
Certificates of many-body properties assisted by machine learning

Borja Requena¹, Gorka Muñoz-Gil¹, Maciej Lewenstein^{1,2}, Vedran Dunjko³, Jordi Tura^{4,5*}

¹ICFO-Institut de Ciències Fòniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels (Barcelona), Spain

²ICREA, Pg. Lluís Companys 23, 08010 Barcelona, Spain

³LIACS, Leiden University, Niels Bohrweg 1, 2333 CA Leiden, Netherlands

⁴Max-Planck-Institut für Quantenoptik, D-85748 Garching

⁵Instituut Lorentz, Leiden University, Niels Bohrweg 2, 2333 CA, Leiden, Netherlands

1 Introduction

Computationally intractable tasks are ubiquitous in physics and optimization [1, 2, 3, 4]. One of the most prominent examples is finding the ground state energy of a many-body quantum system. Whereas variational approaches (i.e., coming up with an ansatz that contains a good approximate to the optimal solution) are the most direct tool to find such a good solution [5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15], they generally suffer from two important drawbacks: (i) non-convexity of the cost function and/or the feasible set and (ii) they provide only a bound to the optimal solution, by construction, without the direct possibility of knowing how close the variational solution is to the global optimum. On the other hand, relaxation techniques, where e.g. a combinatorial problem is approximated by a continuous parameter problem which allows for more efficient solution methods, provide lower bounds to the solutions [16, 17, 18, 19]. Hence, a combination of a good variational ansatz and a good relaxation of the problem provides a bounded interval containing the optimal solution and therefore a means to control the error in the optimization (see Fig. 1a).

In this work, we propose a novel approach combining the power of relaxation techniques with deep reinforcement learning (DRL) [20, 21, 22] to find the best possible relaxations given a limited computational budget. We specifically apply it to the problem of bounding ground state energies. Relaxation techniques are widely used in quantum information processing. In many cases of interest, given an optimization problem, we consider a set of constraints that must be fulfilled by the solution. We can solve the resulting constrained optimization problem through semidefinite programming (SdP). Relaxing the constraints simplifies the problem, but it may yield looser bounds. Some "smart" relaxations, however, provide better approximations to the optimal solution than others of similar complexity [17, 23]. Hence, it becomes paramount to find the best trade-off between accuracy and simplicity. Nevertheless, a successful search often requires leveraging upon specific knowledge about the problem at hand. Conversely, the analysis of an efficient set of constraints is more likely to reveal useful insight about the system's properties [24, 25, 26].

We propose a systematic method to search for an optimal set of constraints within a given computational budget. We depict a scheme in which we distinguish two main parts. On one hand, a DRL-based agent can input a set of constraints into a black box. On the other hand, the black box solves the associated problem relaxation through SdP (Fig. 1b). The agent can strengthen or loosen the constraints obtaining a reward that depends on the quality of the resulting bound and its computational cost. This way, the agent is trained to obtain the best possible bound within a computational budget, while using the least amount of resources. Notice that this is a completely unsupervised process in which the agent is absolutely agnostic to the physical problem at hand.

