

Multi-Agent Board Game Strategy Through Simulation

Andrés Espinosa
University of Florida
andres.espinosa.academic@gmail.com

Cody Flynn
University of Florida
codyflynn358@outlook.com

Jason Li
University of Florida
jason.li1@ufl.edu

Han Mach
University of Florida
hmach@ufl.edu

Maximillian Banach
University of Florida
mbanach@ufl.edu

Cathy Quan
University of Florida
cathyquan@ufl.edu

Brian Magnuson
University of Florida
b.magnuson@ufl.edu

ABSTRACT

Settlers of Catan (*Catan*) is a popular board game featuring expansion, resource management, negotiation, and long-term planning. Unlike traditional turn-based strategy games like *Go* or *Chess*, which have been effectively mastered by AI agents, *Catan* presents a far greater challenge due to its stochastic nature, non-stationary dynamics, and multi-agent interactions—with up to 4 players in the base game. *Catan* also features slow-paced gameplay with frequent low-impact actions, such as unsuccessful trade proposals or inconsequential dice rolls. These elements lead to long decision sequences with limited immediate feedback, which have made it difficult for both AI-based and game-theoretic models to outperform the average human player.

To address this challenge, we leverage two emerging areas of research: Turn-Based Multi-Agent Reinforcement Learning (TMARL) and Graph Neural Networks (GNNs). TMARL provides a learning framework where agents iteratively improve their decision-making through experience in multi-agent environments. Meanwhile, GNNs offer a powerful representation learning approach that captures both local and global structures within a game state.

In this paper, we investigate key factors that influence the effectiveness of TMARL agents in the environment of *Catan*. Specifically, we examine the impact of:

- (1) a GNN-based state-to-action network designed to leverage the board's spatial and relational structure,
- (2) various exploration strategies to navigate extended decision-making phases,
- (3) different replay sampling distributions to prioritize learning from high-impact decisions, and
- (4) an ensemble model which mimics strategies employed by human players.

The effectiveness of these approaches is assessed by comparing the best-performing model against established baseline agents, including Random, Value Function, Game-Theoretic, and Human players.

Our findings provide valuable insights into the challenges and opportunities of applying RL to complex, multi-agent environments. The simultaneous interplay of competing objectives and diplomatic negotiations presents a unique challenge for AI—one that remains largely unsolved. Therefore, our results also have broader implications for the development of decision-support systems in military strategy, where turn-based simulations are commonly used to model future engagements, resource allocation, and coalition dynamics—similar to *Catan*.

ABOUT THE AUTHORS

Andrés Espinosa (AI Team Lead) - Industrial Engineering *summa cum laude* graduate, of the University of Florida with experience in supply chain management, reinforcement learning, and transportation optimization research.

Cody Flynn (AI Team Member) graduated from UF with a Computer Science major and Mathematics minor. He specializes in algorithm design, backend development, and testing pipelines. Cody is working in industry at Greenway Health and will pursue a MS in 2026.

Han Mach (Programming Team Member) is Computer Science alum at the University of Florida, with a minor in Digital Arts and Sciences.

Jason Li (Programming Team Lead) is a Computer Science and Statistics alum of the University of Florida with experience in full-stack website development and app development. He is currently pursuing his PhD at the University of Utah.

Maximillian Banach (Project Manager) is an Eagle Scout & Electrical Engineering & Economics graduate of the University of Florida, currently pursuing his M.S. at UF.

Brian Magnuson (Programming Team Member) is a Computer Science graduate of the University of Florida with a minor in Digital Arts and Sciences.

Cathy Quan (AI Team Member) is a 2025 Computer Science graduate of the University of Florida. She minored in Mathematics and has obtained an AI Certificate.

Multi-Agent Board Game Strategy Through Simulation

Andrés Espinosa
University of Florida
andres.espinosa.academic@gmail.com

Maximilian Banach
University of Florida
mbanach@ufl.edu

Cody Flynn
University of Florida
codyflynn358@outlook.com

Han Mach
University of Florida
hmach@ufl.edu

Cathy Quan
University of Florida
cathyquan@ufl.edu

Jason Li
University of Florida
jason.li1@ufl.edu

Brian Magnuson
University of Florida
b.magnuson@ufl.edu

INTRODUCTION

A wargame, as defined by Dr. Peter Perla, is a model involving people making decisions in a synthetic environment of competition or conflict, in which they see the effects of their decisions on that environment and then get to react to those changes (Perla, 2022). This framework extends naturally to competitive environments like 4X games and strategic board games, which share similar dynamics. The application of artificial intelligence (AI) to complex, multi-agent environments has seen significant advancements in recent years. However, military wargaming remains a particularly challenging domain for AI.

