


Article

# Reinforcement Learning Your Way: Agent Characterization through Policy Regularization

Charl Maree <sup>1,2,\*</sup>  and Christian Omlin <sup>2,†</sup><sup>1</sup> Chief Technology Office, Sparebank 1 SR-Bank, 4007 Stavanger, Norway<sup>2</sup> Center for AI Research, University of Agder, 4879 Grimstad, Norway; christian.omlin@uia.no

\* Correspondence: charl.maree@uia.no

† Current address: Jon Lilletunsvai 9, 4879 Grimstad, Norway.

**Abstract:** The increased complexity of state-of-the-art reinforcement learning (RL) algorithms has resulted in an opacity that inhibits explainability and understanding. This has led to the development of several post hoc explainability methods that aim to extract information from learned policies, thus aiding explainability. These methods rely on empirical observations of the policy, and thus aim to generalize a characterization of agents' behaviour. In this study, we have instead developed a method to *imbue* agents' policies with a characteristic behaviour through regularization of their objective functions. Our method guides the agents' behaviour during learning, which results in an intrinsic characterization; it connects the learning process with model explanation. We provide a formal argument and empirical evidence for the viability of our method. In future work, we intend to employ it to develop agents that optimize individual financial customers' investment portfolios based on their spending personalities.

**Keywords:** explainable AI; multi-agent systems; deterministic policy gradients



**Citation:** Maree, C.; Omlin, C. Reinforcement Learning Your Way: Agent Characterization through Policy Regularization. *AI* **2022**, *3*, 250–259. <https://doi.org/10.3390/ai3020015>

Academic Editor: Gianni D'Angelo

Received: 20 January 2022

Accepted: 22 March 2022

Published: 24 March 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Recent advances in reinforcement learning (RL) have increased complexity which, especially for deep RL, has brought forth challenges related to explainability [1]. The opacity of state-of-the-art RL algorithms has led to model developers having a limited understanding of their agents' policies and no influence over learned strategies [2]. While concerns surrounding explainability have been noted for AI in general, it is only more recently that attempts have been made to explain RL systems [1,3–5]. These attempts have resulted in a wide suite of methods that typically rely on post hoc analysis of learned policies, which give only observational assurances of agents' behaviour. However, it is pivotal that future development of RL methods focus on more fundamental approaches towards inherently explainable RL [1]. We therefore propose an *intrinsic* method of guiding an agent's learning by controlling the objective function; there are two ways of manipulating the learning objective: modifying the reward function and regularizing the actions taken during learning. Whereas the reward function is specific to the particular problem, our ambition is to establish a generic method. We therefore propose a method which regularizes the objective function by minimizing the difference between the observed action distribution and a desired prior action distribution; we thus bias the actions that agents learn. While current methods for RL regularization aim to improve *training performance*—e.g., by maximizing the entropy of the action distribution [6], or by minimising the distance to a prior sub-optimal state-action distribution [7]—our aim is the *characterization* of our agents' behaviours. We extend single-agent regularization to accommodate multi-agent systems, which allows intrinsic characterization of individual agents. We provide a formal argument for the rationale of our method and demonstrate its efficacy in a toy problem where agents learn to navigate to a destination on a grid by performing, e.g., only right turns (under the premise that right turns are considered safer than left turns [8]). There are

several useful applications beyond this toy problem, such as asset management based on personal goals, intelligent agents with intrinsic virtues, and niche recommender systems based on customer preferences.

## 2. Background and Related Work

### 2.1. Agent Characterization

There have been several approaches to characterizing RL agents, with most—if not all—employing some form of post hoc evaluation technique. Some notable examples are:

*Probabilistic argumentation* [9] in which a human expert creates an ‘argumentation graph’ with a set of arguments and sub-arguments; sub-arguments attack or support main arguments which attack or support discrete actions. Sub-arguments are labelled as ‘ON’ or ‘OFF’ depending on the state observation for each time-step. Main arguments are labelled as ‘IN’, ‘OUT’, or ‘UNDECIDED’ in the following RL setting: *states* are the union of the argumentation graph and the learned policy, *actions* are the probabilistic ‘attitudes’ towards given arguments, and *rewards* are based on whether an argument attacks or supports an action. The learned ‘attitudes’ towards certain arguments are used to characterize agents’ behaviour.

*Structural causal modelling (SCM)* [10] learns causal relationships between states and actions through ‘action influence graphs’ that trace all possible paths from a given initial state to a set of terminal states, via all possible actions in each intermediate state. The learned policy then identifies a causal chain as the single path in the action influence graph that connects the initial state to the relevant terminal state. The explanation is the vector of rewards along the causal chain. Counter-explanations are a set of comparative reward vectors along chains originating from counter-actions in the initial state. Characterizations are made based on causal and counterfactual reasons for agents’ choice of action.

