
Residual Q-Learning: Offline and Online Policy Customization without Value

Chenran Li^{1*}, Chen Tang^{1*}, Haruki Nishimura², Jean Mercat²,
Masayoshi Tomizuka¹, Wei Zhan¹

¹ University of California Berkeley, ² Toyota Research Institute, USA
{chenran_li, chen_tang, tomizuka, wzhan}@berkeley.edu
{haruki.nishimura, jean.mercat}@tri.global

Abstract

Imitation Learning (IL) is a widely used framework for learning imitative behavior from demonstrations. It is especially appealing for solving complex real-world tasks where handcrafting reward function is difficult, or when the goal is to mimic human expert behavior. However, the learned imitative policy can only follow the behavior in the demonstration. When applying the imitative policy, we may need to *customize* the policy behavior to meet different requirements coming from diverse downstream tasks. Meanwhile, we still want the customized policy to maintain its imitative nature. To this end, we formulate a new problem setting called *policy customization*. It defines the learning task as training a policy that inherits the characteristics of the prior policy while satisfying some additional requirements imposed by a target downstream task. We propose a novel and principled approach to interpret and determine the trade-off between the two task objectives. Specifically, we formulate the customization problem as a Markov Decision Process (MDP) with a reward function that combines 1) the inherent reward of the demonstration; and 2) the add-on reward specified by the downstream task. We propose a novel framework, *Residual Q-learning (RQL)*, which can solve the formulated MDP by leveraging the prior policy *without knowing the inherent reward or value function of the prior policy*. We derive a family of residual Q-learning algorithms that can realize offline and online policy customization, and show that the proposed algorithms can effectively accomplish policy customization tasks in various environments. Demo videos and code are available on our website: <https://sites.google.com/view/residualq-learning>.

1 Introduction

Imitation learning (IL) is a widely used framework for learning imitative behavior from demonstrations. It is especially appealing for solving complicated real-world tasks where handcrafting reward function is difficult, or when the goal is to mimic expert behavior. For instance, IL has been widely studied for synthesizing human-like agents from large-scale human demonstrations on challenging real-world tasks, such as autonomous driving [4, 7, 2] and robot manipulation [59, 37, 39]. However, when applying the imitative policy in open real-world scenarios, we may encounter different circumstances that impose diverse requirements beyond merely following the behavior in the demonstration, such as reaching particular goals [45], enforcing safety [36], or complying with certain behavior preferences [61]. In this case, we need to *customize* the imitative policy to meet those task-specific objectives. Meanwhile, in many applications, we may still want the customized policy to possess some inherent characteristics of the expert demonstration (e.g., human-like).

*Equal Contribution

To this end, we formulate a new problem setting we refer to as *policy customization*. In the policy customization problem, we are given a pre-trained policy, and the task is to train a new policy that 1) inherits the properties of the prior policy; 2) satisfies some additional requirements imposed by a given downstream task. At first glance, this newly formulated problem may seem to be seamlessly solved by existing reinforcement learning (RL) fine-tuning algorithms [44, 40, 61, 51], where RL is applied to fine-tune a pre-trained policy obtained from behavior cloning (BC) or offline RL [31, 29, 34]. In particular, some RL fine-tuning methods combine different forms of IL objectives with RL to accelerate policy training, e.g., maximizing the policy likelihood on demonstration [44], regularizing the KL-divergence between the fine-tuned and prior policies [40, 61]. Different from the policy customization problem we consider, the IL objectives are introduced heuristically as regularization in RL fine-tuning. Solving policy customization instead requires a *principled* way to design the learning objective to jointly optimize the policy’s performance on the imitative and downstream tasks, which has not been well-studied in the RL fine-tuning literature.

In this work, we propose such a principled approach to interpret and design the policy customization objective with a solid theoretical basis. Formally, we formulate policy customization as a Markov decision process (MDP) whose reward function is a linear combination of 1) the inherent reward of the demonstration; and 2) an add-on reward specified by the downstream task. The reward weight directly determines the trade-off between the imitative and downstream tasks for policy optimization. However, we cannot solve this MDP with existing RL algorithms, since we can only access the imitative policy without knowing its underlying reward.

To tackle this challenge, we propose a novel *Residual Q-Learning (RQL)* framework that can solve the formulated MDP *without knowing the inherent reward or value function of the prior policy*. Specifically, we define a residual Q-function that, together with the log-likelihood of the prior policy, can conveniently construct the maximum-entropy policy solving the target MDP. We show that the residual Q-function can be iteratively learned with an update rule derived from soft Bellman equation [17]. Hence, the proposed RQL method can be applied on top of any value-based maximum-entropy RL algorithms. In particular, we derive residual soft Q-learning and residual soft actor-critic—two model-free policy customization algorithms through RL fine-tuning—which are adapted from soft Q-learning [17] and soft actor-critic [18] respectively. In addition, we introduce a model-based policy customization algorithm adapted from maximum-entropy Monte Carlo Tree Search (MCTS) [55], which enables *zero-shot online* policy customization when a dynamics model is available. In our experiments, we show that the proposed RQL algorithms can effectively customize the policies toward the additional task objective while maintaining their performance on the basic tasks for which the prior policies are trained.

2 Policy Customization with Residual Q-Learning

2.1 Problem Formulation: Policy Customization

We consider a MDP defined by the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, r, p)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, and $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function. Without loss of generality, we introduce the formulation under continuous state and action spaces. The state transition probability function $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, \infty)$ represents the probability density of the next state $s_{t+1} \in \mathcal{S}$ given the current state $s_t \in \mathcal{S}$ and action $a_t \in \mathcal{A}$. While the following formulation applies whether the reward function r is given or not, we are primarily interested in the case when r remains *unknown* and, instead, expert demonstrations are available to indicate desirable behaviors. We assume that we have access to a prior policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, \infty)$ pre-trained with imitation learning [15, 22] methods. In particular, we follow the common practice and model π as a maximum-entropy policy to account for sub-optimal demonstration [60], which is especially important for human demonstration. Specifically, the policy π maximizes the entropy-augmented cumulative reward that is inherent in the demonstrations. The resulting optimal policy follows a Boltzmann distribution [17] as follows:

$$\pi(a|s) = \frac{1}{Z_s} \exp\left(\frac{1}{\alpha} Q^*(s, a)\right), \quad (1)$$

where $Q^*(s, a)$ is the soft Q-function as defined in [17] and it satisfies the soft Bellman equation:

$$Q^*(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim p(\cdot|s, a)} \left[\alpha \log \int_{\mathcal{A}} \exp\left(\frac{1}{\alpha} Q^*(s', a')\right) da' \right], \quad (2)$$

and Z_s is the normalization factor defined as $\int_{\mathcal{A}} \exp\left(\frac{1}{\alpha} Q^*(s, \mathbf{a})\right) d\mathbf{a}$. The temperature coefficient α determines the relative importance of entropy and reward in the maximum-entropy policy [17]. It is straightforward to derive from Eqn. (1) that:

$$Q^*(s, \mathbf{a}) = \alpha \log \pi(\mathbf{a}|s) + \alpha \log Z_s. \quad (3)$$

Instead of directly deploying the policy π , we aim to *customize* it to satisfy some additional requirements for a given downstream task, specified by an add-on reward function $r_R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. Meanwhile, we want the customized policy to inherit the characteristics of the original policy. Formally, we aim to leverage the pre-trained prior policy π to train a new policy for the MDP defined by $\hat{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, \omega r + r_R, p)$. The new reward function is defined as a weighted sum of the basic reward r and the add-on reward r_R so that the customized policy meets both requirements. Under this setting, the main challenge lies in how to find the optimal policy for the new task without knowing the prior reward r but only the imitative policy π .

