ASIA-PACIFIC SOLAR RESEARCH CONFERENCE

PEER REVIEWED ASIA-PACIFIC SOLAR RESEARCH CONFERENCE

Luke Marshall

Development of an Open AI Energy Market Simulator for Deep Learning

Luke Marshall, Anna Bruce[1], and Iain MacGill[3]

[1]*School of Photovoltaic and Renewable Energy Engineering, University of New South Wales, Sydney NSW 2052, Australia*

[2]*Centre for Energy and Environmental Markets and School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney NSW 2052, Australia*

E-mail: luke.marshall@unsw.edu.au

Abstract

Restructured electricity markets such as Australia's National Electricity Market (NEM) feature complex auction mechanisms that are intended to incentivise participants to reveal their true costs so that efficient investment and operational decisions can be made. A major problem in the design and regulation of electricity market mechanisms is that participant incentives are difficult to model; exogenous factors may significantly alter the offers of electricity generators away from pure short-run marginal generation cost. The entrance of new near-zero marginal cost participants such as wind and photovoltaics, as well as low-marginal cost energy storage introduce pricing and inter-temporal complexities that mechanism designers may struggle to fully model. This means that while policymakers may strive to create rules that lead to fair and efficient market mechanisms, it is difficult to anticipate the impact of the transition to clean energy.

Electricity market simulations are one tool that can be used to understand the way participants may behave under electricity market mechanisms. These simulations have historically relied on algorithm designers to impart knowledge of the operation of these markets to create agents that bid in a rational manner; these approaches may be limited by the algorithm designer's understanding of the market mechanism.

This paper details the structure, design and development of a novel electricity market simulation that leverages emerging artificial intelligence (AI) technologies such as deep Q-learning. These technologies have recently surpassed a number of game-playing benchmarks considered critical to developing 'general purpose' strategic game-playing AI. There appears to be a wealth of knowledge and insight that could come from an application of these emerging AI technologies to electricity market simulations. Without a standard set of tools however, research is limited and difficult to replicate. Additionally, there is limited cross-over in domain knowledge between software engineers with the ability to develop such environments and researchers who wish to conduct these experiments, highlighting the need for pre-made, customisable models that require a less detailed technical knowledge of software engineering. Commercial market players can also be expected to explore AI options for maximising market outcomes. There thus appears to be a need for the development of a standard, open-source electricity market environment for AI experimentation among a wide range of stakeholders.

In this paper, an electricity market simulation developed for the Open AI Gym framework is described, which allows researchers to introduce a range of general game-playing AI systems (including deep Q-learning agents) to electricity market simulations. Using this tool, researchers can investigate the impact of changing technology types, rules and pricing outcomes on electricity market participant behaviour, without making assumptions about or approximations of participant bidding strategies.

1. Introduction

Due to the technical complexity of managing secure and reliable power system operation, restructured electricity markets are complex mechanisms and are expected to incentivise strategic behaviour. Good mechanism design should encourage participants to reveal their hidden costs, so that efficient economic decisions (ie. dispatch, investment) can be made (Roughgarden, 2016). A major problem in the design of such mechanisms is that participant incentives are difficult to model; external forces such as subsidies, hedging contracts, portfolio structuring, revenue sufficiency, carbon abatement policy and other tax incentives may vastly alter the hidden cost of energy generation away from pure short-run marginal cost (Sessa et al., 2016). Additionally, the entrance of new near-zero marginal cost participants (ie. photovoltaic and wind generation) and low-marginal cost energy storage, introduce new pricing and inter-temporal complexities that mechanism designers may struggle to fully model, such as large discrepancies between short and long-run marginal cost, consistently negative hidden costs due to the availability of renewable energy certificates and valuation of storage based purely on arbitrage potential. This means that while policymakers may strive to create rules that lead to efficient market mechanisms, their understanding how participants may react to a given set of rules is limited. A tool that allows researchers and policymakers to better model the impact of rule changes in an electricity market mechanism may illuminate unseen complexities and provide better evidence for the effectiveness of proposed modifications.

Traditionally, mechanism designers have relied on two distinct approaches to modelling the behaviour of participants in complex games. The first has been to conduct human-participant experiments, whereby real-life players are encouraged to participate in a simulated market mechanism (Y. Chen & Ledyard, 2010). This methodology has the advantage of allowing researchers access to human-level strategic thinking on an agent-scale, meaning that complex strategies can emerge and be monitored among participants. The complexity of market mechanisms and participants' limited domain knowledge may however mean that the complex adverse behaviour that mechanism designers may be aiming to observe could fail to emerge due to the time and incentive limitations of such a study. This means that while human-participant experiments may serve as a useful starting-point for understanding the function of a market mechanism, they may struggle to illuminate the very edge cases they seek.

