

Vision-Based Generic Potential Function for Policy Alignment in Multi-Agent Reinforcement Learning

Hao Ma^{1,2*}, Shijie Wang^{1,2*}, Zhiqiang Pu^{1,2†}, Siyao Zhao^{1,2}, Xiaolin Ai²

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²Institute of Automation, Chinese Academy of Sciences

{mahao2021, wangshijie2020, zhiqiang.pu, zhaosiyao2023, xiaolin.ai}@ia.ac.cn

Abstract

Guiding the policy of multi-agent reinforcement learning to align with human common sense is a difficult problem, largely due to the complexity of modeling common sense as a reward, especially in complex and long-horizon multi-agent tasks. Recent works have shown the effectiveness of reward shaping, such as potential-based rewards, to enhance policy alignment. The existing works, however, primarily rely on experts to design rule-based rewards, which are often labor-intensive and lack a high-level semantic understanding of common sense. To solve this problem, we propose a hierarchical vision-based reward shaping method. At the bottom layer, a visual-language model (VLM) serves as a generic potential function, guiding the policy to align with human common sense through its intrinsic semantic understanding. To help the policy adapt to uncertainty and changes in long-horizon tasks, the top layer features an adaptive skill selection module based on a visual large language model (vLLM). The module uses instructions, video replays, and training records to dynamically select suitable potential function from a pre-designed pool. Besides, our method is theoretically proven to preserve the optimal policy. Extensive experiments conducted in the Google Research Football environment demonstrate that our method not only achieves a higher win rate but also effectively aligns the policy with human common sense.

Introduction

Multi-agent reinforcement learning (MARL) has achieved great success in solving complex decision-making problems across lots of domains, including real-time strategy games (Vinyals et al. 2019), autonomous driving teams (Xu et al. 2018), robotic systems (Yang and Gu 2004; Gu et al. 2023), and sports games (Liu et al. 2023). However, MARL faces a very challenging issue: it is difficult to learn policies that conform to the real-world cognition of the domain in complex tasks, especially with sparse rewards (Zhu, Wang, and Wang 2018; Dossa et al. 2019).

Although state-of-the-art MARL algorithms can achieve high rewards in complex tasks with sparse rewards, their policies often diverges from human understanding. This discrepancy is reflected in the inconsistency between the



Figure 1: Multi-agent policies trained by MAPPO and HAPPO in GRF 11 vs 11 scenario. (a) MAPPO: The non-ball-holding players of the yellow team do not form reasonable formation to occupy valuable space. (b) HAPPO: The non-ball-holding players of the yellow team lack meaningful positioning and movement to provide effective passing opportunities for the ball-holding player.

learned policies in simulation environments and human common sense of the tasks. For instance, in Google Research Football (GRF) environment (Kurach et al. 2020), researchers typically hope to learn impressive tactical combinations which can inspire tactical choices in the real world. However, existing algorithms tend to hack the environment by finding and exploiting unintended shortcuts or loopholes to achieve higher rewards. (Bolander et al. 2018; Liu et al. 2023), which limits the effectiveness in real-world applications of MARL. As shown in Figure 1, two typical unreasonable policies trained by MAPPO (Yu et al. 2022) and HAPPO (Kuba et al. 2022) illustrate the lack of structured formations and coherent passing coordination.

Therefore, it is a challenging issue that MARL algorithms is hard to learn policies which conform to human common sense. This issue mainly have three reasons: (i) **Limitations of simulation environments.** The complexity and diversity of the real world make it challenging to capture all relevant variables and dynamic changes. Thus, the inability to construct an environment model that is identical to the real world is an unavoidable problem. (ii) **Difficulty in modeling human common sense.** Human common sense is often formed based on their rich experiences and comprehensive cognition, making it difficult to express accurately through simple mathematical formulas or rules. Consequently, we find it challenging to quantify and model human common

*These authors contributed equally.

†Corresponding author: zhiqiang.pu@ia.ac.cn.

sense using manually designed conventional reward functions, especially in complex and long-horizon multi-agent tasks. **(iii) Complexity and dynamic change of goals.** In MARL, the typical objective for agents is to optimize specific and quantifiable performance metrics. However, human decision-making processes in complex tasks often involve dynamically changing goals across multiple stages. The simplified goal settings in traditional MARL fail to capture the complexity and diversity of human decision-making, resulting in a significant gap between the learned policies and human common sense. Some researchers leverage imitation learning (Hussein et al. 2017) or reward shaping (Csikszentmihalyi 2015) to guide policies to align as closely as possible with human cognition. However, these methods have the drawback of relying on high-quality human behavior data and extensive domain knowledge.

Recent researches in foundation models have demonstrated notable achievements across various applications (Xiao et al. 2023; Qian et al. 2023; Wang et al. 2023; Abi Raad et al. 2024). These models are useful because their pre-training encompasses a vast amount of high-quality human behavior data and advanced understanding of various tasks. Therefore, the inherent capabilities of foundation models hold great potential and value in addressing the difficulty in modeling human common sense in complex MARL. Many researchers utilize the rich knowledge in visual-language models (VLM) to generate dense rewards for vision-based reinforcement learning tasks with sparse rewards (Adeniji et al. 2023; Rocamonde et al. 2023; Baumli et al. 2023). However, they primarily focus on single agent performing individual tasks, which does not address complex multi-agent tasks. In addition, to address the issues of complex and dynamically changing goals, researchers are dedicated to designing dense rewards that can capture the complexity of human decision-making, such as intrinsic rewards and potential-based rewards. However, existing designs still lack sufficient expressiveness and fail to model advanced semantics of common sense. Therefore, how to design comprehensive dense rewards that align with the common sense in human decision-making processes remains a significant challenge.