2.2 Residual Q-Learning

Since the reward r is unknown, we cannot directly train the new policy to maximize the total reward through RL. One plausible solution is to learn the inherent reward via inverse RL [12] and then combine the inferred reward with the add-on reward for RL training. However, practical inverse RL algorithms require jointly optimizing the policy and reward function, which could be challenging to yield good performance in certain environments when compared to imitation learning [16]. Instead, we aim to explore a solution without the need to explicitly infer the underlying reward to derive a more flexible policy customization framework. It turns out that we can leverage the imitative policy π to construct the new policy *without knowing π 's underlying reward or value function*. The proposed method can be applied on top of any *value-based maximum-entropy* RL algorithms [17, 18, 55], which rely on the soft Bellman update operator defined as:

$$\hat{Q}_{t+1}(s, \mathbf{a}) = r_R(s, \mathbf{a}) + \omega r(s, \mathbf{a}) + \gamma \mathbb{E}_{s' \sim p(\cdot|s, \mathbf{a})} \left[\hat{\alpha} \log \int_{\mathcal{A}} \exp\left(\frac{1}{\hat{\alpha}} \hat{Q}_t(s', \mathbf{a}')\right) d\mathbf{a}' \right], \quad (4)$$

where \hat{Q}_t is the estimated soft Q-function for $\hat{\mathcal{M}}$ at the t^{th} iteration. Instead of directly applying the soft Bellman update operator to iterate \hat{Q}_t during policy learning, we define a *residual* Q-function as $Q_{R,t} := \hat{Q}_t - \omega Q^*$ and attempt to iterate this residual Q-function during policy learning. We then derive an update rule for $Q_{R,t}$ that does not depend on the reward or value function of π :

$$\begin{aligned} & Q_{R,t+1}(s, \mathbf{a}) \\ &= r_R(s, \mathbf{a}) + \omega r(s, \mathbf{a}) + \gamma \mathbb{E}_{s'} \left[\hat{\alpha} \log \int_{\mathcal{A}} \exp\left(\frac{1}{\hat{\alpha}} \hat{Q}_t(s', \mathbf{a}')\right) d\mathbf{a}' \right] - \omega Q^*(s, \mathbf{a}), \end{aligned} \quad (5a)$$

$$\begin{aligned} &= r_R(s, \mathbf{a}) + \omega Q^*(s, \mathbf{a}) - \omega \gamma \mathbb{E}_{s'} \left[\alpha \log \int_{\mathcal{A}} \exp\left(\frac{1}{\alpha} Q^*(s', \mathbf{a}')\right) d\mathbf{a}' \right] \\ &+ \gamma \mathbb{E}_{s'} \left[\hat{\alpha} \log \int_{\mathcal{A}} \exp\left(\frac{1}{\hat{\alpha}} \hat{Q}_t(s', \mathbf{a}')\right) d\mathbf{a}' \right] - \omega Q^*(s, \mathbf{a}), \end{aligned} \quad (5b)$$

$$\begin{aligned} &= r_R(s, \mathbf{a}) - \omega \alpha \gamma \mathbb{E}_{s'} \log Z_{s'} \\ &+ \gamma \mathbb{E}_{s'} \left[\hat{\alpha} \log \int_{\mathcal{A}} \exp\left(\frac{1}{\hat{\alpha}} (Q_{R,t}(s', \mathbf{a}') + \omega \alpha \log \pi(\mathbf{a}'|s') + \omega \alpha \log Z_{s'})\right) d\mathbf{a}' \right], \end{aligned} \quad (5c)$$

$$\begin{aligned} &= r_R(s, \mathbf{a}) - \omega \alpha \gamma \mathbb{E}_{s'} \log Z_{s'} + \omega \alpha \gamma \mathbb{E}_{s'} \log Z_{s'} \\ &+ \gamma \mathbb{E}_{s'} \left[\hat{\alpha} \log \int_{\mathcal{A}} \exp\left(\frac{1}{\hat{\alpha}} (Q_{R,t}(s', \mathbf{a}') + \omega \alpha \log \pi(\mathbf{a}'|s'))\right) d\mathbf{a}' \right], \end{aligned} \quad (5d)$$

$$= r_R(s, \mathbf{a}) + \gamma \mathbb{E}_{s'} \left[\hat{\alpha} \log \int_{\mathcal{A}} \exp\left(\frac{1}{\hat{\alpha}} (Q_{R,t}(s', \mathbf{a}') + \omega \alpha \log \pi(\mathbf{a}'|s'))\right) d\mathbf{a}' \right], \quad (5e)$$

where Eqn. (5b) results from expressing the reward r with the soft Q-function Q^* based on Eqn. (2). Eqn. (5c) results from the definition of the residual Q-function and Eqn. (3). Please refer to Appendix A for a step-by-step breakdown of the derivation process. In each iteration, we can then define the policy corresponding to the current estimated \hat{Q}_t without computing \hat{Q}_t :

$$\hat{\pi}_t(\mathbf{a}|s) \propto \exp\left(\frac{1}{\hat{\alpha}} (Q_{R,t}(s, \mathbf{a}) + \omega \alpha \log \pi(\mathbf{a}|s))\right). \quad (6)$$

One remaining issue is that we do not have access to the underlying temperature coefficient α of the prior policy π . Note that α is multiplied upon the reward weight ω in Eqn. (5) and Eqn. (6). In practice, the reward weight is a hyperparameter to tune. Hence, we can define $\omega' = \omega\alpha$ and directly treat ω' as the hyperparameter without the need to know the true α .

3 Practical Algorithms

In this section, we present three practical algorithms for policy customization under different settings based on the update rule of residual Q-function in Eqn. (5) and the maximum-entropy policy defined in Eqn. (6). In these algorithms, the goal is to find a policy that solves the MDP \mathcal{M} given a prior policy π and the add-on reward function r_R . We first introduce two model-free algorithms to learn the residual Q-function and customize the policy through additional RL steps, which are adapted from soft Q-learning [17] and soft actor-critic [18] respectively. Then we introduce a model-based policy customization algorithm adapted from maximum-entropy MCTS [55], which enables *zero-shot online* policy customization when a dynamics model is available.