*Corresponding author: tura@lorentz.leidenuniv.nl

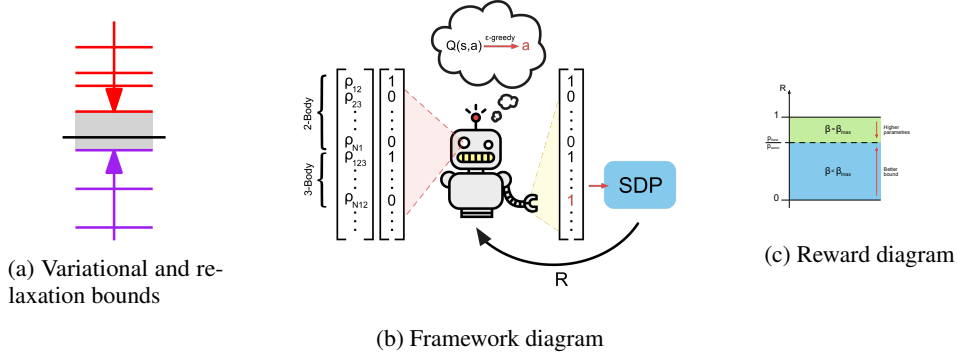


Figure 1: (a) Values of the objective function. In black, the optimal (unknown) value. In red, the different minimums obtained by different variational methods. The smaller their value, the better is the bound. In purple, different minimums obtained by relaxation techniques. The greater their value, the more accurate is their associated certificate. In gray, the uncertainty region where the optimal solution lies, given by the best variational and the best certificate obtained so far. (b) Schematic representation of the reinforcement learning framework. First, the agent observes the state: a 1-hot encoding of the active constraints. Given the observation, it estimates the Q-values associated to the possible actions with the DQN. Then, the agent decides which action to take according to an ϵ -greedy policy, which results into a new set of constraints, bringing the agent to a new state. Finally, the black box solves the SdP associated to the new state, providing the agent with a reward, whose parts are illustrated in (c).

Our aim is to understand up to which extent such a fully unsupervised approach may help us study physical systems inferring, for instance, their phase diagrams.

Our procedure is applied in the context of finding ground state energies of local Hamiltonians. We showcase the validity of the method in various scenarios where the ground states have different properties. Besides, our framework allows the straightforward implementation of other optimization algorithms, such as classical simulated annealing and breadth first search, which we benchmark against the DRL approach. Finally, we study the effect of the so-called transfer learning [27] between Hamiltonians, exploiting the experience obtained by the agent solving previous problems in solving new ones.

We highlight that our methods, while presented on the context of ground state energy approximation, are actually of much broader applicability, ranging, e.g. from entanglement witnesses optimization [28, 29, 30, 31] to device-independent quantum information processing tasks.

2 Methods

In the interest of simplicity, we shall consider a running example. Let us therefore fix an optimization task, which is to find the ground state energy E_0 of a quantum local Hamiltonian $H = \sum_{i=1}^m H_i$. The Hamiltonian H acts on n qubits, and it is a sum of terms H_i , each of which acts on at most $k = O(1)$ qubits. The sum of H has therefore $m = O(\text{poly}(n))$ terms. The support of H_i , denoted $\text{supp}(H_i)$ is the set of qubits where H_i acts non-trivially. The supports of the different H_i may overlap; *i.e.*, $\text{supp}(H_i) \cap \text{supp}(H_j)$ may not be empty.

To find E_0 , one can try to directly construct a quantum state that has E_0 energy with respect to H . Therefore, a first approach is to parameterize a family of quantum states $|\psi(\theta)\rangle$ exploiting some known properties of H . We can safely assume the parameterization yields a valid, normalized quantum state for any value of the parameters θ . Additionally, by construction, $\langle \psi(\theta) | H | \psi(\theta) \rangle \geq E_0$ for all θ . Let us denote

$$\gamma = \min_{\theta} \langle \psi(\theta) | H | \psi(\theta) \rangle, \quad (1)$$

which satisfies $E_0 \leq \gamma$ by construction. An example of such a parameterization would be to describe $|\psi(\theta)\rangle$ as a tensor network contraction, which exploits the locality properties of H , which limit the entanglement present in its ground state [6, 32, 33, 34].

Complexity theory results strongly suggest that finding, or even approximating, the ground state energy of a local Hamiltonian is a hard task, even for a quantum computer [1, 2, 3]. Furthermore, this hardness persists in physically relevant instances [4]. Observe that even if we found the actual solution $|\psi(\boldsymbol{\theta})\rangle$ we cannot prove, solely from that, that it is the global minimum [15].

