




Distributed Asynchronous Policy (DAP) with RL

Lately, applying deep learning to graph-based problems has gained a lot of interest. Graph neural networks, in particular, achieved great success, on a wide range of topics. At the moment all prominent models act in synchronous rounds and fall in the class of Synchronous Distributed Algorithms (send \rightarrow receive \rightarrow compute). In contrast, an Asynchronous Model is event driven (on receive \rightarrow act). Especially in IoT systems it is interesting to act in an asynchronous model. First of all the connections could be more loosely (e.g. no simultaneous wake-up/ leader election) and agents do not have to keep a connection alive for irrelevant actions (energy saving). E.g. the smart fridge does not need to know that the motion sensor is triggered.

The goal of this project is to explore if and how a distributed asynchronous policy can be trained. Identify which obstacles one has to overcome, program and train a model being able to either solve a variety of tasks directly or provide a framework to do so. Clean, general and practical solutions are preferred.



dap

Requirements: Strong motivation, knowledge in deep learning, or a solid background in machine learning. Prior experience in Machine Learning, Reinforcement Learning, Distributed Systems and Graph Neural Networks extremely help full. Creativity and programming skills are necessary.

Interested? Please contact us for more details!

Contact

- Karolis Martinkus: martinkus@ethz.ch, ETZ G95