

Alpha-beta pruning

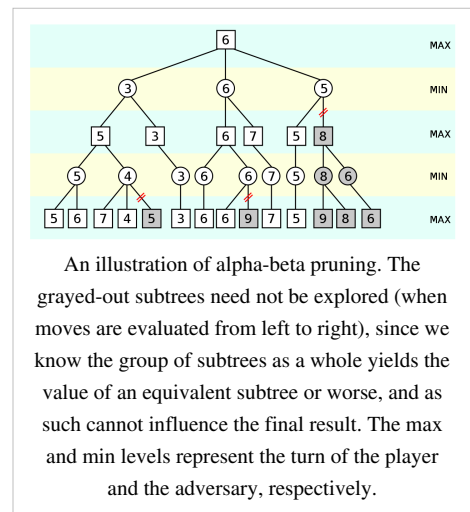
Alpha-beta pruning is a search algorithm which seeks to reduce the number of nodes that are evaluated by the minimax algorithm in its search tree. It is a search with adversary algorithm used commonly for machine playing of two-player games (Tic-tac-toe, Chess, Go, etc.). It stops completely evaluating a move when at least one possibility has been found that proves the move to be worse than a previously examined move. Such moves need not be evaluated further. *Alpha-beta pruning is a sound optimization in that it does not change the score of the result of the algorithm it optimizes.*

History

Allen Newell and Herbert Simon who used what John McCarthy calls an "approximation"^[1] in 1958 wrote that alpha-beta "appears to have been reinvented a number of times".^[2] Arthur Samuel had an early version and Richards, Hart, Levine and/or Edwards found alpha-beta independently in the United States.^[3] McCarthy proposed similar ideas during the Dartmouth Conference in 1956 and suggested it to a group of his students including Alan Kotok at MIT in 1961.^[4] Alexander Brudno independently discovered the alpha-beta algorithm, publishing his results in 1963.^[5] Donald Knuth and Ronald W. Moore refined the algorithm in 1975^[6] [7] and it continued to be advanced.

Improvements over native minimax

The benefit of alpha-beta pruning lies in the fact that branches of the search tree can be eliminated. The search time can in this way be limited to the 'more promising' subtree, and a deeper search can be performed in the same time. Like its predecessor, it belongs to the branch and bound class of algorithms. The optimization reduces the effective depth to slightly more than half that of simple minimax if the nodes are evaluated in an optimal or near optimal order (best choice for side on move ordered first at each node).



With an (average or constant) branching factor of b , and a search depth of d plies, the maximum number of leaf node positions evaluated (when the move ordering is pessimal) is $O(b*b*...*b) = O(b^d)$ – the same as a simple minimax search. If the move ordering for the search is optimal (meaning the best moves are always searched first), the number of leaf node positions evaluated is about $O(b*1*b*1*...*b)$ for odd depth and $O(b*1*b*1*...*1)$ for even depth, or $O(b^{d/2}) = O(\sqrt{b^d})$. In the latter case, where the ply of a search is even, the effective branching factor is reduced to its square root, or, equivalently, the search can go twice as deep with the same amount of computation.^[8] The explanation of $b*1*b*1*...$ is that all the first player's moves must be studied to find the best one, but for each, only the best second player's move is needed to refute all but the first (and best) first player move – alpha-beta ensures no other second player moves need be considered. If $b=40$ (as in chess), and the search depth is 12 plies, the ratio between optimal and pessimal sorting is a factor of nearly 40^6 or about 4 billion times.

Normally during alpha-beta, the subtrees are temporarily dominated by either a first player advantage (when many first player moves are good, and at each search depth the first move checked by the first player is adequate, but all second player responses are required to try and find a refutation), or vice versa. This advantage can switch sides

many times during the search if the move ordering is incorrect, each time leading to inefficiency. As the number of positions searched decreases exponentially each move nearer the current position, it is worth spending considerable effort on sorting early moves. An improved sort at any depth will exponentially reduce the total number of positions searched, but sorting all positions at depths near the root node is relatively cheap as there are so few of them. In practice, the move ordering is often determined by the results of earlier, smaller searches, such as through iterative deepening.

The algorithm maintains two values, alpha and beta, which represent the minimum score that the maximizing player is assured of and the maximum score that the minimizing player is assured of respectively. Initially alpha is negative infinity and beta is positive infinity. As the recursion progresses the "window" becomes smaller. When beta becomes less than alpha, it means that the current position cannot be the result of best play by both players and hence need not be explored further.