Residual Soft Q-Learning. We adapt the soft Q-learning [17] algorithm to a residual soft Q-learning algorithm for offline policy customization (i.e., policy customization via additional RL training steps) when the action space is discrete. The residual algorithm is essentially the same as the soft Q-learning. The only difference is that we use a function approximator parameterized by θ to model the residual Q-function instead of the soft Q-function and denote it as $Q_{R,\theta}(s, \mathbf{a})$. At each iteration, the parameters are updated by taking gradient steps to minimize the temporal-difference (TD) error:

$$J_{Q_R}(\theta) = \mathbb{E}_{(s_t, \mathbf{a}_t, r_{R,t}, s_{t+1}) \sim \mathcal{D}} \left[\frac{1}{2} \left(Q_{R,\bar{\theta}}^{\text{target}}(s_t, \mathbf{a}_t) - Q_{R,\theta}(s_t, \mathbf{a}_t) \right)^2 \right], \quad (7)$$

where \mathcal{D} is the replay buffer and the target residual Q-value $Q_{R,\bar{\theta}}^{\text{target}}(s_t, \mathbf{a}_t)$ is computed as:

$$Q_{R,\bar{\theta}}^{\text{target}}(s_t, \mathbf{a}_t) = r_{R,t} + \gamma \hat{\alpha} \log \sum_{\mathbf{a}' \in \mathcal{A}} \exp \left(\frac{1}{\hat{\alpha}} \left(Q_{R,\bar{\theta}}(s_{t+1}, \mathbf{a}') + \omega' \log \pi(\mathbf{a}' | s_{t+1}) \right) \right), \quad (8)$$

with $\bar{\theta}$ being the target parameters. The above objective function can be derived straightforwardly from the original objective function of soft Q-learning. We provide the detailed derivation in Appendix B for completeness. With a discrete action space, we can compute the policy distribution exactly given the residual Q-function and the prior policy as in Eqn. (6).

Residual Soft Actor-Critic. For tasks with continuous action spaces, we adapt the soft actor-critic [18] algorithm to a residual soft actor-critic algorithm. We aim to train a parameterized residual Q-function $Q_{R,\theta}(s, \mathbf{a})$ and a policy network $\hat{\pi}_\phi(\mathbf{a} | s)$ by alternating between policy evaluation and improvement. In the policy evaluation step, we update the residual Q-function for evaluating the current policy. The loss function is essentially the one in Eqn. (7) but with the target residual Q-value computed with respect to the current policy $\hat{\pi}_\phi$:

$$Q_{R,\bar{\theta}}^{\text{target}}(s_t, \mathbf{a}_t) = r_{R,t} + \gamma \mathbb{E}_{\mathbf{a}' \sim \hat{\pi}_\phi} \left[Q_{R,\bar{\theta}}(s_{t+1}, \mathbf{a}') + \omega' \log \pi(\mathbf{a}' | s_{t+1}) - \hat{\alpha} \log \hat{\pi}_\phi(\mathbf{a}' | s_{t+1}) \right] \quad (9)$$

The formula for the residual Q-value is derived from the modified Bellman update operator of soft actor-critic [18] following a similar procedure as in Eqn. (5). The detailed derivation is attached to Appendix C for completeness. In the policy improvement step, we improve the policy $\hat{\pi}_\phi$ leveraging the current estimated residual Q-value $Q_{R,\theta}(s, \mathbf{a})$ and the prior policy π by minimizing the expected KL-divergence between $\hat{\pi}_\phi$ and the desired maximum-entropy policy given the current $Q_{R,\theta}(s, \mathbf{a})$, which is equivalent to minimizing the loss function below:

$$J_{\hat{\pi}}(\phi) = \mathbb{E}_{s_t \sim \mathcal{D}} \left[\mathbb{E}_{\mathbf{a} \sim \hat{\pi}_\phi} \left[\hat{\alpha} \log \hat{\pi}_\phi(\mathbf{a} | s_t) - Q_{R,\theta}(s_t, \mathbf{a}) - \omega' \log \pi(\mathbf{a} | s_t) \right] \right]. \quad (10)$$

In our experiments, we implement the algorithm based on the practical soft actor-critic algorithm presented in [19], which incorporates auto entropy adjustment and double residual Q-functions [14].

Residual Maximum-Entropy MCTS. The last algorithm we introduce is a model-based policy customization algorithm adapted from the maximum-entropy MCTS [55] algorithm, which we refer to as residual maximum-entropy MCTS. Our primary interest is in applying this algorithm in the online setting with an offline learned dynamics model. MCTS essentially solves an online planning

problem leveraging the prior policy and add-on reward function, which enables *zero-shot online* policy customization. Such a paradigm is particularly useful for problem domains like autonomous driving, where large-scale human driving datasets [11, 58, 53] are available for learning good imitative policies [1, 46] as well as accurate traffic predictive models for online planning [10, 33]. Note that MCTS-based methods inherently assume a discrete action space. In practice, we may extend it to tasks with continuous action space by discretizing the action space [47, 57].

To solve the online planning problem, maximum-entropy MCTS constructs a look-ahead tree \mathcal{T} incrementally via Monte Carlo simulation. In the tree \mathcal{T} , each node $n(\mathbf{s}) \in \mathcal{T}$ corresponds to a state \mathbf{s} . For each action $\mathbf{a} \in \mathcal{A}$, the node $n(\mathbf{s})$ stores an estimated soft Q-value $\hat{Q}(\mathbf{s}, \mathbf{a})$ and a visit count $N(\mathbf{s}, \mathbf{a})$. At each iteration, simulations are conducted starting from the root node of the search tree, which consists of two stages: 1) querying a *tree policy* to select actions and traversing the tree until a leaf node is reached; 2) querying an *evaluation function* (e.g., Monte Carlo simulations leveraging a roll-out policy) at the leaf node to estimate the return-to-go. After simulation, the estimated returns are propagated backward to update all the nodes along the simulated paths. Afterward, the tree is grown by expanding the leaf nodes reached during the simulations. Different from a typical MCTS algorithm [8], maximum-entropy MCTS exploits maximum-entropy policy optimization. Hence, the tree policy used in maximum-entropy MCTS is defined as:

$$\hat{\pi}(\mathbf{a}|\mathbf{s}) = (1 - \lambda_s) \frac{\exp(\frac{1}{\alpha} \hat{Q}(\mathbf{s}, \mathbf{a}))}{\sum_{\mathbf{a}' \in \mathcal{A}} \exp(\frac{1}{\alpha} \hat{Q}(\mathbf{s}, \mathbf{a}'))} + \lambda_s \frac{1}{|\mathcal{A}|}, \quad (11)$$

