THE SOLUTION FOR THE BRANCHING FACTOR OF THE ALPHA-BETA PRUNING ALGORITHM

Judea Pearl
Cognitive Systems Laboratory
School of Engineering and Applied Science
University of California, Los Angeles
Los Angeles, California 90024

ABSTRACT

This paper analyzes $N_{n,d}$, the average number of terminal nodes examined by the α-β pruning algorithm in a uniform game-tree of degree n and depth d for which the terminal values are drawn at random from a continuous distribution. It is shown that $N_{n,d}$ attains the branching factor $\mathscr{R}_{\alpha-\beta}(n) = \xi_n/1-\xi_n$ where ξ_n is the positive root of $x^n+x-1 = 0$. The quantity $\xi_n/1-\xi_n$ has previously been identified as a lower bound for all directional algorithms. Thus, the equality $\mathscr{R}_{\alpha-\beta}(n) = \xi_n/1-\xi_n$ renders α-β asymptotically optimal over the class of directional, game-searching algorithms.

1. INTRODUCTION

The α-β pruning algorithm is the most commonly used procedure in game-playing applications. It serves to determine the minimax value of the root of a tree for which the terminal nodes are assigned arbitrary numerical values [1]. Although the exponential growth of such game-tree searching is slowed significantly by that algorithm, quantitative analyses of its effectiveness have been frustrated for over a decade. One concern has been to determine whether the α-β algorithm is optimal over other game-searching procedures.

The model most frequently used for evaluating the performance of game-searching methods consists of a uniform tree of depth d and degree n, where the terminal positions are assigned random, independent, and identically distributed values. The number of terminal nodes examined during the search has become a standard criterion for the complexity of the search method.

Slagle and Dixon (1969) showed that the number of terminal nodes examined by α-β must be at least $n^{\lfloor d/2 \rfloor} + n^{\lceil d/2 \rceil} - 1$ but may, in the worst case, reach the entire set of n^d terminal nodes [2]. The analysis of expected performance using uniform trees with random terminal values had begun with Fuller, Gaschnig, and Gillogly [3] who obtained formulas by which the average number of terminal examinations, $N_{n,d}$, can be computed. Unfortunately, the formula would not facilitate asymptotic analysis; simulation studies led to the estimate $\mathscr{R}_{\alpha-\beta} \approx (n)^{.72}$.

Knuth and Moore [1] analyzed a less powerful but simpler version of the α-β procedure by ignoring deep cutoffs. They showed that the branching factor of this simplified model is O(n/log n) and speculated that the inclusion of deep cutoffs would not alter this behavior substantially. A more recent study by Baudet [4] confirmed this conjecture by deriving an integral formula for $N_{n,d}$ (deep cutoffs included) from which the branching factor can be estimated. In particular, Baudet shows that $\mathscr{R}_{\alpha-\beta}$ is bounded by $\xi_n/1-\xi_n \le \mathscr{R}_{\alpha-\beta} \le M_n^{1/2}$ where ξ_n is the positive root of $x^n+x-1 = 0$ and M_n is the maximal value of the polynomial $P(x) = \frac{1-x^n}{1-x} \cdot \frac{1-[1-x^n]^n}{x^n}$ in the range $0 \le x \le 1$. Pearl [5] has shown both that $\xi_n/1-\xi_n$ lower bounds the branching factor of every directional game-searching algorithm and that an algorithm exists (called SCOUT) which actually achieves this bound. Thus, the enigma of whether α-β is optimal remained contingent upon determining the exact magnitude of $\mathscr{R}_{\alpha-\beta}$ within the range delineated by Baudet.

This paper now shows that the branching factor of α-β indeed coincides with the lower bound $\xi_n/1-\xi_n$, thus establishing the optimality of α-β over the class of directional search algorithms.

