

THE BANDWIDTH HEURISTIC SEARCH

Larry R. Harris *
Department of Mathematics
Dartmouth College
Hanover, New Hampshire 03755

Abstract

By placing various restrictions on the heuristic estimator it is possible to constrain the heuristic search process to fit specific needs. This paper introduces a new restriction upon the heuristic, called the "bandwidth" condition, that enables the ordered search to better cope with time and space difficulties. In particular, the effect of error within the heuristic is considered in detail.

Beyond this the bandwidth condition quite naturally allows for the extension of the heuristic search to MIN/MAX trees. The resulting game playing algorithm affords many desirable practical features not found in minimax based techniques, as well as maintaining the theoretical framework of ordered searches. The development of this algorithm provides some additional insight to the general problem of searching game trees by showing that certain, somewhat surprising changes in the cost estimates are required to properly search the tree. Furthermore, the use of an ordered search of MIN/MAX trees brings about a rather provocative departure from the conventional approach to computer game playing.

KEYWORDS: Artificial Intelligence, tree searching, game playing, heuristic search, alpha-beta minimax, game trees, MIN/MAX trees.

Introduction

The heuristic search (Hart1, Nilsson) has become an important aspect of artificial intelligence because of its wide area of application (Nilsson2, Pohl3) and because of the theoretical framework on which it is based. This framework is in large part due to various restrictions imposed upon the heuristic that guides the search and the resulting effect on the search algorithm itself.

In order to discuss some of these restrictions it is necessary to introduce the following notation. For a node n of a tree or graph, the following functions are defined as part of the problem.

- k(m,n) = the arc cost from node m to node n.
g(n) = the minimum arc cost from the root to node n.
h(n) = the minimum arc cost from node n to a goal.
f(n)=g(n)+h(n) the minimum arc cost from the root to a goal via node n.

* research partially sponsored by Office of Naval Research contract ONR N00014-73-A-0260

Since some of these functions are not known during the actual search, we must be guided by estimates of these functions.

- g*(n) = an estimate of g(n). NOTE: g'(n)«g(n) for a tree.
h'(n) = an estimate of h(n). This is the "heuristic".
f'(n)=g'(n)+h'(n) an estimate of f(n).

At each step of the heuristic search the most promising node, the one with minimum f'(n) value, is expanded. In this way the heuristic h' orders the search. Using Nilsson's terminology any node which has been expanded is called CLOSED, and any node not yet expanded is labeled OPEN. The open nodes are maintained on a list called the OPEN list. With this notion of the heuristic search, we can see the effect of certain restrictions placed upon the heuristic.

The most important restriction on the heuristic h' is the "admissibility" condition (Hart1, Nilsson2): If h'(n) < h(n) for all nodes n, then the heuristic search will always find a minimal cost (optimal) goal.

Another important restriction on the heuristic is the "consistency" condition (Nilsson^): If h'(m)-h'(n) < k(m,n) then all nodes expanded by the heuristic search have g'(n)=g(n) even if the graph is not a tree. This allows for considerable simplification of the algorithm.

This paper introduces another restriction on the heuristic, which allows the heuristic search to better cope with the practical problems of time and space.

The "Bandwidth" Condition

When using the heuristic search to solve complex problems, both time and space difficulties arise because the number of nodes in the search tree increases exponentially. The time required to search the tree or graph can be reduced only by using a heuristic that more closely estimates h(). It may be impractical to find such a heuristic that never violates the "admissibility" condition. For many problems, an example of which will be given shortly, it is often easier to find a heuristic that estimates the future cost well, but occasionally overestimates it. In these cases it is reasonable to make use of this inadmissible heuristic to order the search and then to choose from one of two alternatives.

1) Bound the extra cost of the resulting goal by the extra cost in overestimating $h()$. We may consider accepting a non-optimal solution, if it is much easier to find than an optimal one.

2) Continue the search beyond the expansion of a goal, reducing the added cost bound until an optimal goal is found.