where $\lambda_s = \epsilon^{|\mathcal{A}|/\log(\sum_{\mathbf{a}} N(\mathbf{s}, \mathbf{a})+1)}$ with ϵ as a hyperparameter. Given a simulated trajectory from the root node to a leaf node, i.e., $\{\mathbf{s}_0, \mathbf{a}_0, \dots, \mathbf{s}_T, \mathbf{a}_T\}$, the Q-values of all the visited nodes are updated with a *softmax backup*:

$$\hat{Q}(\mathbf{s}_t, \mathbf{a}_t) = \begin{cases} r_R(\mathbf{s}_t, \mathbf{a}_t) + \omega r(\mathbf{s}_t, \mathbf{a}_t) + \gamma \hat{R} & t = T - 1, \\ r_R(\mathbf{s}_t, \mathbf{a}_t) + \omega r(\mathbf{s}_t, \mathbf{a}_t) + \gamma \hat{\alpha} \log \sum_{\mathbf{a}' \in \mathcal{A}} \exp\left(\frac{1}{\alpha} \hat{Q}(\mathbf{s}_{t+1}, \mathbf{a}')\right) & t < T - 1, \end{cases} \quad (12)$$

where \hat{R} is the return-to-go estimated by querying an evaluation function on the leaf node $n(\mathbf{s}_T)$. To adapt the algorithm to residual maximum-entropy MCTS, we instead have each node store an estimated residual Q-value $Q_R(\mathbf{s}, \mathbf{a})$. The tree policy is revised as:

$$\hat{\pi}(\mathbf{a}|\mathbf{s}) = (1 - \lambda_s) \frac{\exp(\frac{1}{\alpha} (Q_R(\mathbf{s}, \mathbf{a}) + \omega' \log \pi(\mathbf{a}|\mathbf{s})))}{\sum_{\mathbf{a}' \in \mathcal{A}} \exp(\frac{1}{\alpha} (Q_R(\mathbf{s}, \mathbf{a}') + \omega' \log \pi(\mathbf{a}'|\mathbf{s})))} + \lambda_s \frac{1}{|\mathcal{A}|}. \quad (13)$$

During backward propagation, the residual Q-value is updated as follows:

$$Q_R(\mathbf{s}_t, \mathbf{a}_t) = \begin{cases} r_R(\mathbf{s}_t, \mathbf{a}_t) + \gamma R & t = T - 1, \\ r_R(\mathbf{s}_t, \mathbf{a}_t) + \gamma V_R(\mathbf{s}_{t+1}) & t < T - 1, \end{cases} \quad (14)$$

with $V_R(\mathbf{s}_{t+1}) = \hat{\alpha} \log \sum_{\mathbf{a}' \in \mathcal{A}} \exp\left(\frac{1}{\alpha} (Q_R(\mathbf{s}_{t+1}, \mathbf{a}') + \omega' \log \pi(\mathbf{a}'|\mathbf{s}_{t+1}))\right)$. In particular, we use Monte Carlo simulations to estimate the return-to-go R . The roll-out policy is set as the prior policy π , and R is computed as the cumulative return of the add-on reward r_R over the roll-out trajectories. While the estimated R is biased, it serves as a heuristic to guide the tree search to select nodes with future trajectories collecting high add-on rewards while staying close to the prior policy. The effect of this biased estimation will gradually be mitigated as the search tree expands to the terminal time step.

4 Related Work

Our residual Q-learning framework can be considered a principled approach to combine the objectives of IL and RL. Prior works have proposed different approaches for this purpose. In the IL literature, reward augmentation [9] has been introduced as a general framework to provide additional incentives from prior knowledge into the imitation learning process. Under the Generative Adversarial Imitation Learning (GAIL) framework, a surrogate reward can be conveniently incorporated to augment the learned discriminator for policy optimization [35, 3]. The augmented imitative reward is essentially the same as the total reward in our framework. The difference is that ours aims to further customize a trained imitative policy for downstream tasks, instead of regularizing IL with prior knowledge.

There are also a variety of works that directly design the objective function as a weighted mixture of IL and RL objectives. In some methods, like DQfD [21], a supervised classification loss is combined

with a TD loss to train the Q-function. It regularizes the learned Q-function to assign higher values to demonstrated actions. Other methods, such as DAPG [44] and BC-SAC [36], use a weighted mixture of behavior cloning loss and RL objectives in their loss function. SHAIL [25] formulates the imitation problem as a constrained RL problem with a constraint on the lower bound of policy likelihood on demonstrations. These approaches have shown that IL can help accelerate RL learning, especially under sparse rewards [21, 44], and vice versa [25]. As discussed in Sec. 1, similar objectives are also adopted in some RL fine-tuning algorithms [44, 40, 61, 51]. The pre-trained policy is not only used to initialize RL but also leveraged consistently over the fine-tuning stage as regularization. The introduction of RL has also been shown to effectively robustify behavior cloning [36] for challenging and safety-critical autonomous driving tasks. Conversely, IL regularization is used in offline RL approaches, such as TD3+BC [13] and CQL [31], to prevent overestimating the values on out-of-distribution actions and improve the robustness of offline RL policies. Different from these works, where the additional objective is treated as regularization, our framework formulates the learning task as a new MDP and provides a principled way to interpret and design the combined objective.

TS2C [56] is similar to ours from the perspective of having an imperfect expert—the prior policy in RQL can be considered as an imperfect expert for the overall task—and fusing knowledge extracted from this imperfect expert and RL exploration. The difference is that TS2C aims to train a policy towards a known reward; Thus, a value function can be trained to assert if the expert is reliable. In our case, since we rely on the prior policy to encode the basic task reward implicitly, we cannot use such a value function to govern the fusion between imitation and RL. One work in the literature that aligns with our motivation behind MCTS online customization is Deep Imitative Model [45], which aims to direct an imitative policy to arbitrary goals during online usage. It is formulated as a planning problem solved online to find the trajectory that maximizes its posterior likelihood conditioned on the goal. In contrast, we formulate the additional requirements as rewards instead of goals, which are more flexible to specify. Also, RQL is a unified framework for both offline and online customization.

In curriculum RL, the concept of learning the residuals of Q-functions has been adopted in boosted curriculum RL (BCRL) [28, 24]. Under the BCRL framework, a task curriculum is constructed by decomposing a complex task into a sequence of subtasks with increasing difficulty. For each task of the curriculum, its Q-function is modeled and learned as the sum of residuals, with each residual corresponding to a previous task of the curriculum, which was shown to outperform fitting a single function approximator [50, 28]. Similar to RL fine-tuning, BCRL focuses on expediting RL learning on a target task with the Q-functions learned on the previous tasks. The residual is learned to construct the Q-function for the target task with a *known* reward. Conversely, RQL aims at learning the residual to construct the Q-function for a new task whose reward function combines the *unknown* underlying reward of the expert and the *known* add-on reward.

