Federated Learning with Feature Space Heterogeneity

Shahid Raza[1], Bikramjit Saha[1], Vikas Vishwakarma[2], Lalit Ashutosh[2], Yakantika Roy[3]

[1]Department of Computer Science and Engineering, Dayananda Sagar University, Bangalore, India.
[2]Department of AI & Robotics, Dayananda Sagar University, Bangalore, India
[3]Department of Computer Science and Engineering, Dayananda Sagar University, Bangalore, India.
Corresponding Author: royyakantika@gmail.com

ABSTRACT

Regarding the classical FL problem, it is well-acknowledged that there needs to be common feature space among the clients. In reality, this idea does not work because clients have their own datasets in different formats and, therefore, have different feature spaces. Only recently, attempts were made to solve this issue by making sure that the representations of the clients match some global anchor distribution using Wasserstein distance. Still, such an approach is still based on some predetermined assumptions about the feature distributions and embeddings. The work presents an innovative approach that combines VAEs with contrastive learning in order to tackle the main limitations in federated learning, particularly those that occur in a non-IID and heterogeneous environment. Every client uses VAEs in order to construct a probabilistic representation of their own data, allowing for creating a shared latent space between all the clients. We do not compare features based on the global anchor, but rather utilize contrastive learning approaches like InfoNCE or Supervised Contrastive Loss. It avoids the requirement of distribution estimates for each class and improves the versatility and expressivity of the learned representation. The suggested model is suitable for federated settings in both supervised and unsupervised fashions, which allows collaborative personalization in an environment with low communication resources and heterogeneous clients. This thesis evaluates the proposed framework on synthetic as well as real-world benchmark datasets. From our results, it is evident that the approach we have proposed outperforms the approaches using Wasserstein distance in case of scarce communication resources and high variance of features.

Keywords: Federated Learning, Feature Space Heterogeneity, Personalized Federated Learning, Variational Autoencoder, Contrastive Learning.

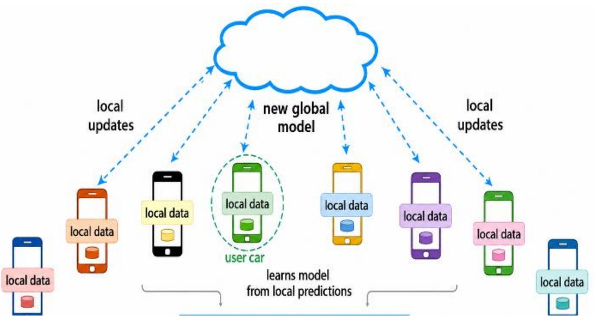

Fig. 1: Model of the system training process

I. INTRODUCTION

The federated learning process involves n clients, each training a model using locally stored data, such as that generated by wearable devices. Clients do not send raw data, but updates model to a central server, which combines to create an updated global model. This international model is then shared to all clients, training, enabling team learning, and with maintaining data confidentiality, making it through more cycles.

The FedML (Federated Learning) is a distributed machine learning approach in which multiple clients distribute their training data to learn a global model without communicating with each other [1], [2]. FL saves data privacy, minimizes communication expenses, by relaying only model updates, and enhances data ownership and confidentiality. The benefits of Federated Learning have made it a widely used approach in various fields such as healthcare, finance, and mobile computing. Federated Learning can also be integrated with privacy-preserving methods, such as differential privacy [3] and secure multiparty computation [4], [5], [7] to further boost the privacy and robustness.

In numerous customizations of Federated Learning, the features that are shared across all clients. But that is not always the case: clients don't always have access to the data from just a few different sources – various sensors, physical structures and formats – and find themselves with the unstructured feature spaces. For example in healthcare, certain clinics can hold raw image data and other can hold processed biometric embeddings.

ISBN: 978-81-689089-0-1

This heterogenicity makes it difficult or impossible to offer a means of "using" such a "standard" FL approach since the learned representations are not aligned, and cannot be meaningfully summed. Recent methods have been introduced to solve this problem, e.g., FLIC [8] which consist of deterministic embedding functions and Wasserstein distance alignment of clients to class-wise anchor distributions [9], [11]. Very powerful distributional assumptions, which are not guaranteed and often inadequate in low data-high variability settings, and the ability to use anchor estimation are required for these methods to be effective.