The complexity of wargames stems from a variety of factors: non-stationary dynamics, the need for high-level long-term planning that coordinates resources and actions across multiple domains, pervasive uncertainty, diplomatic negotiation, strict constraints, and competing objectives. To illustrate these challenges, consider the reinforcement learning (RL) breakthrough of AlphaZero in mastering Chess (Silver et al., 2017a) AlphaZero must learn in a deterministic environment—mastering board representation, understanding game rules, evaluating positions, and grasping opening move sequences which have been studied by experts for centuries. Now, imagine scaling this up to a wargame like Catan, where an RL agent must navigate not only the same complexities but also a constantly changing board, the need to negotiate with multiple players, and the unpredictable dynamics of player alliances and conflicts. The stakes are higher, the environment is more chaotic, and the agent's ability to adapt becomes far more critical. From this, it becomes clear why wargames have remained a formidable challenge for AI. However, solving these challenges holds immense potential for military strategy, diplomacy, and other high-stakes, multi-agent domains.

Beyond its contributions to the RL research community, this paper highlights three key applications of enhanced wargaming RL agents: enriched educational and training simulations, automated adversarial and red-team strategy development, and more robust decision support systems for high-stake strategic environments. First, improved RL agents can elevate educational wargaming by serving as adaptive opponents or allies, enabling trainees to experience a wider variety of realistic scenarios. Second, as RL continues to develop, these agents can expose blind spots and stress-test assumptions in military planning, offering a scalable alternative to traditional red-teaming methods. Finally, high-stakes decision-making environments such as cyber, logistical, or battlefield operations can benefit by integrating RL agents in wargaming simulations. These systems can quickly simulate long-term consequences of strategic choices, adapt to changing conditions in real-time, and provide commanders with actionable insights under uncertainty. To explore these possibilities in a controlled yet sufficiently complex environment, we turn our attention to the strategic board game *Catan*.

Catan offers a uniquely rich environment for studying RL in wargame-like settings: it incorporates many of the core challenges found in military wargames, including strict constraints, stochasticity, and multi-agent negotiation. Unlike purely deterministic games such as chess or Go, *Catan* forces agents to reason under uncertainty, engage in trade and diplomacy with opponents, and make strategic decisions over a longer horizon. Additionally, its widespread popularity and standardized ruleset make it a reproducible and well-understood testbed for AI research.

These characteristics, combined with the availability of the open-source simulator *Catanatron* (Collazo & contributors, 2024), make *Catan* an ideal sandbox for testing RL methods designed for complex, adversarial simulations.

We now present the structure of our investigation and results in applying RL methods to *Catan*. First, in the section titled "Related Work", we discuss related work in the areas of RL and wargames, highlighting both foundational studies and recent advancements. In the section titled "Background", we provide a comprehensive overview of the necessary concepts for the paper. In the section titled "Implementation", we detail the problem formulation and the methodology used to develop and test the RL Agents in the *Catanatron* environment. In the section titled "Results", we present the results of our experiments, comparing the performance of different RL models and their implications to the broader realm of RL in wargames. Finally, in our Conclusion section, we discuss the broader implications of our findings, including future real-world applications in wargame simulations, and conclude with suggestions for future research.

RELATED WORK

Research on artificial agents in strategic environments has evolved significantly over the past decades. Early work focused on rule-based systems and heuristic searches, with notable examples including Minimax and Alpha-Beta pruning used in traditional board games like Chess and Checkers. The advent of deep learning led to a major shift in agent design.

The introduction of Deep Q-Networks (DQN) (Mnih et al., 2013) marked a breakthrough by enabling agents to learn value functions directly from high-dimensional input. This success culminated in AlphaGo (Silver et al., 2017a) and AlphaZero (Silver et al., 2017b), which demonstrated superhuman performance in Go, Chess, and Shogi through self-play and Monte Carlo Tree Search (MCTS). However, these successes were largely confined to deterministic, fully observable environments. In contrast, environments like the board game *Catan* (Teuber, 1995) and the 4X game *Civilization* (Meier & Shelley, 1991) present more complex, stochastic, and partially observable settings. The characteristics of imperfect information, multiple win conditions, and negotiation align more closely with military wargames and diplomatic simulations.

Recent work has begun exploring Reinforcement Learning in multi-agent and highly constrained settings. In (Amato and Shani 2010), the Dyna-Q algorithm adapts to different personality types in *Civilization IV*, inspiring our ensemble approach. In the realm of Pokémon battles, a Q-learning agent underperforms against a custom Minimax algorithm (Kalose et al., 2018). While in *Tribes* (Perez-Liebana et al., 2020), a genetic algorithm yields the best results—though a rule-based agent performs comparably.

Efforts to apply RL to *Catan* reveal numerous challenges. Dr. Henry Charlesworth (Charlesworth, 2022) shows that introducing negotiation dramatically expands the action space, severely impacting agent performance. To mitigate similar issues, Kim and Li (Kim & Li, 2021) disable trading altogether to avoid model degradation. In (McAughan et al., 2018), the problem is simplified drastically by restricting a DQN agent to only the opening phase of the game. Most notably, *Catanatron* employs an Alpha-Beta model which relies on a hand-crafted heuristic value function, which remains the best-performing agent in *Catan* to date for the purposes of our research.

