

Accelerating Reinforcement Learning of Robotic Manipulations via Feedback from Large Language Models

Kun Chu, Xufeng Zhao, Cornelius Weber, Mengdi Li, Stefan Wermter
Knowledge Technology, Department of Informatics, University of Hamburg, Germany
kun.chu@uni-hamburg.de

Abstract: Reinforcement Learning (RL) plays an important role in the robotic manipulation domain since it allows self-learning from trial-and-error interactions with the environment. Still, sample efficiency and reward specification seriously limit its potential. One possible solution involves learning from expert guidance. However, obtaining a human expert is impractical due to the high cost of supervising an RL agent, and developing an automatic supervisor is a challenging endeavor. Large Language Models (LLMs) demonstrate remarkable abilities to provide human-like feedback on user inputs in natural language. Nevertheless, they are not designed to directly control low-level robotic motions, as their pretraining is based on vast internet data rather than specific robotics data. In this paper, we introduce the **Lafite-RL** (Language agent feedback interactive Reinforcement Learning) framework, which enables RL agents to learn robotic tasks efficiently by taking advantage of LLMs’ timely feedback. Our experiments conducted on RL Bench tasks illustrate that, with simple prompt design in natural language, the Lafite-RL agent exhibits improved learning capabilities when guided by an LLM. It outperforms the baseline in terms of both learning efficiency and success rate, underscoring the efficacy of the rewards provided by an LLM.

Keywords: Robots, Large Language Models, Interactive Reinforcement Learning

1 Introduction

Reinforcement Learning (RL) has shown its power in solving sequential decision-making problems in the robotic domain [1, 2], through optimizing control policies directly from trial-and-error interactions with environments. However, there are still several challenges [3], like sample inefficiency and difficulties in specifying rewards, limiting its applications to the field.

Inspired by how we human beings learn skills from more knowledgeable persons such as teachers or supervisors, a potential solution for the above limitations is learning from human expert guidance, so as to inject additional information into the learning process. Human guidance has shown some benefits in terms of providing additional rewards or guidance to accelerate the learning of new tasks, including learning from human demonstrations [4, 5] and feedback [6, 7, 8, 9]. However, collecting sufficient human guidance is time-consuming and costly.

Recently, Large Language Models (LLMs) have shown remarkable abilities to generate human-like responses in the textual domain [10, 11], and their applications have been explored in the robotic domain. While some approaches prompt LLMs to instruct robots in performing tasks [12, 13, 14, 15], they focus on utilizing LLMs’ common-sense knowledge to give high-level advice for employing pre-trained or hard-coded low-level control policies, which requires much data collection or expert knowledge respectively. Since these works do not perform policy learning when executing tasks with LLMs, the robots’ performance highly depends on the LLM’s capabilities and consistent presence

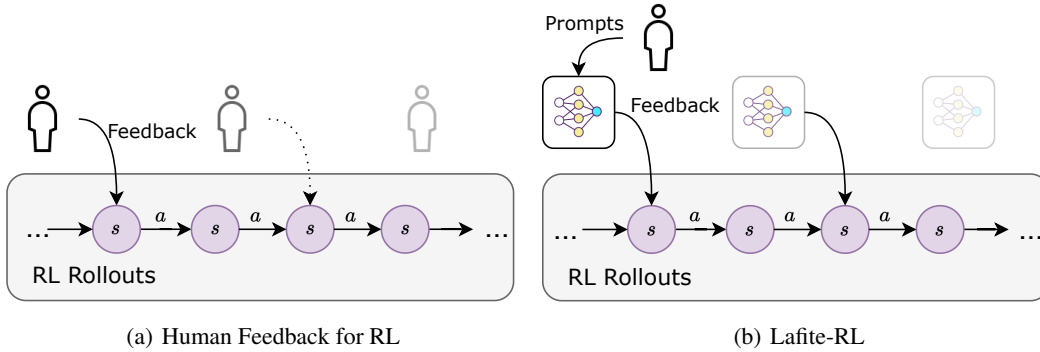


Figure 1: Comparisons of human feedback for RL and Lafite-RL. Normal human feedback for RL requires frequent human interaction during the learning process. In contrast, Lafite-RL allows a human to interact with the LLM only once, where the human prompts the LLM to observe the RL process and provide real-time feedback.

