Contents lists available at ScienceDirect

# **Physics Letters B**

journal homepage: www.elsevier.com/locate/physletb

# BPS spectroscopy with reinforcement learning

Federico Carta<sup>a,b</sup>, Asa Gauntlett<sup>a,c</sup>, Finley Griffin<sup>a,d</sup>, Yang-Hui He<sup>a,e,,,\*</sup>

<sup>a</sup> London Institute for Mathematical Sciences, Royal Institution, London, W1S 4BS, UK

<sup>b</sup> Physics Department, King's College London, Strand, London, WC2R 2LS, UK

<sup>c</sup> Department of Computer Science, University College London, Gower Street, London, WC1E 6BT, UK

<sup>d</sup> Hertford College, University of Oxford, Catte Street, Oxford, OX1 3BW, UK

e Merton College, University of Oxford, Oxford, OX1 4JD, UK

#### ARTICLE INFO

Editor: F. Bo

# ABSTRACT

We apply reinforcement learning (RL) to establish whether at a given position in the Coulomb branch of the moduli space of a 4d  $\mathcal{N} = 2$  quantum field theory (QFT) the BPS spectrum is finite. If it is, we furthermore determine the full BPS spectrum at such point in moduli space. We demonstrate that using a RL model one can efficiently determine the suitable sequence of quiver mutations of the BPS quiver that will generate the full BPS spectrum. We analyse the performance of the RL model on random BPS quivers and show that it converges to a solution various orders of magnitude faster than a systematic brute-force scan. As a result, we show that our algorithm can be used to identify all minimal chambers of a given  $\mathcal{N} = 2$  QFT, a task previously intractable with computer scanning. As an example, we recover all minimal chambers of the SU(2)  $N_f = 4$  gauge theory, and discover new minimal chambers for theories that can be realized by IIB geometric engineering.

#### 1. Introduction and summary

The Seiberg-Witten solution of 4d  $\mathcal{N} = 2$  quantum field theories (QFT)s determines the low energy dynamics on the Coulomb Branch of the moduli space of vacua [1,2]. One crucial feature of the QFT which is not implied by the Seiberg-Witten solution is the spectrum of BPS states at a generic point of the Coulomb Branch, and how such spectrum is modified by moving in moduli space.

Determining the BPS spectrum of theories with  $\mathcal{N} = 2$  supersymmetry is non-trivial and yet seemingly tractable in numerous examples. A subset of 4d  $\mathcal{N} = 2$  theories admits a BPS quiver [3,4], from which the BPS spectrum can be determined via various techniques. In this letter we consider theories whose spectrum are described by a BPS quiver. We remark that of these, the problem becomes even simpler for a subclass of which are dubbed *complete* in [3,4]. Such  $\mathcal{N} = 2$  models are defined by the rather stringent property that even as one varies all parameters such as moduli, couplings and bare masses, the number of hypermultiplets which constitute a basis of the BPS spectrum remains to be equal to the rank of the lattice of electromagnetic and flavour charges  $\Gamma$ .

An example of a complete QFT is that obtained from two coincident M5-branes on a punctured Riemann surface, which gives us the theory constructed by Gaiotto [5], composed of a product of SU(2)

gauge groups with tri-fundamental matter. Various other theories of the class S construction fall into the set of complete QFTs. In addition to these, Cecotti and Vafa [3] identified 11 exceptional cases, related to Dynkin diagrams for affine and elliptic exceptional groups as well as the Derksen-Owen quivers from [6].

Now, as we enter the age of AI [7], it was inevitable that machine learning (ML) methods should enter theoretical physics, especially in high-energy theory, and pure mathematics: this was introduced in [8–11] (for pedagogical introductions, see [12,13] and for recent reviews on this "AI-guided theoretical discovery", see [14,15]).

The use of ML is by now well-established in string theory and related sub-fields of mathematical physics such as conformal field theory [16,17],  $\mathcal{N} = 1$  quiver QFTs [18], cluster algebras [19,20], and BPS indices [21]. It is therefore natural and expedient to see whether AI, and in particular ML methodologies, can help with our present problem of understanding BPS spectra and the classification of  $\mathcal{N} = 2$  gauge theories.

