

# Blood Bowl: The Next Board Game Challenge for AI

Niels Justesen

IT University of Copenhagen  
noju@itu.dk

Sebastian Risi

IT University of Copenhagen  
sebr@itu.dk

Julian Togelius

New York University  
julian@togelius.com

## ABSTRACT

We propose the popular board game *Blood Bowl* as a new challenge for Artificial Intelligence (AI). *Blood Bowl* is a fully-observable, stochastic, turn-based, modern-style board game with a grid-based playing board. At first sight, the game ought to be approachable by numerous game-playing algorithms. However, as all pieces on the board belonging to a player can be moved several times each turn, the *turn-wise* branching factor becomes overwhelming for traditional algorithms. Additionally, scoring points in the game is rare and difficult, which makes it hard to design heuristics for search algorithms or apply reinforcement learning. We present our work in progress on a game engine that implements the core rules of *Blood Bowl* with a forward model and a reinforcement learning interface. We plan to release the engine as open source and use it to facilitate future AI competitions.

## CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; *Machine learning*; *Reinforcement learning*; *Neural networks*; *Game tree search*; • **Software and its engineering** → **Interactive games**;

## KEYWORDS

Artificial Intelligence, Machine Learning, Board Games

### ACM Reference Format:

Niels Justesen, Sebastian Risi, and Julian Togelius. 2018. Blood Bowl: The Next Board Game Challenge for AI. In *Proceedings of Workshop on Tabletop Games (Foundations of Digital Games 2018)*. ACM, New York, NY, USA, Article 4, 2 pages. [https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)

## 1 BLOOD BOWL

*Blood Bowl* is a board game designed by Jervis Johnson in 1986 and is published by Games Workshop. The game is very popular among competitive tabletop gamers with 161,080 recorded tabletop tournament matches<sup>1</sup>.

### 1.1 A Game of Fantasy Football

*Blood Bowl* is a so-called *fantasy football game* that is played on a board of  $26 \times 15$  squares mimicking a football/rugby-like field. Two players each control a team of *Blood Bowl* miniatures and the goal is to score the most touchdowns. We will refer to players as *coaches* and the miniatures as *players*. Each coach can field 11

<sup>1</sup>[http://naf.talkfantasyfootball.org/total\\_for\\_all\\_competitions.html](http://naf.talkfantasyfootball.org/total_for_all_competitions.html)

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*Foundations of Digital Games 2018, August 2018, Malmö, Sweden*

© 2018 Copyright held by the owner/author(s).

ACM ISBN 123-4567-24-567/08/06.

[https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)

players on the board whereafter coaches take turns to move all their players. Players can either move, pass, hand-off, block (attempt to knock down opposing players), blitz (move and block) or foul (stomp on down players) during their *player turn*. When the ball carrier reaches the opponent's end zone their team (coach) scores a point. Determined by their *Movement Allowance* players can move several squares. Blocks, passes, catches and moving adjacently to opposing players require dice rolls to succeed which depends on the player's *Strength*, *Agility*, and *Armor* attributes, as well as skills such as *Dodge*, *Pass*, and *Block*. Failed dice rolls typically end the coach's turn. The game thus requires intelligent risk management and planning every turn. After two halves, with eight turns for both players, the game ends.

The rules have evolved over time and today most players follow the almost identical *Living Rulebook 5<sup>2</sup>* or *6<sup>3</sup>* or the *Blood Bowl 2016 Edition - Death Zone 2* ruleset. A video game adaptation with 3D graphics was released by Cyanide Studios in 2009 which features online play. The video game also includes an AI but it is far from human-level; it presumably follows a set of scripted rules combined with a pathfinding algorithm. FUMBBL (the acronym combines the football term *Fumble* with BBL; *Blood Bowl League*) is a community-driven online league with more than 2,400,000 recorded games. Matches in FUMBBL are played using an unofficial game client with simple 2D graphics<sup>4</sup>.

### 1.2 Characteristics

This section contains a short analysis of *Blood Bowl*'s characteristics using the dimensions defined in Yannakakis and Togelius [6].

*Blood Bowl* has **perfect information** as the board state is fully observable and players have no hidden information. The game has an optional rule that allows players to have secret special play cards but these are, to our knowledge, never used in competitions.

*Blood Bowl* is **stochastic** as most of the interesting actions require dice rolls to succeed. Coaches have a limited number of re-roll tokens they can use to re-roll a failed roll. Experienced *Blood Bowl* players usually start their turn with safe actions that require easy dice rolls, or preferably none at all, and postpone risky actions till the end of their turn.

*Blood Bowl* is a **turn-based** and a **multi-action** game as coaches take turns to move *multiple* players on the board. Another multi-action game that has been the basis for research on AI methods is *Hero Academy* (Robot Entertainment, 2012) [2]. What makes *Blood Bowl* even more complicated is that players can be moved several steps each turn. A *coach turn* thus contains multiple *player turns* that each allows a sequence of actions.