Based on the analyses, the paper investigates **Vision-based GEneric Potential Funtions (V-GEPPF)** to facilitate hierarchical policy alignment in MARL. Policy alignment refers to the process where the learned policies not only pursue performance optimization, but also align with human common sense and cognitive patterns, whose significance lies in enhancing practicality and meaningfulness of the policies.

Specifically, to address the challenge of modelling human sense as a reward in complex MARL, we design a VLM-based generic potential function at the bottom layer. The generic potential function utilizes rich pretrained human knowledge within the VLM to guide policy learning towards alignment with human common sense. Then, recognizing that human decision-making involves dynamically changing goals, our method requires policies to flexibly adapt to uncertainty and change. Meanwhile, temporal information is crucial for understanding and aligning with the complex

cognitive processes of human decision-making. Therefore, based on the characteristics that visual large language models (vLLM) have broader human knowledge and the ability to handle more complex input information than VLM, we design a vLLM-based adaptive skill selection module. The module use more informative instructions, replayed videos, and reflection on the records of last potential function to adaptively select appropriate next potential function from a pre-designed pool, aiming to achieve a more comprehensive alignment with human common sense at the top layer. Furthermore, our method can be theoretically proven to not alter the optimal policy. The primary contributions of this work are as follows:

- We propose a hierarchical vision-based reward shaping method for MARL, focusing on learning policies that better conform to human common sense, rather than simply optimizing for efficiency.
- To overcome the difficulty in modeling human common sense as rewards, we design VLM-based generic potential functions which utilize rich human knowledge from pre-trained data in VLM to guide policy alignment.
- Considering the dynamic objectives from multiple facets of the human decision-making process, we design a vLLM-based adaptive skill selection module to achieve a higher level and more comprehensive consistency between policies and human common sense.

Related Works

Viusal-language models. Foundational models trained on vast amounts of data can acquire broad and transferable representations of various data types, such as images and languages, making them applicable to a variety of downstream tasks (Wei et al. 2022). As type of foundation models, visual-language models (VLM) integrate linguistic and visual signals, playing a crucial role in fields that require processing of both modalities (Wang et al. 2021; Ju et al. 2022). There are diverse applications of VLM in reinforcement learning (RL). For instance, VLM can be used as reward functions (Dang, Edelkamp, and Ribault 2023), promptable representation learners (Chen et al. 2024), and for data augmentation based on hindsight relabeling (Sumers et al. 2023). The applications demonstrate wide applicability and robust capabilities of VLM, providing effective tools and new perspectives for addressing complex RL challenges.

VLM as RL rewards. Currently, a promising application direction is using VLM to generate dense rewards for RL tasks, especially in scenarios with sparse rewards (Fu et al. 2024; Rocamonde et al. 2023; Wang et al. 2024). Cui et al. (Cui et al. 2022) utilize the pretrained CLIP to provide image-based rewards for robotic manipulation tasks. Mahmoudieh et al. (Mahmoudieh, Pathak, and Darrell 2022) successfully apply the fine-tuned CLIP as a language-described reward model for robotic tasks. Sontakke et al. (Sontakke et al. 2024) use VLM in a robotic environment to provide reward signals for RL agents, primarily defining rewards through video demonstrations. However, existing methods primarily focus on solving specific single-agent

tasks, where the design of the reward function tends to be relatively simple and direct. Moreover, these methods are mostly implemented in vision-based environments. In contrast, multi-agent systems involve more complex interactions and collaboration mechanisms, making dense reward guidance more challenging.

Potential-based rewards. In MARL tasks with sparse rewards, the introduction of additional dense rewards to represent domain-specific expert knowledge has become a common practice to accelerate training (Mguni et al. 2021; Ma et al. 2022). However, it soon became apparent that if used improperly, the dense rewards might alter the optimal response, which directly affects the performance of the policies (Simsek and Barto 2006). Ng et al. (Ng, Harada, and Russell 1999) demonstrate that the optimal policy do not change if the dense reward is designed as $F(s, s') = \gamma\phi(s') - \phi(s)$. Devlin et al. (Devlin and Kudenko 2012) further prove that time-varying potential-based rewards $F(s, t, s', t') = \gamma\phi(s', t') - \phi(s, t)$ do not alter the Nash equilibrium in multi-agent problems. Grzes (Grzes 2017) extend the proof to episodic tasks that terminate after a certain final time step N . Chen et al. (Chen et al. 2022) design potential-based rewards in GRF to be positively correlated with the distance between the ball and the goal. Zhang et al. (Zhang et al. 2023) design the rewards that are proportionally linked to the remaining health of a unit in StarCraft II (Samvelyan et al. 2019). However, existing methods are often based on simple formulas or rules, making it difficult to comprehensively represent human common sense.