This above question we address in the present paper, whose main results we summarize here:

1. As a main result, we develop a reinforcement learning (RL) algorithm which efficiently determines a sequence of quiver mutations

\* Corresponding author.

https://doi.org/10.1016/j.physletb.2025.139646

Received 6 February 2025; Received in revised form 30 April 2025; Accepted 4 June 2025

Available online 6 June 2025



Letter



Keywords: BPS spectrum BPS quiver Quiver mutation Reinforcement learning

*E-mail addresses:* federico.carta@bo.infn.it (F. Carta), asagauntlett1@gmail.com (A. Gauntlett), fingriffin2023@outlook.com (F. Griffin), hey@maths.ox.ac.uk, yang-hui.he@merton.ox.ac.uk (Y.-H. He).

<sup>0370-2693/© 2025</sup> The Author(s). Published by Elsevier B.V. Funded by SCOAP<sup>3</sup>. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

that identifies the full BPS spectrum of a complete QFT, in any finite chamber of the moduli space. The code is available on GitHub.

- 2. We apply this algorithm to the SU(2)  $N_f = 4$  theory. We determine the number of inequivalent chambers in the Coulomb Branch for which the BPS spectrum is finite and minimal. To demonstrate the efficiency of our algorithm, we determine the BPS spectrum in all such chambers. These are also included in the GitHub.
- 3. As another example, we identify minimal chambers and corresponding BPS spectra for theories whose BPS quiver is given by the elliptic *E*-type Dynkin diagrams,  $\hat{E}_n$  with n = 6, 7, 8 as well as the theories identified by the Derksen-Owen quivers  $X_6, X_7$  [6].

Our results represent a tremendous improvement to the efficiency and scalability of conventional methods for determining finite chambers and BPS spectra of  $\mathcal{N} = 2$  QFTs. Such a tool should aid in the current problem of the classification of  $\mathcal{N} = 2$  models and, furthermore, the study of their non-perturbative properties. For example, the mutation sequences identified in each case may be used as input data to the Cordova-Shao algorithm [22] to determine their Schur index.

It is also interesting to try and explain our conjecture for the number of minimal chambers, see Fig. 3, for the SU(2)  $N_f = 4$  model from study of the underlying physics or pure mathematics. Finally, our computations may also be extended to compute non-minimal, or even maximal, finite chambers of  $\mathcal{N} = 2$  QFTs.

#### 2. Background

We begin with a brief overview on how the BPS spectrum of complete theories can be derived from the BPS quiver.

### 2.1. BPS quivers of $4D \mathcal{N} = 2$ quantum field theories

Let us consider a 4d  $\mathcal{N} = 2$  QFT with a gauge group of rank r and flavor symmetry algebra of rank f. The moduli space of supersymmetric vacua is split into a Coulomb branch  $\mathcal{M}_C$  and a Higgs branch  $\mathcal{M}_H$ . On a generic vacua  $u \in \mathcal{M}_C$  the gauge group is broken to  $U(1)^r$ . The lattice of electric, magnetic and flavour charges  $\Gamma$  is of dimension 2r + f and is equipped with a linear function

$$Z_u: \Gamma \to \mathbb{C} \,, \tag{1}$$

called the central charge of the theory.

The complex plane, where  $Z_u$  takes values, is split by choice into two half-planes  $\mathbb{C}_{\pm}$ : states with charges  $\gamma \in \Gamma$  such that  $Z_u(\gamma) \in \mathbb{C}_+$  (respectively,  $Z_u(\gamma) \in \mathbb{C}_-$ ) are labelled as particles (respectively, antiparticles) of the theory. Furthermore, the charge lattice is equipped with an antisymmetric inner product

$$\circ: \Gamma \times \Gamma \to \mathbb{Z} \,. \tag{2}$$

At a point *u* on the Coulomb Branch, the  $\mathcal{N} = 2$  superalgebra imposes the mass constraint on particles with charge  $\gamma \in \Gamma$ 

$$M \ge |Z_u(\gamma)| \,. \tag{3}$$

Particles which saturate this inequality are called *BPS particles*. For every point  $u \in \mathcal{M}_C$ , a relevant and important problem is finding the full spectrum of BPS particles, including as well their spin and multiplicity. For a subset of the 4*d*  $\mathcal{N} = 2$  theories, such a question can be answered with the auxiliary tool of a BPS quiver.

