



10 years ESJ  
*Special edition*

## **Reproducing Neural Network Research Findings via Reverse Engineering: Replication of *AlphaGo Zero* by Crowdsourced *Leela Zero***

*Dustin Tanksley, BS*

*Daniel B. Hier, MD*

*Donald C. Wunsch II, PhD*

Applied Computational Intelligence Laboratory  
Missouri University of Science and Technology, USA

[Doi:10.19044/esj.2022.v18n4p61](https://doi.org/10.19044/esj.2022.v18n4p61)

---

Submitted: 10 November 2021  
Accepted: 20 January 2022  
Published: 08 February 2022

Copyright 2022 Author(s)  
Under Creative Commons BY-NC-ND  
4.0 OPEN ACCESS

*Cite As:*

Tanksley D., Hier D.B. & Wunsch D.C. (2022). *Reproducing Neural Network Research Findings via Reverse Engineering: Replication of AlphaGo Zero by Crowdsourced Leela Zero*. European Scientific Journal, ESJ, 18 (4), 61.

<https://doi.org/10.19044/esj.2022.v18n4p61>

---

### **Abstract**

The reproducibility of scientific findings is essential to the integrity of research. The scientific method requires hypotheses to be validated independently by different laboratories. Investigators are expected to provide sufficient information in their publications to permit an objective evaluation of their methods and an independent reproduction of their results. This is particularly true for research supported by public funds, where transparency of both methods and findings represents a return on public investment. Unfortunately, many publications fall short of this standard for various reasons, including a desire to protect intellectual property or national security. The reproducibility of findings is essential in transferring machine learning findings from research into healthcare practice. Fortunately, the internet makes it easier to overcome these limitations by permitting multiple individuals to participate in reproducibility efforts and to crowdsource the reverse engineering of novel software. We present a case study of this capability from neural network research. The success of the crowdsourced project *Leela Zero* to reverse engineer the findings of *AlphaGo Zero* exemplifies the ability to reproduce novel results despite the lack of extensive

computational resources or a detailed description of the initial experimental methods. The implications of this successful reverse engineering effort for future reproducibility of neural network research are discussed.

---

**Keywords:** Reproducibility, neural networks, DeepMind, crowdsourcing, reverse engineering, *Leela Zero*, *AlphaGo Zero*

## Introduction

The reproducibility of scientific research has received increased attention recently. Several studies have shown that some notable findings in highly-reputed journals were not reproducible. Baker (2016) reported that a survey of 1,500 scientists found that "more than 70% of researchers have tried and failed to reproduce another scientist's experiments and more than half have failed to reproduce their own experiments". In addition, 52% of those surveyed agreed that there is a significant crisis of reproducibility of scientific research. Many top-tier journals are now demanding improvements in the description of research methods and the sharing of data to enhance reproducibility. Greene et al. (2017) state that "journals, although they conduct peer review, do not validate each experimental result or claim." They go on to say that

...rigorous secondary data analysis is critical for maintaining the accuracy and efficiency of scientific discovery. As scientists, we make predictions, perform experiments, and generate data to test those predictions. When we ask rigorous questions, we obtain more accurate findings that can prevent harm.

The National Academies of Sciences, Engineering, and Medicine (2019) notes that

when results are produced by complex computational processes using large volumes of data, the traditional methods section of a scientific paper is insufficient to convey the necessary information for others to reproduce the results. Additional data, code, models, and computational analysis are needed.

However, the desire for reproducibility can conflict with the legitimate interests of the researchers who seek to protect the economic value of their work. The pros and cons of sharing results and methods have been vigorously debated in the medical literature. The decision to share detailed methods can

be difficult for industry researchers who are obligated to protect proprietary research funded by for-profit companies. A prominent example of this conflict over the disclosure of methods occurred when DeepMind (a Google company) used neural network research to create an AI system for the game of Go (*AlphaGo*) which outplayed every other human or computer player in the world. Although DeepMind disclosed the general architecture of *AlphaGo*, they did not disclose the network weights, computer code, or training data that would allow replication of this result (Silver et al. 2016).