The second option has been to develop computational economic models that simulate participants within a market mechanism (Rahimiyan & Mashhadi, 2010). The advantages of such an approach are that many more attempts can be taken at simulating scenarios under a given mechanism, and known edge-cases can be explored in detail due to the speed and repeatability of a computer simulation. There are many algorithmic tools that can be used to create complex gaming behaviour in agent-based simulations. Notably in the field known as algorithmic mechanism design, the popular machine learning Q-learning strategy has been used to model agent behaviour (Rahimiyan & Mashhadi, 2010). Such attempts at creating autonomous bidding mechanisms are however limited by the complexity and multi-dimensionality of the environment and action spaces; real-world electricity markets feature actions by tens or hundreds of participants, across multiple price-volume bands, with multiple environmental factors (ie. demand, plant operation) with inter-temporal links; participants in these markets must take multiple actions (ie. submitting many price-volume pairs), taking into consideration future prices and environments. Machine-learning approaches, based on maximising value from previous observations (such as Q-learning) may struggle with the multitude of possible inputs and outputs, due to an inability to infer between similar scenarios (Atienza, 2018) and purely algorithmic approaches (for example the minimax algorithm) (Russell & Norvig, 2016) rely on the software engineer's ability to pre-construct a bidding strategy or meta-strategy, meaning that models cannot fulfil the intent of uncovering otherwise-hidden participant strategies.

2. The Promise of Emerging Artificial Intelligence Approaches

Recent advances in the field of artificial intelligence (AI) and deep learning have however highlighted the potential of computational systems to invent creative gaming strategies that are provably superior to those of human participants. Deep Q-learning, which uses neural networks instead of a simple table to estimate the value of any given action, has recently been applied by Google (Gibney, 2016) to decisively beat the previously dominant alpha-beta-algorithm-based chess engine; the existing system effectively represented the culmination of all of human chess knowledge, and was beaten by a deep Q-learning system that was trained simply by playing chess against itself for 48 hours. Additionally, a similar general purpose deep Q-learning algorithm has been used to beat the number-one ranked human 'go' player, an AI benchmark that was predicted by mainstream AI researchers to be surpassed only in the 2040s (J. X. Chen, 2016).

The promise of 'deep learning' systems is that they may represent the emergence of a general-purpose game-playing AI that can address many classes of game (Silver et al., 2018; Wang et al., 2016). Restructured electricity markets may represent one such class of game.

The Open AI initiative (founded by a number of high profile Silicon Valley figures including Elon Musk and Sam Altman) (Metz, 2016) is an organisation that provides a set of standard games or 'environments' that can be used to train and compare AI models. These 'games' range from physics simulations and strategically challenging games like chess, to Atari video games (Chrabaszcz et al., 2018). Open AI provides an intermediary that allows AI programs to observe, control and interact with these games so that their performance can be repeated, tested and objectively compared (Brockman et al., 2016).

3. An Open AI Electricity Market Simulation Environment

Valuable insights could come from the application of emerging AI technologies (such as deep Q-learning) to electricity market simulations. Without a standard set of tools however, such research is currently limited and, importantly, non-repeatable. One hindrance in the historical development of such tools may have been the limited cross-over in domain knowledge between developers of AI environments and energy economists wishing to use them to gain insight into participant behaviour. Pre-made (but customisable) models that require a less detailed technical knowledge of software engineering may facilitate energy market modellers with the required knowledge to conduct detailed market design experiments. There thus appears to be a need for a standard electricity market environment for AI experimentation.

The remainder of this paper will describe the justification for and architecture of a new electricity market simulator, based on the Open AI platform. This system will allow current and future programmers to effectively 'plug and play' emerging artificial intelligence systems into a standardised electricity market simulation environment, providing energy market modellers with the ability to consistently run models against the latest tools provided by AI research.

The design detailed in this paper has been implemented in the Python programming language and is available for use and modification as an open-source package.

4. The Case for Open-Source Market Deep Learning

Algorithmic game theory, paired with recent advancements in deep learning, appear to have significant potential to disrupt electricity markets. Market power has historically been exercised by participants, at times maliciously, to extract windfall profits from complex designer electricity markets, but opportunities for the exercise of market power have been difficult to detect in advance. The most serious case may be the California electricity crisis of the early 2000's, (Weare, 2003) in which energy trading company Enron used a number of physical and financial strategies to effectively dictate the market price over a number of years, eventually bankrupting a large retail business and imperilling the state budget (Sweeney, 2002). Enron executives explained that the spoils of the market would rightfully go to them because they were 'the smartest guys in the room' (McLean &

Elkind, 2013). In coming years, market power may instead be exercised by 'the smartest AI's in the room.'