It is therefore highly desirable to obtain a bound from the other side; i.e., a value β for which one can prove $E_0 \geq \beta$. This would guarantee $E_0 \in [\beta, \gamma]$ and, therefore, help determine whether it is worth to refine the search depending on $|\gamma - \beta| < \varepsilon$. However, for a proof of the type $E_0 \geq \beta$, constructing examples of the form $|\psi(\boldsymbol{\theta})\rangle$ is not good enough, and we need a proof that is satisfied by, at least, all valid quantum states. Roughly speaking, we might want to find a proof that works on a set containing all quantum states if this makes it simpler. Such a proof is referred to as a certificate, and it is typically obtained by numerical means. Semidefinite programs (SdP) are natural tools to obtain such certificates, which we use in our work.

2.1 SdP-based relaxations

A common technique to construct a relaxation for the local Hamiltonian problem is via the triangle inequality [35, 36, 37, 38]:

$$\min_{\rho} \text{Tr}[\rho H] \geq \sum_i \min_{\rho_i} \text{Tr}[\rho_i \hat{H}_i], \quad (2)$$

where ρ and ρ_i are density matrices acting on the support of H and H_i respectively. Note that i refers to a Hamiltonian term and it has nothing to do with the i -th party. Furthermore, in eq. (2), the \hat{H}_i are sums of local terms H_i of H , grouped so that $\text{supp}(\hat{H}_i)$ is as large as possible while still allowing for computation of their minimal eigenvalue. This obviously depends on the available computational resources.

Let us observe that the RHS in eq. (2) is a sum of minima, where each minimum is carried out independently. Because of this independence it is false in general that different ρ_i are mutually compatible; i.e., that there exists a global state ρ such that each ρ_i is the corresponding partial trace of ρ . The converse is true however: every valid quantum state ρ has an associated set of partial traces ρ_i , but given a set of ρ_i , a global ρ may not exist. This is what proves the inequality eq. (2).


The minimization of the RHS of eq. (2) is equivalent to solving the following semidefinite program:

$$\begin{aligned} \beta_{\emptyset} := \min_{\{\rho_i\}} & \sum_i \text{Tr}[\rho_i \hat{H}_i] \\ \text{s.t.} & \rho_i \succeq 0 \\ & \text{Tr}[\rho_i] = 1 \end{aligned} \quad (3)$$


Since there is no mutual compatibility enforced among the ρ_i , and each is treated independently, the triangle inequality eq. (2) constitutes what we call a trivial relaxation. A natural way to strengthen the relaxation is to impose further restrictions on the collection of possible ρ_i , in such a way that any quantum state would also satisfy them. The most stringent restriction possible is to directly ask that $\{\rho_i\}$ come from a global quantum state. This would be equivalent to finding the value of E_0 , which is QMA-complete [1, 2, 3]. Furthermore, it is strongly connected to solving the so-called quantum marginal problem (QMP), which is also QMA-complete [39]. The QMP has been solved completely in very rare instances, such as the global state being symmetric [19] or for the case of one-body marginals [40]. Nevertheless, the SdP based formulation eq. (3) motivates a hierarchy of relaxations based on solving the QMP up to some degree of compatibility. The degree of compatibility that uses the available computational resources most efficiently is to be determined by the Agent.revte

2.2 Reinforcement learning

The set of constraints the agent can add to eq. (3) correspond to defining reduced density matrices of subsets of qubits and imposing mutual compatibility constraints among them (i.e., relevant partial traces should coincide). Hence, the set of constraints forms a partially ordered set (poset). We frame the optimization problem of finding the best set of constraints within a limited computational budget as a Markov decision process (MDP). In such a scheme we distinguish two parts: on the one hand, a black box that computes the associated SdP relaxation to a set of constraints. On the other hand, an agent that decides which constraints to try, strengthening or loosening them depending on perceived rewards.



(a) Benchmarking allowing half of the 3-body constraints



(b) Benchmarking allowing all the 3-body constraints

Figure 2: Benchmark of the performance of the three optimization algorithms: BFS, MC and DQN. The algorithms are evaluated in: (a) small and (b) big computational budget. The lines are discontinued when the number of state visits needed to hit 95% exceeds the time limit allocated.

