# Sheaf HyperNetworks for Personalized Federated Learning

Bao Nguyen[1] Lorenzo Sani[1], Xinchi Qiu[1], Pietro Liò[1] & Nicholas D. Lane[1]
[1] Department of Computer Science and Technology, University of Cambridge

Traditional Federated learning (FL) solutions aim to produce a global model that encapsulates local data across all clients. However, an ongoing challenge is to tackle the statistical heterogeneity among clients, which degrades the performance of the global model. Personalized Federated learning (PFL) offers an alternative training strategy by delivering unique parameters based on the local client data distribution while incorporating effective parameter sharing across clients.

HyperNetworks (HNs) [1], a class of neural networks that generates parameters for a target network, have emerged as an effective solution for Personalized Federated Learning (PFL) [2]. Through federated optimisation, HNs enable parameter sharing to produce unique and high-performing model parameters for each client. Following this, HNs have been equipped with graph neural networks (GNNs) [3, 4] to leverage the relational topology among clients. This enables the delineation of clients with similar local data distributions, providing insights on the overall degree of statistical heterogeneity and enhancing parameter sharing across comparable clients. This is also important in cross-device settings such as recommender systems, health monitoring, and traffic forecasting, where strong relationships exist between clients that can be used to improve personalization. Although coupling a GNN and HN is effective, this leads to other problems that limit these models' effectiveness in applying PFL. Problems such as over-smoothing and heterophily, where the former results from increased message-passing layers, lead to node embeddings converging to the same representation [5]. As a result, there is a significant loss of signal, which ultimately incurs performance issues for the task objective.

A new area of research, Sheaves in GNNs [6, 7] are objects that equip a graph with abstract algebraic spaces, called stalks, over its nodes and edges and functions, called restriction maps, between them. They have been shown to mitigate ongoing problems with GNNs, such as over-smoothing and heterophily. Thus, to address existing limitations, we introduce Sheaf HyperNetworks (SHN), which combines Cellular Sheaf theory [8] with HNs to improve existing Graph HyperNetwork solutions. Our contributions are threefold: (1) First, we present an extensive analysis of HNs' capacity to discern the local client data distributions and present how a client relation graph can be naturally constructed. This could help in understanding underlying relations between clients in cross-device settings, in which the topology cannot be obtained without infringing on privacy. (2) Second, we propose a Sheaf HyperNetwork model to enhance PFL, allowing deeper hidden client representations to be learnt and topological relationships to be drawn from the client space. We evaluate across many typical cross-device FL tasks such as multi-class classification, traffic, and weather forecasting. Our model outperforms competitive baselines across multiple benchmarks, achieving up to 6.2% higher accuracy and 5.3% lower mean squared error. Furthermore, we show if existing methods adopt the same constructed client relation graph, as part of our methodology, SHN still yields a 2.7% higher accuracy. (3) Lastly we present an extensive ablation study on SHN's capacity to mitigate over-smoothing and densely connected graphs.

## References

[1] David Ha, Andrew Dai, and Quoc V. Le. *HyperNetworks.* 2016. arXiv: 1609.09106 [cs.LG].

[2] Aviv Shamsian et al. *Personalized Federated Learning using Hypernetworks.* 2021. arXiv: 2103.04628 [cs.LG].

[3] Wanyu Lin et al. *Graph-Relational Federated Learning: Enhanced Personalization and Robustness.* 2024.

[4] Zhengyi Xu, Liu Yang, and Shiqiao Gu. "Heterogeneous Federated Learning Based on Graph Hypernetwork". In: *Artificial Neural Networks and Machine Learning – ICANN 2023.* Ed. by Lazaros Iliadis et al. Cham: Springer Nature Switzerland, 2023, pp. 464–476. ISBN: 978-3-031-44213-1.

[5] Chen Cai and Yusu Wang. *A Note on Over-Smoothing for Graph Neural Networks.* arXiv:2006.13318 [cs, stat]. 2020.

[6] Jakob Hansen and Thomas Gebhart. *Sheaf Neural Networks.* 2020. arXiv: 2012.06333 [cs.LG].

[7] Cristian Bodnar et al. *Neural Sheaf Diffusion: A Topological Perspective on Heterophily and Oversmoothing in GNNs.* 2023. arXiv: 2202.04579 [cs.LG].

[8] Jakob Hansen and Robert Ghrist. "Toward a spectral theory of cellular sheaves". In: *Journal of Applied and Computational Topology* 3.4 (2019), pp. 315–358. ISSN: 2367-1734.