

Quantum error correction for the toric code using deep reinforcement learning

Philip Andreasson, Joel Johansson, Simon Liljestrand, and Mats Granath

Department of Physics, University of Gothenburg, SE-41296 Gothenburg, Sweden

We implement a quantum error correction algorithm for bit-flip errors on the topological toric code using deep reinforcement learning. An action-value Q-function encodes the discounted value of moving a defect to a neighboring site on the square grid (the action) depending on the full set of defects on the torus (the syndrome or state). The Q-function is represented by a deep convolutional neural network. Using the translational invariance on the torus allows for viewing each defect from a central perspective which significantly simplifies the state space representation independently of the number of defect pairs. The training is done using experience replay, where data from the algorithm being played out is stored and used for mini-batch upgrade of the Q-network. We find performance which is close to, and for small error rates asymptotically equivalent to, that achieved by the Minimum Weight Perfect Matching algorithm for code distances up to $d = 7$. Our results show that it is possible for a self-trained agent without supervision or support algorithms to find a decoding scheme that performs on par with hand-made algorithms, opening up for future machine engineered decoders for more general error models and error correcting codes.

1 Introduction

Much of the spectacular advances in machine learning using artificial neural networks has been in the domain of supervised learning where deep convolutional networks excel at categorizing objects when trained with big annotated data sets[1–3]. A different but also more challenging type of problem is when there is no a priori solution key, but rather a dynamic environment through which we want to learn to navigate for an optimal outcome. For these types of problems reinforcement learning (RL) [4] combined with deep learning has had great

Mats Granath: mats.granath@physics.gu.se

success recently when applied to problems such as computer and board games[5–8]. The super-human performance achieved by deep reinforcement learning has revolutionized the field of artificial intelligence and opens up for applications in many areas of science and technology.

In physics the use of machine learning has seen a great deal of interest lately[9–13]. The most natural type of application of neural networks is in the form of supervised learning where the deep network can capture correlations or subtle information in real or artificial data. The use of deep reinforcement learning may be less obvious in general as the type of topics addressed by RL typically involve some sort of "intelligent" best strategy search, contrary to the deterministic or statistical models used in physics.

In this paper we study a type of problem where artificial intelligence is applicable, namely the task of finding a best strategy for error correction of a topological quantum code; the potential basic building blocks of a quantum computer. In the field of quantum computing, smart algorithms are needed for error correction of fragile quantum bits[14–17]. Reinforcement learning has been suggested recently as a tool for quantum error correction and quantum control[18–20], where an agent learns to manipulate a quantum device that functions in an imperfect environment and with incomplete information. Under the umbrella term "Quantum Machine Learning" there are also interesting prospects of utilizing the natural parallelization of quantum computers for machine learning itself[21], but we will be dealing here with the task of putting (classical) deep learning and AI at the service of quantum computing.

Due to the inherently fragile nature of quantum information a future universal quantum computer will require quantum error correction.[14–17] Perhaps the most promising framework is to use topological error correcting codes.[22–26] Here, logical qubits consisting of a large set of entangled physical qubits are protected against local disturbances from phase or bit flip errors as logical operations require global changes. Local stabilizer operators, in the form of parity checks on a group of physical qubits, provide a quantum non-demolition

diagnosis of the logical qubit in terms of violated stabilizers; the so-called syndrome. In order for errors not to proliferate and cause logical failure, a decoder, that provides a set of recovery operations for correction of errors given a particular syndrome, is required. As the syndrome does not uniquely determine the physical errors, the decoder has to incorporate the statistics of errors corresponding to any given syndrome. In addition the syndrome itself may be imperfect, due to stabilizer measurement errors, in which case the decoder must also take that into account.

In the present work we consider Kitaev’s toric code[22, 23, 25] which is a stabilizer code formulated on a square lattice with periodic boundary conditions (see Figure 1 and Section 2.1). We will only consider bit-flip errors which correspond to syndromes with one type of violated stabilizer that can be represented as plaquette defects on the lattice (see Figure 2). The standard decoder for the toric code is the Minimum Weight Perfect Matching (MWPM) or Blossom algorithm[27–29] that works by finding the pairwise matching of syndrome defects with shortest total distance, corresponding to the minimal number of errors consistent with the syndrome. The decoder problem is also conceptually well suited for reinforcement learning, similar in spirit to a board game; the state of the system is given by the syndrome, actions correspond to moving defects of the syndrome, and with reward given depending on the outcome of the game. By playing the game, the agent improves its error correcting strategies and the decoder is the trained agent that provides step by step error correction. As in any RL problem the reward scheme is crucial for good performance. The size of the state-action space is also a challenge, to provide the best action for each of a myriad syndromes, but this is exactly the problem addressed by recent deep learning approaches to RL.[5–8]

We find that by setting up a reward scheme that encourages the elimination of the syndrome in as few operations as possible within the deep Q-learning (or deep Q-network, DQN)[6, 7] formalism we are able to train a decoder that is comparable in performance to MWPM. Although the present algorithm does not outperform the latter we expect that it has the potential to be more versatile when addressing depolarizing noise (with correlated bit and phase flip errors), measurement noise giving imperfect syndromes, or varying code geometries. Compared to the MWPM algorithm the RL algorithm also has the possible advantage that it provides step by step correction whereas the MWPM algorithm only provides information on which defects should be paired, making the former more adaptable to the introduction of additional errors.

In concurrent work by Sweke et al. [30] an application of reinforcement learning to error correction of the

toric code was implemented. That work focuses on the important issue of imperfect syndromes as well as depolarizing noise and used an auxiliary ”referee decoder” to assist the RL decoder. In the present work we consider the simpler but conceptually more direct problem of error correction on a perfect syndrome and with only bit flip error. Also in contrast to [30] we study the actual ”toric” code, rather than the code with boundaries. Clearly the toric code will be harder to implement experimentally but nevertheless provides a well understood standard model. It also provides a simplification from the fact that on a torus only the relative positions of syndrome defects are relevant which reduces the state space complexity that decoder agent has to master. By focusing on this minimal problem we find that we can make a rigorous benchmark on the RL decoder showing near optimal performance.

Finding better performing decoders has been the topic of many studies, using methods such as renormalization group[31, 32], cellular automata[33, 34], and a number of neural network based decoders[19, 35–43]. The decoder presented in this paper does not outperform state of the art decoders, it’s value lies in showing that it is possible to use reinforcement learning to achieve excellent performance on a minimal model. Given that deep reinforcement learning is arguably the most promising AI framework it holds prospect for future versatile self-trained decoders that can adapt to different error scenarios and code architectures.

