Multimodal and Privacy-Preserving Machine Learning for Human-Behavior Analysis

Contact:

Mohamed Chetouani
ISIR - CNRS UMR7222 Sorbonne Université
mohamed.chetouani@sorbonne-universite.fr

Context:

Understanding human behavior is critical for developing intuitive, personalized AI systems, particularly in healthcare [Driss-2020]. However, analyzing human behavior involves sensitive data (e.g., audio, video, wearable sensors) that pose privacy challenges, especially in mental healthcare applications like stress detection and mood recognition [Aigrain-2018]. Current machine learning approaches often rely on centralized data collection, leading to limited applicability across tasks and raising ethical concerns [Guerra-Manzanares-2023]. Privacy regulations such as GDPR and the proposed AI Act emphasize the need for transparent, privacy-compliant solutions. This calls for innovative methods that balance privacy protection with the utility of predictive models while supporting diverse tasks through transferable multimodal representations [Zhao-2020].

Objectives:

This internship focuses on advancing privacy-preserving multimodal machine learning for human behavior analysis. The research will explore methods such as adversarial training, deep auto-encoders, and multimodal foundation models, aiming to achieve robust privacy-utility trade-offs and ethical AI deployment in healthcare [Guerra-Manzanares-2023]. A possible research direction will consider the generation of synthetic data using machine learning.

The specific objectives include:

- Develop generative AI models to learn intrinsic multimodal representations of human behavior while preserving privacy.
- Generate synthetic, privacy-compliant data that maintains utility for predictive tasks.
- Evaluate models on publicly available datasets (e.g., audio, video, and physiological data) for tasks like mental state detection and diagnosis.

This thesis is conducted within a collaborative effort between ISIR (SU) and TICLab (International University of Rabat, IUR, Morocco), which is a strategic partner of Sorbonne University.

Profile:

Level: Master 2 / Engineer school
Skills: Python, Machine Learning, Robotics, and Cognitive Science
Duration: 5-6 months

References:

[Aigain-2018} Aigrain, J., Spodenkiewicz, M., Dubuisson, S., Detyniecki, M., Cohen, D., & Chetouani, M. (2018). Multimodal stress detection from multiple assessments. IEEE Trans. Affect. Comput., 9 (4), 491–506.

[Driss-2020] Drissi, N., Ouhbi, S., Garcı-Berna, J. A., Idrissi, M. A. J., & Ghogho, M. (2020). Sensor-based solutions for mental healthcare: A systematic literature review. In International conference on health informatics.

[Guerra-Manzanares-2023] Guerra-Manzanares, A., Lopez, L. J. L., Maniatakos, M., & Shamout, F. E. (2023). Privacy-preserving machine learning for healthcare: Open challenges and future perspectives. In H. Chen & L. Luo (Eds.), Trustworthy machine learning for healthcare (pp. 25–40).
[Zhao-2020] Zhao, B. Z. H., Kaafar, M. A., & Kourtellis, N. (2020). Not one but many tradeoffs: Privacy vs. utility in differentially private machine learning. In Proceedings of the 2020 acm sigsac conference on cloudcomputing security workshop (p. 15–26). Association for Computing Machinery.