5 Experiments

We evaluate the proposed algorithms on four environments selected from different domains: *CartPole* and *Continuous Mountain Car* environments from the OpenAI gym classic control suite [5], and *Highway and Parking* from the *highway-env* environments [32]. In our experiments, we implemented our algorithms upon Stable-Baselines3 [43] and its imitation library [16]. In Sec. 5.1, we provide the configurations of our experiments, including the settings of policy customization tasks in different environments, baselines, and evaluation metrics. In Sec. 5.2, we present and analyze the experimental results of RL offline policy customization. In Sec. 5.3, we demonstrate an example of zero-shot online customization leveraging the residual maximum-entropy MCTS algorithm. In Sec. 5.4, we investigate two representative RL fine-tuning methods under the context of policy customization and compare them with the proposed algorithms. Please refer to the appendices for detailed experiment configurations, implementation details, ablation studies, and additional experiments in MuJoCo [49].

5.1 Experiment Setup

Environments. The detailed configurations of the environments can be found in Appendix D. We mainly describe the policy customization task settings here. In each environment, we design the basic and add-on rewards to illustrate a practical application scenario, where the basic reward corresponds to some basic task requirements and the add-on reward reflects customized specifications such as behavior preference or additional task objectives.

- **CartPole balancing:** In the cartpole environment, the basic task is to balance the pole as long as possible. During policy customization, we demand an additional task that requires the cart to stay at the center of the rack. To validate if the additional objective is satisfied, we monitor the average absolute position error of the cart which is defined as $\bar{e}_x = \sum_{t=1}^T |x_t|/T$, where x_t is the coordinate of the cart with respect to the center of the rack.
- **Continuous mountain car:** In the mountain car environment, the basic task is to climb the mountain with the least energy consumption. During policy customization, we enforce an additional preference to avoid negative actions whenever possible. This corresponds to use cases where certain actions are less favorable. To monitor the change in action distribution, we count the number of negative actions executed during each episode, denoted as n_{neg} .
- **Highway navigation:** In the highway environment, the basic task is to navigate the vehicle safely and efficiently through the traffic. The vehicle can switch to arbitrary lanes to overtake the other vehicles. During policy customization, we enforce an additional preference to stay on the rightmost lane whenever possible. This reproduces driving habits that human drivers may have in real life. For example, drivers may prefer to drive on the rightmost lane to avoid missing the exits or when they want to drive the car slowly. We compute the average normalized lane index over an episode to validate if the desired behavior emerges, which is defined as $\bar{I}_{\text{lane}} = \sum_{t=1}^T I_{\text{lane},t}/2T$, where $I_{\text{lane},t} \in \{0, 1, 2\}$ is the lane index with the leftmost lane as the lane 0 and the rightmost lane as the lane 2.
- **Parking:** In the parking environment, the basic task is to park the car in the target parking slot. During policy customization, we add an additional requirement to avoid touching the boundaries of the parking slots during parking, which is considered good practice to avoid scratching other cars and to leave sufficient space for neighboring parking slots. We compute the non-violation rate to monitor if the requirement is met, which is defined as the percentage of episodes where the constraint is not violated, denoted as $\gamma_{\text{no-violation}}$.

Baselines. In our experiments, we mainly compare the performance of four categories of policies:

- **RL prior policy:** We train an RL policy optimizing the basic reward. It serves as the prior policy for policy customization. For environments with discrete action space (i.e., *CartPole* and *Highway*), we train the RL prior policy using soft Q-learning. For environments with continuous action space (i.e., *Continuous Mountain Car* and *Parking*), soft actor-critic is used. At test time, we also evaluate its performance on the overall task without customization, which serves as a baseline to show the effectiveness of the proposed algorithms.
- **IL prior policy:** We train another IL prior policy that imitates the RL prior with GAIL [22]. As discussed in Sec. 2, the proposed RQL algorithms are primarily motivated to customize imitative policy without knowledge of its inherent reward. The IL prior policy serves as a baseline for comparison with the residual-Q policy customized from the IL prior policy. In GAIL, we use soft Q-learning as the policy learning algorithm when the action space is discrete, and use soft actor-critic when the action space is continuous.
- **Residual-Q customized policies (Ours):** In each environment, we train two residual-Q customized policies leveraging the RL and IL prior policies respectively. They are mainly compared against the corresponding prior policies to validate the effectiveness of policy customization. We apply residual soft Q-learning for environments with discrete action space and apply the residual soft actor-critic for environments with continuous action space. Besides, we evaluate the zero-shot customization ability of the proposed residual maximum-entropy MCTS algorithm in the *Highway* environment.
- **RL with total reward:** We also train an RL policy with the total reward (i.e., $r + \omega r_R$). Note that we do not intend to treat it as a baseline for fair comparison under the policy customization setting. We aim to compare the residual-Q policies with it to validate that the customized policies indeed solve the overall task $\hat{\mathcal{M}}$.

Metric. We are mainly interested in validating the performance of the proposed policy customization methods from two perspectives: 1) whether the policy is customized toward the add-on reward; 2) whether the customized policy still possesses the characteristics of the prior policy, in other words, maintains the same level of performance on the basic task. Hence, we use the average basic reward and the success rate of completing the basic task as metrics to evaluate the policies' performance on

the basic task. Meanwhile, we use the average add-on reward and task-specific metrics introduced above as the metrics for evaluating the policies' performance with respect to the customized objective.

5.2 RL Offline Customization

The experimental results of RL offline customization are summarized in Table 1. In all the environments and with either RL or IL prior policies, the customized policies significantly outperform the prior policies on the add-on task objectives. Meanwhile, the customized policies maintain the same level of performance as the prior policies on the basic tasks. It validates that the proposed residual-Q algorithms can tailor the policies for customized requirements while maintaining their original characteristics. One interesting observation is that the policy customization improves the performance of the IL policy on the basic task in the *Parking* environment. One factor is that GAIL does not perform well in the first place. The *Parking* environment is indeed more challenging than the other environments given the continuous action space and, more importantly, the nature of the task—navigation in tight space under the non-holonomic constraint. Another aspect we think can explain this phenomenon is that the add-on reward *implicitly* guides the policy to improve its basic task performance. As shown in Appendix G, it is difficult for the imitative policy to stop the vehicle accurately at the target parking slot. Since GAIL merely learns to match the trajectory distributions under the expert and learned policies without the knowledge of the actual task objective, it is challenging for the imitative policy to precisely localize such a narrow parking space. In contrast, the add-on boundary violation constraint confines the vehicle trajectory within the parking space once the vehicle approaches the target parking slot. Hence, the policy is able to reach the goal point faster and receives higher returns even on the basic parking task.