*Reward decomposition* [11,12] decomposes the reward into a vector of intelligible reward classes using expert knowledge. Agent characterization is done by evaluation of the reward vector for each action post training.

*Hierarchical reinforcement learning (HRL)* [13,14] divides agents’ tasks into sub-tasks to be learned by different agents. This simplifies the problem to be solved by each agent, making their behaviour easier to interpret, and thereby making them easier to characterize.

*Introspection (interesting elements)* [15] is a statistical post hoc analysis of the policy. It considers elements such as the frequency of visits to states, the estimated values of states and state-action pairs, state-transition probabilities, how much of the state space is visited, etc. Interesting statistical properties from this analysis are used to characterize the policy.

### 2.2. Multi-Agent Reinforcement Learning and Policy Regularization

We consider the multi-agent setting of partially observable Markov decision processes (POMDPs) [16]: for  $N$  agents, let  $\mathcal{S}$  be a set of states,  $\mathcal{A}_i$  a set of actions, and  $\mathcal{O}_i$  a set of incomplete state observations where  $i \in [1, \dots, N]$  and  $\mathcal{S} \mapsto \mathcal{O}_i$ . Agents select actions according to individual policies  $\pi_{\theta_i}(\mathcal{O}_i) \mapsto \mathcal{A}_i$  and receive rewards according to individual reward functions  $r_i(\mathcal{S}, \mathcal{A}_i) \mapsto \mathbb{R}$ , where  $\theta_i$  is the set of parameters governing agent  $i$ ’s policy. Finally, agents aim to maximize their total discounted rewards:

$$R_i(o, a) = \sum_{t=0}^T \gamma r_i(o_t, a_t)$$

where  $T$  is the time horizon and  $\gamma \in [0, 1]$  is a discount factor. For single-agent systems, the deep deterministic policy gradient algorithm (DDPG) defines the gradient of the objective  $J(\theta) = \mathbb{E}_{s \sim p^\mu} [R(s, a)]$  as [17]:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{s \sim \mathcal{D}} \left[ \nabla_{\theta} \mu_{\theta}(a|s) \nabla_a Q^{\mu_{\theta}}(s, a) \Big|_{a=\mu_{\theta}(s)} \right] \quad (1)$$

where  $p^\mu$  is the state distribution,  $\mathcal{D}$  is an experience replay buffer storing observed state transition tuples  $(s, a, r, s')$ , and  $Q^{\mu_\theta}(s, a)$  is a state-action value function where actions are selected according to a policy  $\mu_\theta(\mathcal{S}) \mapsto \mathcal{A}$ . In DDPG, the policy  $\mu$ —also called the *actor*—and the value function  $Q$ —also called the *critic*—are modelled by deep neural networks. Equation (1) is extended to a multi-agent setting; the multi-agent deep deterministic policy gradient algorithm (MADDPG) learns individual sets of parameters for each agent  $\theta_i$  [18]:

$$\nabla_{\theta_i} J(\theta_i) = \mathbb{E}_{o_i, a \sim \mathcal{D}} \left[ \nabla_{\theta_i} \mu_{\theta_i}(a_i | o_i) \nabla_{a_i} Q^{\mu_{\theta_i}}(o_i, a_1, \dots, a_N) \Big|_{a_i = \mu_{\theta_i}(o_i)} \right] \quad (2)$$

where  $o_i \in \mathcal{O}_i$  and the experience replay buffer  $\mathcal{D}$  contains tuples  $(o_i, a_i, r_i, o'_i)$ ,  $i \in [1, \dots, N]$ .

In this work, we further extend MADDPG by adding a regularization term to the actors' objective functions, thus encouraging them to mimic the behaviours specified by simple predefined prior policies. There have been several approaches to regularizing RL algorithms, mostly for the purpose of improved generalization or training performance. In [7], the authors defined an objective function with a regularization term related to the statistical difference between the current policy and a predefined prior:

$$J(\theta) = \mathbb{E}_{s, a \sim \mathcal{D}} [R(s, a) - \alpha D_{KL}(\pi_\theta(s, a) \| \pi_0(s, a))] \quad (3)$$

where  $\alpha$  is a hyperparameter scaling the relative contribution of the regularization term—the Kullback–Leibler (KL) divergence ( $D_{KL}$ )—and  $\pi_0$  is the prior policy which the agent attempts to mimic while maximising the reward. The KL divergence is a statistical measure of the difference between two probability distributions, formally:

$$D_{KL}(P \| Q) = \sum_{x \in X} P(x) \log \frac{P(x)}{Q(x)}$$

where  $P$  and  $Q$  are discrete probability distributions on the same probability space  $X$ . The stated objective of KL regularization is increased learning performance by penalising policies that stray too far from the prior. The KL divergence is often also called the *relative entropy*, with KL-regularized RL being the generalization of entropy-regularized RL ([19]); specifically if  $\pi_0$  is the uniform distribution, Equation 3 reduces to, up to a constant, the objective function for entropy-regularized RL as described in [6]:

$$J(\theta) = \mathbb{E}_{s, a \sim \mathcal{D}} [R(s, a) + \alpha H[\pi_\theta(s, a)]] \quad (4)$$

where  $H(\pi) = -P(\pi) \log(P(\pi))$  is the statistical entropy of the policy. The goal of entropy-regularized RL is to encourage exploration by maximising the policy's entropy and is used as standard in certain state-of-the-art RL algorithms, such as soft actor-critic (SAC) [6]. Other notable regularization methods include *control regularization* where, during learning, the action of the actor is weighted with an action from a sub-optimal prior:  $\mu_k = \frac{1}{1+\lambda} \mu_\theta + \frac{\lambda}{1+\lambda} \mu_{\text{prior}}$  and *temporal difference regularization*, which adds a penalty for large differences in the Q-values of successive states:  $J(\theta, \eta) = \mathbb{E}_{s, a, s' \sim \mathcal{D}} [R(s, a) - \eta \delta_Q(s, a, s')]$ , where  $\delta_Q(s, a, s') = [R(s, a) + \gamma Q(s', a') - Q(s, a)]^2$  [20,21].

While our algorithm is based on regularization of the objective function, it could be argued that it shares similar goals as those of algorithms based on *constrained RL*, namely the intrinsic manipulation of agents' policies towards given objectives. One example of constrained RL is [22], which finds a policy whose long-term measurements lie within a set of constraints by penalising the reward function with the Euclidean distance between the state and a given set of restrictions, e.g., an agent's location relative to obstacles on a map. Another example is [23], which penalises the value function with the accumulated cost of a series of actions, thus avoiding certain state-action situations. However, where constrained RL attempts to *avoid* certain conditions—usually through a penalty based on expert knowledge of the state—regularized RL aims to *promote* desired behaviours, such as choosing default actions during training or maximizing exploration by maximising

action entropy. The advantage of our system is that it does not require expert knowledge of the state-action space to construct constraints; our regularization term is independent of the state, which allows agents to learn simple behavioural patterns, thus improving the interpretability of their characterization.

### 3. Methodology

We regulate our agents based on a state-independent prior to maximize rewards while adhering to simple, predefined rules. In a toy problem, we demonstrate that agents learn to find a destination on a map by taking only right turns. Intuitively, we supply the probability distribution of three actions—left, straight, and right—as a regularization term in the objective function, meaning the agents aim to mimic this given probability distribution while maximising rewards. Such an agent can thus be characterized as an agent that prefers, e.g., right turns over left turns. As opposed to post hoc characterization, ours is an intrinsic method that inserts a desirable characteristic into an agent’s behaviour *during* learning.

#### Action Regularization

We modify the objective function in Equation (4) and replace the regularization term  $H[\pi_\theta(s, a)] = P(\pi) \log(P(\pi))$  with the mean squared error of the expected action and a specified prior  $\pi_0$ :

$$J(\theta) = \mathbb{E}_{s, a \sim \mathcal{D}}[R(s, a)] - \lambda L \quad (5)$$

$$L = \frac{1}{M} \sum_{j=0}^M \left[ \mathbb{E}_{a \sim \pi_\theta} [a_j] - (a_j | \pi_0(a)) \right]^2 \quad (6)$$

where  $\lambda$  is a hyperparameter that scales the relative contribution of the regularization term  $L$ ,  $a_j$  is the  $j^{\text{th}}$  action in a vector of  $M$  actions,  $\pi_\theta$  is the current policy, and  $\pi_0$  is the specified prior distribution of actions, which the agent aims to mimic while maximising the reward. Note that  $\pi_0(a)$  is *independent* of the state and  $(a_j | \pi_0(a))$  is therefore constant across all observations and time-steps. This is an important distinction from previous work, and results in a prior that is simpler to construct and a characterization that can be interpreted by non-experts; by removing the reference to the state space, we reduce the interpretation to the action space only, i.e., in this example the agent either proceeds straight, turns left, or turns right, independent of the locations of the agent and destination. Since this is a special case of Equation (4), it follows from the derivation given in [6].