The state space of the agent consists of states that encode the active constraints. These are vectors with one-hot-encodings of the constraints into reduced density matrices. The actions take the agent from one state to another. These actions flip bits in the state vector thus adding or removing constraints under certain restrictions (e.g. memory/time consumption to obtain the associated certificate). The reward function defines the optimization goal for the agent. It depends on the energy bound β obtained by the black box and its computational cost p . The latter is determined by the number of variables and constraints in the SdP [41, 42]. Without prior knowledge, the agent relies on a set of references that are updated with the state-space exploration, keeping track of the best β_{\max} and worst β_{\min} energy bounds obtained so far. The reward function is

$$R(\beta, p) = \frac{p_{\text{best}}}{p_{\text{worst}}} \cdot \begin{cases} \frac{p_{\text{worst}}}{p} & \text{if } \beta = \beta_{\max} \\ \left(\frac{\beta - \beta_{\min}}{\beta_{\max} - \beta_{\min}} \right)^d & \text{otherwise,} \end{cases} \quad (4)$$

where d is a fixed exponent that controls the shape of the line $(\beta - \beta_{\min})/(\beta_{\max} - \beta_{\min})$. Notice that $p_{\text{worst}} \geq p_{\text{best}}$ and, therefore, $p_{\text{worst}}/p \geq 1$. The optimization problem is solved through reinforcement learning [20, 21, 22]. More precisely, we use Q-learning [43] in the context of function approximation leveraging a deep Q-network (DQN) [44] with double Q-learning [45]. The agent follows an ϵ -greedy policy.

3 Results

We showcase the results with the following problem: finding the ground state energy of the 1D quantum Heisenberg XY model [46] with periodic boundary conditions, which is described by

$$H = \sum_{i=0}^{n-1} J_i (\sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y) + \sum_{i=0}^{n-1} B_i \sigma_i^z. \quad (5)$$

We limit the agent’s memory so it can only allocate half of the possible three-body (neighboring) constraints. We further fix the interactions to be $J_i = i \bmod 3$ and $B_i = 1$ in order to know the optimal set of constraints beforehand (thus to compare different methods with the optimal solution). We benchmark the performance of the DQN-based agent against a breadth first search (BFS) in the space of constraints as well as a Monte-Carlo (MC) optimization for different system sizes (Fig. 2).

- We observe in Fig. 2(a) that for small budget the DQN agent is more efficient than MC and it is more efficient than BFS in all instances. For a bigger budget, allowing for all the 3-body constraints, the difference is even more notable (Fig. 2(b)).
- Finally, we also study the effect of transfer learning within the same Hamiltonian phase, where we observed approximately a 5 times speedup within the same phase, which goes into no gain gradually as one crosses a phase transition and enters a new Hamiltonian phase in a continuous way: Since energy changes continuously, the same certificate in the learnt phase still obtains good rewards in the vicinity of the quantum phase transition.

These results suggest that more advanced AI and RL methods may provide substantial improvements for practical applications of relaxation methods. Particularly, for the problems of finding ground energies of local quantum Hamiltonians which is a quintessential problem in modern Physics.

Broader Impact

- Who can benefit from this research: This work constitutes basic research, thus it can benefit the scientific community as a whole. In particular, the fields of condensed matter, quantum information and optimization.
- Who can be put at disadvantage from this research: Does not apply.
- What are the consequences of failure of the system. Does not apply. The relaxation obtained by the agent is always a valid certificate, by construction, no matter how poorly the agent performs in finding a resource-efficient one.
- Does the task/method leverage biases in the data? Does not apply. The agent is completely agnostic to the model it is learning.

Acknowledgments

ICFO group acknowledges support from ERC AdG NOQIA, Spanish Ministry of Economy and Competitiveness (“Severo Ochoa” program for Centres of Excellence in RD (CEX2019-000910-S), Plan National FIDEUA PID2019-106901GB-I00/10.13039 / 501100011033, FPI), Fundació Privada Cellex, Fundació Mir-Puig, and from Generalitat de Catalunya (AGAUR Grant No. 2017 SGR 1341, CERCA program, QuantumCAT U16-011424 , co-funded by ERDF Operational Program of Catalonia 2014-2020), MINECO-EU QUANTERA MAQS (funded by State Research Agency (AEI) PCI2019-111828-2 / 10.13039/501100011033), EU Horizon 2020 FET-OPEN OPTOLogic (Grant No 899794), and the National Science Centre, Poland-Symfonia Grant No. 2016/20/W/ST4/00314. This work was in part supported by the Dutch Research Council (NWO/OCW), as part of the Quantum Software Consortium programme (project number 024.003.037). This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number 414325145 in the framework of the Austrian Science Fund (FWF): SFB F71. J. T. thanks the Alexander von Humboldt foundation for support.

