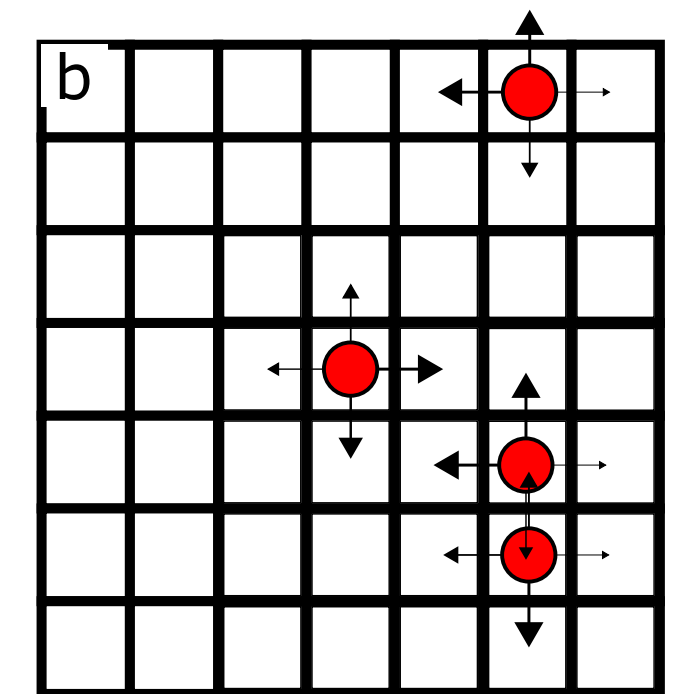
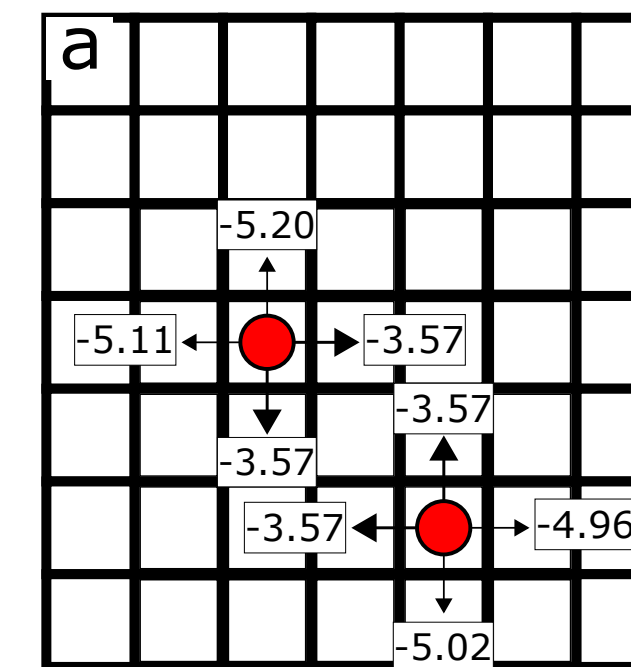
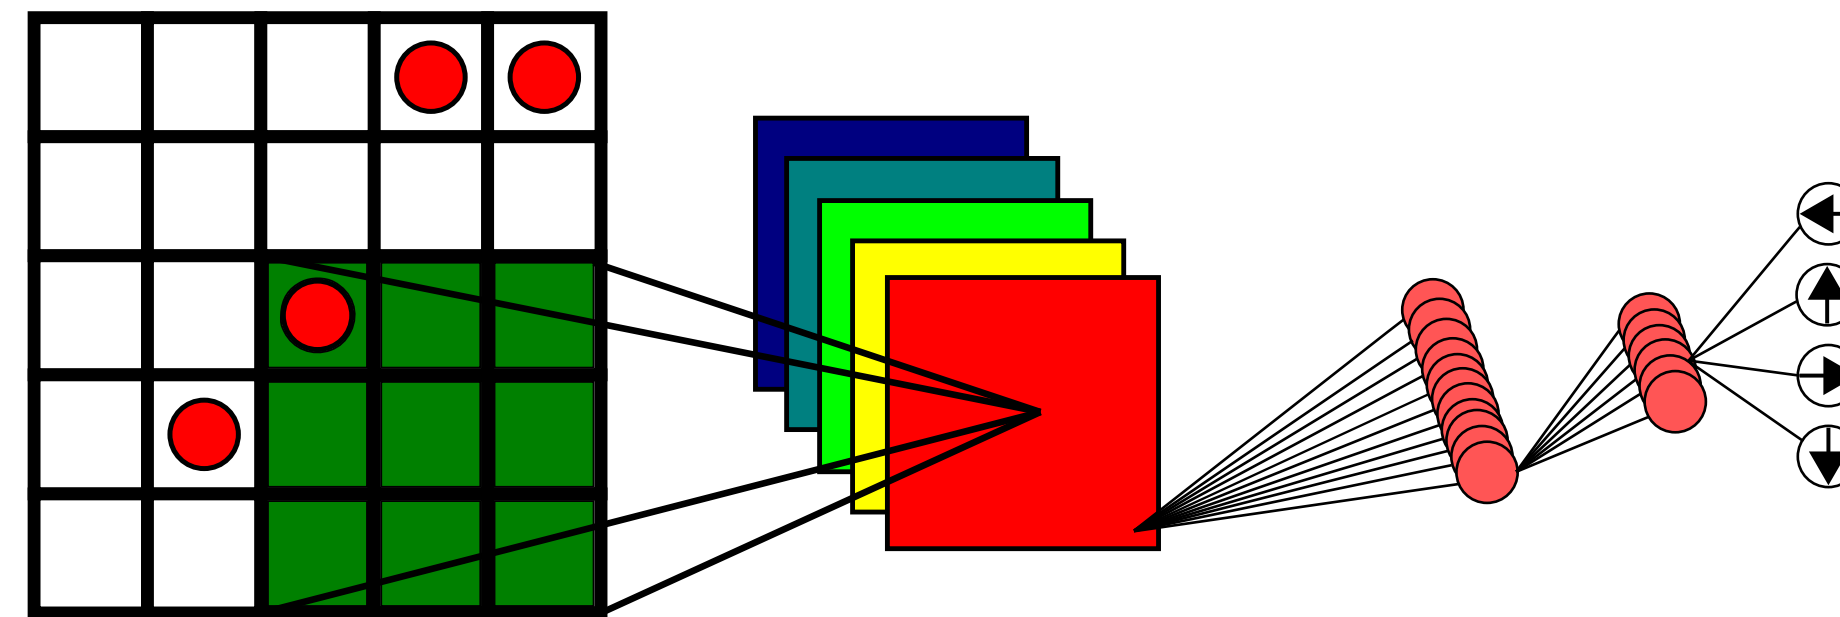
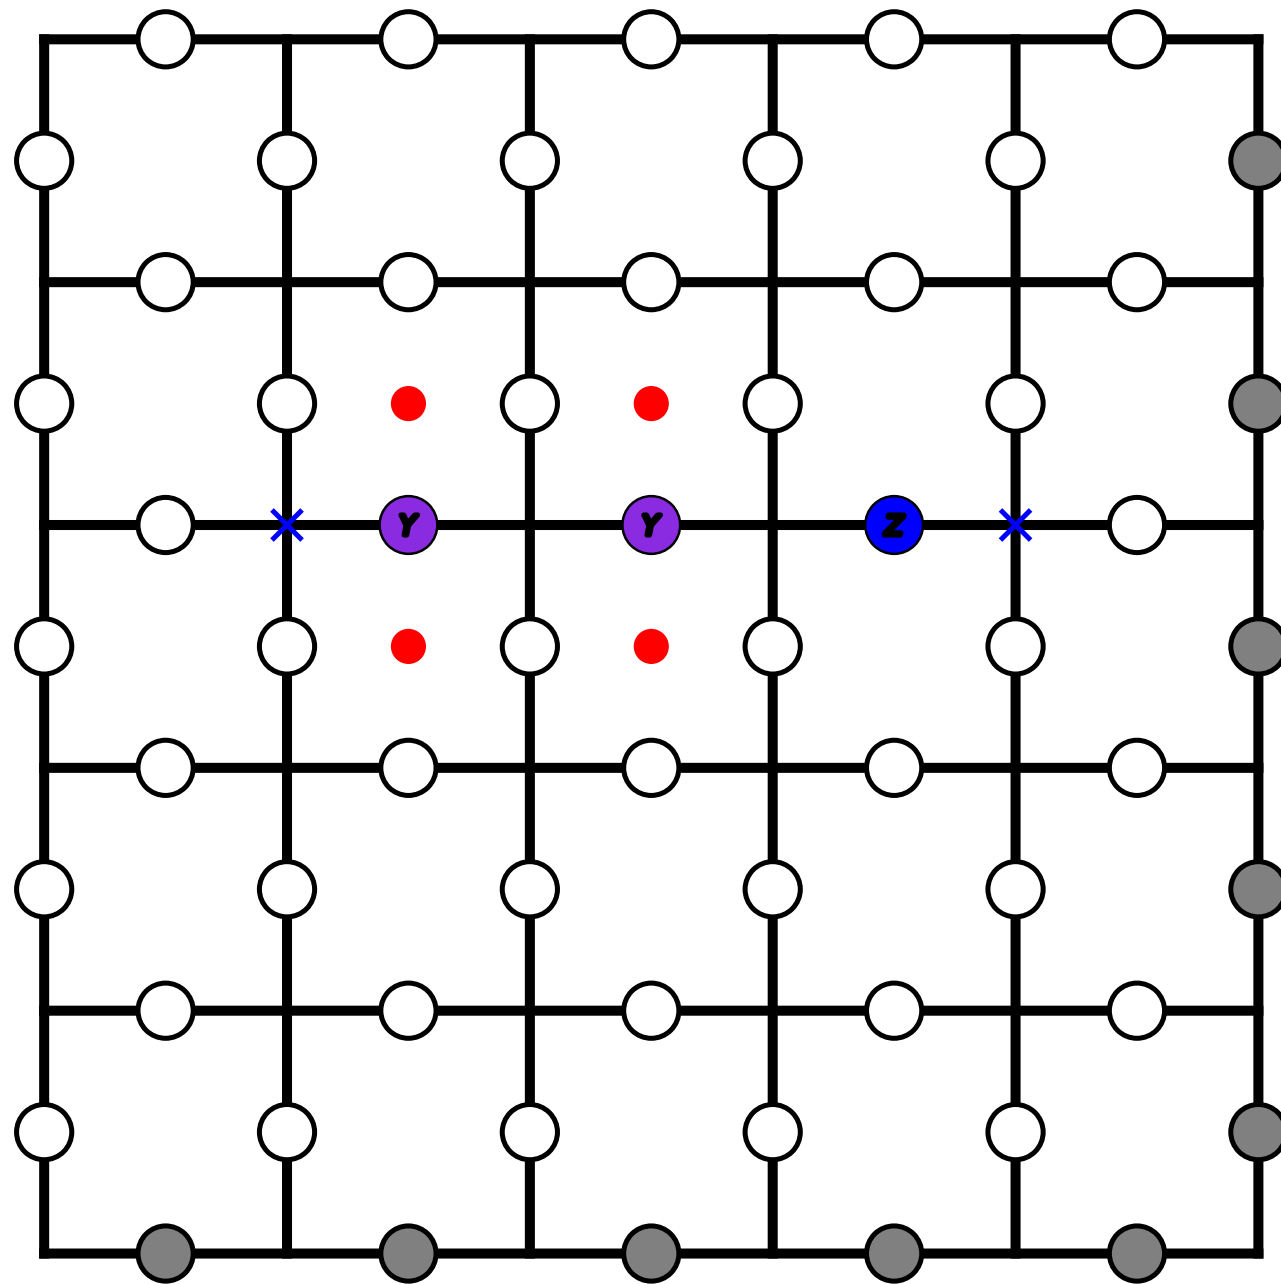


# Quantum error correction for the toric code using deep reinforcement learning

Mats Granath  
Department of Physics  
University of Gothenburg

*Machine Learning for Quantum Technology  
Max Planck Institute for the Science of Light  
May 9th, 2019*



Philip Andreasson, Joel Johansson, Simon Liljestr nd, MG, arXiv:1811.12338

Mattias Eliasson, David Fitzek, MG, in progress

## **Bottom line**

**We do the “simplest” error correction problem for a topological code**

- **Periodic boundary conditions**
- **No measurement noise/perfect syndrome**
- **only bit flip noise (initially, can also do depolarizing)**

**Still challenging for reinforcement learning: deep Q-networks needed**

**Allows for easy benchmark**

Fault-tolerant quantum computation by anyons

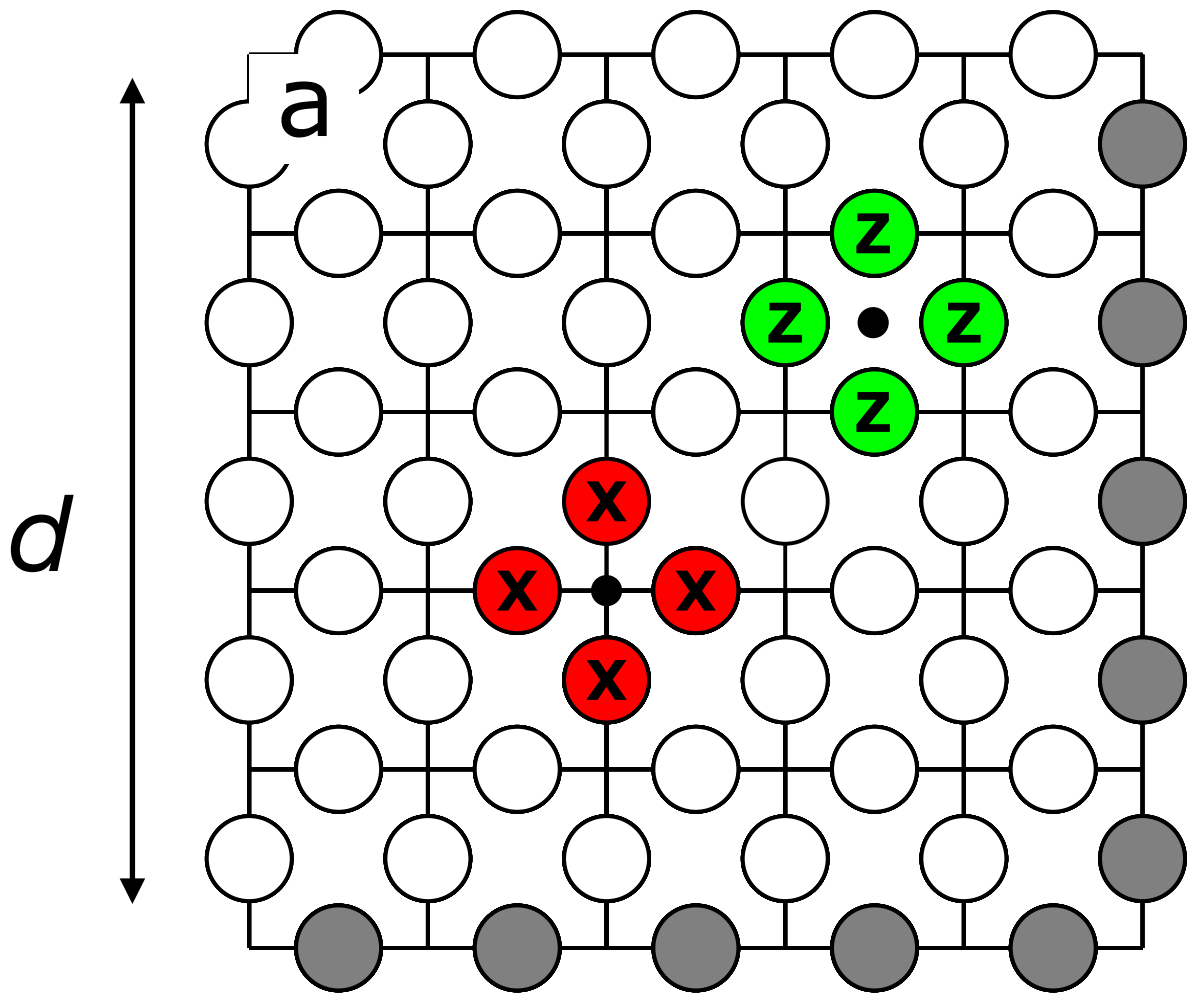
A.Yu. Kitaev\*

L.D. Landau Institute for Theoretical Physics, 117940, Kosygina St. 2, Germany

Received 20 May 2002

Topological quantum memory<sup>a)</sup>

Eric Dennis<sup>b)</sup>  
Princeton University, Princeton, New Jersey 08544  
Alexei Kitaev,<sup>c)</sup> Andrew Landahl,<sup>d)</sup> and John Preskill<sup>e)</sup>  
Institute for Quantum Information, California Institute of Technology,  
Pasadena, California 91125



$$H = - \left( \sum_{\alpha} \hat{P}_{\alpha} \right) - \left( \sum_{\nu} \hat{V}_{\nu} \right)$$

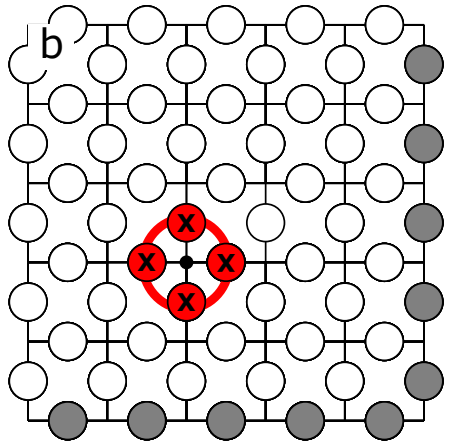
$$\hat{P}_{\alpha} = \prod_{i \in \alpha} \sigma_i^z$$
$$\hat{V}_{\nu} = \prod_{i \in \nu} \sigma_i^x$$