during the interactions each time tasks are executed. Moreover, owing to their limited experience in the continuous 3D robotic world, LLMs are hard to employ for low-level robotic motion control. Direct learning low-level control policies taking advantage of feedback from LLMs is a promising solution, which leads to a research question: *can LLMs understand scene information in robotic space and provide appropriate feedback to supervise reinforcement learning agents?*

To this aim, we propose the **Lafite-RL (Language agent feedback interactive Reinforcement Learning)** framework that enables RL agents to efficiently learn user-specified robotic tasks with LLMs’ real-time feedback, as shown in Fig. 1. Before task learning, a human user with domain knowledge, but not necessarily robot programming knowledge, just needs to design two prompts: one to let the LLM understand the robotic scene, and the other to instruct it on how to evaluate the robot’s behavior. Then, the LLM is able to provide interactive rewards mimicking human evaluative feedback, accelerating the RL process while freeing up human effort during the RL agent’s numerous interactions with the task environment. Besides, the RL agent acquires information from the task environment and is robust to potential mistakes of the LLM. Experiments on RLBench tasks demonstrate our method’s higher learning efficiency compared to the baseline method, which showcases that with easily designed prompts, an LLM’s feedback is effective in accelerating RL in robotic manipulations.

2 Related Works

2.1 Reinforcement Learning from Human Guidance

There are several kinds of human guidance used for improving RL agents’ ability to solve sequential decision-making tasks [16], mainly including human demonstrations [5] and feedback such as evaluative scores [6] and preferences[8].

Learning from human demonstrations. Inverse RL [4, 17, 5] is an important technique of learning from human demonstrations, which focuses on inferring a reward function through observing collected human demonstrations. Then, the RL agent is able to learn policies efficiently with an estimated reward function that succinctly describes the current task. Other works combined imitation learning techniques with Deep RL [18, 19], in which demonstrations are mostly used to initialize the agent’s policy. However, those approaches normally require to collect high-quality demonstrations, which are not applicable to those task environments that are hard to demonstrate, e.g., when robots have a high degree of freedom.

Learning from human feedback. Another set of helpful sources of information comes from diverse human feedback, mainly including evaluative scores [6, 9, 20] and trajectory preferences [8, 21,

22]. Some works allow humans to observe state-action pairs when the agent interacts with the environment following the current policy, and provide feedback either explicitly like via keyboard keys [6, 20] or implicitly like natural language [23, 24], gestures [25], facial expressions [26], or their combinations [27, 28]. Instead of evaluating state-action pairs, learning from human preferences allows human experts to give feedback on a set of trajectories in terms of rankings or preferences [8, 21]. Those forms of feedback are interactively collected and then injected into the RL processes in terms of rewards [6, 8, 24] or policies [29, 20, 30]. While the existence of numerous forms of human feedback, the quality of human feedback is also an important aspect when discussing how to inject external knowledge into the RL process [29, 31]. Recently, learning from human preferences has also shown its power in fine-tuning LLMs to make them generate more human-like responses [32]. Despite the successes shown in the works mentioned above, they require a significant amount of human feedback, which is normally time-consuming and expensive to collect.

2.2 Large Language Models in Robotics

The recent advancements in LLMs exhibit their remarkable abilities to generate complex text on a wide range of topics [11, 33]. Likewise, some works utilize LLMs’ zero-shot reasoning capabilities to guide agents on robotic tasks [12, 13, 14, 15, 34, 35]. With several skills libraries either manually designed by experts or pertained on massive data, these approaches do not perform any model training but instead use LLMs as planners for giving advice in a high-level manner. In this sense, robots’ performance highly depends on the LLM’s capabilities and requires the LLM’s presence each time when the task is performed. Meanwhile, due to limited text corpora in the continuous 3D robotic world, LLMs struggle with direct low-level robotic motion control. In contrast to theirs, our framework aims to prompt the LLM to provide feedback to train RL-based models for use-specified tasks. After the training, the RL-based model is able to perform the task with a higher success rate without LLM’s feedback or response.

