International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(23s) 2026 pp. 894-908
https://doi.org/10.70917/ijcisim-2026-5652
Article

Federated Clinical Meta-Learning for Personalized Treatment-Response Prediction in Heterogeneous Oncology Data

Roma Fayaz[1], Sharyar Wani[2], Nasir Nazir[3]

[123] KICT, International Islamic University Malaysia, Kuala Lumpur, Malaysia.
[1] Email: f.roma@live.iium.edu.my ORCID: 0009-0002-4111-0276
[2] Email: sharyarwani@iium.edu.my ORCID: 0000-0001-6812-0066
[3] Email: nasir.nazir@live.iium.edu.my ORCID: 0009-0009-2632-7216

Abstract: Artificial intelligence can support treatment-response prediction in oncology, but centralized learning is constrained by institutional data heterogeneity and restrictions on sharing patient-level records. This study presents a federated meta-learning framework that combines learnable representations of mixed clinical variables with client-level adaptation and a lightweight personalization mechanism. De-identified oncology records from the Genomic Data Commons are used to simulate a multi-client environment in which raw records remain local and only model updates are aggregated. The framework is intended to improve adaptation under non-independent and identically distributed client data while preserving data locality. The experimental conclusions, including comparative accuracy, area under the receiver-operating-characteristic curve, and communication efficiency, must be finalized from the archived evaluation outputs before submission. The current evidence therefore supports the proposed architecture and evaluation protocol, but not yet a definitive numerical superiority claim.
Keywords: Federated learning, meta-learning, oncology informatics, personalized prediction, privacy-preserving machine learning, treatment response

1. INTRODUCTION

The growing access to large-scale clinical and molecular oncology data has opened up new possibilities of data-driven solutions to assist in the personalization of treatment decisions [1]. The development of machine learning has made it possible to predict complex patient features, tumor profiles, and treatment histories to determine the response to treatment and clinical outcomes [2]. Nonetheless, the oncological data are always heterogeneous, as they are obtained in different hospitals, cohorts, and in various clinical settings, each of which has different rates of patients and treatment regimens [3]. Meanwhile, the privacy policies and the ethical considerations restrict direct exchange of patient-level information between institutions [4]. Such difficulties require computational models that are able to learn on distributed clinical data without violating privacy, and that can be able to condition predictions to specific patients, which is the key to personalized medicine.

The research has investigated machine learning and deep learning methods of predicting treatment response on centralized datasets and has already shown promising results on particular cancer types [5]. Recent studies have added federated learning to overcome the privacy issue by providing an option to learn together without exposing raw patient data [6]. Metalearning techniques have been independently suggested to enhance model generalization and can response quickly to new tasks or new datasets. Despite these developments, the majority of the existing works have significant flaws. First, federated learning algorithms usually depend on global model averaging, which fails to deal with high heterogeneity between institutions. Second, most of the models make use of a simple encoding scheme to treat clinical variables, which restricts their capability to reflect complicated relationship among heterogeneous clinical

features. Third, current strategies are usually aimed at population-level forecasts and do not provide any mechanisms of personalizing them to the patient level, which is necessary to predict treatment responses correctly. Consequently, existing approaches are not capable of generalizing to different clinical settings or cannot be sufficiently used to make personalized decisions.

In an attempt to overcome these drawbacks, this research develops a Federated Clinical Meta-Learning Architecture to predict Personalized Treatment Response in heterogeneous Oncology Data Ecosystems [7]. The suggested system incorporates federated learning and meta-learning, allowing to approach privacy-preserving distributed model training and ensure quick adaptation to new patients and clinical environments. Computer science-wise, the essential novelty is two-fold: (i) an approach to embedding clinical features to learn that offers personalized predictions of heterogeneous clinical variables by training a global model, and (ii) a patient-specific personalization layer, which trains only the final model without re-training the global model. These modules are integrated into a federated meta-learning structure, through which the model is able to extrapolate across institutions and make personalized predictions of treatment responses. Using open clinical oncological data of the Genomic Data Commons (GDC) and emulating a multi-institutional setting, the proposed study shows how advanced machine learning systems can work towards delivering privacy, model generalization, and personalized clinical decision-making services.

A. Problem Statement

The growing data-driven strategies in oncology have facilitated predictive modeling of treatment response on a basis of clinical and molecular patient information [8]. Machine learning methods have demonstrated potential in the analysis of complex oncology data in order to aid in individualized treatment decisions. Nonetheless, the data on oncology is heterogeneous in nature, and it can be obtained in various institutions, which have different patient groups, regimens, and characteristics distributions [9]. Meanwhile, on-the-one hand, stringent privacy controls impede centralized data exchange, reducing the scalability and cooperation capabilities of traditional machine learning systems. The current solutions are either based on centralized learning which undermines the privacy of patients or follows federated learning models that prioritize privacy over heterogeneity of data and personalization at the patient level. In addition, most existing models adopt simplistic methods of feature encoding which do not reflect the latent associations between a variety of clinical variables resulting in poor prediction of treatment responses. The insufficient adaptive capabilities to make predictions in individual patients also restricts clinical applicability of these systems. Such drawbacks may lead to poor predictions, diminished cross-institutional generalization, and lack of support of individualized treatment planning. In order to mitigate these challenges, this research paper seeks to create a federated clinical meta-learning system, which combines learnable clinical feature representations via embeddings and patient-specific personalization. The proposed framework aims to facilitate privacy-protected, heterogeneous learning of data and provide accurate and personalized predictions of response to treatment to distributed oncology data ecosystems.

B. Research Motivation

The reason behind this study is that it is important to establish smart and privacy-aware computational models that can aid in personalized prediction of response to treatment in cancer. Although clinical data in scale are being generated across the institutions, the magnitude of the problem of privacy and heterogeneity of the data makes effective centralized analysis quite difficult. Current machine learning solutions are not able to generalize to various clinical settings and do not personalize at the patient level. The present research is driven by the fact that federated meta-learning with sophisticated feature representation and personalization strategies can help to address those issues and achieve clinically relevant predictive performance.

