Proceedings of the Sixteenth International Symposium on Combinatorial Search (SoCS 2023)

Structural Bias in Heuristic Search (Student Abstract)

Alison Paredes

Department of Computing, Queen's University, Canada
20asp3@queensu.ca

Abstract

In this line of work, we consider the possibility that some fast heuristic search methods introduce structural bias, which can cause problems similar to sampling bias for downstream statistical learning methods. We seek to understand the source of this kind of bias and to develop efficient alternatives. Here we present some preliminary results in developing a variation of canonical A* that can overcome the structural bias introduced by first-in-first-out duplicate detection, which we observed under the condition of variable heuristic error. These results inspire a model of greedy-best-first-search for this problem in the satisficing setting. We hope to apply our approach in a novel planning application–activity selection for agent-based modeling for epidemiology–where planning technology should avoid introducing structural bias if possible.

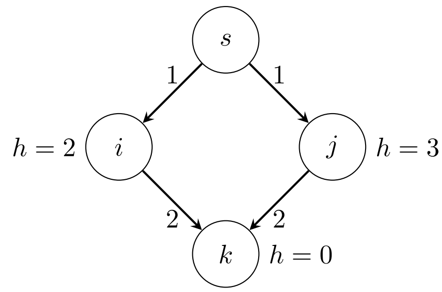

Figure 1: FIFO duplicate selection prefers the plan, among equally cost-optimal plans, with more heuristic error, node i, in the presence of variable heuristic error, where $h^*(i) = h^*(j)$ and $h(i) < h(j)$

Introduction

It is generally accepted that a valid solution to a given planning task should conform to the problem specification. It may be the case, however, that some solutions will remain out of reach of state-of-the-art planning technology without imposing some additional constraints on the solution space. Some such constraints might include a preference for shorter plans as might be imposed by the class of algorithms that use an estimate of the distance to the goal to influence search order (Thayer and Ruml 2011). Other constraints are less obvious such as a node generation order, sometimes known as accidents of implementation (Helmert and Röger 2008), and, although poorly understood, are suspected to contribute to better search speed. Tie-breaking may be one of the most well-known sources of structural bias in heuristic search (Asai and Fukunaga 2016). In this line of work, we question the assumption that any valid plan will do. Our work has applications anywhere one hopes to generate training data via planning technology (Ferber et al. 2022; Ernandes and Gori 2004).

Consider how easy it is to affect which plans get into the training set. We introduce an example caused by duplicate detection in the presence of variable heuristic error. Duplicate detection is a powerful variation of A* (Hart, Nilsson, and Raphael 1968) which achieves exponential improvements in space efficiency in the presence of planning tasks with many different cost-optimal paths to the same node. A typical implementation records the first time a state is expanded and afterward avoids generating any additional search nodes for this state. Assuming a consistent heuristic this procedure does not affect any cost-optimality guarantees on the solution; only which path from the goal node is recovered.

We observe that, in the presence of variable heuristic error, this path can be deterministic even in the presence of random tie-breaking (Figure 1). Assuming an admissible heuristic, although states i and j represent equal cost paths from s, state k will be generated via state i first, put on $OPEN$, and then, since $f(k) < f(j)$, expanded next. When the node representing state j is expanded, its successor node representing state k will be discarded, since k is $CLOSED$. The implication for downstream statistical learning methods is that training data may be skewed toward examples with poorer heuristic estimates.

Problem Definition

Given a classical planning task, which describes a graph composed of nodes n with a cost function assigning a positive value $c(n_i, n_j)$ between each node n_i and its successor n_j. Additionally, define $P^*_{n_i - n_g}$, the set of minimal cost paths from the initial state to a goal node (Pearl 1984), we will refer to this set as simply P^*.

