Diploma Thesis

Microproccessors and Digital Systems Laboratory

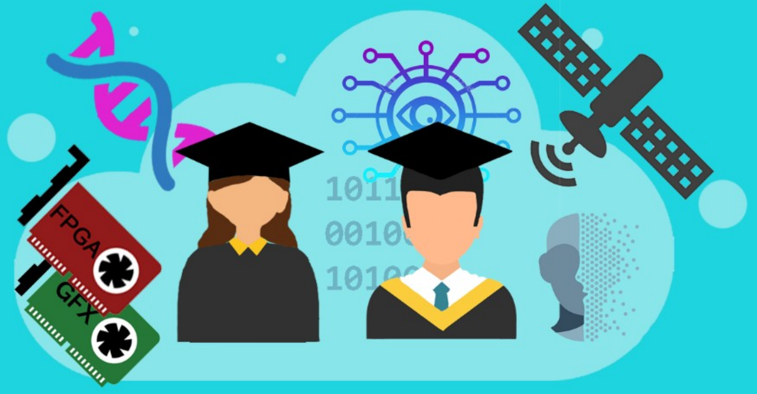

Optimizing Federated Learning Systems through Approximation Techniques

Federated Learning (FL) has emerged as a powerful decentralized machine learning paradigm, allowing multiple clients (e.g., edge devices or remote servers) to collaboratively train models without sharing raw data. This preserves data privacy while enabling large-scale distributed learning. However, Federated Learning introduces several challenges, such as high communication costs, limited computational resources on edge devices, and the need for frequent synchronization between clients and the central server. These issues become even more pronounced in real-world deployments, where network bandwidth and device energy consumption are critical constraints. To address these challenges, Approximate Federated Learning has been proposed, where trade-offs between accuracy, communication efficiency, and energy consumption are explored to make FL more practical and scalable.

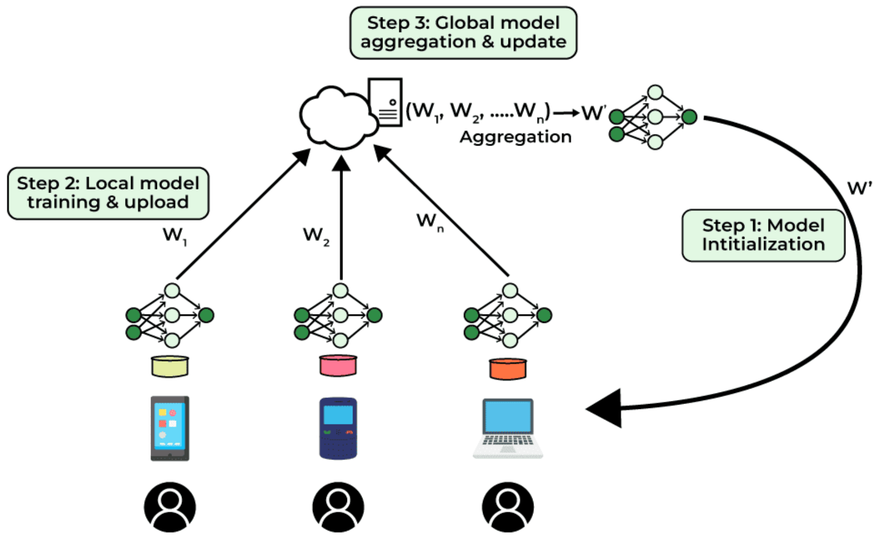

Figure 1: The concept of Federated Learning training

(https://www.geeksforgeeks.org/collaborative-learning-federated-learning/)

This series of theses will explore a broad range of topics aimed at optimizing Approximate Federated Learning systems. The focus includes (but is not limited to): 1) algorithmic improvements, such as developing approximate model updates and compression techniques that reduce communication overhead while maintaining model accuracy; 2) communication protocol advancements designed to minimize synchronization delays, reduce bandwidth usage, and lower energy consumption in distributed environments; 3) hardware/software (HW/SW) co-design solutions that align the computational capabilities of edge devices with the specific demands of Federated Learning. This may involve creating energy-efficient training schemes and optimizing memory usage for on-device training. By addressing these key challenges, the research aims to push the limits of Federated Learning, making it more robust, scalable, and efficient for real-world applications.

RELATED MATERIAL:

[1] https://research.ibm.com/blog/what-is-federated-learning

[2] Qiu, X., Fernandez-Marques, J., Gusmao, P. P., Gao, Y., Parcollet, T., & Lane, N. D. (2022). Zerofl: Efficient on-device training for federated learning with local sparsity. arXiv preprint arXiv:2208.02507.

[3] Mei, Y., Guo, P., Zhou, M., & Patel, V. (2022). Resource-adaptive federated learning with all-in-one neural composition. Advances in Neural Information Processing Systems, 35, 4270-4284.

CONTACT INFORMATION:

Dimosthenis Masouros, Ph.D.: (dmasouros@microlab.ntua.gr)

George Retsinas, Ph.D.

Prof. Dimitrios Soudris: (dsoudris@microlab.ntua.gr)