Searching for Majorana Zero Modes Using Model-free Reinforcement Learning

Makhamisa Senekane[a], **Mhlambululi Mafu**[b] **and Molibeli Taele**[a]

[a]Department of Physics and Electronics, National University of Lesotho, Lesotho

[b]Department of Physics and Astronomy, Botswana International University of Science and Technology, Palapye, Botswana

E-mail: `mc.senekane@nul.ls, mafum@biust.ac.bw bm.taele@nul.ls`

Abstract. Majorana fermions are particles which are their own antiparticles; hence they have zero charge. They are governed by non-Abelian statistics. For a Majorana fermionic operator γ, and the Hamiltonian of a system H, Majorana fermions satisfy fermionic anti-commutation relation (that is, for a pair of Majorana operators γ_i, γ_j; $\{\gamma_i, \gamma_j\} = 2\delta_{ij}$) and a Majorana fermion squares to 1 (that is, $\gamma^2 = 1$). If, in addition to this, the fermionic operator commutes with the Hamiltonian of the system (that is, $[\mathrm{H},\gamma] = 0$), then such an operator is a Majorana zero mode (MZM). Majorana zero modes are Majorana fermions bound to zero energy. MZMs have applications in both topological quantum computation and spintronics. However, Majorana zero modes are yet to be conclusively demonstrated experimentally. In this work, we report the algorithm that searches for MZMs using reinforcement learning. Reinforcement learning is a machine learning paradigm where the learner is a decision-maker (agent) that takes action in an environment and receives rewards or penalties for the actions taken. Results obtained from this work demonstrate the significance of using reinforcement learning in the quest for Majorana zero modes.

1. Introduction

Majorana fermions; which are also affectionately referred to as Majoranas, are named after an Italian physicist Ettore Majorana, who first proposed their existence as a real solution to Dirac's quantum relativistic equation of spin-$\frac{1}{2}$ particles (spinors) [1]. Majoranas are particles which are their own anti-particles, and follow neither the Fermi-Dirac nor the Bose-Einstein exchange statistics [1, 2, 3, 4, 5].

Originally, Majoranas were proposed in the context of particle physics [1, 6]. On the other hand, Majoranas in condensed matter physics exist as exotic quasi-particle excitations which are their own anti-particles. Thus, unlike in particle physics; where Majoranas are proposed as fundamental particles, Majoranas in condensed matter systems are emergent. Such Majoranas (the latter) have applications in topological quantum computation [7] and spintronics [8].

Recently, the search for Majoranas in condensed matter systems is an active area of research, even though experimental observation of Majoranas is still a challenge. most of research in this area is geared towards theoretical predictions of Majoranas in various states of matter. In this manuscript, we propose an algorithm that can be used to search for Majoranas, as a tool that can be used to further investigate the presence of Majoranas in condensed matter systems. This algorithm uses machine learning; which is the sub-field of artificial intelligence [9]. The key contribution of the work reported in this manuscript is the development of a machine learning-based tool (algorithm) that can be used to search for Majoranas.

The remainder of this structured as follows. The next section provides background information on Majorana fermions and machine learning; including reinforcement learning, which is the machine learning paradigm used in the work reported in this manuscript. Section 3 discusses the implementation of model-free reinforcement learning algorithm. This is followed by Section 4, which presents and discusses the results obtained. The last section concludes this manuscript.

2. Background Information

2.1. Bosons, Fermions and Anyons

In three spatial dimensions, particles can exist as either bosons or fermions, depending on exchange statistics involved [7]. For bosons, the wave function is symmetric while that of a fermion is anti-symmetric. Examples of bosons are photons and super-fluids. On the other hand, examples of fermions are Fermi liquids and electrons.

When particles are confined to two spatial dimensions, the boson-fermion dichotomy mentioned above breaks off. Exotic particles called anyons; which are neither bosonic nor fermionic, can be formed is such a configuration (two spatial dimensions) [7, 10, 11]. Majoranas are the simplest manifestations of anyons. Majoranas are Hemitian [6]. Thus, for a Majorana fermionic operator γ:

$$\gamma_j = \gamma_j \dagger . \tag{1}$$

Additionally, Majorana fermionic operators satisfy the anti-commutator relation [11]:

$$\{\gamma_i, \gamma_j\} = 2\delta_{ij}. \tag{2}$$

Furthermore, unlike conventional fermions, Majorana fermions square to 1:

$$\gamma^2 = 1. \tag{3}$$

2.2. Machine Learning

Machine learning is a set of statistical techniques that enable computers to learn from data without being explicitly programmed [9, 12]. It enables a computer to identify patterns in observed data, build models that explain the world, and predict things without having explicit pre-programmed rules and models. A computer is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E [13].