In parallel, the rise of Graph Neural Networks (GNNs) has provided a powerful mechanism for learning over structured data, which is particularly valuable in games like *Catan*, where spatial and relational features are central to strategic reasoning. Relational inductive biases are critical when tasks naturally map to graph representations, such as board games or networked environments (Battaglia et al., 2018). Recent surveys and reviews (Munikoti, Agarwal, et al., 2022) further highlight the potential of GNN-RL integration across wargame-like domains including game simulations, pandemic modeling, and cybersecurity. Notably, graph-based reinforcement learning has been successfully applied to previously unsolved domains, such as influence maximization (Munikoti, Natarajan, & Halappanavar, 2022) traffic signal optimization (Devailly et al., 2022), and cyberattack simulation (Berglund, 2022). These studies showcase GNN's ability to scale and generalize across graph-structured domains.

Despite the breadth of work in strategic agents, few studies have tackled environments that simultaneously exhibit strict constraints, stochasticity, and social complexity. *Catan* embodies these traits, yet existing reinforcement learning approaches often simplify the environment or rely on domain-specific heuristics. At the same time, graph-based RL methods have shown promise in similarly structured domains but remain largely unexplored in social-diplomatic

games. Our research aims to document which components of reinforcement learning align well with wargame-like mechanics—from graph representations of spatial dynamics to policy learning under uncertainty and negotiation.

BACKGROUND

Settlers of Catan

Catan is a popular multiplayer board game centered around resource management, player negotiation, and settlement expansion. It is played on a hexagonal board composed of various resource tiles. Players roll dice for resource acquisition and may trade with others to build roads, settlements, and cities strategically in efforts to be the first person to reach 10 victory points and win the game against other opponents.

Unlike deterministic games like chess or Go, Catan incorporates many stochastic elements. The game board is randomized in every session; resource tiles and their associated number tokens vary for each game, meaning players must adapt to a different environment with every playthrough. Resource availability is determined by dice rolls, which introduce probabilistic variability even though some numbers are more likely than others. Additionally, the game includes elements of diplomacy, as players can propose, accept, or reject trades. These interactions often reflect hidden goals, temporary alliances, or attempts to sabotage opponents. Such dynamics create a rich space for analyzing multi-agent behavior and contribute to the game's expansive action space.

This large action space, combined with the presence of 4 players, results in slow-paced gameplay with frequent low-impact actions, such as unsuccessful trade proposals or inconsequential dice rolls. To succeed, players must think about the actions of all 3 opponents and rationalize the hidden decisions and goals. These elements lead to long decision sequences with limited immediate feedback, which have made it difficult for both AI-based and game-theoretic models to outperform the average human player.

Deep Reinforcement Learning

Deep Reinforcement Learning (DRL) is a subfield of machine learning that integrates reinforcement learning (RL) principles with deep neural networks. In RL, an agent interacts with an environment and learns to select actions, through trial and error, that maximize cumulative future rewards. The agent aims to learn optimal actions in given states to achieve the highest possible long-term rewards over time. DRL extends traditional RL by using deep networks to approximate either value functions such as Deep Q-Learning (DQN), which use neural networks to approximate the Q-value function, or by employing policy gradient methods directly like Proximal Policy Optimization (PPO) and Trust Region Policy Optimization (TRPO), which directly optimize the policy. This integration allows the agent to handle high-dimensional state and action spaces, making DRL suitable to solve more complex problems.

One of the most well-known DRL breakthroughs is AlphaZero. This breakthrough demonstrated that deep networks accompanied by self-play could achieve superior performance against humans in games like chess and Go. However, DRL success has yet to hold in multi-agent environments like Catan, where the presence of other multi-agent interactions and stochasticity introduces instability in training. These challenges also present an opportunity to explore and advance methods to make DRL more effective in environments where strategic reasoning and uncertainty exist.

Graph Neural Networks

Graph Neural Networks (GNNs) are a type of neural network architecture that is designed to process graph-structured data. In a graph, nodes represent entities, while edges represent the relationships between them. Layers in a GNN iteratively update node representations by aggregating features from neighboring nodes, allowing the model to capture local and global structural information within the graph.

In the context of games like Catan, GNN's can be useful to model the game board as a graph. In this graph formulation, elements of the game are the nodes and spatial interactions between them are edges, which allows an agent to develop an understanding about the strategic dependencies more effectively than flat vector input structures. Feedforward neural networks (FNNs), are limited in their ability to encapsulate the relationship between entities and often lose spatial information in the process.

While GNNs have the upper hand in complex environments, they also come with limitations. For instance, they can suffer from issues like over-smoothing, where node representations can become indistinguishable after many layers, and scalability issues when applied to large or dynamic graphs. As a result, it is crucial to carefully design meaningful node and edge features to ensure the GNN can learn informative and distinctive representations.

IMPLEMENTATION

To advance the state of RL in wargaming, our implementation focuses on three core research directions. First, we investigate how wargames can be naturally represented as graphs or sets of graphs, and we provide both mathematical and empirical evaluations of the effectiveness of such representations. Second, we examine the impact of RL exploration strategies on the agent's ability to solve wargame-like tasks. Third, we propose an ensemble RL framework that integrates multiple heuristic sub-models, each specializing in different aspects of the game. By offloading low-level constraint handling and decision making, we allow the RL agent to focus on the long-term aspects of the wargame while still outperforming the individual heuristic sub-models.

