

Towards Multi-agent Reinforcement Learning using Quantum Boltzmann Machines

Tobias Müller, Christoph Roch, Kyrill Schmid and Philipp Altmann

Mobile and Distributed Systems Group, LMU Munich, Germany

Keywords: Multi-agent, Reinforcement Learning, D-Wave, Boltzmann Machines, Quantum Annealing, Quantum Artificial Intelligence.

Abstract: Reinforcement learning has driven impressive advances in machine learning. Simultaneously, quantum-enhanced machine learning algorithms using quantum annealing underlie heavy developments. Recently, a multi-agent reinforcement learning (MARL) architecture combining both paradigms has been proposed. This novel algorithm, which utilizes Quantum Boltzmann Machines (QBMs) for Q-value approximation has outperformed regular deep reinforcement learning in terms of time-steps needed to converge. However, this algorithm was restricted to single-agent and small 2x2 multi-agent grid domains. In this work, we propose an extension to the original concept in order to solve more challenging problems. Similar to classic DQNs, we add an experience replay buffer and use different networks for approximating the target and policy values. The experimental results show that learning becomes more stable and enables agents to find optimal policies in grid-domains with higher complexity. Additionally, we assess how parameter sharing influences the agents' behavior in multi-agent domains. Quantum sampling proves to be a promising method for reinforcement learning tasks, but is currently limited by the Quantum Processing Unit (QPU) size and therefore by the size of the input and Boltzmann machine.

1 INTRODUCTION

Recently, adiabatic quantum computing has proven to be a useful extension to machine learning tasks (Benedetti et al., 2018; Biamonte et al., 2017; Li et al., 2018; Neukart et al., 2017a). Especially hard computational tasks with high data volume and dimensionality have benefitted from the possibility of using quantum devices with manufactured spins to speed-up computational bottlenecks (Neven et al., 2008; Rebertrost et al., 2014; Wiebe et al., 2012).

One specific type of machine learning is Reinforcement Learning (RL), where an interacting entity, called agent, aims to learn an optimal state-action policy through trial and error (Sutton and Barto, 2018). Reinforcement Learning has gained the public attention by defeating the 9-dan Go grandmaster *Lee Sedol* (Silver et al., 2016), which has been thought to be impossible for a machine. In the latest years, reinforcement learning has seen many improvements, gained a large variety of application fields like economics (Charpentier et al., 2020), autonomous driving (Kiran et al., 2020), biology (Mahmud et al., 2018) and even achieved superhuman performance

in chip design (Mirhoseini et al., 2020). Reinforcement Learning has only seen quantum speed-ups for special models (Levit et al., 2017; Neukart et al., 2017a; Neukart et al., 2017b; Paparo et al., 2014). Especially multi-agent domains have rarely been researched (Neumann et al., 2020).

Real-world reinforcement learning frameworks predominantly use deep neural networks (DNNs) as function approximators. Since DNNs are powerful - see the latest prominent example *AlphaFold2* (Jumper et al., 2020) - and can be run efficiently for large datasets on classical computers, deep reinforcement learning is able to tackle complex problems in large data spaces. Hence, there was little need for improvements.

However, since recent work has proved speed-ups for classical RL by leveraging quantum computing (Levit et al., 2017; Neumann et al., 2020) and the application field gets more and more complex, it could be beneficial to explore quantum RL algorithms. These inspiring studies considered Boltzmann machines (Ackley et al., 1985) as function approximator - instead of traditionally used DNNs. Boltzmann machines are stochastic neural networks, which

are mainly avoided due to the fact, that their training times are exponential to the input size. Since finding the energy minimum of Boltzmann machines can be formulated as a "Quadratic Unconstrained Binary Optimization" (QUBO) problem, simulated annealing respectively quantum annealing is well suited to accelerate training time.

Nevertheless, the combination of RL and Boltzmann machines using (simulated) quantum annealing only worked properly for small single-agent environments and reached its limit at a simple 3×3 multi-agent domain. This work proposes an architecture inspired by DQNs (Mnih et al., 2015) to enable more complex domains and stabilize learning by using *experience replay buffer* and separating policy and target networks. We thoroughly evaluate the effects of these augmentations on learning.

Lately, an inspiring novel method to speed-up quantum reinforcement learning for large state and action spaces by proposing a combination of regular NNs and DBMs/QBMs, namely *Deep Energy Based Networks* (DEBNs) was proposed (Jerbi et al., 2020). More specifically, these architectures are constructed with an input layer consisting of action and state units, which are connected with the first hidden layer through directed weights. This is followed by a single undirected stochastic layer. The remaining layers are linked with directed deterministic connections. Lastly, a final output layer returns the negative free energy $-F(s, a)$.

In contrast to QBMs, DEBNs therefore only comprise one stochastic layer, return an output similar to traditional deep neural networks and can be trained through backpropagation. DEBNs also use an experience replay buffer and separate the policy and target network. Additionally, they allow to trade off learning performance for efficiency of computation. Jerbi et al. briefly stated, that QBMs are applicable. Unfortunately, no numerical results were given for purely stochastic, energy-based QBM agents or domains with multiple agents. We aim to build on this.