Background

Partially observable Markov decision process. A multi-agent reinforcement learning problem in vision-based tasks can be formally defined as a partially observable Markov decision process (POMDP) (Puterman 1990). A POMDP is represented as a tuple: $\langle \mathcal{N}, \mathcal{S}, \mathcal{A}, \mathcal{P}, r, \mathcal{G}, \mathcal{O}, \gamma \rangle$, where $|\mathcal{N}| = N$ is the number of agents, $s \in \mathcal{S}$ is the state space, $\gamma \in [0, 1]$ is a discount factor. At each time step t , the agent $i \in \mathcal{N}$ chooses its action $a_t^i \in \mathcal{A}$ based on its observations $o_t^i \in \mathcal{O}$. The actions of all agents constitute the joint action space $a \in \mathcal{A}^N$. The state and joint action make up a joint reward $r(s, a, s')$ based on the state transition function $\mathcal{P}(s'|s, a)$. The team’s objective is to maximize the expected discounted return $G = \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, s_{t+1})$.

CLIP models. VLM (Zhang et al. 2024) is a typical foundation model that can handle sequences incorporating both language inputs $l \in \mathcal{L}^{\leq n}$ and visual inputs $v \in \mathcal{V}^{\leq m}$. \mathcal{L} represents a finite alphabet encompassing strings of length no greater than n , while \mathcal{V} represents the space for 2D RGB images whose length do not exceed m . Contrastive language-image pretraining model (CLIP) (Radford et al. 2021) is a representative VLM which trains by aligning image and text embeddings in the latent space. Specifically, the CLIP model consists of a text encoder τ_L and an image encoder τ_I , both of which are mapped to a same latent space \mathbb{R}^k . Typically, the encoders are trained by minimizing the cosine distance between embeddings for pairs of images and languages.

Methodology

The goal of this paper is to explore how to learn policies that are more aligned with human common sense in complex MARL with sparse rewards. We propose a hierarchical vision-based reward shaping method named V-GEPF, which utilizes rich human knowledge from foundation models to guide policy alignment from different levels. The algorithm framework is described in Appendix A.1.

VLM as Generic Potential Functions

Guiding MARL policies to align with human common sense is a challenging problem. Human common sense is typically formed based on rich experiences and comprehensive cognition of specific tasks, making it challenging to be modelled and quantified through simple formulas or rules, especially in complex and long-horizon multi-agent tasks.

Recent researches have shown effectiveness of additional dense rewards, such as potential-based rewards, to guide policy alignment (Chen et al. 2022; Zhang et al. 2023). However, although these dense rewards are useful in certain tasks, they do have some limitations: **(i) Insufficient expressiveness.** Manually designed dense rewards are overly simplistic, failing to fully capture the multifaceted goals and high-level semantic information of human common sense in complex tasks. **(ii) Limited understanding ability.** Dense rewards often rely on the intuition and experience of experts, which may overlook complex behavioral dynamics and environmental factors in tasks, leading to a lack of in-depth and comprehensive understanding of complex tasks. **(iii) Lack of adaptability and flexibility.** Dense rewards are typically fixed, which are struggle to cope with dynamic changes in the environment. These limitations make it difficult to effectively align multi-agent policies with human common sense in complex decision-making tasks through existing dense rewards.

The VLM possesses inherent exceptional capabilities due to the rich expert knowledge it has acquired from training datasets encompassing various complex tasks. Therefore, we design VLM as generic potential functions to guide policies in exploring directions that align with human common sense.

Instruction-conditioned potential function. To design VLM-based generic potential functions, we first introduce an instruction-conditioned potential function. By incorporating the concept of instructions into the potential function, we lay the groundwork for exploring the alignment of VLM-guided policies with human common sense.

Definition 1. An *instruction-conditioned potential function* $\phi(s|l) : S \times L \rightarrow R$ is a mapping from the current state s to its latent value ϕ , given the language instruction l .

The instruction-conditioned potential function does not directly depend on the agents’ actions, but is determined by the state under the given instructions. The instructions reflect the guidance of human common sense. It is worth noting that when considering trajectories with a finite time step, the final step N has the terminal state, whose potential value

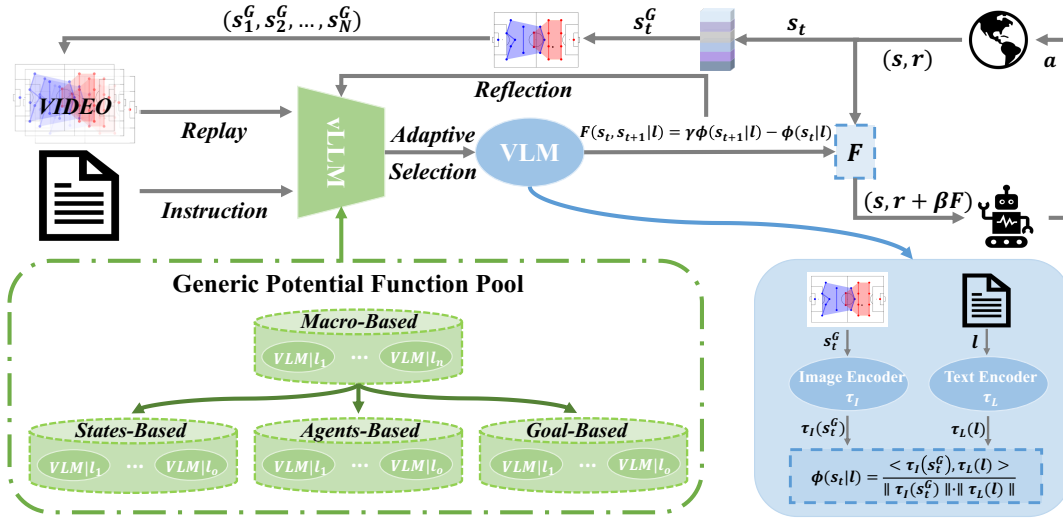


Figure 2: Framework of the V-GEPF. The state’s image representation s_t^G and the human instructions l are input separately into image encoder and text encoder. The cosine distance computed from the outputs is designed as a generic potential function $\phi(s_t|l)$, which is weighted and combined with the original environmental rewards to guide the learning of policies. Furthermore, the video replay of last episode, the initial human instructions, information about potential function pool, and the reflection on the records of last potential function are fed into vLLM to adaptively select appropriate generic potential function at various training phase.

