

Deep Ordinal Reinforcement Learning

Alexander Zap, Tobias Joppen, and Johannes Fürnkranz

TU Darmstadt, 64289 Darmstadt, Germany
alexander.zap@stud.tu-darmstadt.de
{tjoppen, juffi}@ke.tu-darmstadt.de

Abstract. Reinforcement learning usually makes use of numerical rewards, which have nice properties but also come with drawbacks and difficulties. Using rewards on an ordinal scale (ordinal rewards) is an alternative to numerical rewards that has received more attention in recent years. In this paper, a general approach to adapting reinforcement learning problems to the use of ordinal rewards is presented and motivated. We show how to convert common reinforcement learning algorithms to an ordinal variation by the example of Q-learning and introduce Ordinal Deep Q-Networks, which adapt deep reinforcement learning to ordinal rewards. Additionally, we run evaluations on problems provided by the OpenAI Gym framework, showing that our ordinal variants exhibit a performance that is comparable to the numerical variations for a number of problems. We also give first evidence that our ordinal variant is able to produce better results for problems with less engineered and simpler-to-design reward signals.

Keywords: Reinforcement Learning · Ordinal Rewards

1 Introduction

Conventional reinforcement learning (RL) algorithms rely on numerical feedback signals. Their main advantages include the ease of aggregation, efficient gradient computation, and many use cases where numerical reward signals come naturally, often representing a quantitative property. However, in some domains numerical rewards are hard to define and are often subject to certain problems. One issue of numerical feedback signals is the difficulty of *reward shaping*, which is the task of creating a reward function. Since RL algorithms use rewards as direct feedback to learn a behavior which optimizes the aggregation of received rewards, the reward function has a significant impact on the behavior that is learned by the algorithm. Manual creation of the reward function is often expensive, non-intuitive and difficult in certain domains and can therefore cause a bias in the optimal behavior that is learned by an algorithm. Since the learned behavior is sensitive to these reward values, the rewards of an environment should not be introduced or shaped arbitrarily if they are not explicitly known or naturally defined. This can also lead to another problem called *reward hacking*, where algorithms are able to exploit a reward function and miss the intended goal of the environment, caused by being able to receive better rewards through undesired behavior. The use of numerical

rewards furthermore requires *infinite rewards* to model undesired decisions in order to not allow trade-offs for a given state. This can be illustrated by an example in the medical domain, where it is undesirable to be able to compensate one occurrence of *death of patient* with multiple occurrences of *cured patient* to stay at a positive reward in average and therefore artificial feedback signals are used that can not be averaged. These issues have motivated the search for alternatives, such as preference-based feedback signals [13].

In this paper, we investigate the use of rewards on an ordinal scale, where we have information about the relative order of various rewards, but not about the magnitude of the quality differences between different rewards. Our goal is to extend reinforcement learning algorithms so that they can make use of ordinal rewards as an alternative feedback signal type in order to avoid and overcome the problems with numerical rewards.

Reinforcement learning with ordinal rewards has multiple advantages and directly addresses multiple issues of numerical rewards. Firstly, the problem of *reward shaping* is minimized, since the manual creation of the ordinal reward function specifically by the reward ordering often is intuitive and can be done easily without the need of exact specifications for reward values. Even though the creation of ordinal reward values, so-called reward tiers, through the ascending order of feedback signals introduces a naturally defined bias, it omits the largely introduced artificial bias by the manual shaping of reward values. At the same time, ordinal rewards simplify the problem of *reward hacking* because the omission of specific numeric reward values has the effect that any possible exploitation of rewards by an algorithm is only dependent on an incorrect reward order, which can be more easily fixed than the search for correct numerical values. While the use of *infinite rewards* can not be modelled directly, it is still possible to define infinite rewards as highest or lowest ordinal reward tier, and implement policies which completely avoid and encourage certain tiers.

Since the creation of the ordinal reward function is cheap and intuitive, it is especially suitable for newly defined environments since it enables the easy definition of ordinal rewards by ordering the possible outcomes naturally by desirability. Additionally it should be noted that for existing environments with numerical rewards it is possible to extract ordinal rewards from these environments.