Summarized, our contribution is three-fold:

- We provide a Quantum Reinforcement Learning (Q-RL) framework, which stabilizes learning leading to more optimal policies
- Based on single- and multi-agent domains, we provide a thorough evaluation on the effects of an Experience Replay Buffer and an additional Target Network compared to traditional QBM agents
- Additionally, we demonstrate and discuss limitations to the concept

We first describe the preliminaries about reinforcement learning and quantum Boltzmann machines underlying the proposed architectures. Afterwards, the state-of-the-art algorithm and extensions made to it will be explained. We test and evaluate the approach and finally discuss restrictions and potential grounds for future work.

2 PRELIMINARIES

This chapter describes the basics needed to understand our proposed architecture. First, reinforcement learning and the underlying Markov Decision Process will be explained followed by Boltzmann Machines and the process of quantum annealing.

2.1 Reinforcement Learning

We first describe Markov Decision Processes as the underlying problem formulation which is followed by an introduction to reinforcement learning in general. The subsequent sections specify independent and cooperative multi-agent reinforcement learning.

Markov Decision Processes. The problem formulation is based on the notion of Markov Decision Processes (MDP) (Puterman, 1994). MDPs are a class of sequential decision processes and described via the tuple $M = \langle S, A, P, R \rangle$, where

- S is a finite set of states and $s_t \in S$ the state of the MDP at time step t .
- A is the set of actions and $a_t \in A$ the action the MDP takes at time step t .
- $P(s_{t+1}|s_t, a_t)$ is the probability transition function. It describes the transition that occurs when action a_t is executed in state s_t . The resulting state s_{t+1} is chosen according to P .
- $R(s_t, a_t)$ is the reward, when the MDP takes action a_t in state s_t . We assume $R(s_t, a_t) \in \mathbb{R}$

Consequently, the cost and transition function only depend on the current state and action of the system. Eventually, the MDP should find a policy $\pi: S \rightarrow A$ in the space of all possible policies Π , which maximizes the return G_t at state s_t over an infinite horizon via:

$$G_t = \sum_{k=0}^{\infty} \gamma^k \cdot R(s_{t+k}, a_{t+k}), \quad (1)$$

with $\gamma \in [0, 1]$ as the discount factor. This policy is called the optimal policy π .

Reinforcement Learning. Model-free reinforcement learning (Strehl et al., 2006) is considered to search the policy space Π in order to find the optimal policy π^* . The interacting reinforcement learning agent executes an action a_t for every time step $t \in [1, \dots]$ in the MDP environment. In model-free algorithms, the agent acts without any knowledge of the environment and the algorithm only keeps information of the value-function. Therefore, the agent knows its current state s_t and the action space A , but neither the reward nor the next state s_{t+1} of any action a_t in any state s_t .

Consequently, the agent needs to learn from delayed rewards without having a model of the environment. A popular value-based approach to solve this problem is Q-learning (Peng and Williams, 1994). In this approach, the action-value function $Q^\pi : S \times A \rightarrow \mathbb{R}$, $\pi \in \Pi$ describes the accumulated reward $Q^\pi(s_t, a_t)$ for an action a_t in state s_t . The optimal Q-learning function Q^* is approximated by starting from an initial guess for Q and updating the function via:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)] \quad (2)$$

The learned Q-function will eventually converge to Q^* , which then implies an optimal policy. In the traditional experiments a deep neural network is used as a parameterized function approximator to calculate the optimal action for a given state.

Independent Multi-agent Learning. When multiple agents interact with the environment, a fully cooperative multi-agent task can be described as a stochastic game G , defined as in (Foerster et al., 2017) via the tuple $G = \langle S, A, P, R, Z, O, n, \gamma \rangle$, where:

- S is a finite set of states. At each time step t , the environment has a true state $s_t \in S$.
- A is the set of actions. At each time step t each agent ag simultaneously chooses an action $a^{ag} \in A$, forming a joint action $a \in A \equiv A^n$.
- $P(s_{t+1} | s_t, a_t)$ is the probability transition function as previously defined.
- $R(s_t, a_t)$ is the reward as previously defined. All agents share the same reward function.
- Z is a set of observations of a partially or fully observable environment.
- $O(s, ag)$ is the observation function. Each agent draws observations $z \in Z$ according to $O(s, ag)$.
- n is the number of agents identified by $ag \in AG \equiv \{1, \dots, n\}$.
- $\gamma \in [0, 1)$ is the discount factor.

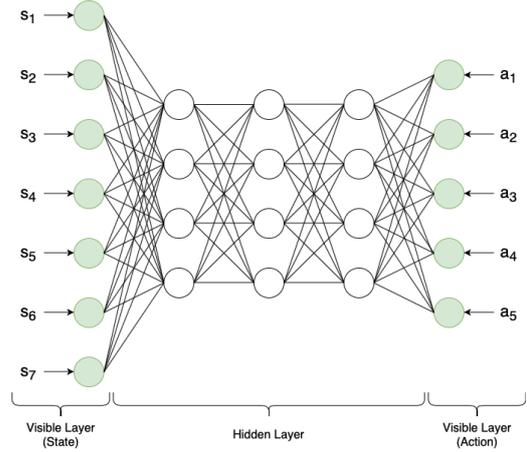


Figure 1: A Deep Quantum Boltzmann Machine with seven input state neurons and five input action neurons. The QBM additionally consists of three hidden layers with four neurons each. The state and action are given as fixed input and the configuration of the hidden neurons are sampled via (simulated) quantum annealing. The weights between two neurons are updated in the Q-learning step as described in section 3.