References

- [1] Julia Kempe and Oded Regev. 3-local hamiltonian is qma-complete. *Quantum Info. Comput.*, 3(3):258–264, May 2003.
- [2] Julia Kempe, Alexei Kitaev, and Oded Regev. The complexity of the local hamiltonian problem. *SIAM Journal on Computing*, 35(5):1070–1097, jan 2006.
- [3] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Communications in Mathematical Physics*, 287(1):41–65, jan 2009.
- [4] Norbert Schuch and Frank Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature Physics*, 5(10):732–735, aug 2009.
- [5] F. Verstraete, D. Porras, and J. I. Cirac. Density matrix renormalization group and periodic boundary conditions: A quantum information perspective. *Physical Review Letters*, 93(22), nov 2004.
- [6] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349:117–158, oct 2014.
- [7] Thomas E. O’Brien, Bruno Senjean, Ramiro Sagastizabal, Xavier Bonet-Monroig, Alicja Dutkiewicz, Francesco Buda, Leonardo DiCarlo, and Lucas Visscher. Calculating energy derivatives for quantum chemistry on a quantum computer. *npj Quantum Information*, 5(1), dec 2019.
- [8] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, feb 2017.
- [9] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4), dec 2019.

- [10] Ivan Glasser, Nicola Pancotti, Moritz August, Ivan D. Rodriguez, and J. Ignacio Cirac. Neural-network quantum states, string-bond states, and chiral topological states. *Physical Review X*, 8(1), jan 2018.
- [11] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm.
- [12] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller. Self-verifying variational quantum simulation of lattice models. *Nature*, 569(7756):355–360, may 2019.
- [13] A. Garcia-Saez and J. I. Latorre. Addressing hard classical problems with adiabatically assisted variational quantum eigensolvers.
- [14] R. Sagastizabal, X. Bonet-Monroig, M. Singh, M. A. Rol, C. C. Bultink, X. Fu, C. H. Price, V. P. Ostroukh, N. Muthusubramanian, A. Bruno, M. Beekman, N. Haider, T. E. O'Brien, and L. DiCarlo. Experimental error mitigation via symmetry verification in a variational quantum eigensolver. *Physical Review A*, 100(1), jul 2019.
- [15] Jakub Czartowski, Konrad Szymański, Bartłomiej Gardas, Yan Fyodorov, and Karol Życzkowski. Ground state energy: Can it really be reached with quantum annealers?
- [16] Miguel Navascués, Stefano Pironio, and Antonio Acín. Bounding the set of quantum correlations. *Phys. Rev. Lett.*, 98:010401, Jan 2007.
- [17] *Semidefinite Optimization and Convex Algebraic Geometry*. SOC FOR INDUSTRIAL & APPLIED M, 2013.
- [18] Matteo Fadel and Jordi Tura. Bounding the set of classical correlations of a many-body system. *Phys. Rev. Lett.*, 119:230402, Dec 2017.
- [19] Albert Aloy, Matteo Fadel, and Jordi Tura. The quantum marginal problem for symmetric states: applications to variational optimization, nonlocality and self-testing.
- [20] Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [21] Yuxi Li. Deep reinforcement learning: An overview.
- [22] Vedran Dunjko and Hans J Briegel. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics*, 81(7):074001, jun 2018.
- [23] S. Pironio, M. Navascués, and A. Acín. Convergent relaxations of polynomial optimization problems with noncommuting variables. *SIAM Journal on Optimization*, 20(5):2157–2180, 2010.
- [24] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Complete family of separability criteria. *Physical Review A*, 69(2), feb 2004.
- [25] Alexia Salavrakos, Remigiusz Augusiak, Jordi Tura, Peter Wittek, Antonio Acín, and Stefano Pironio. Bell inequalities tailored to maximally entangled states. *Phys. Rev. Lett.*, 119:040402, Jul 2017.
- [26] Jędrzej Kaniewski, Ivan Šupić, Jordi Tura, Flavio Baccari, Alexia Salavrakos, and Remigiusz Augusiak. Maximal nonlocality from maximal entanglement and mutually unbiased bases, and self-testing of two-qutrit quantum systems. *Quantum*, 3:198, oct 2019.
- [27] Matthew E. Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *J. Mach. Learn. Res.*, 10:1633–1685, December 2009.
- [28] Barbara M. Terhal. Bell inequalities and the separability criterion. *Physics Letters A*, 271(5):319 – 326, 2000.
- [29] F. Troiani and I. Siloi. Energy as a witness of multipartite entanglement in chains of arbitrary spins. *Physical Review A*, 86(3), sep 2012.
- [30] A. Aloy, J. Tura, F. Baccari, A. Acín, M. Lewenstein, and R. Augusiak. Device-independent witnesses of entanglement depth from two-body correlators. *Physical Review Letters*, 123(10), sep 2019.
- [31] Jean-Daniel Bancal, Nicolas Gisin, Yeong-Cherng Liang, and Stefano Pironio. Device-independent witnesses of genuine multipartite entanglement. *Physical Review Letters*, 106(25), jun 2011.
- [32] F. Verstraete, V. Murg, and J.I. Cirac. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Advances in Physics*, 57(2):143–224, mar 2008.

