Feint Behaviors and Strategies: Formalization, Implementation and Evaluation

The 38th Conference on Neural Information Processing Systems (NeurIPS 2024)

Junyu Liu
Brown University

Xiangjun Peng
The Chinese University of Hong Kong

Feint Behaviors

Widely-used tactics in competitive games:

- Boxing and fencing
- Basketball and football
- Motor racing
- Electronic games
- …

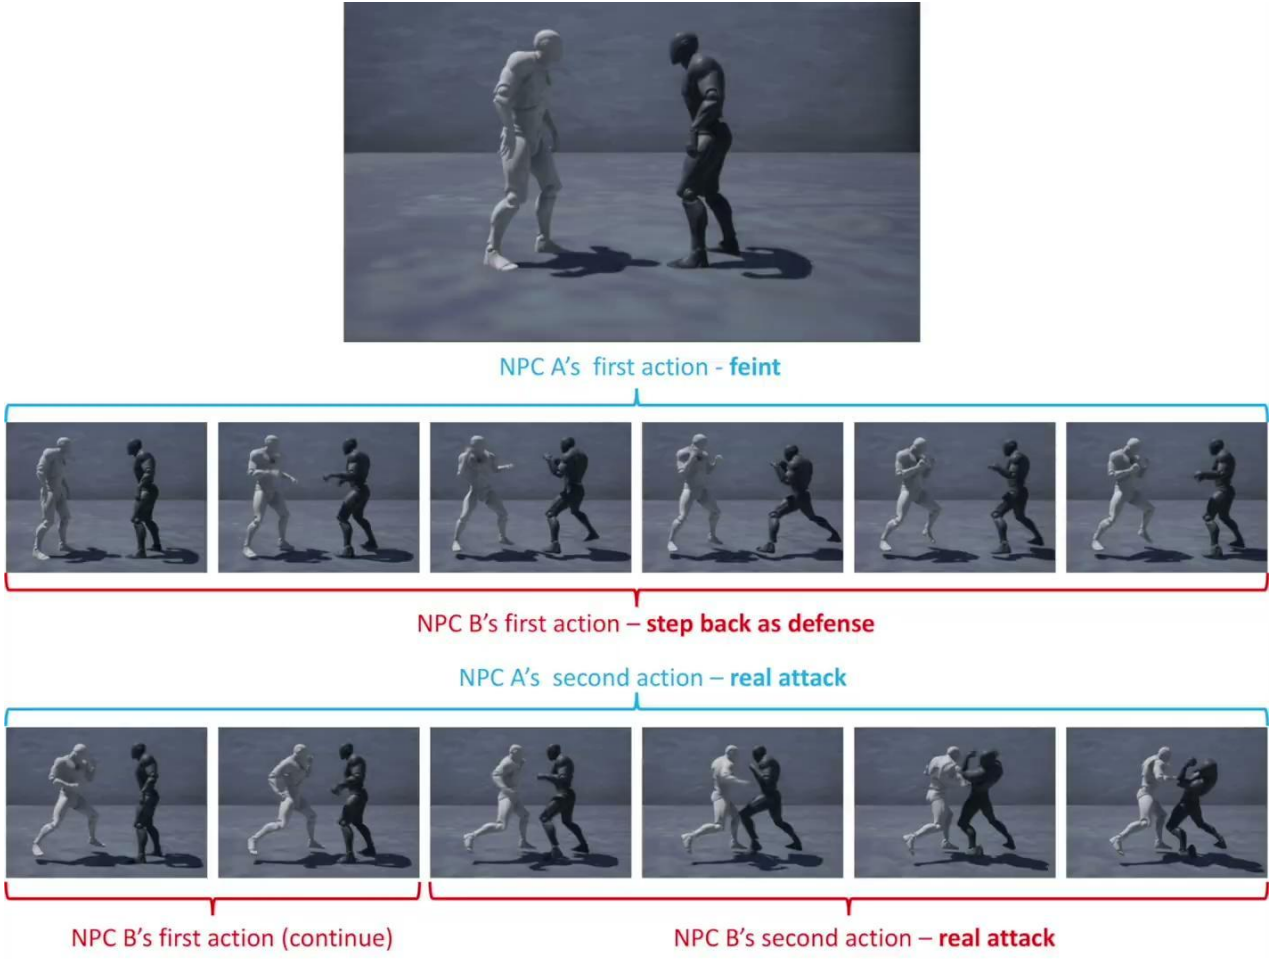

Outline

Background and Motivation

Feint Formalization in Action-Level

Feint Formalization in Strategy-Level

Proof-of-Concept Implementation

Evaluation Overview

Background and Motivation

Action-Level Formalization

- **Existing works:**
  - Animate of Feint behaviors as proof-of-concept of character animation techniques [1].
  - Learn control strategies from motion clips which incorporate nuanced behaviors like Feint behaviors [2].
- **Uniqueness of Feint behaviors:**
  - (Semi-)palindrome
  - Incorporation with other behaviors

Strategy-Level Formalization

- **Existing works:**
  - Neglect Feint behaviors.
  - Assume Feint behaviors are glitches of other behaviors, inducing same impact consideration and strategy learning.
- **Uniqueness of Feint behaviors:**
  - Temporal impacts
  - Spatial impacts
  - Collective impacts

[1] Kevin Wampler, Erik Andersen, Evan Herbst, Yongjoon Lee, and Zoran Popovic. Character animation in two-player adversarial games. *ACM Trans. Graph.*, 29(3):26:1–26:13, 2010.
[2] Jungdam Won, Deepak Gopinath, and Jessica K. Hodgins. Control strategies for physically simulated characters performing two-player competitive sports. *ACM Trans. Graph.*, 40(4):146:1–146:11, 2021.

