

Project 4: Review a Paper

Due [one week after release]

In this project you will take the role of a reviewer considering an academic paper for publication. You may choose one of the following papers to read and review:

- Culberson, Joseph C. and Jonathan Schaeffer. "Pattern databases." *Computational Intelligence* 14.3 (1998): 318-334. ([PDF](#))
 - A principled type of heuristic for A* search that applies to a broad class of problems (applied to the 15-puzzle).
- Dai, Peng, and Eric A. Hansen. "Prioritizing Bellman Backups without a Priority Queue." In *Proceedings of the 19th International Conference on Automated Planning and Scheduling*, p. 113-119, 2007. ([PDF](#))
 - Speeding up Value Iteration by prioritizing important updates
- David, Omid E. et al. "Genetic algorithms for evolving computer chess programs." *IEEE Transactions on Evolutionary Computation* 18.5 (2014): 779-789. ([PDF](#))
 - *Using genetic algorithms to search for chess strategies*
- Gelly, Sylvain, and David Silver. "Achieving Master Level Play in 9 x 9 Computer Go." In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, p. 1537-1540, 2008. ([PDF](#))
 - An early predecessor of the famous AlphaGo agent
- Ginsberg, Matthew L. and William D. Harvey. "Iterative Broadening." In *Proceedings of the 8th National Conference on Artificial Intelligence*, p. 216-220, 1990. ([PDF](#))
 - An uninformed search algorithm that attempts to control the branching factor, rather than the depth of the search
- Lelis, Levi H. S. "Stratified Strategy Selection for Unit Control in Real-Time Strategy Games." In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, p. 3735-3741, 2017. ([PDF](#))
 - An application of search methods to a really big problem
- Sorg, Jonathan, Satinder Singh, and Richard Lewis. "Internal Rewards Mitigate Agent Boundedness." In *Proceedings of the 27th International Conference on Machine Learning*, p. 1007-1014, 2010. ([PDF](#))
 - Shows that if the agent will not perform optimally, one has to be very careful when designing the agent's reward function.

Note that because the paper you are reviewing was already accepted for publication, you can expect it to be of high quality. However, it is very unlikely to be perfect! It is

certainly possible that the paper contain errors or missing details, or could otherwise be improved by more effective writing, additional experiments, better structure, etc. Your goal is to read the paper carefully and critically, noting both what it does well and how it could be stronger. A “careful reading” means reading the paper multiple times and slowing down during (not skipping over!) dense technical details. It means actively interrogating the paper – taking notes, developing questions, then hunting down the answers in the text.

Your review should be 1-2 single-spaced pages and should address the following:

- First, give a 1-paragraph summary of the paper (in your own words) that describes the problem being solved, the approach taken (at a high level), and the key findings.
- What are the main claims of the paper and are they well supported?
 - If a claim is supported empirically, does the experiment provide compelling evidence?
 - If a claim is supported theoretically, is the argument/proof sound/correct?
 - Does the paper discuss *both* the strengths and weaknesses of the approach? How effectively are the strengths/weaknesses illustrated and evaluated?
- Is the paper well-written and clear?
 - Were you able to follow the logic of the paper?
 - Were important concepts clearly introduced?
 - Are both the high-level ideas and the low-level details communicated effectively?
 - Could an informed reader reproduce the results presented in the paper?
- How significant are the contributions of the paper likely to be?
 - Did the paper do a good job of motivating the problem? Are you convinced that it is important?
 - Did the paper do a good job of motivating its approach? Are you convinced that this is a promising direction?
 - Did the paper clearly describe the relationship of this work to existing results?
 - Are you convinced that the findings are likely to change practice, substantially increase understanding, or be built upon in future work?
- Lastly give a brief summary of your review and an overall assessment of the quality of the paper.

Your reviews should be clear and to the point (e.g. explicitly state your point up front, then support it). You should write in full, well-structured sentences and paragraphs, but this is not an essay or a book/movie review. The flow of the document as a whole is not as important as effectively and efficiently getting your points and reasoning across – consider splitting your review into sections with appropriate headings to make it easy to find your comments on the various review

criteria. That said, just as you would with an essay you should draft and revise your reviews to make them as clear, persuasive, and helpful as you can.

Note that for grading I am far more interested in the quality of your writing than your opinions about the paper. If you take a patently absurd position then you may lose points, but generally I will focus on whether your points are well supported and stated clearly and not on whether I agree with your assessment.

A paper review has two main audiences, and you should think about them both as you write (and revise).

One important audience is the authors of the paper. To best reach the authors, you must take a respectful, constructive tone. The authors want to know that you have carefully read their paper and taken the time to understand it as well as you can in order to feel confident that your critique is informed and well-considered. You should point out what the paper does well and any criticisms you make should be specific, well-supported, and accompanied by suggestions for how the paper can be improved.

The other main audience member is the program chair or editor in charge of the paper. This person will receive multiple reviews and must rely on them to make a determination about the paper's fate. Depending on the venue, he or she may not have read the paper as carefully as you (or at all!). So, make sure you state the reasoning behind your critiques and support your arguments with examples/evidence from the paper. That said, an overly long or wordy review can waste the editor's time; give concise assessments of the various aspects of the paper.