State Space

The first step in any RL project is to model the environment as a Markov Decision Process (MDP). The game of *Catan* can clearly be modeled as an MDP through translating board and player information into numeric representations. However, its state space is incredibly high-dimensional and complex due to the game's stochastic nature, different win conditions, and multi-agent dynamics. *Catan* is a multi-agent environment with 2 to 4 players. In this work, we standardize the player count to 4 for both vector and graph-based representations. This assumption allows us to train a single, unified model that can generalize across 1v1, 1v2, and 1v3 testing environments. We also assume the game will always use the standard board configuration of 19 hex tiles, 54 building locations, and 72 road locations.

Vector Representation

The vector-based state representation encodes the game state of *Catan* at any given time step as a flattened numerical array. This approach to modeling the game benefits from easy ingestion by a neural network.

We denote the game state as a vector $\mathbf{s} \in \mathbf{R}^{\mathbf{d}}$, where each entry s_i represents a specific element of the game state. The total dimensionality \mathbf{d} corresponds to the total length of all encoded components, organized as follows:

- $s_0 - s_{287}$: One-hot encoding of edge ownership.
- $s_{288} - s_{503}$: One-hot encoding of settlements.
- $s_{504} - s_{719}$: One-hot encoding of cities.
- $s_{720} - s_{1043}$: One-hot encoding of port locations.
- $s_{1044} - s_{1063}$: One-hot encoding of the robber's position.
- $s_{1064} - s_{1177}$: One-hot encoding of tile resource types.
- $s_{1178} - s_{1367}$: One-hot encoding of dice values assigned to tiles.
- $s_{1368} - s_{1442}$: Other game information outside of the board.

See the appendix for more detailed information.

Graph Representation

The same 1443-length state vector described above can be decomposed into two logical parts: board-specific features (0 - 1367) and global game state features (1368 - 1442). This separation naturally lends itself to a graph-based modeling approach: the board is represented as a graph, while the global state becomes a set of auxiliary features. In this representation, spatial and relational dynamics are explicitly captured. We model the board as a graph with:

- **Nodes**: Representing building locations (settlements and cities).
- **Edges**: Representing roads and connectivity between nodes.
- **Node features**: Encoding ownership, port access, and properties of adjacent tiles.
- **Edge features**: Encoding road ownership.
- **Graph-level features**: Capturing the remaining global game state (e.g., player inventories, development cards, bank resources).

While we maintain fixed assumptions on map size and player count for consistency with the vector representation, one advantage of the graph model is its flexibility. By modeling players and tiles as nodes of different types, the graph-based approach can easily scale to more dynamic or irregular game setups.

Each node embedding \mathbf{x} is structured as:

- \mathbf{x}_0 : Node ID. (used for debugging)
- $\mathbf{x}_1 - \mathbf{x}_4$: One-hot encoding of whether player j has a settlement on this node.
- $\mathbf{x}_5 - \mathbf{x}_8$: One-hot encoding of whether player j has a city on this node.
- $\mathbf{x}_9 - \mathbf{x}_{14}$: One-hot encoding of whether this node has a port of resource type l .
- $\mathbf{x}_{15} - \mathbf{x}_{50}$: One-hot encoding of the resource type of each adjacent tile k .
- $\mathbf{x}_{51} - \mathbf{x}_{110}$: One-hot encoding of the dice value of each adjacent tile k .
- $\mathbf{x}_{111} - \mathbf{x}_{116}$: One-hot encoding of whether the robber is on each adjacent tile k .

Each edge embedding \mathbf{e} is structured as:

- \mathbf{e}_0 : Edge ID. (used for debugging)
- $\mathbf{e}_1 - \mathbf{e}_4$: One-hot encoding of whether player j owns the road on this edge.

The graph-level feature vector \mathbf{y} includes:

- $\mathbf{y}_0 - \mathbf{y}_{74}$: The same global features used in the vector representation, including development cards, bank inventory, dice rolls, available player structures, and current victory points.

Comparison and Motivation

By constructing both vector and graph-based state representations, we aim to highlight the trade-offs of each approach. The vector format provides a compact and dense encoding, ideal for traditional feedforward networks.

However, it fails to capture spatial relationships and topological structure—critical elements in games like *Catan*.

The graph representation preserves these relationships and opens the door to leveraging graph neural networks (GNNs), which are more adept at modeling interactions in structured environments. This approach aligns well with how many wargames—including 4X games like *Civilization* and naval simulations like *Harpoon*. In these games, spatial relationships between players, units, and terrain are fundamental and often too complex to be flattened into a fixed-length vector.

Games with rich, dynamic maps and interactions are more naturally expressed as graphs. Units, tiles, players, and resources can be modeled as heterogeneous nodes, and their interactions as typed edges. As wargames become increasingly complex and agentic, we believe that graph-based state modeling will be a foundational tool for future AI systems in this domain.