The focus of this paper is the technique of using ordinal rewards for reinforcement learning. To this end, we propose an alternative reward aggregation for ordinal rewards, introduce a method for policy determination from ordinal rewards and compare the performance of ordinal reward algorithms to algorithms for numerical rewards. In Section 2, we discuss related work and previous approaches. A formal definition of common reinforcement learning terminology can be found in Section 3. Section 4 introduce reinforcement learning algorithms which use ordinal reward aggregations instead of numerical rewards, and illustrates the differences to conventional approaches. In Section 5 experiments are executed on the framework of OpenAI Gym and common reinforcement learning algorithms are compared to ordinal reinforcement learning.

2 Related Work

The technique of using rewards on an ordinal scale as an alternative to numerical rewards is mainly based on the approach of preference learning (PL) [1]. In contrast to traditional supervised learning, PL follows the core idea of having preferences over states or symbols as labels and predicting these preferences as the output on unseen data instances instead of labelling data with explicit nominal or numerical values.

Recently, there have been several proposals for combining PL with RL, where pairwise preferences over trajectories, states or actions are defined and applied as feedback signals in reinforcement learning algorithms instead of the commonly used numerical rewards. For a survey of such preference-based reinforcement learning algorithms, we refer the reader to [13].

While preference-based RL provides algorithms for learning an agent’s behavior from pairwise comparison of trajectories, [12] presents an approach for creating preferences over multiple trajectories in the order of ascending ordinal reward tiers, thereby deviating from the concept of pairwise comparisons over trajectories. Using a tutor as an oracle, this approach approximates a latent numerical reward score from a sequence of received ordinal feedback signals. This alternative reward computation functions as a reward transformation from the ordinal to the numerical scale and is applicable on top of an existing reinforcement learning algorithm.

Contrary to this approach, we do not use a tutor for the comparison of trajectories but can directly use ordinal rewards as a feedback signal. In order to use environments where numerical feedback already exists without the need for acquiring human feedback about the underlying preferences, we automatically extract rewards on an ordinal scale from existing environments with numerical rewards. To this end, we adapt an approach that has been proposed for Monte-Carlo Tree Search [4] to reinforcement learning.

Furthermore, we handle ordinal rewards in a similar manner as previous approaches by directly using aggregated received ordinal rewards for comparing different options. The idea of direct comparison of ordinal rewards builds on the works of [10], [11], [2] and [4], which provide criteria for the direct comparison of ordinal reward aggregations. We utilize the approach of [4], which transfers the numerical reward maximization problem into a best-choice maximization problem for an alternative computation of the value function for reinforcement learning from ordinal feedback signals. [4] used this idea for adapting Monte Carlo Tree Search to the use of ordinal rewards.

In summary, we automatically transfer numerical feedback into preference-based feedback and propose a new conceptual idea to utilize ordinal rewards for reinforcement learning, which should not be seen as an alternative for the existing algorithms stated above. Hence, we do not compare the performance of our new approach to any of the algorithms that use additional human feedback, but to common RL techniques that use numerical feedback.

3 Markov Decision Process and Reinforcement Learning

In this section, we briefly recapitulate Markov decision processes and reinforcement learning algorithms. Our notation and terminology is based on [8].

3.1 Value function and policy for Markov Decision Process

A *Markov Decision Process* (MDP) is defined as a tuple of (S, A, P, R) with S being a finite set of states, A being a finite set of actions, T being the transition function $S \times A \times S \rightarrow \mathbb{R}$ that models the probability of reaching a state s' when action a is performed in state s , and R being the reward function $S \times A \times S \rightarrow \mathbb{R}$ which maps a reward r from a subset of possible rewards $r \in \{r_1, \dots, r_n\} \subset \mathbb{R}$ to executing action a in state s and reaching s' in the process. For further work we assume that T is deterministic and a transition always has the probability of 0 or 1. Furthermore it is assumed that each action $a \in A$ is executable in any state $s \in S$, hence the transition function is defined for every element in $S \times A \times S$. A policy π is the specification which decision to take based on the environmental state. In a deterministic setting, it is modeled as a mapping $\pi : S \rightarrow A$ which directly maps an environmental state s to the decision a which should be taken in this state. The value function $V_\pi(s)$ represents the expected quality of a policy π in state s with respect to the rewards that will be received in the future. Value functions for numerical rewards are computed by the expectation of the discounted sum of rewards $\mathbb{E}[R]$. The value function $V_\pi(s)$ of a policy π in an environmental state s therefore can be computed by

$$V_\pi(s) = \mathbb{E}[R], \quad R = \sum_{t=0}^{\infty} \gamma^t r_t \quad (1)$$

where R is the discounted sum of rewards when following policy π , γ a discount factor, and r_t the direct reward at time step t . The optimal policy π^* in a state s is the policy with the largest $V_\pi(s)$, which complies with the goal of an RL algorithm to maximize expected future reward.