C. Research Significance

The study is of importance because it touches upon important issues in the application of machine learning to oncology data, especially the data privacy, heterogeneity and personalization. The proposed framework can allow collaborative learning without patient confidentiality loss through the integration of federated learning with meta-learning. Learnable clinical feature embeddings and a personalization layer that is patient-specific, contribute to a

more adaptable model in a wide range of clinical settings, as well as to the increase in accuracy in prediction. The paper has added a generalizable and scalable computational methodology that aids in predicting the response to personalized treatment, which promotes the ability of intelligent decision-support systems to make precise decisions in oncology.

D. Key Contributions

- A federated meta-learning architecture for binary treatment-response prediction from heterogeneous clinical oncology variables while retaining patient records at the client level.
- Learnable embeddings and projections that represent categorical and numerical clinical variables in a shared latent space.
- A lightweight adaptation mechanism intended to personalize predictions without retraining the complete global model.
- An evaluation design covering client heterogeneity, predictive performance, personalization, and communication-round behavior.

E. Paper Organization

The remainder of this paper is organized as follows. Section II reviews federated and personalized learning in healthcare. Section III describes the dataset, preprocessing, feature representation, federated partitioning, meta-learning procedure, and personalization mechanism. Section IV presents the experimental results and limitations. Section V concludes the paper and outlines future work.

2. RELATED WORK

Liu et al. [10] explored federated learning in the context of prediction of treatment response in multicenter non-small cell lung cancer (NSCLC) environments where direct data transfer is limited. The experiment suggested a CNN-based federated learning model whereby it is possible to train the model locally at each hospital and aggregate the parameters centrally without infringing on the privacy of the patients. They utilized CT data that were of 245 patients with NSCLC in four hospitals receiving chemoradiotherapy and simulated and realistic federated setups. Accuracy, AUC, ROC curves, and confusion matrices were used to measure model performance. The findings indicated that the federated model obtained similar or even better AUC values compared to centralized deep learning. Nonetheless, imaging data were studied only and the research did not include personalization mechanisms, which indicates the potential to add more clinical features and patient-specific modeling.

Liu et al. [8] proposed an individualized contrastive representative federated learning system (PCRFed) to overcome

performance reduction due to non-IID image data distributions of multi-center medical images. The approach combines contrastive representation learning and a weighted modelcontrastive loss to dynamically control local client models to the global model to produce client-level personalization without extra communication cost. This framework was tested on two non IID multicenter medical image segmentation datasets with encoder-decoder architectures. Results of the experiment illustrated better global and single-client results as compared to the state-of-the-art FL approaches with an advance of 2.63 percent in mean Dice score of prostate segmentation. The paper, however, concentrated on the tasks of segmentation and did not investigate the field of clinical outcome prediction or heterogeneous integration of clinical features.

Mateus et al. [11] used federated learning to explore data sharing constraints between multimodal biological aging markers. The training of a federated deep learning model was done on three population-based cohorts to estimate BrainAge using brain MRI and compared with locally trained models. Federated approach had much lower age prediction error, which was further improved by cohort age-interval harmonization. To test the relationship between BrainAge and MetaboAge using the metabolomics data, association and survival analyses were performed. The

findings showed direct weak association, but complementary predictive value of mortality risk. The article results in the effectiveness of FL in multimodal analysis and fails to mention individualized clinical prediction tasks or heterogeneous model of treatment-response.

Alsulaimawi [12] presented a prototype, Meta-Federated Learning (Meta-FL), to help solve the major problems of federated learning, namely the heterogeneity of data and the diversity of models in medical settings. The paper presents a meta-aggregator based on an optimization that uses metafeatures of local models to carry out the accuracy-aware aggregation and provide more adaptive and individualized global model updates. Empirically, it was proven that the framework outperforms the traditional FL approaches in terms of accuracy, robustness, and scalability because it was tested on four healthcare-related datasets. Findings also demonstrated that Meta-FL can deliver better results with fewer rounds of communication and emphasize that it is efficient in large-scale and heterogeneous federated healthcare networks, even though it is not directly oriented on personalized clinical outcome prediction problems.

Jia et al. [1] presented a federated learning system augmented with meta-learning to overcome the major limitations, i.e., communication inefficiency, heterogeneity of clients, and non-IID data distributions. The suggested method combines a federated modulator, which finds contextual information based on local data batches and produces modulation parameters in order to adaptively modify the activations of a base model. The personalization strategy is a MAML-based approach that can enhance client-specific adjustment. The tests on various benchmark datasets showed that its convergence is faster and better predictive than traditional FL. Although the framework has been effective in enhancing generalization and efficiency, it is not directly proven to have been effective on clinical or oncology-based treatment response prediction tasks.

Pan et al. [13] presented privacy-aware clinical risk prediction with electronic health records by introducing adaptive federated learning framework to manage the data distribution drift across institutions. It involves breaking down features of inputs into stable, domain-specific, and conditional-irrelevant features, according to how they are related to clinical outcomes, facilitating stronger cross-site learning. The structure was tested on multi-institutional ICU data to predict the onset of sepsis and acute kidney injuries. The experimental findings showed that it had higher predictive performance than the conventional FL baselines, in addition to providing interpretable features. In spite of its ability to work with EHR-based risk prediction, the study is not concerned with oncology data and individual response to treatment modeling.

Feng et al. [9] examined CT-based privacy-preserving disease risk prediction by building a powerful federated learning model to detect high-risk recurrence after gastric cancer surgery in several facilities. The research applied four independent hospitals of multicenter CT data to overcome the data isolation and privacy limitations without the central data sharing. The federated model was shown to be highly predictive in nature with an AUC of 0.710 to 0.869 betweencentres. This was further analyzed to find out the adaptive and common imaging features that lead to prediction. Although the study is effective in the context of recurrence risk stratification in gastric cancer it does not include meta-learning and personalized adaptation to individualized treatment response prediction.