In either case the search process will run faster than an admissible search since the new heuristic is a more accurate estimator of $h(>)$.

The "bandwidth" condition requires that:

$$h(n)-d < h'(n) < h(n)+e$$

The above constraints, one of which is a loosening of the admissibility condition, can be used to provide a theoretically sound method of coping with both the space and time difficulties encountered in solving large problems. We also assume that the bandwidth heuristic is exact for any goal. It is clear that when $e=0$ and $d>h(n)$ for all nodes that the bandwidth condition reduces to the admissibility condition. The bandwidth condition is similar to Pohl's⁴ "bounded error". However, our approach will be to study the relation between this error bound as applied to each node of the search and how it effects the cost of the goal found by the search. We will also demonstrate that bandwidth heuristics can be used effectively to find optimal goals, in spite of the possibility of over-evaluation.

The upper bound allows us to compute the extra cost of a non-optimal goal found by using $h'()$ to guide the search. We can bound this added cost (Harris 5) by noting that some node n^* on the path to the optimal goal is OPEN when the non-optimal goal p was selected for expansion. Thus,

$$\begin{array}{ll} f'(p) < f'(n^*) & p \text{ chosen before } n^* \\ g(p)+h(p) < g(n^*)+h'(n^*) & \text{definition of } f'() \\ g(p)+h(p) < g(n^*)+h'(n^*) & \text{since } h'(p)=h(p)=0 \\ q(p)+h(p) < g(n^*)+h(n^*)+e & \text{bandwidth bound} \\ f(p) < f(n^*)+e & \text{definition of } f() \\ f(p) < f(p^*)+e & n^* \text{ on path to } p^* \end{array}$$

providing the desired bound on the added cost of goal p . Such a goal is called e-optimal and the algorithm e-admissible.

Knowing this bound allows us to use heuristics that may not satisfy the admissibility condition without fear of encountering some degenerate case that could add an arbitrary additional cost to the resulting goal. In this sense bandwidth heuristics are easier to find than admissible heuristics since they are allowed to overestimate the true future cost.

The lower bound of the bandwidth condition allows us to save space by dropping nodes from the open list without surrendering the admissibility of the algorithm. We can eliminate any nodes m from the search if there exists a node q that satisfies the cutoff condition:
 $f(q) < f'(m)-(e+d)$.

It is clear from the bounding argument above that we must insure that the path to the optimal goal will not be dropped from the search tree. Assume that it is; then for some node q we have the following.

$$\begin{array}{ll} f'(q) < f'(n^*)-(e+d) & \text{the drop criterion} \\ f(q)-d < f(n^*)+e-(e+d) & \text{bandwidth bounds} \\ f(q) < f(n^*) & \text{collect terms} \\ f(q) < f(p^*) & n^* \text{ on path to } p^* \end{array}$$

Yielding a contradiction since p^* is an optimal goal. We can determine a bound on the extra cost for a goal p by comparing $f(p)$ to $f'(s)-e$ for all open nodes s , since one of these nodes is the n^* used above. If we are not satisfied with this result, we can continue until the bound is lowered to a desirable limit.

Knowing this bound, we may be satisfied with the e-optimal goal, or we may wish to continue the search for the optimal goal. This can be done since all open nodes will eventually have estimated cost greater than $f(p^*)+e$ where p^* is the minimum cost goal expanded by the search. In either case the use of a bandwidth heuristic speeds the search by better estimating future cost, and allows the dropping of obviously bad open nodes to conserve storage.

For many problems it may be impractical to find a single heuristic that can adequately fulfill both of the bandwidth criteria at one time. In these cases we can define two heuristics, one to satisfy each bound.

$$h(n)-d < h_2'(n) \quad \& \quad h_1'(n) < h(n)+e$$

For other cases we may wish e and d to themselves be functions of the node n . The extension to these other cases is straightforward.