should be set to zero (Wierstra et al. 2008):

$$\phi(s_t|l) = \begin{cases} 0 & \text{if } t = N \\ \phi(s_t|l) & \text{otherwise.} \end{cases} \quad (1)$$

VLM as a generic potential function. The pre-training data of VLM includes a vast repository of high-quality human understanding and cognition related to complex tasks. As a result, VLM has significant potential and value in overcoming the challenges of modeling human common sense in complex MARL tasks.

Contrastive language-image pretraining model (CLIP) (Radford et al. 2021) is a typical efficient VLM which is designed as generic potential functions in this paper. CLIP consists of an image encoder, τ_I , and a text encoder, τ_L . These encoders map images and text to embeddings in the same latent space. They are trained on a dataset containing a large number of image-text pairs by minimizing the cosine distance between the corresponding embeddings. Therefore, images and texts with similar semantics can be mapped to similar embeddings. We first convert the environmental state s_t into an image s_t^G , enabling the CLIP to more intuitively understand the current situational context. Subsequently, the image and the corresponding instruction are fed into the image encoder and text encoder respectively, producing two embeddings $\tau_I(s_t^G)$ and $\tau_L(l)$. These embeddings reflects the semantics of instruction and current state. To align the policy with human common sense, the VLM-based generic potential function is designed as the cosine distance between the state and instruction embeddings:

$$\phi(s_t|l) = \frac{\langle \tau_I(s_t^G), \tau_L(l) \rangle}{\|\tau_I(s_t^G)\| \cdot \|\tau_L(l)\|}. \quad (2)$$

Then, based on Def. 1, we subsequently design **generic potential-based reward** as:

$$F(s_t, s_{t+1}|l) = \gamma\phi(s_{t+1}|l) - \phi(s_t|l), \quad (3)$$

where s_t and s_{t+1} represent the current and next state respectively, γ is the discount factor of MARL to ensure the designed rewards do not alter the optimal policy (detailed proofs can be found in Appendix A.2). With the cross-modal representation capabilities of CLIP, this generic potential-based reward can enhance the agents’ understanding of the semantic of state. Instruction l is set to represent a ideal state that align with human common sense, e.g. ‘the home players show cooperative attacking tactics’. Detailed instructions l are provided in Appendix B.2. The visualization details of states s_t^G are presented in Appendix B.1.

Since there is an original reward r_{env} provided by the environment, the total reward is described as:

$$R(s_t, s_{t+1}) = r_{env}(s_t, s_{t+1}) + \rho \cdot F(s_t, s_{t+1}|l), \quad (4)$$

where ρ represents a scalar coefficient to balance the generic potential-based reward with the environmental reward.

It can be theoretically demonstrated that incorporating the generic potential-based reward ensures that the Nash equilibrium in multi-agent problems remains unchanged. Detailed proofs are presented in Appendix A.2.

However, the decision-making process in complex long-horizon tasks often involves multiple dynamically changing goals, which means that the guiding direction of the policy needs to adaptively adjust at different stages. This is a key characteristic of human decision-making process. In this context, generic potential-based reward with a fixed instruction is inadequate, as it cannot comprehensively and accurately quantify the diversity and dynamics of these complex

goals. Meanwhile, the lack of temporal information limits the understanding and aligning with the complex cognitive processes involved in human decision-making. These limitations hinder our policies’ ability to effectively align with human common sense to some extent. To enhance flexibility and adaptability in our method, we design a vLLM-based adaptive skill selection module.

vLLM-based Adaptive Skill Selection

In complex long-horizon tasks, human decision-making process can adapt to dynamic changes and involve multidimensional objectives. Taking football match as an example, a team’s policies will flexibly adjust according to the progress of the match and the current situation. At the beginning of the match or when the advantage is not significant, a team often employ more aggressive offensive tactics, focusing on coordination among players to achieve rapid passing or multi-player attacks. When holding a significant lead, the team pay more attention to maintaining the stability of their formation and controlling ball possession to preserve their advantages (Harrop and Nevill 2014).

To align the decision-making process of MARL with human cognition, the designed rewards must be capable of adaptively and flexibly adjusting to the uncertain and changing circumstances. Leveraging the advanced understanding and representation capabilities of the visual large language model (vLLM), we propose a vLLM-based adaptive skill selection module. This module dynamically selects different skills (VLMs with pre-defined instructions) at various stages of training, allowing MARL agents to adjust their behaviors based on the evolving game context.

We design a generic potential function pool emulating the thought processes involved in human decision-making. As researched by Simon (Simon 1955), in the decision-making process, humans first attempt to understand the macro environment, then turn their attention to the specific details at the micro-level, which includes the analysis of individual behaviors and the pursuit of goals. Thus, different generic potential functions are designed from macro and micro perspectives to guide policies exploring towards the respective desired directions. At the macro level, the potential functions focus on the overall situation (such as team advantage) in order to set macro objectives in complex environments. At the micro level, we analyze from three perspectives: states-based, agents-based, and goal-based.