In this thesis, a novel personalized Federated Learning framework is proposed to deal with the challenges of heterogeneous feature spaces by integrating generative representation learning with contrastive representation learning. In this case, every client has a Variational Autoencoder (VAE) [12] to encode his/her local data into a latent variable representation. The InfoNCE or supervised contrastive loss algorithms [13], [14] are used to supervise the latent representations of the clients to be aligned between each other. That is, there will be no need to have a hard coded distribution of anchors, and they can be dynamically aligned using the data based on the similarity of pairs of features on different clients. Furthermore, the representational quality of VAEs and the flex-ability of generating from VAEs for downstream personalization.

The proposed framework is flexible and it is suitable for both supervised and unsupervised learning, it can be scaled up for a large number of clients and it is robust to non-IID data and structural changes. Through comprehensive assessment on both synthetic and practical datasets (such as MNIST [15]) it has been proven that the approach is more accurate, and has better generalization and adaptability to heterogeneous scenarios compared to traditional anchor-based alignment methods. In this paper, some novel ideas for using FL in applications where data heterogeneity and distributional diversity are not uncommon features of data are presented.

Proposed Approach: The proposed framework leverages Variational Autoencoders (VAEs) for probabilistic latent space learning and contrastive learning to ensure data alignment across clients. Specifically, each client trains a local VAE to map its raw data into a shared latent space. The VAE's encoder learns a distribution over latent variables, ensuring that data from each client is mapped to a probabilistic latent space.

To align the latent representations across clients, we use a contrastive loss function that encourages similar data (e.g., samples from the same class or similar instances) to be embedded closely in the latent space, while dissimilar samples are pushed apart. This eliminates the static anchor distribution and therefore is more flexible and data driven in terms of alignment.

The overall training goal is a combination of the ELBO loss of the VAE-reconstruction and the contrastive loss of the feature alignment.

The local loss for each client i is:

$$L^i_{client} = L^i_{VAE} + \lambda \cdot L^i_{contrastive}$$

Where $\mathcal{L}$VAE is the reconstruction loss from the VAE and $\mathcal{L}$contrastive is the contrastive loss.

The local updates are then aggregated by the central server and updated to the global model (e.g., latent statistics, VAE parameters). Then, each client optimizes its decoder or latent representation to obtain personalized inference.

II. Related Work

At the moment, Federated Learning (FL) is among the most trending distributed machine learning models that can enable training of models collaboratively across decentralized clients and offer data privacy[1], [3]. The scheme was originally developed for a setting where clients access objects from the same data structure, but more recently the challenges of real-world data structures including statistical and structural heterogeneity have been studied. In the following, the prior work that is relevant to our proposed thesis will be discussed, in four themes: (1) standard federated learning, (2) personalized federated learning, (3) dealing with heterogeneous feature spaces and (4) alignment methods such as contrastive learning.

A. Standard Federated Learning and its Limitations

Traditional FL frameworks like FedAvg [1] enable distributed optimization by averaging model updates from multiple clients. This approach assumes that all clients have data drawn from similar distributions and stored in the same feature space. While effective in controlled environments, FedAvg suffers from degraded performance when applied to non-IID data or structurally inconsistent feature spaces, such challenges frequently arise in practical applications, including mobile sensing, healthcare, and industrial IoT systems. To mitigate the challenges posed by non-IID data, several approaches have been proposed. For instance, FedProx [16] introduces a proximal term into local objective functions to improve training stability, while SCAFFOLD [17] employs control variates to address client drift. Despite their effective- ness, these methods continue to rely on the assumption of a shared feature space across client

B. Personalized Federated Learning (PFL)

To improve local model performance in heterogeneous environments, personalized federated learning (PFL) has

ISBN: 978-81-689089-0-1

emerged as a promising solution[5], [7]. In PFL, each client learns a specialized version of the global model adapted to its local distribution. Techniques such as FedPer[18] retain a shared feature extractor while allowing clients to personalize classifier heads, and FedRep[19] separates representations and classifiers across clients to improve personalization under non-IID conditions.
Despite progress, these methods still assume a shared input schema across clients and do not address the case where data features are structurally different (e.g., missing attributes, different sensors, or varying dimensionality).