The BPS quiver encodes a 1*d* unitary quiver gauge theory living on the worldline of BPS particles that describes their dynamics. It can be defined as follows. Let  $\{\gamma_i\}$  be the charges of the hypermultiplets which form a basis of the BPS spectrum. Then: For each pair of charges in the basis, compute the electric-magnetic inner product γ<sub>i</sub> ∘ γ<sub>j</sub>. If γ<sub>i</sub> ∘ γ<sub>j</sub> > 0, the nodes γ<sub>i</sub> and γ<sub>j</sub> are connected with γ<sub>i</sub> ∘ γ<sub>i</sub> arrows,<sup>1</sup> each pointing from node *j* to node *i*.

From the BPS quiver alone it is possible to determine the BPS spectrum at any point in the moduli space. One way of doing so consists in determining the stability conditions of various quiver representations, each one corresponding to a particle of the QFT. We will not include a detailed discussion of this method, and instead refer the reader to [23]. A much simpler method involves a computation of a sequence of quiver mutations, and is dubbed the *mutation method* [23]. We will review this method in the next subsection.

#### 2.2. Quiver mutation and finite chambers

If at a point  $u \in \mathcal{M}_C$  the BPS spectrum is finite, the spectrum can be fully determined by the so-called mutation method [23]. A *quiver mutation* is an operation <sup>2</sup> that is applied at a node of a BPS quiver Q. After the mutation, a new BPS quiver  $\tilde{Q}$  is produced. The new quiver  $\tilde{Q}$  will have the same number of nodes, although now labelled with different charges  $\tilde{\gamma}_i$ , and a different set of arrows.

In order to find the full BPS spectrum one needs to apply a sequence of quiver mutations (generically acting on different nodes) such that at the end of the process, the final quiver produced will have all charges flipped compared to the original quiver and all arrows inverted. We call such end-point of the mutation method the *antiparticle quiver*. Keeping track of all the charges that appear at all nodes in every intermediate steps is equivalent to determine the full BPS spectrum. We will recall below the rules defining the mutation operation. To define the mutation, let us suppose that we mutate at the node labelled by 1, and whose corresponding charge is  $\gamma_1$ . Let us call Q the original quiver, and  $\tilde{Q}$  the quiver after the mutation. The procedure to derive  $\tilde{Q}$  from Q is then the following:

- The nodes of the mutated quiver Q
   *Q* are the same of the nodes of the quiver Q. Also, for every arrow in the quiver Q, write the correspondent arrow in the quiver Q
   *Q*.
- For each length two path of arrows passing through node 1 in Q, draw a new arrow in Q
   connecting the initial and final node of the two steps path.
- 3. Invert all arrows in  $\tilde{Q}$  which end on node 1.
- If between any two nodes *i*, *j* of the quiver Q̃ there is both an arrow *i* → *j* and one *j* → *i*, delete both of them. Repeat this procedure until no such couple of lines with opposite orientation exist anymore, for all couple of nodes.
- 5. The new basis is given by

$$\widetilde{\gamma}_1 = -\gamma_1$$
 (4)

$$\widetilde{\gamma}_{j} = \begin{cases} \gamma_{j} + (\gamma_{j} \circ \gamma_{1})\gamma_{1} & \text{if } \gamma_{j} \circ \gamma_{1} > 0\\ \gamma_{j} & \text{if } \gamma_{j} \circ \gamma_{1} \le 0. \end{cases}$$
(5)

There are further rules to derive the superpotential of the mutated quiver  $\tilde{Q}$ , but we will not make use of those in this paper.

We have reviewed here the mutation mechanism to determine the BPS spectrum from the BPS quiver. Such a method is of course much

<sup>•</sup> Each element  $\gamma_i$  in the basis corresponds one node of the quiver.

<sup>&</sup>lt;sup>1</sup> If  $\gamma_i \circ \gamma_j < 0$ , then we have no arrows. Note that this convention is slightly different from that in the quiver gauge theory literature where the antisymmetrized adjacency matrix is used so that a negative intersection means an arrow going the other way.

<sup>&</sup>lt;sup>2</sup> In the mathematics literature, this was coined by Fomin-Zelevinsky in the context of cluster algebras [24]. Interestingly, in physics, this was realized independently around the same time [25] as Seiberg duality or toric duality. It was only years later that at a Oberwolfach workshop that the two communities realized that they were talking about the same thing.

Cumulative Reward over Time;  $SU(2) N_f = 4$ 



**Fig. 1.** Cumulative reward ( $\lambda = 0.1$ ) against mutation for a successful training episode on the BPS quiver  $Q_1$  of the SU(2)  $N_f = 4$  SCFT, from Fig. 3, identifying a finite chamber containing 12 states.



