Intelligent Networks | DOI:10.1145/3512343

Privacy-Preserving AI for Future Networks

BY DIEGO PERINO, KLEOMENIS KATEVAS, ANDRA LUTU, EDUARD MARIN, AND NICOLAS KOURTELLIS

TELCO NETWORKS AND systems evolved over the years to deal with novel services. Today, they are highly complex, distributed ecosystems composed of very diverse sub-environments (see Figure 1). They include myriad types of devices, connectivity means, protocols, and infrastructures often managed by different teams with varying expertise and tools, or even different companies.

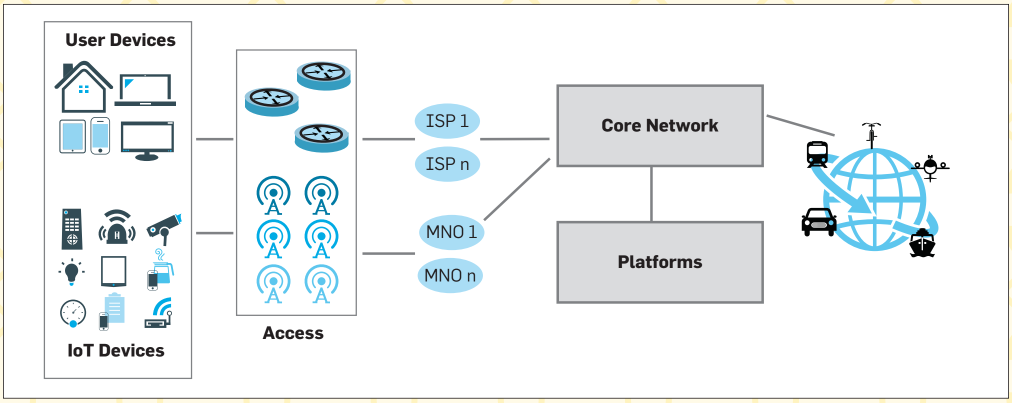

Figure 1. High-level view of the complexity of telcos' networks and systems with a large variety of devices, connectivity means, protocols, and infrastructures.

Traditional network management solutions (for example, network overprovisioning, rule-based systems, reactive approaches) are reaching their limits in dealing with this complex ecosystem. Novel solutions are required to guarantee strict service requirements and effective resource management, especially in cases where entities have a partial view of the system.

AI to the rescue? In the last decade, we have witnessed a growing interest within the networking research community toward artificial intelligence (AI),[a] whose techniques have been applied to a wide range of use cases: network optimization, routing, scheduling algorithms, resource and fault management mechanisms, Quality of Service (QoS) and Quality of Experience (QoE) management, network security and many other tasks.[1] On the industrial side, AI is slowly complementing traditional networking approaches worldwide: Small Medium Enterprises (SMEs) and start-ups are developing AI solutions to deal with specific use cases, traditional networking vendors are evolving their products to support AI tools, and major cloud/software providers are adapting AI tools to be used in the networking domain. This is also the case for telephone service providers (telcos), which are exploring the application of AI algorithms through internal research and innovation (R&I) projects. For instance, our group is working on AI-based approaches in many use-cases focusing on realistic environments and applications (for example, Kattadige et al.[2] and Perino et al.[5]), partially in collaboration with other European partners in the context of European R&I actions.[b]

What about my privacy? AI models and tools are potentially vulnerable, and their usage introduces new attack vectors for telco environments. For instance, membership and property inference or data reconstruction attacks, and adversarial learning, can reveal different aspects of the data used (for example, which specific users' data were used for model training), the values of their data attributes, and even user patterns such as their mobility or browsing behavior. Therefore, the use of AI techniques can impact user privacy more strongly than traditional data analysis methods since AI models can distill information from multiple data sources and infer rich patterns regarding

AI models and tools are potentially vulnerable, and their usage introduces new attack vectors for telco environments.

[a] We refer to artificial intelligence as the set of tools falling in the category of machine learning and deep learning.

[b] CONCORDIA (https://www.concordia-h2020.eu), DAEMON (https://h2020daemon.eu), ACCORDION (https://www.accordion-project.eu), CHARITY (https://www.charity-project.eu/en), SPATIAL (https://spatial-h2020.eu)

their data owners. Further exacerbating the privacy problem, data and model poisoning attacks can even manipulate AI models to take adversarial decisions for targeted users. Thus, recently enforced data privacy regulations (first in EU with GDPR and e-Privacy, since 2020 in U.S. with CCPA, and elsewhere in the world) attempt to mitigate these risks with specific guidelines to data operators/processors using AI methods.

To address this challenge, to follow regulations, and to build systems able to guarantee privacy by design, the R&I community is investigating the use of privacy-preserving AI (PPAI) methods, including techniques such as federated learning (FL), differential privacy, policy-based AI, and trusted execution environments (TEEs).

Our research group is building on these techniques with focus on the complex telco ecosystems. Indeed, the capabilities at the edge and network devices (for example, IoT, phones, routers, antennas) can be used to perform the computation of AI models in a distributed hierarchical fashion without the need of sharing the data, and thus limiting the risk of private information leakage. This could also be done following an "as a service" approach to facilitate AI model building in a collaborative fashion between companies and mitigating attacks against FL model building using TEEs, as shown in our recent works[3,4] and Figure 2). Furthermore, there are major trade-offs between privacy and utility of PPAI,[6] and especially when introducing hierarchies in the FL process.

What's Next?

R&I is still needed to create autonomous, intelligent, and yet fully privacy-preserving and secure telco "networks." Creating customized AI-based approaches that tackle the specific challenges of network-related problems require more work. Particular attention should be devoted to the trade-off between creating complex tools that guarantee the required level of performance and robustness, and tools that network engineers trust and use. Interestingly, for many scenarios, it is still unclear whether AI is superior to traditional approaches, and how to seamlessly complement traditional methods with AI.

