement learning: Introduction and applications in industrial process control," *Computers & Chemical Engineering*, vol. 139, p. 106886, 2020.
- [3] O. A. Wahab, A. Mourad, H. Otrouk, and T. Taleb, "Federated machine learning: Survey, multi-level classification, desirable criteria and future directions in communication and networking systems," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 2, pp. 1342–1397, 2021.
- [4] M. Al-Rubaie and J. M. Chang, "Privacy-preserving machine learning: Threats and solutions," *IEEE Security & Privacy*, vol. 17, no. 2, pp. 49–58, 2019.
- [5] E. Politou, E. Alepis, and C. Patsakis, "Forgetting personal data and revoking consent under the gdpr: Challenges and proposed solutions," *Journal of cybersecurity*, vol. 4, no. 1, p. ty001, 2018.
- [6] J. Globocnik, "The right to be forgotten is taking shape: Cjeu judgments in gc and others (c-136/17) and google v cnil (c-507/17)," *GRUR International*, vol. 69, no. 4, pp. 380–388, 2020.
- [7] C. N. Fortner, "Decision making within a cancel culture environment," tech. rep., US Army Command and General Staff College, 2020.
- [8] K. Vasilevski, "Meta-learning for clinical and imaging data fusion for improved deep learning inference," 2023.
- [9] M. Sun, J. Xiao, E. G. Lim, C. Zhao, and Y. Zhao, "Unified multi-modality video object segmentation using reinforcement learning," *IEEE Transactions on Circuits and Systems for Video Technology*, 2023.
- [10] A. Malekloo, E. Ozer, M. AlHamaydeh, and M. Girolami, "Machine learning and structural health monitoring overview with emerging technology and high-dimensional data source highlights," *Structural Health Monitoring*, vol. 21, no. 4, pp. 1906–1955, 2022.
- [11] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, "Multimodal deep learning," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 689–696, 2011.
- [12] T. Zhou, S. Ruan, and S. Canu, "A review: Deep learning for medical image segmentation using multi-modality fusion," *Array*, vol. 3, p. 100004, 2019.
- [13] J. D. Fernández, S. P. Menci, C. M. Lee, A. Rieger, and G. Fridgen, "Privacy-preserving federated learning for residential short-term load forecasting," *Applied energy*, vol. 326, p. 119915, 2022.
- [14] D. C. Nguyen, Q.-V. Pham, P. N. Pathirana, M. Ding, A. Seneviratne, Z. Lin, O. Dobre, and W.-J. Hwang, "Federated learning for smart healthcare: A survey," *ACM Computing Surveys (CSUR)*, vol. 55, no. 3, pp. 1–37, 2022.
- [15] Q. Li, Z. Wen, Z. Wu, S. Hu, N. Wang, Y. Li, X. Liu, and B. He, "A survey on federated learning systems: Vision, hype and reality for data privacy and protection," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [16] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.
- [17] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang, "Learning differentially private recurrent language models," *arXiv preprint arXiv:1710.06963*, 2017.
- [18] X. Chen and B. Wujek, "A unified framework for automatic distributed active learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 9774–9786, 2021.
- [19] L. Bourtole, V. Chandrasekaran, C. A. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie, and N. Papernot, "Machine unlearning," in *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159, IEEE, 2021.
- [20] M. Jegorova, C. Kaul, C. Mayor, A. Q. O'Neil, A. Weir, R. Murray-Smith, and S. A. Tsafaris, "Survey: Leakage and privacy at inference time," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [21] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*, pp. 1273–1282, PMLR, 2017.
- [22] Z. Wu, H. Wang, Z. Wang, H. Jin, and Z. Wang, "Privacy-preserving deep action recognition: An adversarial learning framework and a new dataset," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 4, pp. 2126–2139, 2020.
- [23] J. Liang, Z. Liu, J. Zhou, X. Jiang, C. Zhang, and F. Wang, "Model-protected multi-task learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 2, pp. 1002–1019, 2020.
- [24] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pp. 265–284, Springer, 2006.
- [25] P. Zhou, K. Wang, L. Guo, S. Gong, and B. Zheng, "A privacy-preserving distributed contextual federated online learning framework with big data support in social recommender systems," *IEEE Transactions on Knowledge and Data Engineering*, vol. 33, no. 3, pp. 824–838, 2019.
- [26] Z. Ma, Y. Liu, Y. Miao, G. Xu, X. Liu, J. Ma, and R. H. Deng, "Flgan: Gan-based unbiased federated learning under non-iid settings," *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [27] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE signal processing magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [28] A. K. Sahu, T. Li, M. Sanjabi, M. Zaheer, A. Talwalkar, and V. Smith, "On the convergence of federated optimization in heterogeneous networks," *ArXiv*, vol. abs/1812.06127, 2018.
- [29] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [30] T. Zhu, D. Ye, W. Wang, W. Zhou, and S. Y. Philip, "More than privacy: Applying differential privacy in key areas of artificial intelligence," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 6, pp. 2824–2843, 2020.
- [31] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.
