Planet Wars RTS AI

IEEE WCCI 2026 Results

Simon Lucas
Director of Game AI Research Group
Queen Mary University of London

Queen Mary
University of London

History

- Galcon (2007-2008)
- Google AI Challenge Planet Wars (2010)
- Planet Wars at Dagstuhl Seminar (2015)
- Ghost in the Cell (2017) on CodingGame
- Fast Planet Wars Variants (Lucas, 2018)
  - Gravity, actuators, parameterised, fast forward models
- Planet-Wars RTS AI (Lucas, 2025)
- Orbit Wars, Kaggle 2026

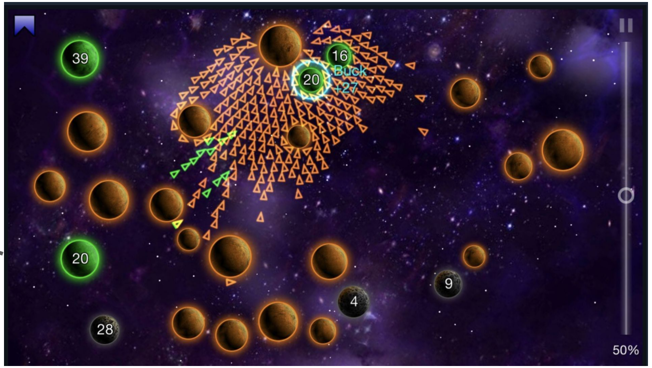

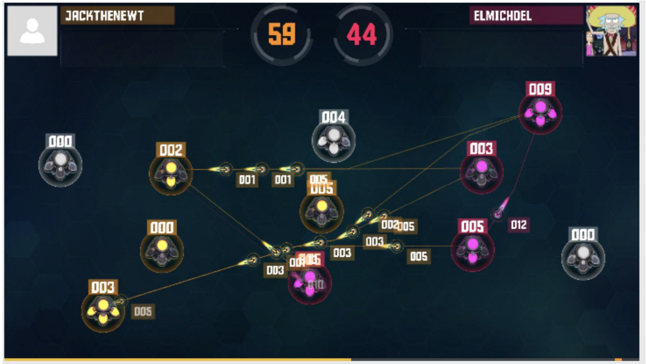

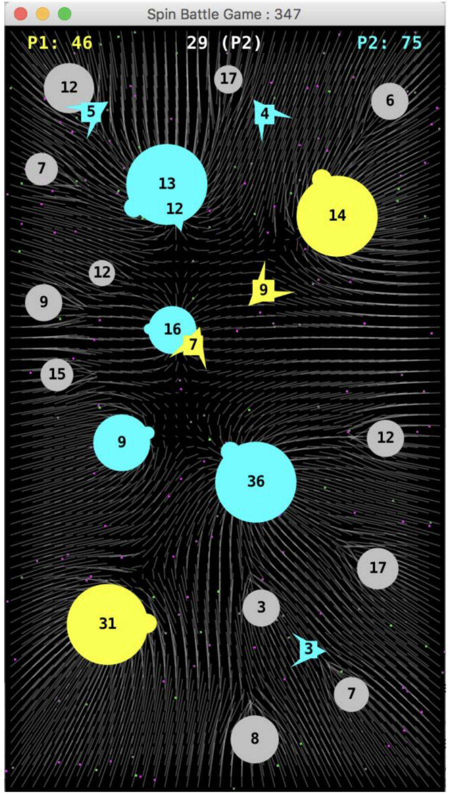

Game Rules

Win by eliminating opponent, or having more **ships on planets** when time runs out

Each planet has a single **transporter**

If you own the planet, and the transporter is at home, you can send ships across space - either to attack opponent or neutral planets, or to defend your own

An invading force **a** will overcome the occupants **b** only if **a > b** then ownership is transferred. Always **|a-b|** ships remaining after battle.

