

SPARQ: Selective Progress-Aware Resource Querying

Anujith Muraleedharan, Anamika J H

Abstract: Human feedback can greatly accelerate robot learning, but in real-world settings, such feedback is costly and limited. Existing human-in-the-loop reinforcement learning (HiL-RL) methods often assume abundant feedback, limiting their practicality for physical robot deployment. In this work, we introduce SPARQ, a progress-aware query policy that requests feedback only when learning stagnates or worsens, thereby reducing unnecessary oracle calls. We evaluate SPARQ on a simulated UR5 cube-picking task in PyBullet, comparing against three baselines: no feedback, random querying, and always querying. Our experiments show that SPARQ achieves near-perfect task success, matching the performance of always querying while consuming about half the feedback budget. It also provides more stable and efficient learning than random querying, and significantly improves over training without feedback. These findings suggest that selective, progress-based query strategies can make HiL-RL more efficient and scalable for robots operating under realistic human effort constraints.

1 Introduction

Robots deployed in the real world must adapt to diverse and dynamic environments while operating under safety and efficiency constraints. HiL-RL has emerged as a powerful paradigm to align robot behavior with human intent by incorporating interactive signals such as evaluative feedback [1], corrective actions [2, 3], and preference comparisons [4]. These approaches accelerate learning in tasks ranging from manipulation to navigation [5, 6]. However, their practicality is limited by the cost of human supervision: attention is a scarce resource, constrained by fatigue, multitasking demands, and operational limitations [6, 7]. Continuous querying for feedback overwhelms supervisors, while too few queries slow adaptation and degrade performance.

To address this tension, interactive strategies have been proposed where robots query humans only when needed. “Human-gated” approaches allow supervisors to intervene when they see fit [3, 8], but require continuous monitoring and cannot scale to multi-robot settings. “Robot-gated” methods shift the responsibility to the agent, enabling it to request feedback when encountering novel or risky states [9, 10]. While these strategies reduce unnecessary interactions, they typically lack explicit mechanisms for managing strict feedback budgets and often struggle in continuous-control tasks where queries must be carefully timed to avoid disrupting smooth execution.

Inspired by how humans naturally allocate their effort rationally, we propose SPARQ (Selective Progress-Aware Resource Querying), a budget-aware HiL-RL method. Rather than relying on continuous or uncertainty-driven feedback, SPARQ monitors task progress and selectively requests help only when learning stagnates or worsens. By explicitly modeling human attention as a limited budget and enforcing cooldowns between queries, SPARQ balances learning efficiency with supervision cost. An overview of our approach is illustrated in Fig. 1. Panel (a) shows the SPARQ-augmented training pipeline, where the agent selectively queries a human oracle only when needed. Panel (b) zooms into the SPARQ decision rule, which determines when queries are triggered based on progress, patience, and budget constraints.

Our contributions are as follows:

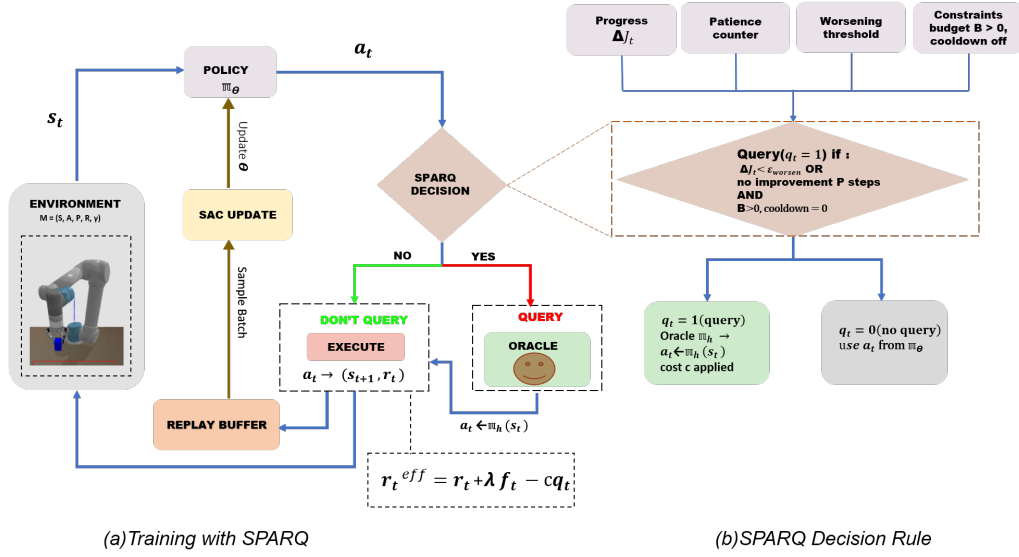


Figure 1: (a) **Training with SPARQ**: the policy π_θ samples a_t , the SPARQ gate decides whether to query the human oracle π_h ; the executed transition (s_t, a_t, s_{t+1}, r_t) is stored with the effective reward $r_t^{\text{eff}} = r_t + \lambda f_t - c q_t$ in replay D , and SAC updates θ . (b) **SPARQ decision rule (zoom-in)**: query when $(\Delta J_t < -\epsilon_{\text{worsen}})$ or (no improvement for P steps), and only if $B > 0$ and cooldown=0; otherwise use a_t from π_θ .

- We propose SPARQ, a resource-rational HiL-RL method that models human attention as a budgeted resource and allocates queries selectively.
- We introduce a progress-aware query rule that triggers feedback requests based on learning stagnation or deterioration, with patience and cooldown to prevent redundancy.
- We empirically evaluate SPARQ on a simulated UR5 cube-picking task, showing that it matches the success rate of full-feedback baselines while using about half the feedback budget, and performs comparably to Random on this task while providing explicit budget control and smoother temporal allocation of queries.

2 Related Work

Human-in-the-loop reinforcement learning (HiL-RL) has been studied extensively through frameworks that leverage explicit human feedback to accelerate agent learning. Early approaches such as TAMER[1] and its extension Deep TAMER[2] enabled agents to learn policies directly from human reward signals, but they required frequent intervention, making them impractical in resource-constrained scenarios. PEBBLE [11] improved upon this by using preference-based feedback to reduce cognitive burden, yet it assumes relatively abundant annotations. Similarly, COACH [8] modeled human feedback as policy-dependent, but lacked scalability to complex, high-dimensional tasks. These methods underscore the benefits of human guidance but fail to explicitly account for query efficiency, often leading to redundant human effort.

On the other hand, imitation learning has explored methods to minimize human supervision while ensuring robust policy transfer. DAgger [12] introduced iterative aggregation of expert demonstrations to mitigate compounding errors, but at the cost of requiring frequent expert corrections. HG-DAgger [13] extended this paradigm by letting a human supervisor decide when to provide demonstrations, reducing the labeling load compared to vanilla DAgger. However, this “human-gated” mechanism still relies on continuous monitoring, limiting scalability in practice. To address

this, ThriftyDagger [5] proposed querying experts only at uncertain states, but its heuristic query strategy lacked robustness in continuous-control settings. More recently, Diff-Dagger [14] incorporated diffusion-based uncertainty modeling to refine query selection, yet its reliance on dense expert demonstrations limits its practical applicability in settings with costly supervision.

Our work differs in two key respects. First, unlike TAMER- and PEBBLE-style methods that assume abundant human feedback, SPARQ explicitly models human attention as a limited resource and enforces a feedback budget. Second, unlike DAGger-style approaches that depend on dense expert demonstrations and uncertainty heuristics, SPARQ introduces a progress-aware query criterion: the agent requests feedback only when learning stagnates or deteriorates, combined with patience and cooldown to avoid redundancy. This budget-aware design enables SPARQ to match the performance of full-feedback baselines while cutting feedback usage nearly in half, offering a scalable pathway for HiL-RL in real-world robotics.

3 Methods

Our objective is to study query policies that enable efficient human-in-the-loop reinforcement learning (HiL-RL) under limited and costly feedback. We first formalize the problem setting as a constrained reinforcement learning task, and then describe four query strategies: three baselines and our proposed approach SPARQ (Selective Progress-Aware Resource Querying).

3.1 Problem Formulation

We model the robot learning task as a discrete-time Markov Decision Process (MDP), $M = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$ with continuous state space $s \in \mathcal{S}$, continuous action space $a \in \mathcal{A}$, unknown transition dynamics P , reward function R , and discount factor $\gamma \in [0, 1]$. The agent’s objective is to learn a policy $\pi_\theta : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes expected discounted return. Unlike standard RL, the reward function R is partially specified and can be augmented by human feedback. We assume access to a human oracle π_h that can provide corrective signals when queried [12, 13]. Each query incurs a cost $c > 0$, reflecting the limited time and attention of the human.

We define the agent’s effective reward at timestep t as:

$$r_t^{\text{eff}} = r_t + \lambda \cdot f_t - c \cdot q_t, \quad (1)$$

where r_t is the environment reward, f_t is the corrective feedback from the oracle π_h when queried, λ is a scaling factor, and $q_t \in \{0, 1\}$ indicates whether a query is made at time t . This formulation follows potential-based reward shaping principles, which preserve policy invariance [15].