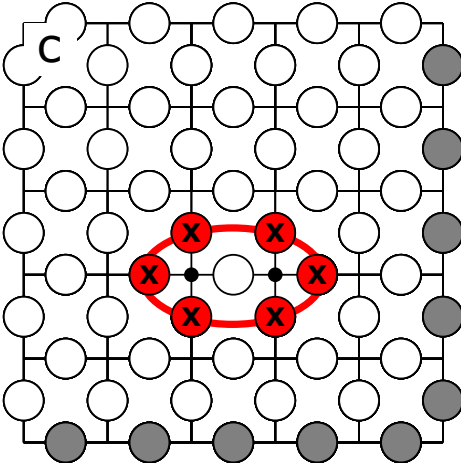
Plaque and Vertex stabilizers (parity checks)

$2d^2$  physical qubits,  $2d^2-2$  independent stabilizers

Ground state

consider:  
 $|\uparrow\uparrow\uparrow \dots\rangle$   
plaquette operator  
ground state:

act with vertex op:  
  
still a plaquette  
ground state

act with two vertex op:  
  
still a plaquette  
ground state

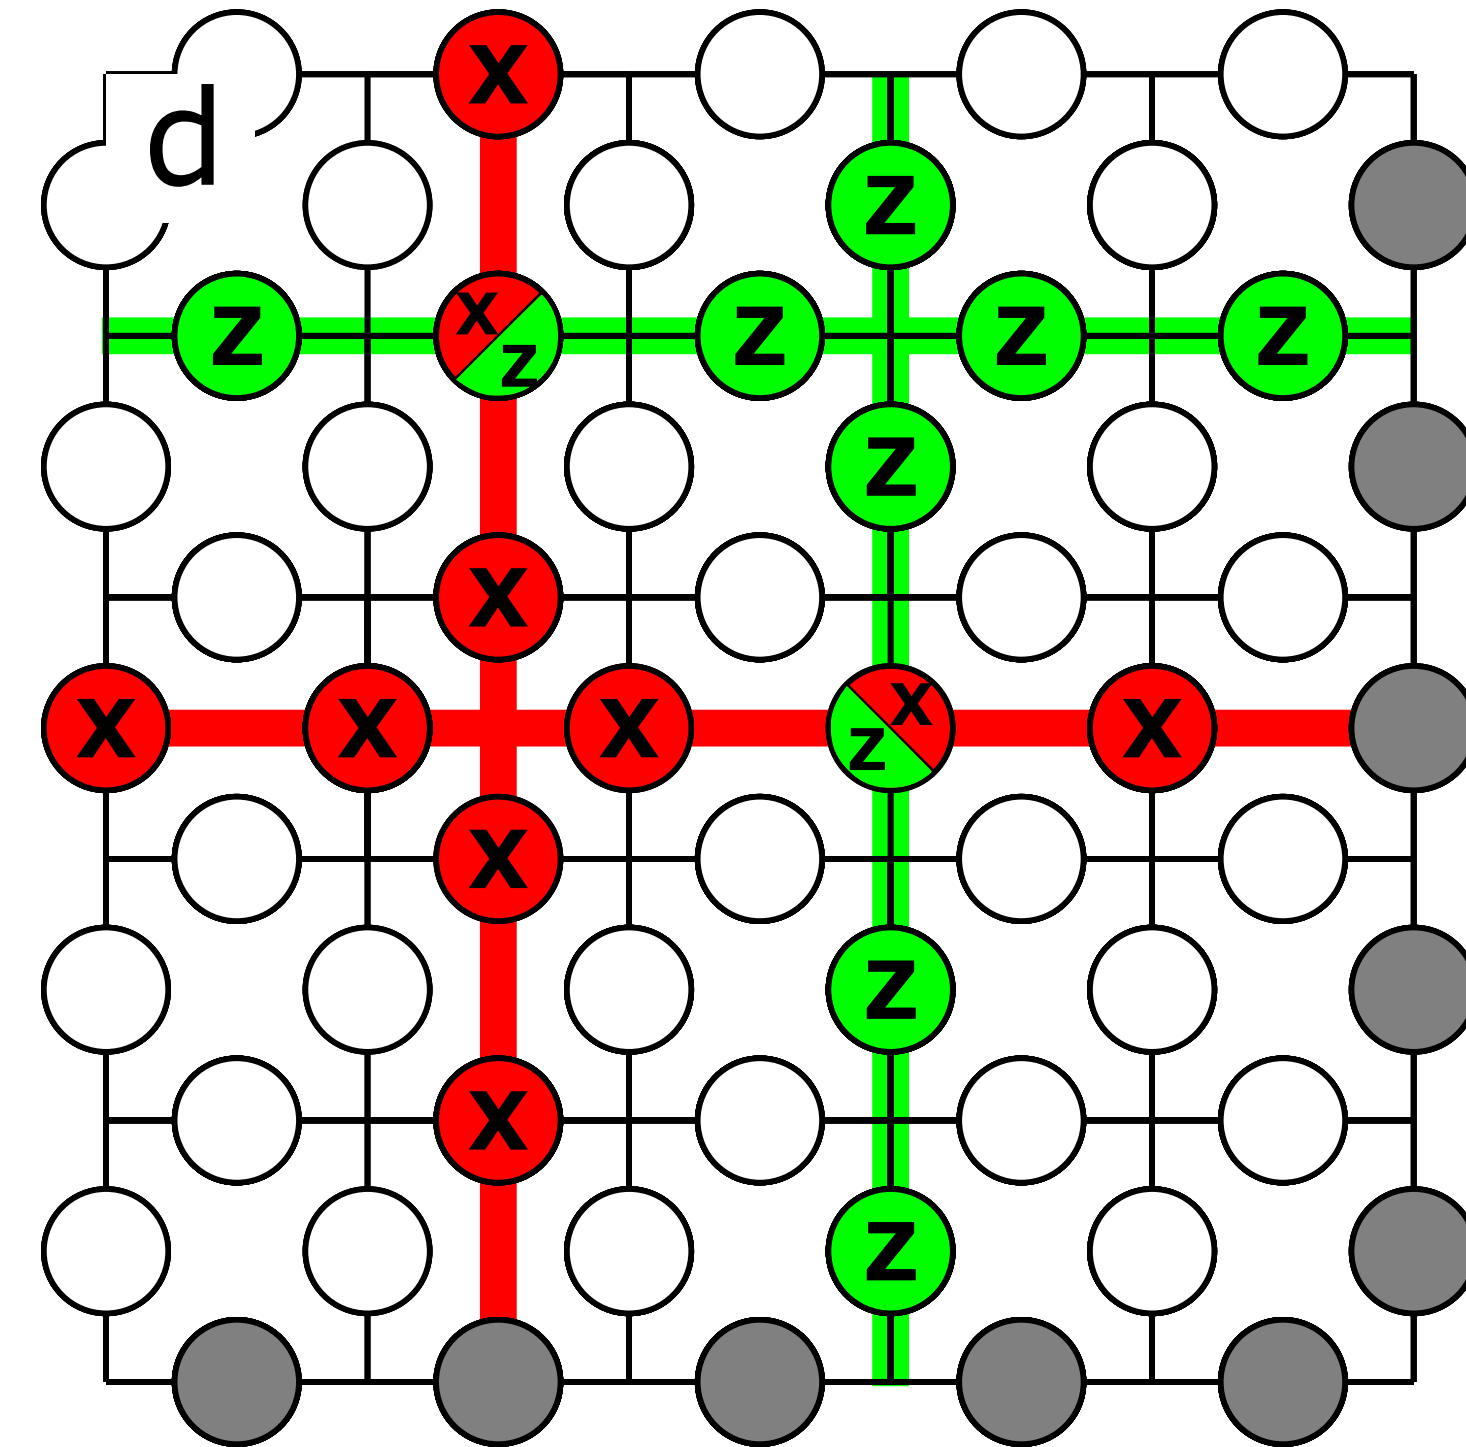
GS is symmetric superposition  
of all *trivial* loops:

$$|GS_0\rangle = \sum_{i \in \text{all trivial loops}} \text{loop}_i |\uparrow\uparrow\uparrow \dots\rangle$$

highly entangled

# Ground state degeneracy

Non-trivial loops (encircling torus)  $X_1, X_2$   
are not products of vertex operators.



Four ground states/The logical qubit

$$\{ |GS_0\rangle, X_1 |GS_0\rangle, X_2 |GS_0\rangle, X_2 X_1 |GS_0\rangle \}$$

Distinguished by  $\pm 1$  eigenvalues of  $Z_1$  and  $Z_2$ .

Corresponding to  $2(d^2-1)$  independent stabilizers on  $2d^2$  physical qubits.

## Topologically protected qubit

Non-trivial loops=Logical bit-flip operators  
Requires at least  $d$  physical bit-flip errors  
**code distance  $d$**

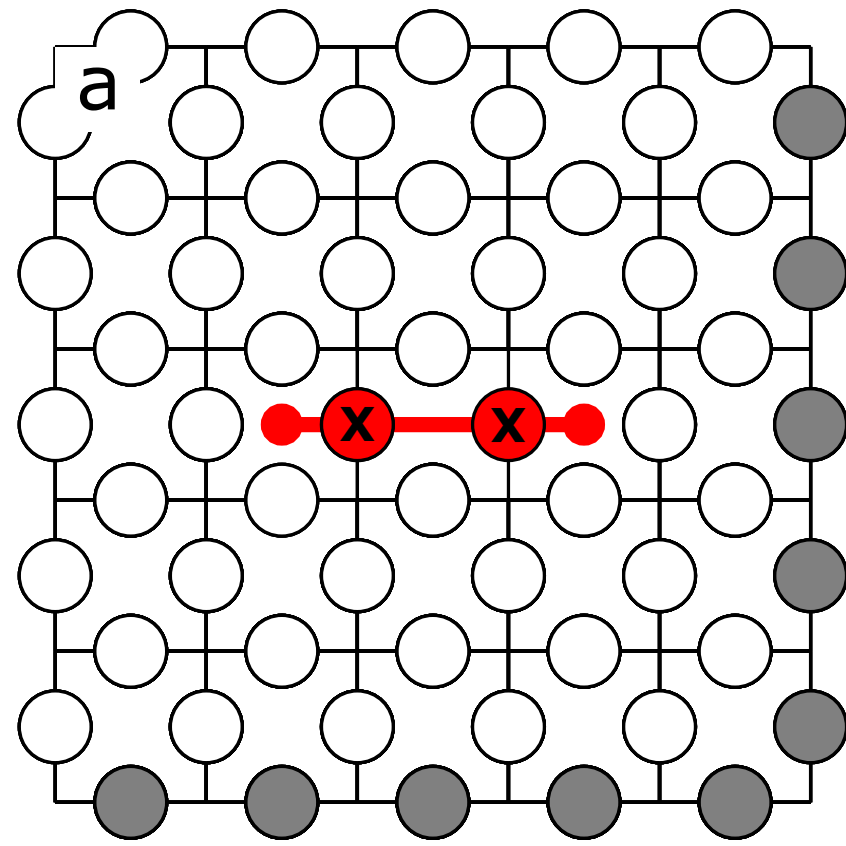
# Error correction

consider bit-flip errors

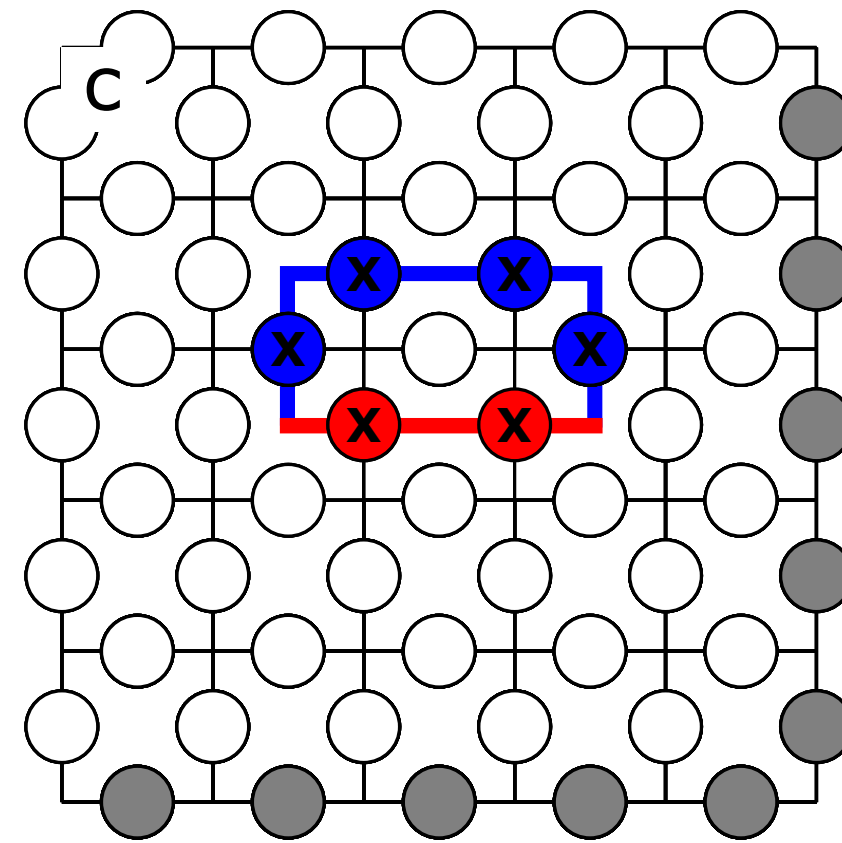
The *syndrome* (defects/bad plaquettes), is quantum non-demolition measurement

**Ex.**

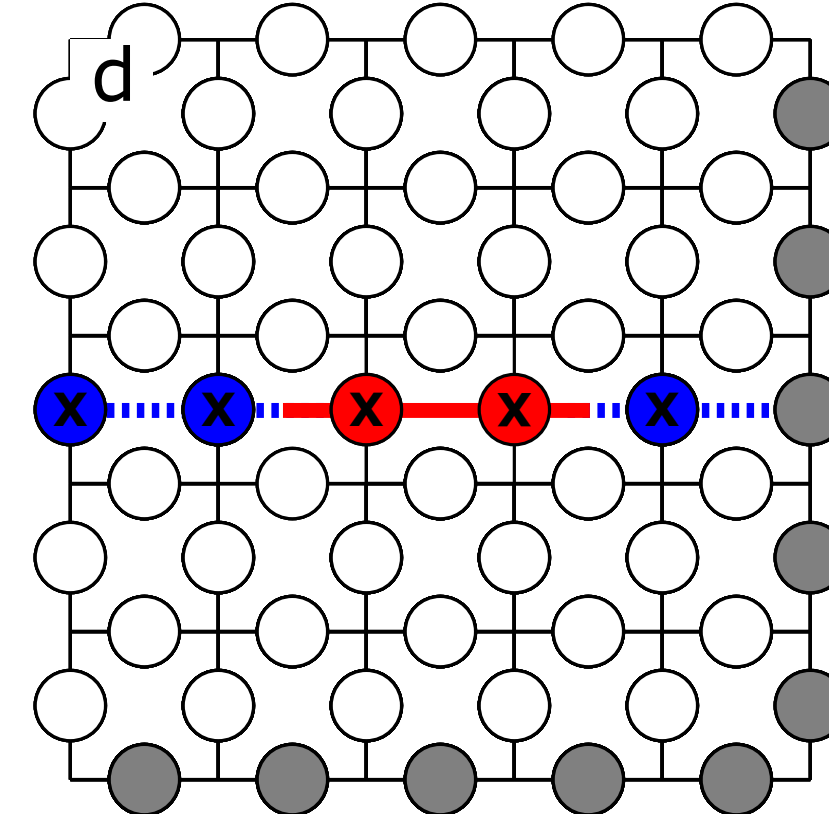
two neighbouring bit flip errors,  
two defects



**proper** error correction  
trivial loop



**failed** error correction  
non-trivial loop



**Standard algorithm to suggest error correcting strings:**  
Minimum Weight Perfect Matching (MWPM)/Blossom

J. Edmonds, 1965

Find shortest total correction string. (Which is the most likely)

# Error models

## Depolarizing

- $(1-p)$  no error
- $p/3$  X
- $p/3$  Y=XZ
- $p/3$  Z

## Uncorrelated

- $(1-p)^2$  no error
- $p(1-p)$  X
- $p^2$  Y=XZ
- $p(1-p)$  Z

Bit- and phase-flip errors  
(i.e. plaquette and vertex errors).  
are independent. Corrected separately.

