

The faculty of Engineering of the Vrije Universiteit Brussel invites you to attend the public defense leading to the degree of

DOCTOR OF ENGINEERING SCIENCES

of **Yiming Chen**

The public defense will take place on **Friday 15th November 2024 at 4pm** in room **Q.D.** (Building Q, VUB Main Campus)

To join the digital defense, please click [here](#)

COMMUNICATION-EFFICIENT AND PRIVACY-PRESERVING
DECENTRALIZED TRAINING OF DEEP LEARNING METHODS

BOARD OF EXAMINERS

Prof. dr. Elisa Gonzalez Boix

Prof. dr. ir. Wendy Meulebroeck

Prof. dr. ir. Hichem Sahli

Prof. dr. ir. Bart Goossens

Prof. dr. Roel Wuyts

PROMOTORS

Prof. dr. ir. Nikolaos Deligiannis

Abstract of the PhD research

The rapid advancement in deep learning models, particularly in Natural Language Processing (NLP) and Computer Vision (CV), has significantly impacted various domains, such as content creation in media, medical diagnostics, and autonomous systems. For instance, the Vision Transformer (ViT) models, especially the ViT-Huge variant, which contains 632 million parameters, demonstrate outstanding performance on image classification benchmarks. Their superior capability is crucial for the development of visual perception systems in next-generation autonomous systems. Yet, such deep learning models demand considerable computational resources and extensive training datasets. Decentralized training frameworks emerge as a viable strategy to mitigate these challenges, which distributes the computational load across multiple edge nodes. This framework not only accelerates the training process but also maintains data privacy, as the training data resides locally at the edge nodes and is not transmitted to a central server. Consequently, it encourages third parties to contribute more sensitive data by ensuring data privacy. Despite the benefits of decentralized training, its wider adoption is hampered by issues such as communication overhead and potential privacy breaches. The exchange of substantial gradient volumes between the central server and edge nodes necessitates high communication rates, which impose bandwidth limitations and latency issues, particularly in scenarios lacking robust internet infrastructure. Moreover, privacy concerns have been amplified by studies demonstrating that malicious servers could deduce sensitive client attributes (e.g., gender, age) from shared gradients. A more alarming threat, Gradient Inversion Attacks (GIAs), can reconstruct clients' training data from gradients, thereby extracting maximum information and posing serious privacy risks. Those privacy risks impede the willingness of third parties to contribute the training data. Taking a step in addressing these challenges, this thesis explores three pivotal domains in decentralized training: (1) efficient communication, (2) privacy leakage assessment, and (3) privacy-preserving gradient-sharing techniques. Our first major contribution is a distributed Adam optimization approach paired with an aggressive gradient sparsification compression strategy tailored for transformer-based models. This approach drastically reduces gradient transmission to just 0.1% of its original size without compromising model efficacy. The second contribution presents a novel GIA that effectively compromises some well-established privacy-preserving gradient-sharing techniques relying on stochasticity (perturbation) during the edge training, thereby exposing their security vulnerabilities. Lastly, we introduce a learned lossy compression approach aid to prevent information leakage, marking our third contribution.