The total expected return under policy π and query policy π_q is:

$$J(\pi, \pi_q) = \mathbb{E} \left[\sum_{t=0}^T \gamma^t (r_t + \lambda f_t - c q_t) \right]. \quad (2)$$

We further define ΔJ_t as a proxy for learning progress, e.g., the change in average episodic return or a task-specific metric such as distance-to-goal.

The goal is to jointly optimize the robot policy and query policy such that the agent learns efficiently while minimizing human effort:

$$(\pi^*, \pi_q^*) = \arg \max_{\pi, \pi_q} J(\pi, \pi_q), \quad (3)$$

subject to a feedback budget

$$\sum_{t=0}^T q_t \leq B, \quad (4)$$

where B represents the maximum number of queries available.

The query policy $\pi_q : \mathcal{S} \rightarrow \{0, 1\}$ thus decides when to request feedback, ideally only when learning progress stagnates or worsens.

3.2 Baselines

All query strategies are trained using Soft Actor-Critic (SAC) [16], chosen for its stability and sample efficiency in continuous control tasks. We compare three baselines against our proposed method.

No Oracle. The no-oracle baseline represents standard reinforcement learning without human intervention. The query policy is identically zero, $\pi_q(s) = 0 \quad \forall s \in \mathcal{S}$, and the effective reward reduces to $r_t^{\text{eff}} = r_t$. This baseline provides a lower bound on sample efficiency.

Random Querying. In this baseline, the agent queries the oracle with fixed probability p at each timestep: $\pi_q(s) \sim \text{Bernoulli}(p)$. This strategy ignores task structure and may waste queries on uninformative states. It serves as a mid-point between no supervision and full supervision.

Always Querying. Here, the agent queries the oracle at every timestep: $\pi_q(s) = 1 \quad \forall s \in \mathcal{S}$. This baseline achieves fast learning but immediately exhausts the query budget and imposes maximal human cost. It functions as an upper bound for performance under unlimited supervision.

3.3 SPARQ: Selective Progress-Aware Resource Querying

We propose a resource-rational query policy that adaptively balances performance improvement against feedback cost. The key idea is to trigger queries selectively, only when learning progress stagnates or worsens.

Progress-Aware Criterion. We also track learning progress using ΔJ_t based on episodic return or distance-to-goal. Queries are triggered when performance worsens beyond ϵ_{worsen} , or when no improvement is observed for P steps: $\Delta J_t < -\epsilon_{\text{worsen}} \vee \text{no improvement for } P \text{ steps}$.

Budget and Cooldown. To respect the constraint $\sum_{t=0}^T q_t \leq B$, we maintain a finite query budget and impose a cooldown period after each query to avoid redundancy. Algorithm 1 summarizes the overall procedure for our proposed method SPARQ.

Practical tuning. We tune three scalars once per domain: patience P (steps without improvement), worsening threshold ϵ_{worsen} (minimum negative ΔJ_t to trigger), and cooldown C (steps to defer

Algorithm 1 SPARQ: Selective Progress-Aware Resource Querying

Require: Initial policy parameters θ , query budget B , thresholds ϵ_{worsen} , P , cooldown C

Ensure: Learned policy π_θ

```

1: Initialize cooldown counter  $c \leftarrow 0$ 
2: for each episode do
3:   Reset environment, obtain  $s_0$ 
4:   for each timestep  $t = 0, \dots, T$  do
5:     Sample  $a_t \sim \pi_\theta(\cdot \mid s_t)$ 
6:     Estimate progress  $\Delta J_t$ 
7:     if  $c = 0$  and  $B > 0$  and  $(\Delta J_t < -\epsilon_{\text{worsen}}$  or no improvement for  $P$  steps) then
8:       Query oracle:  $a_t \leftarrow \pi_h(s_t)$ 
9:       Set  $q_t \leftarrow 1$ , update  $B \leftarrow B - 1$ 
10:      Reset cooldown counter:  $c \leftarrow C$ 
11:    else
12:      Execute  $a_t$ , set  $q_t \leftarrow 0$ 
13:      if  $c > 0$  then
14:        Decrement cooldown counter:  $c \leftarrow c - 1$ 
15:      end if
16:    end if
17:    Observe reward  $r_t$ , construct  $r_t^{\text{eff}} = r_t + \lambda f_t - cq_t$  ▷ Eq. 1
18:    Update  $\pi_\theta$  with SAC using  $r_t^{\text{eff}}$ 
19:    Transition to  $s_{t+1}$ 
20:   end for
21: end for

```

re-queries). A simple guideline is: choose P near $1\text{--}2\times$ the median episode length; set ϵ_{worsen} to the 5–10th percentile of negative $|\Delta J_t|$ magnitudes observed early in training; and pick C so the average query rate tracks the target feedback budget (e.g., adjust C until $\frac{1}{T} \sum_{t=1}^T q_t \approx b/100$ for a desired $b\%$). These rules-of-thumb yielded stable behavior without per-seed retuning in our setting.