Our Work Overview

Action-Level Formalization

- Feint characteristics and templates
- Feint behaviors in game steps

Concrete and Unified Implementation

- Address action and strategy level formalizations
- Can directly incorporate common MARL models and schemes

Strategy-Level Formalization

- Temporal impacts
- Spatial impacts
- Collective impacts

Experiments and Evaluation

- Game reward gains
- Diversity gains
- Overhead

Outline

Background and Motivation

Feint Formalization in Action-Level

Feint Formalization in Strategy-Level

Proof-of-Concept Implementation

Evaluation Overview

Feint Behavior Characteristics and Templates

Most real-world attack behaviors can be divided into 3 stages.

Feint templates can be automatically generated from them.

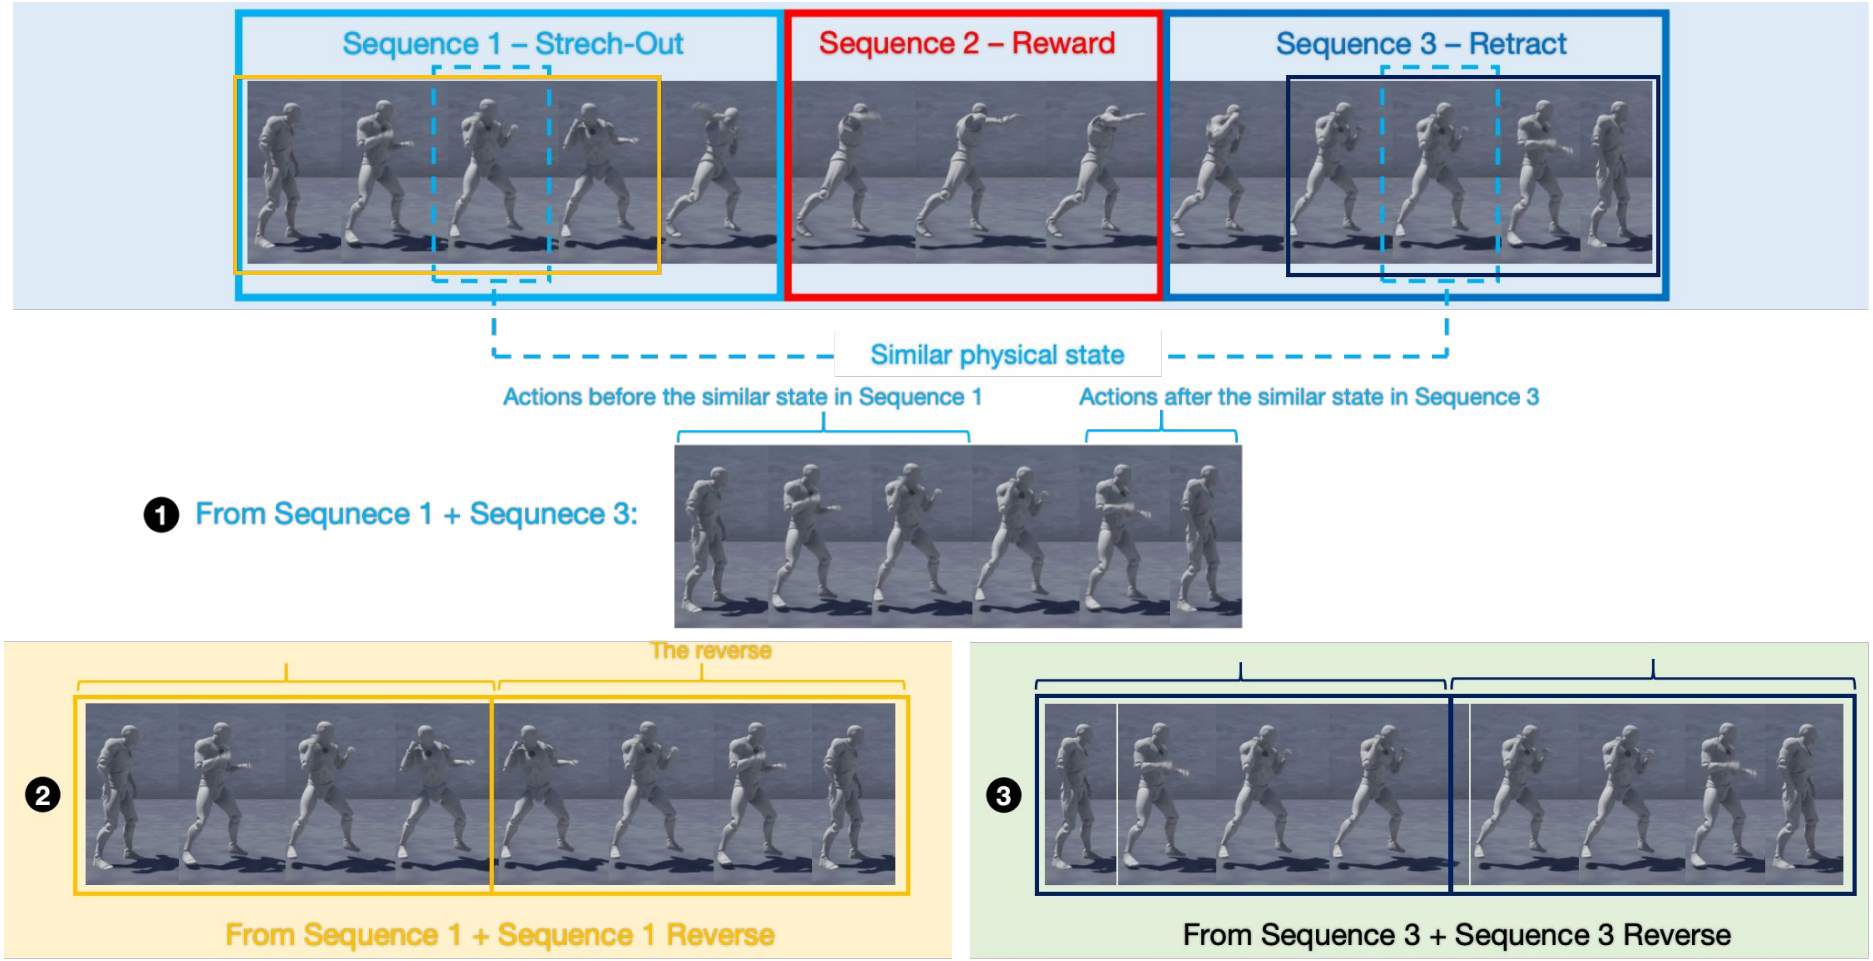
Sequence 1 – Strech-Out
Sequence 2 – Reward
Sequence 3 – Retract
Similar physical state
Actions before the similar state in Sequence 1
Actions after the similar state in Sequence 3
1 From Sequnece 1 + Sequnece 3:
The reverse
2
From Sequence 1 + Sequence 1 Reverse
3
From Sequence 3 + Sequence 3 Reverse