This paper is divided into four parts. First, we have introduced the problem of reproducibility of scientific research. Second, we describe the surprising neural network achievement of DeepMind which created *AlphaGo Zero*, the best Go player in the world. Third, we highlight how crowdsourcing reproduced the performance of *AlphaGo Zero* utilizing minimal resources. Finally, we discuss the implications of this successful replication of the capabilities of *AlphaGo Zero* by the crowdsourced effort *Leela Zero* for the reproducibility of neural network research in the future.

### **Case Study Part I: The development of the World Champion Go-playing engines *AlphaGo* and *AlphaGo Zero***

Go is an ancient Chinese board game that has been played continuously for over 2,500 years. It is played on a 19 x 19 squared grid with black and white-colored stones. An estimated 46 million people worldwide know how to play Go. Traditionally the ability of Go players was graded on the dan/kyu ranking system from 20 kyu (beginner) to 9 dan (professional). Recently, a numerical skill rating system call Elo has been developed with ratings from 100 (beginner) to 2940 (top professional). Ratings above 3,000 are possible but are considered super-human. Elo ratings can be used for human players and computer engines (Go ranks and ratings, n.d.).

As early as 1984, there has been computer versus computer Go challenges. Since at least 1986 there has also been computer versus human Go challenges. In general, computer Go engines could only play at no better than an amateur level (Levinovitz, 2014). Most Go researchers did not expect the development of a champion-level system in the near term. Go is much harder than any of the games solved by previous AI systems. The game tree (all possible moves) is estimated to be between 10575 and 10620, hundreds of orders of magnitude greater than the game tree for Chess which is estimated at around 10120 (Allis, 1994; Allis et al., 1990; Koch, 2016). Fundamental research had been conducted on Go for decades. Silver, Sutton, & Müller (2007) speculated that Go would require an altogether different approach than used for other games. Krikke (2007) foresaw that methods that are successful at determining the value of Go positions might prove useful for image processing, as the analysis of Go positions is highly visual. *AlphaGo* used

these and other insights, notably training a neural network to screen for a smaller subset of moves (Zaman & Wunsch, 1999), training via reinforcement learning in self-play (Zaman, Prokhorov & Wunsch, 1997) and Monte-Carlo-based methods (Silver, 2009). However, none of these early systems were scaled up to a champion level of play.

In 2016, DeepMind surprised the computer Go community by announcing the development of *AlphaGo*. *AlphaGo* was developed with a novel combination of policy neural networks to select moves and value neural networks to evaluate board positions. In addition, *AlphaGo* used supervised learning from human expert games and reinforcement learning from self-play to improve performance. A novel Monte Carlo simulation method was implemented to utilize input from the value and policy networks to plan moves. The human design effort and computational resources for training were unprecedented. DeepMind was given access to thousands of Google tensor processing units (new hardware specifically designed for machine learning calculations) and trained for months. In their Nature paper, the DeepMind team announced that “our program *AlphaGo* achieved a 99.8% winning rate against other Go programs and defeated the human European Go champion by 5 games to 0” (Silver et al., 2016). Later in 2016, *AlphaGo* defeated the world champion Go player Lee Sedol. In 2017, the improved *AlphaGo* Master defeated a world champion Go player Kie Jie.