Owned planets (not **neutral** ones) grow ships in proportion to their size

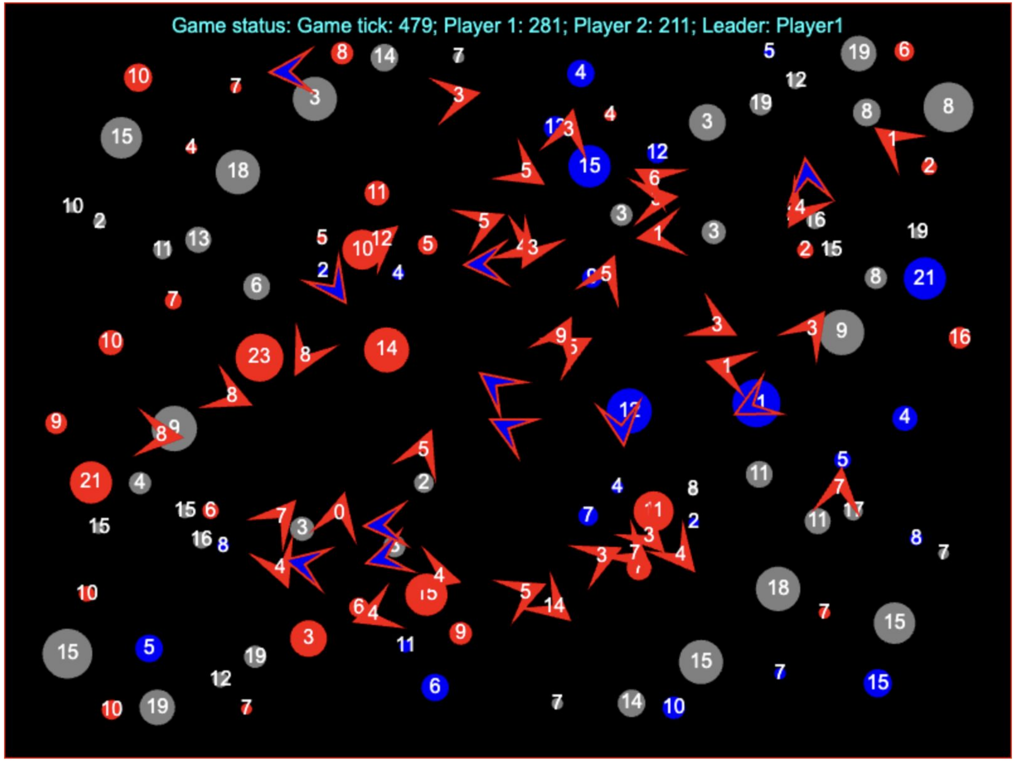

Another Game AI competition?

- Sweet spot for Game AI Challenges:
- RTS - **with real-time enforced**
- Highly configurable
- Hard test for sophisticated AI agents
- Low compute budget
- Challenge for RL agents to cope with high variability in maps and game params
- Open-ended challenge: **depends on opponents**
- **Fast Plannable Forward Model**
    - **1 million steps per second**