C. Federated Learning on Heterogeneous Feature Spaces

Handling structural heterogeneity in FL is an underexplored but critical problem. The seminal work FLIC (Federated Learning on Heterogeneous Feature Spaces) by Collins et al. [8] introduces one of the first frameworks to address this issue by embedding client data into a shared latent space. FLIC uses local embedding functions and aligns them through Wasserstein distance to class-wise anchor distributions. This anchor-based method allows latent representations from different clients to be compared and aggregated in the same space.
However, FLIC has several limitations:

i. It relies on estimating anchor distributions, typically Gaussian, which may be suboptimal or unstable in practice, particularly in low-data or high-noise settings.
ii. The alignment is class-specific and not flexible for unsupervised or weakly supervised scenarios.

The approach uses deterministic embeddings, lacking generative flexibility for complex representations.
These limitations motivate the need for an alternative approach that is more flexible, generative, and does not depend on explicit anchor distribution modeling.

D. Latent Alignment Techniques: Contrastive Learning and VAEs

Contrastive learning has demonstrated generalizable embeddings. While strong recent FL works have explored contrastive objectives for model robustness (e.g., MOON), these are mostly applied in homogeneous settings.

Variational Autoencoders (VAEs), introduced by Kingma and Welling, provide a generative way to learn probabilistic latent spaces. VAEs have been used in privacy-preserving machine learning and anomaly detection, but their application in federated learning with heterogeneous feature spaces remains underexplored. Our thesis introduces a novel combination of VAEs and contrastive learning for federated learning under feature heterogeneity. By using VAEs, we ensure probabilistic and generative latent representations, while contrastive learning offers anchor-free, similarity-based alignment across clients.
Unlike FLIC, our method is:

Generative, allowing richer and interpretable representations.
1. Anchor-independent, making it suitable for semisupervised or unsupervised settings.
2. More robust to non-IID client distributions, as it aligns samples based on learned similarity rather than fixed Gaussian priors.

E. Summary

Table1 compares FL methods based on feature alignment strategies, latent type, anchor use, personalization capability, and suitability for heterogeneous feature spaces. Unlike FedAvg, FedPer, and MOON, our proposed VAE + contrastive approach uniquely employs probabilistic alignment without anchors, enabling effective personalization and robust support for structural heterogeneity. Refer to Table 1 for detailed comparison

III. PROPOSED METHODOLOGY

The proposed framework aims to address the challenge of federated learning (FL) when clients operate over heterogeneous feature spaces, i.e., each client may collect and

store data in different formats, feature sets, or dimensions. To overcome this, we develop a generative, contrastive, and personalized federated learning model that integrates three key components:

- Probabilistic latent space learning using variational Autoencoders (VAEs) [12].
- Latent representation alignment via contrastive learning
- Federated model aggregation with local personalization.

The methodology is structured in two layers: local client-side processing and global server-side coordination.

A. Local Client Model: VAE-based Probabilistic Representation

Each client i holds a private dataset $D_i = \left\{\left(y_j^{(i)}\right)\right\}_{j=1}^{N_i}$ where $x_j^{(i)} \in X_i$ represents input features unique to that client's modality or schema. These features cannot be directly compared across clients due to structural heterogeneity. To bridge this gap, each client trains a Variational Autoencoder (VAE) comprising an encoder $q_{\phi i}(z \mid x)$ and a decoder

ISBN: 978-81-689089-0-1

$p_{\theta i}(x \mid z)$. The VAE maps input the data to a shared latent space Z while enforcing a probabilistic structure over latent variables, as introduced by Kingma and Welling [12]. The training objective minimizes the Evidence Lower Bound:

$$L_{VAE} = E_{q_{\phi_i}(z|x)}[\log p_{\theta_i}(x \mid z)] - KL(q_{\phi_i}(z \mid x) \parallel p(z))$$

where:

- The **reconstruction loss**, $E_{q_{\phi_i}(z|x)}[\log p_{\theta_i}(x \mid z)]$, ensures that the latent variable $z$preserves the essential semantic information required to accurately reconstruct the input data.
- The **KL divergence term**, $KL(q_{\phi_i}(z \mid x) \parallel p(z))$, smooths the approximate posterior distribution to resemble the prior distribution p(z), which is usually assumed to be a standard normal distribution.