We continue by extending our objective function to support a multi-agent setting. From Equation (5) and following the derivation in [18], we derive a multi-agent objective function with  $i \in [1, N]$  where  $N$  is the number of agents:

$$J(\theta_i) = \mathbb{E}_{o_i, a_i \sim \mathcal{D}}[R_i(o_i, a_i)] - \lambda \frac{1}{M_i} \sum_{j=0}^{M_i} \left[ \mathbb{E}_{a \sim \pi_{\theta_i}(o_i, a)} (a_j) - (a_j | \pi_{0_i}(a)) \right]^2 \quad (7)$$

Further, in accordance with the MADDPG algorithm, we model actions and rewards with actors and critics, respectively [18]:

$$\mathbb{E}_{a_i \sim \pi_{\theta_i}(o_i, a_i)} [a_i] = \mu_{\theta_i}(o_i) \quad (8)$$

$$\mathbb{E}_{o_i, a_i \sim \mathcal{D}} [R_i(o_i, a_i)] = Q_{\theta_i}(o_i, \mu_{\theta_1}(o_1), \dots, \mu_{\theta_N}(o_N)) \quad (9)$$

Through simple substitution of Equations (8) and (9) into Equation (7), we formulate our multi-agent regularized objective function:

$$J(\theta_i) = Q_{\theta_i}(o_i, \mu_{\theta_1}(o_1), \dots, \mu_{\theta_N}(o_N)) - \lambda L_i \quad (10)$$

$$L_i = \frac{1}{M_i} \sum_{j=0}^{M_i} \left[ \mu_{\theta_i}(o_i)_j - (a_j | \pi_{\theta_0_i}(a)) \right]^2 \quad (11)$$

Algorithm 1 optimizes the policies of multiple agents given individual regularization constraints  $\pi_{\theta_0_i}$ .

---

**Algorithm 1** Action-regularized MADDPG algorithm.
 

---

```

Set the number of agents  $N \in \mathbb{N}$ 
for  $i$  in  $1, N$  do ▷ For each agent
  Initialize actor network  $\mu_{\theta_{\mu,i}}$  with random parameters  $\theta_{\mu,i}$ 
  Initialize critic network  $Q_{\theta_{Q,i}}$  with random parameters  $\theta_{Q,i}$ 
  Initialize target actor network  $\mu'_{\theta_{\mu,i}}$  with parameters  $\theta_{\mu',i} \leftarrow \theta_{\mu,i}$ 
  Initialize target critic network  $Q'_{\theta_{Q,i}}$  with parameters  $\theta_{Q',i} \leftarrow \theta_{Q,i}$ 
  Set the desired prior action distribution  $\pi_{\theta_0_i}$ 
  Set the number of actions  $M_i \leftarrow |\pi_{\theta_0_i}|$ 
end for
Initialize replay buffer  $\mathcal{D}$ 
Set regularization weight  $\lambda$ 
for  $e = 1, \text{Episodes}$  do
  Initialise random function  $F(e) \sim N(0, \sigma_e)$  for exploration
  Reset environment and get the state observation  $s_1 \mapsto o_{[1..N]}$ 
   $t \leftarrow 1, \text{Done} \leftarrow \text{False}$ 
  while not Done do
    for  $i$  in  $1, N$  do ▷ For each agent
      Select action with exploration  $a_{i,t} \leftarrow \mu_{\theta_{\mu,i}}(o_i) + F(e)$ 
    end for
    Apply compounded action  $a_t$ 
    Retrieve rewards  $r_{[1..N],t}$  and observations  $s'_t \mapsto o'_{[1..N],t}$ 
    Store transition tuple  $\mathcal{T} = (o_t, a_t, r_t, o'_t)$  to replay buffer:  $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{T}$ 
     $t \leftarrow t + 1$ 
    if (end of episode) then
      Done  $\leftarrow \text{True}$ 
    end if
  end while
  Sample a random batch from the replay buffer  $\mathcal{B} \subset \mathcal{D}$ 
  for  $i$  in  $1, N$  do ▷ For each agent
     $\widehat{Q}_i \leftarrow r_{\mathcal{B},i} + \gamma Q'_i(o'_{\mathcal{B},i}, \mu'_1(o'_{\mathcal{B},1}), \dots, \mu'_N(o'_{\mathcal{B},N}))$ 
    Update critic parameters  $\theta_{Q,i}$  by minimising the loss:
    
$$\mathcal{L}(\theta_{Q,i}) = \frac{1}{|\mathcal{B}|} \sum_{\mathcal{B}} \left( Q_{\theta_{Q,i}}(o_{\mathcal{B}}, a_{\mathcal{B},1}, \dots, a_{\mathcal{B},N}) - \widehat{Q}_i \right)^2$$

    Update the actor parameters  $\theta_{\mu,i}$  by minimising the loss: ▷ From Equation (10)
    
$$\widehat{R}_i = \overline{Q_i(o_{\mathcal{B},i}, \mu_1(o_{\mathcal{B},1}), \dots, \mu_N(o_{\mathcal{B},N}))}$$

    
$$\mathcal{L}(\theta_{\mu,i}) = -\widehat{R}_i + \lambda \frac{1}{M_i} \sum_{j=1}^{M_i} \left[ \mu_i(o_{\mathcal{B},i})_j - (a_j | \pi_{\theta_0_i}) \right]^2$$

    Update target network parameters:
    
$$\theta_{\mu',i} \leftarrow \tau \theta_{\mu,i} + (1 - \tau) \theta_{\mu',i}$$

    
$$\theta_{Q',i} \leftarrow \tau \theta_{Q,i} + (1 - \tau) \theta_{Q',i}$$

  end for
end for

