

DareFightingICE LLM Competition 2026

Temporal_Fighter

Guhan Elangovan

College of Information Science and Engineering
Ritsumeikan University

May 15, 2026



Agenda

- Project Goal & Constraints (LLM)
- Related Work
- Proposed Approach
- Experimentation & Results
- Conclusions



1. Project Goal & Competition Constraints

Objective: Develop an LLM-based DareFightingICE agent driven by real-time tactical semantics.

- **Target Leagues:** Standard (Win-rate/HP differential vs MCTS) and Style (Intent-based rushdown/zoning).
- **Strict Constraints:** 16.6ms decision window, Java socket environment, max 7B parameters.



2. Related Work & Research Gaps

- RL agents are mechanically optimal but strategically shallow [Liu et al., 2025].
- ML models often overfit to specific scenarios [Tavares et al., 2025].
- **Critical Gap:** Existing approaches fail to process deep tactical semantics like true intent and momentum [Tavares et al., 2025; Liu et al., 2025].



3. Proposed Approach & Dataset Formatting

Pipeline Features:

- **Architecture:** Llama 3.2 3B Instruct (Unsloth) with 4-bit NF4 quantization.
- **Optimization:** Bare-metal PyTorch inference pipeline.

Dataset Features:

- **State Parsing:** Tactical extraction of Distance, Momentum, and Actions.
- **Formatting:** ChatML tagging (e.g., [MOMENTUM]) to guide classification.

```
{"frame_number":0,"p1":{"hp":400,"x":100,"y":335,"state":"STAND","action":"STAND","graphic_size_x":400,"height":205},"p2":{"hp":400,"x":500,"y":335,"state":"STAND","action":"STAND"}}
```



```
{"messages": [{"role": "system", "content": "You are an elite fighting game AI. Analyze spatial distance, wall proximity, momentum, and state to output the optimal exact action command."}, {"role": "user", "content": "[SPATIAL] Dist: 395, Rel_Y: 0, Wall_Prox: 105. [MOMENTUM] My_Vel: (5, 0), Opp_Vel: (0, 0). [VITALS] My HP: 400, Opp HP: 400. [STATE] My Posture: STAND, Opp Posture: STAND. [OPP ACTION] STAND_GUARD. [PERSONA] Body(W,H): (400, 205)."}, {"role": "assistant", "content": "FORWARD_WALK"}]}
```




4. Experimentation & Current Results

Configuration:

- 50,800 data points (80/20 split), AdamW-8bit, $LR = 2e - 4$.
- **LoRA** ($r = 16, \alpha = 16$): 320 steps convergence ($Loss \approx 0.95$).

Evaluation Plan:

- **Metric:** HP Difference (ΔHP).
- **Test:** Mann-Whitney U Test.



5. Conclusions & Future Work

- **Dataset Poisoning:** Purged $\sim 46\%$ of passive, unmapped states (e.g., CROUCH) to eliminate action hallucinations.
- **Latency:** 4-bit NF4 pipeline achieved efficient GPU inference within constrained VRAM limits.
- **Future Work:** Implement vLLM-based PagedAttention to reach sub-16.6ms real-time targets.

