UniVarFL: Uniformity and Variance Regularized Federated Learning for Heterogeneous Data (Student Abstract)

Sunny Gupta, Nikita Jangid, Amit Sethi

Indian Institute of Technology Bombay, Mumbai, India
(sunnygupta, nikitajangid, asethi)@iitb.ac.in

Abstract

Federated Learning (FL) often suffers from severe performance degradation when faced with non-IID data, largely due to local classifier bias. Traditional remedies such as global model regularization or layer freezing either incur high computational costs or struggle to adapt to feature shifts. In this work, we propose **UniVarFL**, a novel FL framework that emulates IID-like training dynamics directly at the client level, eliminating the need for global model dependency. UniVarFL leverages two complementary regularization strategies during local training: *Classifier Variance Regularization*, which aligns class-wise probability distributions with those expected under IID conditions, effectively mitigating local classifier bias; and *Hyperspherical Uniformity Regularization*, which encourages a uniform distribution of feature representations across the hypersphere, thereby enhancing the model's ability to generalize under diverse data distributions. Extensive experiments on multiple benchmark datasets demonstrate that UniVarFL outperforms existing methods in accuracy, highlighting its potential as a highly scalable and efficient solution for real-world FL deployments, especially in resource-constrained settings.

Code — https://github.com/sunnyinAI/UniVarFL

Introduction

Federated Learning (FL) (McMahan et al. 2017) trains deep neural networks across decentralized devices without central data collection, often matching centralized performance (Verbraeken et al. 2020), but non-IID data induces *client drift*, causing local models to deviate from the global objective.

To mitigate client drift, local models are often regularized against the global model (Acar et al. 2021), while recent studies highlight classifier bias under non-IID data. Techniques such as augmentation (Yang, Zhu, and Chen 2020) or freezing classifiers for orthogonality (Oh, Kim, and Yun 2021) address this bias but incur overhead—FedProx (Li et al. 2020), SCAFFOLD (Karimireddy et al. 2020), and FedDyn (Acar et al. 2021) require layer-wise comparisons, whereas MOON (Li, He, and Song 2021) and FedAlign (Mendieta et al. 2022) demand extra forward passes. Furthermore, many methods are tested mainly under label-shift (Dirichlet-simulated) scenarios, limiting their effectiveness for feature shifts.

To address these issues, we propose **UniVarFL**, a novel framework that emulates IID-like training dynamics using two complementary regularization strategies.

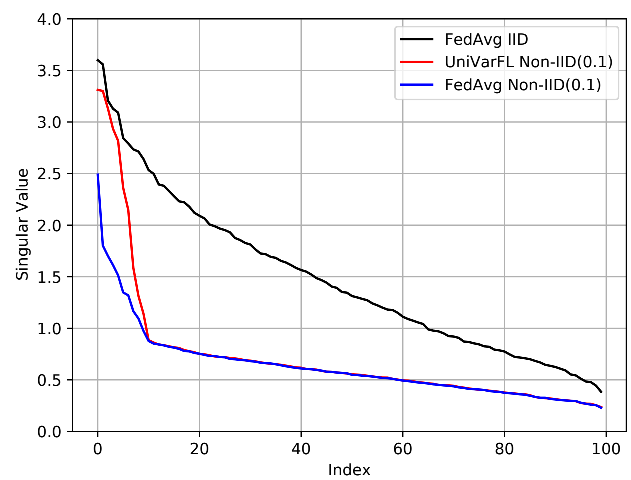

Figure 1: **Singular values of the classifier's final layer.** Plot computed from a model trained on CIFAR-100. Axis labels and tick labels are rendered at a minimum of 8 pt to ensure readability in grayscale printing.

Proposed Method: UniVarFL

UniVarFL addresses two major challenges in Federated Learning (FL) with non-IID data: (i) *classifier variance degeneration*, where classifiers bias toward locally dominant classes, and (ii) *feature imbalance*, where encoder outputs fail to maintain a uniform representation space. By jointly regularizing classifier predictions and feature embeddings, UniVarFL improves generalization and robustness.

Classifier Variance Regularization. In non-IID FL, classifier weights often collapse toward local optima, reflected by rapid singular value decay. To counter this, we introduce a

Method	STL-10 α=0.01/1.0	CIFAR-100 α=0.01/1.0	PACS	HAM10000
FedAvg	28.1±1.3 / 69.1±0.6	51.7±1.9 / 58.3±1.5	62.5±0.4	72.8±0.1
FedProx	26.9±0.7 / 67.6±0.8	50.3±0.4 / 59.2±0.9	58.9±0.6	73.5±0.3
MOON	26.0±0.8 / 69.4±0.4	47.9±1.4 / 60.0±1.0	63.7±1.5	72.9±0.2
Freeze	23.8±1.0 / **72.0±0.1**	51.5±1.4 / 59.3±1.9	61.8±0.1	73.1±0.6
UniVarFL	**28.5±1.4** / 59.9±0.6	**56.9±1.0 / 60.5±0.9**	**66.0±0.9**	**74.6±0.5**

Table 1: Test accuracy (%) under Dirichlet data heterogeneity with $\alpha \in \{0.01, 1.0\}$.