Feint Behavior in Consecutive Game Steps

Standalone Feint behavior is meaningless.
Need to be combined with other behaviors.

2 Constraints:

- Physical constraint
  - Lead to physically plausible follow-up behaviors.
- Effectiveness constraint
  - Enable temporal and spatial advantages.

⇨ **Dual-Behavior Model**

Successful Feint

Agent - **Feint Behavior (Fake punch towards the oppoent's head)**

Opponent's response to the Feint – **a full for defend for its head**

Temporal Advantage

Agent - **Followed-up attack (Hook towards the oppoent's waist)**

Opponent about to finish head defending

Opponent won't able to accomplish a defend for its waist in time

Unsuccessful Feint - Too Short

NPC A's first behavior - Feint (too short)

NPC A's second behavior - real attack

NPC B's first behavior - step back as defense

NPC A's second behavior - real attack (no effective reward since NPC B is still defending)

NPC B's first behavior (continue) - step back as defense

Unsuccessful Feint - Too Long

NPC A's first behavior - Feint (too long)

NPC B's first behavior - step back as defense

NPC A's first behavior (continue) - Feint

NPC A's second behavior - real attack (interrupted by NPC B)

NPC B's second action – real attack (effective reward)

Outline

Background and Motivation

Feint Formalization in Action-Level

Feint Formalization in Strategy-Level

Proof-of-Concept Implementation

Evaluation Overview

Temporal Dimension: Influence Time

Dynamic short-long term:

$$Rew_{short}(\pi_i^{'}, t_0, t_f, t_s, \boldsymbol{\alpha}) = \boxed{\sum_{t=t_0}^{t=t_0+t_f} \alpha_{Feint_t} R^i(s_t, a_t^i, a_t^{-i})} + \boxed{\sum_{t=t_s f}^{t=t_0+t_s} \alpha_{attack_t} R^i(s_t, a_t^i, a_t^{-i})}$$

$$Rew_{long}(\pi_i^{'}, t_0, t_s, T, \boldsymbol{\beta}) = \frac{1}{T} \sum_{t=t_0+t_s+1}^{T} \beta_t R^i(s_t, a_t^i, a_t^{-i})$$

$$\begin{aligned} Rew_{temporal}(\pi_i^{'}, t_0, t_f, t_s, T, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \lambda_{short} Rew_{short}(\pi_i^{'}, t_0, t_f, t_s, \boldsymbol{\alpha}) \\ + \lambda_{long} Rew_{long}(\pi_i^{'}, t_0, t_s, T, \boldsymbol{\beta}) \end{aligned}$$

Spatial Dimension: Influence Range

Model the influence range of Feint behaviors as the **divergence of occupancy measures** (Behavioral Diversity [3]) during policy learning.

Maximize the effective influence range under the influence distribution of Feint.

$$max_{\pi_i'} Rew_{spatial}(\pi_i^{'}, \pi_{-i}, s) = D_f(\rho_{\pi_i', \pi_{-i}}(s) \mid\mid \rho_{\pi_i, \pi_{-i}}(s))$$

[3] Xiangyu Liu, Hangtian Jia, Ying Wen, Yaodong Yang, Yujing Hu, Yingfeng Chen, Changjie Fan, and Zhipeng Hu. Unifying behavioral and response diversity for open-ended learning in zero-sum games. *CoRR*, abs/2106.04958, 2021.

Collective Impacts: Influence Degree

Jointly model the temporal and spatial impacts of Feint behaviors:

- Aggregate the **temporal reward** on **spatial domain** for all agents.
- Aggregate the **spatial reward** on **temporal domain** for all state transitions.

$$Rew_{collective}(\pi_i^{'}, \pi_{-i}) = \mu_1 \sum_{\pi \in \{\pi_i^{'}, \pi_{-i}\}} Rew_{temporal}(\pi, t_0, t_f, t_s, T, \boldsymbol{\alpha}, \boldsymbol{\beta}) + \mu_2 \sum_{s=s_0}^{s_T} Rew_{spatial}(\pi_i^{'}, \pi_{-i}, s)$$

