



Offer #2023-06744

PhD Position F/M Privacy on-demand and Security preserving Federated Generative Networks or Models

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

About the research centre or Inria department

Le centre Inria d'Université Côte d'Azur regroupe 37 équipes de recherche et 8 services d'appui. Le personnel du centre (500 personnes environ) est composé de scientifiques de différentes nationalités, d'ingénieurs, de techniciens et d'administratifs. Les équipes sont principalement implantées sur les campus universitaires de Sophia Antipolis et Nice ainsi que Montpellier, en lien étroit avec les laboratoires et les établissements de recherche et d'enseignement supérieur (Université Côte d'Azur, CNRS, INRAE, INSERM ...), mais aussi avec les acteurs économiques du territoire.

Présent dans les domaines des neurosciences et biologie computationnelles, la science des données et la modélisation, le génie logiciel et la certification, ainsi que la robotique collaborative, le Centre Inria d'Université Côte d'Azur est un acteur majeur en termes d'excellence scientifique par les résultats obtenus et les collaborations tant au niveau européen qu'international.

Assignment

Context

Future sixth-generation (6G) networks will be highly heterogeneous, with the massive development of mobile edge computing inside networks. Furthermore, 6G is expected to support dynamic network environments and provide diversified intelligent services with stringent Quality of Service (QoS) requirements. Various new intelligent applications and services will emerge (including augmented reality (AR), wireless machine interaction, smart city, etc) and will enable tactile communications and Internet of everything (IoE). This will challenge wireless networks in the dimensions of delay, energy consumption, interaction, reliability, and degree of intelligence and knowledge, but also in the dimension of information and data sharing. In turn, 6G networks will be expected about leveraging data at the next step of the new communication system generation. First of all, they will generate large amounts of data much more data than 5G networks: multiple sources as Core, Radio Access Network, OAM, User Equipments (UEs) but also as private and/or personal devices/machines massively connected, data-generator applications as sensing, localization, context-awareness services etc. Besides, unlike today's networks where traffic is almost entirely centralized, most 6G traffic will remain localized and highly distributed. The communication system will not only provide the bits reliably, but more importantly will provide the intelligent data processing through connectivity and resources computing in the devices, the edge, and the cloud in the network. For this, with Artificial Intelligence (AI) and Machine Learning (ML), machines will bring to networks the necessary intelligence very close to the place of action and decision-making and will also make data sharing possible.

Reliable and efficient transmission, data privacy and security are great challenges in data sharing. Specially for 5G advanced and 6G networks data is distributed with the wide deployment of various connected Internet of Thing (IoT) devices, and are generated from many distributed network nodes, e.g., end users, small Base Stations or Distributed Units and the network edge. Also, how to collect/share efficiently data from multiple sources (e.g., sensors or device) up to AI/ML-based Network applications/services of Orchestration and Automation Layer (network management system) in Edge? The models shall be trained, updated regularly and operate in real-time.

Recently, generative models have been demonstrated playing a key role in data sharing while preserving privacy and security. They are able to generate synthetic data which distribution is similar to the original data one. So, instead of sending original data, many applications (medical or financial) use them to transfer data. Generative models are shown be useful in many scenarios such as health and financial applications [VSV+22]. However, the highly distributed architecture in 5G advanced/6G motivates the need for distributed, multi-agent learning for building generative models located at given anchor points of data collection (Edge server or Central Units) inside the RAN/Edge.

Challenges and objectives

We aim to design a communication-efficient and privacy-preserving on demand framework such that the local agents inside RAN/Edge cooperatively generate a synthetic dataset which represents well the global data distribution for model utility. To this end, one can train a generative adversarial network in a federated way [AMR+20], where the agents and the server alternatively minimize the loss function of the

discriminator and the generator. However, deep generative models have a tendency to memorize the training examples which may leak private information [HMDC19, CYZF20]. While, applying the traditional privacy-preserving defense such as differential privacy mechanism [Dwo06] will degrade the generative model's utility and thus influences the synthetic data quality. Moreover, the training requires 500-10000 communication rounds in practice for convergence (see [KMA+21, Table 2]) which is expensive for communication cost. Recently, there is another work [ZCL+22] where the server makes use of all the local trained models to train a generator, which minimizes the communication cost to only one round. However, transferring these local models are extremely dangerous as they can be used to infer the private information on the dataset of devices [FJR15, YGFJ18]. Alternatively, instead of transferring the models as the previous work proposed, the devices can transfer directly the distilled synthetic data which are computed locally [ZPM+20]. However, the quality of the assembled synthetic dataset degrades especially when some agents have just few training samples.

