Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Learning to Play 7 Wonders Duel Without Human Supervision

(Article begins on next page)

18 October 2024

Learning to Play 7 Wonders Duel Without Human Supervision

Giovanni Paolini
University of Bologna
Bologna, Italy
g.paolini@unibo.it
0000-0002-3964-9101

Lorenzo Moreschini
University of Brescia
Brescia, Italy
lorenzo.moreschini@unibs.it
0000-0002-4203-3359

Francesco Veneziano
University of Genoa
Genoa, Italy
veneziano@dima.unige.it
0000-0002-2225-7769

Alessandro Iraci
University of Pisa
Pisa, Italy
alessandro.iraci@unipi.it
0000-0002-3158-3929

Abstract—This paper introduces ZeusAI, an artificial intelligence system developed to play the board game 7 Wonders Duel. Inspired by the AlphaZero reinforcement learning algorithm, ZeusAI relies on a combination of Monte Carlo Tree Search and a Transformer Neural Network to learn the game without human supervision. ZeusAI competes at the level of top human players, develops both known and novel strategies, and allows us to test rule variants to improve the game's balance. This work demonstrates how AI can help in understanding and enhancing board games.

Index Terms—Artificial Intelligence, Reinforcement Learning, 7 Wonders Duel, Board Game Strategy, Board Game Enhancement.

I. INTRODUCTION

7 Wonders Duel (7WD) is an award-winning board game published in 2015, as a 2-player version of the world-famous multi-player game 7 Wonders. As of May 2024, 7WD holds the 19th place in the BoardGameGeek ranking,[1] and is the 8th most played game on Board Game Arena (BGA)[2] with 8.7 million games played since its introduction in 2020. Its presence on BGA has fostered competitive play and the development of deep strategies, also thanks to frequent tournaments as well as seasonal and all-time rankings.

This paper introduces ZeusAI, an AI system capable of matching, and potentially outperforming, top human players at 7WD. To the best of our knowledge, this is the first publicly declared AI that plays 7WD at an expert level. We train ZeusAI using Deep Reinforcement Learning, with no supervision from games played between humans or from human-developed strategies. Our approach is similar to AlphaZero [1], the algorithm developed by DeepMind to achieve super-human performance in games such as Go, Chess, and Shogi.

ZeusAI differs from AlphaZero in two key aspects. First, it has a mechanism to bound the number of random continuations considered when making a move (this is not needed in AlphaZero, as games like Go and Chess have no randomness). Second, it employs a Transformer [2] in place of a Convolutional Neural Network. Transformers have revolutionized

natural language processing and, more recently, the whole field of AI [3]; unlike Convolutional and Recurrent Networks, Transformers do not have a built-in assumption of spatial or temporal proximity in the input, making them better-suited to handle the complex state of a 7WD game.

We are currently in the process of testing ZeusAI against some of the best players in the world. To this end, we have been holding best-of-5 live matches[3] as well as an ongoing series of offline games. ZeusAI has won 26 out of the 38 games (68.4%) played so far against top human players.

ZeusAI rediscovered well-known strategies to play 7WD, but it also frequently makes moves that are counter-intuitive to expert players, suggesting that 7WD is far from solved. We use ZeusAI to derive numerous statistics of interest, such as a ranking of the wonders and progress tokens. We also confirm the known belief that the 1st player holds a significant advantage, estimating a 66.8% win rate. We propose and evaluate several variants that reduce this imbalance to as little as 51.6%, while keeping changes to a minimum; we hope that these variants will gain traction in the 7WD community.

II. THE GAME OF 7 WONDERS DUEL

7WD is a game with no hidden information, and so it admits a theoretically optimal strategy (like Go and Chess) with no dynamics such as bluffing. However, unlike Go and Chess, 7WD has elements of randomness.

In 7WD, 2 players compete to develop their cities and win in one of 3 ways: civilian (by obtaining the most points), scientific (by acquiring 6 distinct scientific symbols), or military (by advancing the conflict pawn to the opponent's capital on the military track). These different victory conditions introduce considerable strategic depth, as players need to carefully balance all of them.

The game components include 10 progress tokens, 12 wonders, and 73 building cards, all granting permanent bonuses or one-time effects. Notably, 5 of the wonders grant the owner an immediate extra turn when constructed.

