

# Bandit-Based Monte-Carlo Planning: Beyond the Game of Go

Olivier Teytaud

(Inria-Futurs) October 9, 2009

The TAO team developed MoGo, the first program which defeated a professional Go player, thanks to bandit-based Monte-Carlo planning (BBMCP; variants are known as UCT or MCTS).

The most interesting point in this victory is that BBMCP can be applied for one player games as well, including stochastic planning. The talk will survey BBMCP, so that all attendees, at the end of the talk, are able to implement it if they want to; we will also present ideas for future work, in particular for partially observable frameworks.