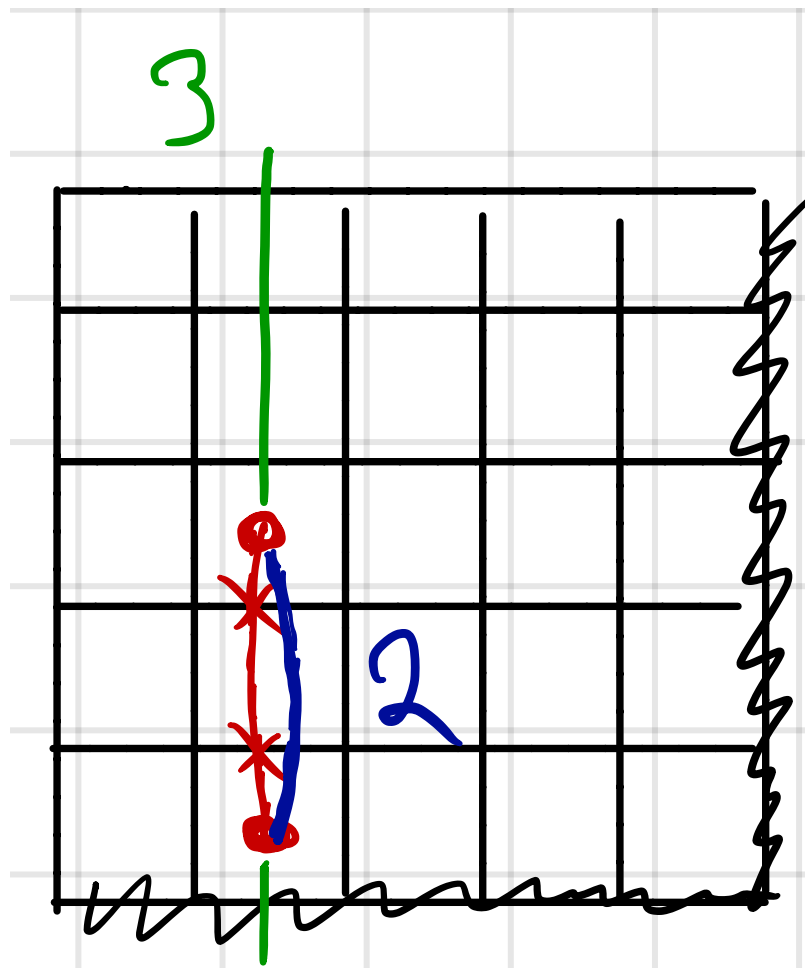
**MWPM is (near) optimal**

# Minimum Weight Perfect Matching Low-p fail rate for bit-flip errors

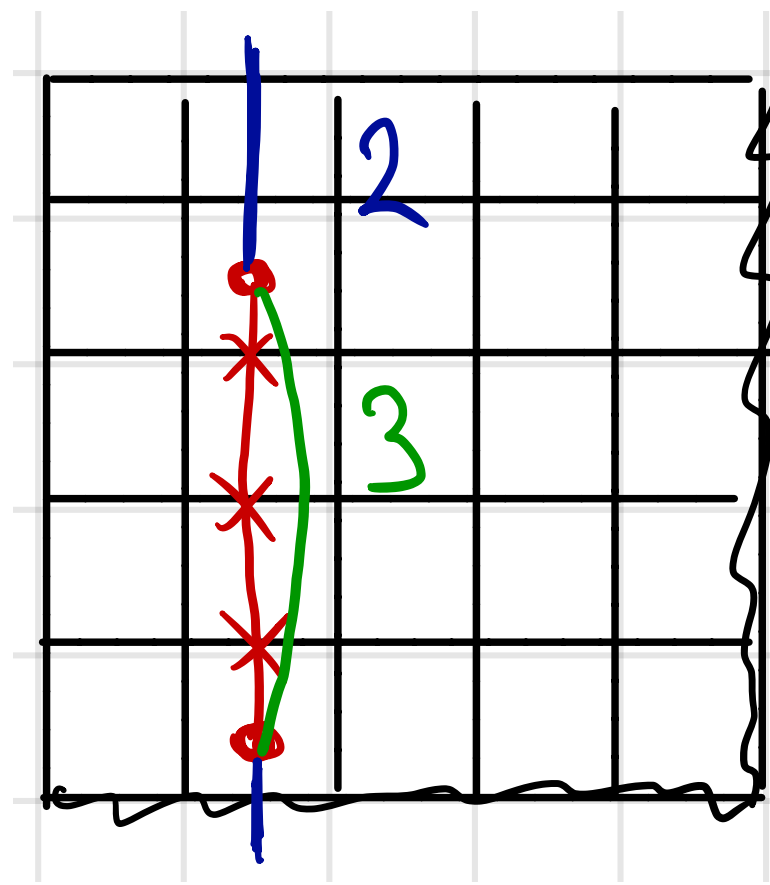
For  $p \rightarrow 0$  we only need to consider error chains with minimal number of errors that can give failed error correction

Consider  $d=5$ :

Two errors is always corrected successfully

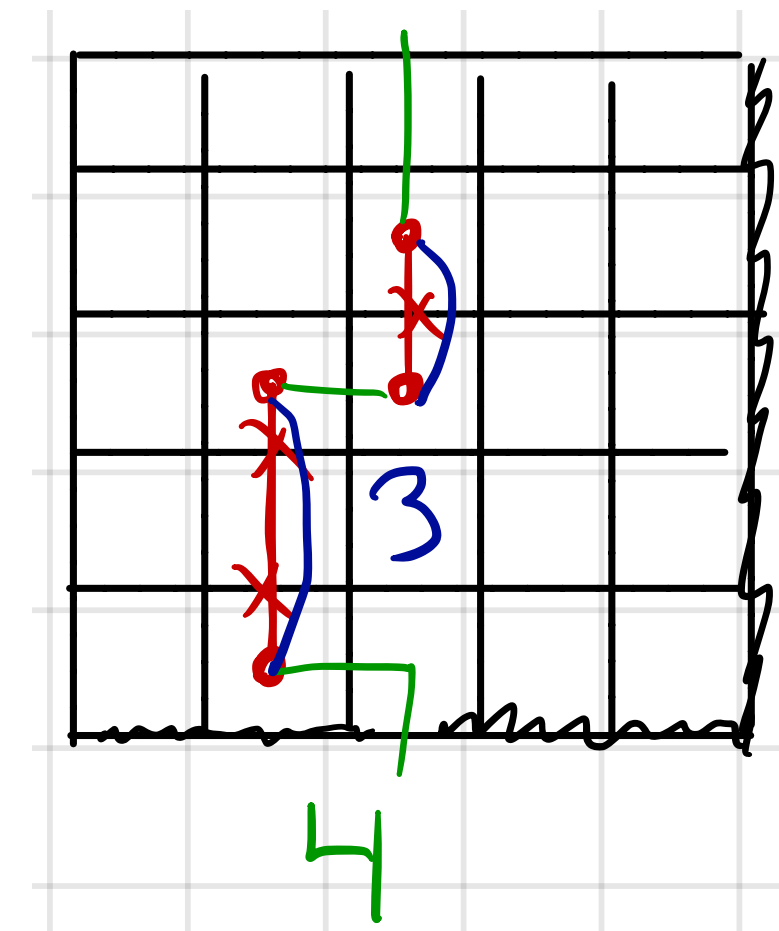


Three errors in a row always gives failed error correction



reward for RL?

Three errors not in a row always gives successful correction



MWPM asymptotic (lowest order in  $p$ ) fail rate is:

$$p_L = 2d \binom{d}{\lceil d/2 \rceil} p^{\lceil d/2 \rceil}$$

# Deep reinforcement learning/Deep Q-learning

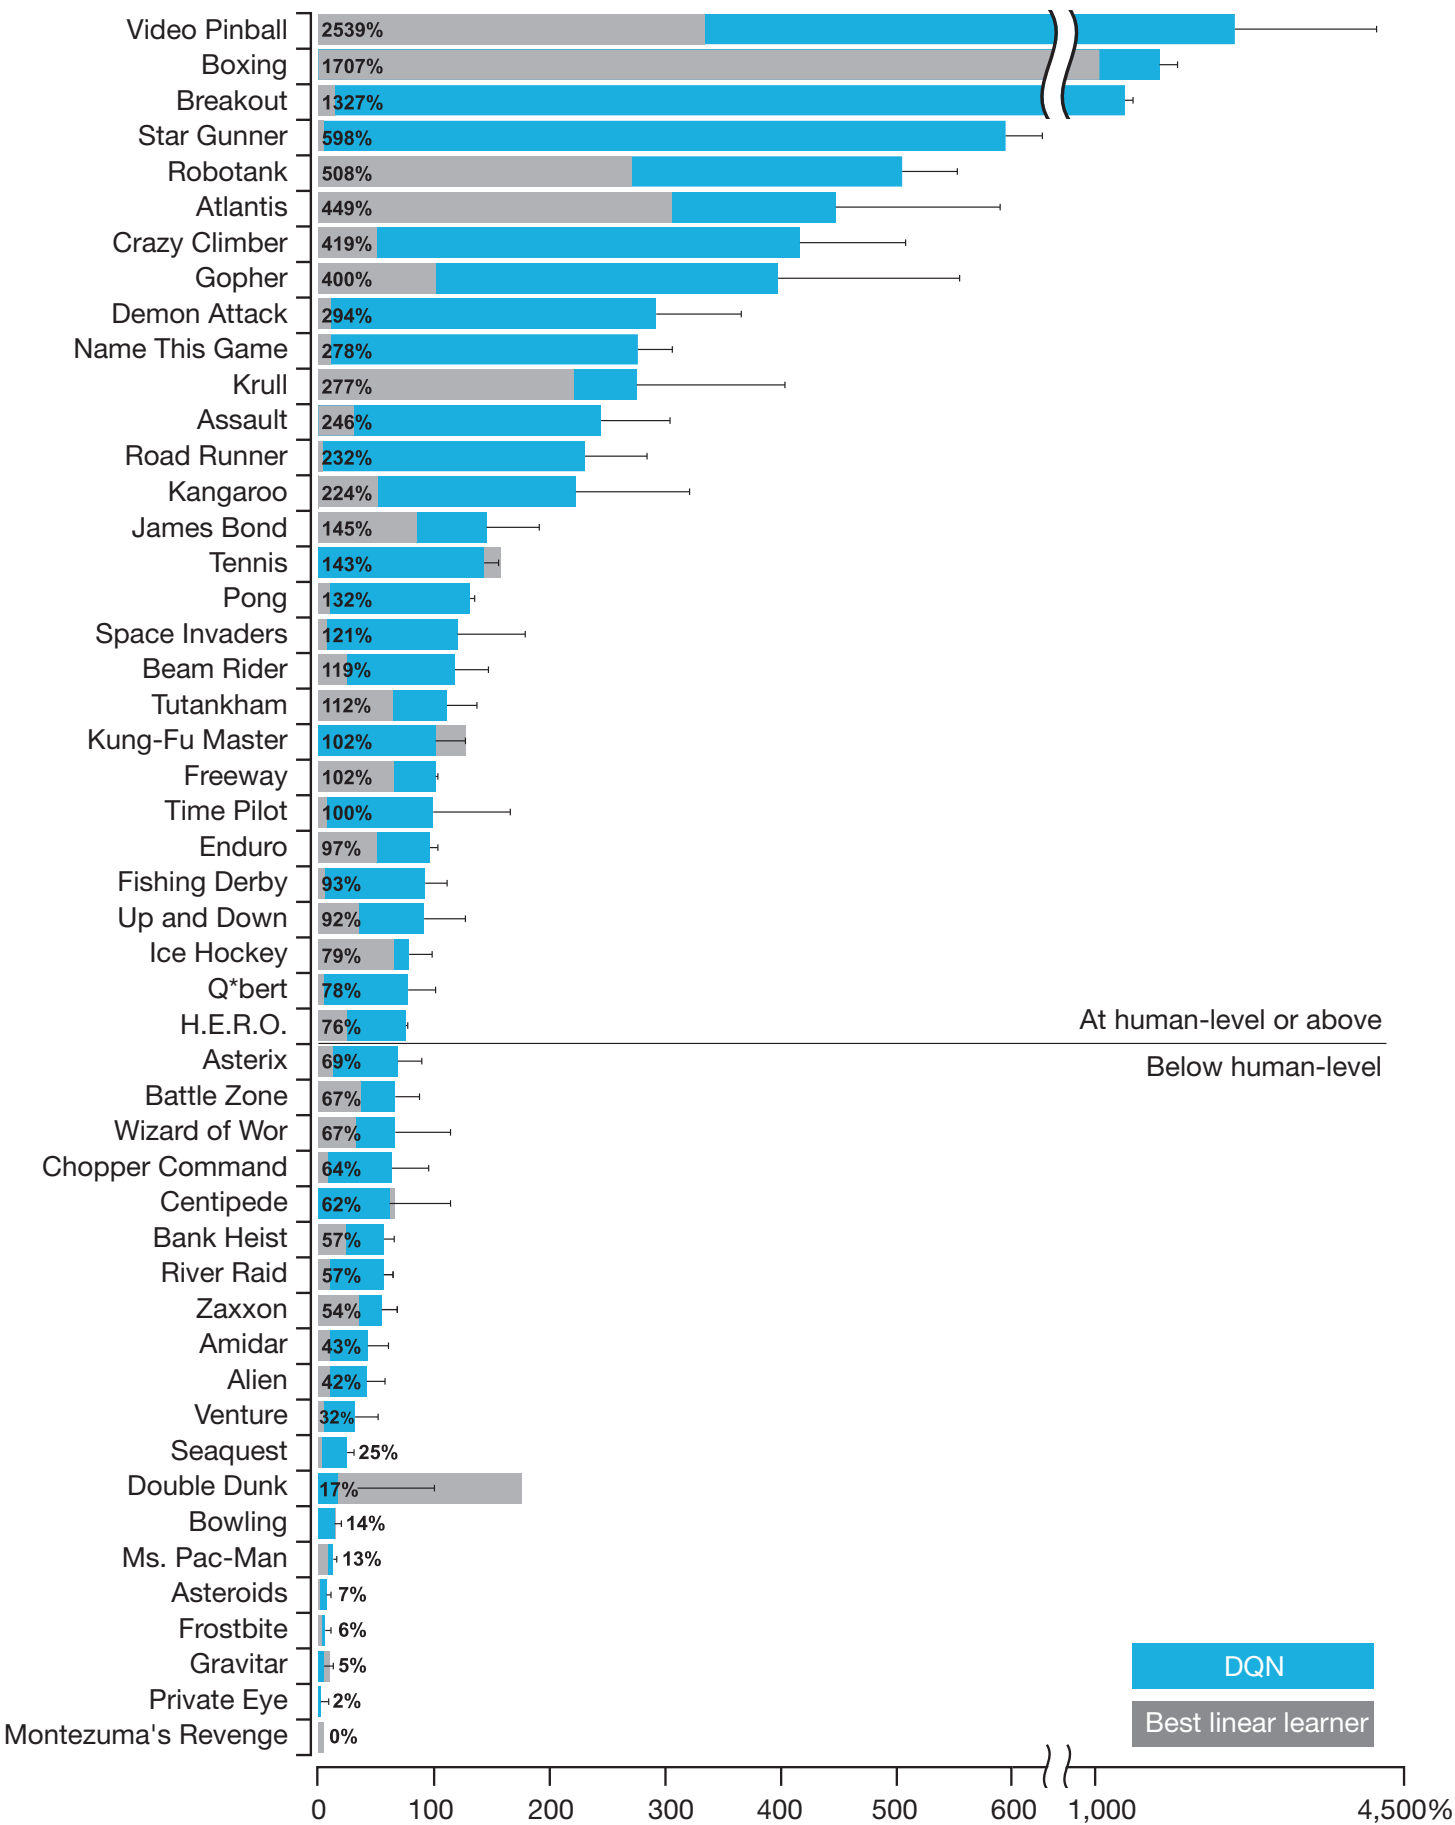
LETTER

2015

doi:10.1038/nature14236

## Human-level control through deep reinforcement learning

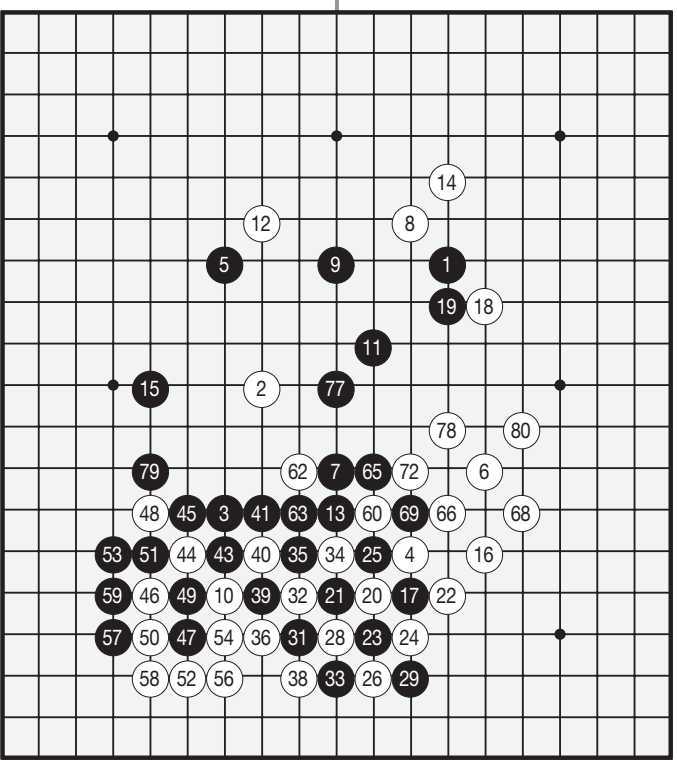
Volodymyr Mnih<sup>1\*</sup>, Koray Kavukcuoglu<sup>1\*</sup>, David Silver<sup>1\*</sup>, Andrei A. Rusu<sup>1</sup>, Joel Veness<sup>1</sup>, Marc G. Bellemare<sup>1</sup>, Alex Graves<sup>1</sup>, Martin Riedmiller<sup>1</sup>, Andreas K. Fiedjeland<sup>1</sup>, Georg Ostrovski<sup>1</sup>, Stig Petersen<sup>1</sup>, Charles Beattie<sup>1</sup>, Amir Sadik<sup>1</sup>, Ioannis Antonoglou<sup>1</sup>, Helen King<sup>1</sup>, Dharshan Kumaran<sup>1</sup>, Daan Wierstra<sup>1</sup>, Shane Legg<sup>1</sup> & Demis Hassabis<sup>1</sup>



2017

## Mastering the game of Go without human knowledge

David Silver<sup>1\*</sup>, Julian Schrittwieser<sup>1\*</sup>, Karen Simonyan<sup>1\*</sup>, Ioannis Antonoglou<sup>1</sup>, Aja Huang<sup>1</sup>, Arthur Guez<sup>1</sup>, Thomas Hubert<sup>1</sup>, Lucas Baker<sup>1</sup>, Matthew Lai<sup>1</sup>, Adrian Bolton<sup>1</sup>, Yutian Chen<sup>1</sup>, Timothy Lillicrap<sup>1</sup>, Fan Hui<sup>1</sup>, Laurent Sifre<sup>1</sup>, George van den Driessche<sup>1</sup>, Thore Graepel<sup>1</sup> & Demis Hassabis<sup>1</sup>



## AlphaStar 2019



# Q-learning

- Agent in an environment described by a **state**  $s$ .
- Agent takes **actions**  $a$  to move between states,  $s \rightarrow s'$ .
- **Reward** (positive or negative)  $r$  is given depending on state/action.
- Agent learns **policy**,  $\pi(s,a)$ , to navigate environment for optimal accumulated reward (return) by exploring.