Outline

Background and Motivation

Feint Formalization in Action-Level

Feint Formalization in Strategy-Level

Proof-of-Concept Implementation

Evaluation Overview

A General Implementation Scheme I

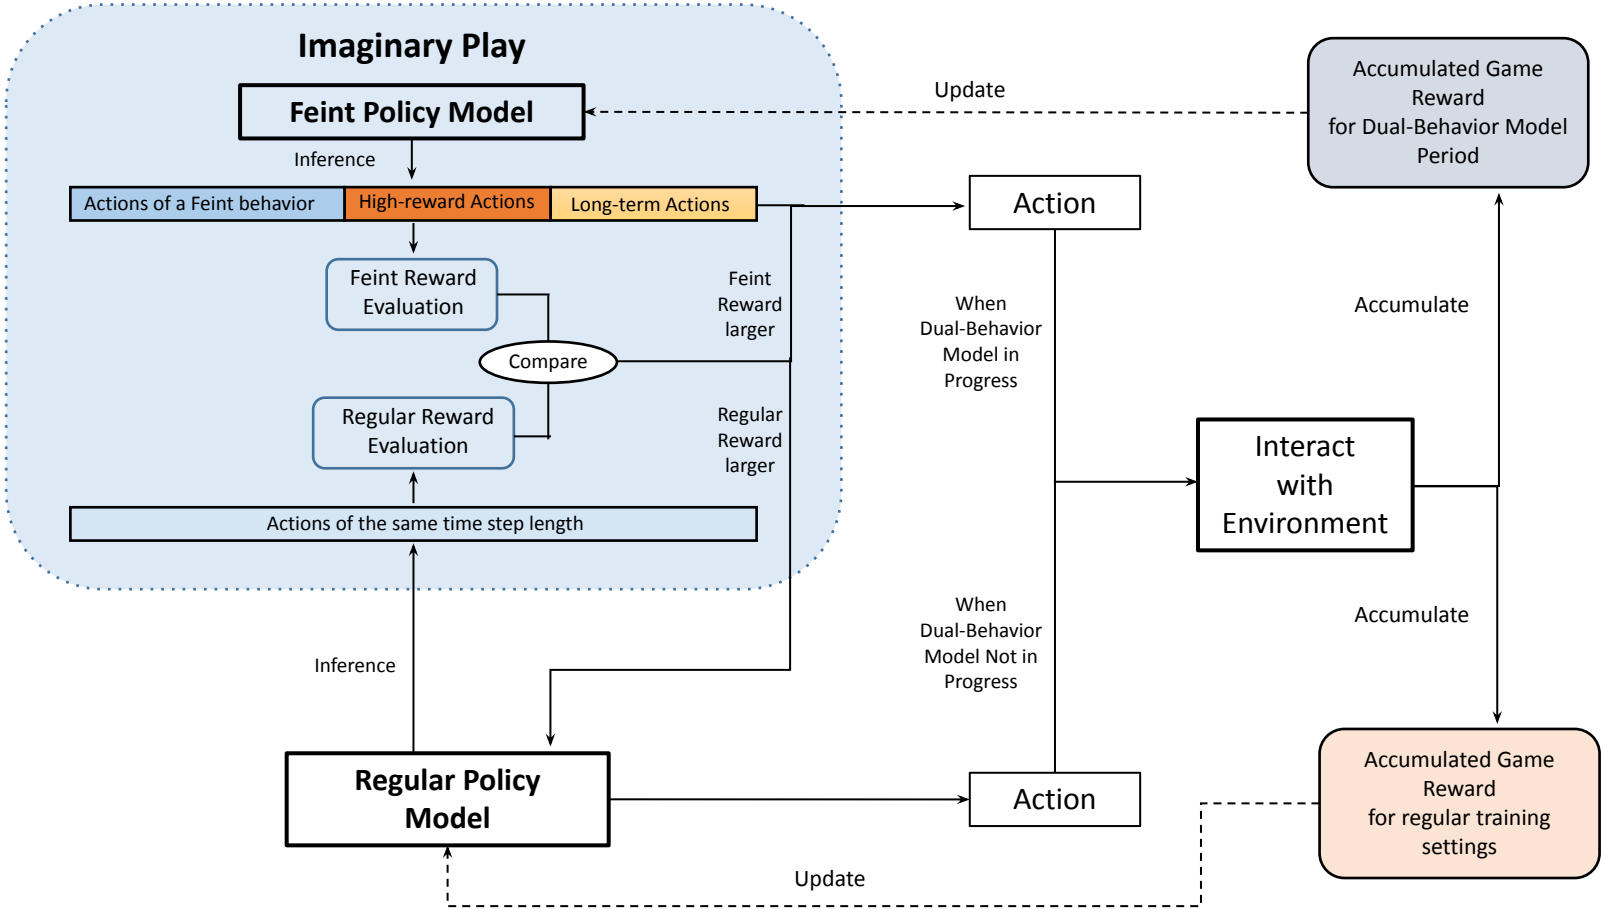

A General Implementation Scheme II

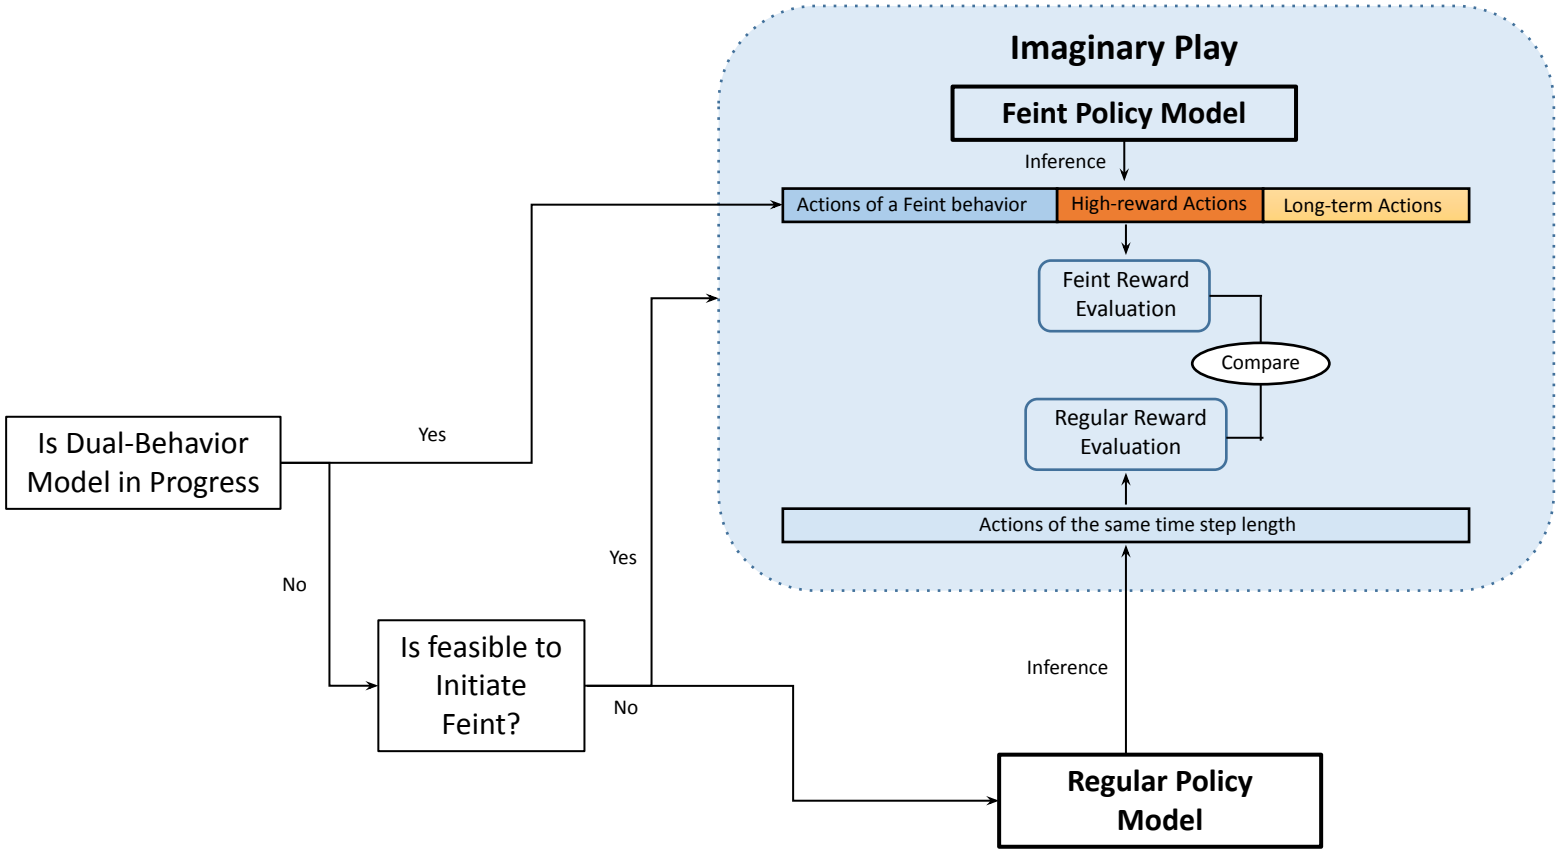

Outline

Background and Motivation

Feint Formalization in Action-Level

Feint Formalization in Strategy-Level

Proof-of-Concept Implementation