Gouthamchand et al. [14] constructed and assessed multicentric clinical and radiomics outcome prediction models in head and neck cancer in a federation learning model. By using the dataset on India and the Netherlands and follow-ups of F.A.I.R. data principles, the research utilized the Vantage6 infrastructure on the paradigm of Personal Health Train to guarantee the privacy of data. In six datasets (294 patients in each) of 1,131 patients, overall survival, distant metastasis, and locoregional recurrence were modeled by correlation-based Feature Selection and LASSO-regularized Cox regression. The findings showed that federated learning can be used to achieve powerful, privacy-preserving survival prediction on heterogeneous cancer cohorts, but the methodology does not involve personalized meta-learning processes.

Haripriya et al. [15] nearer to investigated the privacypreserving medical image classification by combining federated learning with transfer learning in a simulated multi-centre health care setting. The study addressed the issue of data heterogeneity and data privacy using ResNet and VGG16 architectures that had been pre-trained on ImageNet and fine-tuned on tuberculosis chest X-rays, brain tumor MRI, and diabetic retinopathy data sets. One of the

contributions is an adaptive aggregation strategy switching between FedAvg and FedSGD in real-time based on divergence in real-time data between clients. The test outcomes revealed the same rate of classification up to 96.3 percent and a reduction of about 20 percent in execution and competitive F1-scores. Although efficient in image classification, the method is not specific to oncology response to treatment or treatment response prediction and personalized clinical.

Adnan et al. [16] presented a comprehensive privacy and explainable lung cancer prediction model in order to address the difficulties associated with massive data processing, data privacy, lack or light collaboration between institutions, and model interpretability. The research integrated MapReduce to enable scaleable computation, private blockchain to enable data management and handling in a secure and tamper-proof way, federated learning to facilitate model training and no data sharing, and explainable AI to increase clinical trust. Experimental scoring proved to be better in terms of predictive accuracy attaining 98.21 with a low miss rate of 1.79. Though the framework is very strong in terms of scalability, security and interpretability, it is more concerned with the accuracy of detection than the prediction of individual treatment response to different heterogeneous oncology data ecosystems.

TABLE I. SUMMARY ON LITERATURE REVIEW

Author	Key Findings	Advantages	Disadvantages
Liu et al. [10]	Federated CNN model predicts treatment response in NSCLC using CT images across hospitals.	Protects privacy; works across multiple centers.	Uses only images; no personalization.
Liu et al. [8]	Personalized FL using contrastive learning improves performance on non-IID data.	Handles data heterogeneity well.	Focused on segmentation, not treatment response.
Mateus et al. [11]	FL model learns biological age from MRI and metabolomics data.	Supports multimodal learning; privacypreserving.	Not cancerfocused; no treatment prediction.
Alsulaimawi [12]	Meta-FL improves model aggregation in heterogeneous healthcare data.	Efficient and scalable aggregation.	No clinical outcome or oncology focus.
Jia et al. [1]	Meta-learning with FL improves personalization and convergence.	Better generalization; faster learning.	Not tested on clinical oncology data.
Pan et al. [13]	Adaptive FL improves EHRbased risk prediction under data drift.	Interpretable and robust.	Not applied to oncology.
Feng et al. [9]	FL predicts gastric cancer recurrence using CT images.	Good performance; privacy preserved.	No personalization or metalearning.
Gouthamchand et al. [14]	FL-based survival prediction for head and neck cancer.	Secure multicenter modeling.	No personalized deep learning.

Haripriya et al. [15]	Adaptive FL improves medical image classification efficiency.	Faster training; adaptive aggregation.	Not treatmentresponse focused.
Adnan et al. [16]	FL with blockchain and XAI improves lung cancer detection.	Secure and interpretable.	Focuses on detection, not treatment response.

Despite significant progress in federated learning for healthcare and oncology, existing studies primarily focus on privacy-preserving collaboration and performance improvement without adequately addressing personalized treatment response prediction. Most oncology-focused works rely heavily on imaging data, with limited integration of heterogeneous clinical variables such as tumor stage, histology, and treatment regimens. While some approaches handle data heterogeneity through adaptive aggregation or contrastive learning, they lack patient-specific personalization mechanisms. Meta-learningbased FL frameworks improve generalization but are rarely validated on real-world oncology treatment outcomes. Furthermore, existing models often emphasize detection or survival analysis rather than individualized response prediction. The absence of learnable clinical feature representations restricts effective modeling of complex patient variability. Additionally, cross-institutional heterogeneity is typically handled at the model level rather than at the clinical feature level. Collectively, these limitations highlight the need for a federated meta-learning framework that integrates heterogeneous clinical data with patient-specific personalization to enable accurate and privacypreserving treatment response prediction in oncology.

3. METHODOLOGY

The proposed methodology is designed for privacy-preserving treatment-response prediction from heterogeneous oncology data. It comprises data acquisition and preprocessing, mixed-type feature representation, simulated client partitioning, local adaptation, server-side aggregation, and prediction refinement. Because the institutional clients are simulated from a public dataset, the study evaluates algorithmic behavior under distributed partitions; it does not claim deployment across real hospitals or formal privacy guarantees.