4 Experimental Results

We evaluate SPARQ against three baseline query policies—No Oracle, Random Querying, and Always Querying, on a simulated UR5 cube-picking task in PyBullet. The evaluation addresses the following research questions:

- **RQ1:** Can progress-aware querying improve query efficiency under a limited budget?
- **RQ2:** Does SPARQ achieve comparable or better task success rates than baselines with fewer queries?
- **RQ3:** How do different query policies trade off task performance and human effort cost?

4.1 Task Setup

The evaluation task requires the UR5 arm to reach and grasp a cube placed at random positions on a tabletop workspace. The action space consists of continuous (x, y) end-effector target positions, while the observation space is the cube’s (x, y) location. A sparse success reward of +1 is given when the cube is successfully grasped, augmented with potential-based shaping based on distance-to-goal.

All query policies are trained using Soft Actor-Critic (SAC) for 50k timesteps. Hyperparameters such as the feedback budget B , query cost c , and patience P are fixed across all methods.

4.2 Query Efficiency (RQ1)

An effective query policy should minimize unnecessary oracle requests while still providing sufficient supervision to enable learning. We measure query efficiency as the average number of oracle requests per episode relative to the nominal budget.

As shown in Table 1 and Fig. 2b, SPARQ consumes $\sim 13\%$ of the budget, compared to 27% for Always Querying, effectively halving supervision demand. This demonstrates that progress-aware triggering avoids redundant queries that random or unconditional strategies incur. We allow occasional post-convergence queries due to replay shaping and drift checks; in deployment, a stop-when-converged guard (zero queries once success \geq a threshold over a patience window) would eliminate these residual queries.

Method	Success Rate	Budget %	Cost-Adj. Return	Queries / Success	Final Dist
always	1.000 ± 0.000	26.6%	198.4	1.33	0.001 ± 0.000
no_oracle	0.610 ± 0.490	0.0%	91.1	0.00	0.002 ± 0.000
random	1.000 ± 0.000	12.6%	196.6	0.63	0.003 ± 0.001
sparq	1.000 ± 0.000	13.2%	192.4	0.66	0.002 ± 0.001

Table 1: Overall comparison at 50k steps (mean \pm std unless noted). Cost-Adj. Return assumes query cost $c = 0.05$. Final Dist is reported as median \pm MAD end-effector→cube distance at grasp termination.


4.3 Task Performance (RQ2)

We next evaluate task success and learning dynamics. Figure 2a reports success rates and episodic returns across training. SPARQ attains a 100% final success rate, on par with Always Querying and substantially higher than No Oracle (61%). Unlike Random Querying, SPARQ achieves faster

convergence and more stable training, indicating that its progress-based criterion provides timely and informative feedback.

4.4 Trade-off Analysis (RQ3)

Finally, we examine the trade-off between task performance and human effort cost. Table 1 summarizes this balance by reporting success rate, budget usage, cost-adjusted return, query efficiency, and final grasp distance. SPARQ achieves near-optimal success rates while using only about half as many queries as Always Querying. Its cost-adjusted return is nearly identical to the unconstrained baseline, despite consuming substantially fewer human interactions. Moreover, SPARQ’s final grasp distance is comparable to or better than other baselines, indicating that reduced supervision does not compromise execution precision.



(a) Learning curves (success vs. timesteps).

(b) Queries per episode vs. timesteps.

Figure 2: Training dynamics on the UR5 cube-picking task. (a) SPARQ converges to 100% success with far fewer queries than Always. (b) SPARQ’s query rate remains low and stable over training.

4.5 Why does Random perform so well here?

This benchmark is intentionally simple: a low-dimensional state (x, y) for the cube, a shaped progress signal (distance-to-goal), and stable SAC training. Under these conditions, many query placements are similarly helpful, so a fixed-rate Random policy can sit close to the Pareto front. SPARQ’s advantage shows up in budget regularization and temporal smoothness (Fig. 2b): its query rate remains low and steady, avoiding bursts that can fatigue human operators, while achieving the same final success. We view this as evidence that progress-aware gating is a principled default when scaling to higher-dimensional observations and noisier progress proxies.