In *independent* multi-agent learning algorithms, each agent learns from its own action-observation history and is trained independently. This means, every agent simultaneously learns its own Q-function (Tan, 1993).

2.2 Boltzmann Machines

The structure of a Boltzmann machine (BM) (Ackley et al., 1985) is similar to Hopfield networks and can be described as a stochastic energy-based neural network. A traditional BM consists of a set of visible nodes V and a set of hidden nodes H , where every node represents a binary random variable. The binary nodes are connected through real-valued, bidirectional, weighted edges of the underlying undirected graph. The global energy configuration is generally given by the energy level of Hopfield networks. Since clamped BMs fix the assignment of the visible binary variables, these nodes are removed from the underlying graph and contribute as constant coefficients to the associated energy. Therefore the formula, which we aim to minimize, is given as the energy level of Hopfield networks with constant visible nodes:

$$E(h) = - \sum_i w_{ii} v_i - \sum_j w_{jj} h_j - \sum_i \sum_j v_i w_{ij} h_j, \quad (3)$$

with v_i as the visible nodes, h_j as the hidden nodes and weights w .

For this work, we implemented a Deep Boltzmann Machine (DBM) as trainable state-action approximator, which is constructed with multiple hidden layers,

one visible input layer for the state and one visible action input layer. Finally, we modified the DBM to get a Quantum Boltzmann Machine (QBM), where qubits are associated to each node of the network instead of random binary variables (Crawford et al., 2019; Neumann et al., 2020; Levit et al., 2017). A visualization of a QBM for seven state neurons and five input action neurons can be seen in figure 1. For any QBM with $v \in V$ and $h \in H$, the energy function is described by the quantum Hamiltonian \mathcal{H}_v :

$$\mathcal{H}_v = - \sum_{v,h} w_{vh} v \sigma_h^z - \sum_{v,v'} w_{vv'} v v' - \sum_{h,h'} w_{hh'} \sigma_h^z \sigma_{h'}^z - \Gamma \sum_h \sigma_h^x \quad (4)$$

Furthermore, Γ is the annealing parameter, while σ_i^z and σ_i^x are spin-values of node i in the z - and x - direction. Because measuring the state of one direction destroys the state of the other, we follow the architecture of Neumann et al. (2020) (Neumann et al., 2020) and replace all σ_i^x by σ_i^z by using replica stacking based on the Suzuki-Trotter expansion of the Hamiltonian \mathcal{H}_v . The BM is replicated r times in total and connections between corresponding nodes in adjacent replicas are added. By this, we obtain a new effective Hamiltonian $\mathcal{H}_{v=(s,a)}^{eff}$ in its clamped version given by:

$$\mathcal{H}_{v=(s,a)}^{eff} = - \sum_{h \in H} \sum_{k=1}^r \frac{w_{sh}}{r} \sigma_{h,k} - \sum_{h \in H} \sum_{k=1}^r \frac{w_{ah}}{r} \sigma_{h,k} - \sum_{(h,h') \subseteq H} \sum_{k=1}^r \frac{w_{hh'}}{r} \sigma_{h,k} \sigma_{h',k} - \Gamma \sum_{h \in H} \sum_{k=0}^r \sigma_{h,k} \sigma_{h,k+1} \quad (5)$$

For each evaluation of the Hamiltonian, we get a spin configuration \hat{h} . After n_{reads} reads for a fixed combination of s and a , we get a multi-set $\hat{h}_{s,a} = \{\hat{h}_1, \dots, \hat{h}_{n_{reads}}\}$. We average over this multi-set to gain a single spin configuration $C_{\hat{h}_{s,a}}$, which will be used for updating the network. If a node is $+1$ or -1 depends on the global energy configuration:

$$P_{node\ i=1} = \frac{1}{1 + \exp(-\frac{\Delta E_i}{T})}, \quad (6)$$

with T as the current temperature.

Since the structure of Boltzmann Machines are inherent to Ising models, we sample spin values from the Boltzmann distribution by using simulated quantum annealing, which simulates the effect of transverse-field Ising model by slowly reducing the temperature or strength of the transverse field at finite temperature to the desired target value (Levit et al.,

2017). As proven in (Morita and Nishimori, 2008), spin system defined by simulated quantum annealing converges to quantum Hamiltonian. Therefore it is straightforward to use simulated quantum annealing (SQA) to find a spin configuration for $h \in H$ - given $s \in S$ - which minimizes the free energy.

