
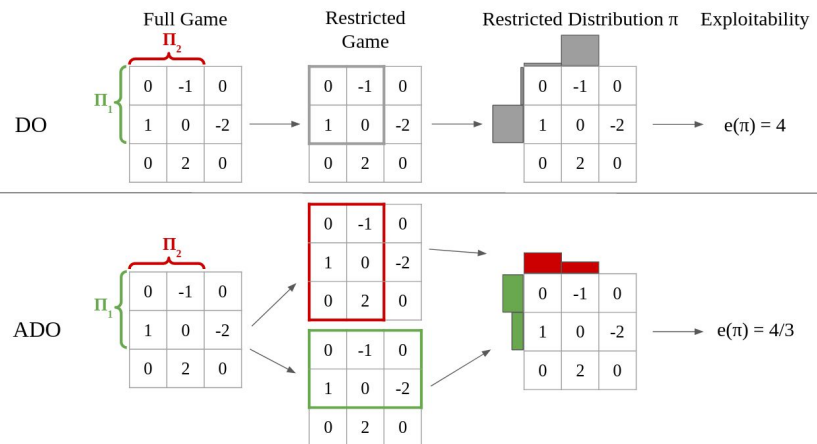


Anytime PSRO for Two-Player Zero-Sum Games



Stephen McAleer, Kevin Wang, JB Lanier, Marc Lanctot, Pierre Baldi, Tuomas Sandholm, Roy Fox

Problem: Double Oracle can increase exploitability before converging

To scale to large games, update restricted distribution via no-regret algorithm against opponent best response



Solution: Create two restricted games where one player is restricted but the other is not (similar to [1]).



1. Zinkevich, M., Bowling, M., and Burch, N. A new algorithm for generating equilibria in massive zero-sum games. In Twenty-Second Conference on Artificial Intelligence (AAAI), 2007