The outline of the paper is the following. In the *Background* section we give a brief but self-contained summary of the main features of the toric code including the basic structure of the error correction and a similar summary of one-step Q-learning and deep Q-learning. The following section, *RL Algorithm*, describes the formulation and training of the error correcting agent. In the *Results* section we shows that we have trained the RL agent up to code distance $d = 7$ with performance which is very close to the MWPM algorithm. We finally conclude and append details of the asymptotic fail rate for small error rates as well as the neural network architecture and the RL and network hyperparameters.

2 Background

2.1 Toric code

The basic construction of the toric code is a square lattice with a spin- $\frac{1}{2}$ degree of freedom on every bond, the physical qubits, and with periodic boundary conditions, as seen in Figure 1.[22, 23]¹ (An alternative rotated lattice representation with the qubits on sites is also com-

¹Figures in this section were inspired by lecture notes [44].




Figure 1: A $d = 5$ toric code with rings indicating the physical qubits and grey showing the periodic boundary conditions. a) Plaquette (green) and vertex (red) stabilizer operators, as products of σ^z and σ^x Pauli matrices. b) A single vertex operator can be represented as a loop flipping the qubits that the it crosses. c) Two neighboring vertex operators make up a larger loop. d) The logical operators $\bar{X}_{1/2}$ (red) and $\bar{Z}_{1/2}$ (green) consist of loops winding the torus and are not representable in terms of products of vertex or plaquette operators.

mon in the literature.) The model is given in terms of a Hamiltonian

$$H = - \sum_{\alpha} \hat{P}_{\alpha} - \sum_{\nu} \hat{V}_{\nu}, \quad (1)$$

where α runs over all plaquettes and ν over all vertices (sites). The stabilizers are the plaquette operators $\hat{P}_{\alpha} = \prod_{i \in \alpha} \sigma_i^z$ and the vertex operators $\hat{V}_{\nu} = \prod_{i \in \nu} \sigma_i^x$, where σ^z and σ^x are the Pauli matrices. (Where, in the σ^z basis, $\sigma^z |\uparrow / \downarrow\rangle = \pm 1 |\uparrow / \downarrow\rangle$ and $\sigma^x |\uparrow / \downarrow\rangle = |\downarrow / \uparrow\rangle$.) The stabilizers commute with each other and the Hamiltonian thus block diagonalizing the latter. On a $d \times d$ lattice of plaquettes $d^2 - 1$ plaquette operators are linearly independent (e.g. it is not possible to have a single -1 eigenvalue with all other $+1$) and correspondingly for the vertex operators. With $2d^2$ physical qubits and $2d^2 - 2$ stabilizers the size of each block is $2^{2d^2} / 2^{2d^2 - 2} = 4$, corresponding in particular to a ground state which is 4-fold degenerate. These are the states that will serve as the logical qubits. (More precisely, given the 4-fold degeneracy it is a qudit or base-4 qubit.)

To derive the ground state consider first the plaquette operator in the σ^z -basis; clearly a ground state must have an even number of each spin-up and spin-down on every plaquette to be a $+1$ eigenstate of each plaquette

operator. Let's consider the state with all spin-up $|\uparrow\uparrow\uparrow \dots\rangle$; acting with a vertex operator on this flips all the spins around the vertex (see Fig. 1b) giving a state still in ground state sector of the plaquette operators as an even number of spins are flipped on the plaquettes surrounding the vertex. (As is also clear from the fact that all the stabilizer operators commute.) The $+1$ eigenstate of that particular vertex operator is thus the symmetric superposition of the two states. A convenient way to express the operation of one or several adjacent vertex operators is in terms of loop traversing the flipped spins. Such loops (fig. 1b-c) generated from products of vertex operators will always be topologically trivial loops on the surface of the torus since they are just constructed by merging the local loop corresponding to a single vertex operator. Successively acting with vertex operators on the states generated from the original $|\uparrow\uparrow\uparrow \dots\rangle$ we realize that the ground state is simply the symmetric superposition of all states that are generated from this by acting with (trivial) loops $|\text{GS}_0\rangle = \sum_{i \in \text{all trivial loops}} \text{loop}_i |\uparrow\uparrow\uparrow \dots\rangle$.

To generate the other ground states we consider the operators \bar{X}_1 and \bar{X}_2 (Fig. 1d) which are products of σ^x corresponding to the two non-trivial loops that wind the torus. (Deformations of these loops just correspond to multiplication by trivial loops and is thus inconsequential.) Correspondingly there are non-trivial loops of σ^z operators \bar{Z}_1 and \bar{Z}_2 . The four ground states are thus the topologically distinct states $\{|\text{GS}_0\rangle, \bar{X}_1|\text{GS}_0\rangle, \bar{X}_2|\text{GS}_0\rangle, \bar{X}_2\bar{X}_1|\text{GS}_0\rangle\}$ distinguished by their eigenvalues of \bar{Z}_1 and \bar{Z}_2 being ± 1 . For a torus with $d \times d$ plaquettes there are $2d^2$ physical qubits and the code distance, i.e. minimum length of any logical operator (\bar{X}_i or \bar{Z}_i), is d .

2.1.1 Error correction

Errors in the physical qubits will take the state out of the ground state sector and thereby mask the encoded state. The task of the error correction procedure is to move the system back to the ground state sector without inadvertently performing a logical operation to change the logical qubit state. A σ^x error on a physical qubit corresponds to a bit-flip error. On the toric code this gives rise to a pair of defects (a.k.a. quasiparticles or anyons) in the form of neighboring plaquettes with -1 eigenvalues of the plaquette stabilizers. Similarly a σ^z error corresponds to a phase-flip error which gives rise to a pair of neighboring -1 defects on two vertices. A $\sigma^y = i\sigma^x\sigma^z$ simultaneously creates both types of defects. A natural error process is to assume that X, Y, Z errors occur with equal probability, so called depolarizing noise. This however requires to treat correlations between X and Z errors and the simpler uncorrelated




Figure 2: Bit-flip errors (red 'X') and possible error correction bit-flips (blue 'X'). (a) Two neighboring errors and the corresponding error chain (red line) and syndrome (red dots). (b) Visualized in terms of the syndrome with error chain and two possible correction chains (blue) as expressed explicitly in (c) and (d). The error chain plus the correction chain in (d) constitutes a non-trivial loop and a logical bit-flip operation (as in Figure 1d), thus a failed error correction, in contrast to the trivial loop in (c).

noise model is often used, which is what we will consider in this work, focusing on bit-flip errors and corresponding plaquette defects. Here X and Z errors occur independently with probability p whereas $Y = XZ$ errors occur with probability p^2 . Correcting independent X and Z errors is completely equivalent (with defects either on plaquettes or on vertices) and it is therefore sufficient to formulate an error correcting algorithm for one type of error. (For actual realizations of the physical qubits the error process may in fact be intermediate between these two cases[45].) Regardless of noise model and type of error an important aspect of the error correction of a stabilizer formalism is that the entanglement of the logical qubit states or its excitations does not have to be considered explicitly as errors act equivalently on all states that belong to the same stabilizer sector.