Several works use LLMs to provide rewards for RL. Yu et al. [35] propose to prompt LLMs to define reward functions with specific parameters to optimize control policies on a variety of robotic tasks. However, their method relies on much manual work by experts, and the reward function is domain-specific and thus can not be applicable to broader domains. In contrast, our framework enables a normal human user to easily design the prompt, and interactive rewards are easy to understand and apply to various task domains. While Lee et al. [36] and Kwon et al. [37] demonstrate the effectiveness of employing LLMs’ feedback in terms of preferences and evaluative scores in the RL context, they are mainly on textual world tasks like text summarizations [36] and negotiation games [37]. In contrast, we aim to enable RL agents to benefit from LLMs’ feedback when training on robotic manipulations, which requires LLMs’ reasoning ability to understand the robot’s behavior in continuous 3D space.

3 Language Agent Feedback Interactive Reinforcement Learning

We formulate the learning agent’s interaction with the task environment in RL as a Markov Decision Process (MDP) in terms of a tuple $(\mathcal{S}, \mathcal{A}, \gamma, \mathcal{T}, \mathcal{R})$, where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $\gamma \in [0, 1]$ is the discount factor, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition probability function, and \mathcal{R} is the reward function that maps the state-action-state pair to a numerical reward $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$, such that at timestep t , the RL agent receives a reward r_t when transiting to state s_{t+1} after executing an action a_t at state s_t . By interacting with the environment, an RL algorithm seeks to learn an optimal policy π that maximizes the accumulated rewards.

Inspired by the practice of letting humans provide evaluative feedback for the agent’s behaviour as in Section 2.1, we propose the **Lafite-RL** framework that enables an LLM to generate real-time interactive rewards like a “human teacher” to accelerate the robot’s RL process. As shown in Fig. 2, with designed prompts, for the request at timestep t , our framework enables the LLM to “observe” the transition $s_t \rightarrow s_{t+1}$ and understand the robot’s motion a_t under the current task requirement




Figure 2: Depiction of proposed Lafite-RL framework. Before learning a task, a user provides designed prompts (see Table 1), including descriptions of the current task background and desired robot’s behaviors, and specifications for the LLM’s missions with several rules respectively. Then, **Lafite-RL** enables an LLM to “observe” and understand the scene information which includes the robot’s past action, and evaluate the action under the current task requirements. The language parser transforms the LLM response into evaluative feedback for constructing interactive rewards.

and generate an evaluation for it. Then, with a language parser, the textual response is transformed to a numerical value and injected into the RL process.

LLM as Scene Observer. We develop a scene descriptor to describe the scene based on the original information from the environment, including coordinates of the gripper and objects, the gripper’s opening/closing state, and the gripper sensor information (if some object is detected or grasped). This description illustrates the movement from the previous state to the current state. We provide a prompt to the LLM that contains a function API: $distance(objectA, objectB)$ since the LLM requires geometric information from the environment. In this way, the LLM is required to understand the robot’s motion under the current task requirement, and then make several function calls that can help itself for further evaluation. Based on the original information from the environment, and some outputs after the function calls, the LLM is prompted to give the final evaluation of the robot’s motion.

LLM as Motion Evaluator. For the evaluation, LLM is prompted to infer which phase the agent is during this period, and then evaluate if the current motion is good or not. A long-horizon manipulation task consists of many steps that belong to several phases, i.e., the gripper is required to approach the desired object, grasp the desired object, and lift the desired object to the given place. Therefore, the LLM needs to estimate which phase it is currently in since the optimal behaviors for each phase are different. Then, the LLM will generate its judgment for this motion in terms of *good move* or *bad move*, which is formed into a numerical score by the language parser. Formally, the reward for the RL agent at time t is a combination of environment reward r_t and scaffolding reward r_{llm} by an LLM:

$$r_t = r_{env} + r_{llm}. \quad (1)$$

The additional reward r_{llm} is parsed from the LLM response, mapped as

$$r_{llm} = \begin{cases} 1, & \text{if good move,} \\ -1, & \text{if bad move,} \\ 0, & \text{else.} \end{cases} \quad (2)$$

The language parser interprets the response of the LLM based on keyword detection. In case it fails to identify the LLM output as *good move* or *bad move*, it will pass zero to the agent.