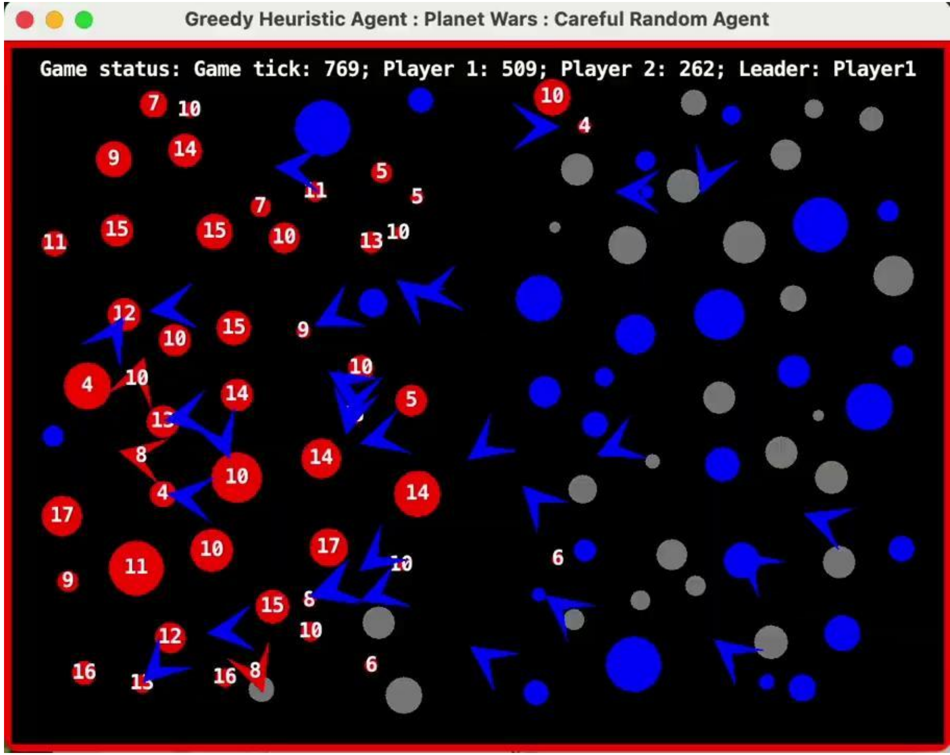

[Video Link]

Some Interesting Research Questions

- Can Claude Code write a good agent?
- Can Evolutionary Program Search Design good agents?
    - LLaMEA
    - LLM4AD
    - EASE
    - DeepEvolve
    - OpenEvolve
    - ShinkaEvaolve
    - … **DIY?**
- RL?
- Graph Neural Networks?
- MCTS, Rolling Horizon Evolution, SFP Variants? (online planning)
- Human-design???

WCCI 2026 Variations

- Number of Planets: 10-30
- Neutral Ratio: 0.25-0.35
- Growth Rate: 0.05-0.2
- Transporter Speed: 2.0-5.0 units/tick

Branching factor:

nSource * nTarget * nShipOptions

E.g. 10 x 20 x 10 = 2000

Often much smaller

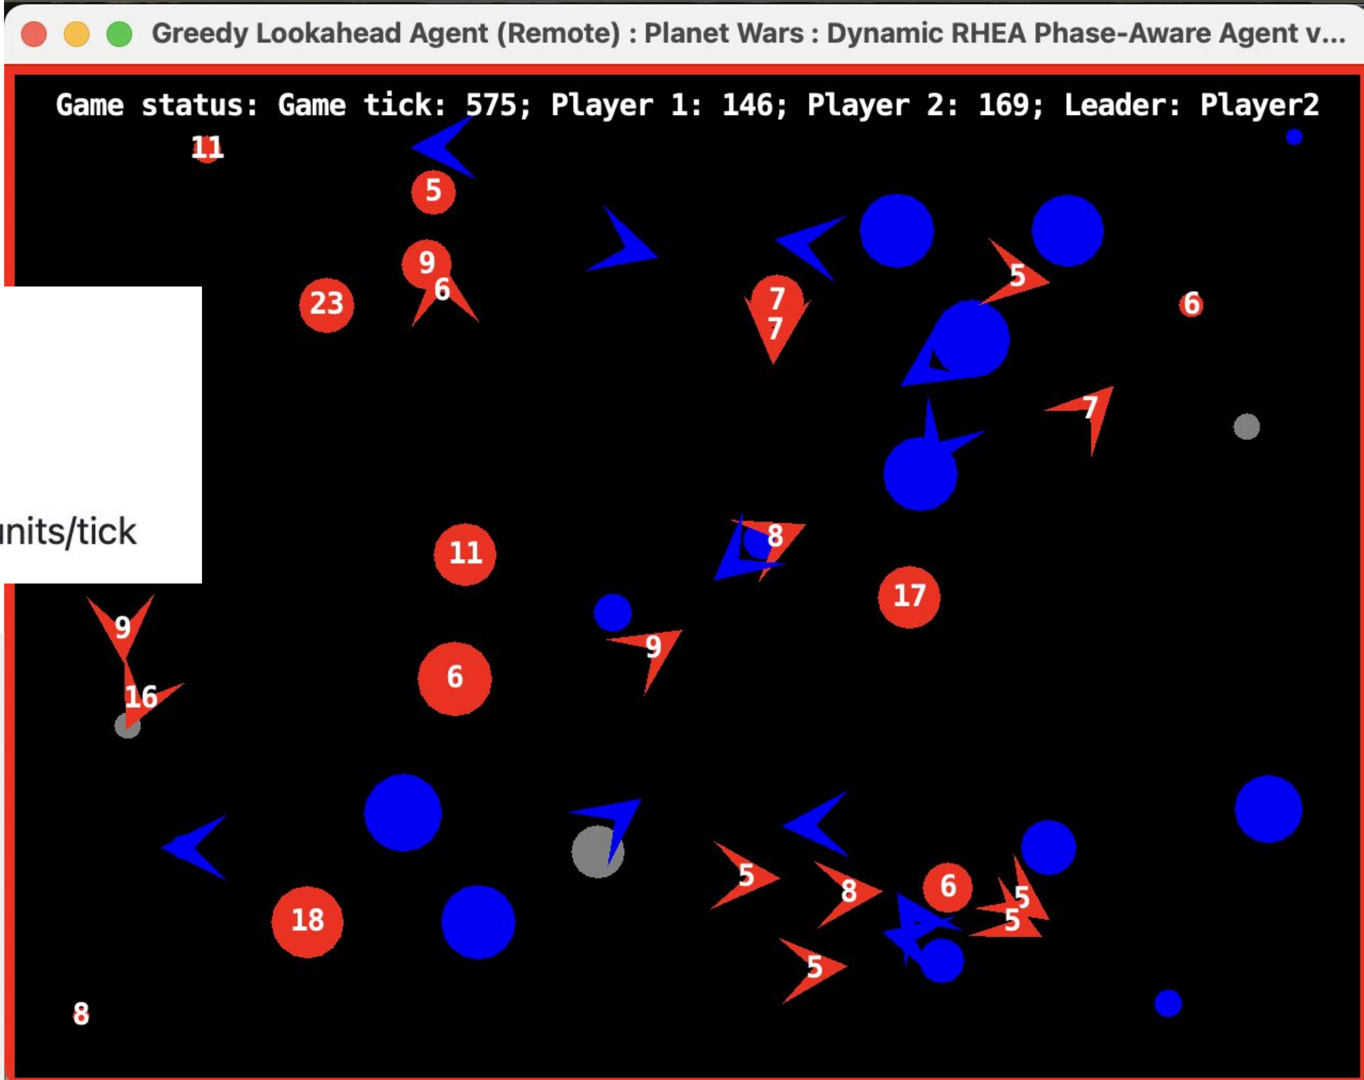

Small number of planets, fast ships

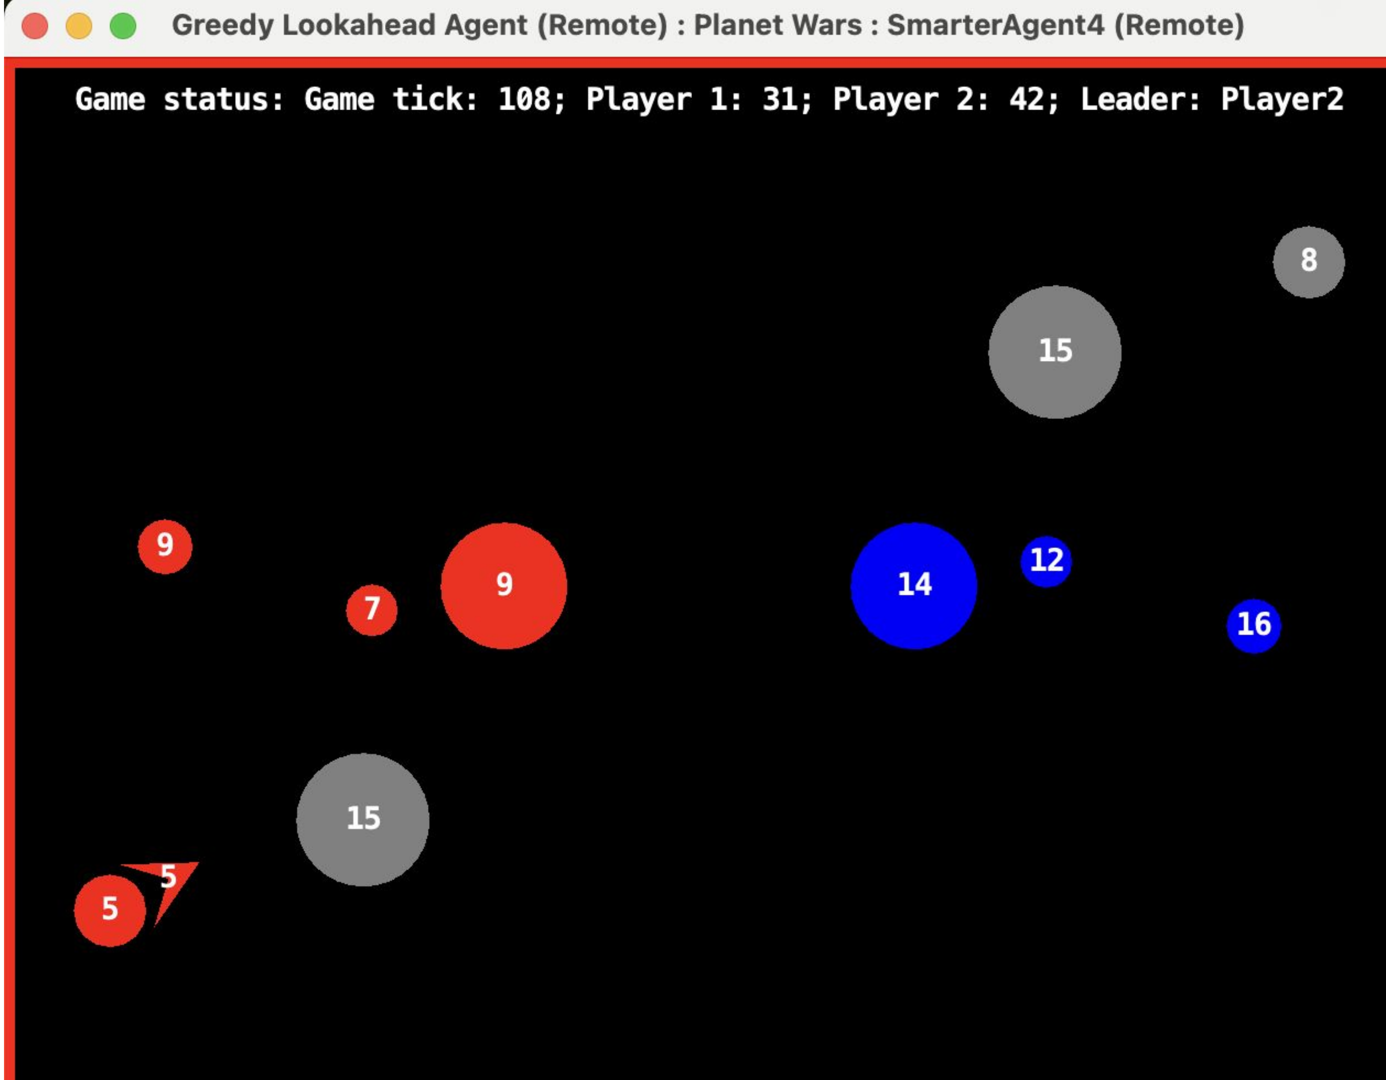

Adversarial Decision Boundary Agent (ADBA) (Vibed with Claude Code)

Full Game Simulation using Forward Model

- Simulates the entire game from current state to end
- Incorporates both player's planned actions
- Models planet ownership changes, ship movements, and growth over time
- Returns final score difference to evaluate position quality

Opponent Response Simulation

- Consider possible opponent moves, and pick our action that works best against the opponent’s best move (Minimax)

Note: Initial version had terrible win-rate: needed (manual) parameter tuning to play well

AAMAS 2026 Ranking (top 12) - TrueSkill(™) League

Planet Wars Spring 2026 — Leaderboard