- States-based skills focus on analyzing microscopic state information, such as the local advantage around a key agent.
- Agents-based skills focus on the behaviors and interactions of the agents to uncover the complex dynamics of their actions and relationships, such as the level of coordination between the agents.
- Goal-based skills aim to maintain a keen focus on specific objectives, such as the status of key targets.

These different types of skills enable the policies to make more refined and targeted choices and adjustments in various contexts, promoting a more comprehensive alignment

with human common sense. We provide a more detailed description of the designed instructions in Appendix B.2.

Furthermore, information on the temporal dimension is crucial for capturing the continuity and evolution of policy learning. Thus, we present the entire episode’s states in a video replay to provide temporal information as input to the vLLM.

Overall, the vLLM-based adaptive skill selection module and the VLM-based generic potential function hierarchically construct a vision-based generic potential function. The vLLM receives the video replay of the last episode, initial human instructions, information about the potential function pool, and the reflection on records of the previous potential function. Using these information, a VLM-based potential function is selected from the pool, providing multi-angle, multi-level guidance for policy alignment. The dialogue examples of this module are provided in Appendix B.3.

Experiments

Experimental Setup

All experiments are implemented in a typical long-horizon multi-agent benchmark Google Research Football (GRF) (Kurach et al. 2020), which simulates real football matches and serves as a widely used benchmark for researchers in football game AI. We evaluate our method on GRF full-field (11 vs 11) tasks with different difficulty levels. In the full-field tasks, all 11 players of home team’s are controlled by RL policies, while the opposing 11 players are managed by a built-in script with adjustable difficulty settings. Each player has a 19-dimensional discrete action space, with all players sharing global state information. This constitutes a vast exploration space. We set the reward to be given only when a goal is scored, and each episode can be up to 3000 steps long, which constitutes a sparse reward problem. The full-field tasks are sufficiently challenging, making it very difficult for MARL to learn good collaborative policies within them.

V-GEPF is implemented based on MAPPO and compared against MAPPO (Yu et al. 2022), IPPO (De Witt et al. 2020), and HAPPO (Kuba et al. 2022). During training, all algorithms set 32 workers to sample training data in parallel, with each worker sampling 3,000 steps per epoch. The training ends after 300 epochs, resulting in approximately 28 million steps. Additionally, V-GEPF adaptively selects skills every 50 epochs. For the VLM, we use CLIP-RN50¹, for the vLLM, we employ MiniCPM-Llama3-V². To further evaluate V-GEPF, we introduce a baseline called MAPPO-xT, which also utilizes potential-based rewards. Using data from the FIFA World Cup 2018³, we calculate expected threat (xT) values for different positions on the field. In MAPPO-xT, the mapping from positions to xT values (the higher the better) serves as a potential function to guide MAPPO’s policy learning. Calculation details of xT can be found in Ap-

¹<https://github.com/openai/CLIP>

²https://huggingface.co/openbmb/MiniCPM-Llama3-V-2_5

³https://figshare.com/collections/Soccer_match_event_dataset/4415000

pendix C.2. The hyperparameters of all algorithms can be found in Appendix C.3.

Careful Comparison to SOTA

We compare the average win rate curves of V-GEPF and state-of-the-art (SOTA) baselines during training. As illustrated in Figure 3(a), in the 11 vs. 11 easy task, V-GEPF outperforms all baselines in terms of win rate, significantly surpassing IPPO. While HAPPO, MAPPO, and MAPPO-xT exhibits an advantage in convergence speed, their final win rates are lower than that of V-GEPF. The slower convergence speed is attributed to the potential-based reward guiding the policy to explore cooperative policy that align with human common sense, e.g. dribbling, keep formation, rather than solely focusing on scoring goals.




Figure 3: Average win rate curves during training in GRF 11 vs 11 scenario.

In the 11 vs. 11 hard task, as illustrated in Figure 3(b), V-GEPF significantly outperforms all baselines in terms of win rate. Both MAPPO, MAPPO-xT, and HAPPO’s average win rate curves plateau around 15 million steps. By visualizing SOTA policies, we observe that they have learned strange policies (as shown in Figure 1): one player dribbles the ball toward the goal while the other players fail to cooperate, instead running toward the opponent’s sideline or congregating at the corner or edge of the opponent’s field. This represents a local optimum that RL can easily discover. However, once a policy is trapped in this local optimum, it becomes challenging to explore alternative policies due to the vast exploration space of this task. In contrast, V-GEPF benefits from the adaptive guidance provided by the vLLM and VLM, enabling the policy to avoid getting stuck in this local optimum. As a result, V-GEPF eventually achieves a higher win rate compared to the baselines.

We further use radar maps for a fine-grained comparison between the policies of MAPPO and V-GEPF. As shown in Figure 4, the six dimensions in the radar chart characterize how well a policy performs beyond the average win ratio, which can comprehensively reflect human common sense. V-GEPF is superior to MAPPO in all aspects, indicating that V-GEPF leads to a policy that is more aligned with human cognition.

