Pathology in heuristic search

Mitja Luštrek

Jožef Stefan Institute
Department of Intelligent Systems
Jamova 39, 1000 Ljubljana, Slovenia
E-mail: mitja.lustrek@ijs.si

Practice shows that game-playing programs using minimax search perform better when searching deeper. Mathematical analyses, however, showed the opposite. This paradox was termed *minimax pathology*. Our real-valued minimax model demonstrates that appropriate modeling of the heuristic error is enough to eliminate the pathology. We examine the conditions under which the pathology appears and explain the mechanism that makes minimax otherwise beneficial. The reasons for the pathology in single-agent search are also addressed.

Keywords: heuristic search, pathology, minimax

1. Introduction

Most game-playing programs use minimax search: they build a game tree, heuristically evaluate the leaves, back up the heuristic values to the root and use them to select a move. Practice shows that the deeper such programs search, the better they play. However, mathematical models intended to explain this formally paradoxically showed that deeper search results in worse play [6,1]. This phenomenon was termed minimax pathology.

The most common explanation for the pathology is that the mathematical models lacked dependence between nearby positions, which is clearly present in games. In the Ph.D. thesis we propose an explanation that does not require dependence and uses real-number position values. Since game-playing programs use discrete values, we also investigate the effect of the number of possible values as well as other parameters on the pathology. Finally we return to the question that started pathology research: why is minimax search beneficial.

Pathology was recently observed in single-agent search as well [2]. We investigate the conditions under which it appears and show that it has a practical impact by observing it in path-finding on maps from commercial computer games.

2. Pathology in minimax search

Our minimax model is a real-valued version of Beal's [1] classic pathological model. We use real numbers for both true and heuristic position values and represent the error of the heuristic evaluation function as Gaussian noise. Beal's model uses losses and wins for position values and represents the heuristic error as the probability of mistaking a loss for a win or vice versa. The heuristic error is independent of the depth of search in both models and they are otherwise equivalent as well.

Despite the similarity between the models, Beal's is pathological and ours is not. The reason lies in the way heuristic error is modeled. Each move causes a limited change in position value, so if the position at the root of the game tree is undecided, positions close to it are also close to undecided, whereas positions lower down can be clearly lost or won. The way we model heuristic error, it is less likely that a clearly lost position is mistaken for a won one or vice versa. Consequently heuristic error, when viewed in terms of losses and wins, is smaller when searching deeper. We showed this is enough to counter the amplification of loss/win error through minimaxing observed by Beal and that no other properties of game trees are required [3].

The pathology is chiefly affected by three parameters. The first one is the number of possible position values. We know that two-valued minimax models tend to be pathological and that our real-valued model is not. If sufficiently few discrete values are used, the pathology also appears. This happens because minimaxing eliminates high and low values from the game tree, so if the initial number of values is small enough, after a few levels of minimaxing only two values are left. The next parameter is the branching factor of the game tree b. Since each level of minimaxing eliminates $b-1$ out of b nodes, their elimination is faster and pathology more likely when b is large. The final parameter is the dependence between nearby positions, which most consider the main reason for the absence of the pathology in games. We introduce dependence in a way that corresponds to chess and show that it prevents the pathology as expected [5].

3. Why is minimax beneficial

Fig. 1 shows probability density functions of the heuristic values of two sibling nodes f_L (lower) and f_H (higher) and their parent f_P (maximum in this case). The curve of f_P is narrower than the curves of f_L and f_H, which means that the parent deviates less from its true value t_P. To understand why, we must consider two cases. (1) The heuristic value of the parent is larger than t_P. The most likely reason is that the heuristic value of the higher sibling is larger than t_P, but it may also be caused by the lower sibling. (2) The heuristic value of the parent is smaller than t_P. The reason is that the heuristic value of the higher sibling is smaller than t_P, but this can be compensated by the heuristic value of the lower sibling being larger than the heuristic value of the higher sibling. So the effect of the lower sibling may be (1) harmful or (2) beneficial. Since its heuristic value is more likely in the area where it is beneficial, it will move the heuristic value of the parent closer to t_P. Every level of minimaxing has this effect, so deeper search moves heuristic values closer to the corresponding true values and thus reduces the error.
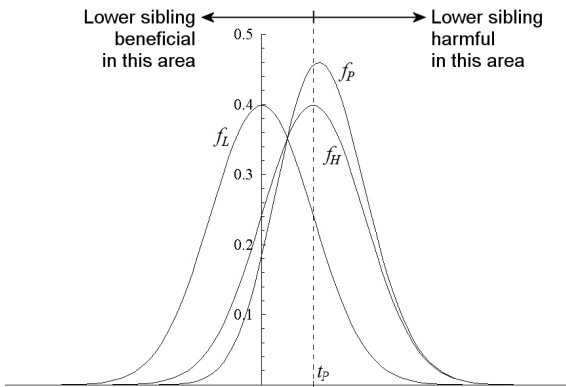

Fig. 1. Explanation of the benefit of minimax search.

4. Pathology in single-agent search

The pathology is affected by the heuristic evaluation function. These functions are usually designed to be admissible (i.e., optimistic) and monotonically non-decreasing (i.e., the heuristic value of a node is not smaller than the heuristic value of its parent). Our experiments on synthetic game trees showed that to avoid the pathology, the evaluation function must foremost be non-decreasing, but it is desirable that it is also admissible. This is somewhat surprising, because more attention is usually paid to admissability.

We also experimented with path-finding on maps from commercial computer games using LRTS algorithm [4]. We found that the pathology is common, but not very strong. It has two causes: (1) the learning component of LRTS and (2) the fact that during path-finding, LRTS searches only every d steps, where d is the depth of search. Learning turned out to be most effective at smaller depths and consequently smaller depths sometimes outperformed larger ones, which is pathological. Searching every d steps also brings different depths closer to each other, because even though searches to smaller depths are less effective, they are performed more frequently, and consequently smaller depths sometimes again outperformed larger ones.

5. Conclusion

In the thesis we provided a new, simple explanation for the minimax pathology and we quantified and explained the three major parameters that affect it. Moreover, we described the mechanism that makes minimax search beneficial. We also contributed to the previously very limited knowledge about the pathology in single-agent search.

The full thesis is available here: `http://dis.ijs.si/mitjal/supplements/PhD.pdf`.

References

[1] D.F. Beal, An analysis of minimax, *Advances in Computer Chess* **2** (1980), 103–109.

[2] V. Bulitko, Lookahead pathologies and meta-level control in real-time heuristic search, *Proceedings of 15th Euromicro Conference on Real-Time Systems* (2003), 13–16.

[3] M. Luštrek, I. Bratko and M. Gams, Why minimax works: An alternative explanation, *Proceedings of International Joint Conference on Artificial Intelligence* (2005), 212–217.

[4] M. Luštrek and V. Bulitko, Lookahead pathology in real-time path-finding, *Proceedings of National Conference for Artificial Intelligence, Learning for Search Workshop* (2006), 108–114.

[5] M. Luštrek, M. Gams and I. Bratko, Is real-valued minimax pathological?, *Artificial Intelligence* **170**(6–7) (2006), 620–642.

[6] D.S. Nau, *Quality of Decision versus Depth of Search on Game Trees*, Ph.D. thesis, Duke University, 1979.