

Report on PhD Thesis:

General Game Description Languages

by Jakub Kowalski

Julian Togelius

1 Summary of Contribution

Jakub Kowalski’s PhD thesis concerns game description languages, i.e. languages that specify elements of games, in particular their rules. The contributions to game description languages include learning rules of games from observing them, generating game rules, translations between different game description languages, compiling languages and additions to existing game description languages. Much of the work builds on the Stanford GDL, used for the original General Game Playing Competition, but several of the contributions also make use of the Simplified Board Games language, and one of the contributions build on the Card Game Description Language.

To put the contributions into context, game description languages are used both in general game playing, where AI game-playing agents are tested on multiple unseen games, and in game generation, where systems are constructed that attempt to algorithmically generate aspects of games; the latter can be seen as a generalization of procedural content generation, the algorithmic creation of game content. The communities working on these problems are largely disjoint, but Kowalski effectively connects to both communities and sets of research questions through the extensive literature survey.

The work on game generation uses an evolutionary algorithm to search for good rulesets in the Simplified Board Games domain. Two different types of fitness functions are used, one based on game features, and the other on Relative Algorithm Performance Profiles, the idea of estimating game skill depth through the performance differential of several algorithms. This idea has previously only been used in a video game domain, but Kowalski shows that it is possible to use this method to achieve good results—arguably better than in the video game domain—in the simplified board game domain.

The translation from the Card Game Description Language, a rather specialized game description language, to the much more generic and commonly used Stanford GDL, is an interesting contribution in that it shows how one game description language can be reduced to another, but also what is lost in terms of compactness and readability through this operation. I have not seen another such translation before, and I think it gives interesting perspectives on

representational choices in game description language design.

The extensions of the Stanford GDL with a real-time element, as well as the language compiler, are as far as I can judge technically solid and novel contributions to that language and the infrastructure around it. However, it should be noted that I am not sufficiently well versed in that language to comprehend all the technical details.

2 Evaluation of Contribution

There is much to commend in Kowalski's thesis. Kowalski has demonstrated deep technical mastery of both several game description languages, a good understanding of artificial intelligence and procedural content generation as applied to games, and proficient use of evolutionary computation. The thesis has provided several novel contributions that other researchers are likely to build on, including in particular the additions to the Stanford GDL and the work on generating Simple Board Games. The work in the thesis is technically sophisticated, well situated with regards to the state of the art of the fields it contributes to, and original. Compared to other PhD theses I have examined, I would say that Kowalski's work shows more technical maturity than most, and also draws on a somewhat wider range of work. The thesis is overall well-written and well-sourced. It is clear that there is much promising follow-up work to be done, which bodes well for the future career of the author of the thesis.

My main criticism is the relative lack of unifying research question and/or interdependency between the various contributions in the thesis. The different chapters are mostly related by all being about game description languages. While each of the chapters contain valuable contributions, I have the sense that the depth of the contributions could have been greater if the author had spent more time exploring fewer of the handful research directions that are taken. I would also have preferred to see the author provide a narrative of how e.g. the language translation and compilation effort could help with game generation or vice versa. However, I do not see this as a serious problem, given the strength of the work in general.

3 Conclusion

It is clear to me that Jakub Kowalski's work is well deserving of a PhD. He has drawn on a wide range of research and made several original contributions to the state of the art within game description languages and their use in both general game playing and procedural content generation. I therefore recommend that the the thesis be accepted.

A handwritten signature in blue ink, reading "Julian Togelius". The signature is fluid and cursive, with a long horizontal stroke extending to the right.

Julian Togelius
New York, April 18, 2017
Associate Professor
Department of Computer Science and Engineering
Tandon School of Engineering, New York University
2 MetroTech Center, 11201 Brooklyn, NY, USA