The **state representation** is especially relevant for deep learning methods. Go and most Atari games are particularly suitable

<sup>2</sup><http://www.bloodbowlonline.com/Rules.shtml>

<sup>3</sup>[http://www.bbaa.org/downloads/NewTeams\\_LRB6.pdf](http://www.bbaa.org/downloads/NewTeams_LRB6.pdf)

<sup>4</sup><https://fumbbl.com/>



Figure 1: The main part of the user interface in FFAI.

for deep learning methods as they have an image, or image-like, state representation as well as a fixed action space, which is not the case for Blood Bowl. Here, the state is represented by players on the board, each with multiple dimensions (player attributes and whether it is standing, knocked down etc.), a dugout for both players with reserve, knocked out, and injured players, weather conditions, and occasionally a dice roll. The state representation thus consists of both multi-layered spatial features and non-spatial features very similar to the representation in SC2LE [5]. Additionally, the **action space** in these two games varies between steps.

### 1.3 Complexity

The action space in Blood Bowl varies between 1 and 395. Sometimes the coach has to select between a few dice results and other times one of 395 squares on the board to kick or pass the ball to. In most situations, the coach has to select one of eight adjacent squares to move a player to or select one of six different action types for a player. For simplicity, we will estimate the average *step-wise* branching factor as 10. With 10 players that can each move 5 squares, depending on several factors, the average *turn-wise* branching factor is approximately  $10^{10 \times 5} = 10^{50}$ . In comparison, the *turn-wise* branching factor in Chess is 30 and 300 in Go.

Long action sequences with sparse scores make both search algorithms and reinforcement learning harder to apply. A game of Blood Bowl consists of approximately  $5 \times 10 \times 32 = 1600$  steps (with high variance), using the numbers previously estimated multiplied by 32 turns. Games usually end with around 0–3 points per team.

## 2 GAME ENGINE

Existing Blood Bowl implementations are closed source and do not have an AI interface. Thus, we are currently developing our own game engine. To avoid legal issues with the trademark owners Games Workshop, our game engine will be named the *Fantasy Football AI* (FFAI) client and will not include any copyrighted artwork or trademarked names. Figure 1 shows a screenshot of the current user interface in FFAI with our own 2D graphics<sup>5</sup>.

FFAI is implemented in Python allowing a simple way to interface with popular machine learning libraries. We considered implementing the engine in C++ with a Python interface on top, but the state updates in Blood Bowl are fairly simple, and thus fast, even in Python. FFAI will implement the Open AI Gym interface [1]

<sup>5</sup>The player icons are copyright protected by Nicholas Kelsch. We are still awaiting permission to use them in FFAI.

with the exception of the variable action space. Similar to SC2LE [5], the observation object will include several spatial feature layers as well as several non-spatial features. Aside from the reinforcement learning interface, the engine itself can be used as a forward model. In our current version, one game step with a randomly sampled action takes on average 130 microseconds on a regular laptop. With our estimated 1,600 steps per game, it will take 0.21 seconds to simulate a complete game with random actions and 0.0065 seconds for a turn.

## 3 AI COMPETITION

Games have proven to be a successful testbed for Artificial Intelligence (AI) and in the last few years deep reinforcement learning has enabled computers to learn how to play games such as Chess [4], Go [4] and Atari games [3]. This has unfortunately led to a common misconception that computers can now play all interesting board games. Based on our brief analysis we believe that Blood Bowl can offer a new exciting type of testbed for AI due to the high complexity of the game, while still resembling many classical board games. We further believe that this testbed will be beneficial for research in areas such as tree search, evolutionary planning, and reinforcement learning. Our plan forward is to organize an annual AI competition using FFAI. The competition will allow all types of methods, including controllers that are scripted, search-based, neural network-based, as well as hybrids combining some of these approaches. We hope that both academics and hobbyist will participate in the competition and aim to write a follow-up paper describing the details of FFAI and the forthcoming competition as well as results of several state-of-the-art methods. Currently, we have tested a random agent in 350,000 games against itself and it never managed to score a point. This is remarkable and shows that proper planning is required to play this game. Games in which random agents never (as in almost never) score points or wins (playing itself) are extremely challenging for many algorithms such as Monte Carlo tree search and Q-learning as they rely on random exploration. Thus, we do not expect vanilla implementations of such algorithms to score any points either. FFAI and our competition will also include versions of the game with smaller boards sizes. Blood Bowl is a relatively simple testbed, which allows exploring new ways to overcome highly complex domains without having to deal with hidden information and real-time decision making.

## REFERENCES

- [1] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. arXiv:arXiv:1606.01540
- [2] Niels Justesen, Tobias Mahlmann, Sebastian Risi, and Julian Togelius. 2017. Playing Multi-Action Adversarial Games: Online Evolutionary Planning versus Tree Search. *IEEE Transactions on Computational Intelligence and AI in Games* (2017).
- [3] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [4] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmarajan Kumaran, Thore Graepel, et al. 2017. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. *arXiv preprint arXiv:1712.01815* (2017).
- [5] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. 2017. StarCraft II: a new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782* (2017).
- [6] Georgios N. Yannakakis and Julian Togelius. 2018. *Artificial Intelligence and Games*. Springer. <http://gameaibook.org>.