Exploration Strategy

To explore the action space effectively, we implement two exploration strategies: epsilon-greedy and Gibbs Softmax.

In the epsilon-greedy approach, the agent selects a random action with a fixed probability ϵ , encouraging exploration, while exploiting the current best-known action with probability $1 - \epsilon$. This method is simple to implement and has been widely used in reinforcement learning due to its straightforward balance between exploration and exploitation. On the other hand, Gibbs Softmax assigns probabilities to actions based on their estimated values using a Softmax function. Actions with higher estimated values are selected more frequently, but less optimal actions are still considered with non-zero probability. This probabilistic approach allows for more nuanced exploration compared to the binary decision-making of epsilon-greedy.

The motivation for comparing these two strategies lies in their differing strengths. Epsilon-greedy is computationally efficient and effective in simpler environments, but it may struggle in heavily constrained action spaces like *Catan*, where most actions are unavailable at any given timestep. Gibbs Softmax, by focusing exploration on high-value actions while still considering alternatives, may better navigate such complex, multi-agent environments. By

evaluating both methods, we aim to identify which strategy is more effective in balancing exploration and exploitation in the context of *Catan's* unique challenges.

Action Space

To introduce the action space, we examine the possible decisions a player can make in *Catan*. The action space in *Catan* shares several important characteristics with those found in many wargames.

The first characteristic is that the action space is tightly constrained. For instance, a typical turn in *Catan* consists of approximately 2-10 discrete actions. Each turn begins with a mandatory "Roll Dice" action to simulate resource distribution and ends with an "End Turn" action to relinquish control. As a result, the action distribution is often highly skewed, with a small subset of actions dominating the majority of agent behavior.

A second characteristic is the presence of both competitive and diplomatic actions, which adds significant variety to the action space. Social actions can dramatically expand the space due to their continuous and combinatorial nature. In *Catan*, for example, trading is a key mechanic, and in theory, a player could propose any quantity of one resource in exchange for any quantity of another. To keep the action space manageable in our simulation, we restrict trades to 1:1 resource exchange.

Finally, wargames are typically characterized by a high degree of stochasticity. The initial game state can vary widely between playthroughs, even though the underlying rules remain constant. In *Catan*, this is reflected in the randomized board layout at the start of each game and the dice rolls used to determine resource distribution in each turn.

Taken together, these characteristics illustrate that the action space in wargames is significantly more complex than in many other types of games. Ultimately, this complexity motivates the integration of heuristic decision-making frameworks with the generalization capabilities of reinforcement learning agents. We propose a comparison between two action spaces: the first samples directly from the complete set of feasible actions, while the second selects from a set of heuristic sub-models designed to guide decision-making.

Direct Actions

In the direct-action formulation, the agent samples from the full set of all legal actions available at a given game state. This includes actions such as building roads, settlements, or cities; buying development cards; trading with the bank; initiating trades with other players; playing development cards; and ending the turn. The space of legal actions varies dynamically based on the player's resources, board position, and game phase, making it nontrivial to enumerate each step.

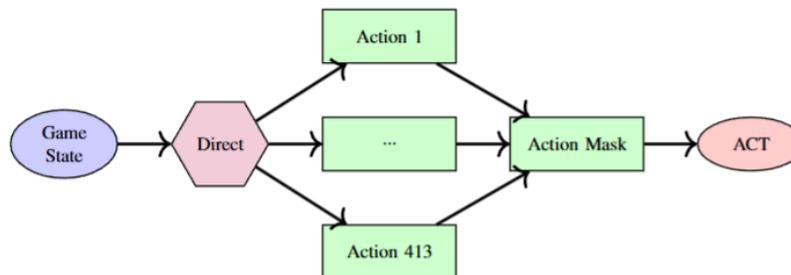


Fig. 1. Overview of the direct model architecture

To ensure tractability, we dynamically generate the legal action set at each timestep and represent it as a discrete action space to mask out infeasible actions. The process of a direct agent selecting a move can be seen above, where the Action Mask is responsible for masking out infeasible actions to ensure the agent does not violate the laws of the game. While this approach gives the agent full control and expressiveness, it suffers from challenges related to sample efficiency and sparse optimal actions—especially during early stages of training where the agent has not yet learned meaningful patterns in action selection.

One thing the current state of RL in wargaming reveals is that the complexities of wargames have made it intractable for RL agents (with the current algorithms and technology) to outperform humans, or even other heuristics. To enable an RL agent to learn how to navigate the long-term strategy of a wargame, while abstracting away individual choices, we propose an ensemble heuristic model in the following section.

Sub-Model Sampling

The heuristic sub-model formulation abstracts the decision-making process by grouping low-level actions into higher-level strategic behaviors. Rather than selecting individual actions directly, the agent chooses from a predefined set of heuristic policies, each designed to encode domain knowledge and target a specific tactical objective.