3 QUANTUM REINFORCEMENT LEARNING

Recently, quantum reinforcement learning algorithms (QRL) using boltzmann machines and quantum annealing of single agent (Crawford et al., 2019) and multi-agent domains (Neumann et al., 2020) for learning grid-traversal policies have been proposed. Although, these architectures were able to learn optimal policies in less time steps compared to classic deep reinforcement learners (DRL), they could only be applied to single-agent or small multi-agent domains. Unfortunately, already 3×3 domains with 2 agents could not be solved optimally (Neumann et al., 2020). QRL seems to be unstable for more complex domains. We intuitively assume that BMs underlie similar instability problems as traditional neural networks. Hence, by correlations present in the sequence of observations and how small updates to the Q-values change the policy, data distribution and therefore the correlations between free energy $F(s_n, a_n)$ and target energy $F(s_{n+1}, a_{n+1})$. Inspired by Deep Q-Networks (Mnih et al., 2015), we propose to enhance the state-of-the-art architecture as described in section 3.1 by adding an experience replay buffer (see section 3.2) to randomize over transitions and by separating the network calculating the policy and the network approximating the target value (see section 3.3) in order to reduce correlations with the target.

3.1 State of the Art

Traditionally, single-agent reinforcement learning using quantum annealing and QBMs is an adaption of Sallans and Hinton (2004) (Sallans and Hinton, 2004) RBM RL algorithm and structured as follows:

Initialization. The weights of the QBM are initialized by setting the weights using Gaussian zero-mean values with a standard deviation of 1.00. The topology of the hidden layers is set beforehand.

Policy. At the beginning of each episode, every agent is set randomly onto the grid and receives its

corresponding observation. At each time step t , every agent i independently chooses an action a_t^i according to its policy π_t^i . To enable exploration, we implemented an ϵ -greedy policy, where the agent acts random with probability ϵ , which decreases by $\epsilon_{decay} = 0.0008$ with each training step until $\epsilon_{min} = 0.01$ is reached. When the agent follows its learned policy, we sweep across all possible actions and choose the action which maximizes the Q-value for state s_t^i . The Q-function of state s and action a is defined as the corresponding negative free-energy $-F$:

$$Q(s, a) \approx -F(s, a) = -F(s, a; w), \quad (7)$$

with w as the vector of weights of a QBM and $F(s, a)$ as:

$$F(s, a) = \langle H_{v=(s,a)}^{eff} \rangle - \frac{1}{\beta} P(s_{t+1}|s_t, a_t) \log P(s_{t+1}|s_t, a_t) \quad (8)$$

Summarized, the agent acts via:

$$\pi = \begin{cases} \text{random}, & \text{if } p \geq \epsilon \\ \text{argmax } Q(s, a), & \text{if } p < \epsilon \end{cases} \quad (9)$$

for $a \in A$ and random variable p .

Weight Update. The environment returns a reward r_{t+1}^i and second state $s_{t+1}^i \leftarrow a_t^i(s_t^i)$ for each agent i . Based on this transition, the QBM is trained. The used update rules are an adaption of the state-action-reward-state-action (SARSA) rule by Rummery et al. (1994) (Rummery and Niranjan, 1994) with negative free energy instead of Q-values (Levit et al., 2017) defined as:

$$\Delta w^{vh} = \mu(r_n(s_n, a_n) - \gamma F(s_{n+1}, a_{n+1}) + F(s_n, a_n)) v \langle \sigma_h^z \rangle \quad (10)$$

$$\Delta w^{hh'} = \mu(r_n(s_n, a_n) - \gamma F(s_{n+1}, a_{n+1}) + F(s_n, a_n)) v \langle \sigma_h^z \sigma_{h'}^z \rangle, \quad (11)$$

with γ as the discount factor and μ as the learning rate. The free energy and configurations of the hidden neurons are gained by applying simulated quantum annealing respectively quantum annealing to the formulation of the effective Hamiltonian $H_{v=(s,a)}^{eff}$ as described in the previous section. At each episode, this process is repeated for a defined number of steps or until the episode ends.

3.2 Experience Replay Buffer

The first extension is a biologically inspired mechanism named experience replay (McClelland et al., 1995; O'Neill et al., 2010). O'Neill et al. (2010)

found, that the human brain stabilizes memory traces from short- to long-term memory by replaying memories during sleep and rest. The reactivation of brain-wide memory traces could underlie memory consolidation. Similar to the human brain, experience replay buffers used in deep Q-networks (DQN) store experienced transitions and provides randomized data during updating neural connections. Hence, correlations of observation sequences are removed and changes in the data distribution are smoothed. Furthermore, due to the random choice of training samples, one transition can be used multiple times to consolidate experiences.

To enable experience replay, at each time step t we store the each agents' experience $e_t = (s_t, a_t, r_t, s_{t+1})$ in a data set $D_t = (e_1, \dots, e_t)$. For every training step, we randomly sample mini-batches from D_t from which to Q-learning updates are performed.

This means, instead of updating the weights on state-action pairs as they occur, we store discovered data and perform training on random mini-batches from a pool of random transitions.

3.3 Policy and Target Network

In order to perform a training step, it is necessary to calculate the policy value $F(s_n, a_n)$ and target value $F(s_{n+1}, a_{n+1})$. Currently, policies and target values are approximated by the same network. Consequently, Q-values and target values are highly correlated. Small updates to Q-values may significantly change the policy, data distribution and target.

To counteract, we separate policy network calculating $F(s_n, a_n)$ from the target network approximating $F(s_{n+1}, a_{n+1})$. Both networks are initialized similarly. The policy network is updated with every training step, whereas the target network is only periodically updated. Every m steps, the weights of the policy network are simply adopted by the target network.