Example: The Traveling Salesman Problem

As an example of the practical value of bandwidth heuristics consider the Traveling Salesman Problem of finding the minimum mileage tour of a set of cities. Since the search tree for an N -city problem has $\{N-DI\}$ nodes it is impractical to find the optimal tour for large N . Recent ad hoc techniques $\{Lin 6\}$ find "good" solutions but can estimate the error only empirically.

The following heuristics can be used to find "good" (e-optimal) solutions. All of the heuristics make use of the reduced mileage matrix defined by the partial tour representing a node. For this problem it is convenient to use two heuristics h_1' & h_2' to fulfill the bandwidth condition. An admissible heuristic h_0' will be used to motivate the definitions of h_1' and h_2' .

FOR A 6-CITY PROBLEM PARTIAL TOUR: 1 3
 CITIES TO LEAVE: 3 2 4 5 6
 CITIES TO ENTER: 2 4 5 6 1

REDUCED MILEAGE MATRIX

	ENTER				
	2	4	5	6	1
L	3	5	8	3	9
E	2	X	4	6	9
A	4	4	X	3	7
V	5	6	3	X	9
E	6	9	7	9	X

$$hO' - \max(3+4+2+3+4, 4+3+3+7+2) = 19$$

An admissible heuristic hO' would be the maximum of the sums of the row and column minima. The heuristic hO' is admissible since any tour must use one element from each row and column and can do no better than the minimum element. Thus $hO'(n) < h(n)$ for all nodes n . A more accurate heuristic would be one that accounts for the likelihood that the minimum elements of all the rows cannot all be used to form a tour. This can be done by taking the weighted average of the two smallest elements in each row.

$$hi = w(3+4+2+3+4) + (1-w)(5+5+3+5+7)$$

This heuristic gives some weight to the alternative choices. If it is indeed possible to use only the minimum elements, then this heuristic will overestimate the true future cost. However for the majority of nodes it will provide a much more accurate estimate of the future cost and thus provide a better guide for the search. It is because this heuristic provides a more realistic estimate of $h()$ that the search algorithm runs faster than the admissible search. Of course the resulting solution may be more costly, but we can bound this added cost by knowing how far we can overestimate a particular node.

The heuristic $h2'$ is an estimate on how big $h(n)$ can get for a particular node n . For the Traveling Salesman Problem, given any partial tour, we need only estimate the cost of the rest of the tour from this point to calculate this bound. Using the reduced cost matrix we could simply sum the major diagonal, or for a lower and therefore better estimate we could take the minimum choice from each city until a tour is complete. Each of these will have $d=0$ since we know the minimum future cost must be lower than or equal to the cost of a particular tour. Typically $h2'$ is only calculated when there is a need to drop nodes.

Thus for each node n , we know that we can do no worse than $h2'(n)$. Knowing this we can drop any node m if there exists a node q such that $f2^*(q) - f1^*(m) < (e+d)$ since any tour through m must cost more than a tour through q .

Searching Game Trees

Given the success of ordered searches in so many domains it seems reasonable to extend the heuristic search to game playing. The advantages of ordering the search should be as apparent in game playing as in the Traveling Salesman Problem. Unfortunately special problems occur which require significant changes in the cost functions in order to maintain certain desirable properties.

We would like to be able to make the optimal move from a given board configuration. This would require being able to prove that we can win or tie from an arbitrary board configuration, since someone must be able to force a win or tie in a 2-person 0-sum game. It is not that we cannot devise algorithms to perform this proof (Nilsson²), but that for any interesting game the algorithms are impractical.

If we develop admissible algorithms for game trees, such as Nilsson's ordered search procedure for AND/OR trees, critics respond that such an algorithm is impractical for any interesting game such as checkers or chess. On the other hand, algorithms that run in a reasonable amount of time are criticized because they don't always make the "best" move, and that they lack the theoretical framework of the admissible algorithms.