A crucial aspect of quantum error correction is that the actual bit-flip errors cannot be measured without collapsing the state into a partial basis and destroying the qubit. What can be measured without destroying the logical qubit are the stabilizers, i.e. for bit-flip error the parity of the plaquette operators. The complete set of incorrect (-1) plaquettes makes up the syndrome of the state. The complete set of bit-flip errors will pro-

duce a unique syndrome as the end-points of strings of bit-flip errors. The converse however is not true, which is what makes the task challenging. In order to do the error correction we need to suggest a number of physical bits that should be flipped in order to achieve the pair-wise annihilation of the defects of the syndrome. Consider a single pair of defects which have been created by a particular chain of errors. (See Figure 2.) The error correction needs to suggest a correction string connecting the two defects. If this is done properly the correction string and the error string form a trivial loop, thus returning the qubit to the original state. If instead the correction string and the error string together make up a non-trivial loop that winds the torus we have eliminated the error syndrome but changed the state of qubit (corresponding to a logical bit-flip), thus failed the task of correcting the error.

For the uncorrelated noise model it can be shown, by mapping to the random bond Ising model, that for $d \rightarrow \infty$ there is a critical threshold $p_c \approx 0.11$ below which the most probable correction chains to complement the error chain will with certainty form trivial loops, while for $p > p_c$ non-trivial loops occur with finite probability.[23] For a finite system, the sharp transition is replaced by a cross-over, as seen in Figure 6, where for increasing d the fraction of successful error correction evolves progressively towards 1 for $p < p_c$, and to $1/4$ (thus completely unpredictable) for $p > p_c$.

For the uncorrelated noise model on the torus the most likely set of error chains between pairs of defects which is consistent with a given syndrome would be one that corresponds to the smallest number of total bit flips, i.e. the shortest total error chain length. Thus, a close to optimal algorithm for error correction for this system is the Minimum Weight Perfect Matching (MWPM) algorithm[27]. (This algorithm is also near optimal for the problem with syndrome errors as long as it is still uncorrelated noise[23, 28].) The MWPM algorithm for the perfect syndrome corresponds to reducing a fully connected graph, with an even number of nodes and with edges specified by the inter-node distances, to the set of pairs of nodes that minimize the total edge length. This algorithm can be implemented efficiently[46] and we will use this as the benchmark of our RL results. In fact, as we will see, the RL algorithm that we formulate amounts to solving the MWPM problem. In this sense the work presented in this paper is to show the viability of the RL approach to this problem with the aim for future generalizations to other problems where MWPM is sub-optimal, such as for depolarizing noise or more general error models.

2.2 Q-learning

Reinforcement learning is a method to solve the problem of finding an optimal policy of an agent acting in a system where the actions of the agent causes transitions between states of the system.[4] The policy $\pi(s, a)$ of an agent describes (probabilistically perhaps) the action a to be taken by the agent when the system is in state s . In our case the state will correspond to a syndrome, and an action to moving a defect one step. The optimal policy is the one that gives the agent maximal return (cumulative discounted reward) over the course of its interaction with the system. Reward r_{t+1} is given when the system transitions from state $s_t \rightarrow s_{t+1}$ such that the return starting at time t is given by $R_t = r_{t+1} + \gamma r_{t+2} + \gamma^2 r_{t+3} + \dots$. Here $\gamma \leq 1$ is the discounting factor that quantifies how we want to value immediate versus subsequent reward. As will be discussed in more detail, in the work presented in this paper a constant reward $r = -1$ will be given for each step taken, so that in practice the optimal policy will be the one that minimizes the number of actions, irrespectively of the value of γ . (Although in practice, even here the value of γ can be important for the convergence of the training.)

One way to represent the cumulative reward depending on a set of actions and corresponding transitions is by means of an action-value function, or Q-function. This function $Q(s, a)$ quantifies the expected return when in state s taking the action a , and subsequently following some policy π . In one-step Q-learning we quantify Q according to $Q(s, a) = r + \gamma \max_{a'} Q(s', a')$, with $s \xrightarrow{a} s'$, which corresponds to following the optimal policy according to our current estimate of Q . In order to learn the value of the Q-function for all states and actions we should explore the full state-action space, with the policy given by taken action a according to $\max_a Q(s, a)$ eventually guaranteed to converge to the optimal policy. However, an unbiased exploration gets prohibitively expensive and it is therefore in general efficient to follow an ϵ -greedy policy which with probability $(1 - \epsilon)$ takes the optimal action based on our current estimate of $Q(s, a)$ but with probability ϵ takes a random action. From what we have learned by this action we would update our estimate for Q according to

$$Q(s, a) \leftarrow Q(s, a) + \alpha [(r + \gamma \max_{a'} Q(s', a')) - Q(s, a)], \quad (2)$$

where $\alpha < 1$ is a learning rate. This procedure is then a trade-off between using our current knowledge of the Q function as a guide for the best move to avoid spending extensive time on expensive moves but also exploring to avoid missing out on rewarding parts of the state-action space.




Figure 3: Structure of the deep Q-network. The input layer is a $d \times d$ matrix corresponding to the "perspective" P , of one defect of the syndrome. (Using translational symmetry on the torus, any defect can be placed at the center.) The output layer gives the action value $Q(P, a, \theta)$ of moving the central defect to any of the four neighboring plaquettes $a = U, D, R, L$, given the current training state of network parameters θ . The hidden layers consist of a convolutional layer (of which a 3×3 filter is indicated on the input layer) and several fully connected layers. (For details, see Appendix.) Successively scanning all defects using the same network gives the full action value function of the syndrome.