3.2 Reinforcement Learning

Reinforcement learning can be described as the task of learning a policy that maximizes the expected future numerical reward. The agent learns iteratively by updating its current policy π after every action and the corresponding received reward from the environment. Furthermore, the agent may perform multiple training sessions, so-called episodes, in the environment. Using the previously defined formalism, this can be expressed as approximating the optimal policy iteratively with a function $\hat{\pi}$, by repeatedly choosing actions that lead to states s with the highest estimated value function $V_{\hat{\pi}}(s)$. In the following section two common reinforcement learning algorithms are introduced.

Q-learning. The key idea of the Q-learning algorithm [9] is to estimate Q-values $Q(s, a)$, which estimate the expected future sum of rewards $\mathbb{E}[R]$ when choosing an action a in a state s and following the optimal policy π^* afterwards. Hence the Q-value can be seen as a measure of *goodness* for a state-action pair (s, a) , and therefore, in a given state s , the optimal policy π^* should select the action a that maximizes this value in comparison to other available actions in that state. The approximated Q-values are stored and iteratively updated in a Q-table. The Q-table is updated after an action a has been performed in a state s and the reward r and the newly reached state s' is observed. The computation of the expected Q-value is done by

$$\hat{Q}(s, a) = r(s, a) + \gamma \max_{a'} Q(s', a') \quad (2)$$

Following this so-called Bellman equation, every previously estimated Q-value is updated with the newly computed expected Q-value with the formula

$$Q(s, a) = Q(s, a) + \alpha[r(s, a) + \gamma \max_{a'} Q(s', a') - Q(s, a)] \quad (3)$$

where α represents a learning rate and γ the discount factor.

Deep Q-Network. The original Q-learning algorithm is limited to very simple problems, because of the explicitly stored Q-table, which essentially memorizes the quality of each possible state-action pair independently. Thus it requires, e.g., that each state-action pair has to be visited a certain number of times in order to make a reasonable prediction for this pair. A natural extension of this method is to replace the Q-table with a learned Q-function, which is able to predict a quality value for a given, possibly previously unseen state-action pair. The key idea behind the Deep Q-Network (DQN) [6, 7] is to learn a continuous function $Q^{DQN}(s)$ in the form of a deep neural network with m input nodes, which represent the feature vector of s , and n output nodes, each containing the Q-value of one action a .

Neural networks can be iteratively updated to fit the output nodes to the desired Q-values. The expected Q-value for a state-action pair is calculated in the same manner as defined in (2) with the difference that the Q-values are now predicted by the DQN, with one output node $Q_a^{DQN}(s)$ for each possible action a . Therefore (2) becomes

$$\hat{Q}_a^{DQN}(s) = r(s, a) + \gamma \max_{a'} Q_{a'}^{DQN}(s') \quad (4)$$

where $Q_a^{DQN}(s)$ represents the Q-value node of action a in state s .

In order to optimize the learning procedure, DQN makes use of several optimizations such as *experience replay*, the use of a separate *target and evaluation network*, and *Double Deep Q-Network*. More details on these techniques can be found in the following paragraphs.

Experience replay. Using a neural network to fit the Q-value of the previously executed state-action pair as described in (4) leads to overfitting to recent experiences because of the high correlation between environmental states across multiple successive time steps, and the property of neural networks to overfit recently seen training data. Instead of only using the previous state-action pair for fitting the DQN, experience replay [5] uses a memory M to store previous experience instances (s, a, r, s') and iteratively reuses a random sample of these experiences to update the network prediction at every time step.

Target and evaluation networks. Frequently updating the neural network, which is simultaneously used for the prediction of the expected Q-value, leads to unstable fitting of the network. Therefore these two tasks, firstly the prediction of the target Q-value for network fitting and secondly the prediction of the Q-value which is used for policy computation, allows for a split into two networks. These two networks are the *evaluation network*, which is used for policy computation, and the *target network*, which is used for predicting the target value for continuously fitting the evaluation network. In order to keep the target network up to date, it is replaced by a copy of the evaluation network every c steps.

Double Deep Q-Network. Deep Q-Networks tend to overestimate the prediction of Q-values for some actions, which may result in an unjustified bias towards certain actions. To address this problem, Double Deep Q-Networks [3] additionally use the target and evaluation networks to decouple the action choice and Q-value prediction by letting the evaluation network choose the next action to be played, and letting the target network predict the respective Q-value.