The use of a bandwidth heuristic allows us to work somewhere between these two extremes. We begin by reducing our aim in two ways; first, by looking for an $(e+d)$ -optimal goal instead of an optimal one and second, by looking only for the first move towards such a goal. The bandwidth heuristic search can achieve these more modest goals.

The standard technique for finding a good move from a current board position in 2-person games is the alpha-beta minimax (Samuel⁷) and its variations (Slagle⁸, Nilsson²). It is helpful to note some of the possible shortcomings of minimaxing to motivate an improved technique. The following detrimental factors arise when minimaxing.

1) The minimax technique implicitly assumes that the opponent is using the same evaluation of board positions as the program when it minimizes on alternate ply levels. This assumption must be made in order to determine what move the opponent will make at each of these levels. It is clear that this is an unwarranted assumption that can lead to serious errors in play.

2) Related to this problem is the fact that minimaxing makes permanent decisions on the basis of comparing board evaluations. If the evaluator mis-orders two boards the minimax will never recover. We are forced to admit that the evaluator must misorder some nodes or else we could play optimally with no tree search at all. In light of the inaccuracy of the board evaluator, it seems unwarranted to make irrevocable decisions based upon it.

3) Minimizing indicates what line of play looks most promising in a particular sub-tree, but does not provide an error estimate of how good the move is with respect to the entire move tree. It is desirable to provide such a definition of a "good" move.

4) since minimizing regenerates a portion of the move tree each time the program is to move there is little consistency in play from move to move. While substantial savings in computation can be obtained by ordering the generation of the tree (alpha-beta), there is little tendency to continue a particular line of attack on the next move because the new search is based solely upon values generated from 2 levels deeper in the tree. Since these values tend to vary somewhat from their ancestors the line of play suggested earlier may not be followed.

The bandwidth heuristic search provides a means of coping with each of these problems inherent to minimizing. Before presenting the algorithm in detail, some definitions are required. First we define the cost functions based upon one cost unit per move. For OPEN nodes of the search tree:

$g'(n) = g(n)$ cost from the root to node n
 $h(n) - d < h'(n) \leq h(n) + e$ estimate of future cost
 $f'(n) = h'(n) + g'(n)$ total estimated cost via node n

For terminal nodes of the search tree:
 $h'(WIN) = 0$ & $h'(LOSS) > N + (e+d)$

From this it is clear that $h'()$ estimates the number of moves to a win. Note that this means a low value is good for the program, the opposite of the minimax evaluator. Special care is taken when estimating losing nodes since $h'(LOSS)$ cannot be bounded. The value of N is arbitrary.

For CLOSED nodes of the search tree these functions are defined recursively.

node n on a MAX level:
 Define node m to be that son of n that is the best estimated move for the opponent. That is, $f'(m) \geq f'(k)$ for all sons k of n.

$$h'(n) = h'(m) \quad g'(n) = g'(m) \quad f'(n) = h'(n) + g'(n)$$

Node n on a MIN level:
 Define node m as the best estimated move for the program. That is, $f'(m) \leq f'(k)$ for all sons k of n.


$$h'(n) = h'(m) \quad g'(n) = g'(m) \quad \underline{f'(n) = h'(n) - g'(n)}$$

This last definition requires further explanation but is absolutely required to extend one of the desirable features of regular heuristic searches to MIN/MAX trees. A detailed discussion of this definition takes place in a following section.

A node n is said to be accessible if all the opponents moves on the path to n represent the best move for him in terms of true cost. That is, for each MIN node m on the path to n, $h(m) > h(b)$ for all brothers b of m.

The Bandwidth Heuristic Search Algorithm

A linked search tree is maintained using the following node specification.



Closed nodes of the search tree typically contain the 3 pointers plus values for $f'(n)$, $g'(n)$ and possibly $h'(n)$. Typically the linked list of brothers is sorted in increasing or decreasing order for MAX or MIN levels respectively. Tip nodes of the search tree must also contain the board configuration or a list of moves from which the board can be constructed.