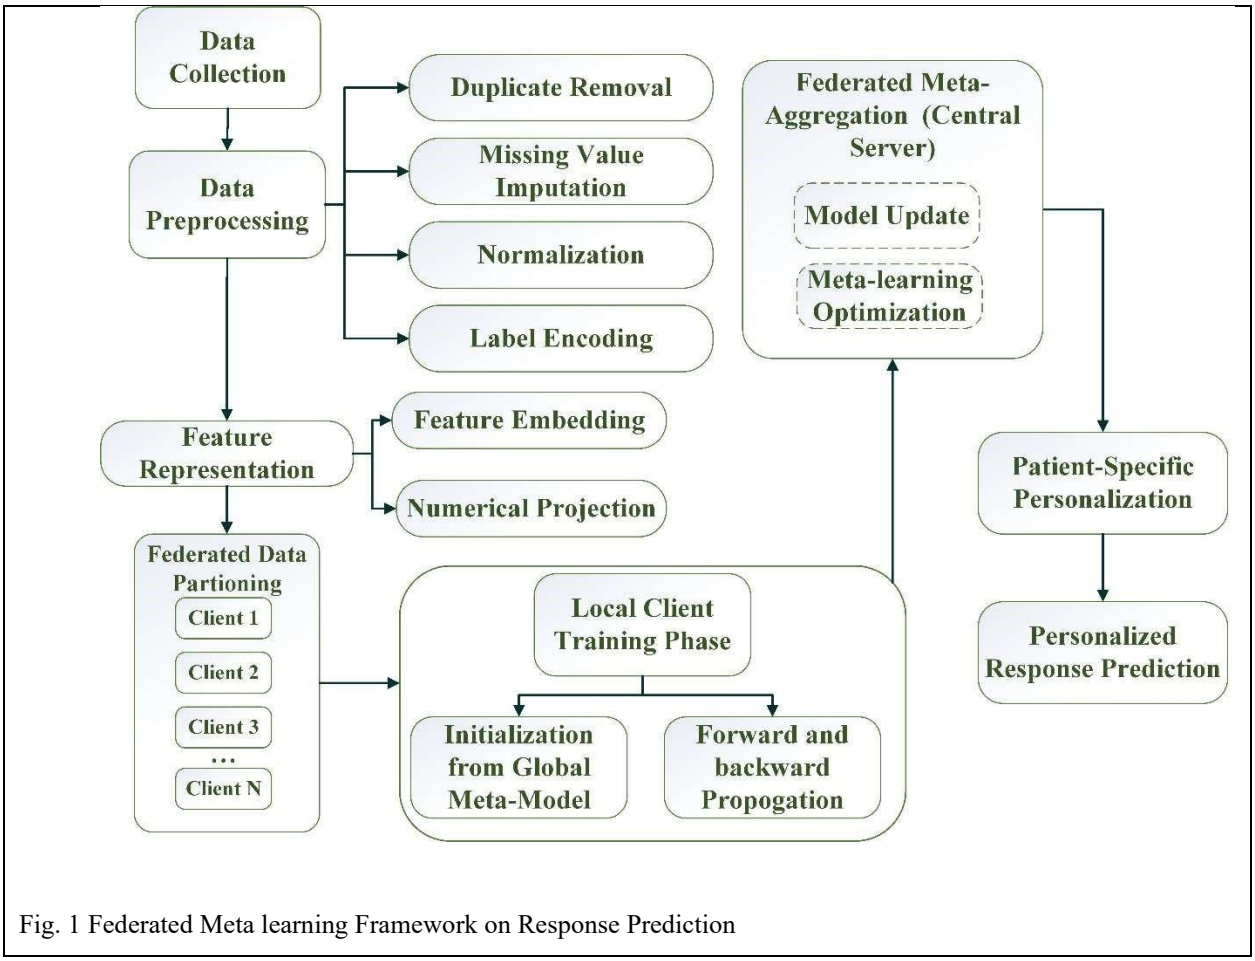

Fig. 1 Federated Meta learning Framework on Response Prediction

A. Data Collection

De-identified clinical oncology records were obtained from the Genomic Data Commons (GDC) Data Portal [17]. Before submission, this subsection must identify the exact GDC project or cancer cohort, download date, inclusion and exclusion criteria, number of unique patients, treatment variables, follow-up fields, and the operational definition of treatment response. These details are necessary for reproducibility and external assessment of the clinical task.

B. Data Preprocessing

Preprocessing of data is an important procedure of assuring reliability and consistency of heterogeneous clinical oncology data across various institutions. Raw clinical data is not directly applicable to federated meta-learning because of the difference in data completeness, scale, and format. Correct preprocessing enhances reduced noise, reduced bias, and model convergence. Preprocessing in this work is concerned with the absence of values, numerical features standardization, the encoding of categorical clinical variables, and outcome label preparation. The steps will guarantee the data are adequate to the privacypreserving federated training and customized predicting treatment responses.

1) Data Cleaning and Duplicate Removal: Large repositories of clinical data might have duplicate or overlapping records of patients. Elimination of duplicate entries ensures that there will be no leakage of data across federated clients and will eliminate learning bias due to repeated samples. This measure makes sure that the patients make only one contribution to the learning process.

$$D\_clean = D \setminus D\_dup, \qquad (1)$$

Where, D is the original clinical dataset, Ddup is duplicated patient records, Dclean represents the cleaned dataset after duplicate removal.

2) Missing Value Imputation: The lack of values in clinical oncology data is typical because of either missing documentation or loss of follow-up. To maintain the integrity of the dataset, the imputation of the numerical features is done with the median, whereas the imputation of the categorical features is done with the mode. This method is resistant to outliers, and it does not distort the underlying data distribution.

x_imp = median(X), if x is a missing numerical value; *(2a)*

x_imp = mode(X), if x is a missing categorical value. *(2b)*

Where, x is missing feature value, X represents the set of observed values for that feature, $median(X)$ represents the median of numerical values, $mode(X)$ represents the most frequent categorical value.

3) Normalization: The age, tumor size, or biomarkers values are clinical numerical features that differ in different scales. The form of normalization that is used is Z-score normalization so that all numerical features have an equal

contribution during federated meta-learning. This enhances the stability and convergence of training in heterogeneous clients.

x_norm = (x - μ) / σ. *(3)*

where x is the original value, μ and σ are the mean and standard deviation estimated from the training portion, and x_norm is the standardized value. Statistics must be fitted without using validation or test records.

x represents the normalized feature value.

Feature encoding for embedding: Each categorical value c_i is mapped to an integer index and then to a trainable d-dimensional vector. Unknown and missing categories require dedicated indices fitted from the training data only.

$$c\_i \rightarrow e\_i \in R^d. \quad (4)$$

Where, c represents a categorical clinical feature, e represents its corresponding embedding vector,

d represents the embedding dimension.