During game setup, 5 progress tokens are randomly revealed, and then the *wonder selection phase* begins, consisting

---

[1]https://boardgamegeek.com/browse/boardgame
[2]https://boardgamearena.com

[3]These matches are streamed live on YouTube at https://www.youtube.com/playlist?list=PLA4iH7nzdJ7sNWxqAn2JxB_zQ0nEdoJeg

Fig. 1. Structure of face-up cards (white) and face-down cards (colored) in ages I, II, and III (from left to right).

Fig. 2. The 3 most frequently built cards by ZeusAI in ages I, II, III, and among the guilds. In every group, the most frequently built card is shown in front.

of two *drafts*. In the first draft, 4 wonders are randomly revealed and selected following the 1^{st}-2^{nd}-2^{nd}-1^{st} player sequence. In the second draft, 4 more wonders are revealed and selected according to the 2^{nd}-1^{st}-1^{st}-2^{nd} sequence. This way, each player obtains 4 wonders to build throughout the game.

Gameplay progresses through 3 ages, each involving a card selection mechanism from a unique structure of face-up and face-down cards (see Fig. 1). A card can only be selected if it is not covered by any other, with face-down cards being revealed once all cards on top of them have been removed. Players collect various types of cards to build their cities: brown and grey cards for resources, yellow for commercial advantages, red for military power, green for scientific symbols, blue and purple for victory points (purple cards are called *guilds*). Some of the cards are shown in Fig. 2.

On each turn, a player must choose a card and either: pay its cost to add it to their city and gain its benefits; use it to activate one of their wonders; or discard it to gain coins (which in turn can be used to purchase resources). When a player collects two identical scientific symbols, they can pick one of the available progress tokens.

Only 7 wonders can be built in total. This rule adds pressure to the players, as it is typically convenient to construct all 4 available wonders before the opponent does. Some of the most critical decisions in a game revolve around finding the correct timing to build the wonders.

III. ZEUSAI

ZeusAI is an adaptation of the celebrated AlphaZero reinforcement learning algorithm [1], combining Monte Carlo Tree Search (MCTS) with a Deep Neural Network that models state values and policy. In this section, we give a technical overview of ZeusAI. Some details are omitted for brevity, and we defer a complete description to a future paper.

A. The Model

In 7 Wonders Duel, unlike Go and Chess, the state does not have a chessboard-like geometry, making Convolutional Neural Networks not well-suited. Instead, we use Transformer Encoder models [2] of different sizes. Our final model has 12 layers and 12 attention heads; it maintains 768-dimensional representations, and the feedforward layers have a hidden dimension of 3072. A state is represented as a set of components (cards, wonders, progress tokens, etc.), each paired with their positional information (e.g., a card can be in a specific location of the age structure, in either player's city, discarded, or used to build a wonder). All components and positions have a learned 768-dimensional embedding, akin to words and positions in Transformer-based language models. No explicit information about the rules, and the costs and effects of game components, is encoded; it is all learned through self-play. Additional tokens encode other aspects of the state, such as the number of coins and the status of the military track.

ZeusAI's Transformer outputs a state value between -1 and $+1$, representing the odds of the current player winning the game, and a probability distribution over all legal actions (the policy), with more promising actions having higher probabilities. The total number of trainable parameters is $\sim$92M.

B. Action Selection

To select an action in a given state, ZeusAI employs the MCTS algorithm. Specifically, the procedure begins with a tree consisting only of the root node, representing the current state. During each simulation, the tree is traversed starting from the root, until a new state is encountered and added; actions are chosen in a way that balances exploration (prioritizing less explored actions), and exploitation (favoring actions considered more promising by the model or that have previously led to high-value outcomes). In some cases, applying an action results in a state where a random event, such as revealing one or more cards, must occur. We refer to these states as *afterstates*. To reduce the branching factor of the search tree, each afterstate is limited to having at most 11 child states; in non-training games, this limit is gradually increased as the number of visits to an afterstate increases. If a simulation reaches an afterstate already at the limit, it continues randomly with one of the existing child states. The limit of 11 is based on the maximum number of child states possible when a single card must be revealed.