2.2.1 Deep Q-learning

For a large state-action space it is not possible to store the complete action-value function. (Disregarding symmetries, for a $d \times d$ system with N_S defects, the state space has size $\binom{d^2}{N_S}$, $\sim 10^{13}$ for $p \approx 10\%$ and $d = 7$.) In deep Q-learning[7], the action-value function is instead represented by a deep neural network with the input layer corresponding to some representation of a state and the output layer corresponding to the value of the possible actions. The idea is that similarities in the value of different regions of the state-action space may be stored in an efficient way by the deep network. Parametrizing the Q-function by means of neural network we write $Q(s, a, \theta)$, where θ represents the complete set of weights and biases of the network. (We use a convolutional network with $\sim 10^6$ parameters for the $d = 7$ problem.) As outlined in more detail in the following sections the latter can be trained using supervised learning based on a scheme similar to one step Q-learning.

3 RL Algorithm

The decoder presented in this paper is a neural network-based agent optimized using reinforcement learning to observe toric code syndromes and suggesting recovery chains for them step by step. The agent makes use of a deep convolutional neural network, or Q-network, (see Fig. 3) to approximate Q values of actions given a syndrome.

In a decoding session, a syndrome S corresponding to the coordinates of N_S defects e_i ($i = 1, \dots, N_S$) is fed to the algorithm as input. The syndrome is the state

of the system as visible to the agent. The syndrome at any time step is that generated by accumulated actions of the agent on the syndrome given by the initial random distribution of bit-flips. There is also a hidden state corresponding to the joint set of initial and agent flipped qubits. After the complete episode resulting in a terminal state with an empty syndrome, an odd number of non-trivial loops (in either X_1 or X_2) indicates a failed error correction. In the algorithm used in this work however, the success/fail information does not play any explicit role in the training, except as external verification of the performance of the agent. Instead reward $r = -1$ is given at every step until the terminal state regardless of whether the error correcting string(s) gave rise to an unwanted logical operation. Taking the fewest number of steps to clear the syndrome is thus the explicit target of the agent, corresponding to actuating the MWPM algorithm. (An alternative formulation with different dependence on γ would be to reward +1 at the terminal step.)

It would seem very natural to base the RL reward scheme on the success/failure information from the hidden state. However, we found it difficult to converge to a good agent based on this, for the following reason: given a particular starting syndrome, consistent with a distribution of different error strings, most of these are properly corrected by the MWPM algorithm whereas a minority are not. As the syndrome is all that the agent sees, it has no chance to learn to distinguish between these two classes, thus trying to use it for training will only obscure the signal. Nevertheless, for future more advanced tasks, such as dealing with noise biased towards bit or phase flips or with spatial variations it will probably be necessary to explore the use of the fail/success information for the reward scheme.

3.1 State-space formulation

Due to the periodic boundary conditions of the code, the syndrome can be represented with an arbitrary plaquette as its center. Centering a defect e_i , we define the *perspective*, P_i , of that defect, consisting of the relative positions of all other defects in the syndrome. The set of all perspectives given a syndrome we define as an *observation*, O , as exemplified in Figure 4. (The syndrome, observation and perspective all contain equivalent information but represented differently.)

The agent will be given the option of moving any defect one plaquette in any direction (left, right, up, or down), corresponding to performing a bit flip on one of the physical qubits enclosing the plaquette containing the defect. Clearly the total number of available actions varies with the number of defects, which is inconvenient if we want to represent the Q-function in




Figure 4: State formulation. The toric code syndrome, defines an "observation" that contains the centralized "perspectives" for each defect.

terms of a neural network. In order for the Q network to have a constant-sized output regardless of how many defects are present in the system, each perspective in the observation is instead sent individually to the Q network. Thus, $Q(P, a, \theta)$ represents the value of moving the central defect $a = L, R, U, D$, given the positions of all other defects specified by the perspective P , for network parameters θ . The network with input and output is represented graphically in Figure 3. The full Q-function corresponding to a syndrome is given by $\{Q(P, a, \theta)\}_{P \in O}$. When the Q value of each action for each defect has been obtained, the choice of action and defect is determined by a greedy policy. The new syndrome is sent to the algorithm and the procedure is repeated until no defects remain.

3.2 Training the neural network

Training of the decoder agent was done using the Deep Q Network (DQN) algorithm [7]. This algorithm utilizes the technique of experience replay in which the experience acquired by the agent is stored as transition tuples in a memory buffer. When updating the Q network (given by parameters θ), a mini-batch of random samples is drawn from this memory buffer. By taking random samples of experience, the temporal correlation of the data is minimized, resulting in a more stable training procedure of the neural network. To further increase the stability of the training, the DQN algorithm makes use of a target Q network (with parameters θ_T) to compute update targets. The target Q network is periodically synchronized with the updated Q network.

A training sequence begins with an acting stage, where a syndrome is sent to the agent, which uses the Q network θ to suggest a defect perspective, P , and an action, a . An ϵ -greedy policy is used by the agent, meaning that it will suggest the action with the highest Q-value with probability $(1 - \epsilon)$. Otherwise a random action is suggested. The action is performed on the

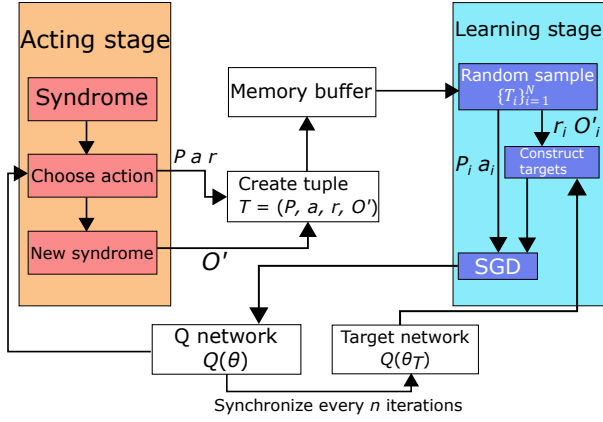


Figure 5: Flowchart of the training procedure. A learning sequence consists of an acting stage followed by a learning stage. (P, a, r) is (perspective, action, reward) corresponding to what defect is acted on with what action and giving what reward. O' is the observation (see Fig. 4) corresponding to the resulting syndrome. A pseudocode representation is given in Algorithm 1. (SGD is Stochastic Gradient Descent.)