Method	STL-10 ($\alpha = 0.01$)		
	$p = 0.1$	$p = 0.5$	$p = 1.0$
FedAvg	17.6±3.3	21.6±1.7	27.4±1.7
FedProx	19.4±2.8	21.9±1.5	26.4±1.6
MOON	20.6±3.5	23.4±1.8	25.9±1.4
Freeze	21.7±3.5	23.6±1.7	23.9±1.4
UniVarFL	**23.4±3.1**	**27.9±1.1**	**28.5±1.4**

Table 2: Test accuracy (%) on STL-10 under different ρ

variance-preserving regularizer:

$$L_V(f_\theta(X)) = \frac{1}{D}\sum_{j=1}^{D} \max(0,\, c - \mathrm{Var}_{x_i \sim X}(f_\theta(x_i)_j)), \quad (1)$$

where D is the number of classes and Var is the variance of class-wise predicted probabilities. The threshold c is derived from one-hot IID labels, anchoring local predictions toward balanced class separability. This prevents overconfidence on locally frequent classes and reduces client drift.

Hyperspherical Energy Regularization. Preserving variance alone does not prevent feature collapse under domain or style shifts. To enforce diversity, we introduce a hyperspherical energy term:

$$L_{HE}(g_\theta(X)) = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n}\left(1 - z_i^\top z_j + \epsilon\right)^{-1}, \quad (2)$$

where $z_i = g_\theta(x_i)$ are normalized encoder outputs and $z_i^\top z_j$ measures cosine similarity. By penalizing highly similar embeddings, this spreads features uniformly across the hypersphere, mitigating redundancy and improving generalization across heterogeneous clients. Unlike contrastive losses, this formulation is fully local and does not require negative sampling, making it lightweight for FL.

Unified Objective. UniVarFL combines both regularizers with cross-entropy loss:

$$L = L_{CE} + \mu L_{HE} + \lambda L_V, \quad (3)$$

where μ, λ weight the auxiliary terms. During each round, clients minimize this loss locally and send updates to the server for aggregation. By constraining both classifier outputs and feature representations, UniVarFL drives local training toward IID-like dynamics. This dual regularization enhances accuracy, stabilizes convergence, and remains computationally efficient, proving effective across both label-shift and feature-shift scenarios in real-world federated learning.

In real-world FL, only a fraction of clients typically join each round; we model this via participation ratio ρ (e.g., $\rho = 0.1$ means 10% clients per round) and evaluate STL-10 over 100 aggregation rounds, with results in Table 2.

Conclusion

In this paper, we introduced **UniVarFL**, a new framework for mitigating non-IID challenges in Federated Learning by directly emulating IID-like training dynamics at the local level. Unlike prior approaches that rely on explicit global models or freeze specific layers, UniVarFL promotes class-wise probability variance through *classifier variance regularization* and encourages a more balanced feature space via *hyperspherical uniformity*. By jointly addressing classifier bias and feature imbalance, our method not only achieves state-of-the-art performance under both label-shift and feature-shift conditions but also maintains low computational overhead.

References

Acar, D. A. E.; Zhao, Y.; Navarro, R. M.; Mattina, M.; Whatmough, P. N.; and Saligrama, V. 2021. Federated learning based on dynamic regularization. *arXiv preprint arXiv:2111.04263*.

Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, 5132–5143. PMLR.

Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10713–10722.

Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2: 429–450.

McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.

Mendieta, M.; Yang, T.; Wang, P.; Lee, M.; Ding, Z.; and Chen, C. 2022. Local learning matters: Rethinking data heterogeneity in federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8397–8406.

Oh, J.; Kim, S.; and Yun, S.-Y. 2021. Fedbabu: Towards enhanced representation for federated image classification. *arXiv preprint arXiv:2106.06042*.

Verbraeken, J.; Wolting, M.; Katzy, J.; Kloppenburg, J.; Verbelen, T.; and Rellermeyer, J. S. 2020. A survey on distributed machine learning. *Acm computing surveys (csur)*, 53(2): 1–33.

Yang, T.; Zhu, S.; and Chen, C. 2020. Gradaug: A new regularization method for deep neural networks. *Advances in neural information processing systems*, 33: 14207–14218.