The bandwidth heuristic search proceeds as follows.

- 1) Select a node p for expansion by following the first son chain to a tip node (the node p).
- 2) If p is a goal node (a Win), stop and make the 1st move towards p.
- 3) If p is a terminal node (a Loss), resign or play towards p expecting to lose.
- 4) Double expand p (on both a MIN and MAX level) using the drop criterion to avoid the needless generation of the grandsons of p (similar to the alpha beta cutoff). Note: a check must be made that a terminal node is not generated on the first expansion of p.
- 5) Evaluate $h'()$ and calculate $f'()$ for all grandsons of p. Add the generated nodes to the search tree.
- 6) Recalculate $f'(p)$ using recursive definition of $f'()$.
- 7) Recalculate $f'()$ for all ancestors of p and alter the brother lists to maintain the order.
- 8) Drop all nodes n (and the subtree below n) if for b a brother of n
 n on a MAX level and $f'(b) > f'(n) + (e+d)$
 n on a MIN level and $f'(b) < f'(n) - (e+d)$
- 9) Go to (1)

The error of the *final* result can be bounded in a manner similar to our earlier argument (Harris⁵). First it is shown that the path to the optimal accessible goal is not dropped from the search tree. The proof indicates that it cannot be dropped on a MAX level because it is optimal, and it cannot be dropped from a MIN level because it is accessible. When a goal p is selected for expansion it was at some level compared to the optimal path. The maximum inaccuracy in this comparison is one bandwidth or $(e+d)$. Thus we can bound the added cost introduced by the imprecision within the heuristic.

Alternating Definition of $f'()$

In the standard heuristic search the purpose of the cost function $g'()$ is to add a breadth-first component to the search. A particular path in the tree will continue to be expanded only if the decrease in $h'()$ outweighs the increase in g' . In order to keep this desirable tendency for the bandwidth search it is necessary to alter the common definitions for $h'()$, $g'()$ and $f'()$ for closed nodes. Consider the following case.



If $f'(n) = g'(n) + h'(n)$ for all levels in the tree then the left-most path would be expanded ad infinitum even though there was not an improvement from the opponent's point of view. We must force a breadth first component to the opponent's choices as well as the program's choices. To do this we must alter the fact that on MIN levels going deeper in the tree makes the opponent's moves look better instead of worse. This is clearly the opposite of the case for the program's choices, and of heuristic searches in general. In fact this breadth first tendency is critical to prove that any heuristic search is admissible (Nilsson²). The new definition $f'(n) = h'(n) - g'(n)$ on MIN levels will enable us to extend this needed tendency to the searching of MIN/MAX trees. The recursive definitions of $h'()$ and $g'()$ allow us to backup the knowledge gained at the tip nodes to the middle of the tree where it can be compared with intersecting paths. If we do not use this alternating definition of $f'()$ the search will tend to investigate only 1 or 2 of the opponent's moves at each level. Given the inaccuracies inherent to the heuristic, it is clear that we must provide for a more conservative search.

Finding Bandwidth Heuristics

Before considering the bandwidth heuristic search any further we must answer two questions which immediately come to mind. First, can we really find a bandwidth heuristic for a game such as chess or checkers? Second, do we really expect the algorithm to halt by expanding a goal each time the computer is to make a move?

The first question should be rephrased to ask "For what values of e & d can we find a bandwidth heuristic for a game such as

chess?" It is clear that for large e & d bandwidth heuristics can easily be found, and in fact the minimax board evaluators satisfy the conditions for some e & d .

For complex games such as chess it is unlikely that heuristics that can be proved to satisfy the bandwidth conditions can be found. In this case the e -admissibility of the algorithm becomes academic. However, the practical advantages of the bandwidth search process remain unaffected. The e and d then become parameters of the search to vary the density of the search tree.