4 Deep Ordinal Reinforcement Learning

In this section, Markov decision processes and reinforcement learning algorithms are adapted to settings with ordinal reward signals. More concretely, we present a method for reward aggregation that fits ordinal rewards and explain how this method can be used in Q-learning and Deep Q-Networks in order to learn to solve environments that return feedback signals on an ordinal scale.

4.1 Ordinal Markov Decision Process

Similar to the standard Markov Decision Process, [10] defines an ordinal version of an MDP as a tuple of (S, A, T, R_o) with the only difference that R_o is the reward function $S \times A \times S \rightarrow \mathbb{N}$ is modified to return ordinal rewards instead of numerical ones. Thus, it maps executing action a in state s and reaching state s' to an ordinal reward r_o from a subset of possible ordinal rewards $r_o \in \{1, \dots, n\} \subset \mathbb{N}$, with n representing the number of ordinal rewards. Whereas a real-valued reward provides information about the qualitative size of the reward, the ordinal scale breaks rewards down to naturally ordered *reward tiers*. These reward tiers solely represent the rank of desirability of a reward compared to all other possible

rewards, which is noted as the ranking position r_o of a reward r in the set of all possible rewards $\{r_1, \dots, r_n\}$. Interpreting the reward signals on an ordinal scale still allows us to order and directly compare individual reward signals, but while the numerical scale allows for comparison of rewards by means of the magnitude of their difference, ordinal rewards do not provide this information.

In order to aggregate multiple ordinal rewards, a distribution to store and represent the expected frequency of received rewards on the ordinal scale is constructed. This distribution is represented by a vector $D(s, a)$, in which $d_i(s, a)$ represents the frequency of receiving the ordinal reward r_i by executing a in s . The distribution vector is defined by

$$D(s, a) = \begin{bmatrix} d_1(s, a) \\ \dots \\ d_n(s, a) \end{bmatrix} \quad (5)$$

Through normalization of distribution vector D , a probability distribution P can be constructed, which represents the expected probability of receiving a reward. The probability distribution is represented by a probability vector $P(s, a)$, in which $p_i(s, a)$ represents the estimated probability of receiving the ordinal reward r_i by executing a in s . Hence the probability vector can be defined by

$$P(s, a) = \begin{bmatrix} p_1(s, a) \\ \dots \\ p_n(s, a) \end{bmatrix} \quad \text{with} \quad \sum_{i=1}^n p_i(s, a) = 1 \quad \text{and} \quad 0 \leq p_i(s, a) \leq 1 \quad (6)$$

Value function for ordinal rewards. While numerical rewards enable the representation of value function $V_\pi(s)$ by the expected sum of rewards, the value function for environments with ordinal rewards needs to be estimated differently. Since ordinal rewards are aggregated in a distribution of received ordinal rewards, the calculation of value function $V_\pi(s)$ in state s can be done based on $P(s, a)$ for action a that is selected by policy π . Hence the computation of the value function can be modeled by the following formula of

$$V_\pi(s) = F(P(s, a)) \quad \text{with} \quad a = \pi(s) \quad (7)$$

The computation of the value function from probability distribution $P(s, a)$ through function F is performed by the technique of *measure of statistical superiority* [4]. This measure computes the probability that action a receives a better ordinal reward than a random alternative action a' in the same environmental state s . This probability can be calculated through the sum of all probabilities of a receiving a better ordinal reward o than a' . Hence the probability of an action a performing better than another action a' can be defined as

$$\mathbb{P}(a \succ a') = \sum_{o=1}^n p_o(s, a) \cdot \left(p_{o<}(s, a') + \frac{1}{2} p_o(s, a') \right)$$

$$\text{with } p_{o<}(s, a) = \sum_{i=1}^{o-1} p_i(s, a)$$

To deal with ties, additionally half the probability of a receiving the same reward tier as a' is added.

The function of the measure of statistical superiority therefore is computed through the averaged winning probability of a against all other actions a' by

$$F(P(s, a)) = \mathbb{E}[\mathbb{P}(a \succ a')] = \frac{\sum_{a'} \mathbb{P}(a \succ a')}{k-1} \quad (8)$$

for k available actions in state s .

Based on (7), the optimal policy π^* can be determined in the same way as for numerical rewards (1) by maximizing the respective value function $V_{\pi}(s)$.