Potential Function Analysis

We select one run of V-GEPF, and record the average value of the potential function per epoch to analyze in detail how the potential function curve changes during training. After

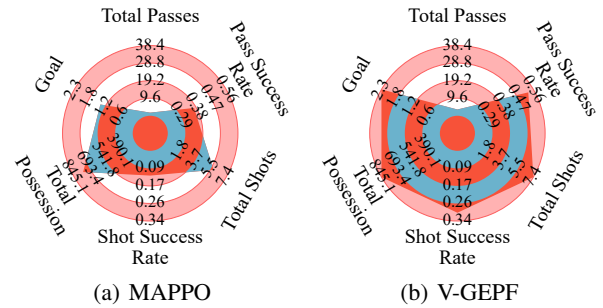


Figure 4: Comparison of policy styles trained with MAPPO and MAPPO enhanced by V-GEPF.

each skill selection, we use the first epoch to calculate the mean and standard deviation of the potential function and normalize it, so the potential function is always zero at the beginning. As shown in Figure 5, each potential function shows an upward trend, which indicates that VLM is able to guiding policy learning given different instruction. From the selection of skill, we find that vLLM can flexibly adapt skills to address the deficiencies of the current policy.

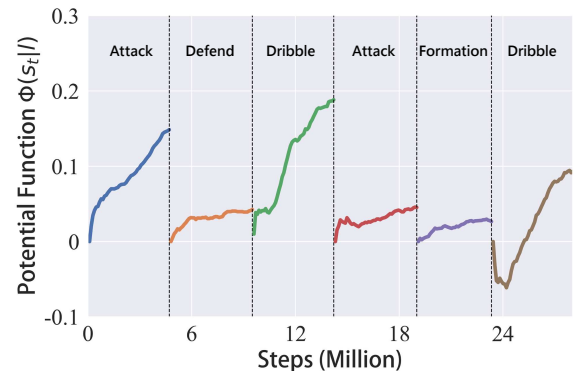


Figure 5: Potential function curves during training. Six VLM-based potential functions are selected sequentially by a vLLM according to replayed videos.

Ablation on Adaptive Skill Selection

To validate the necessity of the adaptive skill selection module, we compare the policies guided by V-GEPF and a fixed VLM-based potential function using radar charts (Figure 7). The charts show that the policies guided by a fixed VLM-based potential function exhibit strong stylistic tendencies, with the policy style aligning well with corresponding instruction. For example, an instruction favoring offence leads to policies with high numbers of shots and shot success rates, while a instruction favoring dribbling results in policies with high pass counts and pass success rates, as this effectively prevents the ball from being intercepted. However, football as a long-horizon complex task, requires a more balanced and adaptive policy. As shown in Figure 4(b), the adaptive skill selection module selects the suitable VLM-based potential function at different training stages, resulting in a




Figure 6: A visual policy analysis of the V-GEPF method in an offensive phase. The visualization shows that the agents have organized positioning and coherent passing combinations, which aligns with human common sense in real football matches.

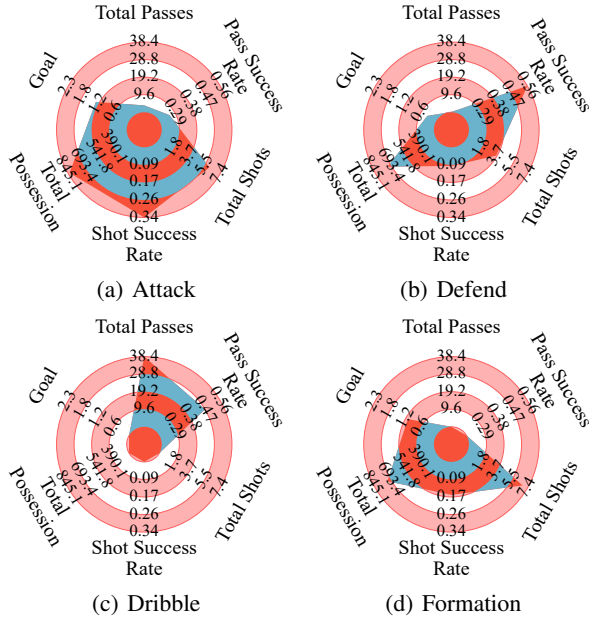


Figure 7: Comparison of the policy styles trained using different VLM-based generic potential functions. Their propensity are attack, defend, dribble, and formation, respectively. The detailed instructions provided to VLM can be found in the Appendix B.2.

more balanced policy style compared to not using the module.

Visualization Analysis

We further analyze the policies learned by V-GEPF through visualization. Our method not only improves the win rate, but also learns policies with standard formation and coherent passing, resulting in effects that align with human common sense in football matches.

In the offensive sequence illustrated in Figure 6, the team executes three distinct passing patterns: transitioning from the backfield to midfield to the forward line, moving from the center to the flanks and back, and switching from long-

distance to short-distance passes.

Specifically, in the first frame, the goalkeeper A initiates a long pass to transfer the ball to right back B. In the second frame, as right back B controls the ball, center back E and right midfielder F move closer to provide support. However, the space created by defensive midfielder C’s forward run poses a greater threat. Thus, B delivers a long ball to the vacant space in the center, targeting C. In the third frame, it is noteworthy that with C advancing into the attack, F positions itself to recover in order to protect against a potential counterattack from the opponents. Center midfielder D prepares to receive the ball, and naturally, C plays a short pass to D. Finally, in the fourth frame, D executes a dribble and takes a shot, resulting in a goal.