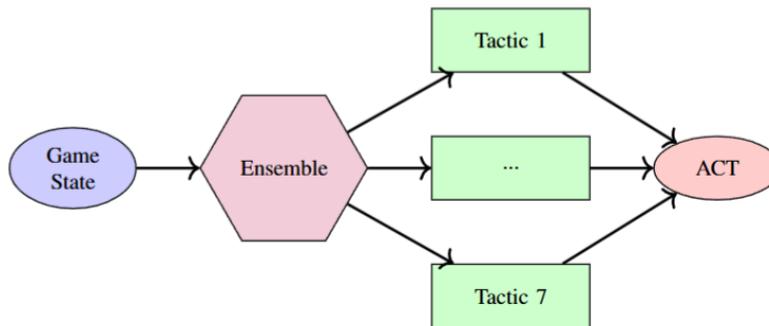


Fig. 2. Overview of the ensemble sub-model decision-making architecture

In our implementation, we employ an ensemble of seven **Value Function** agents—greedy heuristic value function players—each tuned with distinct hyperparameters to specialize in a particular tactic. As illustrated above, the ensemble agent delegates decision-making to one of these sub-models, thereby avoiding the need to explicitly reason about the full space of feasible actions at each timestep. For instance, one sub-model is optimized for road expansion, prioritizing road placement and settlement opportunities, while another emphasizes resource accumulation to enable city development. Once selected, a sub-model autonomously executes a sequence of actions until a predefined stopping condition is met, or control is returned to the agent.

Importantly, the framework is generalizable to any decision-making model that encapsulates narrow, low-correlated tactics, whether learned or rule-based. This hierarchical structure significantly reduces the complexity of the action space, leading to improved sample efficiency and more stable learning dynamics. However, it also introduces a trade-off: the agent's capacity to discover novel or unconventional strategies is constrained by the diversity and expressiveness of the sub-model set.

By comparing this sub-model-based approach with the direct-action formulation, we aim to evaluate the balance between fine-grained control and strategic abstraction in complex, stochastic environments such as Catan.

RESULTS

Several tests were run to compare the performance of different parameters that will be reported in the following subsections. For each table, the win rates are measured over 1000 games against each of the baseline agents in a 4-player environment with one agent of the specified trained model against three agents of a specified baseline model.

Note that for testing whether the following models significantly outperform a given baseline agent, a binomial test was run at the $\alpha = 0.01$ significance level with the alternate hypothesis being that the win rate is greater than 25%. All tests were run in R. When tests were run over 1000 games, the model significantly outperforms a given baseline model if the win rate is $\geq 28.3\%$. When comparing win rates between two models against the same baseline for significant difference, a one-sided Z-test for two proportions was done at the $\alpha = 0.01$ significance level.

Action Space Comparison

We compared two models trained over 1,000 games, differing only in their action space configuration. One model utilized a direct-action selection strategy *Direct4112*, while the other employed a sub-model sampling ensemble approach *Ensemble4111*. The win rates against the three baseline agents are found in Table 1.

Table 1. Comparison of Direct Action vs. Sub-Model Sampling

	Random	Value Function	Alpha-Beta
Direct4112	47.1%	0.6%	0.2%
Ensemble4111	97.6%	26.8%	19.5%

These results indicate that the sub-model sampling strategy used in *Ensemble4111* yields stronger and more generalizable performance across all baseline opponents. In contrast, the direct-action approach in *Direct4112* struggled to compete effectively, particularly against more sophisticated baseline agents.

Representation Comparison

Another model was created to compare the game state representation of agents, particularly the earlier explained graph and vector representation of the environment. These are compared using a GNN model *Ensemble4111* compared to a Feedforward Neural Network (FNN) model *FNN4140* with all other hyperparameters held constant. Performance metrics from this comparison are detailed in Table 2.

Table 2. Comparison of GNN vs. FNN Representation

	Random	Value Function	Alpha-Beta
FNN4140	98.9%	26.2%	16.9%
Ensemble4111	97.6%	26.8%	19.5%

Against **Random** and **Value-Function** models, there is no significant performance difference between the two models. We hypothesize that the lack of performance gain stems from over-smoothing caused by an excessive number of input features in our GNN representation. There is a slight improvement in performance against **Alpha-Beta** agents, although its significance is inconclusive.

Reward Function Comparison

Two models were created to compare the performance of a sparse reward function and a dense reward function. The sparse model consisted of punishing the agent for prolonged games and giving a large reward or punishment corresponding to winning and losing respectively. On the other hand, the dense reward function was constructed to provide rewards encouraging heuristic behaviors. The Performance comparison is detailed in Table 3 with the dense model named *Ensemble4110* the sparse model being *Ensemble4111*.

Table 3. Comparison of Reward Functions

	Random	Value Function	Alpha-Beta
Ensemble4110	99.1%	27.4%	16.1%
Ensemble4111	97.6%	26.8%	19.5%

Against **Random** agents, both models achieved near perfect performance (100% win rate), indicating that both can reliably outperform the simplest models.

Against the **Value Function** baseline, *Ensemble4111* achieved a **26.8%** win rate, compared to 27.4% for *Ensemble4110*, a difference that was not statistically significant. Against the strongest opponent, **Alpha-Beta**, *Ensemble4111*'s win rate rose to **19.5%**, slightly outpacing **16.1%** for *Ensemble4110*, a noticeable improvement.

