Deep reinforcement learning for real-time power grid topology optimization

As we have come to rely on a constant, uninterrupted electricity supply, the transmission of electricity from generators to consumers has become a vital backbone of modern society. Owing to how safely and reliantly power grids usually perform their task, we tend to take them for granted. Power grids are however some of humanities most complex systems, requiring regular human intervention to function well.

In light of rapid climate change, the sustainable development goals, the Paris climate agreement, and dire outlooks by IPCC, recent years have seen The European Union, Japan, South Korea, and more than 110 other countries pledge carbon neutrality by 2050. Living up to such pledges will require drastic changes to generation and consumption of electricity, which in turn will cause new and complex demands on the power grid and its operators.

Such added demands would traditionally be addressed by physically expanding the grid. However, changing regulations and economics means this will no longer do, making other approaches needed. A cheap, promising, and under-exploited mitigation is real-time topology optimization (RTTO). However, beyond the simplest action of line switching, the combinatorial and non-linear nature of RTTO has made all computational approaches infeasible for grids of interesting scale. Deep reinforcement learning (DRL) may be about to change that.

This thesis starts by providing some further background as to why we care about RTTO. Assuming little previous knowledge, it then covers some relevant parts of deep learning and reinforcement learning, before building on this to give a quite thorough explanation of some central DRL concepts. This is followed by a case study on L2RPN 2019 – an RTTO competition in a simulated environment – and the winning submission by Geirina. Finally, some arguments are given for why DRL is a promising approach to the challenge of RTTO.