Additionally, this algorithm can be trivially modified to return an entire principal variation in addition to the score. Some more aggressive algorithms such as MTD(f) do not easily permit such a modification.

Pseudocode

```

function alphabeta(node, depth,  $\alpha$ ,  $\beta$ , Player)
  if depth = 0 or node is a terminal node
    return the heuristic value of node
  if Player = MaxPlayer
    for each child of node
       $\alpha := \max(\alpha, \text{alphabeta}(\text{child}, \text{depth}-1, \alpha, \beta, \text{not}(\text{Player}))$ 
      if  $\beta \leq \alpha$ 
        break (* Beta cut-off *)
    return  $\alpha$ 
  else
    for each child of node
       $\beta := \min(\beta, \text{alphabeta}(\text{child}, \text{depth}-1, \alpha, \beta, \text{not}(\text{Player}))$ 
      if  $\beta \leq \alpha$ 
        break (* Alpha cut-off *)
    return  $\beta$ 
  (* Initial call *)
  alphabeta(origin, depth, -infinity, +infinity, MaxPlayer)

```

Heuristic improvements

Further improvement can be achieved without sacrificing accuracy, by using ordering heuristics to search parts of the tree that are likely to force alpha-beta cutoffs early. For example, in chess, moves that take pieces may be examined before moves that do not, or moves that have scored highly in earlier passes through the game-tree analysis may be evaluated before others. Another common, and very cheap, heuristic is the killer heuristic, where the last move that caused a beta-cutoff at the same level in the tree search is always examined first. This idea can be generalized into a set of refutation tables.

Alpha-beta search can be made even faster by considering only a narrow search window (generally determined by guesswork based on experience). This is known as *aspiration search*. In the extreme case, the search is performed with alpha and beta equal; a technique known as *zero-window search*, *null-window search*, or *scout search*. This is particularly useful for win/loss searches near the end of a game where the extra depth gained from the narrow window and a simple win/loss evaluation function may lead to a conclusive result. If an aspiration search fails, it is straightforward to detect whether it failed *high* (high edge of window was too low) or *low* (lower edge of window

was too high). This gives information about what window values might be useful in a re-search of the position.

Other algorithms

More advanced algorithms that are even faster while still being able to compute the exact minimax value are known, such as Negascout and MTD-f.

Since the minimax algorithm and its variants are inherently depth-first, a strategy such as iterative deepening is usually used in conjunction with alpha-beta so that a reasonably good move can be returned even if the algorithm is interrupted before it has finished execution. Another advantage of using iterative deepening is that searches at shallower depths give move-ordering hints that can help produce cutoffs for higher depth searches much earlier than would otherwise be possible.

Algorithms like SSS*, on the other hand, use the best-first strategy. This can potentially make them more time-efficient, but typically at a heavy cost in space-efficiency.