The loosening of the admissibility condition to the bandwidth condition was originally meant to simplify the task of finding a suitable heuristic for a specific problem, not make it more difficult. The main distinction made by the bandwidth search is that it readily admits that there are inaccuracies in the heuristic, and that we must cope with these in a more subtle way than a single comparison of heuristic values.

The answer to the second question indicates the change in outlook implied by the bandwidth search. The algorithm is not meant to start and stop each move as is the case when minimaxing. The bandwidth heuristic search expands the move tree continuously, even while it is the opponent's turn to move. Of course as actual moves are made the unused sections of the search tree can be discarded. The algorithm only suggests moves when it is forced to stop because of time or space constraints. Under these conditions programs can play as well as current techniques with a near zero response time, since much of the tree search could take place while the opponent is thinking! The expansion of a goal must take place at some time during the game, but probably not before move decisions have to be made.

We must therefore devise a decision procedure for picking a move when the algorithm has not halted, but time or space constraints force us to make a move. The first guess one might have is to make the move towards the node to be expanded next. That is, the accessible node with the minimum f' value. However this is very dangerous since this node could be a misevaluated line of play that would be discarded immediately after its expansion. A safer technique would be to move towards the optimal accessible node as determined by $h'()$ values. This gives a strong bias toward lines of play that have been looked into deeply, since no cost is assigned to a move. It also allows the program to save more of the move tree for use next move, as well as keeping the program from making serious mistakes.

Let us now view the bandwidth search in light of our discussion of minimaxing. The bandwidth search, as would any ordered search, uses heuristic values as a means of guiding the search, not as an end in determining the best move. That is the error due to a misvaluation of two nodes is not critical, as with the minimax based algorithms, since the misvaluation can be corrected after further expansion. The

bandwidth search quite naturally follows forces or obvious lines of play to an arbitrary depth in the move tree. It is not restricted to an a priori depth bound.

Since the accessibility of a node is defined in terms of TRUE cost the bandwidth search does not assume the opponent evaluates boards in a manner similar to the program. The bandwidth search does assume that the opponent is trying to win the game in the fewest moves possible. This is a much more plausible assumption than the one mentioned above for minimaxing.

Since the bandwidth search refines its estimate of accessibility as the tree is expanded the misordering of nodes by the heuristic is not critical. The search may expand along the wrong path for a time but is capable of returning to expand along the correct path. Indeed it must return to examine the better line of play if the misordering is more than a bandwidth from the true values of the nodes.

A "good" move is defined as the first move towards an (e+d)-optimal accessible goal. If the cost of this expanded goal is less than N then the program must win from this position. If the value is greater than N+(e+d) then the program can expect to lose. However, if the value is in the "uncertainty zone" [N,N+(e+d)] as it will be most of the time, then it is unknown who can win and the game should be continued.

In essence, we are saying that the heuristic should be accurate on a coarse scale, but may not necessarily be accurate on a fine scale. That is, for h' 0 values sufficiently far apart the ordering of the heuristic should be the same as the h() ordering. Sufficiently far apart is, of course, defined as one bandwidth (e+d). On a fine scale we admit the inaccuracy of the heuristic and reserve judgement until the nodes have been expanded further. Thus, when h'() values appear in the "uncertainty zone", all is not lost, we must simply continue to expand the search tree until the heuristic values become more exact.

Empirical Results

In order to better isolate the effect of the bandwidth search/ extensive tests were conducted using the games of chess and Four Score, an interesting variant of three dimensional tic-tac-toe. In both cases existing programs that used a fixed ordered alpha-beta minimax were modified to use the bandwidth heuristic search. The new version then competed against the old program as well as human opponents. When two programs competed identical time and space constraints were imposed on each/ and both programs used the identical heuristic evaluator. since all the programs are deterministic it is only possible to play two unbiased games, one with each program going first.