defect, e , corresponding to P , resulting in a reward, r , and a new observation, O' , derived from the resulting syndrome. The whole transition is stored as a tuple, $T = (P, a, r, O')$, in a memory buffer. After this, the training sequence enters the learning stage using (mini-batch) stochastic gradient descent. First, a random sample of transitions, $\{T_i = (P_i, a_i, r_i, O'_i)\}_{i=1}^N$, of a given batch size, N , is drawn with replacement from the memory buffer. (Here the discrete C_4 rotational symmetry of the problem is enforced by including all four rotated versions of the same tuple.) The training target value for the Q-network is given by

$$y_i = r_i + \gamma \max_{P' \in O'_i, a'} Q(P', a', \theta_T), \quad (3)$$

where γ is the discount factor and where the more slowly evolving target network parametrized by θ_T is used to predict future cumulative award. After this, gradient descent is used to minimize the discrepancy between the targets of the sample and the Q network predictions for it, upgrading the network parameters schematically according to $-\nabla_{\theta} \sum_i (y_i - Q(P_i, a_i, \theta))^2$. A new training sequence is then started, and with some specified rate, the weights of the target network, θ_T , are synchronized with the Q network θ . A pseudocode description of the procedure is presented in algorithm 1 and an illustration of the different components and procedures of the training algorithm and how they relate to each other is found in Figure 5.

Algorithm 1 Training the reinforcement learning agent decoder

- 1: **while** syndrome defects remain **do**
 - 2: Get observation O from syndrome \triangleright See figure 4
 - 3: Calculate $Q(P, a, \theta)$ using Q-network for all perspectives $P \in O$.
 - 4: Choose which defect e to move with action a using ϵ -greedy policy
 - 5: $P \leftarrow$ perspective of defect e
 - 6: Perform action a on defect e
 - 7: $r \leftarrow$ reward from taking action a on defect e
 - 8: $O' \leftarrow$ observation corresponding to new syndrome
 - 9: Store transition tuple $T = (P, a, r, O')$ in memory buffer
 - 10: Draw a random sample of transition tuples
 - 11: **for** each transition tuple T_i in sample **do**
 - 12: Construct targets y_i using target network θ_T and reward r_i according to Eqn. 3.
 - 13: **end for**
 - 14: Update Q-network parameters θ
 - 15: Every n iterations, synchronize the target network with network, setting $\theta_T = \theta$
 - 16: **end while**
-

4 Result

Data sets with a fixed error rate of 10% were generated to train the agent to operate on a code of a specified size. The syndromes in a data set is fed one at a time to the agent, which operates on it until no errors remain. The data sets also contain information about the physical qubit configuration (the hidden state) of the lattice, which (as discussed in section 3) is used to check the success rate of the decoder. This is compared to the performance of the MWPM decoder on the same syndromes [46]. The operation of the trained decoder is similar to the cellular automaton decoders [33, 34] in the sense of providing step by step actions based on the current state of the syndrome. This also means that it could be implemented in parallel with the error generation process by continuously adapting to the introduction of new errors.

The proficiency of the well converged agents are shown in figures 6 and 7 as compared to the MWPM performance. Given our specified reward scheme, which corresponds to using as few operations as possible, we achieve near optimal results with a performance which is close to that of the MWPM decoder. For small error rates $p_L \rightarrow 0$ it is possible to derive an exact expression for the MWPM fail rate p_L (see Appendix A and [25, 47]) by explicitly identifying the dominant type




Figure 6: Error correction success rate p_s of the converged agents versus bit-flip error rate p , for system size $d = 3, 5, 7$, and compared to the corresponding results using MWPM (lines). (The MWPM decoder for $d = 30$ is included as a reference for the approach to large d .)

of error string. We have checked explicitly that our Q-network agent is equivalent to MWPM for these error strings and thus gives the same asymptotic performance.

For larger system size $d = 9$ we have only been partially successful, with good performance for small error rates, but sub-MWPM performance for larger error rates. Given the exponential growth of the state space this is perhaps not surprising, but by scaling up the hardware and the corresponding size of the manageable Q-network we anticipate that larger code distances would be achievable within the present formalism.

As a demonstration of the operation of the trained agent and the corresponding Q-network we present in Figure 8 the action values $Q(S, a)$ for two different syndromes. (As discussed previously, $Q(S, a) = \{Q(P, a, \theta)\}_{P \in O}$, where O is the observation, or set of perspectives, corresponding to the syndrome S .) The size of the arrows are proportional to the discounted return R of moving a defect one initial step in the direction of the arrow and then following the optimal policy. In Fig. 8a, the values are written out explicitly. The best (equivalent) moves have a return $R = -3.57$ which corresponds well to the correct value $R = -1 - \gamma - \gamma^2 - \gamma^3 = -3.62$ for following the optimal policy to annihilate the defects in four steps, with reward $r = -1$ and discount rate $\gamma = .95$. Figure 8b shows a seemingly challenging syndrome where the fact that the best move does not correspond to annihilating the two neighboring defects is correctly captured by the

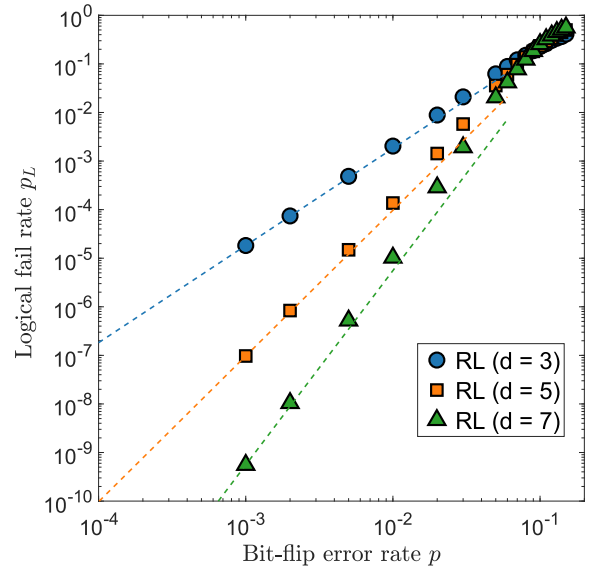


Figure 7: Error correction fail rate $p_L = 1 - p_s$ shown to converge to the known asymptotic MWPM behavior (Appendix A) for small error rates $p \rightarrow 0$. The lines correspond to $p_L \sim p^x$, with $x = \lceil d/2 \rceil = 2, 3, 4$ for $d = 3, 5, 7$ fitted to the lowest p data point.