The same team announced an even more significant result one year later (Silver, Schrittwieser, et al. 2017). They developed a new system, *AlphaGo Zero*, that learned by pure self-play and did not need historical game data for training (Figure 1). Furthermore, it performed better than the original system, which had been renamed *AlphaGo Lee* (in honor of Lee Sedol). *AlphaGo Zero* was trained in 40 days (albeit on state-of-the-art hardware) compared to months of training for *AlphaGo*. *AlphaGo Zero* became the world’s best Go player, outplaying all other human and computer players (Silver & Hassabis 2017). The system was then generalized so that it could learn arbitrary turn-based games with convincing performance. It was able to learn from win-loss data on several games, notably Chess, which it learned after only 9 hours of training (Silver, Hubert, et al., 2017b). In 1000 games against the strongest rule-based Chess engine, Stockfish, *AlphaGo Zero* won 155 games and lost only 6. This result was beyond what many expected, as computer engines had dominated Chess for decades. Even Demis Hassabis, CEO of DeepMind, noted that it was “far from clear that a program of this type could possibly hope to compete with the specialist handcrafted chess engines that had decades of cumulative effort spent on them from some of the best computer scientists and chess grandmasters in the world” (Sadler & Regan, 2019).

Even in the context of the extensive resources expended by DeepMind,

the accomplishments of *AlphaGo* and *AlphaGo Zero* stand out as major milestones in the history of neural networks as well as AI. Public and private investment in the AI field has dramatically increased since their release. There can be little doubt about the financial value of this neural network-related intellectual property. In the three years since the victory over Kie Jie, Google gained 70% in market value (over \$350 billion). Of course, this gain cannot be specifically attributed to the Go victory. Still, the reputation of Google in neural-network-based reinforcement learning contributes to the company's perceived value. The win generated positive media coverage, bolstered prospects of AI-enabled products, and may have influenced national budgets for AI research. South Korea increased AI spending by \$860 million (Zastrow, 2016). China, trying to draw equal to the U.S. by 2020, is expanding its AI industry to \$150 billion by 2030 (Mozur 2017).

It is not surprising that DeepMind did not release the computer code, training data, or network weights for *AlphaGo* or *AlphaGo Zero*. DeepMind had already invested millions of dollars on computing resources for their creation. The labor and training time expended on failed models was significant. Furthermore, previous progress in Go was a factor in the \$500 million purchase of DeepMind by Google in 2014. The potential economic impact of this research justified the expense of acquisition by Google. In the absence of computer code, training data, or network weights the replication of the findings of *AlphaGo Zero* would be a daunting task.

### Timeline of reverse-engineered Reproducibility Studies of *AlphaGo Zero*



Figure 1: A timeline of *AlphaGo* publications and significant events between 2016 and 2019.

### Case Study Part II: The successful replication of performance of *AlphaGo Zero* by crowdsourced *Leela Zero*

When DeepMind published its first paper on *AlphaGo*, interest grew over the new method. It became international news when *AlphaGo* defeated Lee Sedol in March of 2016. One commentator wrote that

Today is a memorable day in human history. In the very first game of Go between the human champion and the computer

program *AlphaGo*, the human champion lost. Pre-match predictions was predominantly in favor of human. Go players unanimously picked human champion Lee Sedol. (Zhe, 2016)

In response to Sedol's defeat, group efforts were initiated to combine computing resources to replicate the success of *AlphaGo Zero*. Initially, replications efforts went nowhere outside of promises for more investments into AI research. Efforts intensified when DeepMind shared that the methods in *AlphaGo Zero* had been successfully applied to two other games, Chess and shogi. Subsequently, two crowdsourced replication efforts were organized under the leadership of Belgian programmer Gian-Carlo Pascutto who had previously authored a Go program called Leela and a Chess program called Sjeng. New programming efforts were started as Leela Go and Leela Chess and GitHub websites were created for both projects. Pascutto (n.d.) writes that "After the release of *AlphaGo Zero* and Mastering the game of Go without human knowledge, and some maths on the computing power required to duplicate that effort, I decided to start an open project to replicate the results - and get them available to the public - named *Leela Zero*."

The success of *AlphaGo Zero* in learning Chess provided a unique opportunity to re-engineer reproducibility. Chess rule-based game engines had already reached the pinnacle of chess play, whereas Go systems had no reliable comparison tool other than the top-ranked humans. Top-ranked Go players could no longer compete with *AlphaGo*. DeepMind claimed only 9 hours of training were needed for chess compared to the 40 days used for Go training. Although training of a replicate system for chess by crowdsourcing would take longer than needed by DeepMind due to less available resources, crowdsourced training for chess would still be faster than crowdsourced training for Go.