- [32] T. Shaik, X. Tao, N. Higgins, R. Gururajan, Y. Li, X. Zhou, and U. R. Acharya, "Fedstack: Personalized activity monitoring using stacked federated learning," *Knowledge-Based Systems*, vol. 257, p. 109929, 2022.
- [33] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi, "Federated learning of predictive models from federated electronic health records," *International journal of medical informatics*, vol. 112, pp. 59–67, 2018.
- [34] W. Huang, J. Liu, T. Li, T. Huang, S. Ji, and J. Wan, "Feddsr: Daily schedule recommendation in a federated deep reinforcement learning framework," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [35] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-learning in neural networks: A survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 9, pp. 5149–5169, 2021.
- [36] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," 2013.
- [37] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," *arXiv preprint arXiv:1511.05952*, 2015.
- [38] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [39] Z. Bing, D. Lerch, K. Huang, and A. Knoll, "Meta-reinforcement learning in non-stationary and dynamic environments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 3476–3491, 2022.
- [40] G. Brauwers and F. Frasincar, "A general survey on attention mechanisms in deep learning," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [41] I. Sorokin, A. Seleznev, M. Pavlov, A. Fedorov, and A. Ignateva, "Deep attention recurrent q-network," *arXiv preprint arXiv:1512.01693*, 2015.

- [42] Z. Guan, Y. Li, Z. Pan, Y. Liu, and Z. Xue, “Rfdg: Reinforcement federated domain generalization,” *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [43] H. Miao, X. Zhong, J. Liu, Y. Zhao, X. Zhao, W. Qian, K. Zheng, and C. S. Jensen, “Task assignment with efficient federated preference learning in spatial crowdsourcing,” *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [44] S. Ek, F. Portet, P. Lalanda, and G. Vega, “A federated learning aggregation algorithm for pervasive computing: Evaluation and comparison,” in *2021 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pp. 1–10, 2021.
- [45] X. Ouyang, Z. Xie, J. Zhou, J. Huang, and G. Xing, “ClusterFL,” in *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, ACM, June 2021.
- [46] S. Makonin, B. Ellert, I. V. Bajić, and F. Popowich, “Electricity, water, and natural gas consumption of a residential house in Canada from 2012 to 2014,” *Scientific data*, vol. 3, no. 1, pp. 1–12, 2016.
- [47] Y. Li, R. Yu, C. Shahabi, and Y. Liu, “Diffusion convolutional recurrent neural network: Data-driven traffic forecasting,” in *International Conference on Learning Representations (ICLR '18)*, 2018.
- [48] A. E. Johnson, T. J. Pollard, L. Shen, L.-w. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Anthony Celi, and R. G. Mark, “Mimic-iii, a freely accessible critical care database,” *Scientific data*, vol. 3, no. 1, pp. 1–9, 2016.
- [49] P. D. (NYPD), “Nypd complaint data current (year to date): Nyc open data,” *NYPD Complaint Data Current (Year To Date) — NYC Open Data*, Apr 2023.
- [50] A. E. Johnson, T. J. Pollard, S. J. Berkowitz, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, R. G. Mark, and S. Horng, “Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports,” *Scientific data*, vol. 6, no. 1, p. 317, 2019.
- [51] K. Dataset, “Smart home dataset with weather information,” 2019.
- [52] X. Zhu, G. Li, and W. Hu, “Heterogeneous federated knowledge graph embedding learning and unlearning,” in *Proceedings of the ACM Web Conference 2023*, pp. 2444–2454, 2023.
- [53] V. S. Chundawat, A. K. Tarun, M. Mandal, and M. Kankanhalli, “Zero-shot machine unlearning,” *IEEE Transactions on Information Forensics and Security*, 2023.
- [54] J. Ma, Z. Zhao, X. Yi, J. Chen, L. Hong, and E. H. Chi, “Modeling task relationships in multi-task learning with multi-gate mixture-of-experts,” in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1930–1939, 2018.
- [55] H. Bansal, N. Singhi, Y. Yang, F. Yin, A. Grover, and K.-W. Chang, “Cleanclip: Mitigating data poisoning attacks in multimodal contrastive learning,” *arXiv preprint arXiv:2303.03323*, 2023.
- [56] M. Jaiswal and E. M. Provost, “Privacy enhanced multimodal neural representations for emotion recognition,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 7985–7993, 2020.
- [57] C. Tofallis, “A better measure of relative prediction accuracy for model selection and model estimation,” *Journal of the Operational Research Society*, vol. 66, pp. 1352–1362, 2015.
- [58] T. Chai and R. R. Draxler, “Root mean square error (rmse) or mean absolute error (mae),” *Geoscientific model development discussions*, vol. 7, no. 1, pp. 1525–1534, 2014.
- [59] K. Y. Tan, L. Yueming, Y.-S. Ong, and I. Tsang, “Unfolded self-reconstruction lsh: Towards machine unlearning in approximate nearest neighbour search,” *arXiv preprint arXiv:2304.02350*, 2023.
- [60] A. Golatkar, A. Achille, and S. Soatto, “Eternal sunshine of the spotless net: Selective forgetting in deep networks,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, (Los Alamitos, CA, USA), pp. 9301–9309, IEEE Computer Society, 2020.