Q-network.

One interesting aspect of the close to MWPM performance of the fully trained agent is the ability of the Q-network to suggest good actions independently of how many defects are in the syndrome. A $d = 7$ system with $p = 10\%$ would start out with a syndrome with maybe 20 defects, which is successively pair-wise reduced down to two and finally zero defects, all based on action-values given by the same Q-network (θ). The network is thus surprisingly versatile and capable, given the enormous reduction of the number of adjustable parameters compared to representing and training the full Q-value function as a table.

5 Conclusions

In conclusion, we have shown how to implement deep reinforcement learning for quantum error correction on the toric code for moderate size systems using uncorrelated bit-flip (or phase-flip) noise. By training an agent to find the shortest paths for the error correction chains we are able to achieve accuracy close to that using a Minimum Weight Perfect Matching decoder. In order to accomplish this we used the deep Q-network formalism that encodes the action-value function by means of an artificial neural network.[6, 7] The construction also made good use of the translational invariance on the torus to be able to efficiently reduce the state space representation.




Figure 8: Action value function produced by the Q-network for two different syndromes and code distance $d = 7$. The magnitude of the arrows indicate the expected return from taking a next step along the arrow and after that following the optimal policy. The optimal policy for the next move corresponds to the action with the biggest arrow(s). In (a) the expected return is written out explicitly, where the best moves are consistent with the constant reward of -1 per step and discounting rate $\gamma = 0.95$ used in the training.

For future work it will be interesting to see how the formalism generalizes to more advanced noise models, imperfect measurements, as well as more general topological codes. Work in progress[48] indicates that the formalism is in fact readily extended to handle depolarizing noise on the toric code by allowing for the full set of X , Y , and Z qubit actions. By learning to account for correlations between plaquette and vertex errors super-MWPM performance can be achieved. Also using larger and better adapted convolutional networks allow for somewhat larger system sizes to be addressed. Nevertheless, given the exponential growth of the state-action space it is clear that going much beyond the code distances presented in this paper will require parallelization of the training[49] as well as massive networks using state of the art hardware, similarly to what is used to achieve super-human performance for board games and computer games.[7, 8]

In the longer perspective the main potential of a deep reinforcement learning approach to quantum error correction lies in the fact that is arguably the most promising implementation of AI. Future developments in that area thus opens up also for powerful and flexible machine engineered quantum decoders.

Acknowledgements

We thank Niklas Forsström, Gustav Karlsson, and Elias Hannouch for contributing to the early stages of this work. We also thank Austin Fowler for valuable discussions. Source code can be found at this url: <https://github.com/phiandre/ToricCodeRL>

References

- [1] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems 25*, pages 1097–1105, 2012.
- [2] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436, 2015. DOI: [10.1038/nature14539](https://doi.org/10.1038/nature14539).
- [3] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [4] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [5] Gerald Tesauro. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3):58–68, 1995. URL <https://link.galegroup.com/apps/doc/A16764437/AONE?u=googlescholar&sid=AONE&xid=f888cd62>.
- [6] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013. URL <https://arxiv.org/abs/1312.5602>.
- [7] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015. DOI: [10.1038/nature14236](https://doi.org/10.1038/nature14236).
- [8] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017. DOI: [10.1038/nature24270](https://doi.org/10.1038/nature24270).
- [9] Louis-François Arsenault, Alejandro Lopez-Bezanilla, O Anatole von Lilienfeld, and Andrew J Millis. Machine learning for many-body physics: the case of the anderson impurity model. *Physical Review B*, 90(15):155136, 2014. DOI: [10.1103/PhysRevB.90.155136](https://doi.org/10.1103/PhysRevB.90.155136).
- [10] Evert PL Van Nieuwenburg, Ye-Hua Liu, and Sebastian D Huber. Learning phase transitions by confusion. *Nature Physics*, 13(5):435, 2017. DOI: [10.1038/nphys4037](https://doi.org/10.1038/nphys4037).
- [11] Juan Carrasquilla and Roger G Melko. Machine learning phases of matter. *Nature Physics*, 13(5):431, 2017. DOI: [10.1038/nphys4035](https://doi.org/10.1038/nphys4035).
- [12] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural

- networks. *Science*, 355(6325):602–606, 2017. DOI: [10.1126/science.aag2302](https://doi.org/10.1126/science.aag2302).
- [13] Xun Gao and Lu-Ming Duan. Efficient representation of quantum many-body states with deep neural networks. *Nature communications*, 8(1):662, 2017. DOI: [10.1038/s41467-017-00705-2](https://doi.org/10.1038/s41467-017-00705-2).
- [14] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A*, 52:R2493–R2496, Oct 1995. DOI: [10.1103/PhysRevA.52.R2493](https://doi.org/10.1103/PhysRevA.52.R2493).
- [15] A. M. Steane. Error correcting codes in quantum theory. *Phys. Rev. Lett.*, 77:793–797, Jul 1996. DOI: [10.1103/PhysRevLett.77.793](https://doi.org/10.1103/PhysRevLett.77.793).
- [16] Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.
- [17] Barbara M Terhal. Quantum error correction for quantum memories. *Reviews of Modern Physics*, 87(2):307, 2015. DOI: [10.1103/RevModPhys.87.307](https://doi.org/10.1103/RevModPhys.87.307).
- [18] Alexey A Melnikov, Hendrik Poulsen Nautrup, Mario Krenn, Vedran Dunjko, Markus Tiersch, Anton Zeilinger, and Hans J Briegel. Active learning machine learns to create new quantum experiments. *Proceedings of the National Academy of Sciences*, 115(6):1221–1226, 2018. DOI: [10.1073/pnas.1714936115](https://doi.org/10.1073/pnas.1714936115).
- [19] Thomas Fösel, Petru Tighineanu, Talitha Weiss, and Florian Marquardt. Reinforcement learning with neural networks for quantum feedback. *Phys. Rev. X*, 8:031084, Sep 2018. DOI: [10.1103/PhysRevX.8.031084](https://doi.org/10.1103/PhysRevX.8.031084).
- [20] Marin Bukov, Alexandre G. R. Day, Dries Sels, Phillip Weinberg, Anatoli Polkovnikov, and Pankaj Mehta. Reinforcement learning in different phases of quantum control. *Phys. Rev. X*, 8:031086, Sep 2018. DOI: [10.1103/PhysRevX.8.031086](https://doi.org/10.1103/PhysRevX.8.031086).
- [21] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195, 2017. DOI: [10.1038/nature23474](https://doi.org/10.1038/nature23474).
- [22] A Yu Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003. DOI: [10.1016/S0003-4916\(02\)00018-0](https://doi.org/10.1016/S0003-4916(02)00018-0).
- [23] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43(9):4452–4505, 2002. DOI: [10.1063/1.1499754](https://doi.org/10.1063/1.1499754).
- [24] Robert Raussendorf, Jim Harrington, and Kovid Goyal. Topological fault-tolerance in cluster state quantum computation. *New Journal of Physics*, 9(6):199, 2007. DOI: [10.1088/1367-2630/9/6/199](https://doi.org/10.1088/1367-2630/9/6/199).
- [25] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, 2012. DOI: [10.1103/PhysRevA.86.032324](https://doi.org/10.1103/PhysRevA.86.032324).
- [26] Julian Kelly, Rami Barends, Austin G Fowler, Anthony Megrant, Evan Jeffrey, Theodore C White, Daniel Sank, Josh Y Mutus, Brooks Campbell, Yu Chen, et al. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature*, 519(7541):66, 2015. DOI: [10.1038/nature14270](https://doi.org/10.1038/nature14270).
- [27] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17(3):449–467, 1965. DOI: [10.4153/CJM-1965-045-4](https://doi.org/10.4153/CJM-1965-045-4).
- [28] Austin G Fowler. Minimum weight perfect matching of fault-tolerant topological quantum error correction in average $\mathcal{O}(1)$ parallel time. *Quantum Information and Computation*, 15(1&2):0145–0158, 2015. URL <http://dl.acm.org/citation.cfm?id=2685188.2685197>.
- [29] Sergey Bravyi, Martin Suchara, and Alexander Vargo. Efficient algorithms for maximum likelihood decoding in the surface code. *Phys. Rev. A*, 90:032326, Sep 2014. DOI: [10.1103/PhysRevA.90.032326](https://doi.org/10.1103/PhysRevA.90.032326).
- [30] Ryan Sweke, Markus S Kesselring, Evert PL van Nieuwenburg, and Jens Eisert. Reinforcement learning decoders for fault-tolerant quantum computation. *arXiv preprint arXiv:1810.07207*, 2018. URL <https://arxiv.org/abs/1810.07207>.
- [31] Guillaume Duclos-Cianci and David Poulin. Fast decoders for topological quantum codes. *Physical review letters*, 104(5):050504, 2010. DOI: [10.1103/PhysRevLett.104.050504](https://doi.org/10.1103/PhysRevLett.104.050504).
- [32] Guillaume Duclos-Cianci and David Poulin. Fault-tolerant renormalization group decoder for abelian topological codes. *Quantum Info. Comput.*, 14(9&10):721–740, July 2014. ISSN 1533-7146. URL <http://dl.acm.org/citation.cfm?id=2638670.2638671>.
- [33] Michael Herold, Earl T Campbell, Jens Eisert, and Michael J Kastoryano. Cellular-automaton decoders for topological quantum memories. *npj Quantum Information*, 1:15010, 2015. DOI: [10.1038/npjqi.2015.10](https://doi.org/10.1038/npjqi.2015.10).
- [34] Aleksander Kubica and John Preskill. Cellular-automaton decoders with provable thresholds for topological codes. *arXiv preprint arXiv:1809.10145*, 2018. URL <https://arxiv.org/abs/1809.10145>.
- [35] Giacomo Torlai and Roger G. Melko. Neural decoder for topological codes. *Phys. Rev. Lett.*, 119:030501, Jul 2017. DOI: [10.1103/PhysRevLett.119.030501](https://doi.org/10.1103/PhysRevLett.119.030501).
- [36] Stefan Krastanov and Liang Jiang. Deep neural network probabilistic decoder for stabilizer

- codes. *Scientific reports*, 7(1):11003, 2017. DOI: [10.1038/s41598-017-11266-1](https://doi.org/10.1038/s41598-017-11266-1).
- [37] Savvas Varsamopoulos, Ben Criger, and Koen Bertels. Decoding small surface codes with feedforward neural networks. *Quantum Science and Technology*, 3(1):015004, 2017. DOI: [10.1088/2058-9565/aa955a](https://doi.org/10.1088/2058-9565/aa955a).
- [38] Paul Baireuther, Thomas E O’Brien, Brian Tarasinski, and Carlo WJ Beenakker. Machine-learning-assisted correction of correlated qubit errors in a topological code. *Quantum*, 2:48, 2018. DOI: [10.22331/q-2018-01-29-48](https://doi.org/10.22331/q-2018-01-29-48).
- [39] Nikolas P Breuckmann and Xiaotong Ni. Scalable neural network decoders for higher dimensional quantum codes. *Quantum*, 2:68, 2018. DOI: [10.22331/q-2018-05-24-68](https://doi.org/10.22331/q-2018-05-24-68).
- [40] Christopher Chamberland and Pooya Ronagh. Deep neural decoders for near term fault-tolerant experiments. *Quantum Sci. Technol.*, 3:044002, 2018. DOI: [10.1088/2058-9565/aad1f7](https://doi.org/10.1088/2058-9565/aad1f7).
- [41] Nishad Maskara, Aleksander Kubica, and Tomas Jochym-O’Connor. Advantages of versatile neural-network decoding for topological codes. *Phys. Rev. A*, 99:052351, May 2019. DOI: [10.1103/PhysRevA.99.052351](https://doi.org/10.1103/PhysRevA.99.052351).
- [42] Xiaotong Ni. Neural network decoders for large-distance 2d toric codes. *arXiv preprint arXiv:1809.06640*, 2018. URL <https://arxiv.org/abs/1809.06640>.
- [43] Ye-Hua Liu and David Poulin. Neural belief-propagation decoders for quantum error-correcting codes. *Phys. Rev. Lett.*, 122:200501, May 2019. DOI: [10.1103/PhysRevLett.122.200501](https://doi.org/10.1103/PhysRevLett.122.200501).
- [44] Dan Browne. Topological codes and computation a lecture course given at the university of innsbruck. 2014. URL <http://bit.do/topological>.
- [45] David K. Tuckett, Stephen D. Bartlett, and Steven T. Flammia. Ultrahigh error threshold for surface codes with biased noise. *Phys. Rev. Lett.*, 120:050505, Jan 2018. DOI: [10.1103/PhysRevLett.120.050505](https://doi.org/10.1103/PhysRevLett.120.050505).
- [46] Vladimir Kolmogorov. Blossom v: a new implementation of a minimum cost perfect matching algorithm. *Mathematical Programming Computation*, 1(1):43–67, 2009. DOI: [10.1007/s12532-009-0002-8](https://doi.org/10.1007/s12532-009-0002-8).
- [47] Austin G Fowler. Optimal complexity correction of correlated errors in the surface code. *arXiv preprint arXiv:1310.0863*, 2013. URL <https://arxiv.org/abs/1310.0863>.
- [48] Mattias Eliasson, David Fitzek, and Mats Granath. *In preperation*, 2019.
- [49] Dan Horgan, John Quan, David Budden, Gabriel Barth-Marón, Matteo Hessel, Hado Van Hasselt,

and David Silver. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018. URL <https://arxiv.org/abs/1803.00933>.