```

---

## 4. Experiments

### 4.1. Empirical Setup


We created a toy problem in which one or more agents navigate a  $6 \times 6$  grid through a set of three actions:  $\mathcal{A}_1 = \text{turn left}$ ,  $\mathcal{A}_2 = \text{go straight}$ , and  $\mathcal{A}_3 = \text{turn right}$ . Every new episode randomly placed a set of destinations in the grid  $D_i$ ,  $i \in [1, N]$ —one for each of  $N \geq 1$  agents—with initial agent locations  $L_{i,0} = (3, 0)$ . Rewards were the agents' Euclidean

distances from their destinations  $R_{i,t} = \|D_i - L_{i,t}\|_2$  where  $L_{i,t}$  is the location of agent  $i$  at time-step  $t$ . Finally, agents' observations were the two-dimensional distances to their destinations:  $\mathcal{O}_{i,t} = D_i - L_{i,t}$ . An episode was completed when either both agents had reached their destinations or a maximum of 50 time steps had passed.

We ran two sets of experiments, one for a single-agent setting and one for a dual-agent setting. We sized all networks in these two settings with two fully connected feed-forward layers; the single agent networks had 200 nodes in each layer, while the dual-agent networks had 700 nodes in each layer. Actor networks had a softmax activation layer, while the critic networks remained unactivated. Our training runs consisted of 3000 iterations and we tuned the hyperparameters using a simple one-at-a-time parameter sweep. We used training batches of 256 time-steps and sized the replay buffers to hold 2048 time-steps. In each iteration, we collected 256 time steps and ran two training epochs. We tuned the learning rates to 0.04 for the actors and 0.06 for the critics, the target network update parameters  $\tau$  to 0.06, and the discount factors  $\gamma$  to 0.95. We specified the regularization coefficient  $\lambda = 2$ , the regularization prior for the single-agent setting as  $\pi_0 = [P(\mathcal{A}_1), P(\mathcal{A}_2), P(\mathcal{A}_3)] = [0.0, 0.6, 0.4]$ , and the regularization priors for the dual-agent setting as  $\pi_{0,1} = [0.0, 0.6, 0.4]$  and  $\pi_{0,2} = [0.4, 0.6, 0.0]$ . This meant that the single agent was regularized to not take any left turns, while slightly favouring going straight above turning right. For the dual agents, agent 1 was to avoid left turns while agent 2 was to instead avoid right turns; we did this to demonstrate the characterization of the agents as preferring either left or right turns while navigating to their destinations. We conducted three experiments to explore the effects of the regularization prior, using the single-agent system for brevity, with the regularization priors  $\pi_0 \in \{[0.4, 0.2, 0.4], [0.33, 0.33, 0.33], [0.1, 0.5, 0.4]\}$ .

#### 4.2. Results

In the single agent setting of our toy problem, we used our algorithm to encourage an agent to prefer right turns over left turns; we used a regularization prior  $\pi_0 = [0.0, 0.6, 0.4]$  to regulate the probability of *left actions* to 0.0, *straight actions* to 0.6, and *right actions* to 0.4. Figure 1 shows three different trajectories that demonstrate such an agent's behaviour for destinations which lie either to the left, straight ahead, or to the right of the agent's starting location. As expected, the agent never turned left and always took the shortest route to its destination given its constraints.