In Table 1, we also compare the customized policies against the RL policies trained with the total reward function, named *RL Full Policy* in the table. In environments with discrete action spaces, the residual soft-Q learning policies behave very similarly to the RL full policies. Meanwhile, the residual soft actor-critic agents have different reward patterns compared to the RL full policies. We think the behavior difference is mainly due to the approximation errors in value and policy networks. The proposed residual-Q framework relies on the prior policy to acquire knowledge of the basic task. Thus, a suboptimal prior will affect the performance of the customized policy, especially on the basic task. It also explains the performance difference between the policies customized from RL and IL priors. Inspired by the observation mentioned above on IL policy customization, we think a practical solution to mitigate this issue is partially incorporating some basic task objectives into the add-on reward if available. It is reasonable for many practical use cases even with an imitative prior. While it is difficult to precisely describe the underlying reward function of human experts, we can comfortably assume that they follow some obvious commonsense objectives (e.g., collision avoidance). In this way, our method can be considered a more principled framework for the general synergy between IL and RL beyond policy customization where the reward functions are completely decomposed.

5.3 MCTS Online Customization

We choose the *Highway* environment to demonstrate the zero-shot online customization ability of residual maximum-entropy MCTS. We implement residual maximum-entropy MCTS based on the MCTS planner provided by the *highway-env* environment. The planner leverages the ground-truth dynamics model of the environment for roll-out simulation. The results are summarized in Table 1. MCTS online customization is able to achieve similar performance as the RL offline counterpart.

5.4 Ablation Study: Comparison with RL Fine-tuning

In our RQL framework, we formulate policy customization as an MDP whose reward function balances the utilities of the imitative and add-on tasks. It provides the theoretical foundation to determine the synergy between the IL and RL objectives in a principled manner. As discussed in Sec. 1 and Sec. 4, different approaches have been proposed in the RL fine-tuning literature to combine IL and RL objectives. We investigate two representative methods under the context of policy customization. We compare them against RQL and summarize the experimental results in this subsection. Please refer to Appendix F for the complete ablation study.

Greedy Reward Decomposition. One common approach is regularizing the divergence between the trained RL and prior policies during policy updates. It has been widely used in offline RL algorithms

to address the distributional shift issue [30, 54], and Advantage Weighted Actor-Critic (AWAC) [40] demonstrates the benefits of such policy regularization in stabilizing and accelerating RL fine-tuning. As shown in AWAC [40], when using KL-divergence to measure the distance between policies, the regularized optimal policy becomes a maximum-entropy policy with the advantage weighted by the prior policy. If we directly adapt it to solve the policy customization problem, the customized policy is essentially defined as the solution to the following optimization problem at each policy update step:

$$\tilde{\pi}_t = \arg \max_{\tilde{\pi} \in \Pi} \mathbb{E}_{\mathbf{a} \sim \tilde{\pi}(\cdot|\mathbf{s})} \left[\tilde{Q}_{R,t}(\mathbf{s}, \mathbf{a}) - \hat{\alpha} \log \tilde{\pi}(\mathbf{a}|\mathbf{s}) \right] - \lambda \mathcal{D}_{\text{KL}}(\tilde{\pi}(\cdot|\mathbf{s}) \parallel \pi(\cdot|\mathbf{s})), \quad (15)$$

where $\tilde{Q}_{R,t}$ is the soft Q-function of the MDP with only the add-on reward $r_R(\mathbf{s}, \mathbf{a})$, which is defined iteratively with the following update rule:

$$\tilde{Q}_{R,t+1}(\mathbf{s}, \mathbf{a}) = r_R(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_{\mathbf{s}' \sim p(\mathbf{s}'|\mathbf{a})} \left[\mathbb{E}_{\mathbf{a}' \sim \tilde{\pi}_t} \left[\tilde{Q}_{R,t}(\mathbf{s}', \mathbf{a}') - \hat{\alpha} \log \tilde{\pi}_t(\mathbf{a}'|\mathbf{s}') \right] \right]. \quad (16)$$

The closed-form solution to the optimization problem is given as:

$$\tilde{\pi}_t(\mathbf{a}|\mathbf{s}) \propto \exp \left(\frac{1}{\hat{\alpha} + \lambda} \left(\tilde{Q}_{R,t}(\mathbf{s}, \mathbf{a}) + \lambda \log \pi(\mathbf{a}|\mathbf{s}) \right) \right). \quad (17)$$

We can reduce our RQL policy to the same form by simply replacing the residual Q-function in Eqn. (6) with the soft Q-function $\tilde{Q}_{R,t}$. The resulting policy solves the optimization problem in Eqn. (15) but with an entropy weight of $\hat{\alpha} - \lambda$ in the objective function. If we consider it for the policy customization problem, it is essentially equivalent to heuristically estimating the optimal Q-function as the sum of the optimal Q-functions corresponding to the basic and add-on rewards. Hence, the synthesized policy is not the optimal solution to the target MDP but a greedy approximation [52, 26]. We compare RQL with this greedily customized policy on the *Continuous Mountain Car* and *Parking* environments. The results can be found in Appendix F.1. We found that this greedy customized policy performed much worse than RQL in the most challenging *Parking* environment. Compared to the prior policy, greedy customization compromises the policy’s success rate on the basic parking task to customize them towards satisfying the add-on objective. However, the greedy policy still has a significantly lower non-violation rate compared to the RQL policy. It validates that our framework is a more principled method to jointly optimize the IL and RL objectives for policy customization than divergence-based regularization.

Augmenting Reward with Policy KL-Divergence. Alternatively, some works [54, 61] have explored adding the policy divergence to the reward function. It has been shown to outperform policy regularization in offline RL setting [54]. To adopt this line of approaches for policy customization, we formulate the reward function of the target MDP as the sum of the add-on reward and a policy divergence regularization term. In particular, we follow [61] to augment the add-on reward with a penalty whose expectation regularizes the KL-divergence between the customized and prior policies. Formally, the reward function is defined as:

$$\check{r}_t(\mathbf{s}, \mathbf{a}) = r_R(\mathbf{s}, \mathbf{a}) - \beta \log \frac{\check{\pi}_t(\mathbf{a}|\mathbf{s})}{\pi(\mathbf{a}|\mathbf{s})}, \quad (18)$$

where $\check{\pi}_t$ denotes the customized policy after the t^{th} iteration of policy update, and β is a hyperparameter to balance the add-on reward and the policy regularization term. We can then apply standard RL algorithms to customize the policy toward this reward function. However, we found that it made the learning task much more difficult in our policy customization problem—as shown in Appendix F.2, soft actor-critic, for example, failed to find a customized policy that at least performed well on the basic parking task. We believe that one crucial factor contributing to the failure of divergence-based regularization is that the basic and add-on rewards are mostly *orthogonal* in our setting. The add-on reward only encodes additional task requirements; thus, we heavily rely on the prior policy to embed the desirable behavior on the basic task. Merely regularizing the KL-divergence is therefore insufficient—being close to the prior policy in terms of probabilistic distance does not necessarily imply behavior that accomplishes the basic task objective. Besides, incorporating the policy likelihood into the reward function makes it more challenging to update the Q-function in practice. For example, in Appendix F.3, we show that our RQL algorithm is theoretically equivalent to directly solving an MDP whose reward function is a weighted sum of the add-on reward and the log-likelihood of the prior policy. However, similar to the case of divergence-augmented reward, soft actor-critic failed to find a solution in *Parking* environment.