Lafite-RL provides an interface for human users’ prompt engineering and RL agents’ training with LLMs’ continual real-time feedback. Table 1 shows snippets of prompts designed in Lafite-RL. Instead of explicitly designing the reward function, the human user describes the requirements of the task in natural language and guides the LLM to understand the scene information and evaluate the robot’s motion accordingly.

Table 1: Snippets of the two prompts required by Lafite-RL for learning one task. The missions of Scene Observer (●) and Motion Evaluator (●) are highlighted with different colors in the prompts.

Prompt 1	Prompt 2
<p>Task Background:</p> <p>A robot arm with a gripper on the work table is accomplishing the task of ...</p> <p>Your mission:</p> <p>We have the following observation of the robot's motion, and you are required to understand the scene under the current task requirement ...</p> <p>Rules:</p> <ol style="list-style-type: none"> 1. Generate several function calls concurrently, specifying the object names. 2. Only utilize the provided 'distance' function API. Creating or calling other functions is prohibited. <p>Now respond with the following observation ...</p>	<p>Task Background:</p> <p>A robot arm with a gripper on the work table is accomplishing the task of ...</p> <p>Your mission:</p> <p>We have the following observation of the robot's motion and some further information, and you are required to evaluate the robot's motion ...</p> <p>Rules:</p> <p>You must clearly give your judgements in terms of 'good move' or 'bad move', and your explanations.</p> <p>Now respond with the following observation ...</p> <p>Distance information:</p>

4 Experiments

We evaluate **Lafite-RL** to determine the extent to which an LLM can assist RL for robotic tasks in terms of evaluative feedback. For the LLM setup, we choose Vicuna-13B v1.5 [38] as the underlying LLM model, which is a lightweight open-source model fine-tuned on the LLaMA-2 [39] model. For higher throughput and faster response time, we adopt the technique of vLLM [40] as an augmentation of our local LLM. For the task learning setup, we chose a state-of-the-art RL algorithm, DreamerV3 [41] as the underlying RL algorithm, and conducted experiments on RLbench [42], a suite of large-scale benchmark and learning environments designed for robot manipulation tasks. Tasks in RLbench are viewed as challenging for RL agents, as their rewards are sparse, i.e., the agent only obtains scores on the goal state and zero otherwise. Our environmental setup consisted of a Franka Emika Panda robot with one camera each on the wrist, front, and overhead sides,




Figure 3: Illustration of two RLbench tasks for evaluation: (a) Push buttons, and (b) Take umbrella out of umbrella stand.