**Figure 1.** Three trajectories of a single agent in the navigation problem. The starting locations are consistently (3,0), and the destinations are marked by red circles. During learning, the agent received a regularization prior  $\pi_0 = [0.0, 0.6, 0.4]$ , where the values in  $\pi_0$  correspond to the probabilities of the actions turn left, go straight, and turn right, respectively. While the agent in (a) makes a series of right turns to reach its destination on the left, the agent in (b) needs not turn, and the agent in (c) follows the shortest path involving a single right turn.

Figure 2 shows three additional experiments which illustrate an agent's behaviour given different regularization priors. In Figure 2a, we used the prior  $\pi_0 = [0.4, 0.2, 0.4]$ , which biased the agent towards taking turns rather than going straight. This agent consistently followed a zig-zag approach to the target, using the same number of steps compared




to a direct path with a single turn. This is an interesting observation, as an unregulated agent would typically take a direct path, as shown in Figure 2b. This agent was regulated with a uniform prior  $\pi_0 = [0.33, 0.33, 0.33]$ , which resulted in a similar strategy as that of an unregulated agent, but with the added benefit of increased exploration as discussed in [6]—entropy regularization uses the uniform distribution for  $\pi_0$ . In Figure 2c, we used the prior  $\pi_0 = [0.1, 0.5, 0.4]$  which assigns a low probability for taking left turns. In this experiment we specifically chose a destination to the immediate left of the starting location; other destinations allowed the agent to take the preferred right turns, whereas an immediate left turn to this shown destination proves that the agent does take this action in special cases. These behaviours are also observed in the multi-agent setting which, for brevity, we do not show here.



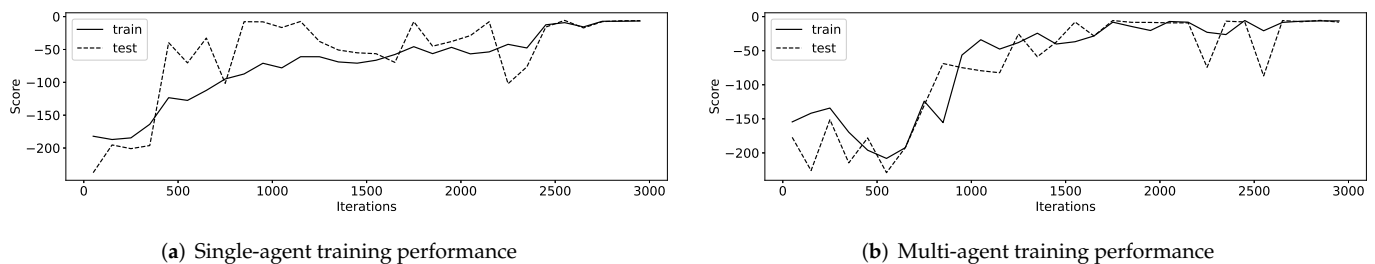
**Figure 2.** Three trajectories of agents trained with various regularization priors. In (a), the agent is biased towards taking turns and follows a zig-zag trajectory towards the destination ( $\pi_0 = [0.4, 0.2, 0.4]$ ). In (b), the agent's regularization prior is uniform—which equally favours all actions—and it follows the shortest path to the destination ( $\pi_0 = [0.33, 0.33, 0.33]$ ). In (c), the agent is allowed to take left turns with a low probability ( $\pi_0 = [0.1, 0.5, 0.4]$ ); we specifically chose the shown destination to encourage the agent to make a left turn.

Figure 3 shows the same grid navigation problem, but this time for a multi-agent setting. Here, we used two agents with different regularization terms to constrain their actions to (1) right turns only and (2) left turns only; the first agent's regularization prior  $\pi_{0,1} = [0.0, 0.6, 0.4]$  specified a probability of 0.0 for the left action, 0.6 for the straight action, and 0.4 for the right action, while the second agent's regularization prior  $\pi_{0,2} = [0.4, 0.6, 0.0]$  specified a probability of 0.4 for the left action, 0.6 for the straight action, and 0.0 for the right action. Clearly, the two agents have learned different strategies in the navigation problem. In Figure 3, it is clear that the two agents consistently took the shortest path to their respective destinations while adhering to their individual constraints. We therefore characterize them as agents that preferred to take right and left turns, respectively. Crucially, this is an *intrinsic* property of the agents imposed by the regularization of the objective function. This separates our method of intrinsic characterization from post hoc characterization techniques.