Table 1: Experimental Results of Residual-Q Policy Customization

Env.	Policy	Basic Task		Add-on Task	
		Succ. Rate	Basic Reward	\bar{e}_x	Add-on Reward
CartPole	RL Prior Policy	100%	427.21 \pm 2.25	0.15 \pm 0.11	-31.11 \pm 22.82
	RL Customized	100%	425.73 \pm 2.74	0.05 \pm 0.05	-11.18 \pm 10.12
	IL Prior Policy	100%	422.17 \pm 4.19	0.19 \pm 0.12	-38.57 \pm 24.26
	IL Customized	100%	423.37 \pm 4.28	0.07 \pm 0.06	-14.12 \pm 12.09
	RL Full Policy	100%	425.27 \pm 3.10	0.05 \pm 0.04	-10.72 \pm 7.43
Env.	Policy	Succ. Rate	Basic Reward	n_{neg}	Add-on Reward
Cont. Mt. Car	RL Prior Policy	100%	95.78 \pm 0.64	42.26 \pm 1.29	-4.23 \pm 0.13
	RL Customized	100%	95.61 \pm 0.43	37.90 \pm 0.75	-3.79 \pm 0.07
	IL Prior Policy	100%	92.84 \pm 2.35	64.40 \pm 21.03	-6.44 \pm 2.10
	IL Customized	100%	94.41 \pm 0.06	41.08 \pm 1.01	-4.11 \pm 0.10
	RL Full Policy	100%	95.14 \pm 0.21	27.86 \pm 3.87	-2.79 \pm 0.39
Env.	Policy	Succ. Rate	Basic Reward	\bar{I}_{lane}	Add-on Reward
Highway	RL Prior Policy	97.22%	44.67 \pm 2.43	0.46 \pm 0.28	9.17 \pm 5.54
	RL Customized	97.19%	42.23 \pm 1.55	0.97 \pm 0.05	19.42 \pm 0.99
	RL ME-MCTS	98.00%	42.71 \pm 1.65	0.96 \pm 0.07	19.19 \pm 1.38
	IL Prior Policy	96.04%	40.03 \pm 0.02	0.50 \pm 0.40	10.06 \pm 8.07
	IL Customized	93.00%	40.36 \pm 0.34	0.99 \pm 0.01	19.88 \pm 0.24
	IL ME-MCTS	97.00%	40.68 \pm 0.46	0.99 \pm 0.01	19.85 \pm 0.30
RL Full Policy	96.35%	41.98 \pm 1.53	0.97 \pm 0.01	19.49 \pm 1.24	
Env	Policy	Succ. Rate	Basic Reward	$\gamma_{\text{no-violation}}$	Add-on Reward
Parking	RL Prior Policy	99.09%	-7.08 \pm 2.65	57.61%	-1.55 \pm 3.71
	RL Customized	98.73%	-7.60 \pm 3.07	96.09%	-0.03 \pm 0.20
	IL Prior Policy	70.91%	-12.37 \pm 9.05	11.27%	-4.94 \pm 5.52
	IL Customized	83.59%	-7.93 \pm 4.84	66.07%	-0.32 \pm 0.73
	RL Full Policy	99.34%	-6.81 \pm 2.39	74.24%	-0.38 \pm 0.72

The evaluation results are computed over 4000 episodes for model-free policies and 100 episodes for maximum-entropy MCTS (ME-MCTS). The statistics with \pm are in the format of mean \pm std.

6 Limitations

As discussed in Sec. 5.2, one limitation of our current algorithms is that the performance of policy customization is *bottlenecked* by the imitative prior policy. Policy customization not only requires it to perform well on the trajectory distribution under the imitative prior policy itself but, more importantly, we need the imitative policy to accurately indicate desired behavior with respect to the basic task objective on the trajectory distribution under the *customized* policy. Hence, without careful design, our algorithms would suffer from the distributional shift issue. While, in practice, this issue can be mitigated by partially incorporating some basic task objectives into the add-on reward if available, a more principled and fundamental solution is to obtain a diverse imitative prior that is best suited for customization. Generative models, such as normalizing flows [48] and, more recently, diffusion models [23, 6, 20], have been shown promising in constructing such a diverse behavior prior. Besides, incorporating human experts in the loop has also been shown to be effective in improving the generalizability of imitation learning [27, 38, 41, 33]. We will investigate their applications in the policy customization problem in future work. Also, we only demonstrate MCTS online customization with an ideal perfect dynamics model in our current experiments. In future work, we will combine residual maximum-entropy MCTS with offline learned predictive models to achieve zero-shot online customization in more practical scenarios.

References

- [1] M. Bansal, A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. In *Proceedings of Robotics: Science and Systems*, Freiburg/Breisgau, Germany, June 2019.
- [2] R. Bhattacharyya, B. Wulfe, D. J. Phillips, A. Kuefler, J. Morton, R. Senanayake, and M. J. Kochenderfer. Modeling human driving behavior through generative adversarial imitation learning. *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [3] R. P. Bhattacharyya, D. J. Phillips, C. Liu, J. K. Gupta, K. Driggs-Campbell, and M. J. Kochenderfer. Simulating emergent properties of human driving behavior using multi-agent reward augmented imitation learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 789–795. IEEE, 2019.
- [4] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [5] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [6] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [7] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 4693–4700. IEEE, 2018.
- [8] R. Coulom. Efficient selectivity and backup operators in monte-carlo tree search. In *Computers and Games: 5th International Conference, CG 2006, Turin, Italy, May 29-31, 2006. Revised Papers 5*, pages 72–83. Springer, 2007.
- [9] P. Englert, N. A. Vien, and M. Toussaint. Inverse kkt: Learning cost functions of manipulation tasks from demonstrations. *The International Journal of Robotics Research*, 36(13-14):1474–1488, 2017.
- [10] J. L. V. Espinoza, A. Liniger, W. Schwarting, D. Rus, and L. Van Gool. Deep interactive motion prediction and planning: Playing games with motion prediction models. In *Learning for Dynamics and Control Conference*, pages 1006–1019. PMLR, 2022.
- [11] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9710–9719, 2021.
- [12] J. Fu, K. Luo, and S. Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- [13] S. Fujimoto and S. S. Gu. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems*, 34:20132–20145, 2021.
- [14] S. Fujimoto, H. Hoof, and D. Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pages 1587–1596. PMLR, 2018.
- [15] D. Garg, S. Chakraborty, C. Cundy, J. Song, and S. Ermon. Iq-learn: Inverse soft-q learning for imitation. In *Advances in Neural Information Processing Systems*, volume 34, pages 4028–4039, 2021.
- [16] A. Gleave, M. Taufeque, J. Rocamonde, E. Jenner, S. H. Wang, S. Toyer, M. Ernestus, N. Belrose, S. Emmons, and S. Russell. imitation: Clean imitation learning implementations. *arXiv preprint arXiv:2211.11972*, 2022.
- [17] T. Haarnoja, H. Tang, P. Abbeel, and S. Levine. Reinforcement learning with deep energy-based policies. In *International conference on machine learning*, pages 1352–1361. PMLR, 2017.
- [18] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.

