A neural network based policy iteration algorithm with global H^2-superlinear convergence for stochastic games on domains

Author and Presenter: Yufei **Zhang** (University of Oxford, UK)

Co-authors: Kazufumi **Ito** (North Carolina State University, USA)
Christoph **Reisinger** (University of Oxford, UK)

Abstract

In this work, we propose a class of numerical schemes for solving semilinear Hamilton-Jacobi-Bellman-Isaacs (HJBI) boundary value problems which arise naturally from exit time problems of diffusion processes with controlled drift. We exploit policy iteration to reduce the semilinear problem into a sequence of linear Dirichlet problems, which are subsequently approximated by a multilayer feedforward neural network ansatz. We establish that the numerical solutions converge globally in the H^2-norm, and further demonstrate that this convergence is superlinear, by interpreting the algorithm as an inexact Newton iteration for the HJBI equation. Moreover, we construct the optimal feedback controls from the numerical value functions and deduce convergence. The numerical schemes and convergence results are then extended to oblique derivative boundary conditions. Numerical experiments on the stochastic Zermelo navigation problem and the perpetual American option pricing problems are presented to illustrate the theoretical results and to demonstrate the effectiveness of the method.

References

[1] K. Ito, Ch. Reisinger and Y. Zhang: *A neural network based policy iteration algorithm with global H^2-superlinear convergence for stochastic games on domains*, Preprint, arXiv:1906.02304, 2019.