2. ANALYSIS

Our starting point is Baudet's formula for $N_{n,d}$:

Theorem 1: (Baudet [4], Theorem 4.2)

Let $f_0(x) = x$ and, for $i = 1, 2, \ldots$, define:

$$f_i(x) = 1-\{1-[f_{i-1}(x)]^n\}^n,$$

$$r_i(x) = \frac{1-[f_{i-1}(x)]^n}{1-f_{i-1}(x)},$$

$$s_i(x) = \frac{f_i(x)}{[f_{i-1}(x)]^n},$$

$$R_i(x) = r_1(x) \times \ldots \times r_{\lceil i/2 \rceil}(x) ,$$

$$S_i(x) = s_1(x) \times \ldots \times s_{\lfloor i/2 \rfloor}(x) .$$

The average number, $N_{n,d}$, of terminal nodes examined by the α-β pruning algorithm in a uniform game-tree of degree n and depth d for which the bottom values are drawn from a continuous distribution is given by:

$$N_{n,d} = n^{\lfloor d/2 \rfloor} + \int_0^1 R_d'(t)\ S_d(t)\ dt \tag{1}$$

The difficulty in estimating the integral in (1) stems from the recursive nature of $f_i(x)$ which tends to obscure the behavior of the integrand. We circumvent this difficulty by substituting for $f_0(x)$ another function, $\phi(x)$, which makes the regularity associated with each successive iteration more transparent.

The value of the integral in (1) does not depend on the exact nature of $f_0(x)$ as long as it is monotone from some interval [a, b] onto the range [0, 1]. This is evident by noting that by substituting $f_0(x) = \phi(x)$ the integral becomes:

$$\int_{x=a}^{b} \frac{dR_d[\phi(x)]}{dx} S_d[\phi(x)]dx = \int_{\phi=0}^{1} \frac{dR_d(\phi)}{d\phi}\ S_d(\phi)\ d\phi$$

which is identical to that in (1). The significance of this invariance is that, when the terminal values are drawn from a continuous distribution, the number of terminal positions examined by the α-β procedure does not depend on the shape of that distribution. Consequently, $f_0(x)$, which represents the terminal values' distribution, may assume an arbitrary form, subject to the usual constraints imposed on continuous distributions.

A convenient choice for the distribution $f_0(x)$ would be a characteristic function $\phi(x)$ which would render the distributions of the minimax value of every node in the tree identical in shape. Such a characteristic distribution indeed exists [6] and satisfies the functional equation:

$$\phi(x) = g[\phi(ax)] \tag{2}$$

where:

$$g(\phi) = 1-(1-\phi^n)^n \ , \tag{3}$$

and a is a real-valued parameter to be determined by the requirement that (2) possesses a non-trivial solution for $\phi(x)$. This choice of $\phi(x)$ renders the functions $\{f_i(x)\}$ in Theorem 1 identical in shape, save for a scale factor. Accordingly we can write:

$$f_i(x) = \phi(x/a^i) \tag{4}$$

$$r_i(x) = r(x/a^i) \tag{5}$$

$$s_i(x) = s(x/a^i) \tag{6}$$

where:

$$r(x) = \frac{1-[\phi(x)]^n}{1-\phi(x)} \tag{7}$$

and:

$$s(x) = \frac{1-\{1-[\phi(x)]^n\}^n}{[\phi(x)]^n} \tag{8}$$

Equation (2), known as Poincare Equation [7], has a non-trivial solution $\phi(x)$ with the following properties [6]:

i) $\phi(0) = \xi_n$ (9)

where ξ_n is the root of $x^n+x-1 = 0$

ii) $a = \frac{1}{g'(\xi_n)} = \left[\frac{\xi_n}{n(1-\xi_n)}\right]^2 < 1$ (10)

iii) $\phi'(0)$ can be chosen arbitrarily, e.g., $\phi'(0) = 1$

iv) $x(\phi) = \lim_{k\to\infty} a^k[g^{-k}(\phi)-\xi_n]$

$\phi(x) \underset{x\to\infty}{\approx} 1-(n)^{-n/n-1} \exp[-(x)^{-\ln n/\ln a}]$

$\phi(x) \underset{x\to-\infty}{\approx} (n)^{-1/n-1} \exp[-(x)^{-\ln n/\ln a}]$

However, only properties (9) and (10) will play a role in our analysis. Most significantly, parameter a, which is an implicit function of n, remains lower than 1 for all n.

Substituting equations (4), (5), and (6) into (1) and considering, without loss of generality, the case where d is an even integer, d = 2h, we obtain:

$$N_{n,d} = n^h + \int_{x=-\infty}^{\infty} \pi_h(x) \left(\sum_{i=1}^{h} \frac{r_i'(x)}{r_i(x)} \right) dx \qquad (11)$$

where:

$$\pi_h(x) = \prod_{j=0}^{h-1} p(x/a^i) \ , \qquad (12)$$

$$p(x) = r(x)\ s(x) = P[\phi(x)] \ , \qquad (13)$$

and:

$$P(\phi) = \frac{1-\phi^n}{1-\phi} \cdot \frac{1-(1-\phi^n)^n}{\phi^n} \ . \qquad (14)$$

