

# AN EXPONENTIALLY CONVERGING PARTICLE METHOD FOR THE MIXED NASH EQUILIBRIUM OF CONTINUOUS GAMES

**Guillaume Wang**  
EPFL, Switzerland  
guillaume.wang@epfl.ch

We consider the problem of computing mixed Nash equilibria of two-player zero-sum games with continuous sets of pure strategies and with first-order access to the payoff function. This problem arises for example in game-theory-inspired machine learning applications, such as distributionally-robust learning. In those applications, the strategy sets are high-dimensional and thus methods based on discretisation cannot tractably return high-accuracy solutions.

In this paper, we introduce and analyze a particle-based method that enjoys guaranteed local convergence for this problem. This method consists in parametrizing the mixed strategies as atomic measures and applying proximal point updates to both the atoms' weights and positions. It can be interpreted as a time-implicit discretization of the “interacting” Wasserstein-Fisher-Rao gradient flow.

We prove that, under non-degeneracy assumptions, this method converges at an exponential rate to the exact mixed Nash equilibrium from any initialization satisfying a natural notion of closeness to optimality. We illustrate our results with numerical experiments and discuss applications to max-margin and distributionally-robust classification using two-layer neural networks, where our method has a natural interpretation as a simultaneous training of the network's weights and of the adversarial distribution.

*Joint work with Lénaïc Chizat (EPFL, Switzerland).*