3.4 Multi-agent Quantum Reinforcement Learning

In this work, we explore independent quantum learning in cooperative and non-cooperative settings. The explicit requirement for cooperation is communication (Binmore, 2007). We enable communication via parameter sharing as proposed by Foerster et al. (2016) (Foerster et al., 2016). In this case, every agents' transition is stored in a centralized experience replay buffer and only one BM is trained. Each agent receives its own observation and the centralised network approximates the agents' Q-value independently. Whereas in non-cooperative settings, every

agent keeps and updates its own BM solely with its own experiences without any information exchange. The policy and weight updates are performed as described in the previous section.

4 EVALUATION

4.1 Domain

To evaluate our approach, we implemented a discrete $n \times m$ multi-agent grid-world domain with i deterministic rewards and i agents. At every time step t each agent independently chooses an action from action space $A = \{up, down, left, right, stand\ still\}$ depending on the policy π . More specifically, the goal of every agent is to collect corresponding balls while avoiding obstacles (e.g. walls and borders) and penalty states (e.g. pits and others' balls). The environment size, number of agents, balls and obstacles can be easily modified. Reaching a target location is rewarded by a value of 220, whereas penalty states are penalized by -220 and an extra penalty of -10 is given for every needed step. An agent is done, when all its corresponding balls were collected. Consequently, we consider the domain as solved, when every agent is done. The main goal lies in efficiently navigating through the grid. Two example domains can be seen in figure 2.

The starting position of all agents are chosen randomly at the beginning of each episode whereas the locations of their goals are fixed. The observation is one-hot-encoded and divided into two layers. One layer describes the agents' position and its goal and the other layer details the position of all other agents and their goals. This observation is issued as input for the algorithm. Therefore, the input shape is $n \times m \times 2$. To assess the learned policies, we use the accumulated episode rewards as quality measure.

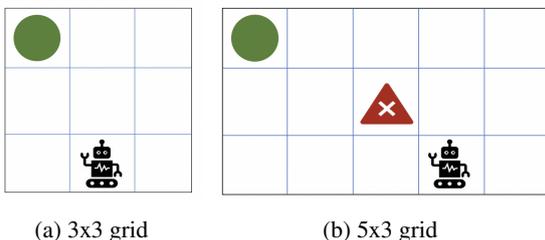


Figure 2: Example figures of two single-agent domains. Picture a) shows a 3×3 grid domain with one reward, whereas b) illustrates a bigger 5×3 grid domain with an additional penalty state.

4.2 Single-agent Results

First, we evaluate how adding an experience replay buffer (ERB) and separating policy and target network influences the learning process and performance of a single agent. We started by running the traditional Q-RL algorithm as proposed by Neumann et al. (2020) (Neumann et al., 2020) including their parameter setting. Then, we only added an experience replay buffer (ERB) respectively solely the target network. Finally, we extended the original algorithm with a combination of both, an ERB and target network. The resulting rewards on running all four architecture on the 3×3 domain (see figure 2) can be seen in figure 3 a) and the corresponding learned policy in figure 3 b). All graphs have been averaged over ten runs. The traditional Q-RL agent without any extensions (blue line) learns unstable with occasional high swings down to -1700 and -1000 reward points. Extended versions seem to show less outliers. This observation gets more evident, when conducting the same experiment on a bigger 5×3 environment. As seen in figure 3 c) - d) the achieved rewards of non-extended agents (blue) collapses frequently. The ERB (black) respectively target network (green) alone stabilize learning, but the combination of both (red) yields smoothest training curve. Hence, these enhancements are getting more important with bigger state space and more complex environments.

After training, we evaluate the resulting policies for 100 episodes without further training. The average rewards of ten test-runs on the 3×3 domain can be seen in figure 3 b). As already described, an agent is rewarded +220 points for reaching its goal and -10 for each taken step. So, when considering an optimal policy, the agent would be awarded +190 for the 3×3 domain (respectively +170 for 5×3) if the agent is spawned furthest from its goal and +220 for the best starting position. Assuming the starting positions over all episodes are distributed evenly, the optimal median reward would be at +205 for the 3×3 domain and +195 for the 5×3 environment.

The traditional QBM agent shows multiple outliers and a higher spread of rewards throughout the evaluation episodes compared to the other architectures. As it can be seen, adding only one of the extensions leads to a better median reward and a seemingly optimal policy is gained through a combination of both. Again, this observation gets more distinct with bigger domains, see figure 3 d). Even though ERB or target network alone significantly enhance the median reward, the plots still show outliers. The combined architecture is free of outliers with less interquartile range and lower overall span indicating re-