A Small error rate

As discussed by Fowler *et al.*[25, 47] the likely operating regime of surface code is in the limit of small error rate $p \ll 1$. In addition, in the limit $p \rightarrow 0$ we can derive an exact expression for the rate of logical failure under the assumption of MWPM error correction, thus providing a solid benchmark for our RL algorithm. Such expressions were derived for the surface code in [47] and here we derive the corresponding expression for bit-flip errors in the toric code.

Consider first the case of code distance d with $d \in \text{odd}$, which is what we have assumed in the present work. (Using odd d gives an additional simplification of the Q-learning set-up from the fact that any plaquette can be considered the center of the lattice.) As a reminder, the error formulation we use is that every physical qubit has a probability p of bit-flip error, and probability $1 - p$ of no error. (In contrast to [47] we don’t consider σ^y errors, which would give rise to both bit-flip and phase-flip errors.) For very low p , we only need consider states with the minimal number of bit-flip errors that may cause a logical failure. One can readily be convinced (from a few examples) that such states are ones where a number $\lceil d/2 \rceil$ (e.g. $\lceil 7/2 \rceil = 4$) of errors are placed along the path of the shortest possible non-trivial (logical) loops. The latter are d sites long, and on the torus there are $2d$ such loops. For such a state MWPM will always fail, because it will provide a correction string which has $\lceil d/2 \rceil$ bit-flips rather than the $\lceil d/2 \rceil$ flips needed to make a successful error correction. The former correction string, together with the initial error string, will sum to one of the non-trivial (shortest length) loops and give rise to a logical bit-flip. The fail-rate p_L , i.e. the fraction of logical fails of all generated syndromes, is thus to lowest order in p and for odd d given by

$$p_L = 2d \binom{d}{\lceil d/2 \rceil} p^{\lceil d/2 \rceil}. \quad (4)$$

Here $2d$ is the number of shortest non-trivial loops, $\binom{d}{\lceil d/2 \rceil}$ is the number of ways of placing the errors on such a loop, and $p^{\lceil d/2 \rceil}$ is the lowest order term in the probability $(p^{\lceil d/2 \rceil} (1 - p)^{2d - \lceil d/2 \rceil})$ of any particular state with $\lceil d/2 \rceil$ errors.

Considering d even (for reference), the corresponding minimal fail scenario has $d/2$ errors on a length d loop. Here the MWPM has a 50% chance of constructing either a non-trivial or trivial loop, thus giving the asymptotic fail rate $p_L = d \binom{d}{d/2} p^{d/2}$.

B Network architecture and training parameters

The reinforcement learning agent makes use of a deep convolutional neural network to approximate the Q values for the possible actions of each defect. The network (see Fig. 3) consists of an input layer which is $d \times d$ matrix corresponding to a perspective (binary input, 0 or 1, with 1 corresponding to a defect), and a convolutional layer followed by several fully-connected layers and an output layer consisting of four neurons, representing each of the four possible actions. All layers have ReLU activation functions except the output layer which has simple linear activation.

Table 1: Network architecture $d=5$. FC=Fully connected

#	Type	Size	# parameters
0	Input	5x5	
1	Conv.	512 filters; 3x3 size; 2-2 stride	5 120
2	FC	256 neurons	524 544
3	FC	128 neurons	32 896
4	FC	64 neurons	8 256
5	FC	32 neurons	2 080
6	FC (out)	4 neurons	132
			573 028

Table 2: Network architecture $d=7$.

#	Type	Size	# parameters
0	Input	7x7	
1	Conv.	512 filters; 3x3 size; 2-2 stride	5 120
2	FC	256 neurons	1 179 904
3	FC	128 neurons	32 896
4	FC	64 neurons	8 256
5	FC	32 neurons	2 080
6	FC (out)	4 neurons	132
			1 228 388

The network architecture is summarized in Table 1 and 2. We also included explicitly a count of the number of parameters (weights and biases) to emphasize the huge reduction compared to tabulating the Q-function. The latter requires of the order $\binom{d^2}{N_S}$ entries, for N_S defects, where N_S will also vary as the syndrome is reduced, with initially $N_S \sim 4pd^2$ as each isolated error creates a defect pair and there are $2d^2$ physical qubits.

In Figure 9 we also provide an example of the initial




Figure 9: Early training convergence of the Q network agent. Success rate P_s versus number of iterations. One iteration corresponds to annihilating all the defects of a single syndrome. (The very early below 1/4 success rate is an artifact of using a max count for the number of error correcting steps for the validation.)

convergence of the algorithm for lattice size $d \times d$, with $d = 3, 5, 7$. Here, each iteration corresponds to solving one syndrome and making the corresponding number of mini-batch training sessions from the experience buffer, as explained in section 3.2. A constant set of syndromes is used for the testing so that fluctuations correspond to actual performance variations of the agent.

In Table 3 we list the hyperparameters related to the Q-learning and experience replay set-up, as well as the neural network training algorithm used. The full RL algorithm is coded in Python using Tensorflow and Keras for the Q-network. A single desktop computer was used, with training converging over a matter of hours (for $d = 3$) to days (for $d = 7$).

Table 3: Hyperparameters

Parameter	Value
discount rate γ	0.95
reward r	-1/step; 0 at finish
exploration ϵ	0.1
max steps per syndrome	50
mini batch size, N	32
target network update rate	100
memory buffer size	1 000 000
optimizer	'Adam'
learning rate	0.001
beta ₁	0.9
beta ₂	0.999
decay	0.0