Evaluation Overview

Methodology

Game Scenarios:

- Complex Boxing Game
- Action space: Body movements
- Behaviors: Full set of Mixamo [4].
- Scenarios:
  - 1-vs-1 combat.
  - 3-vs-3 combat.

MARL Models:

MADDPG [5], MASAC [6], MATD3 [7], MAD3PG [8]

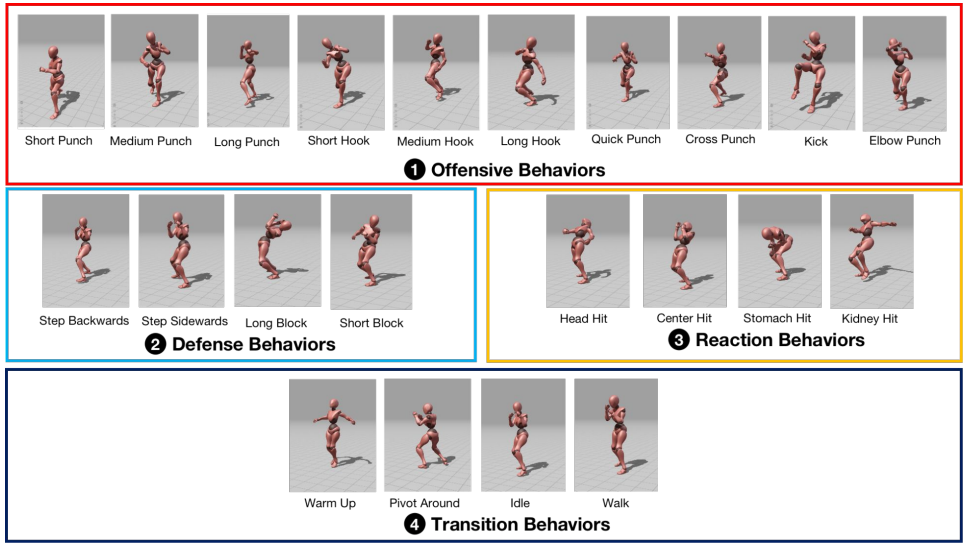

[5] Kevin Wampler, Erik Andersen, Evan Herbst, Yongjoon Lee, and Zoran Popovic. Character animation in two-player adversarial games. *ACM Trans. Graph.*, 29(3):26:1–26:13, 2010.