4.2 Transformation of existing numerical rewards to ordinal rewards

If an environment has pre-defined rewards on a numerical scale, transforming numerical rewards $r \in \{r_1, \dots, r_n\} \subset \mathbb{R}$ into ordinal rewards $r_o \in \{1, \dots, n\} \subset \mathbb{N}$ can easily be done by translating every numerical reward to its ordinal position within all possible numerical rewards. This way the lowest possible numerical reward is mapped to position 1, and the highest numerical reward is mapped to position n , with n representing the number of possible numerical rewards. This transformation process simply results in removing the metric and semantic of distances of rewards but keeping the order.

4.3 Ordinal Reinforcement Learning

In Section 4.1, we have shown how to compute a value function $V_{\pi}(s)$ and defined the optimal policy π^* for environments with ordinal rewards. This can now be used for adapting common reinforcement learning algorithms to ordinal rewards.

Ordinal Q-learning. For the adaptation of the Q-learning algorithm to ordinal rewards, we do not directly update a Q-value $Q(s, a)$ that represents the quality of a state-action pair (s, a) but update the distribution $D(s, a)$ of received ordinal rewards. The target distribution is computed by adding the received ordinal reward i (represented through unit vector e_i of length n) to the distribution $D(s', \pi^*(s'))$ of taking an action in the new state s' according to the optimal policy π^* . The previous distribution $D(s, a)$ is updated with the target distribution by interpolating both values with learning rate α , which can be seen in the formula

$$D(s, a) = D(s, a) + \alpha[e_i(s, a) + \gamma D(s', \pi^*(s')) - D(s, a)] \quad (9)$$

In this adaptation of Q-learning¹, the expected quality of state-action pair (s, a) is not represented by the Q-value $Q(s, a)$ (3) but by the function $F(P(s, a))$ (8) of the probability distribution $P(s, a)$, which is derived from the iteratively updated distribution $D(s, a)$.

¹ This technique of modifying the Q-learning algorithm to deal with rewards on an ordinal scale can analogously be applied to other Q-table based reinforcement learning algorithms like *Sarsa* and *Sarsa- λ* [14]




Figure 1: Example of an array of ordinal deep neural networks for DQN for reward distribution prediction

Ordinal Deep Q-Network. Because ordinal rewards are aggregated by a distribution instead of a numerical value, the neural network is adapted to predict distributions $D(s, a)$ instead of Q-values for every possible action. Hence for one action the network does not predict a 1-dimensional Q-value, but predicts an n -dimensional reward distribution with n being the length of the ordinal scale. Since this distribution has to be computed for each of k actions, the adaptation of the Deep Q-Network algorithm to ordinal rewards requires a differently structured neural network. Contrary to the original Deep Q-Network where one network simultaneously predicts k Q-values for all actions, the structure of the ordinal DQN consists of an array of k neural networks, from which every network computes the expected ordinal reward distribution $D(s, a)$ for one separate action a . In a deep neural network for the prediction of distributions every output node of the network computes one distribution value $d_i(s, a)$. The structure of neural networks used for the prediction of distributions can be seen in Figure 1.

The prediction of the ordinal reward distributions $D(s, a)$ for all actions can afterwards be normalized to a probability distribution and used in order to compute the value function $V_\pi(s)$ through the measure of statistical superiority as has been previously defined in (7). Once the value function and policy have been evaluated, the ordinal variant of the DQN algorithm follows a similar procedure as ordinal Q-learning and updates the prediction of the reward distribution for (s, a) by fitting $D_a^{DQN}(s)$ to the target reward distribution:

$$\hat{D}_a^{DQN}(s) = e_{r_o}(s, a) + \gamma D_{\pi^*(s')}^{DQN}(s') \quad (10)$$

The main difference in the update step between ordinal Q-learning (9) and ordinal DQN consists of fitting the neural network of action a for input s to the expected reward distribution by backpropagation instead of updating a Q-table entry (s, a) . Additional modifications to the ordinal Deep Q-Network in form of experience replay, the split of the target and evaluation network and the usage of a Double DQN are done in a similar fashion as described with the standard DQN algorithm in Section 3.2. These modifications can be seen in the following paragraphs.

Experience replay. A memory M is used to sample multiple saved experience elements (s, a, r_o, s') randomly and replay these previously seen experiences by fitting the ordinal DQN networks to the samples of earlier memory elements.

