MAX PLANCK INSTITUTE FOR INTELLIGENT SYSTEMS
Learning and Dynamical Systems Group (Dr. Michael Muehlebach)

Master Thesis Project

Efficient computation of Nash equilibria

The project is concerned with the efficient computation of Nash equilibria in zero-sum games, that is,

$$\inf_x \sup_y f(x,y),$$

where $f: \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is a smooth function. A fundamental difficulty arises due to the fact that a standard gradient-descent/gradient-ascent algorithm not necessarily converges to local solutions of the above minimax problem. Instead, it might converge to other stationary points, such as saddle points, which are not properly aligned (i.e., although the gradient along x and y vanishes, the saddle point does not correspond to a local minimum along the x-coordinate and a local maximum along the y-coordinate). A solution to this problem has been proposed in [1], where the authors have designed a mechanism that exploits second-order curvature information for avoiding undesirable equilibria. However, at each iteration, the resulting algorithm requires the inversion of the Hessian, which can become prohibitive for large problem sizes.

The thesis project extends the work of [1] and aims at designing algorithms that obtain second-order information via the history of past gradient evaluations. In that way, the computation and inversion of the Hessian could potentially be avoided, improving the scalability of the algorithm to large problem sizes.

Computing Nash equilibria in zero-sum games (or equivalently solving minimax problems of the above form) has many applications ranging from economics and control theory to machine learning. In machine learning, solutions of minimax problems lie at the heart of generative adversarial networks or can be used to robustify learning systems, for example against adversaries or misspecifications of the data-generating process.

[1] Eric Mazumdar, Michael I. Jordan, and S. Shankar Sastry, "*On finding local Nash equilibria (and only local Nash equilibria) in zero-sum continuous games*", arXiv:1901.00838, 2019, https://arxiv.org/abs/1901.00838

Learning and Dynamical Systems Group (https://lds.is.mpg.de/)
The Learning and Dynamical Systems Group is part of the Max Planck Institute for Intelligent Systems in Tübingen, Germany. Part of our research aims at exploiting analogies between dynamical systems and mathematical optimization for gaining insights and constructing new optimization algorithms that are intuitive (due to their physical/mechanical analogy) and scale well to high-dimensional problems. For example, we showed how to rigorously understand and generalize Nesterov acceleration by drawing on analogies to physical mass-spring-damper systems.

Prerequisites
Strong analytical skills and programming experience (Python, MATLAB, C/C++ or similar). Background in machine learning, control theory, statistics, or mathematical optimization is a plus.

Contact
If you have any questions do not hesitate to contact us. When applying for a project, please include your CV, bachelor's and master's transcripts, and a one-page letter of motivation describing your research interests and educational background.

Dr. Michael Muehlebach, michael.muehlebach@tuebingen.mpg.de