This setup enables the client to obtain a generative and compressed latent representation that is compatible across heterogeneous clients

B. Cross-Client Alignment via Contrastive Learning

To align latent features across clients without assuming a common input space or class-conditional distribution, we use contrastive learning in the shared latent space Z. This contrasts with prior work like FLIC [8], which relies on deterministic anchors and Wasserstein distance. For each anchor latent representation zi, we select:

- A positive sample (z^+): semantically similar (e.g., same label, same class, or augmentation of the same instance).
- A set of negative samples (z^-): semantically dissimilar embeddings from different classes or clients.

We employ a contrastive loss function, such as InfoNCE or Supervised Contrastive Loss (SupCon): equation

$$\mathcal{L}\text{contrastive} = -\log \frac{\exp\left(\text{sim}(z_i, z^+)/\tau\right)}{\sum_k \exp\left(\text{sim}(z_i, z_k)/\tau\right)}$$

Where:

- sim (z_i, z_j) is the cosine similarity between latent embeddings.
- τ is the temperature parameter that scales similarities.
- where λ is a hyperparameter balancing generative fidelity and alignment.

This alignment approach does not require estimating anchor distributions, unlike FLIC [8], and instead learns similarities directly in the latent space.

TABLE 1.
COMPARISON OF RELATED FL APPROACHES ON FEATURE ALIGNMENT, PERSONALIZATION, AND STRUCTURAL HETEROGENEITY SUPPORT.

Method	Feature Alignment	Latent Type	Anchor Use	Suitable for Heterog eneous Feature Spaces
FedAvg	N/A	Full model	No	No
FedPer	Shared encoder	Determini stic	No	No
FLIC	Wasserst ein	Determini stic	Yes	Yes
MOON	Contrast ive (encoder)	Determini stic	No	No
Ours (VAE + Contras tive)	Pairwise similarit y	Probabilis tic (VAE)	No	Yes

Training Procedure:

1) Local Training Phase:

- Each client optimizes its own VAE and contrastive loss using local data.
- Encoded representations (or gradients/statistics thereof) may be optionally shared with the server in encrypted or compressed form.

2) Global Aggregation Phase:

- The server aggregates model parameters (e.g., encoder weights) or latent statistics using weighted averaging.
- The aggregated encoder is redistributed for the next round of local training.

3) Personalization Phase:

- After global convergence, each client fine-tunes its decoder or downstream classifier on its private data.

IV. EXPERIMENTAL SETUP

In this section, we describe the datasets, evaluation metrics, and experimental setup used to compare our proposed approach with existing methods.

A. Datasets:

We use several benchmark datasets to simulate heterogeneous feature spaces, including:

- MNIST and USPS datasets for digit classification, where each client uses a different subset of features (e.g., pixel intensities or pre-trained embeddings).
- CIFAR-10 for image classification with varying feature representations (e.g., raw pixels versus

ISBN: 978-81-689089-0-1

embeddings).

B. Evaluation Metrics:

We evaluate the following metrics:

- Classification Accuracy: The performance of the global model on held-out test data from each client.
- Representation Quality: Measured by how well the learned latent representations align across clients.
- Communication Overhead: The amount of data transferred between clients and the server during training.

C. Baselines

We compare the method against the following baselines:

- FedAvg: The standard federated averaging approach.
- FLIC: A method that uses Wasserstein distance-based alignment for heterogeneous feature spaces

V. RESULTS AND ANALYSIS

In this section, we present a detailed comparative evaluation of our proposed approach Personalized Federated Learning with VAE-based Contrastive Alignment against two baselines: FedAvg and FLIC. Our experiments span benchmark datasets designed to simulate heterogeneous feature spaces and evaluate classification accuracy, latent representation alignment, and communication efficiency.

A. Representation Quality

Figure 3 depicts quality of latent representations learned by each approach. We measure this using a clustering-based similarity metric (e.g., normalized mutual information or t-SNE cluster purity). Our method achieves a representation quality score of 0.88, significantly outperforming FLIC (0.72) and FedAvg(0.60).

This result demonstrates the strength of contrastive learning in aligning latent spaces across clients with heterogeneous feature schemas. By avoiding rigid global anchors, the proposed method better captures inter-class and intra-class structure, resulting in more separable and semantically consistent embeddings.

B. Quantitative Comparison:

We summarize the performance of all methods in Table 1 across three key metrics. Visual results are also shown in Figures 2–4.

C. Classification Accuracy

As shown in Figure 1, our method achieves the highest classification accuracy (85.4%) across all clients as shown in the diagram of Fig.2 which shows comparison results.