#	Agent	μ	σ	μ – 3σ	Matches	Updated
1	teamtitansagentv3-256baf2	34.062	0.605	32.248	16040	2026-04-06 08:48
2	gnn_cont999_128_top64-97050fd	32.351	0.517	30.799	13900	2026-04-06 08:48
3	nanaboshi08271-5ea4531	29.686	0.572	27.968	11480	2026-04-06 08:48
4	nanaboshi08281-8df146e	29.145	0.528	27.559	11780	2026-04-06 08:48
5	nanaboshi08280-c2aa68b	28.693	0.515	27.149	11380	2026-04-06 08:48
6	nanaboshi08268-54e82c1	28.629	0.521	27.067	11140	2026-04-06 08:48
7	nanaboshi08270-b40ea88	28.473	0.506	26.954	11340	2026-04-06 08:48
8	gnn_cont999_200_top4-5027e43	27.598	0.502	26.091	11674	2026-04-06 08:48
9	gnn_cont999_200-a31c911	27.038	0.510	25.509	11734	2026-04-06 08:48
10	gnn_cont999_128_base-9531e0d	26.843	0.515	25.299	10160	2026-04-06 08:48
11	adversarial-dba-ba51c9e	23.311	0.501	21.807	8320	2026-04-06 08:48
12	galacticarmada-ecd6ae8	22.116	0.522	20.550	8720	2026-04-06 08:48