- [33] Norbert Schuch, Ignacio Cirac, and David Pérez-García. PEPS as ground states: Degeneracy and topology. *Annals of Physics*, 325(10):2153–2192, oct 2010.
- [34] Norbert Schuch and J. Ignacio Cirac. Matrix product state and mean-field solutions for one-dimensional systems can be found efficiently. *Physical Review A*, 82(1), jul 2010.
- [35] P. W. Anderson. Limits on the energy of the antiferromagnetic ground state. *Physical Review*, 83(6):1260–1260, sep 1951.
- [36] Rolf Tarrach and Roser Valent. Exact lower bounds to the ground-state energy of spin systems: The two-dimensional $S=1/2$ antiferromagnetic heisenberg model. *Physical Review B*, 41(13):9611–9613, may 1990.
- [37] Anushya Chandran, Dagomir Kaszlikowski, Aditi Sen(De), Ujjwal Sen, and Vlatko Vedral. Regional versus global entanglement in resonating-valence-bond states. *Physical Review Letters*, 99(17), oct 2007.
- [38] Fabien Alet, Daniel Braun, and Grégoire Misguich. Comment on regional versus global entanglement in resonating-valence-bond states. *Physical Review Letters*, 101(24), dec 2008.
- [39] Yi-Kai Liu. Consistency of local density matrices is QMA-complete. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 438–449. Springer Berlin Heidelberg, 2006.
- [40] M. Walter, B. Doran, D. Gross, and M. Christandl. Entanglement polytopes: Multiparticle entanglement from single-particle information. *Science*, 340(6137):1205–1208, jun 2013.
- [41] S. Arora, E. Hazan, and S. Kale. Fast algorithms for approximate semidefinite programming using the multiplicative weights update method. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. IEEE, 2005.
- [42] Fernando G.S.L. Brandao and Krysta M. Svore. Quantum speed-ups for solving semidefinite programs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, oct 2017.
- [43] Christopher John Cornish Hellaby Watkins. Learning from delayed rewards. 1989.
- [44] 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):529–533, 2015.
- [45] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. *arXiv preprint arXiv:1509.06461*, 2015.
- [46] Elliott Lieb, Theodore Schultz, and Daniel Mattis. Two soluble models of an antiferromagnetic chain. *Annals of Physics*, 16(3):407–466, 1961.