Finally, Figure 4 shows typical curves of training and testing returns for both the single-agent and multi-agent systems across 3000 training iterations. The agents clearly demonstrate a good learning response with steadily increasing returns both in training and testing and, while training performance is naturally slightly dependant on random initial conditions, there is no significant difference in convergence time between the single-agent and multi-agent systems.



**Figure 3.** Four sets of trajectories for a dual-agent environment in the navigation problem. The first agent—labelled ‘right turns’—received a regularization prior  $\pi_{0,1} = [0.0, 0.6, 0.4]$  while the second agent—labelled ‘left turns’—received a regularization prior  $\pi_{0,2} = [0.4, 0.6, 0.0]$ , where the values in  $\pi_{0,i}$  correspond to the probabilities of the actions turn left, go straight, and turn right. In (a) both agents’ destinations are on the left, but only the agent regularized to prefer left turns actually turns left while the other agent completes a series of right turns to reach its destination. In (b) both agents’ destinations are located such that they have to perform a series of turns according to their regularization priors. In (c) both agents’ destinations are located such that they perform their preferential turn—either to the left, or to the right—according to their regularization priors. In (d) one agent’s destination is straight ahead and it needs to not turn; the regularization prior clearly allows for such a strategy.



**Figure 4.** Training and testing returns for two typical training runs: the single-agent system in (a) and the multi-agent system in (b). In both cases, the learning processes clearly followed steady increases in returns and convergence happened roughly in the same number of iterations.

## 5. Conclusions and Direction for Future Work

Our objective was the *intrinsic* characterization of RL agents. To this end, we investigated and briefly summarized the relevant state-of-the-art in explainable RL and found that these methods have typically been relying on post-hoc evaluations of a learned policy. Policy regularization is a method that modifies a policy; however, it has typically been employed to enhance training performance which does not necessarily aid in policy characterization. We therefore adapted entropy regularization from maximizing the entropy in the policy to minimizing the mean squared difference between the expected action and a given prior. This encourages the agent to mimic a predefined behaviour while maximizing its reward during learning. Finally, we extended MADDPG with our regularization term. We provided a formal argument for the validity of our algorithm and empirically demonstrated its functioning in a toy problem. In this problem, we characterized two agents to follow different approaches when navigating to a destination in a grid; while one agent performed only right turns, the other performed only left turns. We conclude that our fundamentally sound algorithm was able to imbue our agents’ policies with specific characteristic behaviours. In future work, we intend to use this algorithm to develop a set of financial advisors that will optimize individual customers’ investment portfolios according to their individual spending personalities [24]. While maximising portfolio values, these



agents may prefer, e.g., property investments over crypto currencies, which are analogous to right turns and left turns in our toy problem.

**Author Contributions:** Conceptualization, C.M. and C.O.; methodology, C.M.; software, C.M.; validation, C.M.; formal analysis, C.M.; investigation, C.M.; resources, C.M.; data curation, C.M.; writing—original draft preparation, C.M.; writing—review and editing, C.O. and C.M.; visualization, C.M.; supervision, C.O.; project administration, C.M.; funding acquisition, C.M. and C.O.; All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was partially funded by The Norwegian Research Foundation, project number 311465.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** Not applicable.

**Conflicts of Interest:** The authors declare no conflict of interest. The funders had no role in the design of the study, in the collection, analyses, or interpretation of data, in the writing of the manuscript, or in the decision to publish the results.