Overall, these results suggest that sparse reward signals, despite offering less frequent feedback, may more effectively guide the agent toward high-level winning strategies—particularly against stronger, more strategic opponents. Furthermore, sparse models are simpler to construct in most environments. Additionally, sparse reward functions are often simpler to design and implement across a wide range of environments.

Training Opponent Comparison

To evaluate the impact of training diversity on agent performance, we compared two agents trained over 2000 games: *Ensemble4111*, which was trained exclusively against the best performing agents—**Alpha-Beta** agents, and *Catanica*, which was trained against a mixture of **Value Function** and **Alpha-Beta** agents. This experiment aims to assess whether exposure to a broader range of strategies during training leads to better generalization and robustness. Table 2 compares their performance across multiple baseline opponents to determine how training diversity influences overall performance.

Table 4. Comparison of Training Opponents

	Random	Value Function	Alpha-Beta
Ensemble4111	97.6%	26.8%	19.5%
Catanica	98.4%	26.9%	22.7%

Against both **Random** and **Value Function** agents, both models perform nearly the same. However, *Catanica* exhibits an improvement in win rate against **Alpha-Beta**. This suggests that diversifying the training set of opponents can lead to greater generalization and robustness. However, these opponents are also the models that the two reinforcement learning agents were trained against, so this does not yield conclusive evidence, and would require further investigation.

CONCLUSION

In this work, we investigated various approaches to improving the performance of reinforcement learning (RL) agents in the wargame environment of *Catan*. Through a series of experiments comparing different action space configurations, state representations, reward functions, and training strategies, we found that certain strategies, such as sub-model sampling and sparse reward signals, display strong signs of promise. Specifically, the sub-model sampling approach demonstrated superior performance compared to direct action strategies, while sparse reward functions were more effective in guiding agents toward long-term strategic play, particularly against stronger baseline models.

We also explored the impact of training diversity on agent performance, showing that exposure to a broader range of strategies can lead to improved generalization. However, further investigation is required to confirm these findings in more diverse settings.

Applications in Defense and Training

While this study was conducted on a tactical board game, a clear path exists for adapting MARL to similar defense and training simulations. For non-real-time Tactical Decision Games (TDGs), a similar approach can be used to design and train opponent agents to be used for officer training purposes. For instance, distinct agent archetypes, such as offensive, defensive, deceptive, or economy-of-force, could be implemented by either hardcoded rules or learned behaviors, and then combined as modular tactics within a sub-model sampling framework. This would yield a dynamic opponent system that is greater than the sum of its parts, exposing trainees to a richer variety of adversarial strategies.

For real-time simulations, MARL can be extended to define continuous actions (e.g., relocating to the nearest defensible area or executing a complex maneuver) through macros, allowing the agent to perform these behaviors without learning every individual action sequence from scratch.

Constraints

A major constraint of this project was the lack of prior agents capable of playing *Catan*. While games such as Chess and Go have long served as focal points of machine learning research, only a handful of prior works have addressed *Catan*, and even fewer were publicly available. As a result, training the agent required a mix of adapting prior models and developing new ones specifically for this study. By publishing our work and making it available for others, we work to alleviate this constraint.

Additionally, we did not have time to diversify our sub-models, all of which were built using value functions. It is entirely feasible to train specialized RL agents for individual sub-tasks, such as trading, then integrate them into the sub-model sampling structure, which appears promising for future improvements. With more qualitative data, we could also identify situations or move types underrepresented in training and adjust agents to address these gaps.

Future Work

There are several exciting avenues for future research. One limitation in our study was an investigation centered on human players in wargames, which may reveal unique insights into the strengths and weaknesses of current strategies. Additionally, further theoretical exploration into the characteristics of wargames that make them suitable for graph-based representations could lead to more effective model designs.

Another interesting area of research is the potential for subdividing complex wargames into smaller sub-games, each tackled by specialized agents. These agents could use a mix of reinforcement learning or value functions to further reduce the action space and allow the reinforcement learning algorithms to work with the critical decision problems within their sub-games rather than learning the entire game. Additionally, by focusing on a smaller sub-game, we reduce the over-smoothing problem that typically affects large GNNs. Evaluating the performance of such models, particularly in multi-agent environments, could offer valuable insights into how to scale RL agents effectively for more complex games.

Because of the model's nature of choosing a strategy, it is trivial to log which strategy is chosen on a given turn. This functionality could be used to explain the model's actions, which can contribute to the interpretability of the model's choices. It is particularly difficult to create explainable models, which can be critical for high-stakes decision making. In other contexts, ensemble models like Catanica can be applied to robustly explain what choice the model makes at a given time.