Artificial intelligence technology is available publicly (Silver et al., 2018) and it seems likely that soon major electricity industry participants may implement it to improve their strategies, due to its relatively low cost relative to potential competitive advantage. In order to minimise rent seeking behaviour and inefficient outcomes, it appears to be pertinent that regulators and policymakers have equal or better tools to understand market power than the companies they are regulating, yet this does not appear to have been recognised in the current process of market reform for Australia's electricity industry.

Whether the application of AI to electricity market modelling follows a closed or open development trajectory presents a critical dichotomy. 'Open' AI means that source code and knowledge about the implementations of artificial intelligence research are publicly known and available. 'Closed' AI means the opposite, that private firms and institutions keep discoveries about AI secret, in order to exploit the technology for commercial advantage. Many commercial entities employ a hybrid model, keeping specific commercially sensitive discoveries as trade secrets, while sharing generalized discoveries with online developer communities. Aspects of artificial intelligence that are kept secret are done so because they enable private firms to exert a market advantage or market power within their domain. In the context of an electricity auction that should in theory provide the most socially beneficial set of operational and investment decisions ("National Electricity (NSW) Law," 1997), closed AI would appear to provide the opportunity for individual participants to exert market power by means of advanced strategic knowledge of market design deficiencies; which represents an information asymmetry market failure. By contrast, an open AI approach would mean that all participants would have access to the same strategic knowledge, and be more likely to respond to market conditions in a way that would prevent their competitors from exerting market power to their detriment.

In general, a mechanism could not function properly if participants are aware of flaws or loopholes that the designers have not discovered or could not foresee. It is therefore imperative that mechanism designers have access to the best set of tools available for predicting participant behaviour and detecting opportunities for the exercise of market power. With this access, regulators can be proactive rather than reactive in redesigning the rules of the market.

Why not keep these advanced techniques secret among regulators? There appear to be two significant arguments against such a move. The first is that by restricting access to the latest tools, opportunities may emerge, as previously discussed, for one firm to gain an unfair competitive advantage via the procurement of artificial intelligence. The second is that for a mechanism to be publicly perceived as fair, the reasoning behind its design must appear to be logical; if policymakers cannot provide evidence for market rulemaking, there may be significant and indefensible opposition to any changes.

5. Renewable Energy and the Growing Need for Market Simulation Tools

It is important to ensure that non-dispatchable renewable energy can be effectively integrated into electricity industries in order to facilitate decarbonisation of the energy sector. Variable renewable generators (wind and photovoltaics) are currently the least cost new-build generators in many parts of the world (IRENA) and are projected to occupy increasing market share (BNEF, AEMO ISP). However, these non-dispatchable variable renewables have a number of attributes that were not considered in the original designs of restructured electricity markets. Specifically, renewables, with a near-zero short run marginal cost of generation (Elliston et al., 2016) are expected to perform in a market designed to take bids based on incremental marginal cost. Additionally, non-dispatchable renewables are unable to delay generation, and as such appear to be required to act as price-takers. This means that in high-penetration renewable scenarios, variable renewable generators without energy storage capacity may not be able to participate in the market in a similar manner to traditional fossil fuelled or hydroelectric generation. This may present a significant barrier to the efficient functioning of the market in high-penetration renewable scenarios, where the majority of generators are acting as price-takers. A computational economics approach will allow for the modelling of such scenarios, to better understand the concentrations of renewable resources under which these phenomena may occur.

If the integration of renewables cannot be successfully managed, deployment could potentially be halted or stalled. It is important that regulators and policymakers have the tools to anticipate and develop policy solutions to upcoming problems with renewable integration.

6. Open AI Principles

Games simulated within the Open AI framework follow a step-based strategy (Brockman et al., 2016). Agents (decision-making entities based on artificial intelligence algorithms) are provided with an 'observation' of the game state or environment. In the case of an electricity market, this may consist of existing bids, electricity demand and the state of an agent's power plant resources. Based on the observation, the agent takes an 'action', which is passed to the environment and modifies the game state, as in Figure 1 below (Atienza, 2018). The agent is then passed a new observation showing how the game state has changed (e.g. resultant prices) as a result of the action taken, as well as a reward to indicate whether or not the action helped the agent reached its goal (for example, in the electricity market context, the revenue from the sale of electricity in the market). This process is repeated many thousands of times. In this way, agents are given the opportunity to 'learn' the game independent of pre-programmed rules to define participation strategies.

