Magic: The Gathering as a Benchmark for Strategic Intelligence in Artificial Systems: The Project cEDH Framework

Richard Lam
Independent Researcher
January 2026

***Abstract*—Games have historically served as milestones for progress in artificial intelligence. Chess demonstrated search efficiency, Go highlighted deep reinforcement learning, and Poker advanced reasoning under hidden information. These environments remain constrained by fixed rules and limited multi-agent dynamics.**

This paper proposes competitive Magic: The Gathering, specifically the multiplayer Commander format known as competitive Elder Dragon Highlander (cEDH), as a candidate benchmark for evaluating broader forms of strategic intelligence. The domain combines imperfect information, stochastic dynamics, symbolic rule interactions, evolving strategic metas, and multi-agent competition within a single environment.

We introduce Project cEDH, a deterministic simulation framework supporting experimentation with AI agents in this domain. Structural comparison with existing benchmarks and a mathematical framing of decision complexity suggest that consistent performance in cEDH would require integrating probabilistic reasoning, symbolic manipulation, opponent modeling, and long-horizon planning. These properties indicate the domain may serve as a meaningful testbed for evaluating general decision-making systems.

***Index Terms*—Artificial Intelligence, Game Complexity, Multi-Agent Systems, Reinforcement Learning**

I. Introduction

Games provide controlled environments for evaluating artificial intelligence systems. Success in chess and Go demonstrated advances in search and learning [1], [2], while Poker introduced reasoning under uncertainty [3], [4]. Most benchmark games involve fixed rules, two-player competition, and relatively stable strategy spaces.

Real-world decision problems often involve multiple agents, uncertainty, symbolic rule interactions, and changing environments. Competitive Magic: The Gathering, particularly the multiplayer Commander format, exhibits these properties simultaneously. Unlike traditional benchmarks, the domain requires agents to reason about both structured rules and emergent strategic behavior. We argue that this combination makes it a useful candidate benchmark for evaluating more general forms of strategic reasoning and adaptive decision-making.

II. Structural Properties of the Domain

Magic combines several dimensions of complexity rarely present together in existing AI benchmarks.

Imperfect information: hidden hands, unknown draw order, and incomplete opponent knowledge.

Stochastic progression: randomized draws require planning over distributions of outcomes.

Symbolic rule interactions: thousands of card effects interact through layered timing and conditional logic.

Multi-agent competition: games typically involve three to four players, introducing coalition dynamics and non-zero-sum incentives.

Evolving strategy space: new cards and archetypes continually reshape optimal play.

These factors produce a decision environment closer to complex stochastic multi-agent systems than deterministic board games. In particular, the interaction between symbolic rules and probabilistic state transitions creates planning problems that are neither purely combinatorial nor purely statistical, but require integration of both reasoning modes.

III. Mathematical Framing of Complexity

Let S denote the state space and A the action set. Multiplayer Magic is naturally modeled as a partially observable stochastic game.

$$|S| \approx \sum_{d \in D} \prod_{i=1}^{n} |Z_{i,d}| \tag{1}$$

where D represents deck configurations and $Z_{i,d}$ represents zone permutations for player i. Because both terms are large, the effective state space expands combinatorially.

$$b = f(|H|, |B|, |M|, |R|) \tag{2}$$

where H denotes hand size, B battlefield state, M available mana, and R reactive options. This produces irregular search trees that differ from classical deterministic games.

The expected planning horizon also varies dynamically with board interaction density. In highly interactive states, decision consequences may propagate through multiple triggered abilities and responses, effectively deepening the search depth required for optimal play.

TABLE I
REPRESENTATIVE AI BENCHMARK PROPERTIES

Game	Information	Agents	Rule Space
Chess	Perfect	2	Fixed
Go	Perfect	2	Fixed
Poker	Imperfect	2–6	Fixed
cEDH	Imperfect	3–4	Expanding

IV. BENCHMARK CONTEXT

Traditional benchmarks tend to isolate single dimensions of complexity. Magic integrates multiple dimensions simultaneously, suggesting that performance in this environment may better reflect broader strategic capability.

V. PROJECT cEDH FRAMEWORK

Project cEDH implements a deterministic simulation engine supporting reproducible experimentation with AI agents. Card effects are represented in structured form, allowing consistent rule execution and detailed logging. The framework enables large-scale self-play, cross-agent evaluation, and systematic analysis of decision behavior under controlled conditions.

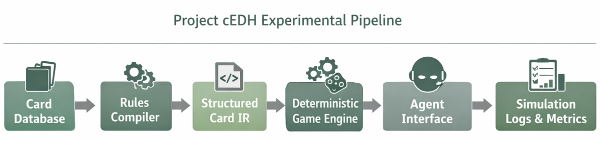

Fig. 1. Project cEDH experimental pipeline. Card definitions are compiled into structured representations, the deterministic engine produces reproducible game states, and AI agents interact through a standardized decision interface. Logged outcomes enable large-scale evaluation and training.

VI. EXPERIMENTAL METHODOLOGY

Evaluation proceeds in three stages. Baseline heuristic agents establish reference performance and verify rule fidelity. Learning agents are trained through large-scale self-play across varied deck distributions, allowing observation of emergent strategies and adaptation patterns. Final evaluation compares agent populations against synthetic opponents and experienced human players using metrics such as win rate, decision stability, strategic diversity, and robustness across differing game states.

This staged approach allows both algorithmic performance and behavioral characteristics to be measured, enabling comparison between agent architectures as well as analysis of how strategic competence develops over training.

VII. CONCLUSION

An agent capable of consistent success in this domain must integrate probabilistic reasoning, symbolic rule interpretation, opponent modeling, and adaptive planning. These requirements extend beyond those of classical benchmarks and resemble capabilities expected in more general decision systems.

Competitive Magic therefore offers a promising candidate environment for evaluating broader forms of strategic intelligence. By providing a deterministic experimental platform while preserving structural complexity, Project cEDH enables systematic study of decision-making in environments that more closely resemble real-world strategic interaction.

ACKNOWLEDGMENT

The author acknowledges the use of ChatGPT to assist with formatting and drafting initial content.

REFERENCES

[1] M. Campbell, A. Hoane, and F. Hsu, "Deep Blue," *Artificial Intelligence*, 2002.
[2] D. Silver et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, 2016.
[3] N. Brown and T. Sandholm, "Superhuman AI for multiplayer poker," *Science*, 2019.
[4] M. Moravčík et al., "DeepStack: Expert-level artificial intelligence in heads-up no-limit poker," *Science*, 2017.
[5] MIT Technology Review, "Magic: The Gathering and computational complexity," 2020.