Target and evaluation networks. In order to prevent unstable behavior by using the same networks for the prediction and updating step, we use separate evaluation networks to predict reward distributions for the policy computation, and use target networks to predict the target reward distributions which are used for fitting the evaluation networks continuously.

Double Deep Q-Network. The neural networks of ordinal DQN tend to overestimate the prediction of the reward distributions for some actions, which may result in an unjustified bias towards certain actions. Therefore, in order to determine the next action to be played by π^* , the measure of statistical superiority is computed based on the reward distributions predicted by the evaluation networks. Afterwards the prediction of the reward distribution for this action is computed by the respective target network.

5 Experiments and Results

In the following, the standard reinforcement algorithms described in Section 3.2 and the ordinal reinforcement learning algorithms described in Section 4.3 are evaluated and compared in a number of testing environments.²

5.1 Experimental setup

The environments which are used for evaluation are provided by OpenAI Gym,³ which can be viewed as a unified toolbox for our experiments. All environments expect an action input after every time step and return feedback in form of the newly reached environmental state, the direct reward for the executed action, and the information whether the newly reached state is terminal. The environments that the algorithms were tested on were *CartPole* and *Acrobot*.⁴

Policies of the reinforcement learning algorithms were modified to use ϵ -greedy exploration [8], which encourages early exploration of the state space and increases exploitation of the learned policy over time. In the experiments the maximum exploitation is reached after half of the total episodes. In order to directly compare the standard and the ordinal variants of reinforcement learning algorithms, the quality of the learned policy and the computational efficiency are investigated across all environments with varying episode numbers. Information

² The source code for the implementation of the experiments can be found in <https://github.com/az79nefy/OrdinalRL>.

³ For further information about OpenAI visit <https://gym.openai.com>.

⁴ Further technical details about the environments *CartPole* and *Acrobot* from OpenAI can be found in <https://gym.openai.com/envs/CartPole-v0/> and <https://gym.openai.com/envs/Acrobot-v1/>.




Figure 2: CartPole scores of standard and ordinal Q-learning for 400 and 10000 episodes

about the quality of the learned policy is derived from the sum of rewards over a whole episode (score) or the win rate while the efficiency is measured by real-time processing time. Additionally to the standard variant with unchanged rewards, the performance of standard Q-learning algorithms is tested with changed rewards in order to simulate the performance on environments where no optimal reward engineering has been performed. It should be noted that the modifications of the rewards is performed under the constraints of remaining existing reward order, therefore not changing the transformation to the ordinal scale. The change of rewards (CR) from the existing numerical rewards $r \in \{r_1, \dots, r_n\}$ is performed for all rewards by the calculation of $r_{CR,i} = \frac{r_i - \min(r)}{100}$.

The parameter configuration of the Q-learning algorithms is learning rate $\alpha = 0.1$ and discount factor $\gamma = 0.9$. The parameter configuration of the Deep Q-Network algorithm is learning rate $\alpha = 0.0005$ and discount factor $\gamma = 0.9$. As for the network specific parameters, the *Adam* optimizer is used for the network fitting, the target network is getting replaced every 300 fitting updates, the experience memory size is 200000 and the replay batch size is 64.

5.2 Experimental results

The results of the comparison between numerical and ordinal algorithms for the CartPole- and Acrobot-environment in terms of score, win rate and computational time are shown and investigated in the following. This comparison is performed based on the averaged results from 10 and respectively 5 independent runs of Q-learning and Deep Q-Network on the environments.

Q-learning. In Figure 2 the scores for the CartPole-environment over the course of 400 and 10000 episodes can be seen which were played by an agent using the ordinal (orange) as well as the standard Q-learning algorithm, with (red) and without (blue) modified rewards. Additionally the individual dots in this figure represent the scores achieved by the respective algorithms by using the optimal policy instead of ϵ -greedy exploration. The evaluation of these scores shows that the ordinal variant of Q-learning performs better than the standard variant with engineered rewards for 400 episodes and reaches the optimal score of 200 quicker




Figure 3: Comparison of value function margin for best action of standard and ordinal Q-learning for 400 and 10000 episodes of CartPole




Figure 4: Acrobot win rates of standard and ordinal Q-learning for 400 and 10000 episodes

for 10000 episodes. Additionally the use of ordinal rewards significantly outperforms the standard variant with modified rewards for both episode numbers. Therefore it can be seen that ordinal Q-learning is able to learn a good policy better than the standard variants for the CartPole-environment.