Machine leaning can be divided into four paradigms, namely, supervised learning [9, 12, 13], unsupervised learning [9, 12], reinforcement learning [9, 13, 14], and generative adversarial network (GAN) [15]. Reinforcement learning is the machine learning paradigm that was used in the work reported in this manuscript. It is concerned with multi-stage decision making with the goal of maximizing the reward. Furthermore, this paradigm involves a learning agent who interacts with the environment, takes actions based on the state of the environment, and in return gets the cumulative reward from the environment. An example of a reinforcement learning algorithm is Q-learning [14]. Figure 1 shows a schematic diagram of a typical reinforcement learning setting.

Elements of reinforcement learning include [14]:

- policy: a mapping from perceived states of the environment to actions to be taken when in those states
- reward signal: defines the goal in a reinforcement learning problem. It defines what are the good and bad events for the agent

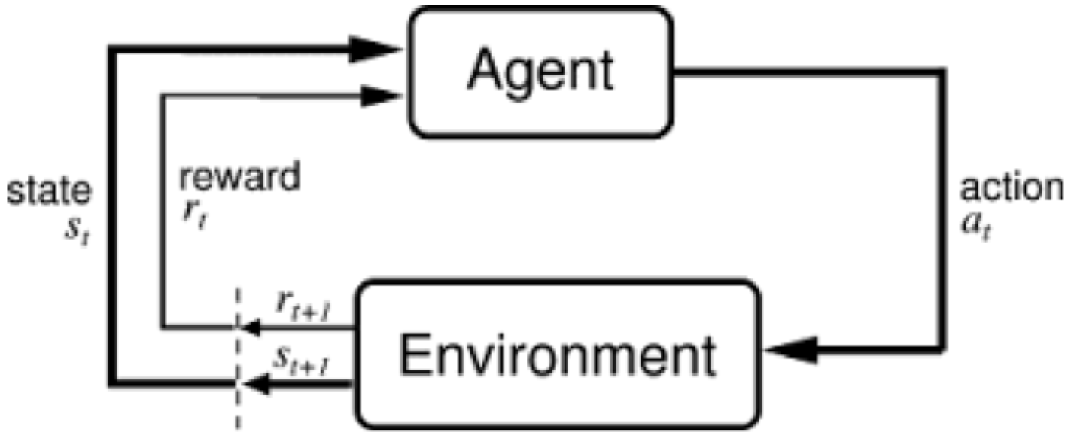

Figure 1. A typical setting for reinforcement learning paradigm. An agent interacts with the environment by taking an action based on the state of the environment. In return, the environment rewards the agent. The ultimate goal is the maximization of cumulative reward by the agent.

- value function: defines what is good in the long run. It is a state is the total amount of reward an agent can expect to accumulate over the future, starting from that state
- model(optional): mimics the behavior of the environment. Therefore, model-based reinforcement learning models dynamics of the environment while on the other hand, model-free reinforcement learning does not

3.

The model-free reinforcement learning algorithm was implemented using the following:

- HP-630 laptop running Ubuntu 17.10 Ubuntu operating system
- Python 3.6 programming language[1]
- Tensorflow framework[2] as a back-end
- Keras[3]
- deep Q-network using a three-layer multi-layer perceptron (MLP) artificial neural network
- OpenAI gym[4] was used for testing the algorithm

Data used in this work was collected from literature, where existence of Majoranas in certain states of matter was proposed. Furthermore, randomly generated states were created.Then using data collected, a reinforcement learning algorithm was developed, using the procedure depicted in Figure 2.