Algorithm Performance on Random Cyclic BPS Quivers

Fig. 2. Plot of the average convergence rate to a finite chamber against number of nodes of a random BPS quiver for a smart walk (SW) and reinforcement learning model (RL).

more efficient than the method of studying representation theory of the BPS quiver. However, crucially, for more complicated quivers guessing by eye the right sequence of mutations is very challenging, and also an extremely inefficient problem to be treated by a computer code which uses a brute-force algorithm. We will solve both of these issues by using Reinforced Machine Learning (RL) to compute the BPS spectrum in finite chambers. In the following sections, we outline the RL algorithm we use.

# 3. Methods

In this work, we use a Proximal Policy Optimisation (PPO) algorithm to address the problem of finding finite chambers of  $\mathcal{N} = 2$  gauge theories. PPO is a model-free reinforcement learning (RL) involving an optimal policy  $\pi^* : S \to \mathcal{A}$  that maps any state  $s \in S$  to an action  $a \in \mathcal{A}$ . Given a state  $s_t$  at discrete timestep t, the PPO algorithm selects an action  $a_t$ , computes a reward and moves to a subsequent state  $s_{t+1}$ . The policy  $\pi$  is updated based on the clipped objective function and

is represented by a deep neural network (NN). The NN we used for the optimisation procedure consists of three hidden layers of 252, 504 and 252 neurons with ReLU activation. The NN was trained using the Adam optimiser and the learning rate  $\alpha$  and discount factor  $\delta$  were set to 0.5 and 0.995, respectively. Further details of RL and PPO can be found in Appendix A.

In our investigation, training is initialized with an *n*-node BPS quiver  $s_0 \in S$ , with S the space of BPS quivers equivalent to  $s_0$  by quiver mutation. A state is represented by an  $n \times 2n$  matrix ( $M \mid G$ ), where M is the adjacency matrix of the underlying digraph, and G, the 'gamma matrix', encodes quiver labels as linear combinations of BPS hypermultiplets { $\gamma_i$ }, i.e.

node label 
$$i = \sum_{j} G_{ij} \gamma_j$$
;  $G_{ij} \in \mathbb{Z}$ . (6)

We illustrate the matrices M and G with the following:

**Example 1.** Let  $s_0$  be the BPS quiver  $Q_1$  for the SU(2)  $N_f = 4$  gauge theory [23] from Fig. 3. Training is initialized with the data structure

The adjacency matrix M adopts the standard convention, where an arrow from nodes  $i \rightarrow j$  is represented by  $M_{ij} = +1$ ,  $M_{ji} = 0$ . The gamma matrix G of the initial quiver  $s_0$  is always set to the identity matrix. Here, the *i*-th node label is simply the hypermultiplet  $\gamma_i$  alone. In general, intermediate quivers will have off-diagonal elements in G, representing bound states of  $\{\gamma_i\}$ . Now, after the first step in the mutation, there is a bound state  $\gamma_5 + \gamma_6$ , so G has a 1 in the entry (5, 6).

We define the action space A as all possible nodes that can be mutated on at a given time-step. Illegal nodes are those with labels that contain negative coefficients of the hypermultiplets, or that have already been mutated on in the action history. For this reason, the action space is dynamic, which makes the use of a PPO algorithm especially suitable,<sup>3</sup> as it demonstrates notable stability with action masking compared to other RL algorithms [26].

Finally, we define a reward function R(s) that provides feedback for beneficial actions and a terminal reward (+1) when the antiparticle quiver is reached. A fitness function  $f : S \rightarrow [0, 1]$  quantifies the proximity of a given intermediate state  $s_t$  to the antiparticle quiver. We define two distinct fitness functions,  $f_M$  and  $f_G$ , which separately evaluate this proximity in terms of the graph structure and node labels, respectively. The reward function is then

$$R(s) = 1 - \lambda f_M(s) - (1 - \lambda) f_G(s),$$
(8)

where  $\lambda \in [0, 1]$  is a tunable parameter representing the relative weighting of structural and label similarities in the reward logic.

**Example 2.** Fig. 1 shows the cumulative reward for a successful episode on the BPS quiver of the SU(2)  $N_f = 4$  SCFT from Example 1, identifying the 12-state chamber

$$\gamma_{3}, \gamma_{4}, \gamma_{2}, \gamma_{1} + \gamma_{3} + \gamma_{4}, \gamma_{2} + \gamma_{5}, \gamma_{2} + \gamma_{6}, \gamma_{1} + \gamma_{3}, \gamma_{1} + \gamma_{4}, \gamma_{2} + \gamma_{5} + \gamma_{6}, \gamma_{1}, \gamma_{5}, \gamma_{6}$$

$$(9)$$

in agreement with the result from [23].

