

Probabilistic Boolean Decision Trees and the Complexity of Evaluating Game Trees

Michael Saks
Avi Wigderson

Abstract

The Boolean Decision tree model is perhaps the simplest model that computes Boolean functions; it charges only for reading an input variable. WE study the power of randomness (vs. both determinism and non-determinism) in this model, and prove separation results between the three complexity measures.

These results are obtained via general and efficient methods for computing upper and lower bounds on the probabilistic complexity of evaluating Boolean formulae in which every variable appears exactly once (AND/OR tree with distinct leaves). These bounds are shown to be exactly tight for interesting families of such tree functions.

We then apply our results to the complexity of evaluating game trees, which is a central problem in \$AI\$. These trees are similar to Boolean tree functions, except that input variables (leaves) may take values from a large set (of valuations to game positions) and the AND/OR nodes are replaced by MIN/MAX nodes. Here the cost is the number of positions (leaves) probed by the algorithm.

The best-known algorithm for this problem is the alpha-beta pruning method. As a deterministic algorithm, it will in the worst case have to examine all positions. Many papers studied the expected behavior of alpha-beta pruning (or uniform trees) under the unreasonable assumption that position values are drawn independently from some distribution. We analyze a randomized variant of alpha-beta pruning, show that it is a considerably faster than deterministic one in worst case, and prove it optimal for uniform trees.