5 Conclusion

We introduced SPARQ, a resource-rational query framework for budget-aware human-in-the-loop reinforcement learning (HiL-RL). SPARQ monitors learning progress and selectively queries for feedback only when performance stagnates or worsens, ensuring that limited supervision is allocated where it is most impactful. In experiments on a simulated UR5 cube-picking task, SPARQ achieved 100% final success, matching the unconstrained Always Querying baseline while consuming about half the supervision budget. It also maintained comparable cost-adjusted returns and demonstrated precise grasp accuracy, underscoring that targeted feedback can deliver both efficiency and performance. While these results highlight SPARQ’s potential for scalable HiL-RL, our evaluation is limited to a single simulated domain. A multi-seed robustness study and deployment on real-world robots remain important future directions. Extending SPARQ to incorporate richer forms of human input (e.g., preferences, demonstrations) and to coordinate feedback across multiple robots are natural next steps toward practical, resource-rational HiL-RL.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback and the CoRL workshop organizers for providing a venue to share this work.

References

- [1] W. B. Knox and P. Stone. Interactively shaping agents via human reinforcement: The tamer framework. In *K-CAP*, 2008. URL <https://www.cs.utexas.edu/~bradknox/papers/icdl08-knox.pdf>.
- [2] G. Warnell, N. R. Waytowich, V. J. Lawhern, and P. Stone. Deep tamer: Interactive agent shaping in high-dimensional state spaces. In *AAAI*, 2018. URL <https://arxiv.org/abs/1709.10163>.
- [3] R. Pérez-Dattari, C. E. Celemin, J. Ruiz-del Solar, and J. Kober. Interactive learning with corrective feedback for policy shaping in deep reinforcement learning. *arXiv:1810.00466*, 2018. URL <https://arxiv.org/abs/1810.00466>.
- [4] P. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning from human preferences. In *NeurIPS*, 2017. URL <https://arxiv.org/abs/1706.03741>.
- [5] R. Hoque, J. Luo, B. Thananjeyan, A. Balakrishna, M. Tan, J. E. Gonzalez, and K. Goldberg. Thriftydagger: Budget-aware novelty and risk gating for interactive imitation learning. In *CoRL*, 2022. URL <https://proceedings.mlr.press/v164/hoque22a/hoque22a.pdf>.
- [6] D. R. Honeycutt, M. Nourani, and E. D. Ragan. Soliciting human-in-the-loop user feedback for interactive machine learning reduces user trust and impressions of model accuracy. *arXiv preprint arXiv:2008.12735*, 2020. URL <https://doi.org/10.48550/arXiv.2008.12735>.
- [7] J. Heo, J. Park, H. Jeong, K. J. Kim, J. Lee, E. Yang, and S. J. Hwang. Cost-effective interactive attention learning with neural attention processes. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pages 3758–3767, 2020. URL <https://proceedings.mlr.press/v119/heo20a/heo20a.pdf>.
- [8] C. E. Celemin, J. Ruiz-del Solar, and J. Kober. Coach: Learning continuous actions from human feedback. In *IROS*, 2019.
- [9] C. Wirth, R. Akrou, G. Neumann, and J. Fürnkranz. A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136):1–46, 2017. URL <https://jmlr.org/papers/v18/16-634.html>.
- [10] T. Kaufmann, P. Weng, V. Bengs, and E. Hüllermeier. A survey of reinforcement learning from human feedback. *arXiv preprint arXiv:2312.14925*, 2023. URL <https://arxiv.org/abs/2312.14925>.
- [11] Y. Lee, L. Smith, P. Abbeel, A. Srinivas, and M. Laskin. PEBBLE: Feedback-efficient interactive reinforcement learning via replay buffer relabeling. In *NeurIPS*, 2021. URL <https://arxiv.org/abs/2106.05091>.
- [12] S. Ross, G. Gordon, and J. A. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011.
- [13] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. In *Robotics: Science and Systems (RSS)*, 2019. URL <https://arxiv.org/abs/1810.02890>.

- [14] S.-W. Lee and Y.-L. Kuo. Diff-dagger: Uncertainty estimation with diffusion policy for robotic manipulation. *arXiv preprint arXiv:2410.14868*, 2024. URL <https://arxiv.org/abs/2410.14868>.
- [15] A. Y. Ng, D. Harada, and S. J. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, 1999.
- [16] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *ICML*, 2018.