When generating self-play games for training, we run MCTS for up to 1k simulations; the action is selected randomly, weighted by the number of visits. In non-training games, we run up to 5k simulations and deterministically select the action with the highest number of visits.

C. Training

We initially generated 35k games using three simple rule-based policies: the first policy selects actions at random, with a preference for not discarding; the second and third policies attempt to build random green and red cards, which are required to achieve scientific and military victories respectively,

Fig. 3. Distribution of victory types and winners (1st or 2nd player) in self-play games throughout training.


Fig. 4. Wonder preference in self-play games throughout training.

and otherwise behave as the first policy. We trained the value head of our Transformer model on all states from these games, assigning a target value based on the game's outcome ($+1$ for a win, -1 for a loss, and 0 for a draw).

Subsequent self-play games were generated with ZeusAI. Every $\sim$3k games, the Transformer model was re-trained on the states from the most recent $\sim$100k games. Specifically, the policy head was trained to predict the visit distribution of the legal actions at the MCTS root node. We generated a total of $\sim$420k self-play games for training, which would take 9 years of real-time training for a human player on BGA.

IV. RESULTS

A. Analysis of the Training Games

Fig. 3 illustrates the evolution of victory types and winners in the self-play games generated by the various iterations of ZeusAI, as a function of the cumulative number of games up to each iteration. Draws only occurred in about 0.1% of the games and are not shown.

ZeusAI quickly learned to exploit the 1st player advantage, but it took $\sim$100k games for the scientific victory to be "discovered" as a viable strategy, due to the very sparse reward obtained when collecting a diverse set of science symbols. The discovery of the scientific victory also marked the moment


Fig. 5. Wonder preference of ZeusAI (top) and human players (bottom).

when we stopped being able to consistently beat ZeusAI and felt confident that it had reached a competent level of play.

Fig. 4 shows the evolution of ZeusAI's wonder preference throughout training, computed averaging a numerical value based on whether the wonder was selected first (100%), second or third (50%), or fourth (0%) during the draft. The Pyramids and the Great Lighthouse became less preferred as ZeusAI learned to exploit wonders offering more strategic benefits.

B. Games Against Human Players

We turned to the BGA community to further assess ZeusAI's level of play. Between January and April 2024, we organized best-of-5 challenges against 3 of the top 7WD players on BGA. These events were streamed live on YouTube and accompanied by our analysis of the games (the link is provided in Section I). Two of the challenges ended with ZeusAI's victory (3–2 and 3–1) and one with a defeat (2–3). We also started a separate series of games against one of the three players, and ZeusAI is currently leading 18–6. Overall, ZeusAI has won 26 out of all the 38 games played so far against top players, confirming the impression that it has achieved a very high level of play.

C. Strategic Insights

Victory types: After training ZeusAI, we generated 10k self-play games to extract statistics and insights. The distribution of victory types is 61.7% civilian / 21.4% scientific / 16.9% military, which is not far from the 58.0% / 25.6% / 16.4% distribution compiled from games played by good players on BGA [4].

Wonder selection: Fig. 5 compares ZeusAI's wonder preference with that of human players on BGA [4]. The 5 extra-turn wonders rank as ZeusAI's favorite, and their preference rates are nearly identical to humans. On the other hand, the Pyramids are almost always selected last. Among the remaining 6 wonders, ZeusAI favors the Statue of Zeus and dislikes the Great Lighthouse. While human experts recognize the strength of the Statue of Zeus, such a strong preference seems in disagreement with how humans currently play 7WD.

Building wonders in age I: ZeusAI builds an extra-turn wonder 52% of the times when there are 2 cards left in age I, and 12% of the times when there are 4 cards left. This is a known tactic that allows a player to obtain more cards in age I and also start age II. ZeusAI frequently employs this tactic even with a single extra-turn wonder left, while human players tend to be more conservative.

Fig. 6. Progress tokens ranked by frequency of selection in ZeusAI games.

