

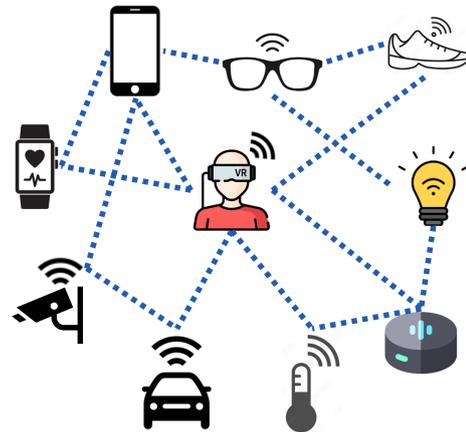


## Topics in Distributed and Decentralized Artificial Intelligence

Are you interested in **distributed systems** and **machine learning**, and would like to work on a project that combines **both** worlds? Here you will find a list of project proposals for **Bachelor**, **Master**, and **Semester** Theses. The chosen project's expectations will depend on the **type** of thesis and the topic can be **refined** together with the candidate depending on their preferences and strengths.

### Introduction

Networks composed of a vast number of devices capable of **computing**, **sensing**, and **communication**, distributed over wide areas form the ideal substrate for **decentralized** intelligent applications in fields ranging from medical, environmental, industrial, and several others. In this scenario, without necessarily relying on centralized orchestration or trusted peers, user devices can offer their resources (compute, communication, data, etc.) to the system, dramatically increasing its **scalability**, **resiliency**, **democracy**, and the global amount of information used for the **collaborative** model training and inference (e.g., with Federated Learning). However, users might require **confidentiality** for some local data, need to **move** in space, and have **constrained resources** (e.g., energy, communication, computation).



Therefore, users should be motivated to participate in such system by various **incentives**: i.e., guaranteeing and optimizing (1) the **privacy** of their local data and (2) their application's **task performance** by granting them use of a suitable share of the system's resources (e.g., computing and model capabilities). In such a system, user devices can continuously monitor the system state and optimize its operation by regulating the parameters of their local distributed algorithm (e.g., with Reinforcement Learning).

### Topics

Such general and complex scenario presents pressing challenges that motivate several research direction, including (among others):

- **Resource-aware Distributed Split Learning.**

Design distributed or decentralized online methods (e.g., Multi-Agent Reinforcement Learning) to partition, offload, and execute model training and/or inference across a set of (possibly mobile) nodes owning a time-varying amount of limited resources. Nodes might require agreeing on the model split and be incentivized to participate. Nodes might cooperate, compete, or be adversarial. Global system's goals are optimizing

model task performance and latency (training, inference), while local agent's goals is maximizing their own benefit under constraints.

- **Heterogeneous Adaptive Distributed Model Architecture Search.**

Design online decentralized methods to adapt the architecture of distributed (e.g., Federated) model on multiple cooperating devices, to maximize task performance (e.g., accuracy, latency) while considering node available resources (e.g., compute, energy, communication).

- **Large-scale Federated Model Selection with Opportunistic Domains**

Design decentralized online methods to build a global model by cooperation of an arbitrarily large number of compute-capable wireless mobile sensors efficiently. Node can communicate via long-range and short-range links. The system operation depends on node resource availability, user mobility, and spatial data distribution correlations, which should be estimated over time, following uncertainty minimization. Each node adapts communication and computation parameters to build models relevant to a geographical area via direct opportunistic communication, combined through an optimized topology of longer-range links.

- **Asynchronous Decentralized Knowledge Compression and Exchange**

Design online decentralized methods to optimally compress local model information, which may be relevant to other peers' models, and to share it with them by optimally leveraging the currently available communication budget and network topology. Methods may use information theory or statistical learning to guide the compression scheme and the dissemination procedure. Compression can also be learned.

- **Incentives for Decentralized Intelligence**

Design decentralized protocols that allow a user to prove their contribution (e.g., in terms of data, compute, or communication) towards a shared model to other peer nodes, who allow the user to benefit from it suitably (e.g., in terms of model capability, usage time, or compute offered by other peers).

- Do you have **another idea** for a research project that combines distributed or decentralized systems and artificial intelligence? **Let's discuss!**

For all research directions, you are free to select any real-world application of your choice, which fits the overarching scenarios and requirements: multiple agents, resource constraints, and online decisions.

### **Interested? Contact us for more information**

- Dr. Antonio Di Maio: [adimaio@ethz.ch](mailto:adimaio@ethz.ch), ETZ G97

*Last update: October 2025*