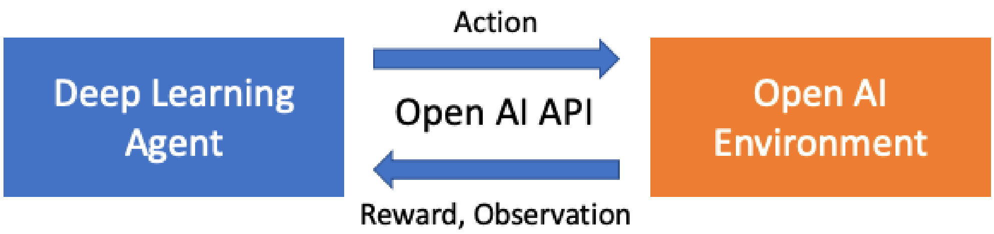

Figure 1: Open AI Action-Observation Loop

7. Architecture

The development of an Open AI Electricity Market Simulator (OAEMS) presents a number of significant software engineering challenges.

Adversarial learning, whereby multiple agents can learn a game by playing against each other, is a desired property of the environment. This presents a challenge because the observation-action loop of the existing Open AI framework requires that a new observation and reward be passed to the agent after each bid is submitted. This means that each agent needs to receive a reward for its last action before making a new decision, and in an electricity market, a bid must necessarily be received from each participant before the reward (ie. revenue) can be determined. Thus, a single agent model is not able to represent multiple participants in an electricity market, and a multi-agent environment is required.

Multi-agent environments do not yet appear to be common within the Open AI framework (Brockman et al., 2016). The necessity that multiple agents make simultaneous decisions presents a host of software design problems, specifically that many deep learning tools such as Keras and PyTorch do not fully support multiple models running simultaneously in different computational threads, and that these models by themselves require vast computational resources that restrict the number of agents that can be run on a single computer. One solution to this problem is a network communications-based approach, whereby a server runs a shared simulation and communicates with a number of agents via a network protocol such as websockets or Zero-MQ (Hintjens, 2013). This way, agents need not be running or training on the same machine, agents can utilise vastly different learning technologies, and may even include human players interacting via a user-interface.

By abstracting the model and effectively simulating real-life communication between participants in existing electricity markets, a more flexible and useful agent-based electricity market simulation has been developed and is described in this paper. The simulator may also be able to transcend the framework of AI-based research and be utilised for other future research purposes.

The architecture of the OAEMS is shown below in Figure 2. A deep learning agent submits an action to an 'environment' via the common Open AI API. The environment translates the action vector to a standard format electricity bid, which is then sent over the network (i.e. LAN / internet) to the market simulation server. Once all bids have been collected (potentially from multiple AI agents), dispatch is decided by this server and the result (in terms of winning bids, spot price, volumes dispatched, next demand etc.) is passed back across the network to the Open AI environment. The Open AI environment translates these results to an observation vector and 'reward' score, which are subsequently passed via the Open AI API to the deep learning agents. This process is repeated many thousands of time in order to train the AI by exposure to various scenarios and allow for reinforcement-based learning.