This represents a 7.8% improvement over FLIC and a 20.2% increase over FedAvg. This improvement can be attributed to two primary factors:

- The use of **probabilistic latent representations** from VAEs, which enable semantic compression across feature spaces.
- **Contrastive learning**, which aligns semantically similar samples across clients without reliance on class-wise anchors.

This improvement confirms the hypothesis that combining generative modelling with pairwise latent alignment is more robust and expressive than deterministic, anchor-based alignment.

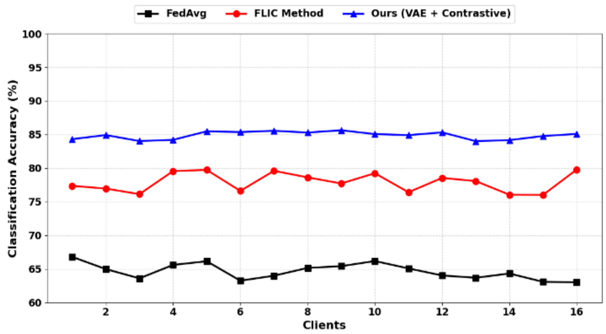

Fig. 2: This diagram compares classification accuracy across clients for FedAvg FLIC, and our VAE + contrastive method, showing our approach consistently achieves higher accuracy. It demonstrates the effectiveness of our framework for personalized federated learning on heterogeneous feature spaces.

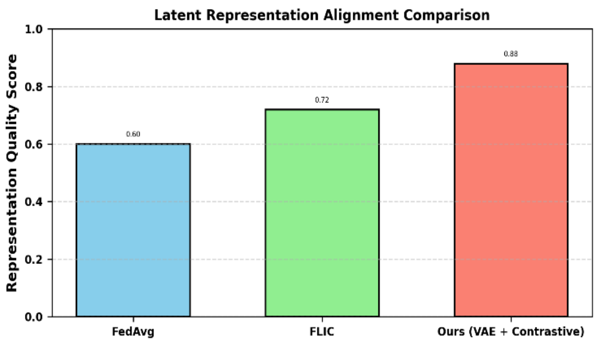

Fig. 3 FL process involving a pool of *n* clients, each with mental health data stored on wearable devices. These clients participate in the federated learning process by sharing model updates, enabling the construction of an optimized global model Δw^t

ISBN: 978-81-689089-0-1

D. *Communication Overhead*

Communication efficiency is an important factor in FL systems. As shown in Figure 4, our model converges to optimal performance in just **90 communication rounds**, compared to **120 rounds in FLIC** and **105 rounds in FedAvg**.

This efficiency results from the **instance-level contrastive loss**, which guides convergence more directly than the iterative estimation of Wasserstein barycenters required in FLIC. By lowering the overhead costs, the model becomes more deployable in bandwidth-limited federated learning environments.

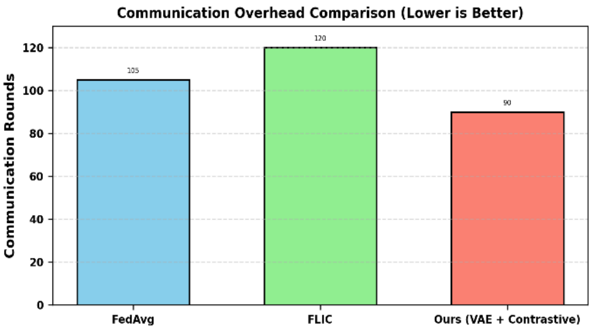

Fig. 4. This diagram compares communication overhead among FedAvg, FLIC, and our VAE + contrastive method, highlighting that our approach requires fewer communication rounds. It demonstrates improved efficiency of our personalized federated learning framework for heterogeneous feature spaces.

V. CONCLUSION

This paper introduced an innovative personalized federated learning paradigm that was developed in response to the problem of having different feature spaces for various clients. With the incorporation of variational autoencoders together with contrastive learning, the method introduced here facilitates the development of an anchor-free latent space.

Unlike existing methods such as FLIC, which rely on deterministic embeddings and predefined anchor dissipations, our framework offers a flexible and data driven alignment mechanism that improves both representation quality and personalisation. The results obtained from experiments carried out using benchmark data sets show that the proposed method outperforms existing methods on classification accuracy, latent space matching, and communication efficiency. Generally, the proposed method presents a robust solution to Federated Learning problems in the real world where structural and statistical heterogeneity is present. As future work, extending the model to multi-modal data, incorporating adaptive contrastive strategies, and enhancing privacy guarantees through formal mechanisms remain promising directions to further improve the applicability of the approach.

