FAIR – Continual Learning at the Edge

Background

Edge AI systems are characterised by many devices located in the network that are in charge of training AI models using only the data they are exposed to, possibly collaborating together, according to decentralised AI paradigms. Since such a pervasive scenario might be highly dynamic, the type and complexity of the learning task that devices have to face might evolve, new data could become available over time, and new devices could join the collaborative task, starting from local models of largely different accuracies. Therefore, Edge devices should be able to update and adapt their AI models at the same pace. Continuous Learning (CL) is based on the idea of learning continuously and adaptively about the external world, enabling autonomous and incremental development of increasingly complex skills and knowledge. In the context of machine learning, this means being able to seamlessly update the prediction model to account for different tasks and data distributions and reuse and retain helpful knowledge and skills over time. CL at the edge poses radically new challenges with respect to traditional CL, as edge devices - in addition to learning novel patterns based on new data, also have to (i) "continuously forget" patterns that are not relevant anymore to optimise local resource usage, and; (ii) collaborate with other devices entering and leaving the collaboration dynamically, which provide local models of different "maturity".

Topic description

In this line of work, the aim is to develop continual learning methods suitable for devices with limited computational resources. The work will initially evaluate the methodologies of the state of the art in continual learning, along with their impact from the perspective of real applicability on Edge devices. The idea is to investigate how to design new algorithms and methods that allow Edge devices to continuously update their AI models to make them learn new patterns that might show over time, but also to "forget" abilities that are not important anymore because they might be related to previous requirements/conditions or, more importantly, for optimising the devices' resources making room for new and more timely abilities to be learnt. Considering that in a pervasive Edge AI scenario, the computational capacity of devices is limited, this topic can be addressed both from the "standalone device" point of view as well as from the distributed learning perspective by proposing new methods for collaborative continual learning.

Type of prospect positions

We plan to open positions at the level of Researcher on this topic.

Funding and partnerships

Candidate profile

Ideal candidates should have or are about to obtain an MSc in Computer Science, Computer Engineering, Mathematics, Physics, or closely related disciplines and a proven track record of excellent University grades. Preferably, the MSc should be in one of the relevant research areas: Artificial Intelligence, Pervasive Systems.

Contacts

Andrea Passarella   andrea.passarella@iit.cnr.it

Scholar profile   https://scholar.google.com/citations?user=sesKnygAAAAJ

Marco Conti   marco.conti@iit.cnr.it

Scholar profile   https://scholar.google.com/citations?user=KniFTD0AAAAJ

Lorenzo Valerio   lorenzo.valerio@iit.cnr.it

Scholar profile   https://scholar.google.com/citations?user=CjoX4voAAAAJ

Chiara Boldrini   chiara.boldrini@iit.cnr.it

Scholar profile   https://scholar.google.com/citations?user=rCrzxbsAAAAJ