#### 4. Results

To demonstrate the capability of our RL model to find a finite chamber for a generic theory, we established a dataset of cyclic digraphs, ranging from 3 - 10 nodes, from which random BPS quivers are drawn. We then trained the RL algorithm on a random selection of these quivers, and recorded the convergence rate of the procedure as the number of time-steps to identify a finite chamber for a given quiver. A negative penalty is applied to episodes that exceed a maximum length or label complexity to encourage further exploration by the agent. As a baseline for performance evaluation, we use a smart walker (SW) algorithm that combines a random walk search with an exhaustive scan. This algorithm is equivalent to the RL algorithm with learning disabled. A log plot of the average convergence rate against number of nodes is shown in Fig. 2.

Some BPS quivers do not have finite chambers, which are identified by allowing the SW algorithm to run until it has exhausted all possible



Fig. 3. The four BPS quivers of the SU(2)  $N_f = 4$  SCFT and the number of minimal chambers identified by the RL model.

mutation sequences. Of the quivers with finite chambers, the RL algorithm achieves a 100% accuracy and, moreover, is orders of magnitude more efficient than the SW algorithm. These results mean that the RL algorithm can efficiently identify finite chambers even for large BPS structures, where traditional exhaustive scans are intractable as the number of possible sequences scales as ~ (number of nodes)<sup>length of sequence</sup> [23].

For this reason, Fig. 2 includes projections for the SW algorithm on 9 and 10 node quivers at  $\mathcal{O}(10^9)$  and  $\mathcal{O}(10^{10})$  time-steps, respectively. This would correspond to a computing time of ~ 10 days and ~ 1 year, respectively, on a home machine. We report an average of  $\mathcal{O}(10^4)$  and  $\mathcal{O}(10^5)$  time-steps with the RL algorithm, or a computing time of ~ 5 minutes and ~ 20 minutes, for 9 and 10 node quivers respectively.

Furthermore, the RL algorithm is inclined to identify minimal chambers of the theory due to the discount factor  $\delta$  in the *Q*-function, see Appendix A. For instance, the 12-state chamber identified in Example 2 is indeed a minimal chamber of the theory. This is not a feature of the SW model, where the length of the chamber is needed a priori, and there is no guarantee that any finite chamber identified is a minimal one.

#### 4.1. Counting minimal chambers

As discussed in [27], it is interesting to find all the minimal chambers for all the BPS quivers of a given theory. To further showcase the capabilities of our model, we offer the original computation of all minimal chambers of the SU(2)  $N_f = 4$  gauge theory.

Minimal chambers for this theory have been identified many times [4,23,27,28], but the scale of an exhaustive computational scan means that a complete census is absent in the literature. For the SU(2)  $N_f = 4$  theory, the Coulomb branch is spanned by four distinct BPS quivers, see Fig. 3, corresponding to the four possible triangulations of the fourth-punctured sphere [3,23]. These can be verified using Keller's mutation Java applet [29], as well as using our code. We note that for each finite

<sup>&</sup>lt;sup>3</sup> Our experiments showed that Deep-Q Learning proves much less stable with dynamic action masking in this context.



**Fig. 4.** Plot of the total number of distinct minimal chambers of SU(2)  $N_f = 4$  found against the number of steps taken for a random walk search using the RL model with PPO.

chamber, there exists a set of equivalent chambers obtained from the symmetries of the quiver. One such symmetry is permutation of topologically equivalent nodes (formally an automorphic equivalence in graph theory). For  $Q_1$  in Fig. 3, the nodes are arranged into classes

$$\{\gamma_1\}, \{\gamma_2\}, \{\gamma_3, \gamma_4, \gamma_5, \gamma_6\}$$
 (10)

which introduces a 4! = 24-fold degeneracy to each chamber found for this quiver. The quivers  $Q_2$  and  $Q_3$  in Fig. 3 also have global symmetries that introduce additional degeneracies, see Appendix B. In the counting procedure, the RL model does not identify these as distinct chambers, and so the results in Fig. 3 and Fig. 4 should be interpreted as the number of essentially unique minimal chambers of each quiver.

