UNIVERSITY OF ALBERTA
EDMONTON, CANADA

Why Hex

Chao Gao, cgao3@ualberta.ca

The transformation brought by AlphaGo and its successors has led notable AI researcher to believe that future research in two-player alternate-turn zero-sum perfect-information games is inconsequential [3]. However, AlphaGo [15], AlphaGo Zero [17] and AlphaZero [16] are essentially heuristic algorithms without theoretical guarantee on success. They contain a number of important hyperparameters and how they affect the overall performance is only meagerly understood. Public re-implementations such as Leela Zero [12] and OpenELF Go [18] are towards the investigation of this phenomenon.

The fundamental difference that sets Hex asides from many other classic board games is its strong mathematical structure [6], which has enabled much research in Hex being presented in an exact rather than heuristic manner. Some examples are the proof that there is no draw [11], the identification of dead, dominated and inferior cells [7, 10], and the computation of connection strategies [1]. The accumulation of these mathematical knowledge, combined with sophisticated search, has enabled computer programs that to solve Hex openings in board sizes up to 10×10, whose state spaces are already far beyond the limit of any simple brutal-force search. For example, the number of states for 9×9 and 10×10 Hex are respectively 10^{37} and 10^{46}; see Table **??**.

However, *mathematical knowledge* accumulation becomes more and more difficult as it has to invoke more and more complicated reasoning. On the other side, algorithms that learn with deep neural networks have shown great capacity in acquiring *heuristic knowledge* from data. These two types of knowledge are fundamentally different and are arguably complementary: given the seemingly intractable problem, the mathematical knowledge states what we can *at least* identify, while the later represents what we can be guessed *at most* after seeing a number of noised observations. Both of them have their merits and limitations. For example, the continual identification of inferior cells [2, 5, 9, 10] and the development of H-search [1, 8, 13]

Table 0.1: Status of solved Hex board sizes. For 10×10, only 2 openings are solved. For other smaller board sizes, *all* openings have been solved.

Board size	status	year	method	computation time
6×6	solved, all	2000	DFS [19]	seconds
7×7	solved, all	2003	DFS [5]	minutes
8×8	solved, all	2008	DFS [9]	hours
9×9	solved, all	2013	parallel FDFPN [14]	months
10×10	solved, 2	2013	parallel FDFPN [14]	months

since the 2000s have quickly led to feasible computer solutions for board sizes from 6×6 to 10×10. See Table 0.1 for a summarization. However, it is unlikely that 11×11 Hex can be solved if no overwhelmingly larger amount of pruning due to inferior cells analysis or H-search is introduced, i.e., although the pruning they brought is often exponential, the state-space complexity inevitably grows at a faster rate which itself is also exponential [1]. The deep neural networks expressed knowledge, although can be probably correct with the help of look-ahead tree search, does not surely prune anything before doing search. Yet, guessing guided look-ahead search in a state-space graph faces another challenge: the solution graph itself could be intractably large which implies that even an error-free guessing technique is employed, verification could still be infeasible. Such an observation highlights the use of state-abstraction method in informed guess-based forward search. In Hex, strategy decomposition is one technique of this sort. Given the observation that advancement in machine learning have enabled more and more accurate heuristic guidance, together with existing exact knowledge computation techniques, we conjecture that a promising future direction for solving Hex is to search decomposition-based solutions — this is arguably how human solves Hex positions [20].

In summary, Hex is a game that has interested a number of mathematicians and computer scientists since its invention; its graph-theoretical, combinatorial, game theoretic, and artificial intelligence aspects are perpetual incentives for attracting more research in the future. Despite grand successes, deep learning techniques have been questioned by the lack of reasoning [4]; as a domain where reasoning is ubiquitous and of utmost importance, Hex could be a valuable domain for pushing machine learning research to incorporate reasoning techniques as well.

References

[1] Vadim V Anshelevich. A hierarchical approach to computer Hex. *Artificial Intelligence*, 134(1):101–120, 2002.

[1] For a state-space of b^d, the effect of these pruning is a constant reduction of b and d, but as board size increases, both b and d increase

[2] Yngvi Björnsson, Ryan Hayward, Michael Johanson, and Jack van Rijswijck. Dead cell analysis in Hex and the shannon game. In *Graph Theory in Paris*, pages 45–59. Springer, 2006.

[3] Murray Campbell. Mastering board games. *Science*, 362(6419):1118–1118, 2018.

[4] Adnan Darwiche. Human-level intelligence or animal-like abilities? *Commun. ACM*, 61(10):56–67, 2018.

[5] Ryan Hayward, Yngvi Björnsson, Michael Johanson, Morgan Kan, Nathan Po, and Jack van Rijswijck. Solving 7x7 Hex: Virtual connections and game-state reduction. In *Advances in Computer Games*, pages 261–278. Springer, 2004.

[6] Ryan B Hayward and Bjarne Toft. *Hex: The Full Story.* CRC Press, 2019.

[7] Ryan B Hayward and Jack Van Rijswijck. Hex and combinatorics. *Discrete Mathematics*, 306(19-20):2515–2528, 2006.

[8] Philip Henderson, Broderick Arneson, and Ryan B Hayward. Hex, braids, the crossing rule, and XH-search. In *Advances in Computer Games*, pages 88–98. Springer, 2009.

[9] Philip Henderson, Broderick Arneson, and Ryan B Hayward. Solving 8x8 Hex. In *Proc. IJCAI*, volume 9, pages 505–510. Citeseer, 2009.

[10] Philip Thomas Henderson. *Playing and solving the game of Hex*. PhD thesis, University of Alberta, 2010.

[11] John F Nash. Some games and machines for playing them. Technical report, Rand Corporation, 1952.

[12] Gian-Carlo Pascutto. Leela Zero. `https://github.com/leela-zero/leela-zero`. Accessed: 2019-08-18.

[13] Jakub Pawlewicz, Ryan Hayward, Philip Henderson, and Broderick Arneson. Stronger Virtual Connections in Hex. *IEEE Transactions on Computational Intelligence and AI in Games*, 7(2):156–166, 2015.

[14] Jakub Pawlewicz and Ryan B Hayward. Scalable parallel DFPN search. In *International Conference on Computers and Games*, pages 138–150. Springer, 2013.

[15] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

[16] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.

[17] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.

[18] Yuandong Tian, Jerry Ma, Qucheng Gong, Shubho Sengupta, Zhuoyuan Chen, James Pinkerton, and Larry Zitnick. ELF OpenGo: an analysis and open reimplementation of alphazero. In *ICML*, pages 6244–6253, 2019.

[19] Jack Van Rijswijck. Computer Hex: Are bees better than fruitflies? Master's thesis, University of Alberta, 2002.

[20] Jing Yang, Simon Liao, and Mirek Pawlak. On a decomposition method for finding winning strategy in hex game. In *Proceedings ADCOG: Internat. Conf. Application and Development of Computer Games (ALW Sing, WH Man and W. Wai, eds.), City University of Honkong*, pages 96–111, 2001.