Using equations (5) and (7), it can be easily shown that $r_i'(x)/r_i(x)$ satisfies:

$$\frac{r_i'(x)}{r_i(x)} \le \frac{n(n-1)}{2} \quad \phi'(x/a^{i-1})\ 1/a^{i-1} \qquad (15)$$

and consequently, (11) becomes:

$$N_{n,d} \le n + \frac{n(n-1)}{2} \int_{-\infty}^{\infty} \pi_h(x)\ [\sum_{i=1}^{h} \phi'\ (x/a^{i-1})\ 1/a^{i-1}]\ dx \qquad (16)$$

We now wish to bound the term $\pi_h(x)$ from above. An examination of $p(x) = P[\phi(x)]$ (equations (13) and (14)) reveals that p(x) is unimodal in x, $p(0) = [\xi_n/1-\xi_n]^2$, and that p(x) lies above the asymptotes $p(-\infty) = p(+\infty) = n$. Moreover, the maximum of $P(\phi)$ occurs below $\phi = \xi_n$ and, consequently, p(x) attains its maximum, M_n, below x = 0.

At this point, were we to use the bound $\pi_h(x) \le M_n{}^h$ in (16), it would result in $N_{n,d} < n^h + \frac{n(n-1)h}{2} M_n{}^h$ and lead to Baudet's bound $\mathscr{R}_{\alpha-\beta} \le M_n{}^{1/2}$. Instead, a tighter bound can be established by exploiting the unique relationships between the factors of $\pi_h(x)$.

<u>Lemma 1</u>: Let $x_0 < 0$ be the unique negative solution of $p(x_0) = p(0)$. $\pi_h(x)$ attains its maximal value in the range $a^{h-1}x_0 \le x \le 0$.

<u>Proof</u>: Since $p(x)$ is unimodal we have $p(x) < p(0)$ and $p'(x) > 0$ for all $x < x_0$. Consequently, for all $x < x_0$, any decrease in the magnitude of $|x|$ would result in increasing $p(x)$, i.e., $p(cx) > p(x)$ for all $0 \le c < 1$. Now Consider $\pi_h(ax)$:

$$\pi_h(ax) = p(x/a^{h-2})\, p(x/a^{h-3}) \ldots p(x)\, p(ax)$$
$$= \pi_h(x)\, p(ax)/p(x/a^{h-1}) \;;$$

for all x' satisfying $x'/a^{h-1} < x_0$ we must have $p(ax') > p(x'/a^{h-1})$ (using $c=a^h<1$) and $\pi_h(ax') > \pi_h(x')$, implying that $\pi_h(x')$ could not be maximal. Consequently, for $\pi_h(x')$ to be maximal, x' must be in the range $x_0 a^{h-1} \le x' \le 0$.

<u>Lemma 2</u>: $\pi_h(x)$ can be bounded by:

$$\pi_h(x) \le A(n)\, [p(0)]^h \tag{17}$$

where $A(n)$ is a constant multiplier independent on h.

<u>Proof</u>: Since $p(x)$ is continuous, there exists a constant α such that $p(x) \le p(0) - \alpha x$ for all $x \le 0$. Consequently, using Lemma 1, we can write:

$$\max_x \pi_h(x) = \max_{a^{h-1}x_0 \le x \le 0} \pi_h(x) \le \max_{a^{h-1}x_0 \le x \le 0} \prod_{i=0}^{h-1} (p(0) - \alpha x/a^i)$$

$$\le [p(0)]^h \max_{a^{h-1}x_0 \le x \le 0} \exp\left(\sum_{i=0}^{h-1} - \frac{\alpha x}{a^i p(0)} \right)$$

$$= [p(0)]^h \exp\left[\frac{-\alpha x_0}{p(0)}\, a^{h-1} \sum_{i=0}^{h-1} 1/a^i \right]$$

$$\le [p(0)]^h \exp\left[\frac{-\alpha x_0}{p(0)\,(1-a)} \right]$$

Selecting $A(n) = \exp\left[\frac{-\alpha x_0}{p(0)\,(1-a)} \right]$ proves the Lemma.

Theorem 2: The branching factor of the α-β procedure for a uniform tree of degree n is given by:

$$\mathcal{R}_{\alpha-\beta} = \frac{\xi_n}{1-\xi_n} \tag{18}$$

where ξ_n is the positive root of the equation $x^n+x-1 = 0$.

