

Strategic Sorcery: Automated Planning for ‘Magic: The Gathering’

Nicholas Tillo*, Raksha Rehal*, Christian Muise

Queen’s University
Kingston, Canada

{ tillo.nicholas, raksha.rehal, christian.muise }@queensu.ca

Abstract

‘Magic: The Gathering’ (MTG) is a strategic card game that includes complex interactions and competitive play, using a catalogue of over 20,000 unique cards. In this demonstration, we explore how numeric planning can be used to tackle an interesting subset of the game mechanics. Strong gameplay strategies are produced using the ENHSP planner, and our work serves as both an interesting testbed for numeric planners and a foundation for more elaborate strategic card-game settings modelled in PDDL.

1 Introduction

Many trading card games, such as Magic: The Gathering (MTG)¹, provide rich and complex play experiences; they often challenge new players with game states that require nuance and experience to navigate. While MTG games can often be modelled in a simple state space representation, the resulting complexity arising from the actions taken in any single state can make it challenging for players to determine the best strategies to guarantee optimal play. Automating this process using planners can provide valuable insights into the best strategies that may be employed during gameplay and the mechanisms of the game itself (an example intermediate state of the game is shown in Figure 1). It is this motivation that leads us to explore how we might model MTG using planning technology.

The goal of creating an abstract planning model of MTG is to give answers on optimizing play and allowing new players to receive streamlined feedback. The cards demonstrate a variety of complex features that may be tricky to grasp and work with, especially for beginners to the game. Hence, this model aims to teach newcomers what actions are the best to take while using high-level abstractions of core MTG gameplay. It also offers a rich setting for planners to be tested on.

Because much of the game involves numeric aspects (including card strength, player health, and creature



Figure 1: The starting state of the fourth problem file, showing hand, board state, and health.

attack levels), we opted to use the numeric fragment of PDDL 2.1 (Haslum et al. 2019). In this demo, we cover some of the high-level ideas of the model, along with specific examples of the game dynamics in PDDL.

2 A Model for MTG

We modelled a simplified version of MTG in PDDL in order to generate an optimized sequence of actions such that each move will lead the player to win the game. The model showcases a simple adversarial game between (1) the player and (2) the opponent. A ‘win’ for the player is modelled as a PDDL goal that requires the planner to get the opponent’s life total to 0.

For clarity, the following is a summary of the simplified model that we will be representing: each agent will start with a predetermined integer “life-total” (e.g. 20), and the goal of the game will be for the player to get their opponent’s life total to zero. We will assume that the player has not played any cards yet, and that the opponent already has some cards on the board. There are 4 phases of play for the player; “main”, “attack”,

*These authors contributed equally.

¹<https://magic.wizards.com/>

```

1 (:action attack_opponent
2   :parameters (?creature - creature)
3   :precondition
4     (and
5       (> (current_life_total_player) 0)
6       (is_owned_by_player ?creature)
7
8       ; checking that it is able to tap
9       (not (is_tapped ?creature))
10      (not (is_summoning_sick ?creature))
11
12      (current_phase attack)
13      (not (is_in_hand ?creature))
14      (not (is_dead ?creature))
15    )
16   :effect
17     (and
18       (when
19         (not (has_vigilance ?creature))
20         (is_tapped ?creature)
21       )
22       (is_attacking ?creature)
23       (not (is_blocked ?creature))
24     )
25 )

```

Figure 2: PDDL snippet for the action to attack an opponent.

colour, *spell_effect*, and *phase* will have associated constants that allow us to apply non-variables that represent implicit game rules. 90

The opponent will be represented as having no cards in hand, no lands played, and starting with all their creatures on their board. The opponent is not able to affect the game state at all, outside of performing routine attacks and blocks as much as possible during their delegated phases. Routine attacks are defined as the opponent attacking with every creature possible on their turn, and routine blocking is defined as blocking each creature possible. In addition, it is assumed that neither player has any cards remaining in their deck for simplicity. It will be a very one-sided fight since the goal of this approach is to model a sequence of actions on the player’s end using the core gameplay features of MTG. Hence, the planner will attempt to create a path to victory for the player from the given game state. This is a simplified model of the 2-player aspect to the game, but still provides a compelling setting for planners. 95 100 105

Figure 2 showcases one such action that a correct path may take. This action is a representation of the player’s creatures attacking the opponent. The planner may decide whether or not to attack if the corresponding preconditions are true – those being that the player is not dead and the creature is ready to attack (i.e. not in hand, not summoning sick, and not dead). The effects of a creature attacking the opponent is that it gains the status of ‘attacking’, and it either taps or remains untapped if it has the keyword ‘vigilance’. Additionally, this action specifies that the creature is not yet blocked after they are declared to be attacking. This action is just one example of the scope of the model. 110 115 120