Label construction: Treatment response is formulated as a binary outcome, with responders assigned y = 1 and non-responders assigned y = 0. Before submission, the manuscript must state the exact GDC fields, assessment window, censoring rules, and exclusions used to derive this label.

whereas those who do not respond will be termed nonresponders. This formulation is consistent with the personalized treatment response prediction goals.

y = 1 for a responder and y = 0 for a non-responder. *(5)*

Where, $y$represents the treatment response label, 1 indicates a positive treatment response, 0 indicates no or poor treatment response.

C. Feature Representation Using Learnable Embeddings

Mixed categorical and numerical clinical variables are represented in a common latent space. The exact variables included in the final experiment, their missingness, cardinalities, embedding dimensions, and preprocessing decisions must be reported in a reproducibility table.

In the case of categorical clinical features, the categorical traits were represented as embedded vectors of dense and lowdimensional vectors using an embedding lookup table. Let c denote the i -th categorical feature, then its embedding representation is defined as:

$$e\_i = E\_i[c\_i]. \quad (6)$$

where $\mathbf{E} \in \mathbb{R}| |\times$ represents the learnable embedding matrix for the i-th categorical feature, $| C |$denotes the number of unique categories, and $d$is the embedding dimension. The optimization during the training of these embeddings aims to encode the clinically important similarities between the dissimilar categories, e.g., similarities between the stage of adjacent tumours.

In the case of numerical clinical features, the normalized features were mapped to the same latent embedding space by a linear mapping.

$$e\_j = W\_j x\_j + b\_j. \quad (7)$$

where $\mathbf{W} \in \mathbb{R} \times$ is a learnable weight matrix and $\mathbf{b} \in$

$\mathbb{R}$ is a bias vector. This projection such that both numerical and categorical features play an equal part in a common representation space, which allows both heterogeneous data to be learned together.

Each of the individual feature embeddings was then concatenated to create a single patient-level feature representation:

$$z = e\_1 \oplus e\_2 \oplus \ldots \oplus e\_n. \quad (8)$$

where $\oplus$ denotes vector concatenation and n is the total number of clinical features. This joint embedding vector, z, is the input of the federated meta-learning model of predicting lung cancer.

The embedding design is intended to reduce sparse one-hot representations and enable joint optimization of heterogeneous variables. Interpretability, however, is not conferred by embeddings alone and would require a separate validated attribution analysis.

D. Federated Data Partitioning

The processed dataset was divided into disjoint client partitions to simulate a decentralized healthcare setting. The manuscript must report the number of clients, partitioning rule, class prevalence per client, sample counts, random seed, and the method used to induce or measure non-IID heterogeneity. Simulated partitions should not be described as actual participating hospitals.

Let the complete dataset be denoted as:

$$D = \cup\_(k=1)^K D\_k, \text{ with } D\_i \cap D\_j = \emptyset \text{ for } i \neq j. \quad (9)$$

Where, $\mathcal{D}$represents the global clinical dataset, Krepresents the total number of federated clients (hospitals), $\mathcal{D}$ represents the local dataset stored at the k-th client, and $\mathcal{D} \cap \mathcal{D} = \emptyset$for $i \neq j$, ensuring no data overlap between clients.

A local model is trained on the private data of each client. The local optimization goal at client kis is formulated as:

$$\theta\_k^* = arg\ min\_(\theta\_k) |D\_k|^(-1) \Sigma\_(x,y \in D\_k) \ell(f\_(\vartheta\_k)(x), y). \quad (10)$$

Where, θ represents the local model parameters at client k, $\mathscr{L}$denotes the local loss function, f is the predictive model,

x represents the patient-specific clinical feature embeddings, $y$represents the treatment response label, and $\ell(\cdot)$denotes the classification loss function.

The clients do not send raw data but after local training, they send model updates to the central server. The expression of global aggregation is formulated as:

$$\theta^(t+1) = \Sigma\_(k=1)^K (|D\_k|/|D|) \theta\_k^(t). \quad (11)$$

where θ^(t+1) is the global parameter vector after communication round t, θ_k^(t) is the locally trained vector from client k, and |D_k|/|D| weights each client by its number of training records.

This design keeps records within their simulated client partitions during training. It reduces raw-data exchange, but it does not by itself establish resistance to gradient leakage, membership inference, or other privacy attacks. Accordingly, the paper uses "data-local" or "privacy-preserving architecture" cautiously and does not claim formal privacy guarantees.

E. Federated Meta-Learning Framework

The federated meta-learning procedure seeks a global initialization that can be adapted to heterogeneous clients using a limited number of local updates. A complete submission must specify the base network, support/query construction, inner-loop steps, inner and outer learning rates, aggregation rule, optimizer, communication rounds, early-stopping criterion, and all random seeds.

During the local adaptation phase, each client updates the shared meta-parameters using its private dataset:

$$\theta'_k = \theta - \alpha \nabla_\theta L_k^{support}(\theta). \quad (12)$$

Where, θ represents the global meta-model parameters, θ denotes the adapted model parameters for client k, $\alpha$is the inner-loop learning rate, and $\mathscr{L}(\theta)$is the local loss function computed using client k's clinical data.

After local adaptation, the meta-update step aggregates the adapted client models to update the global initialization:

$$\theta \leftarrow \theta - \beta \Sigma_{(k=1)}^{K} \nabla_\vartheta L_k^{query}(\vartheta'_k). \quad (13)$$

Where, β denotes the meta-learning rate,

K represents the total number of participating clients, and $\nabla \mathscr{L}(\theta)$captures how well the adapted model performs on client-specific data.

The objective of federated meta-learning is to learn an optimal initialization that minimizes the expected loss across all clients:

$$\min_\theta E_{k\sim p(K)} [L_k^{query}(U_k(\theta))]. \quad (14)$$

Where, $\mathcal{D}$ represents the local dataset at client k,f denotes the client-adapted predictive model, $x$corresponds to patientlevel clinical feature embeddings, y represents the treatment response label, and $\ell(\cdot)$is the classification loss function.