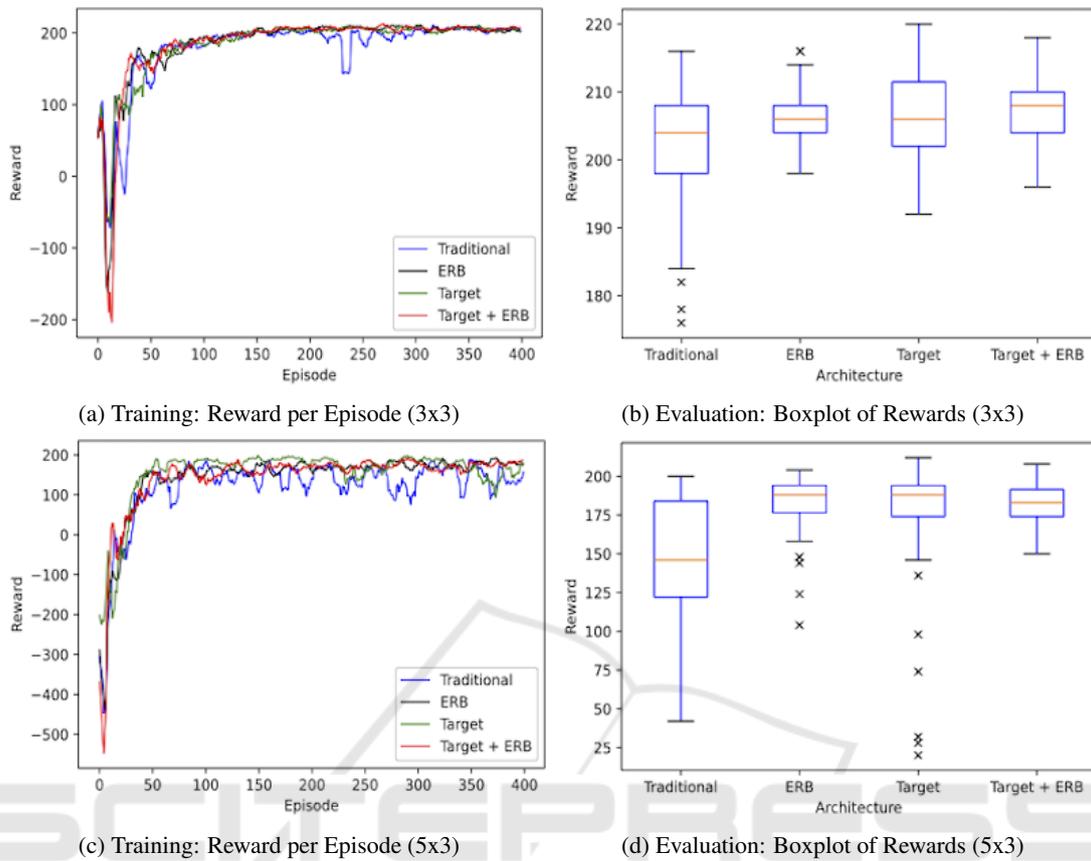


Figure 3: Performance of a single agent with different architectures. a) Shows the gained reward per episode on a 3×3 domain of different architectures, whereas b) displays the corresponding achieved rewards of the learned policy on 400 test episodes. c) illustrates the reward of the same experiment on a 5×3 domain and d) the corresponding learned policy of test episodes.

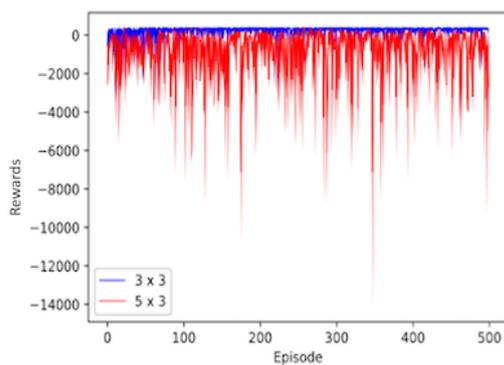
duced variance of training performance and nearly optimal policy. In summary, alleviating data correlation and the problems non-stationary distributions by randomly sampling previous transitions and separating target and policy network increases stable learning leading to robust and more optimal policies. Comparing the results for 3×3 with the 3×5 gridworld, a correlation of impact through the extensions and input size can be suspected.

4.3 Multi-agent Results

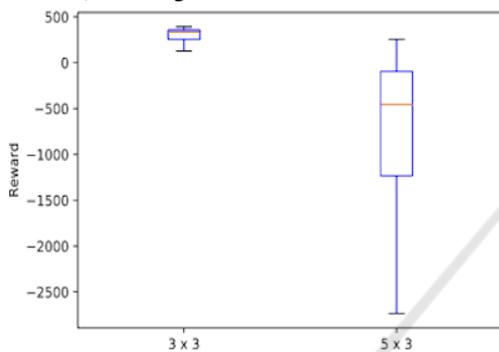
Traditional Q-RL was limited so 2×2 multi-agent domains and bigger domains could not be solved rationally (Neumann et al., 2020). This section explores, if the proposed architecture enables multi-agent reinforcement learning. We modify the known environments by adding one agent and one corresponding goal. If an agent picks up the others goal, it is penalized with -220. The averaged results over 10 runs can be seen in figure 4.

The graphs suggest, that 3×3 domain (blue) can be solved in contrast to the bigger environment (red). Looking at figure 4 b), the median reward of the learned policy on the smaller domain is around +350, which is near optimum. Unfortunately, the bigger domain could not be solved with a median reward of -450. Additionally, the 5×3 learning curve does not seem to converge. Therefore, we can conclude, that it is possible to solve bigger domains with the proposed architecture, but Q-RL with ERB and extra target network still fails in somewhat larger multi-agent domains.