[6] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6379–6390, 2017.

[7] Shariq Iqbal and Fei Sha. Actor-attention-critic for multi-agent reinforcement learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2961–2970. PMLR, 2019.

[6] Gabriel Barth-Maron, Matthew W. Hoffman, David Budden, Will Dabney, Dan Horgan, Dhruva TB, Alistair Muldal, Nicolas Heess, and Timothy P. Lillicrap. Distributed distributional deterministic policy gradients. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.

[8] Johannes Ackermann, Volker Gabler, Takayuki Osa, and Masashi Sugiyama. Reducing overestimation bias in multi-agent domains using double centralized critics. *arXiv preprint arXiv:1910.01465*, 2019.

Game rewards Increase in 1-vs-1 Game

Effectively improve the actual game rewards.

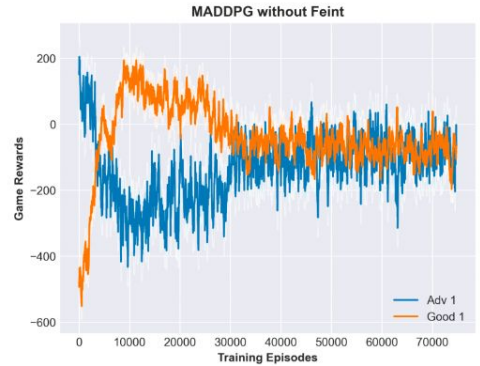

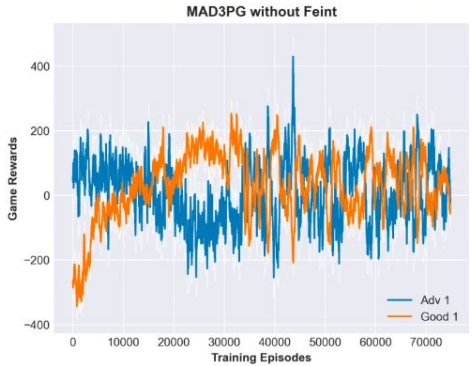

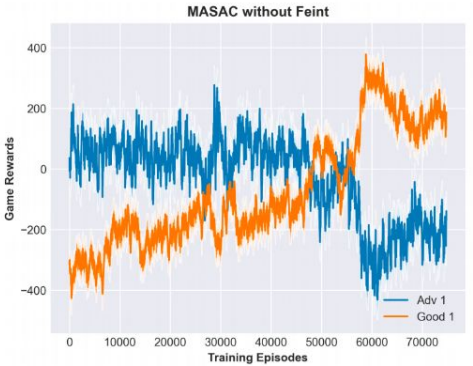

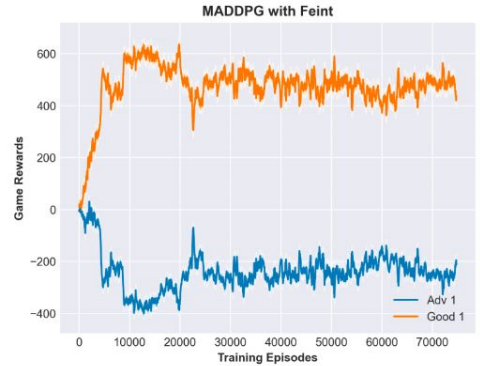

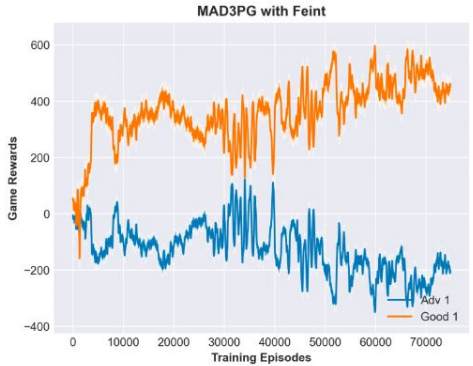

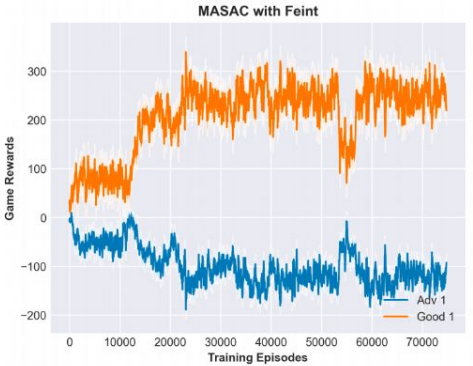

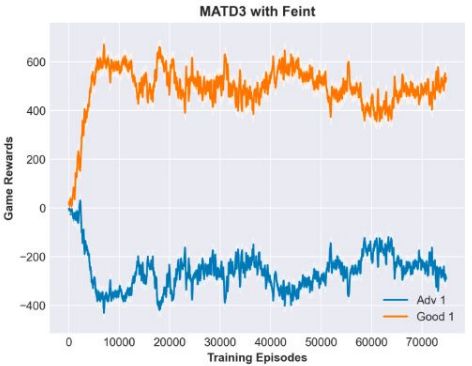

Game rewards Increase in 3-vs-3 Game

Effectively improve the actual game rewards.

