Gossip Learning in Decentralised Clinical Trials: Towards Mobile Parkinson's Speech Biomarkers

Speech-based biomarkers offer a non-invasive and mobile-compatible pathway for continuous Parkinson's disease (PD) monitoring. Yet, training robust models requires diverse patient data distributed across clinical sites such as in the case of the Parkinson's Progression Markers Initiative (PPMI), a longitudinal observational multi-centre natural history study undertaken across more than 50 countries globally. Such data is subject to strict privacy regulations and data sovereignty requirements that prevent it from being centralised. While Federated Learning (FL) enables collaborative model training without raw data exchange, conventional FL frameworks such as Flower retain a central aggregation server. This introduces a trust assumption and a single point of failure at odds with fully decentralised personal data infrastructures.

To address this, we investigate a fully decentralised approach to PD speech biomarker learning across distributed clinical data silos that does not depend on the presence of a centralised point of coordination. We apply Glow, a gossip learning framework built on Flower that removes the central aggregation authority and enables peer-to-peer model exchange across configurable network topologies. Treating each clinical site as an independent node, we use this framework to train classifiers on PD speech datasets previously used for cross-language biomarker detection. We compare three training regimes: (i) a centralised clustering baseline with pooled data, (ii) standard server-mediated federated learning, and (iii) fully decentralised gossip learning via Glow.

Our aim is to characterise the accuracy-privacy trade-off of full decentralisation. Specifically, we investigate whether gossip-based convergence preserves the diagnostic signal present in centralised approaches, and how genuine clinical data heterogeneity across sites interacts with peer-to-peer model propagation. We will present and report a full analysis of the gossip-based results outlined above at the Symposium. Ultimately, this work contributes to a deeper understanding of how fully decentralised ML pipelines can be deployed for privacy-sensitive health monitoring on edge and mobile infrastructure.