Proof: Substituting (17) in (16) yields:

$$N_{n,d} \le n^h + \frac{n(n-1)}{2} \; A(n) \; [P(0)]^h \int_{-\infty}^{\infty} \sum_{i=0}^{h-1} (1/a^i) \; \phi'(x/a^i) \; dx$$

$$= n^h + \frac{n(n-1)}{2} \; A(n) \; [P(0)]^h \; h$$

Finally, using $p(0) = (\xi_n/1-\xi_n)^2 > n$, we obtain:

$$\mathcal{R}_{\alpha-\beta} = \lim \; (N_{n,d})^{1/2h} \le \xi_n/1-\xi_n \tag{19}$$

This, together with Baudet's lower bound $\mathcal{R}_{\alpha-\beta} \ge \xi_n/1-\xi_n$, completes the proof of Theorem 2.

Corollary: The α-β procedure is asymptotically optimal over the class of directional game-searching algorithms.

The corollary follows from (18) and the fact that $\xi_n/1-\xi_n$ lower bounds the branching factor of any directional algorithm [5].

3. CONCLUSIONS AND OPEN PROBLEMS

The asymptotic behavior of $\mathcal{R}_{\alpha-\beta}$ is $O(n/\log n)$, as predicted by Knuth's analysis [1]. However, for moderate values of n ($n \le 1000$) $\xi_n/1-\xi_n$ is fitted much better by the formula $(.925)n^{.747}$ (see Figure 4 of reference [5]) which vindicates the simulation results of Fuller et al. [3]. This approximation offers a more meaningful appreciation of the pruning power of the α-β algorithm. Roughly speaking, a fraction of only $(.925)n^{.747}/n \approx n^{-1/4}$ of the legal moves will be explored by α-β. Alternatively, for a given search time allotment, the α-β pruning allows the search depth to be increased by a factor $\log n/\log \mathcal{R}_{\alpha-\beta} \approx 4/3$ over that of an exhaustive minimax search.

The establishment of the precise value of $\mathscr{R}_{\alpha-\beta}$ for continuous-valued trees, together with a previous result that $\mathscr{R}_{\alpha-\beta} = n^{1/2}$ for almost all discrete-valued trees [5], resolve two major uncertainties regarding the asymptotic behavior of α-β. However, the global optimality of α-β remains an unresolved issue. Naturally, the focus of attention now turns to non-directional algorithms, raising the question of whether any such algorithm exists which exhibits a branching factor lower than $\xi_n/1-\xi_n$.

Recently, Stockman [8] has introduced a non-directional algorithm which examines fewer nodes than α-β. The magnitude of this improvement has not been evaluated yet, and it is not clear whether the superiority of Stockman's algorithm reflects a reduced branching factor or merely a marginal improvement at low h's which disappears on taller trees. The latter seems more likely.

Notably, the problem of determining the existence of an algorithm superior to α-β can be reduced to the simpler problem of finding a superior algorithm for searching a standard bi-valued tree, i.e., a tree for which the terminal nodes are assigned the value 1 and 0 with probability ξ_n and $1-\xi_n$, respectively [5]. Unfortunately, even this reduced problem currently seems far from solution.

REFERENCES

[1] Knuth, D. E., and Moore, R. N. An analysis of alpha-beta pruning. Artificial Intelligence 6 (1975), 293-326.

[2] Slagle, J. R., and Dixon, J. K. Experiments with some programs that search game trees. Journal of the ACM 2 (1969), 189-207.

[3] Fuller, S. H., Gaschnig, J. G., and Gillogly, J. J. An analysis of the alpha-beta pruning algorithm. Department of Computer Science Report, Carnegie-Mellon University (1973).

[4] Baudet, G. M. On the branching factor of the alpha-beta pruning algorithm. Artificial Intelligence 10 (1978), 173-199.

[5] Pearl, J. Asymptotic properties of minimax trees and game-searching procedures. Artificial Intelligence 14 (1980), 113-138.

[6] Pearl, J. A space-efficient on-line method of computing quantile estimates. UCLA-ENG-CSL-8018, University of California, Los Angeles (1980). To appear in Journal of Algorithms.

[7] Kuczma, M. Functional Equations in a Single Variable. Polish Scientific Publishers, Warszawa (1968), p. 141.

[8] Stockman, G. A minimax algorithm better than alpha-beta? Artificial Intelligence 12 (1979), 179-196.