4. Results and Discussion

A multi-layer perceptron artificial neural network's accuracy was found to be 94%. This high accuracy underlines the utility of using MLP for approximating Q function, instead of just using Q-table; which is static. Additionally, other neural network configurations that were attempted

[1] www.python.org
[2] www.tensorflow.org
[3] www.keras.io
[4] www.gym.openai.com

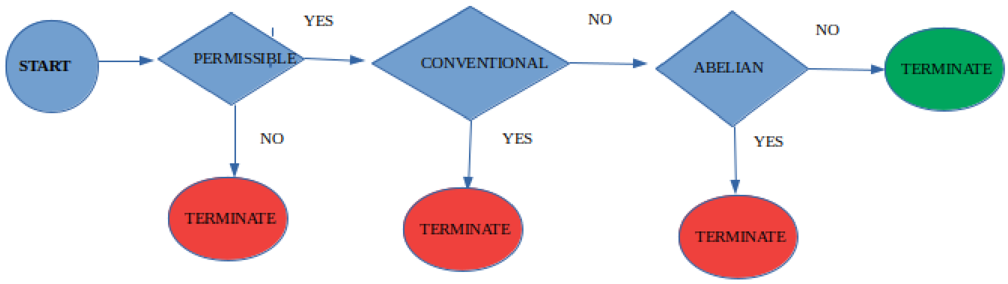

Figure 2. A sequence of decision-making steps that the model-free reinforcement learning algorithm takes to search for Majoranas. First, it tests whether the state of matter is physically permissible or not. Then it tests whether the state is conventional or topological. Further, it tests whether exchange statistics are Abelian or non-Abelian.

performed worse than the three-layer MLP, rendering them less useful than the three-layer MLP for Q function approximation.

Furthermore, average rewards were plotted against the number of episodes, for different values of ϵ[5]. The maximum number of episodes was set to 100. Figures 3 and 4 show the results obtained for ϵ=0.9 and ϵ=0.5 respectively.

5. Concluding Remarks

We have reported the implementation of model-free reinforcement learning scheme that can be used to search for Majoranas. This scheme uses deep Q network with a three-layer multi-layer perceptron for Q function approximation. The results obtained underline the utility of using this technique as a was of fast-tracking discovery of Majoranas in condensed matter systems. However, due to limited computational capabilities, collection of data relied heavily on what is reported in the literature. Future work envisages the use of Monte Carlo technique to generate random states of matter. Additionally, future work will focus on the use of GAN scheme instead of reinforcement learning paradigm. GAN paradigm is still relatively recent, and it provides a promising alternative to reinforcement learning.

Acknowledgments

Makhamisa Senekane and Benedict Molibeli Taele acknowledge support from NUL Research and Conferences Committee. Mhlambululi Mafu thanks his colleagues at BIUST for their unwavering support and constructive discussions.

References

[1] Wilczek F 2009 *Nature Physics* **5** 614
[2] Sarma S D, Freedman M and Nayak C 2015 *npj Quantum Information* **1** 15001
[3] Lutchyn R, Bakkers E, Kouwenhoven L, Krogstrup P, Marcus C and Oreg Y 2017 *Nature Reviews Materials* **3** 52–68
[4] Beenakker C 2013 *Annu. Rev. Condens. Matter Phys.* **4** 113–136
[5] Alicea J 2012 *Reports on progress in physics* **75** 076501
[6] Leijnse M and Flensberg K 2012 *Semiconductor Science and Technology* **27** 124003

