

From the very dawn of the field, search with value functions was a fundamental concept of computer games research. Turing's chess algorithm from 1950 was able to think two moves ahead, and Shannon's work on chess from 1950 includes an extensive section on evaluation functions to be used within a search. Samuel's checkers program from 1959 already combines search and value functions that are learned through self-play and bootstrapping. TD-Gammon improves upon those ideas and uses neural networks to learn those complex value functions — only to be again used within search. The combination of decision-time search and value functions has been present in the remarkable milestones where computers bested their human counterparts in long standing challenging games — DeepBlue for Chess and AlphaGo for Go.

Until recently, this powerful framework of search aided with (learned) value functions has been limited to perfect information games. As many interesting problems do not provide the agent perfect information of the environment, this was an unfortunate limitation. This thesis introduces the reader to sound search for imperfect information games.