CONTRIBUTIONS

This paper makes the following key contributions:

- **Formal MDP formulation expansion** of a popular wargame, Catan, to work in both vector and graph representations, enabling a direct head-to-head comparison of flat vs. structured state encodings in a 4-player, stochastic, partially observable setting.
- **A novel ensemble sub-model sampling architecture (Catanica)**, which decomposes Catan's complex action space into specialized heuristics and achieves far better generalization than direct action policies (e.g. 98.4% vs. 47.1% vs. Random, **26.8%** vs. 0.6% vs. Value Function).
- **An empirical evaluation of different state representations**, showing that an oversmoothed GNN-based model performs similarly to a feed-forward network on the same MDP (e.g. 26.8% vs. 26.2% vs. Value Function, 16.1% vs. 19.5% vs. Alpha-Beta).
- **A study of reward-signal design**, contrasting sparse end-of-game feedback against dense heuristic rewards, and demonstrating that sparse rewards yield significantly higher long-term strategic performance (raising Alpha-Beta win rates from 12.0% to **19.5%**).
- **An analysis of training-opponent diversity**, showing that mixing Value Function and Alpha-Beta opponents during training leads to more robust agents (e.g. **22.7 %** vs. 19.5% vs. Alpha-Beta) than single-opponent curricula.
- **Contribution to the open source Catanatron**, including our entire pipeline, enabling the community to build on our MDP definitions, GNN models, and ensemble framework.

ACKNOWLEDGEMENTS

We gratefully acknowledge Dr. Brian Stensrud for originating the project idea, for his continued correspondence, and for his invaluable support throughout this work. Our thanks also go to Dr. Jörg Peters for his role as IPPD coach and for his thoughtful feedback over the course of the year. We appreciate the University of Florida's IPPD program for uniting our team and fostering the interdisciplinary collaboration that made this project possible. Finally, we thank CAE USA for their generous sponsorship and support.

REFERENCES

- Amato, C., & Shani, G. (2010). High-level reinforcement learning in strategy games, 1, 75–82.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V. F., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gülçehre, Ç., Song, H. F., Ballard, A. J., Gilmer, J., Dahl, G. E., Vaswani, A., Allen, K. R., Nash, C., Langston, V., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M. M., Vinyals, O., Li, Y., & Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks. *CoRR*, *abs/1806.01261*. <http://arxiv.org/abs/1806.01261>
- Berglund, S. (2022). Modelling cyber security of networks as a reinforcement learning problem using graphs: An application of reinforcement learning to the meta attack language.
- Charlesworth, H. (2022). Learning to play Settlers of Catan with deep reinforcement learning. *Settlers RL GitHub Pages*. <https://settlers-rl.github.io/>

- Collazo, B., & contributors. (2024). *Catanatron: Settlers of Catan bot and simulator*. <https://github.com/bcollazo/catanatron>
- Devailly, F.-X., Larocque, D., & Charlin, L. (2022). IG-RL: Inductive graph reinforcement learning for massive-scale traffic signal control. *IEEE Transactions on Intelligent Transportation Systems*, 23(7), 7496–7507. <http://dx.doi.org/10.1109/TITS.2021.3070835>
- Kalose, A., Kaya, K., & Kim, A. (2018). Optimal battle strategy in Pokémon using reinforcement learning. *Stanford University AA228 Course Project Report*. <https://web.stanford.edu/class/aa228/reports/2018/final151.pdf>
- Kim, C. C., & Li, A. Y. (2021). RE-L Catan: Evaluation of deep reinforcement learning for resource management under competitive and uncertain environments. https://cs230.stanford.edu/projects_fall_2021/reports/103176936.pdf
- McAughan, P., Krishnakumar, A., Hahn, J., & Kulkarni, S. (2018). QSettlers: Deep reinforcement learning for Settlers of Catan. <https://akrishna77.github.io/QSettlers/>
- Meier, S., & Shelley, B. (1991). *Sid Meier's Civilization* [Board game]. <https://civilization.2k.com/>
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. A. (2013). Playing Atari with deep reinforcement learning. *CoRR*, abs/1312.5602. <http://arxiv.org/abs/1312.5602>
- Munikoti, S., Agarwal, D., Das, L., Halappanavar, M., & Natarajan, B. (2022). Challenges and opportunities in deep reinforcement learning with graph neural networks: A comprehensive review of algorithms and applications. <https://arxiv.org/abs/2206.07922>
- Munikoti, S., Natarajan, B., & Halappanavar, M. (2022). GRAMER: Graph meta reinforcement learning for multi-objective influence maximization. <https://arxiv.org/abs/2205.14834>
- Perez-Liebana, D., Hsu, Y.-J., Emmanouilidis, S., Khaleque, B., & Gaina, R. (2020). Tribes: A new turn-based strategy game for AI research. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 16(1), 252–258. <https://doi.org/10.1609/aiide.v16i1.7438>
- Perla, P. (2022). Wargaming and the cycle of research and learning. *Scandinavian Journal of Military Studies*.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T. P., Simonyan, K., & Hassabis, D. (2017). Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *CoRR*, abs/1712.01815. <http://arxiv.org/abs/1712.01815>
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., & Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*, 550, 354–. <http://dx.doi.org/10.1038/nature24270>
- Teuber, K. (1995). *The Settlers of Catan* [Board game]. <https://www.catan.com>