This objective is intended to improve adaptation to unseen client distributions. Whether it improves generalization must be demonstrated on held-out clients or a clearly separated test protocol, with repeated runs and uncertainty estimates.

F. Federated Meta-Learning Framework with Patient-Level Personalization

The architecture combines a shared meta-model with a lightweight prediction-refinement module. The term "patient-specific personalization" is retained provisionally: the final paper must explain what patient-specific support information is available at inference and demonstrate that the outcome label is never used during adaptation.

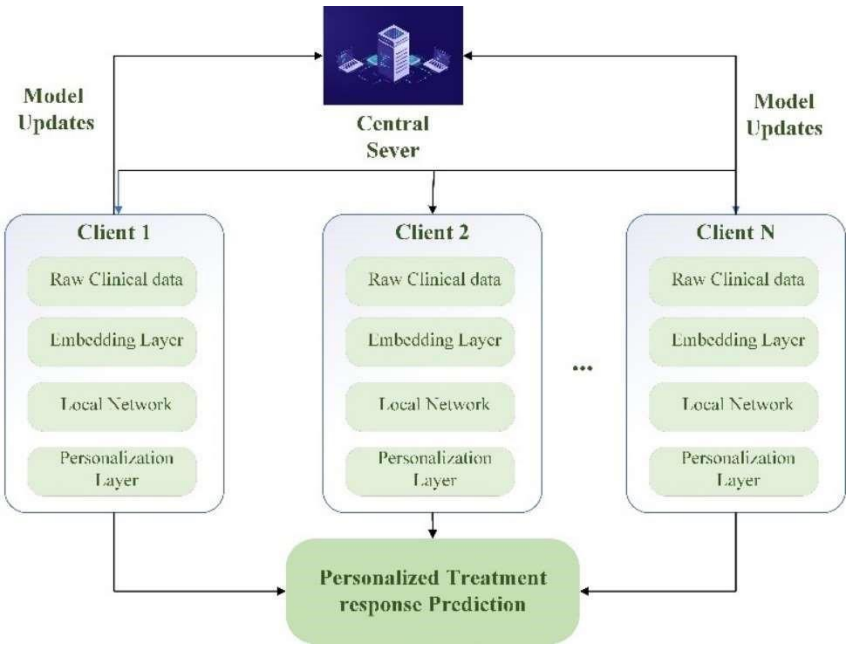

Fig. 2 Federated Meta-Learning Framework

1) Global Meta-Model Learning: In the proposed framework, a global meta-model is learned through iterative collaboration among multiple federated clients, each representing a distinct healthcare institution. During each training round, local models are trained independently on institution-specific clinical data using the shared initialization provided by the meta-model. Only model parameters or gradients are transmitted to the central server, ensuring that sensitive patient data remain localized and privacy-preserving.

The central server aggregates the received updates to refine the global meta-model, which is optimized to capture knowledge that generalizes across diverse clinical settings. Unlike conventional federated learning, the meta-learning objective focuses on learning an initialization that can be rapidly adapted to new clients or patients with minimal additional training. This design enables robustness against non-IID data distributions and institutional variability, which are common challenges in multicenter oncology datasets.

2) Patient-specific personalization layer: The refinement module operates on the learned clinical representation and shared model output. Its inputs, trainable parameters, adaptation data, optimization objective, and inference procedure must be specified. If no repeated labelled observations are available for an individual patient, the mechanism should be described as feature-conditioned prediction rather than patient-level fine-tuning.

3) Treatment-response prediction: The final output is a binary probability of response. The decision threshold must be selected using validation data and then fixed before test evaluation. Report discrimination, calibration, class-sensitive metrics, and confidence intervals rather than accuracy alone.

Algorithm 1. Federated Meta-Learning with Prediction Refinement

Input: Disjoint client datasets {D_k}; communication rounds T; inner-loop rate α; outer-loop rate β.

1:Initialize shared parameters θ and feature-embedding parameters.
2: For each communication round t = 1, …, T, distribute θ to the selected clients.
3: At each client k, split local training data into support and query subsets.
4: Adapt θ on the support subset to obtain θ′_k using the inner-loop update.
5: Evaluate θ′_k on the query subset and transmit only the required update.
6: Aggregate client meta-gradients or parameter updates and update θ using β.
7: Select all hyperparameters and the classification threshold using validation data only.
8: Apply the frozen global model and feature-conditioned refinement module to the test records.

Output: Test-set response probabilities and binary predictions.
The final manuscript must replace the unspecified quantities with the settings used in the archived implementation and must document the support/query construction and privacy safeguards.

IV. Results and Discussion

Fig.5

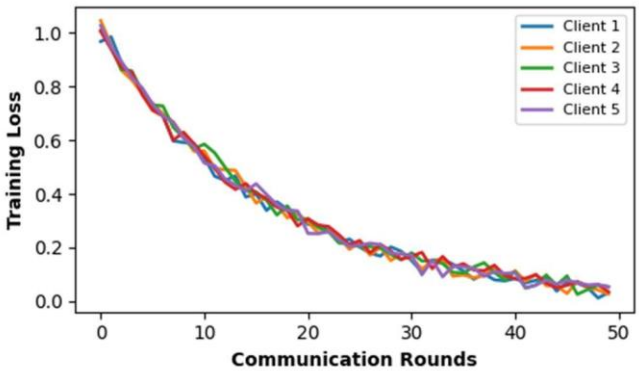

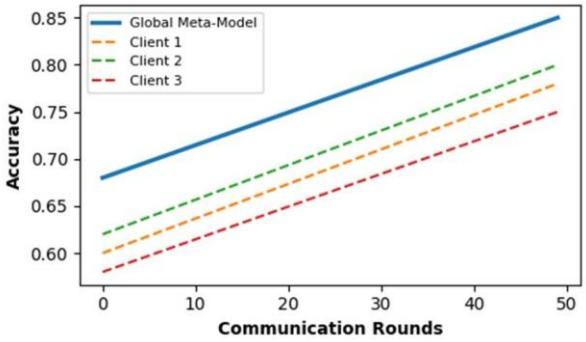

Fig. 5 Global Meta-Model vs Local Models

illustrates the change in the distributions of clinical

EDITORIAL VERIFICATION NOTICE—The numerical results and plots in this section are provisional. Replace them only with values regenerated from the archived experiment outputs before submission.