## References

1. Heuillet, A.; Couthouis, F.; Díaz-Rodríguez, N. Explainability in deep reinforcement learning. *Knowl. Based Syst.* **2021**, *214*, 1–24. [[CrossRef](#)]
2. García, J.; Fernández, F. A Comprehensive Survey on Safe Reinforcement Learning. *J. Mach. Learn. Res.* **2015**, *16*, 1437–1480.
3. Barredo Arrieta, A.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; Garcia, S.; Gil-Lopez, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **2020**, *58*, 82–115. [[CrossRef](#)]
4. Wells, L.; Bednarz, T. Explainable AI and Reinforcement Learning: A Systematic Review of Current Approaches and Trends. *Front. Artif. Intell.* **2021**, *4*, 1–48. [[CrossRef](#)] [[PubMed](#)]
5. Gupta, S.; Singal, G.; Garg, D. Deep Reinforcement Learning Techniques in Diversified Domains: A Survey. *Arch. Comput. Methods Eng.* **2021**, *28*, 4715–4754. [[CrossRef](#)]
6. Harnoja, T.; Tang, H.; Abbeel, P.; Levine, S. Reinforcement Learning with Deep Energy-Based Policies. In Proceedings of the 34th International Conference on Machine Learning (ICML), Sydney, NSW, Australia, 6–11 August 2017.
7. Galashov, A.; Jayakumar, S.; Hasenclever, L.; Tirumala, D.; Schwarz, J.; Desjardins, G.; Czarnecki, W.M.; Teh, Y.W.; Pascanu, R.; Heess, N. Information asymmetry in KL-regularized RL. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
8. Lu, J.; Dissanayake, S.; Castillo, N.; Williams, K. *Safety Evaluation of Right Turns Followed by U-Turns as an Alternative to Direct Left Turns—Conflict Analysis*; Technical Report, CUTR Research Reports 213, University of South Florida, Scholar Commons, Tampa, FL, USA, 2001.
9. Riveret, R.; Gao, Y.; Governatori, G.; Rotolo, A.; Pitt, J.V.; Sartor, G. A probabilistic argumentation framework for reinforcement learning agents. *Auton. Agents Multi-Agent Syst.* **2019**, *33*, 216–274. [[CrossRef](#)]
10. Madumal, P.; Miller, T.; Sonenberg, L.; Vetere, F. Explainable Reinforcement Learning Through a Causal Lens. *arXiv* **2019**, arXiv:1905.10958v2.
11. van Seijen, H.; Fatemi, M.; Romoff, J.; Laroche, R.; Barnes, T.; Tsang, J. Hybrid Reward Architecture for Reinforcement Learning. *arXiv* **2017**, arXiv:1706.04208.
12. Juozapaitis, Z.; Koul, A.; Fern, A.; Erwig, M.; Doshi-Velez, F. Explainable Reinforcement Learning via Reward Decomposition. In Proceedings of the International Joint Conference on Artificial Intelligence. A Workshop on Explainable Artificial Intelligence, Macao, China, 10–16 August 2019.
13. Beyret, B.; Shafti, A.; Faisal, A. Dot-to-dot: Explainable hierarchical reinforcement learning for robotic manipulation. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Macau, China, 3–8 November 2019; pp. 5014–5019.
14. Marzari, L.; Pore, A.; Dall’Alba, D.; Aragon-Camarasa, G.; Farinelli, A.; Fiorini, P. Towards Hierarchical Task Decomposition using Deep Reinforcement Learning for Pick and Place Subtasks. *arXiv* **2021**, arXiv:2102.04022.
15. Sequeira, P.; Yeh, E.; Gervasio, M. Interestingness Elements for Explainable Reinforcement Learning through Introspection. *IUI Work.* **2019**, *2327*, 1–7.
16. Littman, M.L. Markov Games as a Framework for Multi-Agent Reinforcement Learning. In Proceedings of the Eleventh International Conference on Machine Learning, New Brunswick, NJ, USA, 10–13 July 1994; pp. 157–163.

17. Lillicrap, T.P.; Hunt, J.J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; Wierstra, D. Continuous control with deep reinforcement learning. In Proceedings of the 4th International Conference on Learning Representations (ICLR) (Poster), San Juan, Puerto Rico, 2–4 May 2016.
18. Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; Mordatch, I. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS, Long Beach, CA, USA, 4–9 December 2017.
19. Ziebart, B.D. Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy. Ph.D. Thesis, Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA, USA, 2010.
20. Cheng, R.; Verma, A.; Orosz, G.; Chaudhuri, S.; Yue, Y.; Burdick, J.W. Control Regularization for Reduced Variance Reinforcement Learning. *arXiv* **2019**, arXiv:1905.05380.
21. Parisi, S.; Tangkaratt, V.; Peters, J.; Khan, M.E. TD-regularized actor-critic methods. *Mach. Learn.* **2019**, *108*, 1467–1501. [[CrossRef](#)]
22. Miryoosefi, S.; Brantley, K.; Daume III, H.; Dudik, M.; Schapire, R.E. Reinforcement Learning with Convex Constraints. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019; Volume 32, pp. 1–10.
23. Chow, Y.; Ghavamzadeh, M.; Janson, L.; Pavone, M. Risk-Constrained Reinforcement Learning with Percentile Risk Criteria. *J. Mach. Learn. Res.* **2015**, *18*, 1–51.
24. Maree, C.; Omlin, C.W. Clustering in Recurrent Neural Networks for Micro-Segmentation using Spending Personality (In Print). In Proceedings of the 2021 IEEE Symposium Series on Computational Intelligence (SSCI), Orlando, FL, USA, 4–7 December 2021; pp. 1–5.