Although DeepMind did not provide source code or neural network weights, DeepMind did give enough details so that a similar system could be constructed with certain assumptions. A replicate system (Leela Chess Zero) was built and released in late 2017 (LCZero, n.d.). The interested community was invited to contribute to the code and the computational effort of training. Many community members suggested improvements for the system on a public forum. At the same time, the GitHub project was expanded to fix issues with the architecture and upgrade it for more efficient computations. Interested people could download a client, which would run simulations on their computers. The results would be sent to a central repository where a new network would be trained after tens of thousands of games. If the new network performed better than the old one, clients would be updated with the new network and continue to run simulations. Around this time, Google Colabs introduced free GPU resources, allowing people to use their free computing

power from Colabs to help contribute if they lacked a powerful computing system (Chess.com, 2018; Colaboratory, n.d.). Training on diverse systems of much lower capacity took longer. But the system gradually moved from slow, weak play to grandmaster level, beating its first human grandmaster in April 2018 (Lystad 2018), and in January 2019, it managed to win the TCEC Cup 2, playing on par with the best rule-based chess agents (LCZero, 2019).

The same crowdsourced approach to training Leela Chess was used with Leela Go. On the *Leela Zero* website Pascutto (Zero n.d.) remarked

Gimme the weights. Recomputing the *AlphaGo Zero* weights will take 1700 years on commodity hardware. One reason for publishing this program is that we are running a public, distributed effort to repeat the work. Working together, and especially when starting on a smaller scale, it will take less than 1700 years to get a good network....

Pascutto was not that hopeful initially about crowdsourcing for *Leela Zero*. He remarked “realistically there will be maybe 10 people joining, 80 if we're very lucky” (Pascutto, 2017). In fact, more than 500 volunteers joined the crowdsourcing effort and contributed to finding the network weights in less than a year. Training of *Leela Zero* for Go began on November 10, 2017, with an Elo rating of 0 after 0 games. By December 14, 2017, *Leela Zero* had been trained on 1,035,000 games and had an Elo rating of 3,077 (higher than most Go expert players). By February 15, 2020, *Leela Zero* had 21,709,096 accumulated games and a super-human Elo rating of 16,726 (Zero n.d.).

Facebook has released an open-source version of their Go AI system named ELF OpenGo, which has defeated top human players 20-0 (Tian & Zitnik, 2018). One of the leading AI companies in China, Tencent, has released an open-source version of its Phoenix Go program, winner of the 2018 World AI competition. This pursuit of independent replications of novel research ideas in AI is a laudable trend.

## **Discussion and Conclusions**

The replication of neural network research findings is a daunting task due to the resource demands of numerous complex simulations. Furthermore, the software development effort to reproduce results is significant. Open-source software and crowdsourcing of computational resources can overcome these challenges. While this approach to reproducibility is difficult in some cases, it should not be dismissed. Advances in high-performance computing and software neural network model development have lowered entry barriers so that an approach that seemed impractical in one year may become feasible the next. The methods described in this case study may become increasingly

common. Several important lessons can be learned from the successful *Leela Zero* project that replicated the findings of *AlphaGo Zero*:

- 1) The project benefited from the leadership of a skilled and experienced programmer. Pascutto had already had experience writing chess and Go programs.
- 2) The project benefited from the availability of the published system architecture described in the first Nature article (Silver et al. 2016).
- 3) Although training data was not available, the use of self-play for training effectively overcame that limitation.
- 4) Participation of volunteers was greater than expected. Pascutto (2018) remarks “I estimated we would get perhaps 10 or so computer Go enthusiasts to run the client. In fact, it's been generally over 500! Similarly, there have been some very high-quality code contributions as well.”
- 5) The use of large numbers of crowdsourced low-powered community computers compensated for the lack of high-powered dedicated computers.
- 6) *Leela Zero* reached expert-level performance in a shorter than expected time of about 30 days.
- 7) The project made effective use of GitHub for software downloads and data uploads.
- 8) A skilled programmer was able to write working code despite the decision of the original developers to withhold the release of their code.