Fig. 3. Provisional federated training loss across communication rounds.

Figure 3 presents the recorded training-loss trajectories across communication rounds. The plot should be regenerated directly from archived logs, and the reported aggregation setting, number of clients, and stopping rule should accompany it. Decreasing training loss alone does not establish test-set generalization.

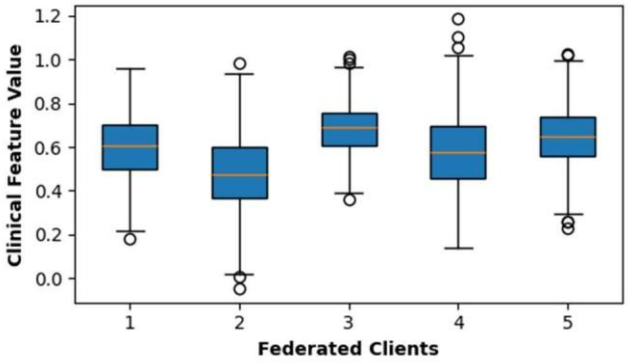

Fig. 4. Provisional client-wise distribution of an unspecified clinical feature.

Figure 4 visualizes between-client variation for one plotted clinical feature. The final manuscript must name the feature, units, sample size per client, and partitioning procedure; one illustrative distribution is insufficient to quantify overall non-IID heterogeneity.

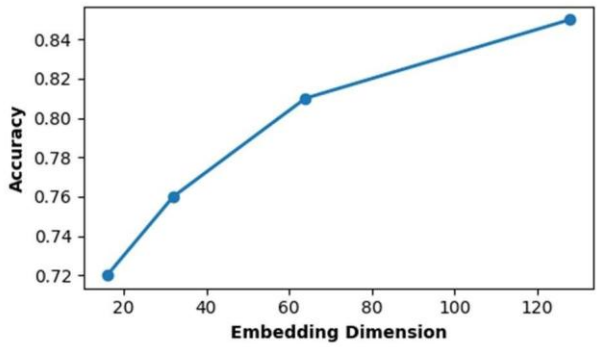

Fig. 6. Provisional validation accuracy by embedding dimension.

Figure 6 shows provisional accuracy values for several embedding dimensions. The final version should state the validation protocol, repeated-run variability, selected dimension, and whether all other hyperparameters were held constant.

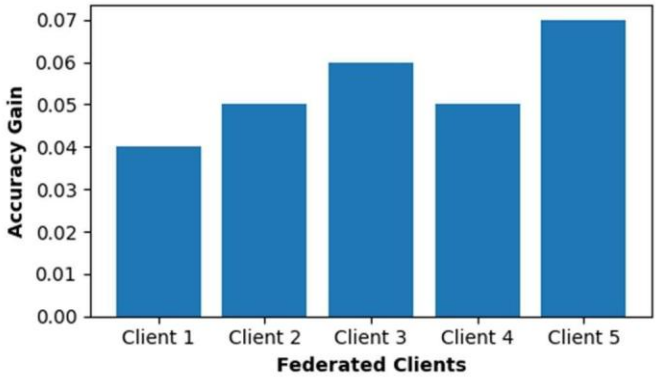

Fig. 7. Provisional client-wise accuracy difference after prediction refinement.

Figure 7 shows provisional client-wise accuracy differences attributed to the refinement mechanism. These values require paired baseline comparisons, sample counts, confidence intervals, and statistical testing before they can support a personalization claim.

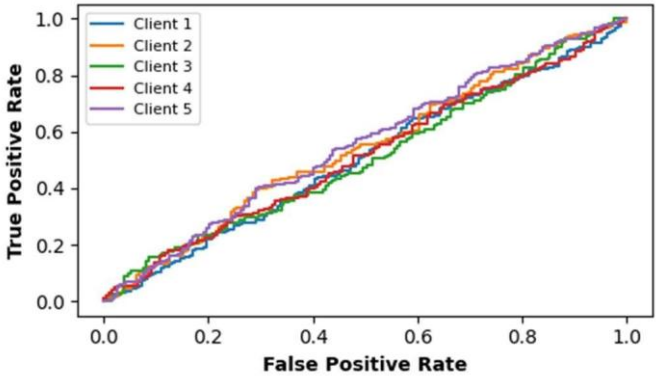

Fig. 8. Provisional client-wise ROC curves requiring regeneration and AUC verification.

Figure 8 is not currently consistent with the manuscript's claim of strong discrimination: the displayed curves remain close to the diagonal. The curves and client-wise AUC values must be regenerated from the saved labels and predicted probabilities. No positive ROC-AUC conclusion should be made until that audit is complete.

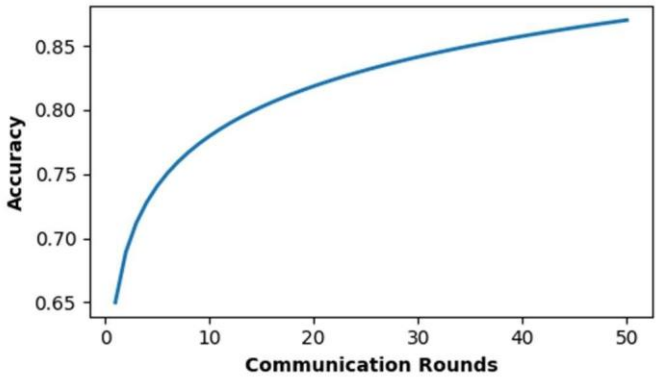

Fig. 9. Provisional accuracy across communication rounds; communication cost is not shown.

Figure 9 shows provisional accuracy as a function of communication rounds. Communication efficiency cannot be inferred from accuracy alone; the final analysis should also report transmitted bytes or parameters, participating-client count, local epochs, runtime conditions, and the comparison protocol.