Lastly, we explore if the cooperation method of parameter sharing enhances quantum multi-agent reinforcement learning. With parameter sharing no explicit communication is necessary since only one centralized entity is trained and shared between the agents. More specifically, the experience of every agent is stored in a centralized ERB. At each training step, one QBM is trained with a randomized sample from the ERB similar to the single-agent case.



a) Learning Process on Both Domains



b) Learned Policy on Both Domains

Figure 4: Performance of two agents on the 3×3 (blue) respectively 3×5 (red) domain. Figure a) shows the learning process over 500 episodes, whereas figure b) displays the learned policy over 100 testing episodes.

Both agents use this network to independently calculate their Q-values based their observation. By this, we additionally smooth the data distribution hoping to achieve a more general policy and not two specific policies adjusted to particular observations.

The results with and without parameter sharing are illustrated in figure 5. Unfortunately, parameter sharing seems to have a negative effect on the small 3×3 domain. In this case, the agents seem to have learned a worse policy with this adaption. Rewards on the bigger environments have increased. However, the 5×3 domain can still not be considered solved. Hence, parameter sharing is sub-optimal for the evaluated use case.

The complexity of the task and size of the input did not increase, so this observation is counter-intuitive. Since the centralized entity is simultaneously learning two independent behaviors, it might be possible that in this case two independently optimal action-state probability distributions (as learned without parameter sharing) cancel out each other when learned together. To proof this assumption, more experiments must be conducted.

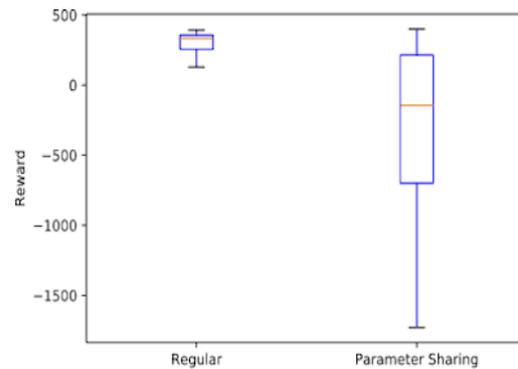
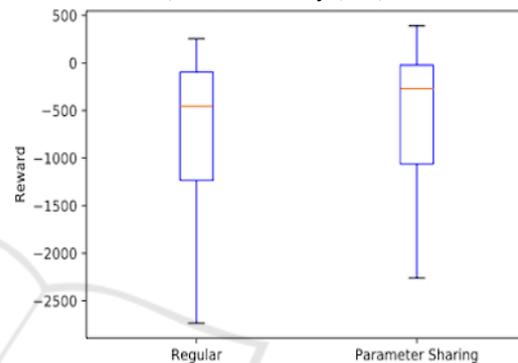
a) Learned Policy (3×3)b) Learned Policy (5×3)

Figure 5: Performance of a two agents with and without parameter sharing. a) Shows the gained reward of the learned policy of 100 testing episodes on a 3×3 domain, whereas b) displays the same experiment on the bigger environment.

5 DISCUSSION

In summary, adding an ERB and additional target network alleviates data correlation and the problem of non-stationary distribution resulting in stabilized learning and a more optimal policies. With the proposed architecture, we were able to solve bigger environments compared to traditional MARL using QBMs. However, this architecture is still limited to relatively small domains.

Even though it is possible to coordinate a single agent in the 5×3 domain and multiple agents in a smaller domain. The question remains why the 5×3 multi-agent domain fails. The QBM-agent receives an input of 15 neurons on 5×3 single-agent domain since only one input layer is needed. When adding more agents to the environment, there is another input layer necessary in order to distinguish between the acting agent and other opposing agents. Hence, the 3×3 multi-agent domain returns an observation size of 18 and bigger multi-agent domain of size 30. The input are considered in the QUBO formulation,

which therefore increases. Hence, simulated quantum annealing is applied to a bigger formulation. A bigger formulation demands more qubits, which may limit the accuracy, variation and stability of the quantum annealing algorithm. This is only an assumption and needs to be examined more closely. Neumann et al. (2020) also already stated, that Q-RL is limited by the current Quantum Processing Unit (QPU) size. However, with the extension of an Experience Replay Buffer and Target Network, we are able to stabilize learning and therefore may reduce the needed QPU size compare to previous approaches.

Quantum sampling has been proven to be a promising method to enhance reinforcement learning tasks to speed-up learning in relation to needed time steps (Neumann et al., 2020). Further work concerning the relation between QPU size and domain complexity (respectively state input) would be needed to strictly determine current limitations.

ACKNOWLEDGEMENTS

This work was funded by the BMWi project PlanQK (01MK20005I).