- [19] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- [20] P. Hansen-Estruch, I. Kostrikov, M. Janner, J. G. Kuba, and S. Levine. Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573*, 2023.
- [21] 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.
- [22] J. Ho and S. Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, volume 29, 2016.
- [23] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.
- [24] S. Jauhri, J. Peters, and G. Chalvatzaki. Robot learning of mobile manipulation with reachability behavior priors. *IEEE Robotics and Automation Letters*, 7(3):8399–8406, 2022.
- [25] K. Judah, A. Fern, P. Tadepalli, and R. Goetschalckx. Imitation learning with demonstrations and shaping rewards. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, 2014.
- [26] Z. Juozapaitis, A. Koul, A. Fern, M. Erwig, and F. Doshi-Velez. Explainable reinforcement learning via reward decomposition. In *IJCAI/ECAI Workshop on explainable artificial intelligence*, 2019.
- [27] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8077–8083. IEEE, 2019.
- [28] P. Klink, C. D’Eramo, J. Peters, and J. Pajarinen. Boosted curriculum reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [29] I. Kostrikov, A. Nair, and S. Levine. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*, 2021.
- [30] A. Kumar, J. Fu, M. Soh, G. Tucker, and S. Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. *Advances in Neural Information Processing Systems*, 32, 2019.
- [31] A. Kumar, A. Zhou, G. Tucker, and S. Levine. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.
- [32] E. Leurent. An environment for autonomous driving decision-making. *GitHub repository*, 2018.
- [33] C. Li, T. Trinh, L. Wang, C. Liu, M. Tomizuka, and W. Zhan. Efficient game-theoretic planning with prediction heuristic for socially-compliant autonomous driving. *IEEE Robotics and Automation Letters*, 7(4):10248–10255, 2022.
- [34] J. Li, C. Tang, M. Tomizuka, and W. Zhan. Dealing with the unknown: Pessimistic offline reinforcement learning. In *Conference on Robot Learning*, pages 1455–1464. PMLR, 2022.
- [35] Y. Li, J. Song, and S. Ermon. Infogail: Interpretable imitation learning from visual demonstrations. *Advances in Neural Information Processing Systems*, 30, 2017.
- [36] Y. Lu, J. Fu, G. Tucker, X. Pan, E. Bronstein, B. Roelofs, B. Sapp, B. White, A. Faust, S. Whiteson, et al. Imitation is not enough: Robustifying imitation with reinforcement learning for challenging driving scenarios. *arXiv preprint arXiv:2212.11419*, 2022.
- [37] A. Mandlekar, D. Xu, R. Martín-Martín, S. Savarese, and L. Fei-Fei. Learning to generalize across long-horizon tasks from human demonstrations. *arXiv preprint arXiv:2003.06085*, 2020.
- [38] A. Mandlekar, D. Xu, R. Martín-Martín, Y. Zhu, L. Fei-Fei, and S. Savarese. Human-in-the-loop imitation learning using remote teleoperation. *arXiv preprint arXiv:2012.06733*, 2020.
- [39] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *Conference on Robot Learning*, pages 1678–1690. PMLR, 2022.
- [40] A. Nair, A. Gupta, M. Dalal, and S. Levine. Awac: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

- [41] Z. Peng, Q. Li, C. Liu, and B. Zhou. Safe driving via expert guided policy optimization. In *Conference on Robot Learning*, pages 1554–1563. PMLR, 2022.
- [42] A. Raffin. RL baselines3 zoo. *GitHub repository*, 2020.
- [43] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.
- [44] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017.
- [45] N. Rhinehart, R. McAllister, and S. Levine. Deep imitative models for flexible inference, planning, and control. In *International Conference on Learning Representations*, 2020.
- [46] O. Scheel, L. Bergamini, M. Wolczyk, B. Osiński, and P. Ondruska. Urban driver: Learning to drive from real-world demonstrations using policy gradients. In *Conference on Robot Learning*, pages 718–728. PMLR, 2022.
- [47] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- [48] A. Singh, H. Liu, G. Zhou, A. Yu, N. Rhinehart, and S. Levine. Parrot: Data-driven behavioral priors for reinforcement learning. *arXiv preprint arXiv:2011.10024*, 2020.
- [49] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
- [50] S. Tosatto, M. Pirotta, C. d’Eramo, and M. Restelli. Boosted fitted q-iteration. In *International Conference on Machine Learning*, pages 3434–3443. PMLR, 2017.
- [51] I. Uchendu, T. Xiao, Y. Lu, B. Zhu, M. Yan, J. Simon, M. Bennice, C. K. Fu, C. Ma, J. Jiao, S. Levine, and K. Hausman. Jump-start reinforcement learning. In *NeurIPS 2021 Robot Learning Workshop, RSS 2022 Scaling Robot Learning Workshop*, 2022.
- [52] H. Van Seijen, M. Fatemi, J. Romoff, R. Larocche, T. Barnes, and J. Tsang. Hybrid reward architecture for reinforcement learning. *Advances in Neural Information Processing Systems*, 30, 2017.
- [53] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, et al. Argoverse 2: Next generation datasets for self-driving perception and forecasting. *arXiv preprint arXiv:2301.00493*, 2023.
- [54] Y. Wu, G. Tucker, and O. Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- [55] C. Xiao, R. Huang, J. Mei, D. Schuurmans, and M. Müller. Maximum entropy monte-carlo planning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [56] Z. Xue, Z. Peng, Q. Li, Z. Liu, and B. Zhou. Guarded policy optimization with imperfect online demonstrations. In *The Eleventh International Conference on Learning Representations*, 2022.
- [57] W. Ye, S. Liu, T. Kurutach, P. Abbeel, and Y. Gao. Mastering atari games with limited data. *Advances in Neural Information Processing Systems*, 34:25476–25488, 2021.
- [58] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clause, M. Naumann, J. Kummerle, H. Konigshof, C. Stiller, A. de La Fortelle, et al. Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps. *arXiv preprint arXiv:1910.03088*, 2019.
- [59] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5628–5635. IEEE, 2018.
- [60] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.
- [61] D. M. Ziegler, N. Stiennon, J. Wu, T. B. Brown, A. Radford, D. Amodei, P. Christiano, and G. Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.