TABLE II. PROVISIONAL BASELINE COMPARISON—REQUIRES SAME-PROTOCOL VERIFICATION

Method	Status
FedProx	Pending verification
FedAvg	Pending verification
Personalized federated baseline	Pending verification
Proposed federated meta-learning	Pending verification

Table II is retained as a verification placeholder. Every baseline must be implemented on the same dataset, client partitions, preprocessing pipeline, train/validation/test protocol, and random seeds. Values taken from unrelated publications must not be presented as direct experimental comparisons.

4. Discussion

The current figures suggest that the optimization procedure converges and that the proposed components merit further evaluation. They do not yet establish clinical utility, privacy guarantees, or superiority over federated baselines. Final conclusions require a single reconciled test result, client-wise and aggregate metrics, ROC-AUC values calculated from archived probabilities, calibration analysis, repeated runs, confidence intervals, and a leakage audit.

5. CONCLUSION AND FUTURE WORK

This study presents a federated meta-learning architecture for treatment-response prediction from heterogeneous clinical oncology variables. The design combines data-local client training, learnable mixed-type feature representations, meta-learning-based adaptation, and a lightweight prediction-refinement mechanism. Its principal potential is to support collaborative model development without centralizing raw records. However, the present manuscript requires completion of the experimental audit before it can claim improved predictive performance, communication efficiency, patient-level personalization, or clinical applicability.

The final evaluation should reconcile all reported accuracy values, regenerate every figure from archived outputs, compare baselines under one controlled protocol, quantify uncertainty, and test for target leakage. These steps are prerequisites for a defensible journal submission.

Future work may extend the framework to genomic, radiomic, and imaging modalities; evaluate external institutional cohorts; incorporate calibrated uncertainty and explainability analyses; and study formal privacy

mechanisms such as differential privacy or secure aggregation. Prospective clinical interpretation should be considered only after external validation.

REFERENCES

1. Z. Jia, T. Zhou, Z. Yan, J. Hu, and Y. Shi, "Personalized meta-federated learning for IoT-enabled health monitoring," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 43, no. 10, pp. 3157–3170, 2024.
2. K. Kourou, T. P. Exarchos, K. P. Exarchos, M. V. Karamouzis, and D. I. Fotiadis, "Machine learning applications in cancer prognosis and prediction," Computational and Structural Biotechnology Journal, vol. 13, pp. 8–17, 2015, doi: 10.1016/j.csbj.2014.11.005.
3. M. Field et al., "Federated learning survival model and potential radiotherapy decision support impact assessment for non–small cell lung cancer using real-world data," Clinical Oncology, vol. 36, no. 7, pp. e197– e208, 2024.
4. S. N. A. Shah, J. Aalam, and R. Parveen, "Deep Learning Approaches to Enhance Lung Cancer Diagnosis Using Next Generation Sequencing: State of the Art," Arch Computat Methods Eng, Sep. 2025, doi: 10.1007/s11831-025-10357-x.
5. M. Jamalinia and R. Weiskirchen, "Advances in personalized medicine: translating genomic insights into targeted therapies for cancer treatment," Ann Transl Med, vol. 13, no. 2, pp. 18–18, Apr. 2025, doi: 10.21037/atm-25-34.
6. D. Navarro-Garcia et al., "Real-world radiology data for artificial intelligence-driven cancer support systems and biomarker development," ESMO Real World Data and Digital Oncology, vol. 8, p. 100120, Jun. 2025, doi: 10.1016/j.esmorw.2025.100120.
7. A. Vettoruzzo, M.-R. Bouguelia, and T. Rögnvaldsson, "Personalized federated learning with contextual modulation and metalearning," in Proceedings of the 2024 SIAM International Conference on Data Mining (SDM), SIAM, 2024, pp. 842–850.
8. S. Liu et al., "PCRFed: personalized federated learning with contrastive representation for non-independently and identically distributed medical image segmentation," Visual Computing for Industry, Biomedicine, and Art, vol. 8, no. 1, p. 6, 2025.
9. B. Feng et al., "Robustly federated learning model for identifying high-risk patients with postoperative gastric cancer recurrence," Nature Communications, vol. 15, no. 1, p. 742, 2024.
10. Y. Liu, J. Huang, J.-C. Chen, W. Chen, Y. Pan, and J. Qiu, "Predicting treatment response in multicenter non-small cell lung cancer patients based on federated learning," BMC cancer, vol. 24, no. 1, p. 688, 2024.
11. P. Mateus et al., "Multi-Cohort Federated Learning Shows Synergy in Mortality Prediction for MRI-Based and Metabolomics-Based Age Scores," Journal of Healthcare Informatics Research, vol. 9, no. 4, pp. 606–628, 2025.
12. Z. Alsulaimawi, "Meta-FL: A novel meta-learning framework for optimizing heterogeneous model aggregation in federated learning," arXiv preprint arXiv:2406.16035, 2024.
13. W. Pan, Z. Xu, S. Rajendran, and F. Wang, "An adaptive federated learning framework for clinical risk prediction with electronic health records from multiple hospitals," Patterns, vol. 5, no. 1, 2024.
14. V. Gouthamchand et al., "Federated prediction models and external validation for radiotherapy outcomes in oropharyngeal cancer using FAIR clinical and radiomics data," medRxiv, pp. 2025–01, 2025.
15. R. Haripriya, N. Khare, M. Pandey, and S. Biswas, "A privacyenhanced framework for collaborative Big Data analysis in healthcare using adaptive federated learning aggregation," Journal of Big Data, vol. 12, no. 1, p. 113, 2025.
16. K. M. Adnan et al., "Secure and interpretable lung cancer prediction model using mapreduce private blockchain federated learning and XAI," Scientific Reports, vol. 15, no. 1, p. 35693, 2025.
17. "GDC Analysis Center." Accessed: Jan. 09, 2026. [Online]. Available: https://portal.gdc.cancer.gov/analysis_page?app=