For this experiment, the RL algorithm is adjusted to apply a negative penalty upon visiting a known chamber, which encourages the agent to further explore the moduli space and identify new minimal chambers. We run the searches until no new chambers are found in  $\mathcal{O}(10^6)$  timesteps, which occurs after  $\mathcal{O}(10^6 \sim 10^7)$  time-steps. The total identified for each quiver is shown in Fig. 3, up to symmetries, and the full lists of these chambers are in the GitHub. Each chamber can also be manually checked using the interactive quiver plot. The total number of minimal chambers found by the agent against total time-steps is shown in Fig. 4.

#### 4.2. Exceptional complete theories

The  $\mathcal{N} = 2$  models engineered from punctured Riemann surfaces [5] are all but finitely many of the complete theories with BPS quivers. We now divert our attention to identifying minimal chambers of the following exceptional complete theories, which are not constructed by the triangulation of a Riemann surface but rather by type IIB geometric engineering [3,4]:

- The three theories identified with the elliptic *E*-type Dynkin diagrams,  $\hat{E}_n$  with n = 6, 7, 8 [30].
- The two theories identified by the Derksen-Owen quivers,  $X_6$  and  $X_7$  [6].

The RL algorithm identifies minimal chambers of these theories in  $\mathcal{O}(10^2 \sim 10^3)$  timesteps, which are shown with their BPS quivers in Fig. 5. In agreement with results from [4], we find that the  $X_7$  quiver has no finite chamber.

#### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Yang-Hui He reports financial support was provided by Leverhulme Trust. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgements

FC would like to thank the Banff International Research Station and Cornell University for hospitality during the days in which this project was completed. AG would like to thank Alex Bloch for useful discussions on reinforcement learning. FC and YHH are supported by a Leverhulme Trust Research Project (Grant No. RPG-2022-145). FG was supported by the Crankstart Scholarship during the early stages of the project.

#### Appendix A. Reinforcement learning

#### A.1. Proximal policy optimisation

The dynamics of the intelligent agent are determined by a Markov decision process (MDP), formally a tuple  $(S, \mathcal{A}, \{P(s, a)\}, \gamma, R)$  where

- *S* is the set of states in the environment.
- $\mathcal{A}$  is the set of actions permissible to the agent.
- P(s, a) are the state transition probabilities. For each state  $s \in S$  and action  $a \in A$ , this describes the distribution over what states the agent will transition to if the action *a* is taken in state *s*.
- $\gamma \in [0, 1)$  is the discount factor.
- $R : S \times A \rightarrow \mathbb{R}$  is the reward function.

The MDP evolves as follows: starting with an initial state  $s_0 \in S$ , the agent selects some action  $a_0 \in A$ , drawn according to  $P(s_0, a_0)$ . This causes a transition from  $s_0$  to some new state  $s_1 \in S$ . The agent then picks another action  $a_1$  drawn according to  $P(s_1, a_1)$ , which leads to  $s_2 \in S$ . The process continues iteratively, with the agent choosing subsequent actions  $a_0, a_1, a_2, \ldots$ , to navigate a trajectory of states:

$$s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} \dots$$
 (A.1)

The goal of any reinforcement learning algorithm is to find a policy  $\pi: S \to A$  that maximises the so-called *Q*-function

$$Q^{\pi}(s_0) = E \Big[ R(s_0, a_0) + \gamma R(s_1, a_1) + \gamma^2 R(s_2, a_2) + \dots \Big| \pi \Big],$$
(A.2)

which is the expected cumulative reward when the agent takes actions according to some policy  $\pi$ , i.e.,  $a_t = \pi(s_t)$ .

Proximal Policy Optimisation (PPO) improves the policy  $\pi$  by parameterising it with a neural network and introducing a clipped objective function to constrain policy updates. This stabilises training by ensuring gradual improvements. This robustness is particularly advantageous when combined with action masking to exclude domain-violating actions dynamically. Domain-violating actions are identified by the constraints of the environment, and we customise the model-free PPO algorithm to mask these actions dynamically.

#### A.2. $\epsilon$ -greedy strategy

Although PPO generates a distribution over actions for each state, we further incorporate an  $\epsilon$ -greedy exploration scheme. In the language of the MDP tuple  $(S, \mathcal{A}, \{P(s, a)\}, \gamma, R)$ , this approach modifies the effective policy distribution  $\pi(s)$  from which actions are sampled. Concretely, with probability  $\epsilon$ , the agent chooses a random valid action (one that is not masked out by domain constraints), and with probability  $1 - \epsilon$ ,



Fig. 5. The 5 exceptional complete  $\mathcal{N} = 2$  theories and their BPS quivers, and their minimal chamber as identified by the RL model. In agreement with the literature, we find that the  $X_6$  Derksen-Owen quiver has no finite chamber.