A new game state can be defined for solving by the creation of a new problem file that outlines the desired game state, which can be run using the ENHSP 2020 planner (Scala et al. 2016) in conjunction with the original domain file. The output will either produce an optimal plan, or no plan at all if there is no such optimal plan. Figure 1 depicts one such problem file showing the cards that start in the player’s hand and the current board state of the game. 125 130

3 Discussion

The model was able to reliably find present optimal mid-length plans for the player’s victory if the opponent played using the level-zero strategy (Stahl 1993) that was encoded for them – i.e., just attacking and defending wherever possible in a deterministic way. The outputs produced by the planner are able to be analyzed for their strategic meaning; for example, the main priority for the planner while in a disadvantaged state was to protect the player’s life total. It used the player’s creatures to block even if this meant that the creature would die. In contrast, when met with an advantaged state, the planner decided to forgo blocking, and instead use the creature to attack, trading the player’s life for damage on the opponent. This mimics the infamous MTG adage “health is a resource”. This unique 135 140 145

“block”, and “end”. The player can play 1 “land” on each of their turns. Lands will “tap” for 1 “mana” of a certain colour, and the player can use their total mana per turn to play “creatures” or “spells”. Once a card is “tapped”, it effectively becomes unusable for any further actions, until the beginning of the player’s next main phase; at which point everything that was tapped on their previous turn becomes untapped. The player can play creatures only on their main phase, and each creature will have 3 aspects to them: their power (integer), their toughness (integer) an optional keyword that has a certain effect (string). Creatures cannot “tap” the turn they enter as they have “summoning sickness” for that turn. In the attack phase, any creature that is not tapped or summoning sick may attack the opponent, with the intention of dealing damage equal to its power to the opponent’s life-total. To defend oneself, the opponent may assign an untapped creature to “block” one other attacking creature (from the opposing agent’s end). If blocked, the attacking creature assigns damage equal to its power to the blocking creature’s toughness, and vice-versa, destroying creatures with negative toughness. After these attacking and blocking phases, the player will progress to the end phase, where the turn will be “cleaned up” for the player in order to prepare for their next main phase. After the opponent attacks and processes damage, the turn passes back to “main” where the player will take their next turn in a similar fashion, once again. 60 65 70 75 80 85

From a planning perspective, the defined object types are *colour*, *land*, *creature*, *spell_effect*, *player*, and *phase*.

way of thinking elevates the player’s strategies, causing them to explore other possible actions. Thus, the plans produced by ENHSP aligned with verifiable strategic lines of play, even when the problems themselves were extremely simplified. The same strategies and thought-processes output by the planner can be applied with confidence to high-complexity game states. Hence, the outputted plans illustrated realistic and feasible approaches that could be analyzed and studied to better help new players make complex decisions. In spite of the complexity of the mechanisms of MTG, this paper has shown that planning models are increasingly capable of solving difficult problems.

One limitation of the model is the hard-coded routine attacks and blocks for the opponent. The fact that the opponent attacks and blocks using every creature each turn restricts the number of meaningful situations. This factor may affect the model, and a more dynamic enemy choice structure defining when the opponent should attack or not could be implemented.

Future work that can be done to extend analysis includes utilizing metrics in PDDL, e.g. to minimize the damage taken by the player. This would reflect the different types of play-styles and approaches to MTG, highlighting the various ways one may strategize their victory. Another interesting gameplay component to implement in a planning setting is the ability for players to draw from decks of cards. This would have a multitude of effects on the factors that play into the strategies and paths explored by the planner.

4 Demonstration

During the live demonstration, we will show the full aspects of the PDDL model, along with actual cards that were modelled, explicitly displaying four major problem files, including the state depicted in Figure 1. Live solving and demonstrated solutions will be done using the VSCode PDDL plugin with ENHSP-2020 planner (Scala et al. 2016). The application of the generated plans will be showcased through a combination of the Cockatrice application² and physical MTG cards that will be used as a visual example to show the corresponding actions in MTG.

References

- Haslum, P.; Lipovetzky, N.; Magazzeni, D.; and Muise, C. 2019. *An Introduction to the Planning Domain Definition Language*. Morgan & Claypool. ISBN 9781627058759.
- Scala, E.; Haslum, P.; Thiébaux, S.; and Ramírez, M. 2016. Interval-Based Relaxation for General Numeric Planning. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, 655–663. IOS Press.
- Stahl, D. O. 1993. Evolution of smartn players. *Games and Economic Behavior*, 5(4): 604–617.

²<https://cockatrice.github.io/>