**Q-function (action-value fcn)  $Q(s,a)$  quantifies expected return from taking action  $a$  in state  $s$  and subsequently following the optimal policy.**

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

$\gamma < 1$  is discounting factor, better to get reward now than later

Explore to get reward and learn  $Q \Rightarrow$  optimal policy

Difficult if big world with many states and actions

Use Artificial Neural Network to represent Q-function

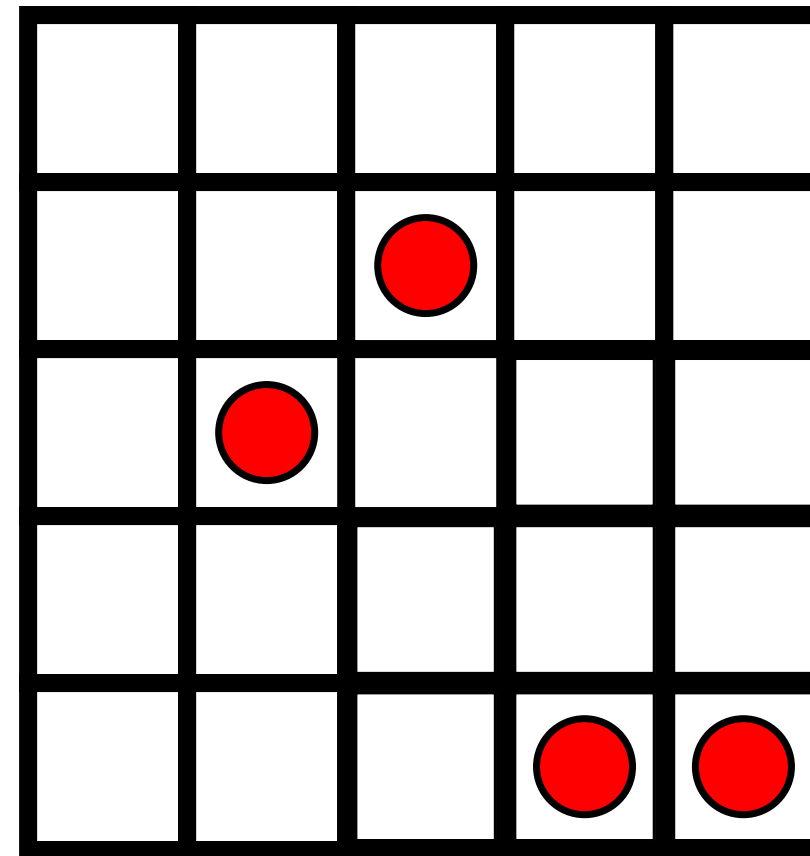
**Deep Q-learning**

# Q-learning for the toric code

**state** is a syndrome

**action** is a bitflip=cardinal move of defect

**reward**,  $r=-1$  per move (i.e. we aim to learn MWPM)



**State space is very big**

number of ways of placing  $N_s$  defects on  $d^2$  sites:

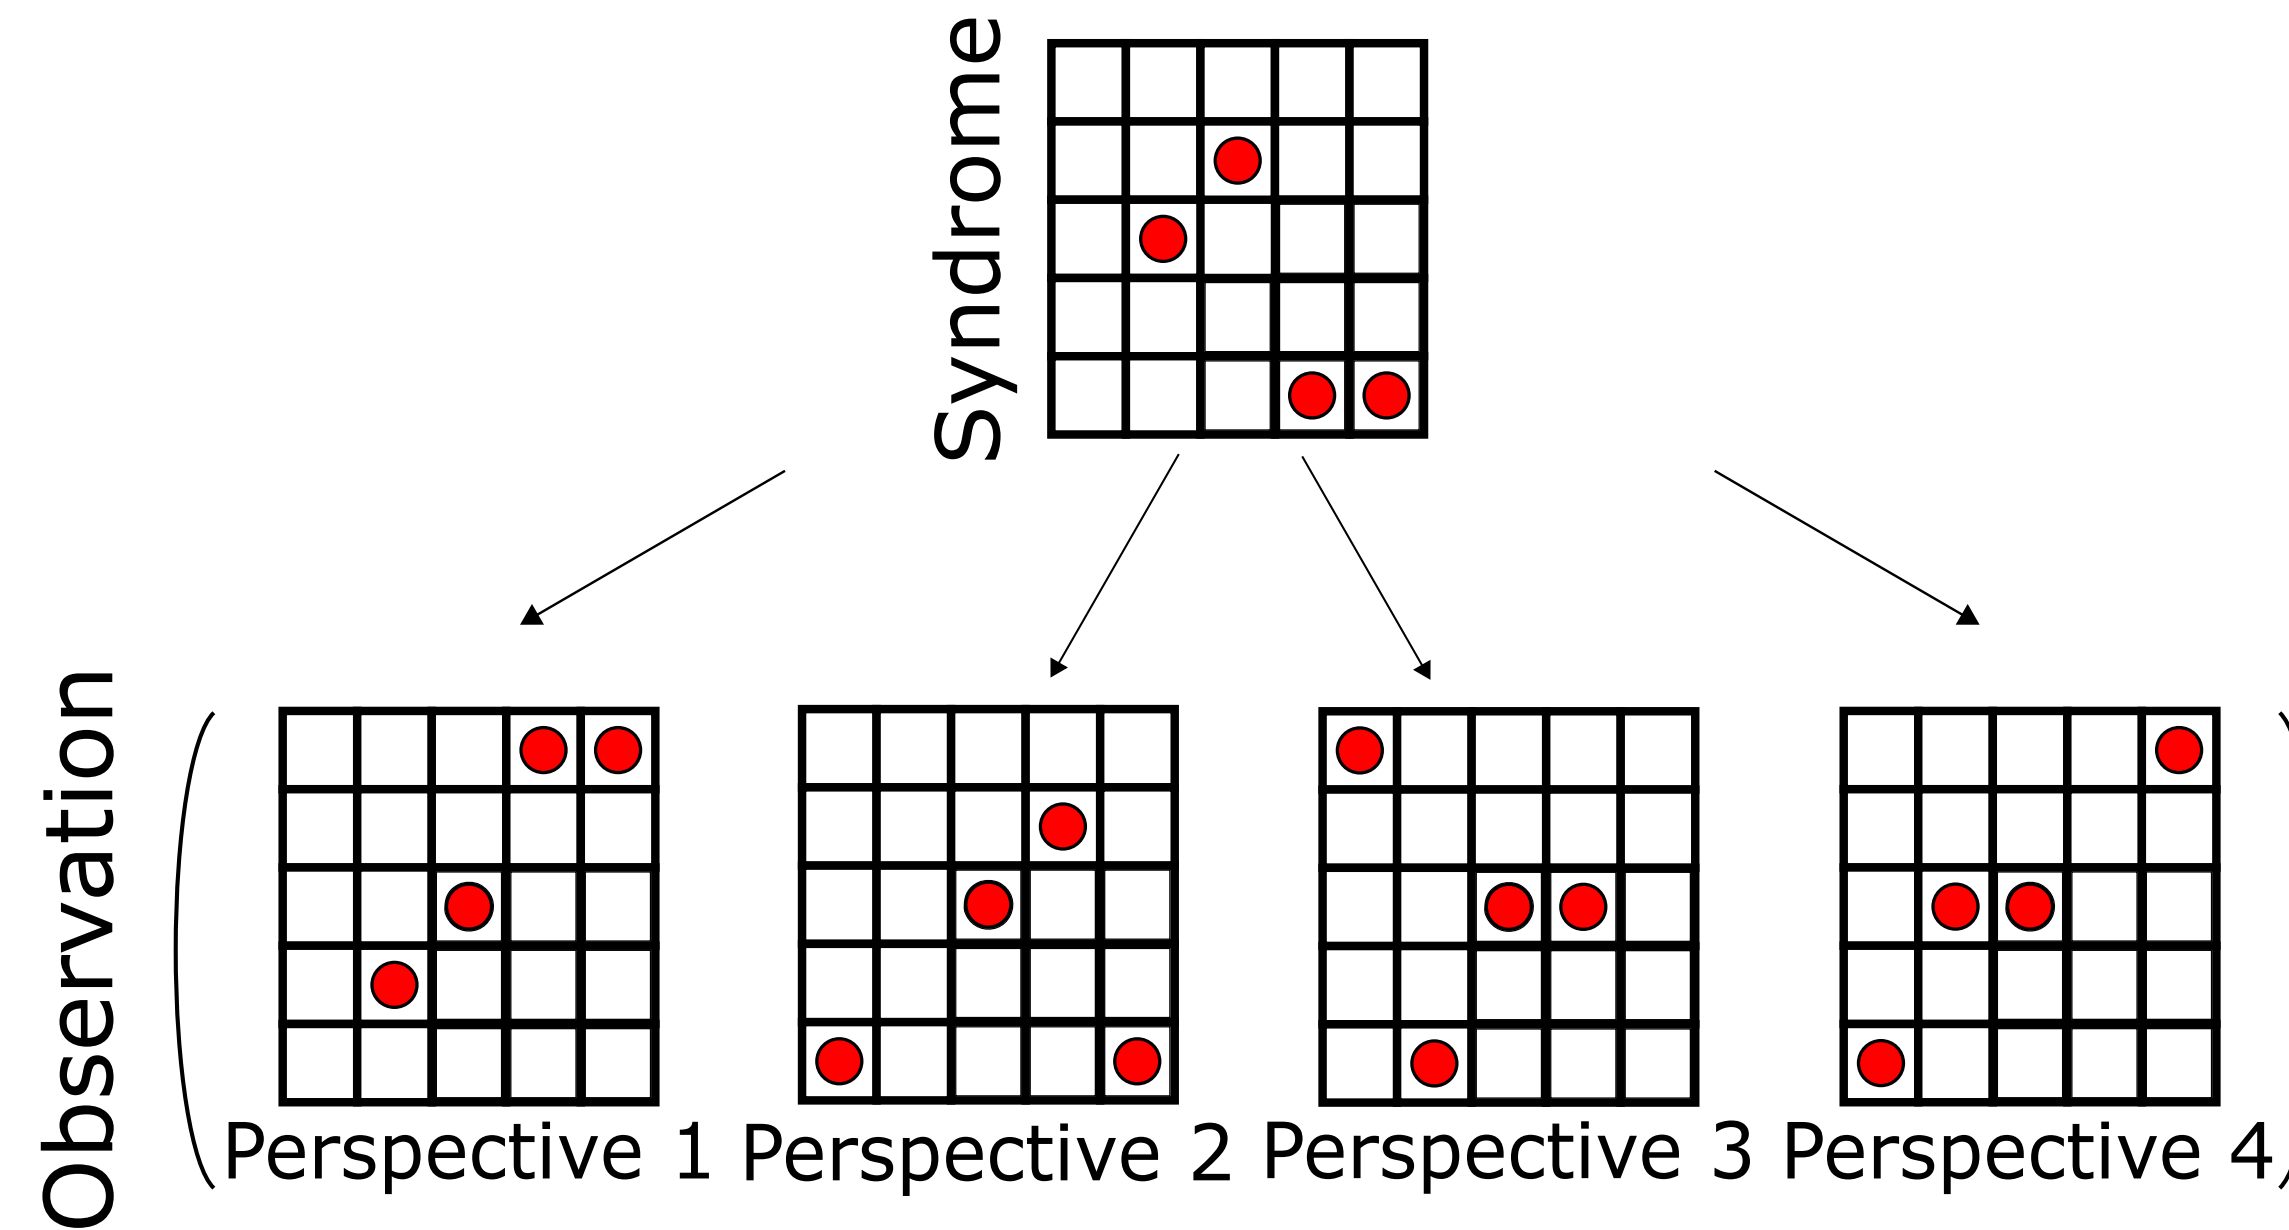
$$\binom{d^2}{N_s} \approx \binom{49}{20} \sim 10^{13}$$

for  $d=7$  and  $p=10\%$

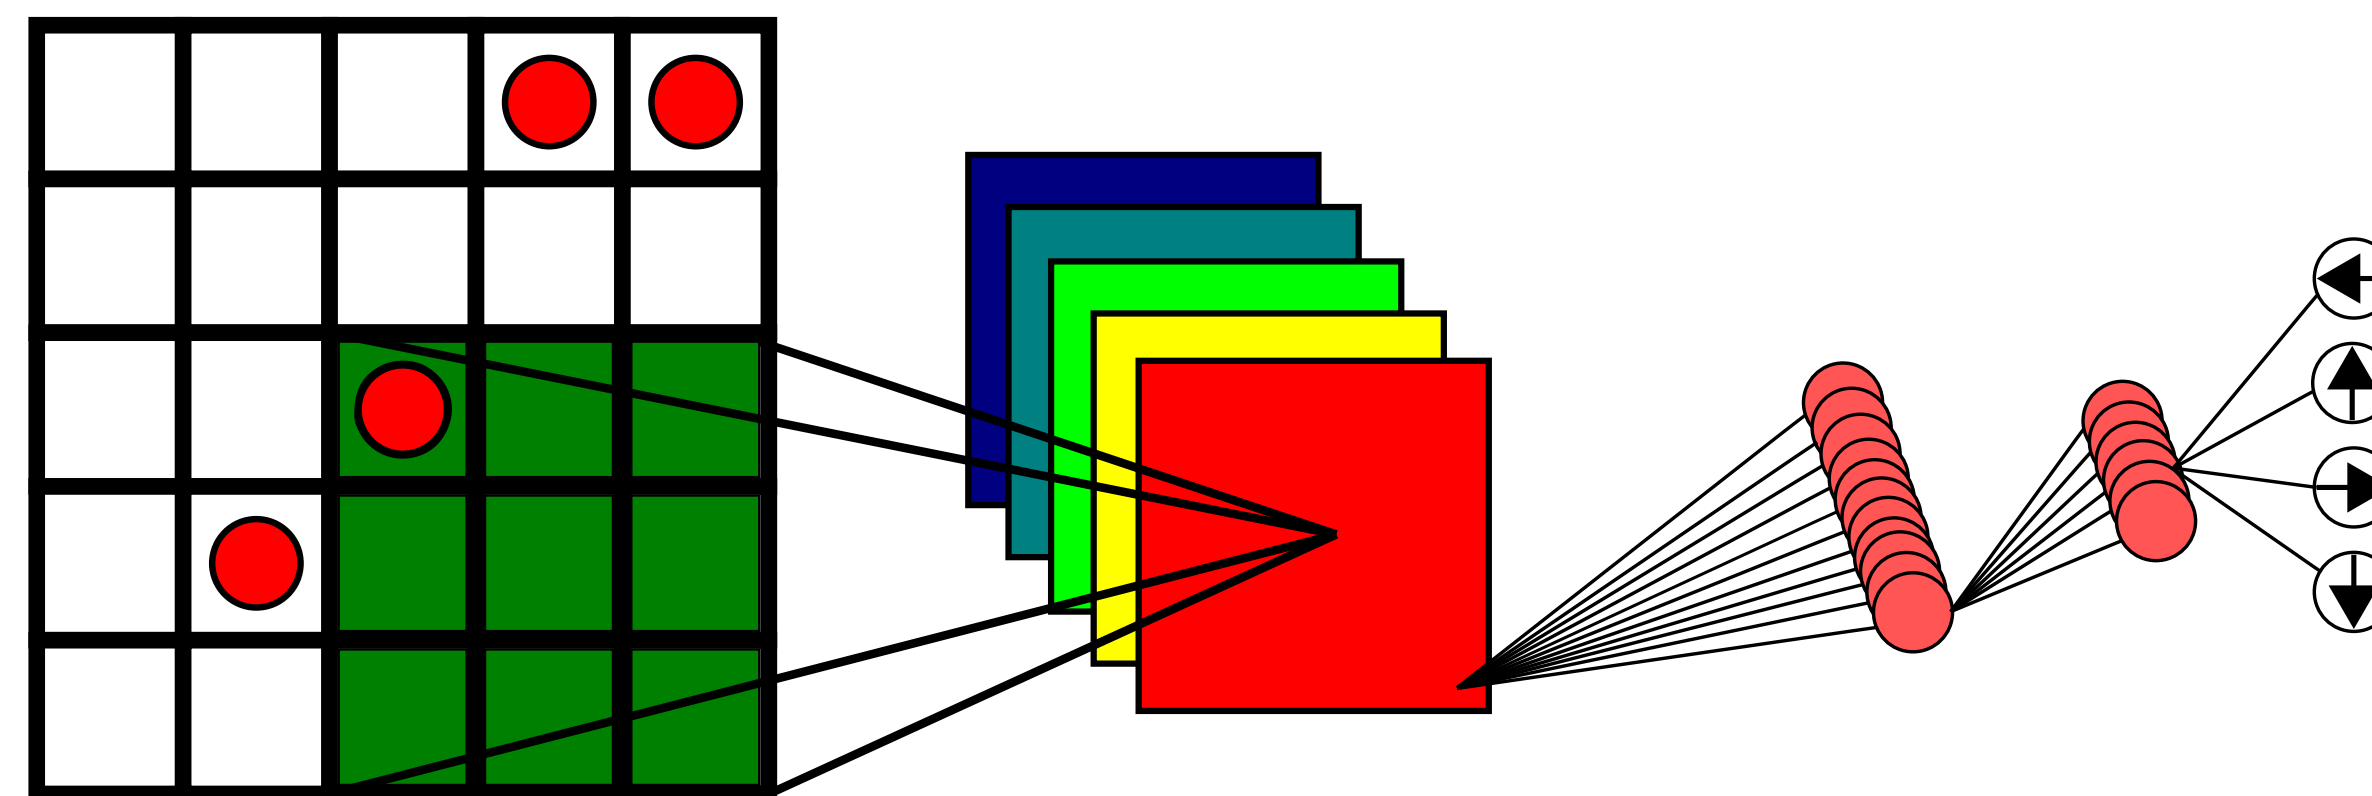
**Use deep Q-learning**

# Efficient implementation of Q-network

Use translational and rotational symmetry  
to center each defect.