In order to explain the difference of learned behavior between the standard and ordinal variant, the average relative difference of Q-values $Q(s, a)$ and respectively measure of statistical superiority functions $F(P(s, a))$ for the two possible actions were plotted and compared in Figure 3 for standard (blue) and ordinal (orange) Q-learning. It can be seen for both episode numbers that the policy which is learned by ordinal RL through the measure of statistical superiority converges to a difference of 0, meaning that the function $F(P(s, a))$ converges to similar values for both actions. This can be interpreted as the policy learning to play safely and rarely entering any critical states where this function would indicate strong preference towards one action (e.g. in big angles). On the other side it can be seen for 400 episodes that common RL does not converge towards similar Q-values for the actions over time and therefore a policy is learned that enters critical states more often. It should be noted that the Q-value differences for standard Q-learning converges to 0 for evaluations with more episodes and a safe policy is eventually learned as well.




Figure 5: Comparison of value function margin for best action of standard and ordinal Q-learning for 400 and 10000 episodes of Acrobot

Table 1: Computation time comparison of standard and ordinal Q-learning for varying episode numbers

Number of episodes	CartPole		Acrobot	
	Standard	Ordinal	Standard	Ordinal
400	2.10 s	4.17 s	35.74 s	52.85 s
2000	10.07 s	24.86 s	174.38 s	266.40 s
10000	67.29 s	130.09 s	855.15 s	1258.30 s
50000	354.52 s	667.87 s	4149.78 s	6178.76 s

In Figure 4 the win rates from the Acrobot-environment were plotted over the course of 400 and 10000 episodes similarly as the scores for the CartPole-environment and it can be seen for low episode numbers that while the policy learned by the standard variant of Q-learning with unchanged rewards performs better than the policy learned by the ordinal variant, changing the numerical values of rewards yields the same performance as the ordinal variant. But for high episode numbers it should be noted that the ordinal variant reaches a similar performance as the standard variant with a win rate of 0.3 after 10000 episodes and clearly outperforms the win rate of the standard Q-learning algorithm with CR.

Similar as for the CartPole-environment, the F - and Q -function margins of the best actions over the course of 400 and 10000 episodes were compared in Figure 5 and yield different observations for the standard and ordinal variants, and it can be therefore be concluded that the learned policies differ. While the ordinal variant decreases the relative margin of $F(P(s, a))$ of the best action and therefore learns a policy which plays safely, the standard variant learns a policy which maximizes the Q-value margin of the best action and therefore follows a policy which enters critical states more often. While the standard variant learns a good policy quicker, it should be noted that both policies perform comparably after many episodes despite the policy differences.




Figure 6: CartPole scores of standard and ordinal DQN for 160 and 1000 episodes

As can be seen in Table 1, using the ordinal variant results in an additional computational load by a factor between 0.8 and 1.2 for CartPole and 0.5 for Acrobot. The additionally required computational capacity is caused by the computation of the measure of statistical superiority which is less efficient than computing the expected sum of rewards. This factor could be reduced by using the iterative update of the function *measure of statistical superiority* described in [4].

Deep Q-Network. In Figure 6 the scores achieved in the CartPole-environment by the ordinal as well as the standard Deep Q-Network, with and without CR, can be seen over the course of 160 and 1000 episodes. For 160 episodes it can be seen that ordinal DQN as well as the standard variant without CR converge to a good policy reaching an episode score close to 150. Contrary to this performance, modified rewards negatively impact standard Q-learning and therefore its performance is significantly worse, not reaching a score above 100. Additionally for low episode numbers it should be noted that the policy learned by the ordinal variant of Deep Q-Network is able to achieve good scores faster than the standard variant, matching the observation made for the Q-learning algorithms. The evaluation for 1000 episodes shows that the performances of standard, with and without CR, and ordinal DQNs are comparable.

Figure 7 plots the win rate of Deep Q-Network algorithms for the Acrobot-environment over the course of 160 and 1000 episodes. For 160 episodes standard DQN with engineered rewards performs better than the ordinal variant, but loses this quality once the rewards are modified. For high episode numbers it can be seen that the ordinal variant is comparable to the standard algorithm without CR and solves the environment with a win rate of close to 1.0, but clearly outperforms the standard DQN with modified rewards which is only able to achieve a win rate of 0.6. It should be noted that all variants of DQN are able to learn a better policy than their respective Q-learning algorithms, achieving a higher win rate in less than 160 episodes.