TABLE I
ESTIMATED WIN RATE OF THE 1ST PLAYER IN THE COIN VARIANT

		2nd player coins			
		7	8	9	10
1st player coins	7	66.8%	65.5%	63.4%	58.8%
	6	63.7%	62.4%	59.2%	53.6%
	5	62.5%	58.3%	59.5%	54.4%

ZeusAI also appears to build wonders earlier than human players usually would. For instance, in age I, ZeusAI builds an average of 0.83 wonders as 1st player and 0.49 as 2nd player.

Cards: Fig. 2 shows the cards most frequently built by ZeusAI in each age (which are not necessarily the cards that are built earlier). In age I, the Tavern is the most built card, with production cards (brown and gray) being built more often than the remaining yellow cards. Among resources, papyrus is preferred to glass, and wood is preferred to stone and clay. ZeusAI often builds blue cards, even in age I and age II, to try to establish an early lead in victory points.

Progress tokens: Fig. 6 shows all progress tokens sorted by how frequently ZeusAI selects them. As expected, the top 3 positions are held by Theology, Law, and Strategy, which are considered to be the most powerful tokens in the game.

D. Balancing 7 Wonders Duel

It is known in the community that the 1st player has a significant advantage in 7WD. Self-play games by ZeusAI confirm this belief: out of 10k games, the 1st player had a 66.8% win rate (see also Fig. 3). This is substantially more than the estimate of 55.7% based on 605 games between human players on BGA [4], further supporting the idea that ZeusAI is stronger than humans. In the rest of this section, we propose rule variants that bring the win rate closer to 50%.

Coin variant: A natural way to balance the game is to change the amount of coins given to each player during the setup. We tried several options, estimating the win rate by generating 2k self-play games with the new setup, and computing the fraction of games won by the 1st player (see Table I, where the 95% confidence intervals have a margin of error of $\sim$2% for setups other than 7-7).

Even a coin difference as large as 5 in favor of the 2nd player is not enough compensation. Indeed, a 1-coin change roughly corresponds to a 2.7% change in the win rate.

Wonder selection variants: The other variants we propose consist of altering the wonder selection phase, so that the advantage of starting age I is balanced by weaker wonders.

To estimate the win rate in each variant, we cannot directly generate self-play games since ZeusAI is not able to handle the new rules. However, we can exploit the fact that afterstates reached at the end of the wonder selection phase are valid afterstates for the standard 7WD rules. Specifically, we use ZeusAI's Transformer to estimate the value of all such possible afterstates; then, for each variant, we use minimax to compute the optimal strategy during wonder selection (with the new rules), where the outcome is given by the value of the resulting afterstate. With this method, we estimate a 67.6% win rate for the 1st player with the standard rules, which is fairly close to our aforementioned estimate of 66.8%.

In these variants, the player making the first decision is not necessarily the same as the player starting age I. Therefore, we refer to the two players as *Aristotle* and *Cleopatra*. We report the estimated win rate for Cleopatra.

- **Variant 1** (60.0% win rate): After seeing the first 4 wonders, Aristotle decides whether to start the first draft; the other player starts the second draft. Cleopatra starts age I. This variant is known within the 7WD community.
- **Variant 2** (54.6% win rate): In both the first and second draft, after seeing the 4 wonders, Aristotle decides whether to start the draft. Cleopatra starts age I. The win rate is comparable to the expected score for white in Chess, deemed acceptable for competitive play.
- **Variant 3** (51.6% win rate): Aristotle sees all 8 available wonders; he assigns 4 to the player starting age I and the remaining 4 to the other player. Cleopatra then decides whether to start age I. This can be considered an "expert variant", as optimal play requires an advanced understanding of the strengths of the wonders.

V. CONCLUSION AND FUTURE WORK

In this paper, we have described our ongoing efforts in building, evaluating, and using ZeusAI to understand and improve 7WD. Going forward, we intend to further assess ZeusAI's strengths and weaknesses, discover deeper insights into 7WD strategy, and playtest ideas to balance the game.

REFERENCES

[1] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis, "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play," *Science*, vol. 362, no. 6419, pp. 1140–1144, 2018.
[2] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
[3] OpenAI, "Introducing ChatGPT," https://openai.com/blog/chatgpt, 2022, accessed 2024-04-24.
[4] Le007n, "7WD in numbers," https://forum.boardgamearena.com/viewtopic.php?t=25369, 2022, accessed 2024-04-24.