Further R&I is also required to design PPAI mechanisms suitable for these environments, including hierarchical FL, on-device training optimizations, or adaptive mechanisms able to deal with heterogeneous computing environments.

References

1. Casas, P. Two decades of AI4NETS—AI/ML for data networks: Challenges and research directions. In *Proceedings of IEEE/IFIP NOMS*, 2020.
2. Kattadige, C., Raman, A., Thilakarathna, K., Lutu, A., Perino, D. 360NorVic: 360-degree video classification from mobile encrypted video traffic. In *Proceedings of ACM NOSSDAV*, 2021.
3. Kourtellis, N., Katevas, K. and Perino, D. FLaaS: Federated learning as a service. In *Proceedings of the ACM CoNext Distributed ML*, 2020.
4. Mo, F., Haddadi, H., Katevas, K., Marin, E., Perino, D. and Kourtellis, N. Privacy-preserving federated learning with TEEs. In *Proceedings of ACM MobiSys*, 2021 (best paper).
5. Perino, D., Yang, X., Serra, J., Lutu, A. and Leontiadis, I. Experience: Advanced network operations in (Un)-connected remote communities. In *Proceedings of ACM MobiCom*, 2020.
6. Zhao, B.Z.H., Kaafar, M.A., Kourtellis, N. Not one but many trade-offs: Privacy vs. utility in differentially private machine learning. In *Proceedings of ACM CCSW*, 2020.

Diego Perino is Director of Research at Telefonica Research, Spain.

Kleomenis Katevas is a Research Scientist at Telefonica Research, Spain.

Andra Lutu is Senior Research Scientist at Telefonica Research, Spain.

Eduard Marin is a Research Scientist at Telefonica Research, Spain.

Nicolas Kourtellis is Senior Research Scientist at Telefonica Research, Spain.

Copyright held by authors/owners.
Publication rights licensed to ACM.

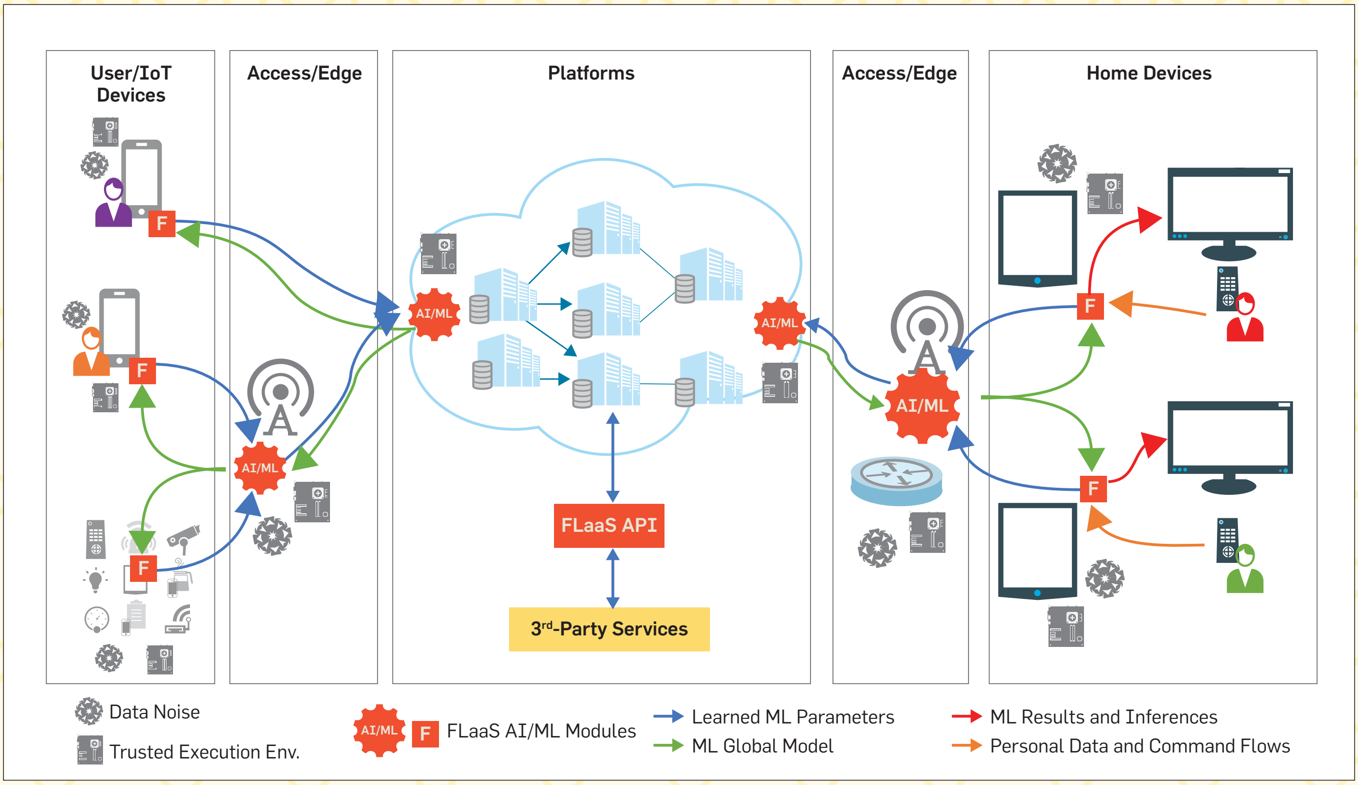

Figure 2. Example of Privacy Preserving Federated Learning in a telco network. Federated Learning is provided in an "as a service" approach with potential attacks mitigated by the usage of TEEs, data noise, and hierarchical model building.