Additionally, it should be noted that the use of the ordinal variant of DQN adds an additional computational factor between 0 and 0.5 for the CartPole-environment and 1.0 for the Acrobot-environment, as can be seen in Table 2.




Figure 7: Acrobot win rates of standard and ordinal DQN for 160 and 1000 episodes

Table 2: Computation time comparison of standard and ordinal DQN

Number of episodes	CartPole		Acrobot	
	Standard	Ordinal	Standard	Ordinal
160	1520.01 s	2232.48 s	3659.44 s	7442.49 s
400	6699.69 s	7001.79 s	9678.80 s	19840.88 s
1000	15428.41 s	15526.84 s	23310.36 s	47755.90 s

Since the evaluation of the ordinal Deep Q-Network algorithm shows comparable results to the standard DQN with engineered rewards and furthermore outperforms the standard variant with modified rewards, it can be concluded that the conversion of the Deep Q-Network algorithm to ordinal rewards is successful. Therefore it has been shown that algorithms of deep reinforcement learning can as well be adapted to the use of ordinal rewards.

6 Conclusion

In this paper we have shown that the use of ordinal rewards for reinforcement learning is able to reach and even improve the quality of standard reinforcement learning algorithms with numerical rewards. We compared RL algorithms for both numerical and ordinal rewards on a number of tested environments and demonstrated that the performance of the ordinal variant is mostly comparable to the learned common RL algorithms that make use of engineered rewards while being able to significantly improve the performance for modified rewards.

Finally, it should be noted that ordinal reinforcement learning enables the learning of a good policy for environments without much effort to manually shape rewards. We hereby lose the possibility of reward shaping to the same degree that numerical rewards would allow, but therefore gain a more simple-to-design reward structure. Hence, our variant of reinforcement learning with ordinal rewards is especially suitable for environments that do not have a natural semantic of numerical rewards or where reward shaping is difficult. Additionally this method enables the usage of new and unexplored environments for RL only

with the specification of an order of desirability instead of the needed effort of manually engineering numerical rewards with sensible semantic meaning.

Acknowledgements

This work was supported by DFG. Calculations for this research were conducted on the Lichtenberg high performance computer of the TU Darmstadt.

References

1. Fürnkranz, J., Hüllermeier, E. (eds.): Preference Learning. Springer-Verlag (2011)
2. Gilbert, H., Weng, P.: Quantile reinforcement learning. CoRR **abs/1611.00862** (2016)
3. Hasselt, H.v., Guez, A., Silver, D.: Deep Reinforcement Learning with Double Q-Learning. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. pp. 2094–2100. AAAI’16, AAAI Press (2016)
4. Joppen, T., Fürnkranz, J.: Ordinal Monte Carlo Tree Search. CoRR **abs/1901.04274** (2019)
5. Lin, L.J.: Reinforcement Learning for Robots Using Neural Networks. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, USA (1992), uMI Order No. GAX93-22750
6. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.A.: Playing Atari with Deep Reinforcement Learning. CoRR **abs/1312.5602** (2013)
7. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M.A., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529–533 (2015)
8. Sutton, R.S., Barto, A.G.: Reinforcement learning - an introduction. Adaptive computation and machine learning, MIT Press, second edn. (2018)
9. Watkins, C.J., Dayan, P.: Q-learning. *Machine Learning* **8**, 279–292 (1992)
10. Weng, P.: Markov decision processes with ordinal rewards: Reference point-based preferences. In: Proceedings of the 21st International Conference on Automated Planning and Scheduling (ICAPS-11). AAAI Press, Freiburg, Germany (2011)
11. Weng, P.: Ordinal Decision Models for Markov Decision Processes. In: Proceedings of the 20th European Conference on Artificial Intelligence (ECAI-12). pp. 828–833. IOS Press, Montpellier, France (2012)
12. Weng, P., Busa-Fekete, R., Hüllermeier, E.: Interactive Q-Learning with Ordinal Rewards and Unreliable Tutor. In: Proceedings of the ECML/PKDD-13 Workshop on Reinforcement Learning from Generalized Feedback: Beyond Numeric Rewards (2013)
13. Wirth, C., Akrou, R., Neumann, G., Fürnkranz, J.: A Survey of Preference-Based Reinforcement Learning Methods. *Journal of Machine Learning Research* **18**(136), 1–46 (2017)
14. Zap, A.: Ordinal Reinforcement Learning. Master’s thesis, Technische Universität Darmstadt (2019), To appear