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.

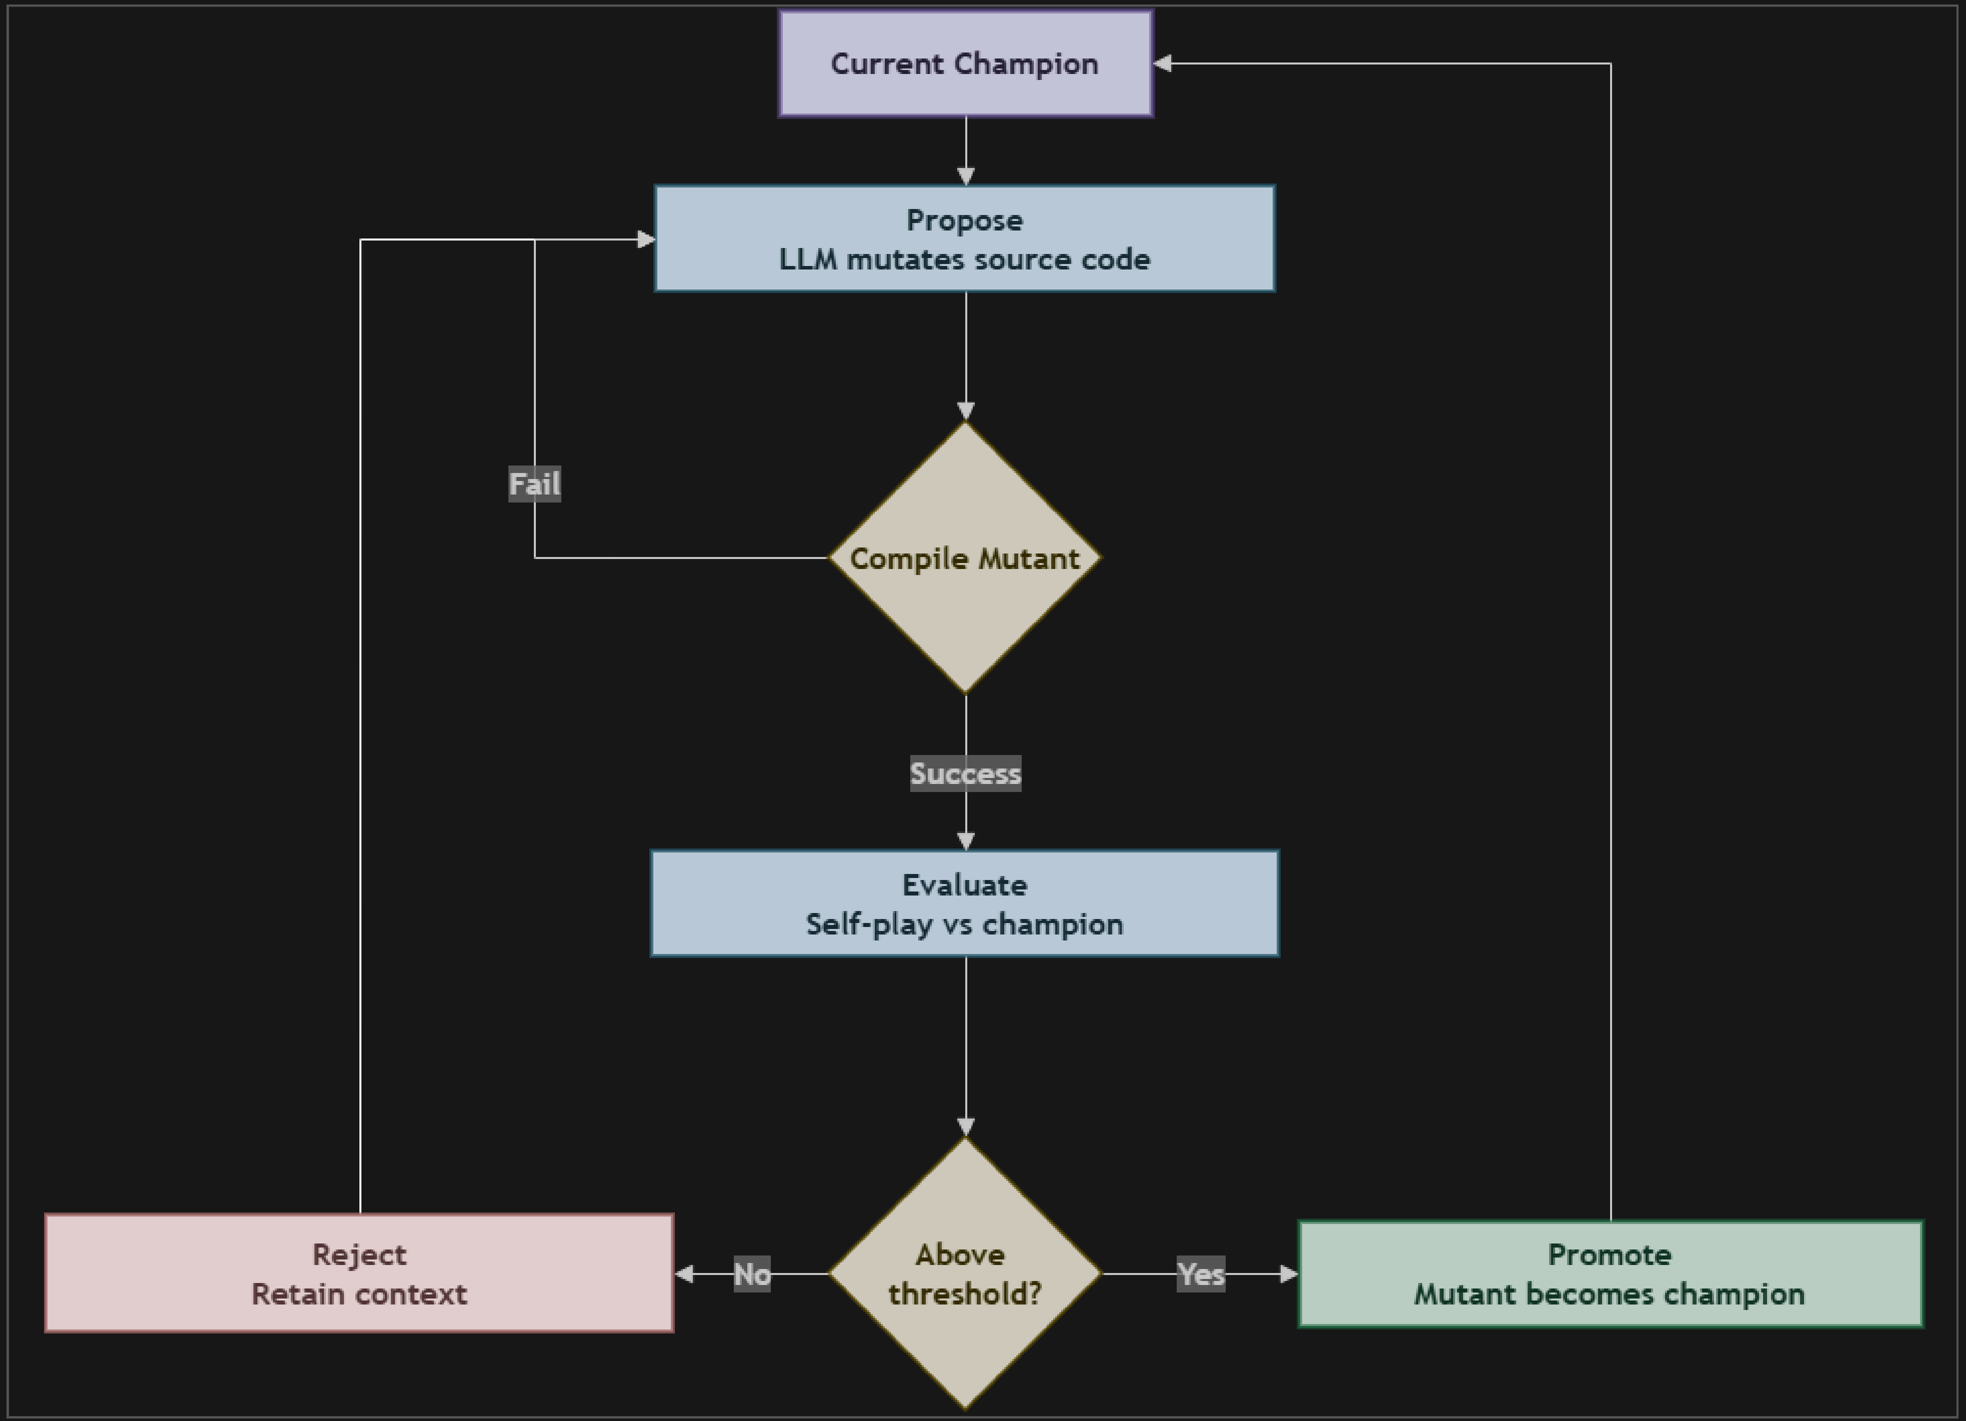

Agent Overview

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.

