

Demos for a Course in Single-Agent Heuristic Search

Nathan R. Sturtevant

Department of Computing Science
Alberta Machine Intelligence Institute (Amii)
University of Alberta, Canada
nathanst@ualberta.ca

Abstract

This paper describes educational material available from a graduate course on single-agent search. The material currently includes 25 interactive demos that illustrate different algorithms and concepts in heuristic search. Coupled with each demo is material from one or more recorded lectures that relate to the work shown in the demo. Together these can be used as a reference or for teaching many foundational topics in single-agent heuristic search.

A Course in Single-Agent Heuristic Search

This abstract describes material from a course in Single-Agent Heuristic Search which is available from <http://www.movingai.com/SAS/>. We cover the motivation and history of the material, the broad topics covered, and how we have used them when teaching our own classes.

History

We have always valued the use of visual materials when doing research and teaching courses. For our graduate classes, this began with exporting static videos from research visualizations. These videos showed algorithms like A*, and also demonstrated different tie-breaking rules and the impact of different heuristics. However, these videos are static, and thus are not as effective for learning as it would be to interact with the underlying simulation.

So, in 2014 we began building demos explicitly for use in a general AI class, an effort that was expanded to include the majority of a graduate heuristic search course in 2016. When grading our final exam that year, we discovered that students did very well on all questions that came from material with demos, and did very poorly on the few topics which weren't covered by online demos. This motivated further development of the material for future offerings of the course.

At that point we could use these demos in our class while interacting with students, but they still weren't available for students to explore themselves. However, in 2018 we integrated the `emscripten` compiler to build these demos (written in C++) into javascript so they could run on the web. Since then we have been expanding and polishing the demos each time our graduate course is taught.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Finally, with the COVID pandemic, we were forced to teach our 2021 graduate course online. We then decided to record all of our lectures, making 48 videos of varying length available both on YouTube and alongside many of the demos. These demos and videos were used as part of a flipped classroom (Tucker 2012) for our course.

Material

As of now, we have built 25 interactive demos for teaching concepts in Single-Agent Heuristic Search. The main demo visualization for one such demo can be seen in Figure 1. Overall, these demos fall into several categories. The first is sample domains, where we currently have the pancake puzzle and the sliding-tile puzzle implemented. These domains are useful for exploring the difficulty of finding optimal solutions on combinatorial puzzles.

The next set of demos are on foundational algorithms in heuristic search such as Dijkstra's Algorithm (Dijkstra et al. 1959), A* (Hart, Nilsson, and Raphael 1968), DFID and IDA* (Korf 1985), and Weighted A* (Pohl 1970). The IDA* demo visualizes the entire 3x2 sliding tile puzzle and clearly shows the behavior of the algorithm on the search tree, including the impact of transpositions. The A* demos are available on both grid maps and general graphs, and can be used to analyze behavior step-by-step.

A third set of demos focus on the heuristics and constraints which can be used to improve search. Heuristics including pattern databases (Culberson and Schaeffer 1996; Felner et al. 2007) for exponential domains, and various embeddings such as differential heuristics (Felner, Sturtevant, and Schaeffer 2009) and FastMap (Cohen et al. 2018). Constraints include methods like reach (Goldberg, Kaplan, and Werneck 2006), bounding boxes (Rabin and Sturtevant




Figure 1: Sample visualization of Weighted A* using a piecewise priority function (Chen and Sturtevant 2021).

2016; Hu et al. 2021), or the canonical orderings used in JPS (Harabor and Grastien 2011; Sturtevant and Rabin 2016)

A final set of demos arise from our research; these demos were often built as part of our research papers and then adapted into demos for use in our course or in poster presentations. This includes numerous demos on the theory and practice of bidirectional search (Chen et al. 2017; Eckerle et al. 2017; Sturtevant and Felner 2018), recent variants of Weighted A* (Chen and Sturtevant 2021), new algorithms like IBEX (Helmert et al. 2019), and work on abstraction and refinement that we’ve implemented for the games industry (Sturtevant 2007; Sturtevant et al. 2019)

Use in a Flipped Classroom

We have most recently used this material as part of a flipped classroom, where students watch course videos before coming to class, and then interact with the material during lecture. Each lecture is paired with one or more tasks that students must complete or questions for them to explore during class time, which we can then discuss. These are tasks like “create a graph where a node is generated with a higher g-cost before it is expanded with a lower g-cost” that require a deeper understanding of algorithmic behavior. The depth of the questions we can ask vary based on each demo, but we find that these interactions bring a deeper understanding of the material than what is gained from just lectures.