We will first compare the above-mentioned existing methods for synthetic dataset generation, in terms of their trade-offs on model accuracy, data similarity, communication cost, model compression and privacy. Then, to expose their privacy vulnerability, we will design computational-efficient attacks, for both passive and active adversary cases. Finally, we will design a framework with better trade-off for the task.

Teams and supervision

- INRIA : COATI (Frédéric Giroire, Chuan Xu), EPIONE (Marco Lorenzi)
- NOKIA: Bell Labs Core Research
- 3-years PhD to be hosted in Sophia Antipolis. The doctoral student will be supervised by his academic supervisor and his industrial supervisor.

Skills

The candidate should have a solid mathematical background, good programming skills and previous experience with PyTorch or TensorFlow. He/She should also be knowledgeable on machine learning, especially generative neural networks, and have good analytical skills. We expect the candidate to be fluent in English.

References

[AMR+20] Sean Augenstein, H Brendan McMahan, Daniel Ramage, Swaroop Ramaswamy, Peter Kairouz, Mingqing Chen, Rajiv Mathews, et al. Generative models for effective ml on private, decentralized datasets. ICLR, 2020.

[CYZF20] Dingfan Chen, Ning Yu, Yang Zhang, and Mario Fritz. Gan-leaks: A taxonomy of membership inference attacks against generative models. In Proceedings of the 2020 ACM SIGSAC conference on computer and communications security, pages 343–362, 2020.

[Dwo06] Cynthia Dwork. Differential privacy. In Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10–14, 2006, Proceedings, Part II 33, pages 1–12. Springer, 2006.

[FJR15] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pages 1322–1333, 2015.

[HMDC19] Jamie Hayes, Luca Melis, George Danezis, and Emiliano De Cristofaro. LOGAN: membership inference attacks against generative models. Proc. Priv. Enhancing Technol., 2019(1):133–152, 2019.

[KMA+21] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. Foundations and Trends® in Machine Learning, 14(1–2):1–210, 2021.

[VSV+22] Rohit Venugopal, Noman Shafqat, Ishwar Venugopal, Benjamin Mark John Tillbury, Harry Demetrios Stafford, and Aikaterini Bourazeri. Privacy preserving generative adversarial networks to model electronic health records. Neural Networks, 153:339–348, 2022.

[YGFJ18] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In 2018 IEEE 31st Computer Security Foundations Symposium (CSF), pages 268–282. IEEE, 2018.

[ZCL+ 22] Jie Zhang, Chen Chen, Bo Li, Lingjuan Lyu, Shuang Wu, Shouhong Ding, Chunhua Shen, and Chao Wu. Dense: Data-free one-shot federated learning. In Advances in Neural Information Processing Systems, 2022.

[ZPM+ 20] Yanlin Zhou, George Pu, Xiyao Ma, Xiaolin Li, and Dapeng Oliver Wu. Distilled one-shot federated learning. ArXiv, abs/2009.07999, 2020.

Main activities

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking (after 6 months of employment) and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage

Remuneration

Durée: 36 mois

Localisation: Sophia Antipolis, France

Rémunération: 2082€ brut mensuel (année 1 & 2) et 2190€ brut mensuel (année 3)

General Information

- **Theme/Domain** : Networks and Telecommunications
System & Networks (BAP E)
- **Town/city** : Sophia Antipolis
- **Inria Center** : [Centre Inria d'Université Côte d'Azur](#)
- **Starting date** : 2024-01-01
- **Duration of contract** : 4 years
- **Deadline to apply** : 2024-06-30

Contacts

- **Inria Team** : [COATI](#)
- **PhD Supervisor** :
Giroire Frédéric / Frederic.Giroire@inria.fr

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

Warning : you must enter your e-mail address in order to save your application to Inria. Applications must be submitted online on the Inria website. Processing of applications sent from other channels is not guaranteed.

Instruction to apply

Defence Security :

This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST). Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :

As part of its diversity policy, all Inria positions are accessible to people with disabilities.