The policy demonstrates that the V-GEPF encourages agents to exhibit cooperative behaviors that align with human cognition. Firstly, agents can establish positioning that adheres to the required formation lines. Secondly, during the passing process, agents are capable of making appropriate supporting runs (such as C, E, and F) which provide the receiver with more options, as well as protective runs (such as F) which prevent counterattacks in the event of an unsuccessful offensive play. Furthermore, agents can actively seek to pass into areas that pose a greater offensive threat rather than simply dribbling the ball individually.

Conclusion

This paper investigates how to align MARL policies with human common sense in complex, long-horizon tasks. While reward shaping methods have proven effective in guiding policies, modeling common sense as a reward remains a challenge. To address this issue, we propose a vision-based reward shaping method called V-GEPF. V-GEPF employs a hierarchical framework that leverages vision information effectively. At the bottom layer, a VLM is used as a generic potential function to efficiently give real-time guidance. At the top layer, a vLLM reflects on the current policy and adjusts the subsequent guidance accordingly. Through experiments, we demonstrate the effectiveness of our proposed method compared to SOTA methods in both improving win rates and aligning with human common sense.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grant 62322316, the Beijing Nova Program under Grant 20220484077 and 20230484435, and the Strategic Priority Research Program of Chinese Academy of Science under Grant No. XDA27030204.

References

- Abi Raad, M.; Ahuja, A.; Barros, C.; Besse, F.; Bolt, A.; Bolton, A.; Brownfield, B.; Buttimore, G.; Cant, M.; Chakera, S.; et al. 2024. Scaling instructable agents across many simulated worlds. *arXiv e-prints*, arXiv–2404.
- Adeniji, A.; Xie, A.; Sferrazza, C.; Seo, Y.; James, S.; and Abbeel, P. 2023. Language reward modulation for pretraining reinforcement learning. *arXiv preprint arXiv:2308.12270*.
- Baumli, K.; Baveja, S.; Behbahani, F.; Chan, H.; Comanici, G.; Flennerhag, S.; Gazeau, M.; Holsheimer, K.; Horgan, D.; Laskin, M.; et al. 2023. Vision-Language Models as a Source of Rewards. *arXiv preprint arXiv:2312.09187*.
- Bolander, T.; Engesser, T.; Mattmüller, R.; and Nebel, B. 2018. Better eager than lazy? How agent types impact the successfulness of implicit coordination. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*.
- Chen, M.; Pu, Z.; Pan, Y.; and Yi, J. 2022. Knowledge transfer from situation evaluation to multi-agent reinforcement learning. In *International Conference on Neural Information Processing*, 3–14. Springer.
- Chen, W.; Mees, O.; Kumar, A.; and Levine, S. 2024. Vision-language models provide promptable representations for reinforcement learning. *arXiv preprint arXiv:2402.02651*.
- Csikszentmihalyi, M. 2015. Intrinsic rewards and emergent motivation. In *The hidden costs of reward*, 205–216. Psychology Press.
- Cui, Y.; Niekum, S.; Gupta, A.; Kumar, V.; and Rajeswaran, A. 2022. Can foundation models perform zero-shot task specification for robot manipulation? In *Learning for dynamics and control conference*, 893–905. PMLR.
- Dang, X.; Edelkamp, S.; and Ribault, N. 2023. CLIP-Motion: Learning reward functions for robotic actions using consecutive observations. *arXiv preprint arXiv:2311.03485*.
- De Witt, C. S.; Gupta, T.; Makoviichuk, D.; Makoviychuk, V.; Torr, P. H.; Sun, M.; and Whiteson, S. 2020. Is independent learning all you need in the starcraft multi-agent challenge? *arXiv preprint arXiv:2011.09533*.
- Devlin, S. M.; and Kudenko, D. 2012. Dynamic potential-based reward shaping. In *11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, 433–440. IFAAMAS.
- Dossa, R. F. J.; Lian, X.; Nomoto, H.; Matsubara, T.; and Uehara, K. 2019. A human-like agent based on a hybrid of reinforcement and imitation learning. In *2019 international joint conference on neural networks (IJCNN)*, 1–8. IEEE.
- Fu, Y.; Zhang, H.; Wu, D.; Xu, W.; and Boulet, B. 2024. FuRL: Visual-Language Models as Fuzzy Rewards for Reinforcement Learning. *arXiv preprint arXiv:2406.00645*.
- Grzes, M. 2017. Reward shaping in episodic reinforcement learning. In *Sixteenth International Conference on Autonomous Agents and Multiagent Systems*, 565–573. ACM.
- Gu, S.; Kuba, J. G.; Chen, Y.; Du, Y.; Yang, L.; Knoll, A.; and Yang, Y. 2023. Safe multi-agent reinforcement learning for multi-robot control. *Artificial Intelligence*, 319: 103905.
- Harrop, K.; and Nevill, A. 2014. Performance indicators that predict success in an English professional League One soccer team. *International Journal of Performance Analysis in Sport*, 14(3): 907–920.
- Hussein, A.; Gaber, M. M.; Elyan, E.; and Jayne, C. 2017. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2): 1–35.
- Ju, C.; Han, T.; Zheng, K.; Zhang, Y.; and Xie, W. 2022. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, 105–124. Springer.
- Kuba, J.; Chen, R.; Wen, M.; Wen, Y.; Sun, F.; Wang, J.; and Yang, Y. 2022. Trust Region Policy Optimisation in Multi-Agent Reinforcement Learning. In *ICLR 2022-10th International Conference on Learning Representations*, 1046. The International Conference on Learning Representations (ICLR).
- Kurach, K.; Raichuk, A.; Stańczyk, P.; Zajac, M.; Bachem, O.; Espeholt, L.; Riquelme, C.; Vincent, D.; Michalski, M.; Bousquet, O.; et al. 2020. Google research football: A novel reinforcement learning environment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 4501–4510.
- Liu, B.; Pu, Z.; Pan, Y.; Yi, J.; Liang, Y.; and Zhang, D. 2023. Lazy agents: a new perspective on solving sparse reward problem in multi-agent reinforcement learning. In *International Conference on Machine Learning*, 21937–21950. PMLR.
- Ma, Z.; Wang, R.; Li, F.-F.; Bernstein, M.; and Krishna, R. 2022. Elign: Expectation alignment as a multi-agent intrinsic reward. *Advances in Neural Information Processing Systems*, 35: 8304–8317.
- Mahmoudieh, P.; Pathak, D.; and Darrell, T. 2022. Zero-shot reward specification via grounded natural language. In *International Conference on Machine Learning*, 14743–14752. PMLR.
- Mguni, D. H.; Jafferjee, T.; Wang, J.; Slumbers, O.; Perez-Nieves, N.; Tong, F.; Yang, L.; Zhu, J.; Yang, Y.; and Wang, J. 2021. Ligs: Learnable intrinsic-reward generation selection for multi-agent learning. *arXiv preprint arXiv:2112.02618*.
- Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, 278–287.
- Puterman, M. L. 1990. Markov decision processes. *Handbooks in Operations Research and Management Science*, 2: 331–434.