Conclusions

This paper describes the single-agent heuristic search demos available from our web site. These videos and demos are free to use. We hope this paper will raise awareness of this resource so many more can benefit from our efforts.

Acknowledgments

This work was funded by the Canada CIFAR AI Chairs Program. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). The work was supported by NSF grants #1551406 and #1815660.

References

Chen, J.; Holte, R. C.; Zilles, S.; and Sturtevant, N. R. 2017. Front-to-End Bidirectional Heuristic Search with Near-Optimal Node Expansions. *International Joint Conference on Artificial Intelligence (IJCAI)* .

Chen, J.; and Sturtevant, N. R. 2021. Necessary and Sufficient Conditions for Avoiding Reopenings in Best First Sub-optimal Search with General Bounding Functions. *AAAI Conference on Artificial Intelligence* .

Cohen, L.; Uras, T.; Jahangiri, S.; Arunasalam, A.; Koenig, S.; and Kumar, T. S. 2018. The FastMap Algorithm for Shortest Path Computations. In *IJCAI*.

Culberson, J. C.; and Schaeffer, J. 1996. Searching with pattern databases. In *Conference of the Canadian Society for Computational Studies of Intelligence*, 402–416. Springer.

Dijkstra, E. W.; et al. 1959. A note on two problems in connexion with graphs. *Numerische mathematik* 1(1): 269–271.

Eckerle, J.; Chen, J.; Sturtevant, N.; Zilles, S.; and Holte, R. 2017. Sufficient Conditions for Node Expansion in Bidirectional Heuristic Search. In *International Conference on Automated Planning and Scheduling (ICAPS)*.

Felner, A.; Korf, R. E.; Meshulam, R.; and Holte, R. C. 2007. Compressed pattern databases. *Journal of Artificial Intelligence Research* 30: 213–247.

Felner, A.; Sturtevant, N.; and Schaeffer, J. 2009. Abstraction-based heuristics with true distance computations. *Symposium on Abstraction, Reformulation and Approximation* 9.

Goldberg, A. V.; Kaplan, H.; and Werneck, R. F. 2006. Reach for A*: Efficient point-to-point shortest path algorithms. In *2006 Proceedings of the Eighth Workshop on Algorithm Engineering and Experiments (ALENEX)*, 129–143. SIAM.

Harabor, D.; and Grastien, A. 2011. Online graph pruning for pathfinding on grid maps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25.

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.; Lattimore, T.; Lelis, L. H. S.; Orseau, L.; and Sturtevant, N. R. 2019. Iterative Budgeted Exponential Search. *International Joint Conference on Artificial Intelligence (IJCAI)* .

Hu, Y.; Harabor, D.; Qin, L.; and Yin, Q. 2021. Regarding Goal Bounding and Jump Point Search. *Journal of Artificial Intelligence Research* 70: 631–681.

Korf, R. E. 1985. Depth-first Iterative-Deepening: An Optimal Admissible Tree Search. *Artificial Intelligence* 27: 97–109.

Pohl, I. 1970. Heuristic search viewed as path finding in a graph. *Artificial intelligence* 1(3-4): 193–204.

Rabin, S.; and Sturtevant, N. R. 2016. Combining Bounding Boxes and JPS to Prune Grid Pathfinding. In *AAAI Conference on Artificial Intelligence*.

Sturtevant, N. R. 2007. Memory-Efficient Abstractions for Pathfinding. In *Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 31–36.

Sturtevant, N. R.; and Felner, A. 2018. A Brief History and Recent Achievements in Bidirectional Search. In *AAAI Conference on Artificial Intelligence*.

Sturtevant, N. R.; and Rabin, S. 2016. Canonical Orderings on Grids. *International Joint Conference on Artificial Intelligence (IJCAI)* 683–689.

Sturtevant, N. R.; Sigurdson, D.; Taylor, B.; and Gibson, T. 2019. Pathfinding and Abstraction with Dynamic Terrain Costs. In *Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*.

Tucker, B. 2012. The flipped classroom. *Education next* 12(1): 82–83.