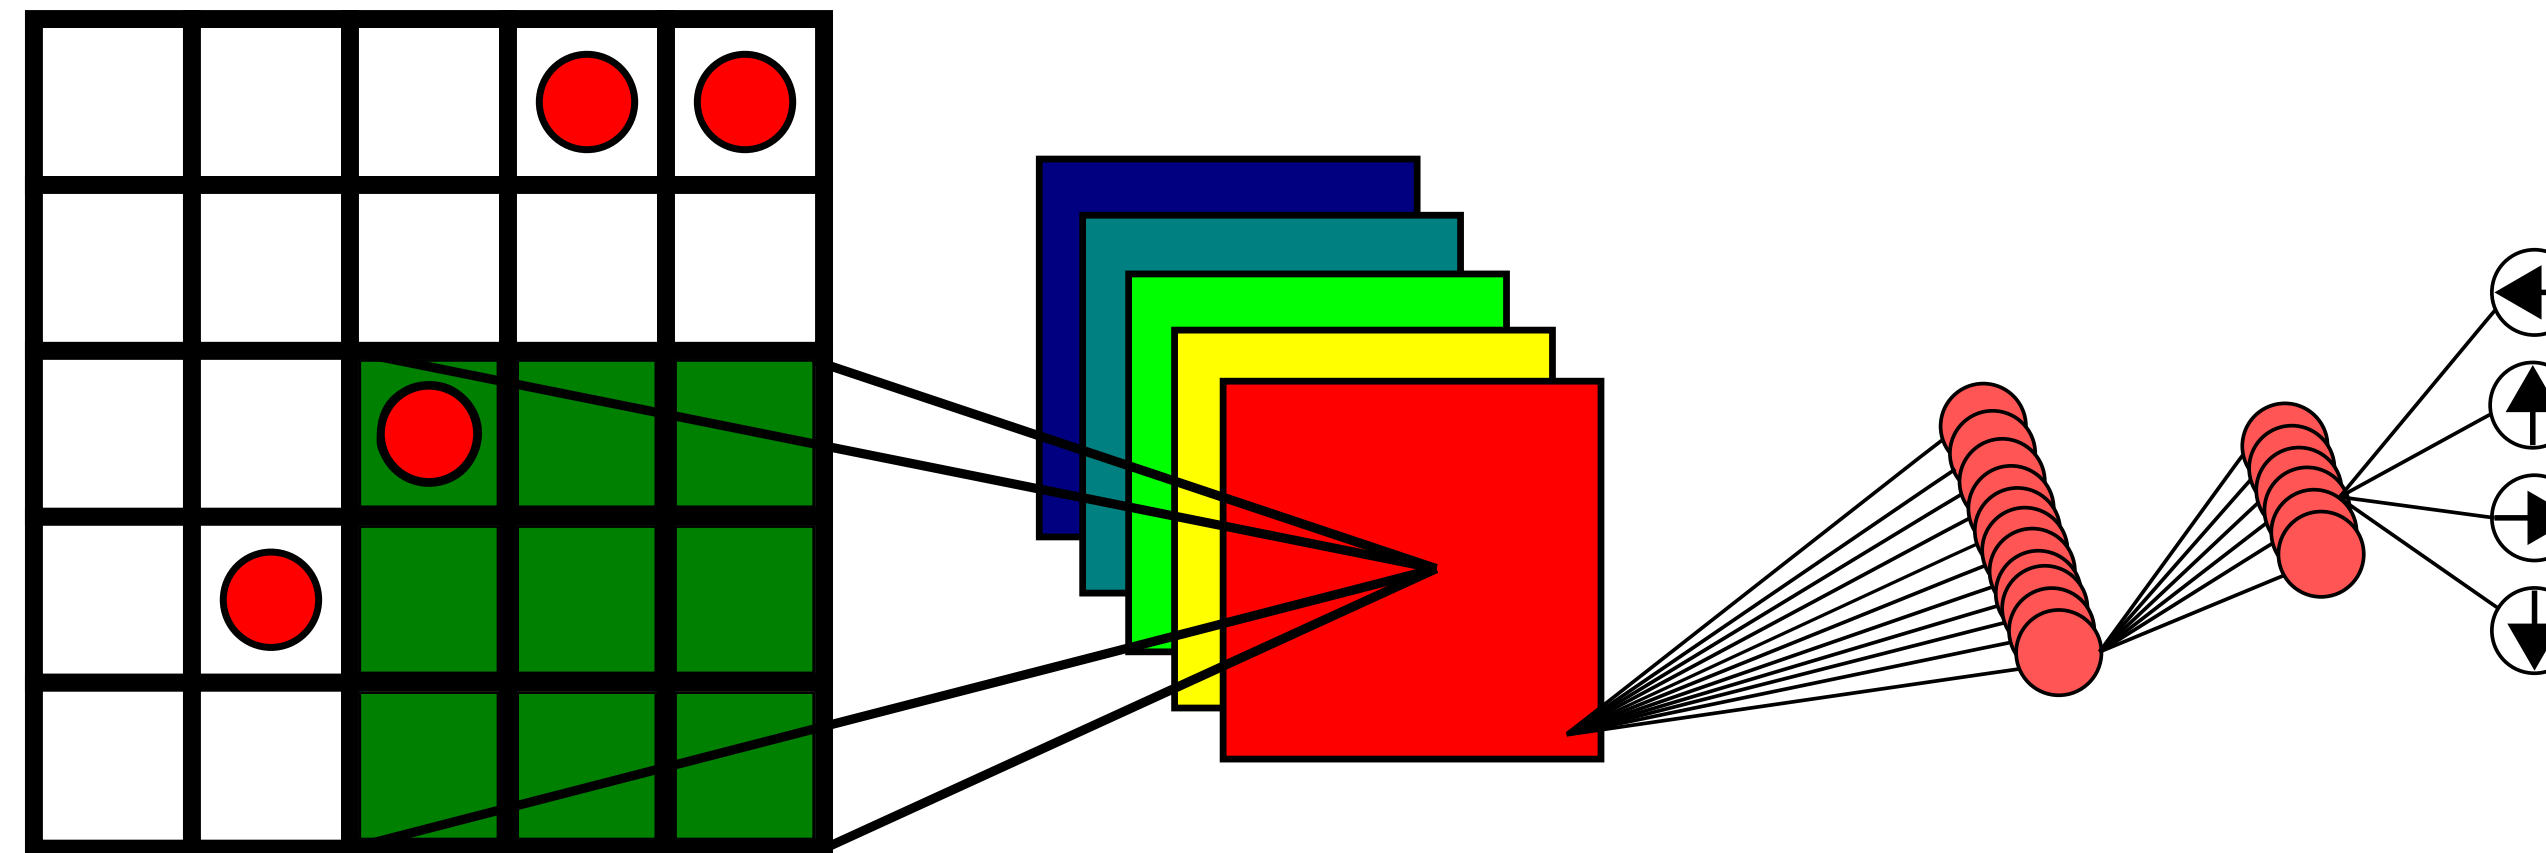


## Convolutional NN



# Deep Q-network

Network gives Q-values for the 4 movements of the **central** defect.  
Crucial simplification, fixed number (4) actions, and doesn't have to learn about boundaries.



## Experience replay is crucial for training

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

Significant reduction in number of parameters.  
Size of state space for d=7, and N<sub>s</sub>=20 defects (10% error)

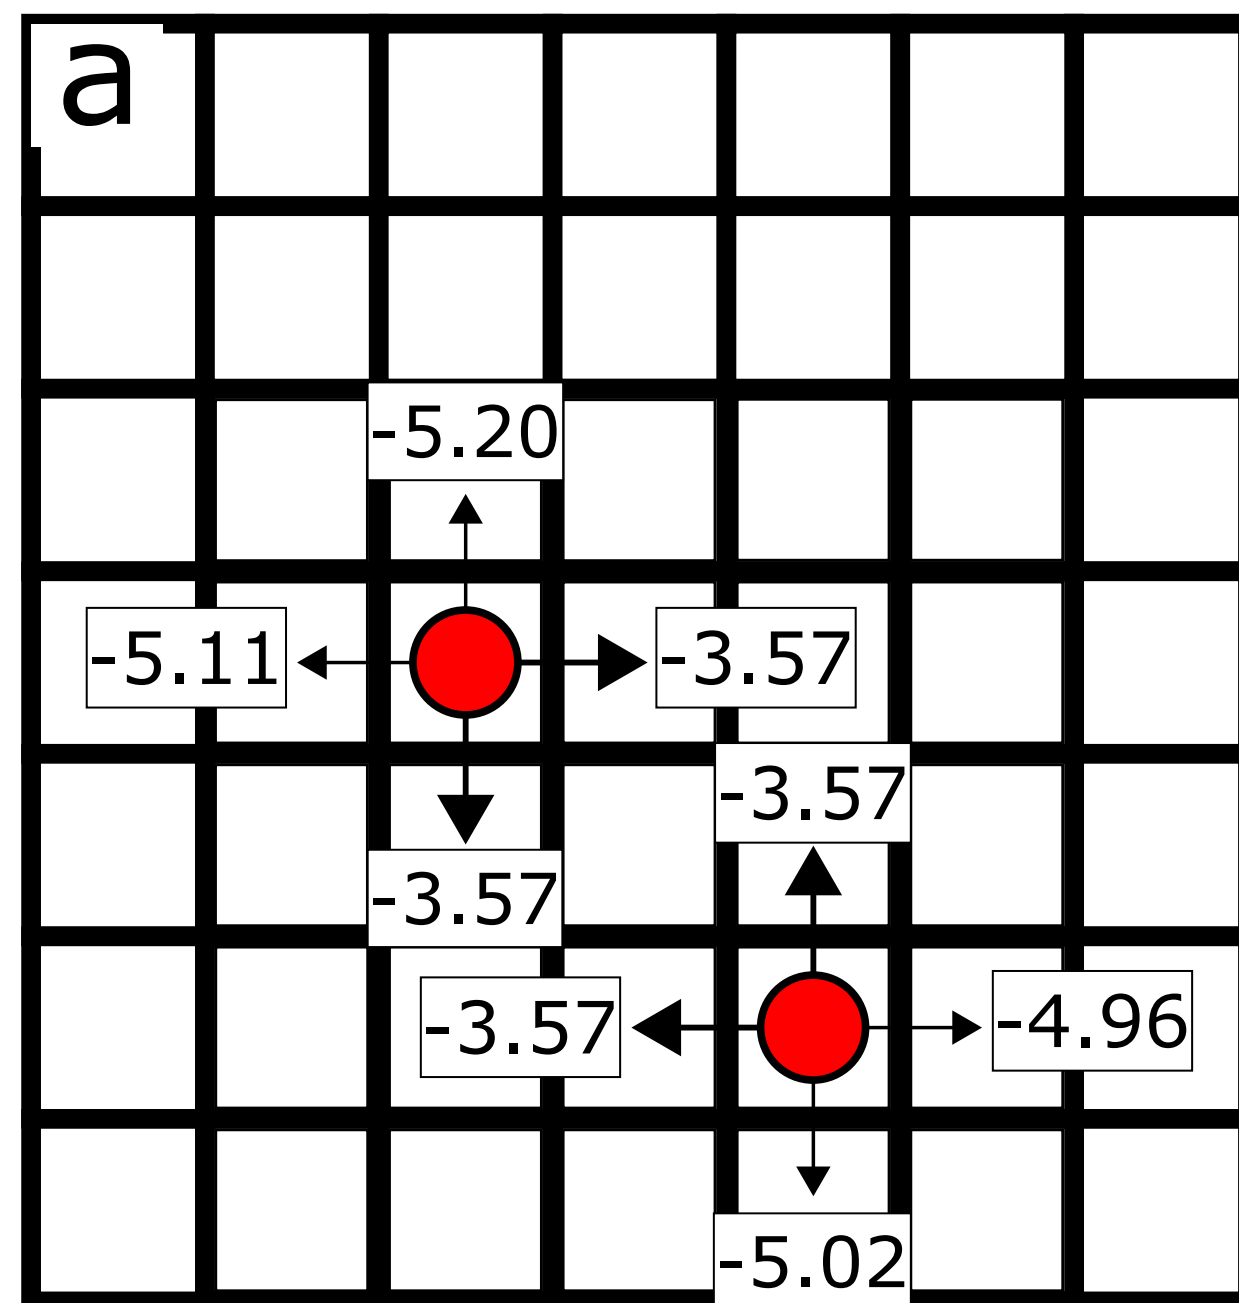
$$\binom{d^2}{N_s} \approx \binom{49}{20} \sim 10^{13}$$

# Results. Converged Q-network.

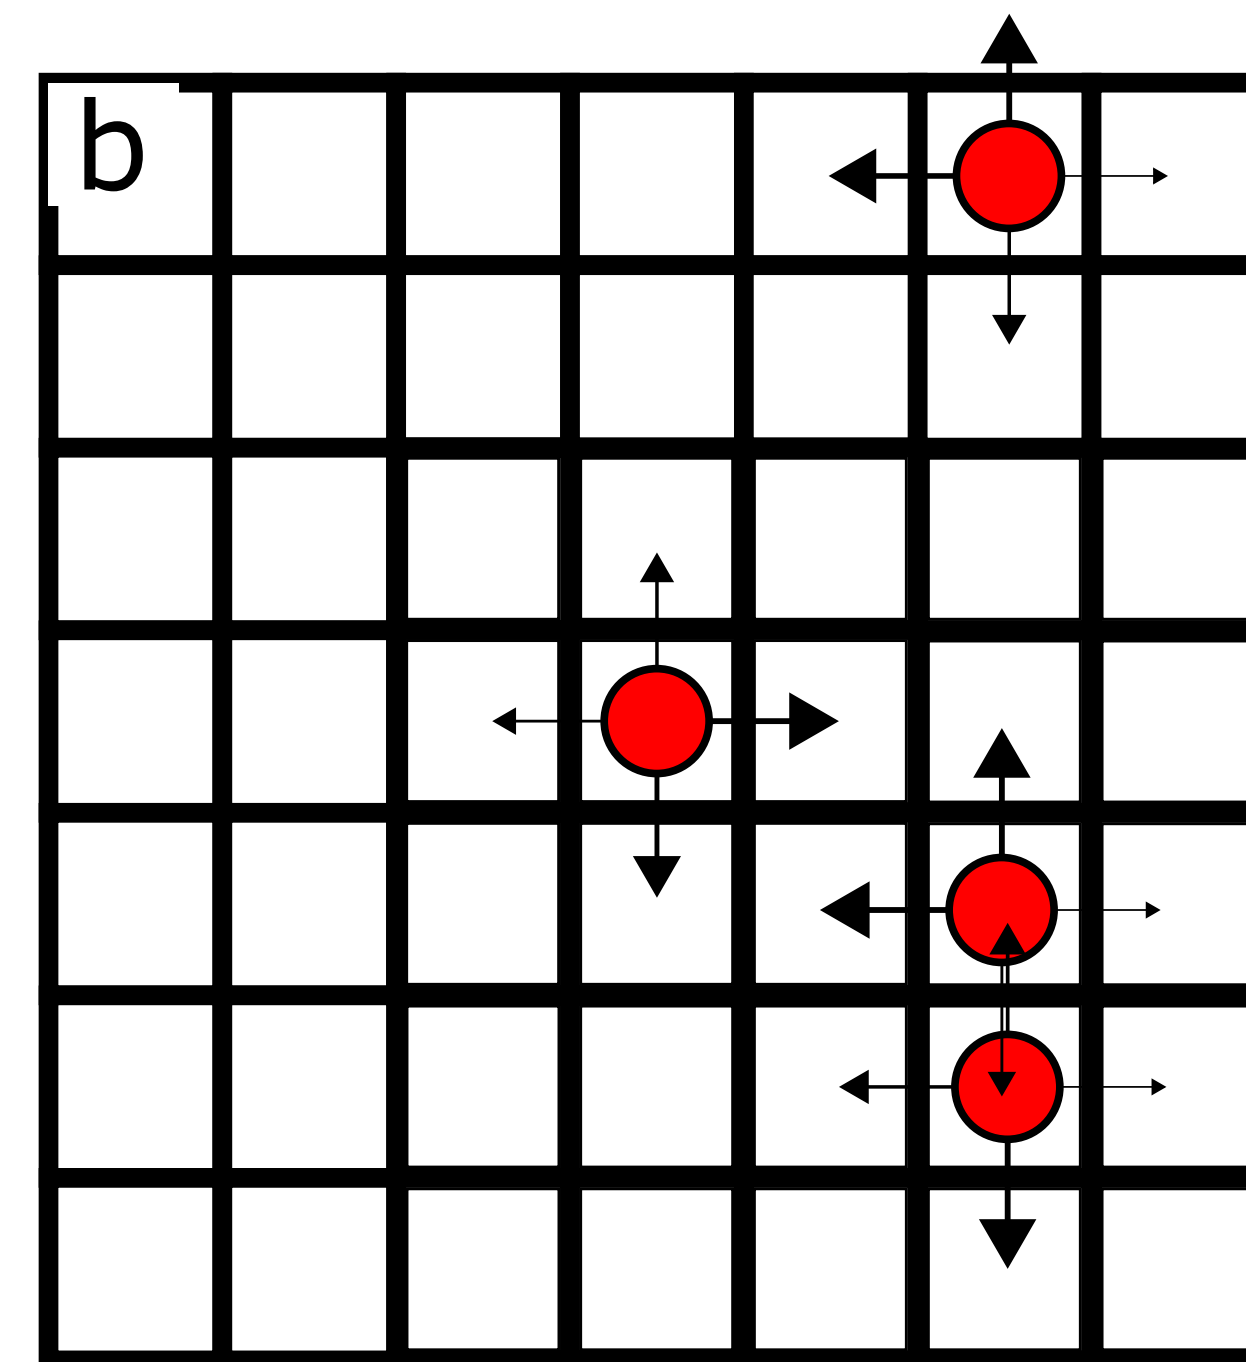
Examples:

Large arrow=Large Q-value for that action

4-steps to elimination



Shortest total path (MWPM)



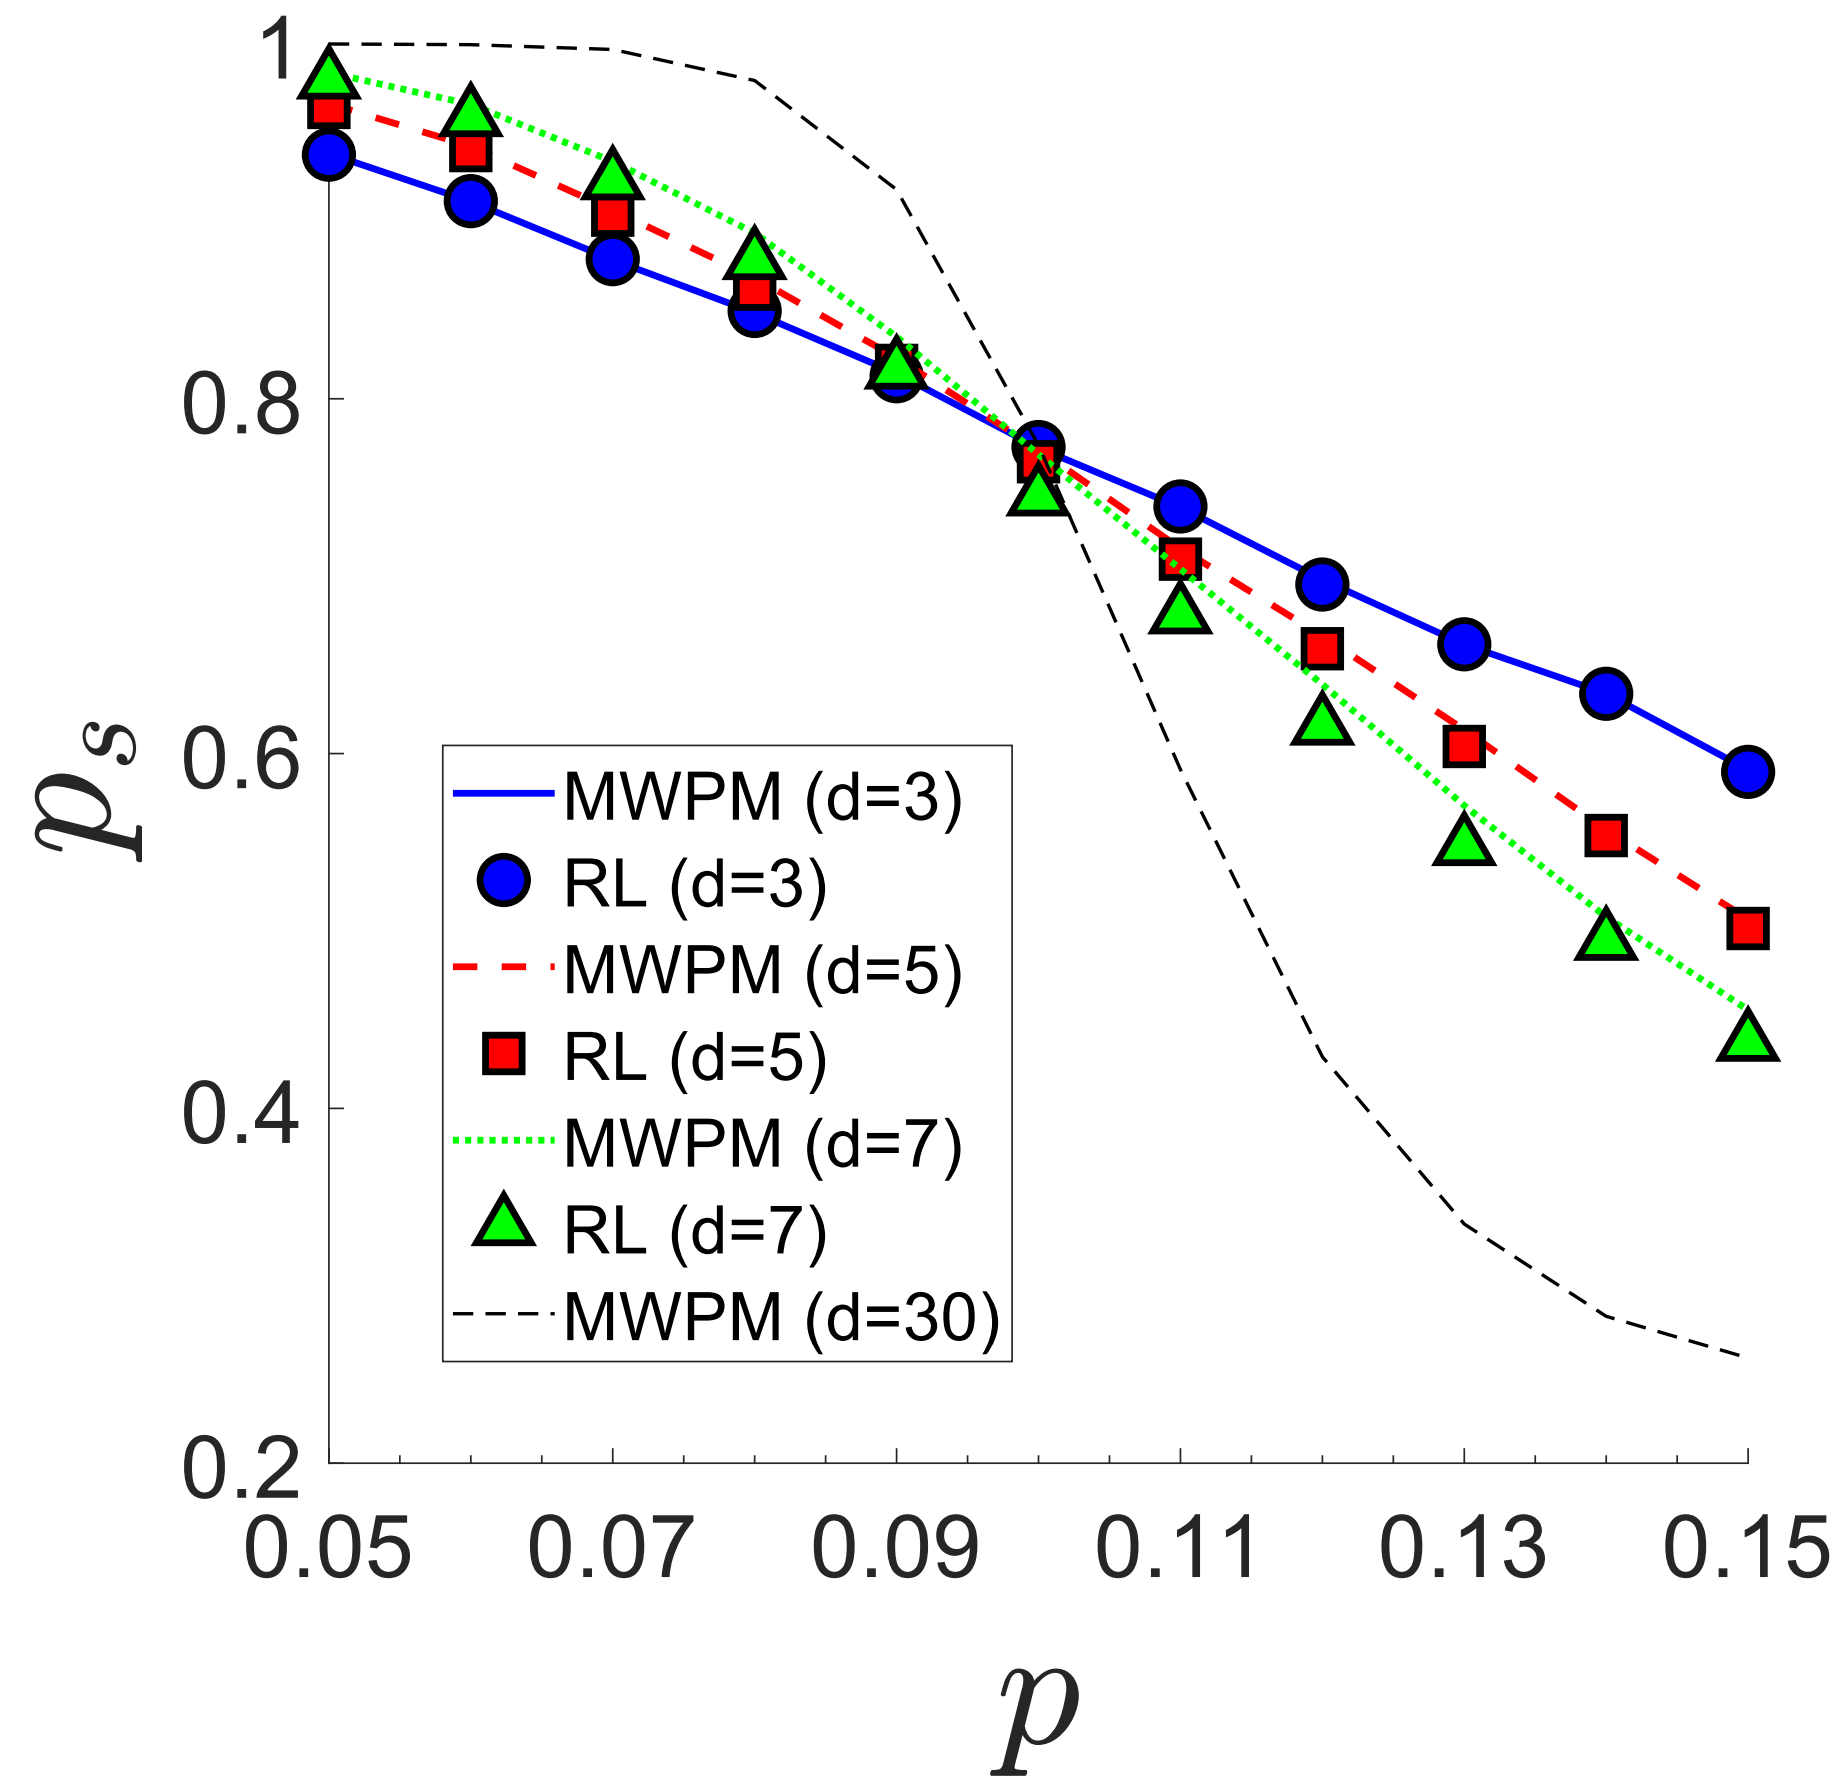
$$R = -1 - \gamma - \gamma^2 - \gamma^3 = -3.62$$