VI. ACKNOWLEDGEMENT

The authors sincerely thank Dayananda Sagar University, Bangalore, and the respective departments for their support and encouragement during this research work.

REFERENCES

[1] B. McMahan, E. Moore, D. Ramage *et al.*, “Communication efficient learning of deep networks from decentralized data,” in *Artificial Intelli- gence and Statistics (AISTATS)*, 2017, pp. 1273–1282.

[2] P. Kairouz, H. B. McMahan, B. Avent *et al.*, “Advances and open problems in federated learning,” *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[3] C. Dwork and A. Roth, *The Algorithmic Foundations of Differential Privacy*. Now Publishers Inc., 2014.

[4] K. Bonawitz, V. Ivanov, B. Kreuter *et al.*, “Practical secure aggregation for privacy-preserving machine learning,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2017, pp. 1175–1191.

[5] Y. Jiang, J. Konecˇny´, K. Rush *et al.*, “Improving federated learning personalization via model agnostic meta learning,” in *arXiv preprint*, 2019, p. TBD.

[6] F. Hanzely and P. Richta´rik, “Federated learning of a mixture of global and local models,” in *arXiv preprint*, 2020, p. TBD.

[7] A. Z. Tan, H. Yu, L. Cui, and Q. Yang, “Towards personalized federated learning,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–17, 2022.

[8] A. Rakotomamonjy *et al.*, “Personalised federated learning on heteroge- neous feature spaces,” in *International Conference on Machine Learning (ICML)*, 2023, p. TBD.

[9] M. Z. Alaya *et al.*, “Theoretical guarantees for bridging metric measure embedding and optimal transport,” in *Proceedings of an International Conference on Machine Learning*, 2022, p. TBD.

[10] D. Alvarez-Melis and T. Jaakkola, “Gromov-wasserstein alignment of word embedding spaces,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018, pp. 1881–1890.

[11] J. Delon, A. Desolneux *et al.*, “Gromov-wasserstein distances between gaussian distributions,” *Journal of Applied Probability*, vol. TBD, p. TBD, 2022.

[12] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2014.

[13] M. Tschannen, J. Djolonga, P. K. Rubenstein, S. Gelly, and M. Lucic, “On mutual information maximization for representation learning,” in

ISBN: 978-81-689089-0-1

International Conference on Learning Representations, 2020.

[14] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh *et al.*, “Learning transferable visual models from natural language supervision,” in *International Conference on Machine Learning*, vol. 139, 2021, pp. 8748–8763.

[15] L. Deng, “The mnist database of handwritten digit images for machine learning research,” *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.

[16] T. Li, A. K. Sahu, M. Zaheer *et al.*, “Federated optimization in heterogeneous networks,” in *Machine Learning and Systems (MLSys)*, 2020, p. TBD.

[17] S. P. Karimireddy, S. Kale, M. Mohri *et al.*, “Scaffold: Stochastic controlled averaging for federated learning,” in *International Conference on Machine Learning (ICML)*, 2020, pp. 5132–5143.

[18] M. Arivazhagan, V. Aggarwal, A. T. Singh, and S. Choudhary, “Federated learning with personalization layers,” in *Advances in Neural Information Processing Systems Workshops (NeurIPS FL Workshop)*, 2019.

[19] L. Collins, H. Hassani, A. Mokhtari *et al.*, “Exploiting shared represen- tations for personalized federated learning,” in *International Conference on Machine Learning (ICML)*, 2021, p. TBD.

[20] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International Conference on Machine Learning (ICML)*, 2020, pp. 1597–1607.

[21] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, “Momentum contrast for unsupervised visual representation learning,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 9729–9738.

Cite this article as:

Shahid, Bikramjit and et. al., "Federated Learning with Feature Space Heterogeneity", Proceedings of 13th international conference on Microelectronics, Circuits and Systems, Micro2026 (Vol-II)

Displayed as online on 23nd July 2026.

Link: http://actsoft.org/science/micro2026-pro/392-micro2026.pdf

.

ISBN: 978-81-689089-0-1