AAMAS 2026 Distinct Leading Entries (26th May)

- **TeamTitans**:
  - **Human**-designed heuristics together with **forward planning**
- **GNN:** graph neural networks
- Nanaboshi:
  - DeepRL using MLP
- Adversarial DBA
  - Vibed with **Claude Code, tuned by me** - some smart heuristics and opponent responses
- Galactic Armada
  - **LLaMEA evolved agent**
  - Heuristic

- Planning approach was clear winner: if you have a plannable model, use it!

Anti

PLANET WARS COMPETITION ENTRY

University of Pretoria

Aren Repko, Alec Watson, Daniel Geerdink

Supervisors

Prof. Nelishia Pillay, Dr. Thambo Nyathi

21 June 2026

Our Strategy

Guided Baseline Design

Our baseline was seeded via guided prompting, with us injecting game rules, strategic concepts, and architecture proposals. The resulting agent already outperformed all provided baselines, ensuring the evolutionary search operated in a competitive regime from iteration one, and avoiding wasted compute on weak early generations.

Parameter Optimisation

We performed parameter optimisation on this initial baseline. Once code evolution began we left parameter tuning to the LLM, allowing it to co-evolve both heuristic structure and parameter values together.

Structural Code Evolution

A (1+1) LLM-driven hill climber iteratively mutates the current best agent's source code, with a single structural change proposed per iteration, which is accepted only if it beats the current champion in self-play by a fixed win-rate threshold.

Context Injection

Each mutation was guided by the history of previous iterations.

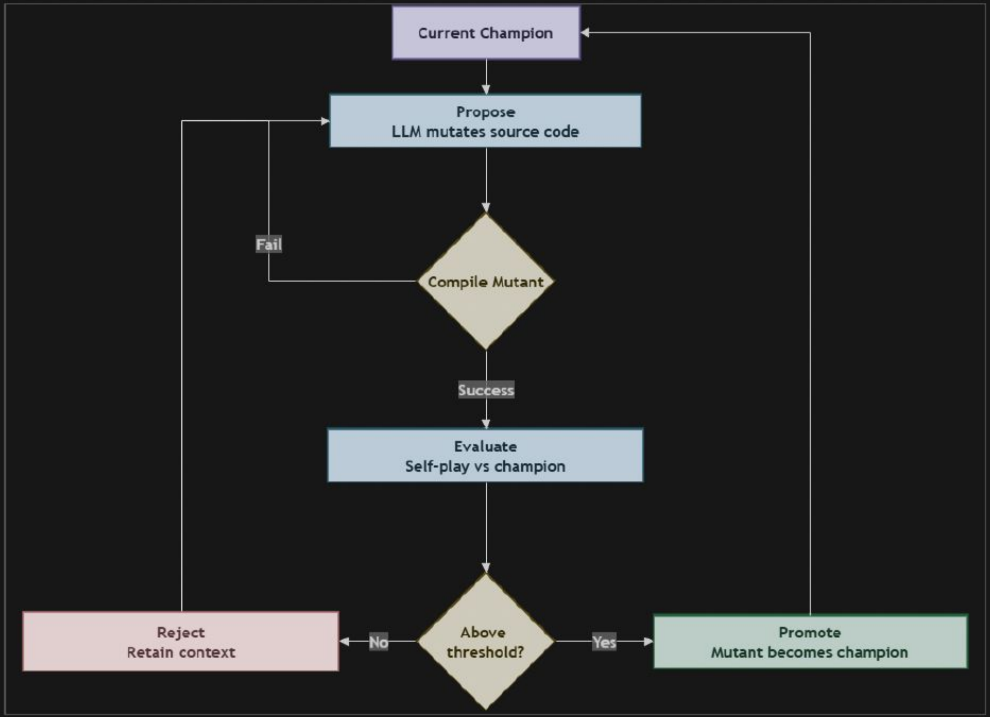

Agent Overview

Lookahead Search

Generates candidate fleet dispatches, each constrained by a per-planet survival calculation, ranks them with a multi-factor heuristic, then validates the top candidates through forward simulation within a fixed 45ms budget.