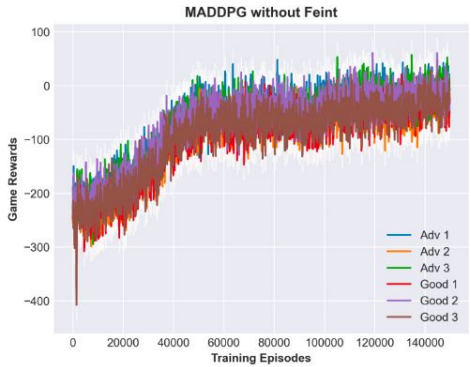

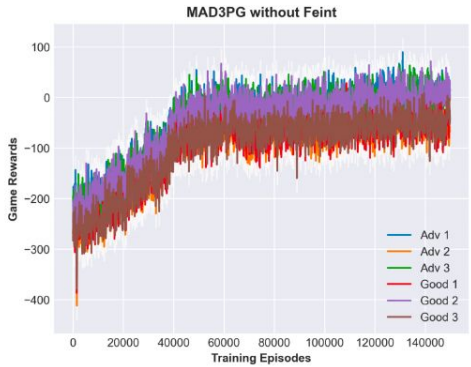

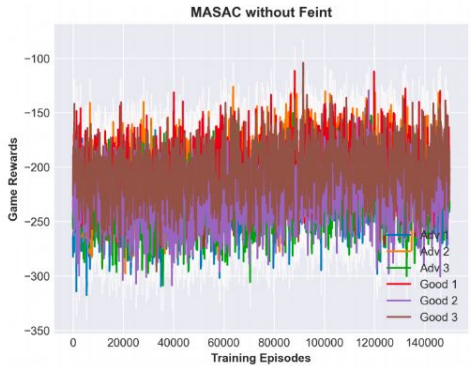

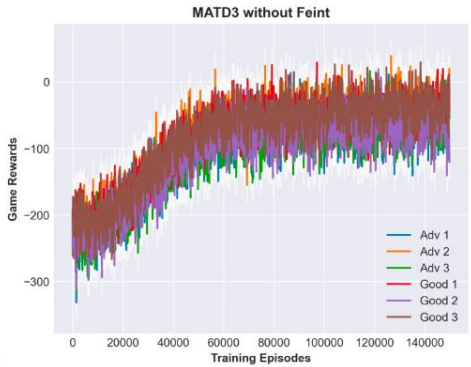

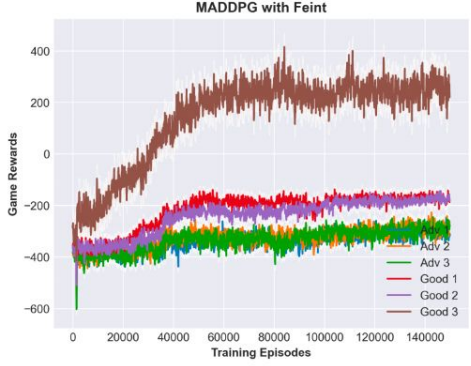

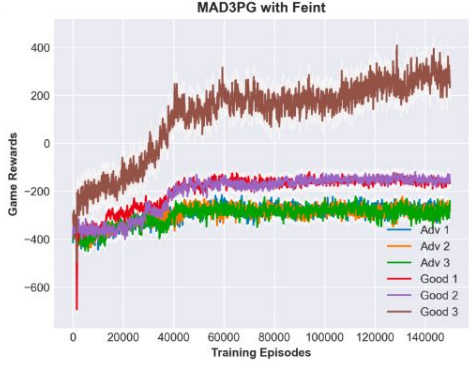

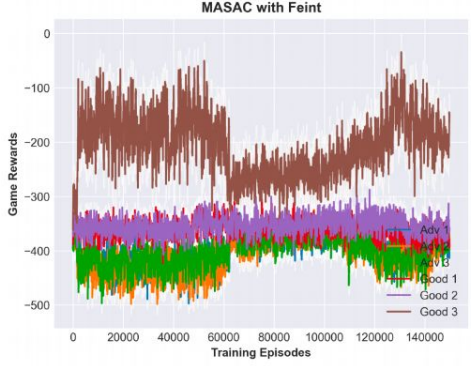

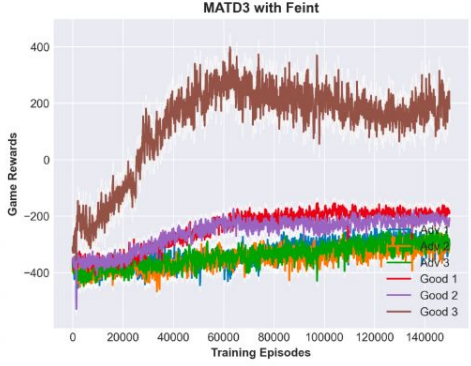

Other Results in the Paper

Diversity gain induced by Feint behaviors.