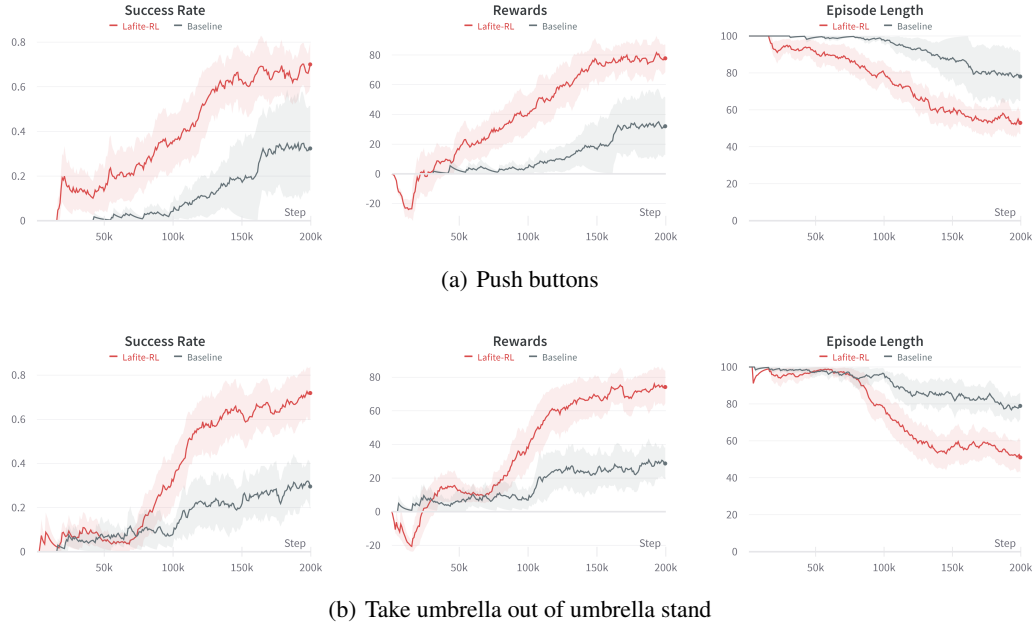


Figure 4: Learning curves of Lafite-RL (red curves) and the baseline without interactive feedback (gray curves) during training on two RL Bench tasks.

respectively. The action space is 4-dimensional, including delta values in x , y , and z directions as well as the manipulation of the gripper, and its observation space combines the visual space of the three camera inputs and the linear space of the gripper’s low-dimensional state.

We choose two RL Bench tasks in experiments, depicted in Fig. 3, *Push buttons* and *Take umbrella out of umbrella stand*. Positions of buttons or umbrellas are randomly initialized for each episode. In all experiments, we set the maximum episode length as 100, the maximum training steps as $2e5$, the reward scale for the task’s goal state to 100, and the frequency of querying the LLM as every 4 timesteps. We choose the middle-size version of Dreamerv3 and keep other hyperparameters as defaults in [41]. Each experiment was repeated three times with different random seeds.

Table 2: Success rates when evaluating different methods on two RL Bench tasks for 100 episodes.

Task Success Rates	Lafite-RL	Baseline
Push buttons	70.3%	41.6%
Take umbrella out of umbrella stand	68.7%	26.1%

In order to examine the influence of the interaction of the LLM, we report the learning curves regarding rewards, success rates, and episode length during the training process in Fig. 4. Tab. 2 shows the success rate of each RL policy over 100 evaluation episodes after learning.

We observe that, for the *Push buttons* task, success rates performed by Lafite-RL witness a jump at 20k steps and then continue to rise, while the rewards drop at the beginning, nadir at around 12k steps, and then continue to rise. In the case of the “take umbrella” task, the RL agent initially faces challenges in its learning process, mainly due to the complexity of the task. Later Lafite-RL agent has a substantial increase compared to the baseline, indicating that, especially for a long-horizon task that requires a larger number of steps to accomplish (cf. Fig. 4 right column), the guidance of an LLM proves to be particularly advantageous. In contrast, the Baseline did not perform well in two tasks in terms of slow-rising speed and low success rates and rewards. By the end of the training, Lafite-RL is able to achieve episode to be an average of 50 in both tasks, while the Baseline can only get around 80, which demonstrates the effectiveness of the control policy learned by Lafite-RL.

During the evaluations after training, as shown in Table 2, Lafite-RL shows a much higher success rate on 100 episodes of evaluations, which showcases that after training with LLMs’ continual real-time feedback, RL agents are able to perform the task well independently.

5 Discussion

Our experiments have shown that an LLM can successfully provide interactive feedback, similarly as a human can do it, to significantly enhance RL. There are other strategies to accelerate RL, however, each incurs costs or relies on preconditions to be satisfied, for example, RLHF requires a human to be continually present; imitation learning requires trajectories being performed veridically beforehand; reward shaping requires suitable knowledge of the solution strategy. When these techniques are unavailable, or as a complementary technique, Lafite-RL offers a new alternative. However, it incurs its own costs, such as the need to prompt an LLM and frequent time-consuming interactions between the agent and the LLM. In our implementation, the agent requests LLM feedback at fixed-time steps to trade off costly feedback vs. cheaper feedback-less phases. More sophisticated strategies could let the agent dynamically request feedback only when it is expected to be particularly useful [43]. At the same time, to use the LLM more efficiently, in addition to evaluative feedback as in this study, more informative feedback from LLMs like action instructions will be needed to guide the RL process.

In the current Lafite-RL framework, a motion description is made based on the coordinates of objects in the environment, which might not be accessible in some environments, and other information like the gripper’s angles or velocities is not involved. More spatial-related tools can be designed to serve as APIs to let LLMs better understand the scene in complex task environments. Currently, several vision-language foundation models are being developed [44, 45]. Their use can augment the LLM’s understanding of the environment with more useful information, so as to provide better judgment and instructions. These visual capabilities will aid in the transfer from simulation to the real world.