Adaptive Depth

Simulation horizon is extended dynamically, with the top candidates receiving deeper rollouts if the time budget allows.

Threat Response

Detects incoming enemy fleets and reinforces at-risk planets, abandoning those that cannot be defended.

Expansion

Targets neutral planets weighted by growth rate, proximity, and strategic position.

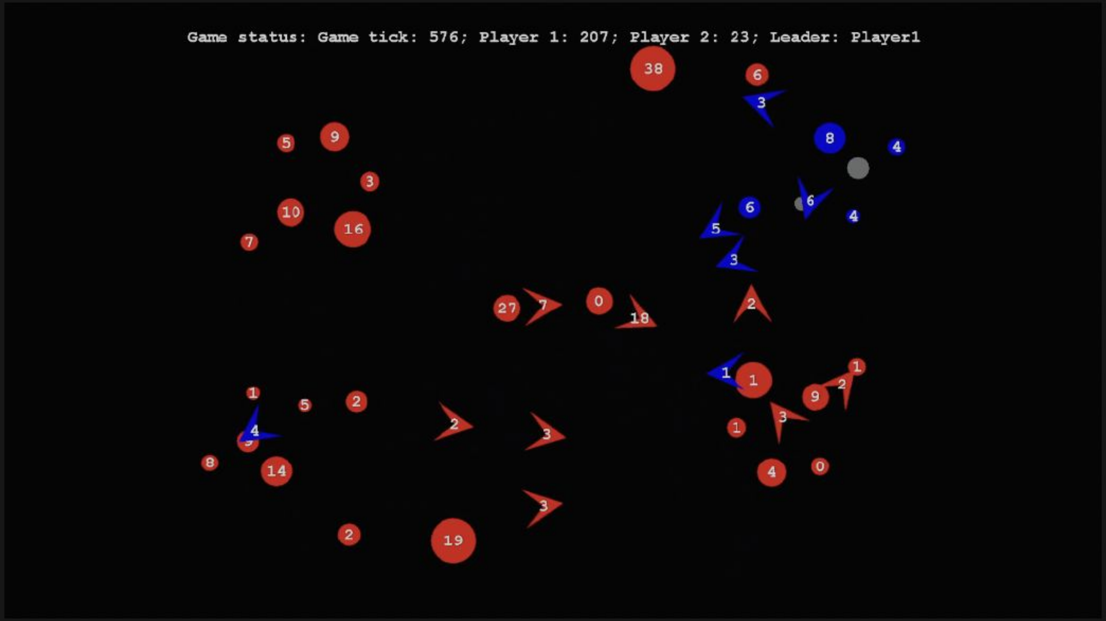

Results

Baseline	Games	Win Rate
DoNothingAgent	1000	99.6%
PureRandomAgent	1000	100%
BetterRandomAgent	1000	100%
CarefulRandomAgent	1000	100%
SimpleEvoAgent	1000	99.6%
GreedyHeuristicAgent	1000	99.9%

Near-Perfect Win Rate

Our agent achieved a 99.85% aggregate win rate across all 6000 games, losing only 9 times total.

Toughest Baselines

The SimpleEvoAgent and DoNothingAgent proved to be the toughest baselines, accounting for 8 of the 9 loses.

Zero Draws

Our agent recorded 0 draws across all baseline matchups.

Analysis and Insights

Training

Training was done through self-play in which we observed 20 promotions over 558 iterations (~3.6% acceptance rate). Additionally as the agent became better the average game length increased dramatically, slowing each evaluation cycle.

Behaviour

The agent's behaviour during games can be described as a careful agent at the beginning of a game, building a lead, sacrificing where necessary or causing maximum damage with doomed troops, and once it has established a lead it aggressively attacks.

Performance and Parameters

Our agent exhibited consistent and strong performance over a wide range of game variations and opponents.

Future Research

Training Time Improvements

Increasing training time with agent quality is a large bottleneck to continued improvement. Eliminating or curbing the increase in training time would result in stronger agents with less resources. Proposed approaches include incentivising aggressive play to shorten games, using statistical game-state breakdowns for early termination, or training a neural network judge to evaluate positions without full rollouts.

