



**KTH Computer Science
and Communication**

OPPONENT RECORD

Thesis compiled by

Rikard Blixt and Anders Ye

Title of thesis:

Reinforcement learning AI to Hive

Opponent:

Bassam Alfarhan (alfarhan@kth.se)

Course:

Degree Project in Computer Science, First Level (dkand13)

Was it easy to understand the underlying purpose of the project? Comments.

The purpose of the project was defined under the “Purpose and problem statement” subsection of the report in a rather informal fashion. The problem statement, however, was not defined explicitly (with a question mark) as one would expect.

The purpose is to:

- 1) “Create a strong AI without strategies”
- 2) “Is it possible to make a strong AI by playing two AI against each other for a long time”

Do you consider that the report title justly reflects the contents of the report?

The report goes through what reinforcement learning is and how it is used to implement an AI that plays Hive. Therefore I believe the title describes the content (more precisely, the purpose of the report), although they are not directly connected.

How did the author describe the project background? Was there an introduction and general survey of this area?

There were clearly no previous research in the area the report discusses. That is, however, not stated explicitly. The background subsection of the report presents the game rules, which is good in a way and bad in another. It is good, since the reader, like myself, might not know what the game is about. At the same time, it is bad, since it is irrelevant to some extent.

To what degree did the author justify his/her choice of method of tackling the problem?

- 1) *The method the authors choose to teach the Artificial Intelligence (reinforcement learning) is stated in the problem statement, the authors did not discuss nor justify why they did choose this method.*
- 2) *The method the authors choose to find out if it is possible to use reinforcement learning is by implementing it and studying the result. This too is not discussed nor justified by the authors.*

Did the author discuss the extent to which the prerequisites for the application of such a method are fulfilled?

- 1) *Yes, the authors stated clearly that the basic requirements of reinforcement learning are to have a set of cases, a set of actions and rules to go from one state to another.*
- 2) *They even stated that they had to implement the game in order for the Artificial Intelligence to be tested.*

Is the method adequately described?

The method used to solve the problem were briefly described, it was understandable and kind of informal as well. I would say it was vague as it was not explicitly stated.

Has the author set out his/her results clearly and concisely?

The results were presented in graphs and tables that shows some of the AI states/behaviours and the improvement of the AI respectively. All of which is well explained.

Do you consider the author's conclusions to be credible?

With respect to the result the authors received (and presented) from the experiment; the conclusion is totally reasonable (and even expected); and therefore also credible. The authors even discussed the result in a good way, which increases the conclusions credibility.

What is your opinion of the bibliography? What types of literature are included? Do you feel they are relevant?

The references stated in the bibliography seems to be relevant, but also questionable as 4 of the total 7 references are links to [Wikipedia] articles, furthermore, no source criticism are provided. Apart from that, the literature are all presented in a consistent fashion.

Which sections of the report were difficult to understand?

*Linguistically, most of the content presented in the report were written in a rather informal (friendly) way, which makes it very easy to read and follow through.
Content wise, some of the informality were misleading, like the "Carrot and Stick" approach, which I still does not know what it means.*

Other comments on the report and its structure.

- 1) I think that the method should be described explicitly, it would be better to dedicate a subsection for the method.*
- 2) The background should present previous research in the subject; clearly there were not any. The current background describes the game rules which is totally essential for any one who does not know how Hive is played (like myself), however, I think it should be included as an appendix or “introduction to the game”-subsection instead.*
- 3) The Artificial Intelligence (AI) subsection describes what AI, Machine learning and Reinforcement learning is; these are three separate topics that, I think, should each be addressed in a subsection.*
- 4) Some of the parts, as in the Implementation section starts with what I think is background related; No previous implementations of the game, etc.*
- 5) Colour coding is good, however, it is somehow cryptic when the report is printed in black and white.*

What are the stronger features of the work/report?

- 1) The result is presented and described very clearly (apart from the numbers below the graph that merged toward the end)*
- 2) Even though I wrote that the game rules should perhaps be included in an appendix; the rules are very well described.*
- 3) The structure is well laid out, apart from the parts I have pointed out earlier.*
- 4) The conclusion is presented in a clear fashion and is easy to find.*

What are the weaker features of the work/report?

- 1) The content was to a high extent presented in an informal fashion. The authors did refer to them self a lot throughout the report.*
- 2) Some parts should be explicitly highlighted; like the method. It might be obvious that the method used is “implementing an Artificial Intelligence using reinforcement learning”. However, I think, it should be highlighted despite that.*
- 3) No source criticism, which is essential as some of the sources are questionable.*

What is your estimation of the news value of the work?

The news value is both high and low as the work is probably one of the first to treat this question for Hive, however, this question have been studied earlier for other board games with similar branching factor. The same result could have been achieved by studying the branching factor of the game and comparing it to other with similar branching factor.

This work is original since the authors choose to start with the problem from scratch instead of interpreting similar board games, which could have been easier (and might yield better results).

Summarize the work in a few lines.

The purpose of the work is to find out if it is possible to make an Artificial Intelligence (AI) that plays the board game Hive using reinforcement learning. That is, by letting the AI play the game and gain experience by itself, which is done by implementing the game and letting two AI (which also is implemented by the authors) play against each other for a while. Statical data from the AI game play is analysed thereafter to show that; it is infeasible, since it would take too long to make the AI reasonably good at the game.

Questions to authors:

- 1. Did you consider that the AI might get stuck in a cycle in which it weight all the actions equally and never become good even if it were able to run indefinitely?*
- 2. Did you consider that the efficiency of the implementation? Your conclusion were mainly based upon the time it takes to teach the Artificial Intelligence, which might be improved considerably by the implementation.*
- 3. Why did you chose (reinforcement learning) instead of, for example, (artificial neural networks) which is proven for similar board games with high branching factors?*
- 4. What is the Carrot and Stick approach? It was stated in a couple of places in the report but never explained!*
- 5. Why is there no source criticism, do you consider [Wikipedia], for example, to be trust worthy?*