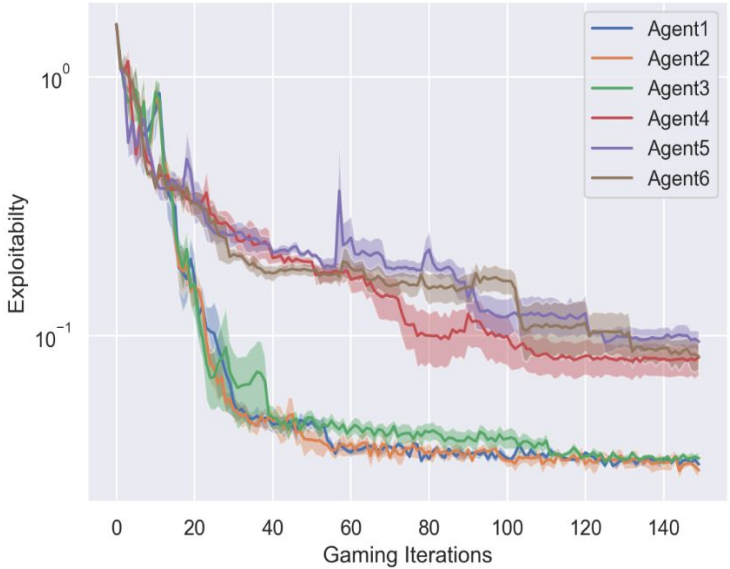

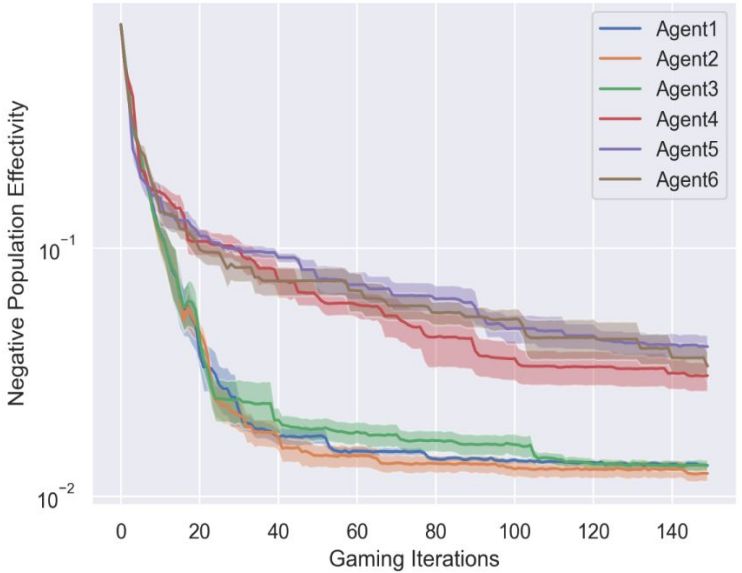

Other Results in the Paper

Neglectable overheads induced by our formalization.

1 VS 1

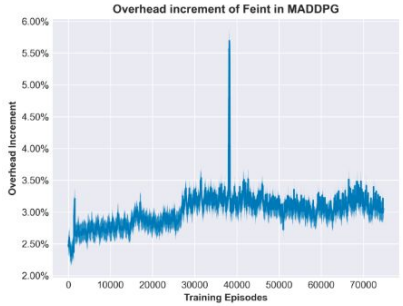

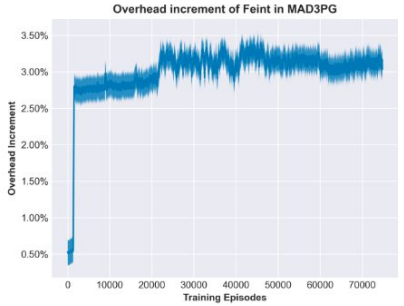

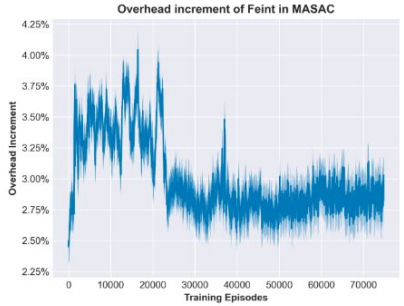

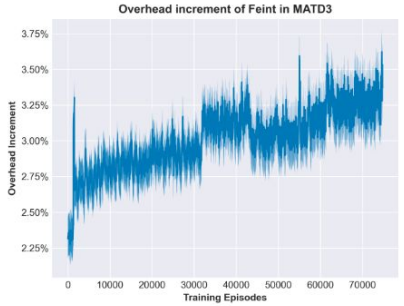

3 VS 3

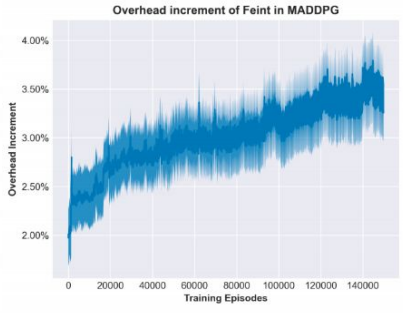

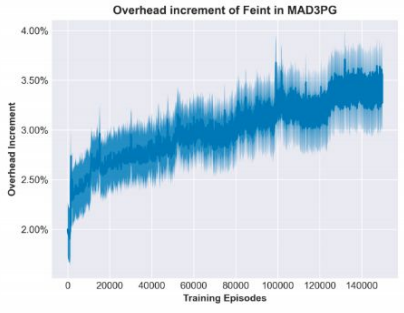

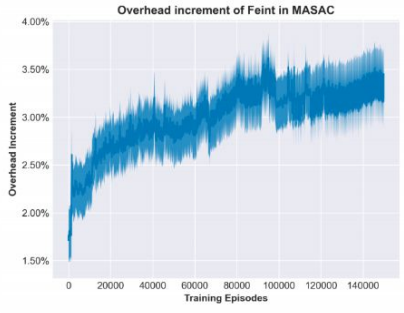

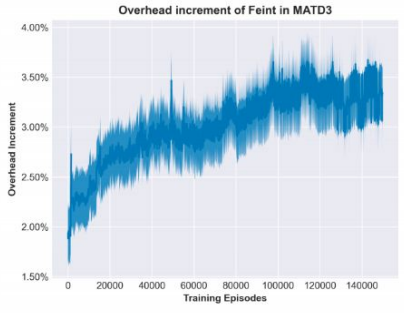

Summary

Action-Level Formalization

- **Feint characteristics and templates** - Palindrome-directed Templates
- **Feint behaviors in game steps** - Dual-behavior Model

Strategy-Level Formalization

- **Temporal impacts** - temporal advantage
- **Spatial impacts** - Maximize the spatial diversity impacts
- **Collective impacts** - collectively aggregate the temporal and spatial impacts in multi-agent environments.

Concrete and Unified Implementation

- Address action and strategy level formalizations.
- Can directly incorporate common MARL models and schemes.

Experiments and Evaluation

- **Game reward gains** - improvements
- **Diversity gains** - improvements
- **Overhead** - negligible