- Qian, C.; Cong, X.; Yang, C.; Chen, W.; Su, Y.; Xu, J.; Liu, Z.; and Sun, M. 2023. Communicative agents for software development. *arXiv preprint arXiv:2307.07924*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rocamonde, J.; Montesinos, V.; Nava, E.; Perez, E.; and Lindner, D. 2023. Vision-language models are zero-shot reward models for reinforcement learning. *arXiv preprint arXiv:2310.12921*.
- Samvelyan, M.; Rashid, T.; De Witt, C. S.; Farquhar, G.; Nardelli, N.; Rudner, T. G.; Hung, C.-M.; Torr, P. H.; Foerster, J.; and Whiteson, S. 2019. The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*.
- Simon, H. A. 1955. A behavioral model of rational choice. *The quarterly journal of economics*, 99–118.
- Simsek, O.; and Barto, A. G. 2006. An intrinsic reward mechanism for efficient exploration. In *Proceedings of the 23rd international conference on Machine learning*, 833–840.
- Sontakke, S.; Zhang, J.; Arnold, S.; Pertsch, K.; Bıyık, E.; Sadigh, D.; Finn, C.; and Itti, L. 2024. Roboclip: One demonstration is enough to learn robot policies. *Advances in Neural Information Processing Systems*, 36.
- Sumers, T.; Marino, K.; Ahuja, A.; Fergus, R.; and Dasgupta, I. 2023. Distilling internet-scale vision-language models into embodied agents. *arXiv preprint arXiv:2301.12507*.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *nature*, 575(7782): 350–354.
- Wang, G.; Xie, Y.; Jiang, Y.; Mandlekar, A.; Xiao, C.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2023. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Wang, Y.; Sun, Z.; Zhang, J.; Xian, Z.; Biyik, E.; Held, D.; and Erickson, Z. 2024. Rl-vlm-f: Reinforcement learning from vision language foundation model feedback. *arXiv preprint arXiv:2402.03681*.
- Wang, Z.; Yu, J.; Yu, A. W.; Dai, Z.; Tsvetkov, Y.; and Cao, Y. 2021. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*.
- Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; Yogatama, D.; Bosma, M.; Zhou, D.; Metzler, D.; et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Wierstra, D.; Schaul, T.; Peters, J.; and Schmidhuber, J. 2008. Episodic reinforcement learning by logistic reward-weighted regression. In *International Conference on Artificial Neural Networks*, 407–416. Springer.
- Xiao, Z.; Zhang, D.; Wu, Y.; Xu, L.; Wang, Y. J.; Han, X.; Fu, X.; Zhong, T.; Zeng, J.; Song, M.; et al. 2023. Chain-of-Experts: When LLMs Meet Complex Operations Research Problems. In *The Twelfth International Conference on Learning Representations*.
- Xu, Z.; Lyu, Y.; Pan, Q.; Hu, J.; Zhao, C.; and Liu, S. 2018. Multi-vehicle flocking control with deep deterministic policy gradient method. In *2018 IEEE 14th International Conference on Control and Automation (ICCA)*, 306–311. IEEE.
- Yang, E.; and Gu, D. 2004. Multiagent reinforcement learning for multi-robot systems: A survey. Technical report, tech. rep.
- Yu, C.; Velu, A.; Vinitzky, E.; Gao, J.; Wang, Y.; Bayen, A.; and Wu, Y. 2022. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35: 24611–24624.
- Zhang, J.; Huang, J.; Jin, S.; and Lu, S. 2024. Vision-language models for vision tasks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhang, T.; Liu, Z.; Yi, J.; Wu, S.; Pu, Z.; and Zhao, Y. 2023. Multiexperience-Assisted Efficient Multiagent Reinforcement Learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zhu, M.; Wang, X.; and Wang, Y. 2018. Human-like autonomous car-following model with deep reinforcement learning. *Transportation research part C: emerging technologies*, 97: 348–368.