(

it selects an action according to the PPO policy. This ensures sufficient exploration of the state-action space during training, mitigating premature convergence to suboptimal deterministic policies and helping the agent discover higher-value trajectories.

#### A.3. Dynamic action masking

Dynamic action masking is used to ensure that the agent never attempts actions outside the domain constraints defined by the environment. At each state s, a subset of actions in  $\mathcal{A}$  may be prohibited due to physical laws, or human-imposed restrictions. Before the policy network computes the action probabilities for a given state s, a mask  $\mathbf{m}(s)$ is generated, where  $\mathbf{m}(s)$  is a binary vector of the same dimension as  $\mathcal{A}$ . Each entry of  $\mathbf{m}(s)$  corresponds to an action in  $\mathcal{A}$ , taking the value 1 if the action is valid and 0 otherwise. This vector is dynamically constructed by querying the underlying environment constraints relevant to the current state. Formally, if  $\mathscr{C}$  is the vector of logits for each action, the masked logits are given by

# $\boldsymbol{\ell}' = \boldsymbol{\ell} \otimes \mathbf{m}(s),$

where  $\otimes$  denotes element-wise multiplication. Logits associated with invalid actions are thus set to zero (or a sufficiently negative value if working in log space), ensuring that subsequent sampling of actions is limited to valid entries only.

#### Appendix B. Graph automorphisms

An automorphism of a directed graph G = (V, E) is a permutation  $\sigma$  of the vertex set V that preserves edge-vertex connectivity. That is, if the vertex pair (u, v) is connected by a directed edge, the pair

$$\sigma(u), \sigma(v)) \tag{B.1}$$

is also connected by an edge with the same direction. BPS quivers are directed graphs, with the vertex set V the set of node labels  $\{\gamma_i\}$ . For each finite chamber of the quiver, there exists a set of equivalent chambers obtained through such automorphisms. The degeneracy of each chamber is therefore equal to the order of the automorphism group of the underlying quiver Aut(G). The automorphism groups of the quivers from Fig. 3 are shown in Fig. B.6 with the following considerations:

- Pairs of nodes that are automorphically equivalent are identified with the same colour. A quiver G that contains n such equivalent nodes contains a factor S<sub>n</sub> in Aut(G).
- 2. Global symmetries of the graph *G* represent automorphisms where entire classes of nodes are interchanged simultaneously. These contribute factors of  $\mathbb{Z}_m$  to the group Aut(*G*), with *m* the order of the symmetry.

The quivers  $Q_1$  and  $Q_4$  contain no such global symmetries. The quivers  $Q_2$  and  $Q_3$  have a third order rotational symmetry  $\mathbb{Z}_3$ , and a reflection symmetry  $\mathbb{Z}_2$ , respectively. Our code utilities the networkx package for Python to automatically detect these automorphisms for a given quiver.



Fig. B.6. The four BPS quivers of the SU(2)  $N_f=4~{\rm SCFT}$  and their automorphism groups.

#### Data availability

Data will be made available on request.

#### References

- N. Seiberg, E. Witten, Electric magnetic duality, monopole condensation, and confinement in N = 2 supersymmetric Yang-Mills theory, Nucl. Phys. B 426 (1994) 19–52, Erratum: Nucl. Phys. B 430 (1994) 485–486, arXiv:hep-th/9407087.
- [2] N. Seiberg, E. Witten, Monopoles, duality and chiral symmetry breaking in N = 2 supersymmetric QCD, Nucl. Phys. B 431 (1994), arXiv:hep-th/9408099.
- [3] S. Cecotti, C. Vafa, Classification of complete N=2 supersymmetric theories in 4 dimensions, arXiv:1103.5832, 3 2011.
- [4] M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, C. Vafa, BPS quivers and spectra of complete N = 2 quantum field theories, Commun. Math. Phys. 323 (2013) 1185–1227, arXiv:1109.4941.