[5] ϵ in this case is for the ϵ-greedy algorithm used

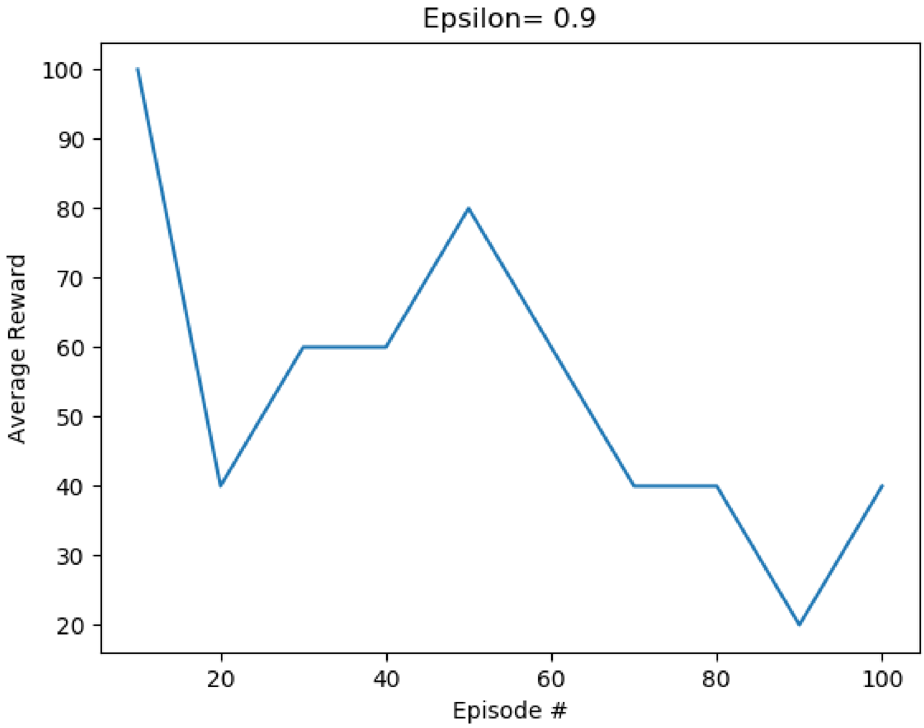

Figure 3. Average rewards versus number of episodes for ϵ=0.9.

[7] Nayak C, Simon S H, Stern A, Freedman M and Sarma S D 2008 *Reviews of Modern Physics* **80** 1083
[8] Wu C T, Anderson B M, Hsiao W H and Levin K 2017 *Physical Review B* **95** 014519
[9] Russell S J and Norvig P 2016 *Artificial intelligence: a modern approach* (Malaysia; Pearson Education Limited,)
[10] Wilczek F 1990 *Fractional statistics and anyon superconductivity* vol 5 (World scientific)
[11] Stern A 2008 *Annals of Physics* **323** 204–249
[12] Marsland S 2011 *Machine learning: an algorithmic perspective* (Chapman and Hall/CRC)
[13] Mohri M, Rostamizadeh A and Talwalkar A 2012 *Foundations of machine learning* (MIT press)
[14] Sutton R S, Barto A G *et al.* 2018 *Reinforcement learning: An introduction* (MIT press)
[15] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 *Advances in neural information processing systems* pp 2672–2680

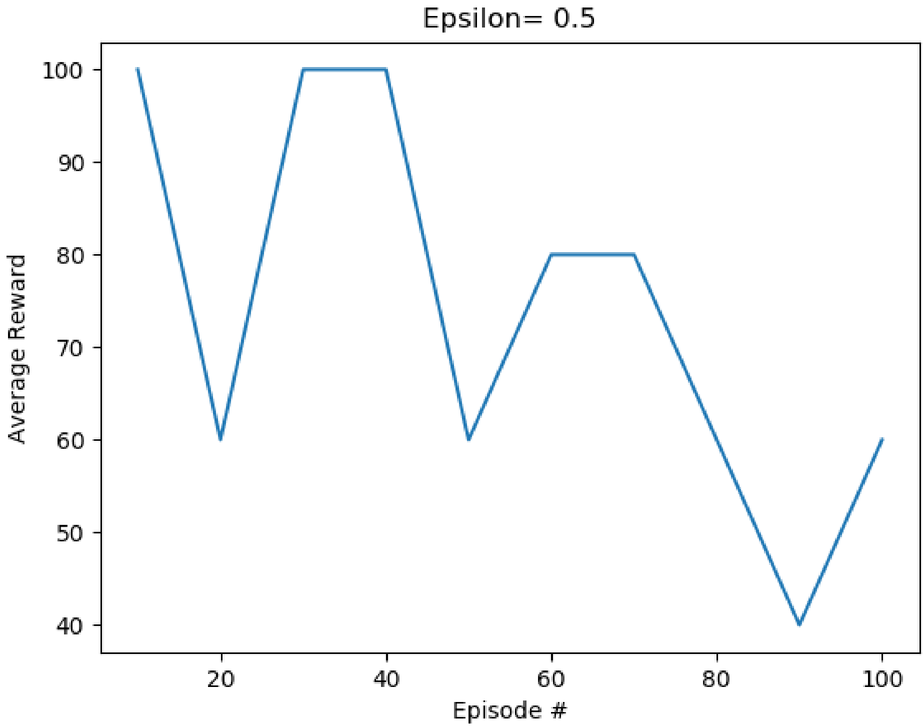

Figure 4. Average rewards versus number of episodes for ϵ=0.5.