References

- George T. Heineman, Gary Pollice, and Stanley Selkow (2008). "Chapter 7:Path Finding in AI". *Algorithms in a Nutshell*. O'Reilly Media. pp. 217–223. ISBN 978-0-596-51624-6.
- [1] McCarthy, John (LaTeX2HTML 27 November 2006). "Human Level AI Is Harder Than It Seemed in 1955" (<http://www-formal.stanford.edu/jmc/slides/wrong/wrong-sli/wrong-sli.html>). . Retrieved 2006-12-20.
- [2] Newell, Allen and Herbert A. Simon (March 1976). "Computer Science as Empirical Inquiry: Symbols and Search" (http://archive.computerhistory.org/projects/chess/related_materials/text/2-3.Computer_science_as_empirical_inquiry/2-3.Computer_science_as_empirical_inquiry.newell_simon.1975.ACM.062303007.pdf) (PDF). *Communications of the ACM, Vol. 19, No. 3*. . Retrieved 2006-12-21.
- [3] Edwards, D.J. and Hart, T.P. (4 December 1961 to 28 October 1963). "The Alpha-Beta Heuristic (AIM-030)" (<http://hdl.handle.net/1721.1/6098>). Massachusetts Institute of Technology. . Retrieved 2006-12-21.
- [4] Kotok, Alan (XHTML 3 December 2004). "MIT Artificial Intelligence Memo 41" (http://www.kotok.org/AI_Memo_41.html). . Retrieved 2006-07-01.
- [5] Marsland, T.A. (<http://www.cs.ualberta.ca/~tony/>) (May 1987). "Computer Chess Methods (PDF) from Encyclopedia of Artificial Intelligence. S. Shapiro (editor)" (<http://www.cs.ualberta.ca/~tony/OldPapers/encyc.mac.pdf>) (PDF). J. Wiley & Sons. pp. 159–171. . Retrieved 2006-12-21.
- [6] * Knuth, D. E., and Moore, R. W. (1975). "An Analysis of Alpha-Beta Pruning" ([http://www.eecis.udel.edu/~ypeng/articles/An Analysis of Alpha-Beta Pruning.pdf](http://www.eecis.udel.edu/~ypeng/articles/An%20Analysis%20of%20Alpha-Beta%20Pruning.pdf)). *Artificial Intelligence Vol. 6, No. 4*: 293–326. .
 - Reprinted as Chapter 9 in Knuth, Donald E. (2000). *Selected Papers on Analysis of Algorithms* (<http://www-cs-faculty.stanford.edu/~knuth/aa.html>). Stanford, California: Center for the Study of Language and Information - CSLI Lecture Notes, no. 102. ISBN 1-57586-212-3. OCLC 222512366. .
- [7] Abramson, Bruce (June 1989). "Control Strategies for Two-Player Games" (<http://www.theinformationist.com/pdf/constrat.pdf/>). *ACM Computing Surveys, Vol. 21, No. 2* 21: 137. doi:10.1145/66443.66444. . Retrieved 2008-08-20.
- [8] Russell, Stuart J.; Norvig, Peter (2003), *Artificial Intelligence: A Modern Approach* (<http://aima.cs.berkeley.edu/>) (2nd ed.), Upper Saddle River, New Jersey: Prentice Hall, ISBN 0-13-790395-2,

External links

- http://www.emunix.emich.edu/~evett/AI/AlphaBeta_movie/sld001.htm
- http://sern.ualgary.ca/courses/CPSC/533/W99/presentations/L1_5B_McCullough_Melnyk/
- http://sern.ualgary.ca/courses/CPSC/533/W99/presentations/L2_5B_Lima_Neitz/search.html
- <http://www.maths.nott.ac.uk/personal/anw/G13GAM/alphabet.html>
- <http://chess.verhelst.org/search.html>
- <http://www.frayn.net/beowulf/index.html>
- <http://hal.inria.fr/docs/00/12/15/16/PDF/RR-6062.pdf>
- Minimax (with or without alpha-beta pruning) algorithm visualization - game tree solving (Java Applet) (<http://wolfey.110mb.com/GameVisual/launch.php?agent=2>)

Article Sources and Contributors

Alpha-beta pruning *Source:* <http://en.wikipedia.org/w/index.php?oldid=417219192> *Contributors:* Albmont, Alexrakia, AllUltima, AndrewHowse, Arvindn, Asqueella, Avatar, Bad Romance, Bart Massey, Beland, Betacommand, Bfinn, Blahedo, BlueNovember, Bumbulski, CharlesGillingham, Circular17, Cobi, Dav!dB, David Haslam, Delirium, Doradus, Dsnpost, Eptalon, Eric Le Bigot, Euryalus, Evand, Fresheneesz, Giftlite, GregorB, Hadal, Heineman, Hiram 99, IanOsgood, Intgr, Janzert, Jecraig@yahoo.com, Jez9999, KJS77, Libor Vilímek, Lightmouse, Luke Gustafson, Makaimc, Maschelos, MathsIsFun, Mav, Melcombe, Michael Hardy, Michael Slone, Msikma, Nelhage, Newfraferz87, Nomis80, Onorem, Pete.Hurd, Pitel, Plattler01, Qwertyus, Rasmus Faber, Rjwilmsi, ScudLee, Sophus Bie, Stephen B Streater, Stephen Morley, Steven Jones, The Anome, Thesilverbail, Thijswijs, Truehalley, Veganfanatic, Willtron, Wolfkeeper, Zophar1, 135 anonymous edits

Image Sources, Licenses and Contributors

Image:AB pruning.svg *Source:* http://en.wikipedia.org/w/index.php?title=File:AB_pruning.svg *License:* GNU Free Documentation License *Contributors:* Original uploader was Jez9999 at en.wikipedia

License

Creative Commons Attribution-Share Alike 3.0 Unported
<http://creativecommons.org/licenses/by-sa/3.0/>