The replication of the findings of *AlphaGo Zero* by *Leela Zero* is encouraging. The independent verification of significant neural network research results should be fostered. Given the accelerating scale of investments in neural networks by diverse stakeholders, we hold that many important contributions will not be described in sufficient detail to allow easy confirmation by the scientific community. However, this should not deter scientists from attempting the reproducibility of the most important findings, with or without the assistance of the original investigators. Confirmation of results by independent investigators lends credibility to novel research findings. Crowdsourced reverse engineering can engage a broad community of researchers in confirming or refuting novel results. The successful reverse engineering of DeepMind *AlphaGo Zero* has demonstrated the feasibility of this approach.

### **Acknowledgments**

Portions of this work were previously made available at arXiv.org as [arXiv:1909.03032](https://arxiv.org/abs/1909.03032). Partial support for this research was received from the



Missouri University of Science and Technology Intelligent Systems Center, The Missouri University of Science and Technology Chancellor's Distinguished Fellowship, the U.S. Department of Education Graduate Assistance in Areas of National Need Program, the Mary K. Finley Missouri Endowment, the National Science Foundation, the Lifelong Learning Machines program from DARPA/Microsystems Technology Office, and the Army Research Laboratory (ARL) under Cooperative Agreement Number W911NF-18-2-0260. The views and conclusions contained in this document are those of the authors. They should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright notation.

### **Disclosures**

The authors have no conflicts of interest to disclose.

### **References:**

1. Allis, L. Victor. (1994). Searching for solutions in games and artificial intelligence.
2. Allis, L.V., Herik, H. J., & Herschberg, I. S. (1990). Which Games Will Survive? Retrieved from <https://books.google.com/books?id=PIZAGwAACAAJ>
3. Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. Nature News, 533(7604), 452. <https://doi.org/10.1038/533452a>
4. Chess.com (2018). Contributing to Leela Chess Zero. Creating the Caissa of Chess engines. Retrieved August 19, 2019, from <https://www.chess.com/forum/view/general/contributing-to-leela-chess-zero-creating-the-caissa-of-chess-engines>
5. CBInsights. (2019). China Is Starting To Edge Out The U.S. In A.I. Investment—C.B. Insights Research. (2019, February 12). Retrieved August 6, 2019, from <https://www.cbinsights.com/research/china-artificial-intelligence-investment-startups-tech/>
6. Chess.com (2018). Contributing to Leela Chess Zero. Creating the Caissa of Chess engines. Retrieved August 19, 2019, from <https://www.chess.com/forum/view/general/contributing-to-leela-chess-zero-creating-the-caissa-of-chess-engines>
7. Colaboratory (n.d.). Frequently asked questions. Retrieved August 19, 2019, from <https://research.google.com/colaboratory/faq.html>
8. Go ranks and ratings (n.d.) Wikipedia. Retrieved as [https://en.wikipedia.org/wiki/Go\\_ranks\\_and\\_ratings](https://en.wikipedia.org/wiki/Go_ranks_and_ratings).