$$\gamma = 0.95$$

(semi-) quantitatively correct Q-values

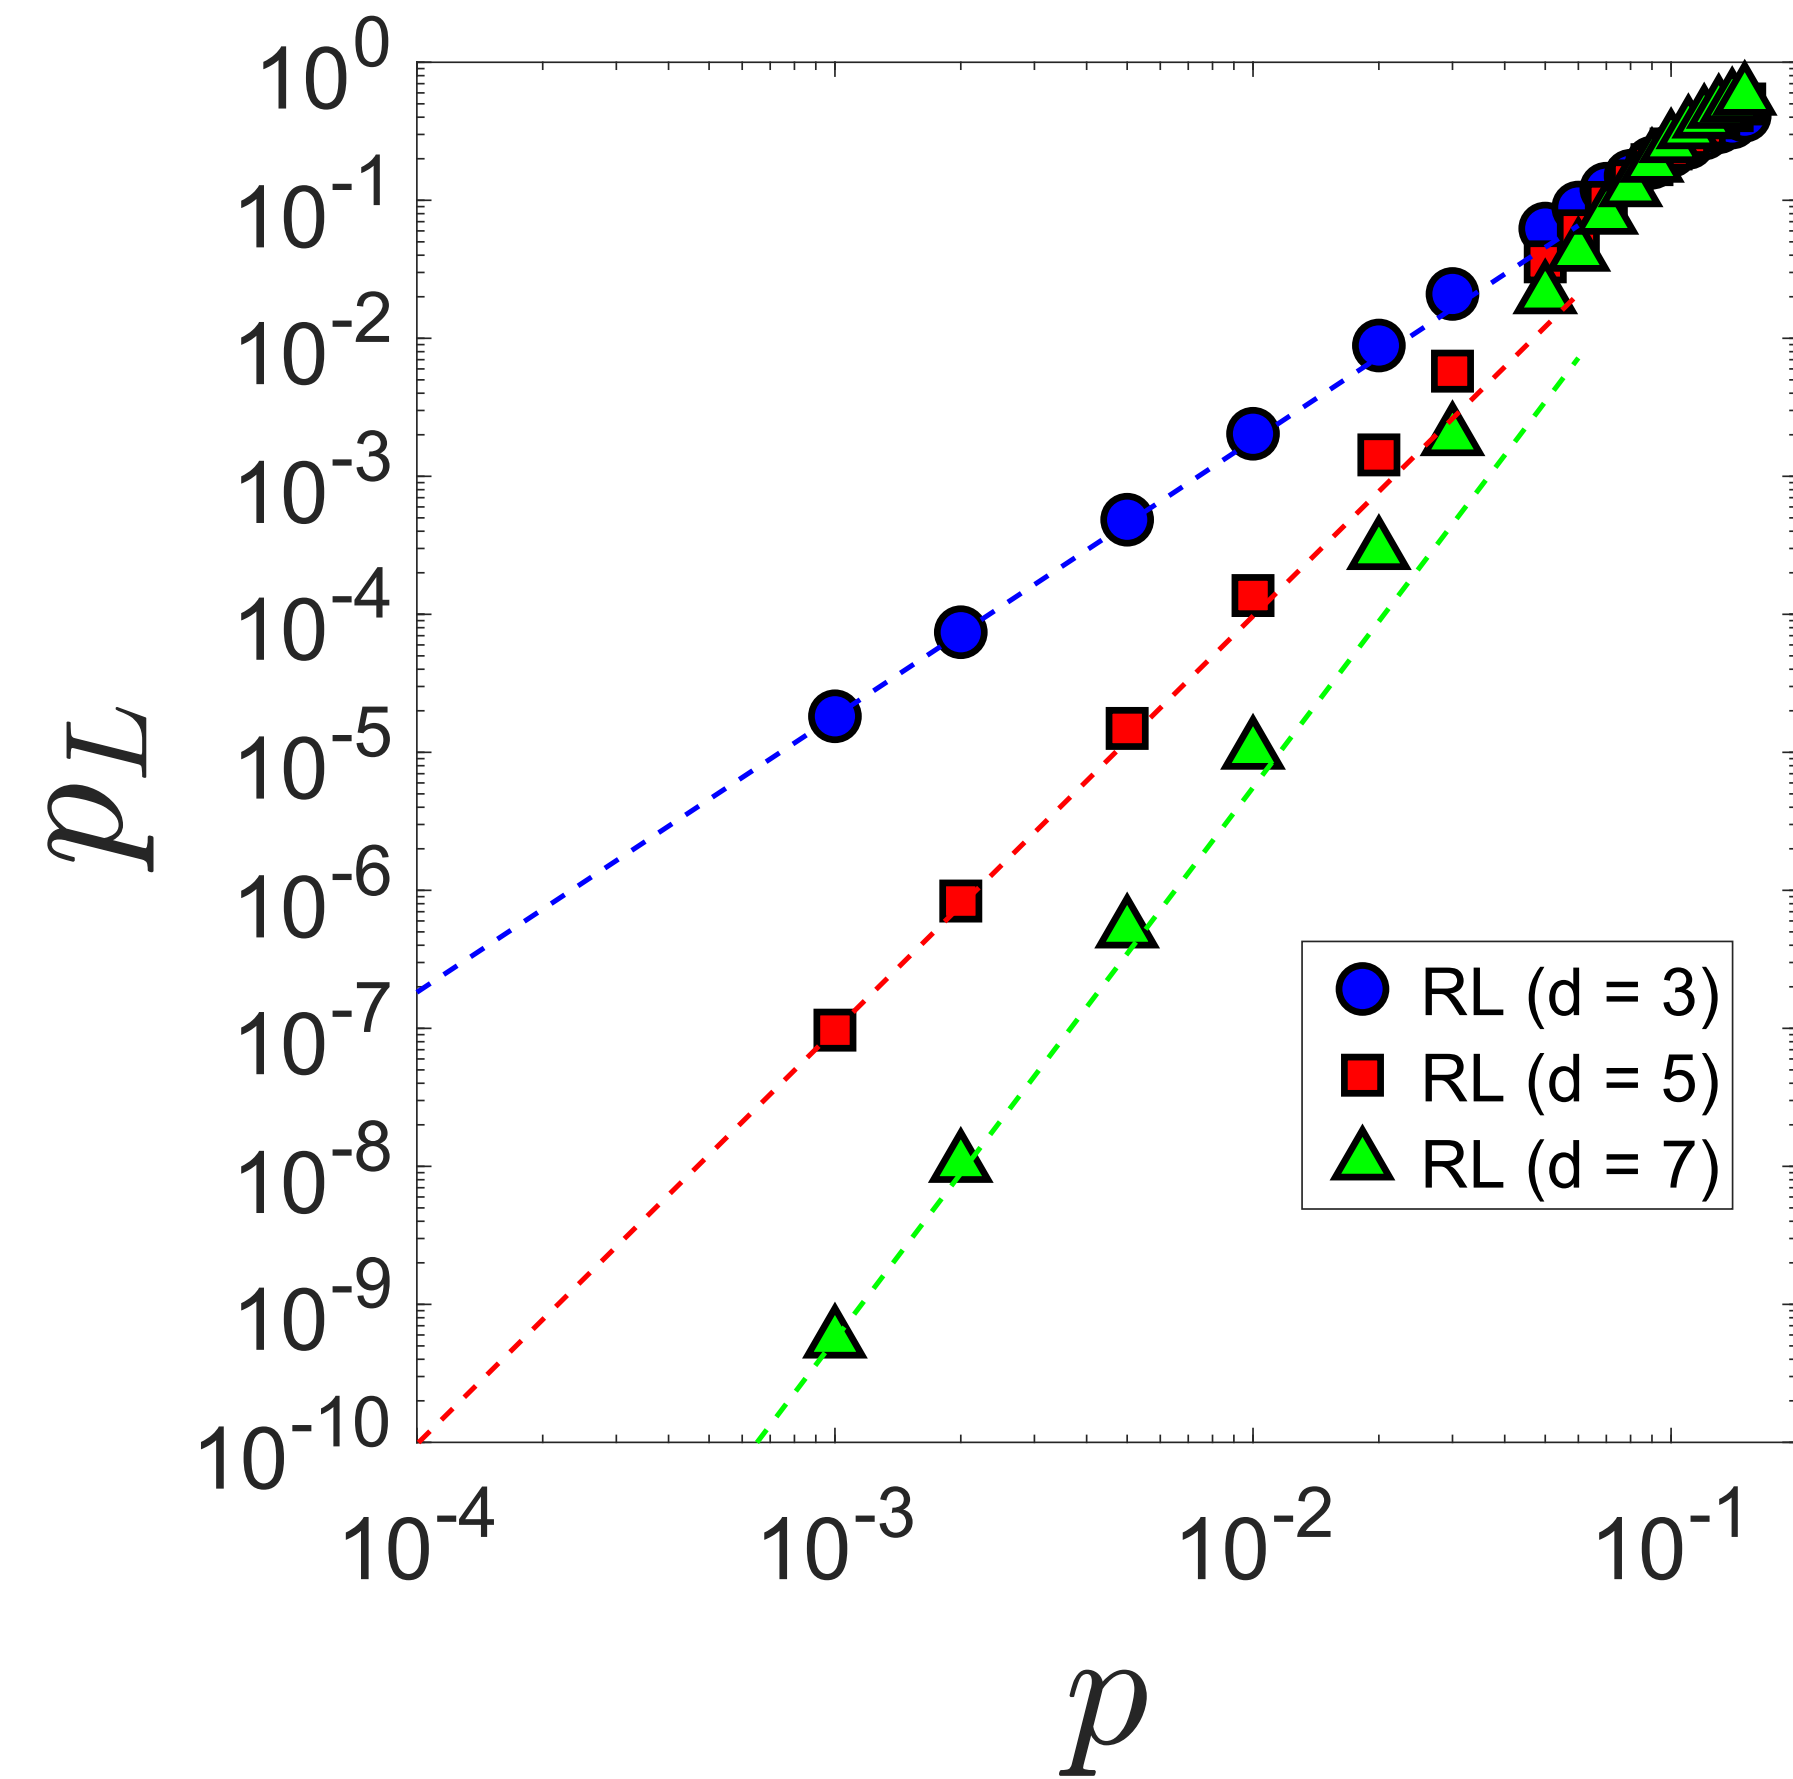
# Results

**Logical success-rate, large error rates  
close to MWPM**



**bit flip error rate**

**Logical fail-rate, small error rates  
identical to MWPM**



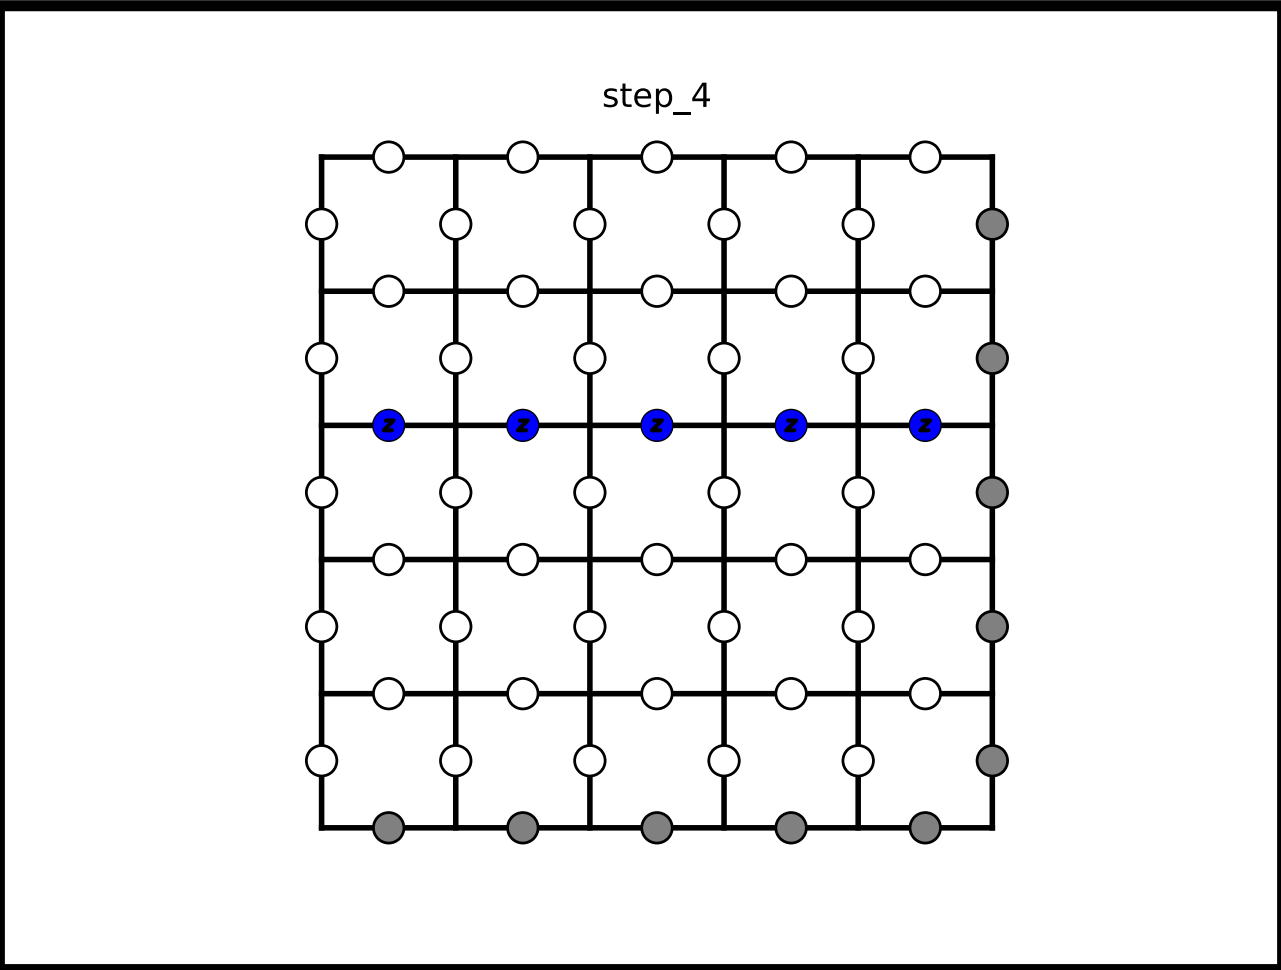
Fits asymptotic form for small  $p$ :

$$p_L = 2d \binom{d}{\lceil d/2 \rceil} p^{\lceil d/2 \rceil}$$

# Depolarizing noise, work in progress

## Example syndrome

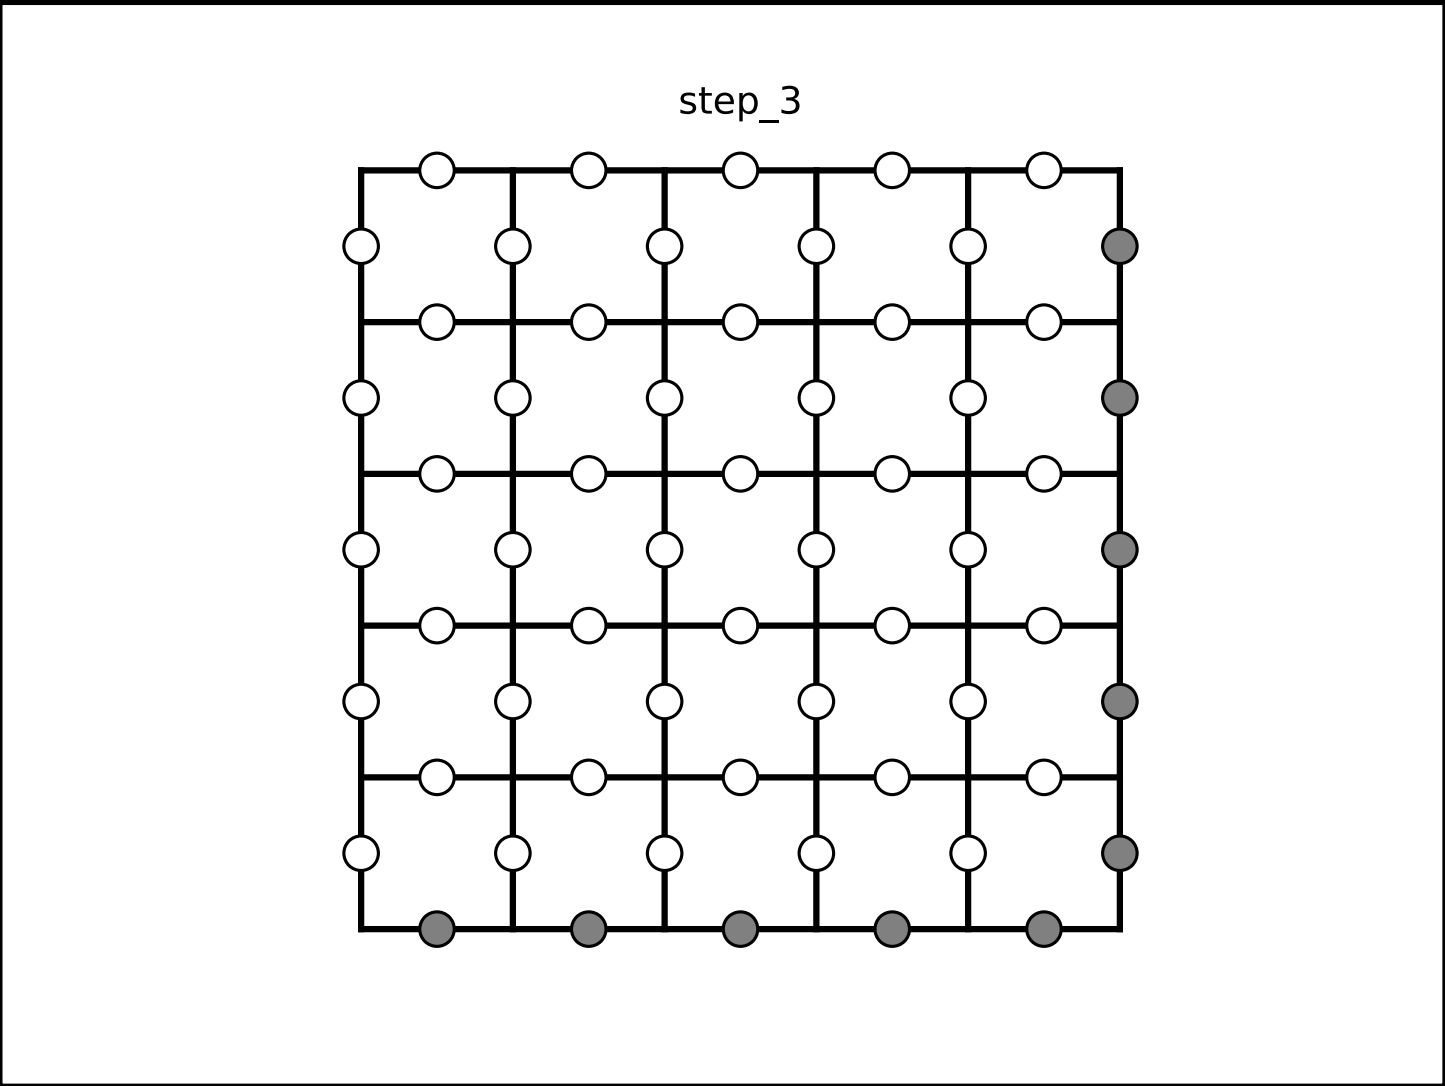
MWPM



logical phase-flip

Reinforcement trained solver

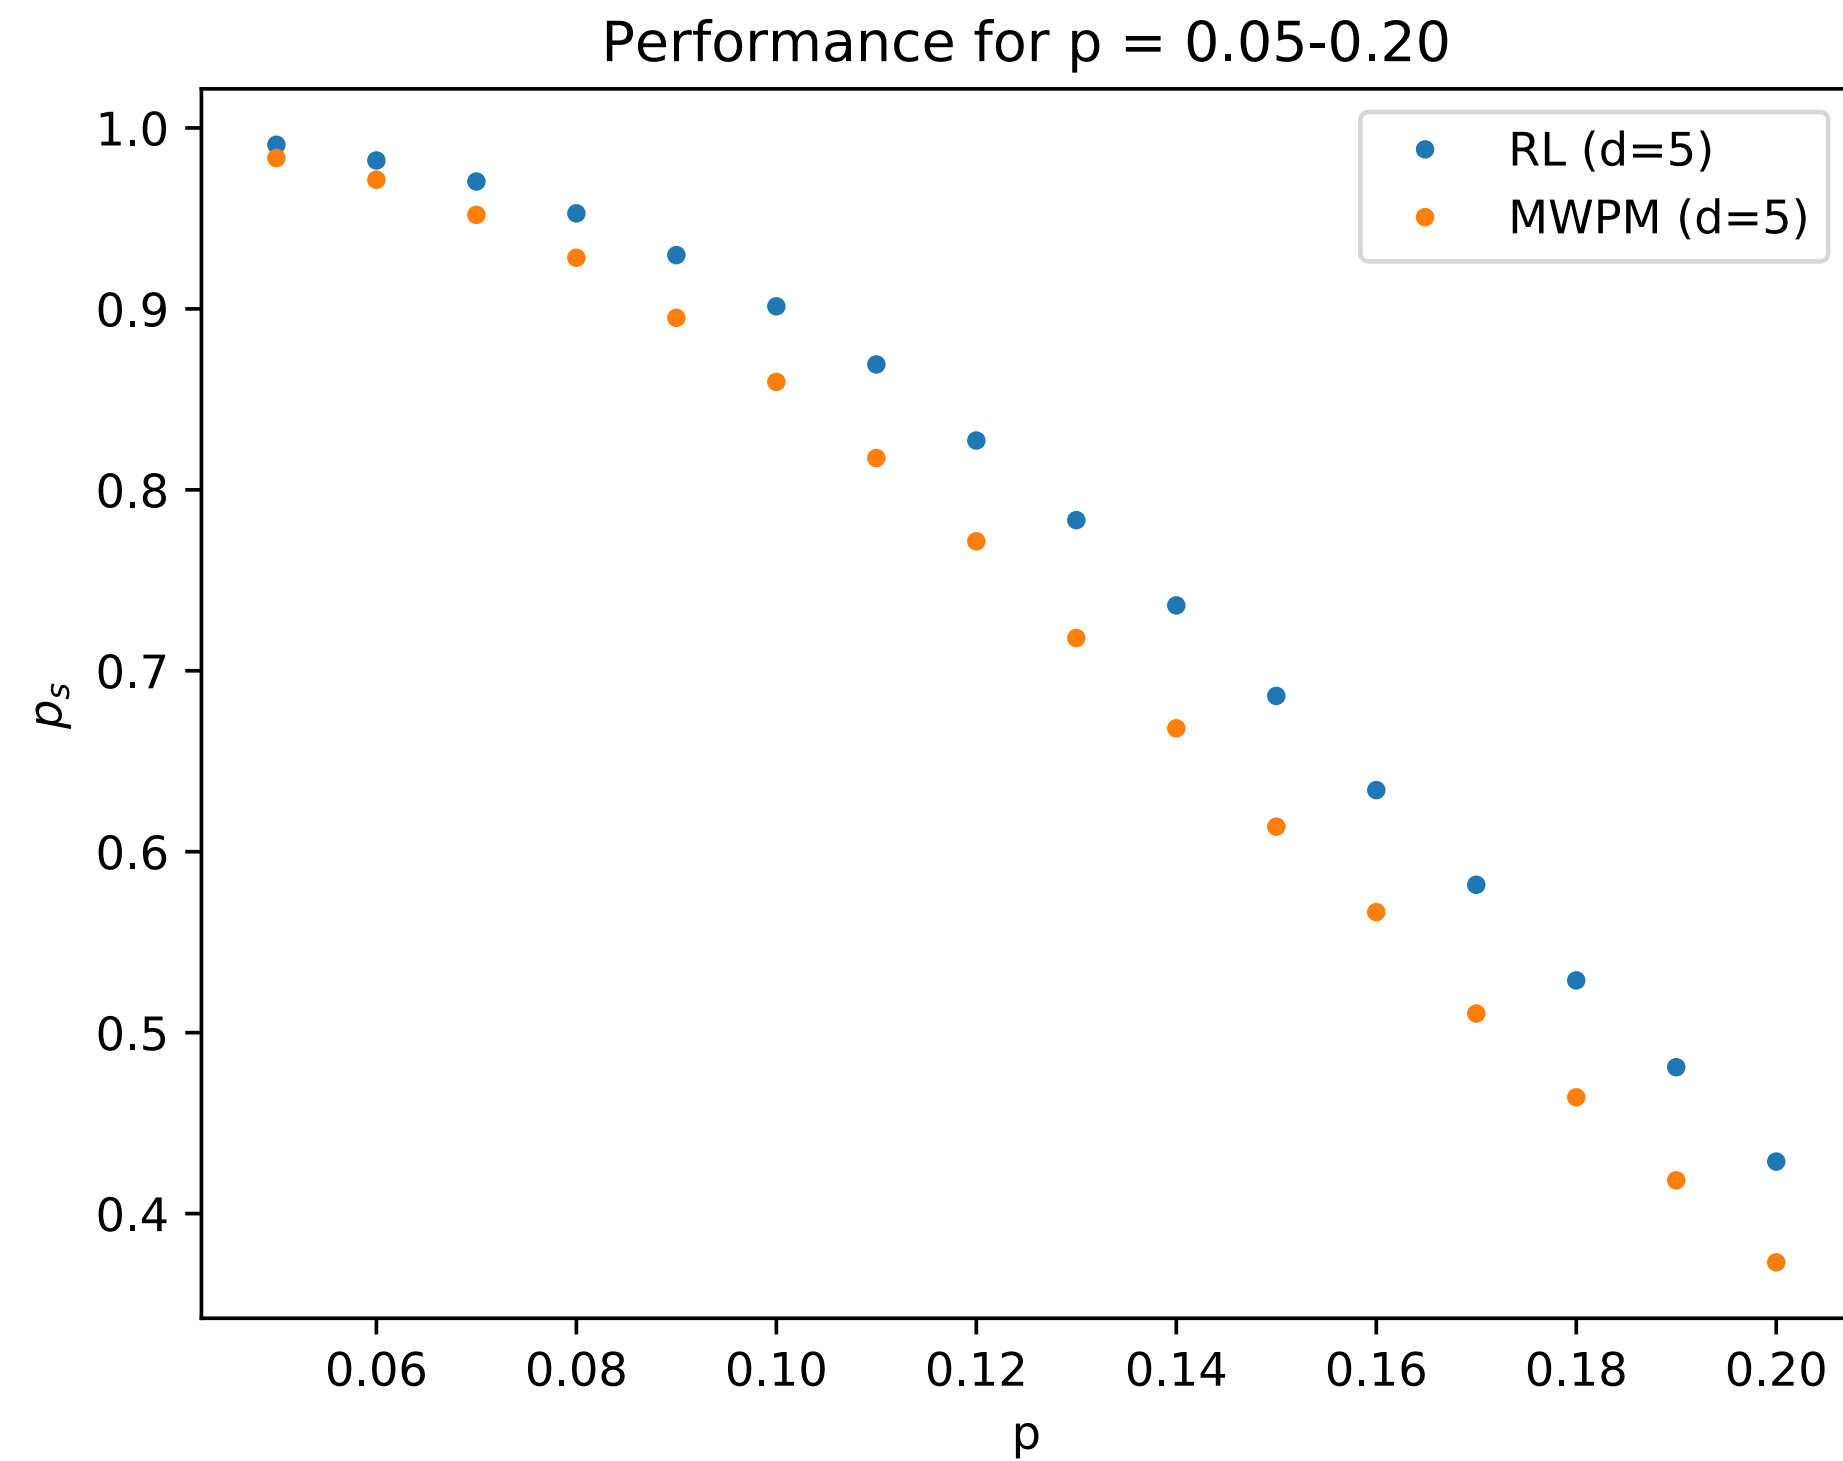
reward=annihilation of complete syndrome + small intermediate reward