REFERENCES

- Ackley, D. H., Hinton, G. E., and Sejnowski, T. J. (1985). A learning algorithm for boltzmann machines. *Cognitive Science*, 9(1):147 – 169.
- Benedetti, M., Realpe-Gómez, J., and Perdomo-Ortiz, A. (2018). Quantum-assisted helmholtz machines: A quantum–classical deep learning framework for industrial datasets in near-term devices. *Quantum Science and Technology*, 3(3):034007.
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., and Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671):195–202.
- Binmore, K. (2007). *Game Theory: A Very Short Introduction*. Oxford University Press.
- Charpentier, A., Elie, R., and Remlinger, C. (2020). Reinforcement learning in economics and finance.
- Crawford, D., Levit, A., Ghadermarzy, N., Oberoi, J. S., and Ronagh, P. (2019). Reinforcement learning using quantum boltzmann machines.
- Foerster, J. N., Assael, Y. M., de Freitas, N., and Whiteson, S. (2016). Learning to communicate with deep multi-agent reinforcement learning.
- Foerster, J. N., Farquhar, G., Afouras, T., Nardelli, N., and Whiteson, S. (2017). Counterfactual multi-agent policy gradients. In *AAAI*.
- Jerbi, S., Trenkwalder, L. M., Nautrup, H. P., Briegel, H. J., and Dunjko, V. (2020). Quantum enhancements for deep reinforcement learning in large spaces.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Tunyasuvunakool, K., Ronneberger, O., Bates, R., Žídek, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Potapenko, A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Steinegger, M., Pacholska, M., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P., and Hassabis, D. (2020). High accuracy protein structure prediction using deep learning. In *Fourteenth Critical Assessment of Techniques for Protein Structure Prediction (Abstract Book)*, 14.
- Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Sallab, A. A., Yogamani, S., and Pérez, P. (2020). Deep reinforcement learning for autonomous driving: A survey.
- Levit, A., Crawford, D., Ghadermarzy, N., Oberoi, J. S., Zahedinejad, E., and Ronagh, P. (2017). Free energy-based reinforcement learning using a quantum processor.
- Li, R. Y., Di Felice, R., Rohs, R., and Lidar, D. A. (2018). Quantum annealing versus classical machine learning applied to a simplified computational biology problem. *npj Quantum Information*, 4(1).
- Mahmud, M., Kaiser, M. S., Hussain, A., and Vassanelli, S. (2018). Applications of deep learning and reinforcement learning to biological data. *IEEE Transactions on Neural Networks and Learning Systems*, 29(6):2063–2079.
- McClelland, J., McNaughton, B., and O'Reilly, R. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102:419–57.
- Mirhoseini, A., Goldie, A., Yazgan, M., Jiang, J., Songhori, E., Wang, S., Lee, Y.-J., Johnson, E., Pathak, O., Bae, S., Nazi, A., Pak, J., Tong, A., Srinivasa, K., Hang, W., Tuncer, E., Babu, A., Le, Q. V., Laudon, J., Ho, R., Carpenter, R., and Dean, J. (2020). Chip placement with deep reinforcement learning.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A., Veness, J., Bellemare, M., Graves, A., Riedmiller, M., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518:529–33.
- Morita, S. and Nishimori, H. (2008). Mathematical foundation of quantum annealing. *Journal of Mathematical Physics*, 49(12):125210.
- Neukart, F., Compostella, G., Seidel, C., von Dollen, D., Yarkoni, S., and Parney, B. (2017a). Traffic flow optimization using a quantum annealer.
- Neukart, F., Dollen, D. V., Seidel, C., and Compostella, G. (2017b). Quantum-enhanced reinforcement learning for finite-episode games with discrete state spaces.
- Neumann, N., Heer, P., Chiscop, I., and Phillipson, F. (2020). Multi-agent reinforcement learning using simulated quantum annealing.
- Neven, H., Denchev, V. S., Rose, G., and Macready, W. G. (2008). Training a binary classifier with the quantum adiabatic algorithm.

- O'Neill, J., Pleydell-Bouverie, B., Dupret, D., and Csicsvari, J. (2010). Play it again: Reactivation of waking experience and memory. *Trends in neurosciences*, 33:220–9.
- Paparo, G. D., Dunjko, V., Makmal, A., Martin-Delgado, M. A., and Briegel, H. J. (2014). Quantum speedup for active learning agents. *Physical Review X*, 4(3).
- Peng, J. and Williams, R. J. (1994). Incremental multi-step q-learning. In Cohen, W. W. and Hirsh, H., editors, *Machine Learning Proceedings 1994*, pages 226 – 232, San Francisco (CA). Morgan Kaufmann.
- Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition.
- Rebentrost, P., Mohseni, M., and Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13).
- Rummery, G. and Niranjan, M. (1994). On-line q-learning using connectionist systems. *Technical Report CUED/F-INFENG/TR 166*.
- Sallans, B. and Hinton, G. E. (2004). Reinforcement learning with factored states and actions. *J. Mach. Learn. Res.*, 5:1063–1088.
- Silver, D., Huang, A., Maddison, C., Guez, A., Sifre, L., Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–489.
- Strehl, A. L., Li, L., Wiewiora, E., Langford, J., and Littman, M. L. (2006). Pac model-free reinforcement learning. In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, pages 881–888, New York, NY, USA. ACM.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA.
- Tan, M. (1993). Multi-agent reinforcement learning: Independent vs. cooperative agents. In *In Proceedings of the Tenth International Conference on Machine Learning*, pages 330–337. Morgan Kaufmann.
- Wiebe, N., Braun, D., and Lloyd, S. (2012). Quantum algorithm for data fitting. *Physical Review Letters*, 109.