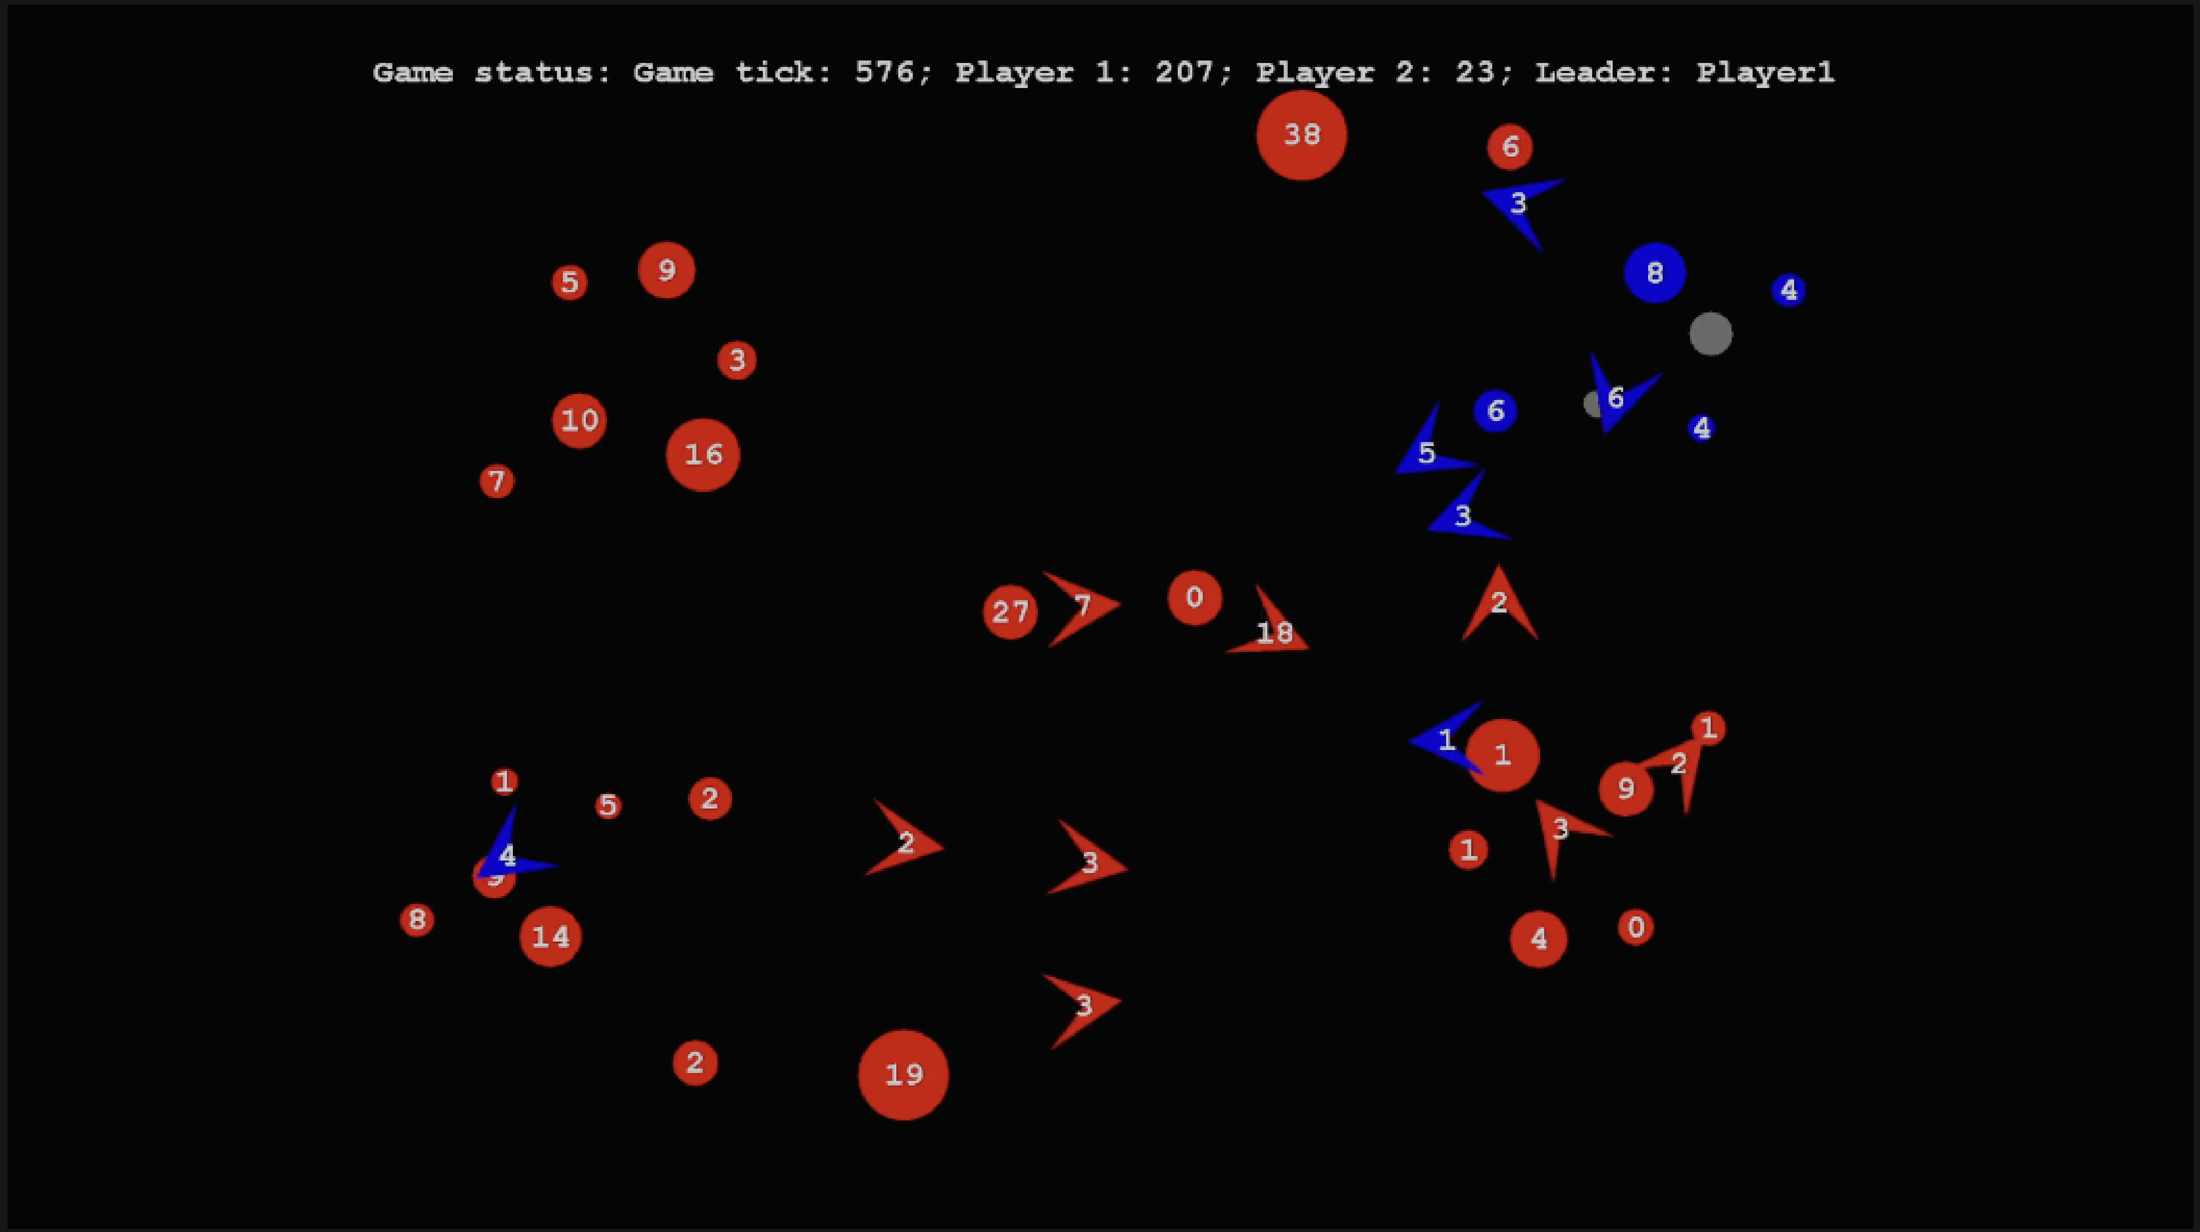

Results

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.