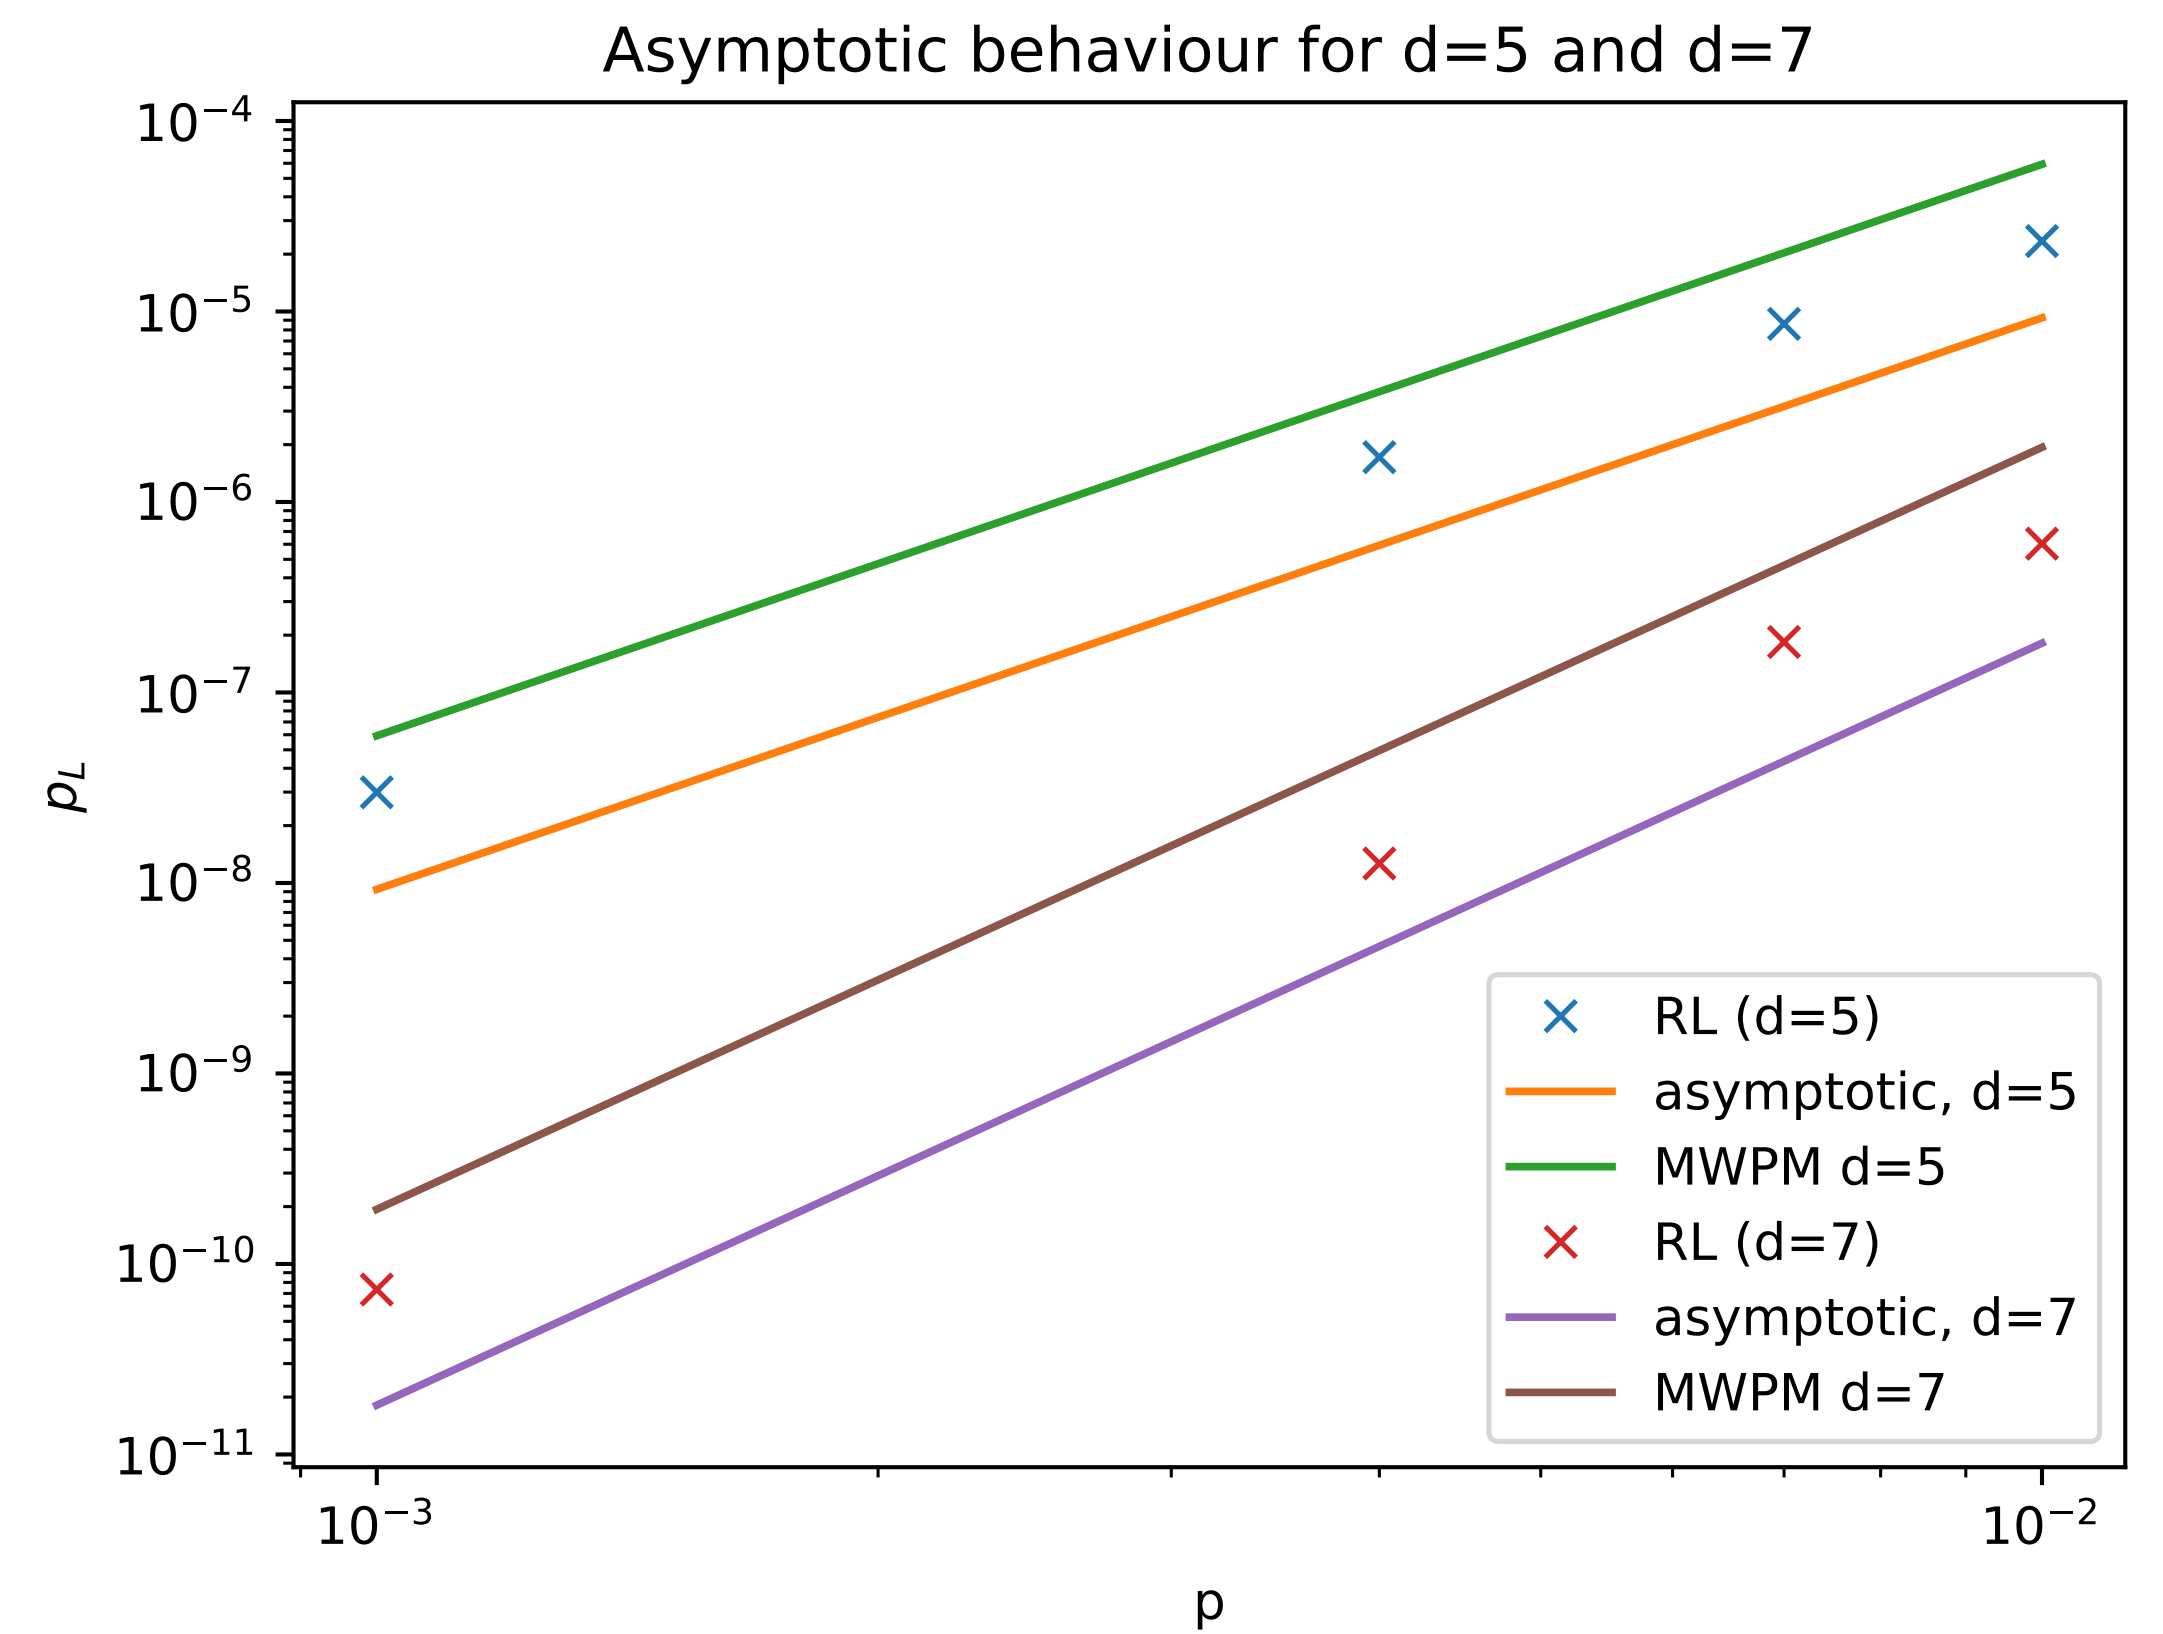
No logical operation

The agent can use Y to take advantage of correlations between bit-flip and phase-flip errors

# Preliminary performance of RL solver for depolarizing noise



## Outperforms MWPM



# Deep Q-networks

## distance 5 code

Layer (type)	Output Shape	Param #
Conv2d-1	[-1, 128, 5, 5]	2,432
Conv2d-2	[-1, 128, 5, 5]	147,584
Conv2d-3	[-1, 120, 5, 5]	138,360
Conv2d-4	[-1, 111, 5, 5]	119,991
Conv2d-5	[-1, 104, 5, 5]	104,000
Conv2d-6	[-1, 103, 5, 5]	96,511
Conv2d-7	[-1, 90, 5, 5]	83,520
Conv2d-8	[-1, 80, 5, 5]	64,880
Conv2d-9	[-1, 73, 5, 5]	52,633
Conv2d-10	[-1, 71, 5, 5]	46,718
Conv2d-11	[-1, 64, 3, 3]	40,960
Linear-12	[-1, 3]	1,731
Total params: 899,320		

trained on desktop GPU for 5 hours  
(using PyTorch)

## distance 7 code

Layer (type)	Output Shape	Param #
Conv2d-1	[-1, 200, 7, 7]	3,800
Conv2d-2	[-1, 190, 7, 7]	342,190
Conv2d-3	[-1, 189, 7, 7]	323,379
Conv2d-4	[-1, 160, 7, 7]	272,320
Conv2d-5	[-1, 150, 7, 7]	216,150
Conv2d-6	[-1, 132, 7, 7]	178,332
Conv2d-7	[-1, 128, 7, 7]	152,192
Conv2d-8	[-1, 120, 7, 7]	138,360
Conv2d-9	[-1, 111, 7, 7]	119,991
Conv2d-10	[-1, 104, 7, 7]	104,000
Conv2d-11	[-1, 103, 7, 7]	96,511
Conv2d-12	[-1, 90, 7, 7]	83,520
Conv2d-13	[-1, 80, 7, 7]	64,880
Conv2d-14	[-1, 73, 7, 7]	52,633
Conv2d-15	[-1, 71, 7, 7]	46,718
Conv2d-16	[-1, 64, 5, 5]	40,960
Linear-17	[-1, 3]	4,803
Total params: 2,240,739		

trained on desktop GPU for 12 hours

Unnecessarily deep?

# Conclusions

**Deep Q-learning works well for error correction on *toric* code.  
Can match or even outperform MWPM (for moderate code distance)**

**But, does require quite deep Q-networks**

**Periodic boundaries important for our approach.**

## **Future challenges:**

- **Larger code distances**
- **Improve reward scheme, use actual success or failure of error correction**
- **Include syndrome measurement error. (R. Sweke et al, arXiv:1810.07207)**
- **Surface code with boundaries. (Tougher due to lack of translational invariance)**

Philip Andreasson, Joel Johansson, Simon Liljestrand, Mats Granath, arXiv:1811.12338

Mattias Eliasson, David Fitzek, MG, in progress