- [5] D. Gaiotto, N = 2 dualities, J. High Energy Phys. 08 (2012) 034, arXiv:0904.2715.
- [6] H. Derksen, T. Owen, New graphs of finite mutation type, Electron. J. Comb. 15 (2008), arXiv:0804.0787.
- [7] H.A. Kissinger, E. Schmidt, D. Huttenlocher, The Age of AI: and Our Human Future, Hachette UK, 2021.
- [8] Y.-H. He, Deep-learning the landscape, Phys. Lett. B 774 (2017) 564–568, arXiv: 1706.02714.
- [9] D. Krefl, R.-K. Seong, Machine learning of Calabi-Yau volumes, Phys. Rev. D 96 (6) (2017) 066014, arXiv:1706.03346.
- [10] J. Carifio, J. Halverson, D. Krioukov, B.D. Nelson, Machine learning in the String landscape, J. High Energy Phys. 09 (2017) 157, arXiv:1707.00655.
- [11] F. Ruehle, Evolving neural networks with genetic algorithms to study the String landscape, J. High Energy Phys. 08 (2017) 038, arXiv:1706.07024.
- [12] Y.-H. He, The Calabi–Yau Landscape: From Geometry, to Physics, to Machine Learning, Lecture Notes in Mathematics, 2021, arXiv:1812.02893.
- [13] F. Ruehle, Data science applications to string theory, Phys. Rep. 839 (2020) 1-117.
- [14] S. Gukov, J. Halverson, F. Ruehle, Rigor with machine learning from field theory to the Poincaré conjecture, Nat. Rev. Phys. (2024) 1–10.
- [15] Y.-H. He, AI-driven research in pure mathematics and theoretical physics, Nat. Rev. Phys. 6 (9) (2024) 546–553, arXiv:2405.19973.
- [16] H.-Y. Chen, Y.-H. He, S. Lal, M.Z. Zaz, in: Machine Learning Etudes in Conformal Field Theories, 6 2020, arXiv:2006.16114.
- [17] G. Kantor, V. Niarchos, C. Papageorgakis, Conformal bootstrap with reinforcement learning, Phys. Rev. D 105 (2022), arXiv:2108.09330.
- [18] J. Bao, S. Franco, Y.-H. He, E. Hirst, G. Musiker, Y. Xiao, Quiver mutations, Seiberg duality, and machine learning, Phys. Rev. D 102 (8) (2020), arXiv:2006.10783.
- [19] M.-W. Cheung, P.-P. Dechant, Y.-H. He, E. Heyes, E. Hirst, J.-R. Li, Clustering cluster algebras with clusters, Adv. Theor. Math. Phys. 27 (3) (2023) 797–828, arXiv:2212. 09771.
- [20] J. He, H. Jenne, H. Chau, D. Brown, M. Raugas, S.C. Billey, H. Kvinge, Machines and mathematical mutations: using GNNs to characterize quiver mutation classes, in: The 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24, 2024.
- [21] S. Gukov, R.-K. Seong, Machine learning BPS spectra and the gap conjecture, Phys. Rev. D 110 (4) (2024) 046016, arXiv:2405.09993.
- [22] C. Córdova, S.-H. Shao, Schur indices, bps particles, and Argyres-Douglas theories, J. High Energy Phys. 2016 (1) (2016), arXiv:1506.00265.
- [23] M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi, C. Vafa, N = 2 quantum field theories and their BPS quivers, Adv. Theor. Math. Phys. 18 (1) (2014) 27–127, arXiv:1112.3984.
- [24] S. Fomin, A. Zelevinsky, Cluster algebras I: foundations, J. Am. Math. Soc. 15 (2) (2002) 497–529.
- [25] B. Feng, A. Hanany, Y.-H. He, D-brane gauge theories from toric singularities and toric duality, Nucl. Phys. B 595 (2001) 165–200, arXiv:hep-th/0003085.
- [26] C.-Y. Tang, C.-H. Liu, W.-K. Chen, S.D. You, Implementing action mask in proximal policy optimization (PPO) algorithm, ICT Express 6 (3) (2020) 200–203.
- [27] D. Xie, BPS spectrum, wall crossing and quantum dilogarithm identity, Adv. Theor. Math. Phys. 20 (3) (2016) 405–524.
- [28] T. Brustle, G. Dupont, M. Perotin, On maximal Green sequences, J. Algebra 580 (580) (2012) 399–422, arXiv:1205.2050.
- [29] B. Keller, Quiver mutation in Java, https://webusers.imj-prg.fr/~bernhard.keller/ quivermutation/, 2006.
- [30] K. Saito, Extended affine root systems I (Coxeter transformations), Publ. Res. Inst. Math. Sci. 21 (1) (1985) 75–179.