Definition 1 *A solution in this new setting is a set of mini-*

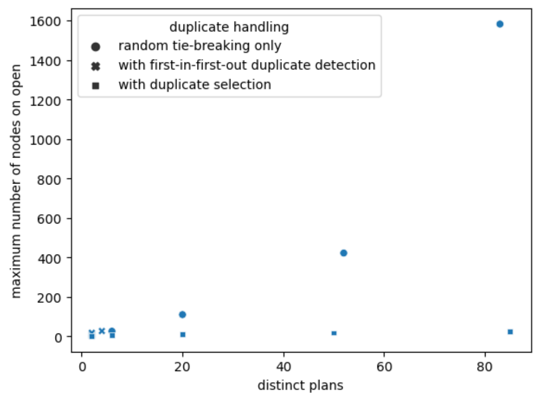

Figure 2: Time-efficiency of A* with our new strategy vs. FIFO duplicate detection

mal cost paths π in P^*, and for all plans in the solution set, $P(\pi) = 1/|P^*|$

A* with Random Duplicate Selection

We address the problem of sampling-bias introduced by duplicate detection in A* by a few simple modifications to the original algorithm (Hart, Nilsson, and Raphael 1968). The most significant change is instead of recording a single parent pointer for each search node, we maintain a list of parent pointers. Additionally we maintain a CLOSED list. On expansion we add the node to CLOSED. On generation we check if the state is on CLOSED already. If it is and the path to this new node is no worse than the node on CLOSED then we *update the closed node's parent list* and discard the successor. A solution is recovered by a random walk from the goal node to the initial state via our new parent lists. We achieve a uniform distribution over all possible plans in P* by sampling parents in proportion to the number of paths they lead to. Figure 2 demonstrates our approach on Gridworld domain with synthetic heuristic error.

A Model for the Greedy Approach

The graph our variation of A* discovers during search is, in the satisficing setting, the reachability graph from the start node to the goal node. This idea inspires our approach to fair planning for greedy-best first search. If we had access to the reachability graph under each node tied on OPEN then we could select the node which undoubtedly leads to more distinct paths, but we do not have access to the reachability graph during forward search so we must consider what kind of estimate might be useful. In this case, knowledge of the branching factor might be helpful. Combined with an estimate of solution depth, the branching factor could be used to estimate the number of paths under a given node. This estimate could be refined during search.

A Novel Application

Agent-based modeling is a method for describing the behavior of complex systems which are either too difficult to model using equation-based methods or are not yet understood well enough, by modeling each individual agent in the system, their state, and their behavior. It has been used broadly in the study of infectious diseases. In epidemiology, the state-of-the-art in activity selection in the absence of real-world observations is a hand-tuned model (Xia 2014). Planning may be a good alternative. Planning technology, however, that always excludes specific perfectly good plans for arbitrary reasons could confound the analysis and interpretation of these models.

Acknowledgements

Queen's University is located on territory populated by the Wendake-Nionwentsïo, Mississauga, and Haudenosaunee nations, which was shared with the British Crown in the Crawford Purchase of 1783; remote work took place in New Hampshire, located in the ancestral homeland of the Abenaki nation. I want to thank J. Benton and J. Frank for their feedback on early iterations of these ideas, and M. Ramirez for his helpful comments on the penultimate draft of this abstract.
Advisor: Christian Muise

References

Asai, M.; and Fukunaga, A. S. 2016. Tiebreaking Strategies for A* Search: How to Explore the Final Frontier. In Schuurmans, D.; and Wellman, M. P., eds., *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 673–679.

Ernandes, M.; and Gori, M. 2004. Likely-Admissible and Sub-Symbolic Heuristics. In Mántaras, R. L. d.; and Saitta, L., eds., *Proceedings of the 16th Eureopean Conference on Artificial Intelligence*, 613–617. IOS Press.

Ferber, P.; Geißer, F.; Trevizan, F. W.; Helmert, M.; and Hoffmann, J. 2022. Neural Network Heuristic Functions for Classical Planning: Bootstrapping and Comparison to Other Methods. In Kumar, A.; Thiébaux, S.; Varakantham, P.; and Yeoh, W., eds., *Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling, ICAPS 2022*, 583–587. AAAI Press.

Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.

Helmert, M.; and Röger, G. 2008. How Good is Almost Perfect? In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 2*, 944–949. AAAI Press.

Pearl, J. 1984. *Heuristics - intelligent search strategies for computer problem solving*. Addison-Wesley series in artificial intelligence. Addison-Wesley.

Thayer, J. T.; and Ruml, W. 2011. Bounded suboptimal search: A direct approach using inadmissible estimates. In *Proceedings of the Twenty Sixth International Joint Conference on Artificial Intelligence*, 674–679,.

Xia, H. 2014. *Modeling, Analysis and Comparison of Large Scale Social Contact Networks for Epidemic Studies*. Ph.D. thesis, Virginia Polytechnic Institute and State University.