From an experimental perspective, it is desirable to conduct more experiments with a larger variety of tasks, as well as evaluate the model with respect to generalization to model-free RL algorithms. Towards deploying the model, studies with non-expert users as well as sim-to-real or training in the real world would also be further steps.

6 Conclusion

In this work, we introduce Lafite-RL, a framework aimed at enabling non-expert users to design prompts that instruct LLMs to give timely feedback during RL agents’ training in robotic tasks, so as to accelerate learning while freeing up human effort. Experimental results show that with easily designed prompts, an LLM’s feedback is effective in speeding up reinforcement learning interactively. Since LLMs are known to be skilled in high-level actions, our results in the low-level robotic manipulation domain demonstrate the potential of LLMs, where we have not used them directly for control but to guide an agent’s learning. In summary, our approach shows substantial potential for utilizing LLMs for interactively enhancing reinforcement learning.

Acknowledgments

We gratefully acknowledge support from the China Scholarship Council (CSC) and the German Research Foundation (DFG) under the project Crossmodal Learning (TRR 169).

References

- [1] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017.
- [2] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, et al. Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning (CoRL)*, pages 651–673. PMLR, 2018.
- [3] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine. How to train your robot with deep reinforcement learning: lessons we have learned. *The International Journal of Robotics Research*, 40(4-5):698–721, 2021.
- [4] A. Y. Ng, S. Russell, et al. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, volume 1, page 2, 2000.
- [5] S. Arora and P. Doshi. A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence*, 297:103500, 2021.
- [6] W. B. Knox and P. Stone. Interactively shaping agents via human reinforcement: The TAMER framework. In *Proceedings of the Fifth International Conference on Knowledge Capture*, pages 9–16, 2009.
- [7] A. Wilson, A. Fern, and P. Tadepalli. A bayesian approach for policy learning from trajectory preference queries. *Advances in Neural Information Processing Systems*, 25, 2012.
- [8] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems*, 30, 2017.
- [9] G. Warnell, N. Waytowich, V. Lawhern, and P. Stone. Deep TAMER: Interactive agent shaping in high-dimensional state spaces. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [10] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.
- [11] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- [12] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, et al. Do as I can, not as I say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [13] S. Vemprala, R. Bonatti, A. Bucker, and A. Kapoor. ChatGPT for robotics: Design principles and model abilities. *Microsoft Auton. Syst. Robot. Res*, 2:20, 2023.
- [14] A. Z. Ren, B. Govil, T.-Y. Yang, K. R. Narasimhan, and A. Majumdar. Leveraging language for accelerated learning of tool manipulation. In *Conference on Robot Learning (CoRL)*, pages 1531–1541. PMLR, 2023.
- [15] X. Zhao, M. Li, C. Weber, M. B. Hafez, and S. Wermter. Chat with the environment: Interactive multimodal perception using large language models. *arXiv preprint arXiv:2303.08268*, 2023.
- [16] R. Zhang, F. Torabi, G. Warnell, and P. Stone. Recent advances in leveraging human guidance for sequential decision-making tasks. *Autonomous Agents and Multi-Agent Systems*, 35(2):31, 2021.