Adaptive Artificial Intelligence

We would like to investigate the possibility of both adapting strategy to the opposing agent as the game moves along, as well as modelling their moves, instead of assuming static moves for forward models.

League

#	Agent	μ	σ	μ – 3σ	Matches	Updated
1	anti-1-8c8bf1c	38.210	0.684	36.159	5140	2026-06-23 09:09
2	anti-0-b1b8324	34.435	0.631	32.541	4420	2026-06-23 09:09
3	teamtitansagentv3-256baf2	33.957	0.543	32.330	17460	2026-06-23 09:09
4	gnn_cont999_128_top64-97050fd	32.378	0.543	30.750	16600	2026-06-23 09:09
5	nanaboshi08281-8df146e	29.807	0.530	28.216	13380	2026-06-23 09:09
6	nanaboshi08280-c2aa68b	28.658	0.535	27.054	13140	2026-06-23 09:09
7	nanaboshi08271-5ea4531	28.417	0.540	26.796	13040	2026-06-23 09:09
8	nanaboshi08270-b40ea88	27.785	0.529	26.198	12760	2026-06-23 09:09
9	nanaboshi08268-54e82c1	27.561	0.537	25.950	12600	2026-06-23 09:09
10	gnn_cont999_200-a31c911	27.446	0.523	25.876	13094	2026-06-23 09:09
11	gnn_cont999_200_top4-5027e43	27.374	0.531	25.782	12994	2026-06-23 09:09
12	gnn_cont999_128_base-9531e0d	26.355	0.520	24.796	11460	2026-06-23 09:09
13	adversarial-dba-ba51c9e	25.965	0.519	24.408	9900	2026-06-23 09:09
14	planetwarriors-agent-v2-1-2efb60d	24.968	0.527	23.387	4740	2026-06-23 09:09
15	genetic_v1-04b0974	24.145	0.484	22.693	740	2026-06-23 09:09
16	galacticarmada-ecd6ae8	24.114	0.504	22.603	10100	2026-06-23 09:09

Round Robin: Top 7 distinct from 37 entries

*Cell shows **row agent's win rate** vs column agent (games played in brackets). Overall = mean of per-opponent win rates (each opponent weighted equally).*

Agent (rank↓)	anti-1-8c8bf1c	anti-0-b1b8324	teamtitansagentv3…	gnn_cont999_128_t…	nanaboshi08271-5e…	adversarial-dba-b…	galacticarmada-ec…	Overall
anti-1-8c8bf1c (μ−3σ=36.2)	—	76.0% (200)	80.5% (200)	82.8% (180)	97.2% (180)	98.1% (160)	99.4% (160)	**89.0%**
anti-0-b1b8324 (μ−3σ=32.5)	24.0% (200)	—	57.0% (200)	60.5% (200)	94.4% (180)	94.2% (120)	89.2% (120)	**69.9%**
teamtitansagentv3-256baf2 (μ−3σ=32.3)	19.5% (200)	43.0% (200)	—	63.8% (680)	91.6% (860)	98.1% (360)	98.0% (820)	**69.0%**
gnn_cont999_128_top64-97050fd (μ−3σ=30.7)	17.2% (180)	39.5% (200)	36.2% (680)	—	74.5% (660)	83.6% (640)	96.7% (640)	**58.0%**
nanaboshi08271-5ea4531 (μ−3σ=26.8)	2.8% (180)	5.6% (180)	8.4% (860)	25.5% (660)	—	80.8% (360)	91.7% (520)	**35.8%**
adversarial-dba-ba51c9e (μ−3σ=24.4)	1.9% (160)	5.8% (120)	1.9% (360)	16.4% (640)	19.2% (360)	—	50.6% (360)	**16.0%**
galacticarmada-ecd6ae8 (μ−3σ=22.6)	0.6% (160)	10.8% (120)	2.0% (820)	3.3% (640)	8.3% (520)	49.4% (360)	—	**12.4%**

Summary

- Congrats to the IEEE WCCI 2026 winners! - this is not an easy challenge!
- *But maybe we ain't seen nothing yet?*
- Set it for your student projects!
- Takeaways: **LLM-Evolutionary Search + Self-Play + HPO works**
- Room for human innovation
- Simulation is a super-power
- **Games are Living Benchmarks**
- **Next Edition: IEEE CoG 2026**