For chess the bandwidth search noticeably improved the play of the program and easily won both games, playing white once and black once. At no time in either game


did the minimax program mount a serious offensive attack although its defense was acceptable. In order to more accurately assess the improvement in chess play a local chess master, rated at 2200, was asked to play each program until he had some feel for their ratings. His estimate was 1100-1200 for the mininax program, and 1400-1500 for the bandwidth search program.

The bandwidth search fared well in Four Score competition, winning while moving first and second. When playing first it was able to effect an 8 move combination in the middle game, which was beyond the horizon of the minimax, Four Score has no standardized rating scheme, thus no meaningful numeric comparison is possible.

Since it is known that minimax chess programs can exceed the 1200 rating of our program, it is clear that the chess heuristic used in this experiment is not the best possible. Since the bandwidth search makes more extensive use of the heuristic than does minimaxing, it seems intuitive, although unsupported, that the bandwidth search should improve at least as fast as the minimax search, for every improvement in the heuristic.

Double Expansion

It is possible to define an ordered search along the sane lines as the Bandwidth Search without using double expansion. The following case demonstrates a more subtle effect that would cause a single expansion algorithm to order the search incorrectly.



When the heuristic is overly sensitive to the last move, the single expansion search can actually reverse the initial ordering of the nodes. Move [2] will be considered next since it was expanded last, even though move [1] is the better play as initially indicated by the heuristic. This problem becomes more acute as the number of possible moves increases, since the entire list of moves can be reversed on a MIN level before progressing beyond the next MAX level.

Conclusion

The "bandwidth" conditions for the heuristic search provide a convenient means for coping with the practical constraints of time and space, as well as for maintaining the admissibility of the search.

For game trees the bandwidth conditions lead to an effective search procedure that has many desirable features over the conventional minimax techniques. Moreover, the bandwidth search brings about a totally different view of computer game playing. Instead of regenerating sections of the move tree each time the program is to move, the bandwidth search expands the move tree over a continuum of time. This allows for more coherent play from move to move, as well as the capability of "thinking" while the opponent is making his move. The bandwidth search minimizes the critical dependence upon the heuristic and eliminates the need for fixed, a priori limitations of the search process.

Another important aspect of this approach to game playing is the capability of the program to interact with a human partner in ordering the search of the move tree. Much of these areas remain open for future research.

The author wishes to acknowledge Professor Steve Garland and the following Dartmouth undergraduates for their help in the development and the testing of the bandwidth search: Warren Hon Monetgomery, Dave Chenerow, Dexter Knzen, and Steve Poulsen.

References

- Hart, P.E., Nilsson, N.J., Raphael, B., "A Formal Basis for the Heuristic Determination of Minimum Cost Paths", IEEE Transactions of Systems Sciences and Cybernetics, Vol SSC-4, No. 2, pp 100-107
2. Nilsson, N.J., Problem Solving Methods in Artificial Intelligence, New York: McGraw-Hill
 3. Pohl, I., "Heuristic Search Viewed as Path Finding in a Graph", Artificial Intelligence I, American Elsevier
 4. Pohl, I., "First Results on the effect of error in Heuristic Search", Machine Intelligence 5, 1969, pp 219-236
 5. Harris, L.R., "A Model for Adaptive Problem Solving Applied to Natural Language Acquisition", Technical Report TR-133, Computer Science Dept., Cornell University
 6. Lin, S., Kernighan, B.W., "A Heuristic Algorithm for the Traveling Salesman Problem", Bell Telephone Laboratories Computer Science Technical Report No. 1, April in, 1972
 7. Samuel, A.L., "Some Studies in Machine Learning using the game of Checkers", IBM Journal, v 3, no. 3, pp 210-229
 8. Slagle, J.R., "Game Trees, M & N Minimaxing and the M & N alpha-beta procedure", Lawrence Radiation Laboratory AI Report Mo. 3, November, Livermore, Calif.
 9. Lawler, E., Wood, D., "Branch & Bound Methods: A Survey", Operations Research (4) 14, July-August 1966, pp 699-715