- [17] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey, et al. Maximum entropy inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.
- [18] M. Vecerik, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. Heess, T. Rothörl, T. Lampe, and M. Riedmiller. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817*, 2017.
- [19] T. Hester, M. Vecerik, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, D. Horgan, J. Quan, A. Sendonaris, I. Osband, et al. Deep q-learning from demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [20] J. MacGlashan, M. K. Ho, R. Loftin, B. Peng, G. Wang, D. L. Roberts, M. E. Taylor, and M. L. Littman. Interactive learning from policy-dependent human feedback. In *International Conference on Machine Learning*, pages 2285–2294. PMLR, 2017.
- [21] K. Lee, L. M. Smith, and P. Abbeel. Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pages 6152–6163. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/lee21i.html>.
- [22] X. Wang, K. Lee, K. Hakhamaneshi, P. Abbeel, and M. Laskin. Skill preferences: Learning to extract and execute robotic skills from human feedback. In *Conference on Robot Learning (CoRL)*, pages 1259–1268. PMLR, 2022.
- [23] E. C. Williams, N. Gopalan, M. Rhee, and S. Tellex. Learning to parse natural language to grounded reward functions with weak supervision. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4430–4436, 2018. doi:10.1109/ICRA.2018.8460937.
- [24] P. Goyal, S. Niekum, and R. J. Mooney. Using natural language for reward shaping in reinforcement learning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 2385–2391. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi:10.24963/ijcai.2019/331. URL <https://doi.org/10.24963/ijcai.2019/331>.
- [25] P. M. Yanik, J. Manganelli, J. Merino, A. L. Threath, J. O. Brooks, K. E. Green, and I. D. Walker. A gesture learning interface for simulated robot path shaping with a human teacher. *IEEE Transactions on Human-Machine Systems*, 44(1):41–54, 2013.
- [26] R. Arakawa, S. Kobayashi, Y. Unno, Y. Tsuboi, and S.-i. Maeda. DQN-TAMER: human-in-the-loop reinforcement learning with intractable feedback. *arXiv preprint arXiv:1810.11748*, 2018.
- [27] F. Cruz, G. I. Parisi, J. Twiefel, and S. Wermter. Multi-modal integration of dynamic audiovisual patterns for an interactive reinforcement learning scenario. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 759–766. IEEE, 2016.
- [28] S. Trick, F. Herbert, C. A. Rothkopf, and D. Koert. Interactive reinforcement learning with bayesian fusion of multimodal advice. *IEEE Robotics and Automation Letters*, 7(3):7558–7565, 2022.
- [29] S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, and A. L. Thomaz. Policy shaping: Integrating human feedback with reinforcement learning. *Advances in Neural Information Processing Systems*, 26, 2013.
- [30] D. Arumugam, J. K. Lee, S. Saskin, and M. L. Littman. Deep reinforcement learning from policy-dependent human feedback. *arXiv preprint arXiv:1902.04257*, 2019.

- [31] F. Cruz, S. Magg, Y. Nagai, and S. Wermter. Improving interactive reinforcement learning: What makes a good teacher? *Connection Science*, 30(3):306–325, 2018.
- [32] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [33] J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. H. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus. Emergent abilities of large language models. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL <https://openreview.net/forum?id=yzkSU5zdwD>. Survey Certification.
- [34] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9493–9500, 2023.
- [35] W. Yu, N. Gileadi, C. Fu, S. Kirmani, K.-H. Lee, M. G. Arenas, H.-T. L. Chiang, T. Erez, L. Hasenclever, J. Humplik, et al. Language to rewards for robotic skill synthesis. *arXiv preprint arXiv:2306.08647*, 2023.
- [36] H. Lee, S. Phatale, H. Mansoor, K. Lu, T. Mesnard, C. Bishop, V. Carbune, and A. Rastogi. RLAIFF: Scaling reinforcement learning from human feedback with AI feedback. *arXiv preprint arXiv:2309.00267*, 2023.
- [37] M. Kwon, S. M. Xie, K. Bullard, and D. Sadigh. Reward design with language models. In *The 11th International Conference on Learning Representations*, 2023.
- [38] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, I. Stoica, and E. P. Xing. Vicuna: An open-source chatbot impressing GPT-4 with 90%* ChatGPT quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- [39] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [40] W. Kwon, Z. Li, S. Zhuang, Y. Sheng, L. Zheng, C. H. Yu, J. E. Gonzalez, H. Zhang, and I. Stoica. Efficient memory management for large language model serving with PagedAttention. *arXiv preprint arXiv:2309.06180*, 2023.
- [41] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- [42] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison. RL Bench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2):3019–3026, 2020.
- [43] L. Torrey and M. Taylor. Teaching on a budget: Agents advising agents in reinforcement learning. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1053–1060, 2013.
- [44] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [45] Z. Gan, L. Li, C. Li, L. Wang, Z. Liu, J. Gao, et al. Vision-language pre-training: Basics, recent advances, and future trends. *Foundations and Trends® in Computer Graphics and Vision*, 14(3–4):163–352, 2022.