9. Greene, C. S., Garmire, L. X., Gilbert, J. A., Ritchie, M. D., & Hunter, L. E. (2017). Celebrating parasites. *Nature Genetics*, 49(4), 483–484. <https://doi.org/10.1038/ng.3830>
10. Kharpal A. (2017). China wants to be a \$150 billion world leader in A.I. by 2030. (2017, July 21). Retrieved August 6, 2019, from <https://www.cnbc.com/2017/07/21/china-ai-world-leader-by-2030.html>
11. Koch C (2016). How the Computer Beat the Go Master. Retrieved at <https://www.scientificamerican.com/article/how-the-computer-beat-the-go-master/>
12. Krikke, J. (2007). The Challenge of Go. *IEEE Intelligent Systems*, 22(01), 7–8. <https://doi.org/10.1109/5254.747899>
13. LCZero. (n.d.). Retrieved February 23, 2019, from <http://lczero.org/>
14. LCZero (2019). Leela won the TCEC CUP! Retrieved February 23, 2019, from <http://blog.lczero.org/2019/02/leela-won-tcec-cup.html>
15. Levinovitz, A. (2014). The Mystery of Go, the Ancient Game That Computers Still Can't Win. Retrieved at <https://www.wired.com/2014/05/the-world-of-computer-go/>
16. Lystad, T.A. (2018) GM Andrew Tang Defends Humanity Against Leela Chess Zero. Retrieved August 19, 2019, from Lichess.org website: <https://lichess.org/blog/WtzZAyoAALvE8ZSQ/gm-andrew-tang-defends-humanity-against-Leela-chess-zero>
17. Mozur, P. (2017) Beijing Wants A.I. to Be Made in China by 2030. *The New York Times*. (2017, July 20). Retrieved August 6, 2019, from <https://www.nytimes.com/2017/07/20/business/china-artificial-intelligence.html>
18. National Academies of Sciences, Engineering, and Medicine. (2019). *Reproducibility and replicability in science*. National Academies Press. <https://doi.org/10.17226/25303>
19. Pascutto GC. (2017). [Computer-go] Zero performance. Retrieved at <http://computer-go.org/pipermail/computer-go/2017-October/010307.html>
20. Pascutto GC (2018). Gian-Carlo Pascutto - The man behind LeelaZero. Retrieved at <https://www.eurogofed.org/?id=205>
21. Sadler, M., & Regan, N. (2019). *Game Changer: AlphaZero's Chess Strategies and the Promise of A.I.* Retrieved from <https://books.google.com/books?id=KhhivwEACAAJ>
22. Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., ... & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587), 484-489. Retrieved from <http://arxiv.org/abs/1712.01815>

23. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T. and Lillicrap, T., (2017). Mastering chess and shogi by self-play with a general reinforcement learning algorithm. arXiv preprint arXiv:1712.01815.
24. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Bolton, A. (2017b). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354.
25. Silver, D., Sutton, R. S., & Müller, M. (2007, January). Reinforcement Learning of Local Shape in the Game of Go. In *IJCAI* (Vol. 7, pp. 1053-1058)
26. Tian, Y. & Zitnik, K. (2018). Facebook Open Sources ELF OpenGo. Retrieved at <https://research.facebook.com/blog/2018/05/facebook-open-sources-elf-opengo/>.
27. Zaman, R., Prokhorov, D., & Wunsch, D. C. (1997). Adaptive critic design in learning to play the game of Go. *Proceedings of International Conference on Neural Networks (ICNN'97)*, 1, ICNN'97.1.1. <https://doi.org/10.1109/ICNN.1997.611623>
28. Zaman, R., & Wunsch, D. C. (1999). T.D. Methods applied to the mixture of experts for learning 9/spl times/9 Go evaluation function. *IJCNN'99. In IJCNN'99 Joint Conference on Neural Networks. Proceedings (Cat. No.99CH36339)*, 6, 3734–3739 vol.6. <https://doi.org/10.1109/IJCNN.1999.830746>
29. Zastrow, M. (2016). South Korea trumpets \$860-million AI fund after *AlphaGo* 'shock'. *Nature*. <https://doi.org/10.1038/nature.2016.19595>
30. Zero (n.d.) Retrieved at <https://zero.sjeng.org/>
31. Zhe, L. (2016) [Game 1] Lee Sedol's Strategy and *AlphaGo*'s weakness. Retrieved at <https://massgoblog.wordpress.com/2016/03/11/lee-sedols-strategy-and-AlphaGos-weakness/comment-page-1/#comment-25>