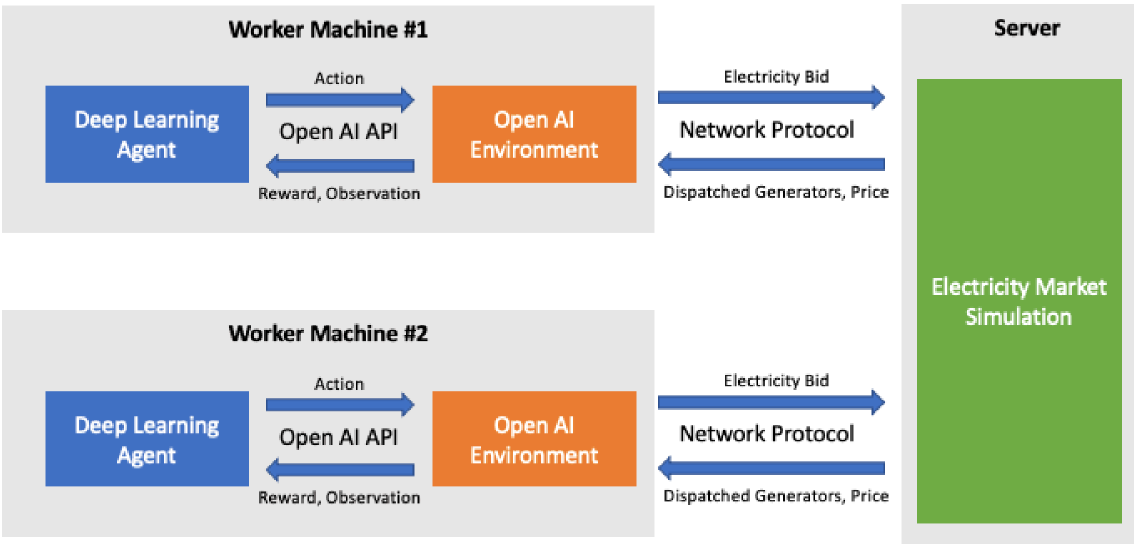

Figure 2: High-level OAEMS Architecture

8. Market Simulations

The modular nature of the OAEMS means that a number of different electricity market scenarios can be swapped out and rewritten as needed. The current release features two simulation options: a sandbox simulation for experimenting with pure market design concepts and equilibria, and a historical bid simulation that incorporates real-world generator bids from the Australian National Electricity Market, and can be used to test potential responses to this competitive environment by additional generators whose historical bid data is not included.

8.1. Sandbox Simulation

The field of mechanism design utilises a range of theoretical constructs that help illuminate the properties of a given market. These constructs often involve assumptions that may not hold in real-world situations, but may be useful in first approximations to determine the impact of rule changes, for example the assumption that bidder preferences are identically and independently distributed, or that true ‘blind’ bidding can occur (Roughgarden, 2016). Mechanism designers simplify market models using these assumptions so that inferences about utility-maximising participant behaviour can be made, and the impacts of rule tweaks can be understood in simple terms. In order to produce a tool that can accompany and test the theoretical framework of the mechanism design field, a ‘sandbox’ market simulation has been developed. This simulation is not intended to model a real-world electricity market mechanism (for example Australia’s NEM), but allows experimenters to easily tweak rules and test basic theoretical assumptions without the complexities of region-specific constraints and regulations.

8.2. AEMO Historical Bid Simulation

In the sandbox simulation, all participants within the market mechanism are expected to be autonomous bidders responding to a range of environmental variables. While theoretical value-maximising bids can be illuminating, generators in real markets are likely to be exposed to a range of incentives that may or may not be captured within a simplified simulation. This means that strategies exhibited by these agents may not map appropriately to existing operational market mechanisms. It may thus be illuminating to have some agents in an electricity market simulation that are bidding not based on a specific deep learning-based value maximiser, but are rather based on an approximation of historical bids in the NEM. In these scenarios, single or multiple deep-learning based agents can then respond to real-world bidding strategies as well as learning through competition with one another.

To this end, a historical bid-based market simulation has been developed, using data from the Australian Electricity Market Operator (AEMO) via the open-source NEMOSIS data tool (Gorman et al., 2018). Real-world half-hourly day-ahead price-volume pairs were extracted from disparate publicly available market datasets by use of the NEM Jenga bid stack explorer tool and python library, which provides a simple interface for parsing bidding data from the NEM (Marshall, 2019). Components of the Jenga library were used to give the simulation access to these bids. In each simulated settlement period, alongside autonomous bidding by AI-based agents, bids for each existing generator in the Australian NEM are also submitted. The amount of knowledge that each AI-agent has around these bids can be easily configured, so that blind or perfect-foresight bidding can be simulated.

It is important to note that while real-world bids are used in this program, it is not intended to be an accurate representation of Australia’s NEM. There are several aspects of electricity flow in the NEM which are difficult to simulate in a manner that is fast enough for deep learning. Deep learning requires that agents are exposed to many tens of thousands of different scenarios from which to build an understanding of market dynamics. For each dispatch period, the Australian Electricity Market Operator utilises NEMDE, a linear solver that incorporates thousands of physical constraints to solve dispatch across Australia’s eastern electricity network. Unfortunately, at this time, linear solvers such as NEMDE are too slow to efficiently train deep learning models, as they take many minutes to provide a dispatch answer to any given set of bids. The deep learning simulator thus provides an approximation so that agents can be trained in a manageable amount of time.

9. Remote Cluster Deployment

OAEMS can feasibly run on any computer, however for processor-intensive tasks such as deep learning, it is recommended that cloud computing clusters are used, so that processing can be performed on faster hardware. To this end, OAEMS has been configured to run on the cloud deep learning platform FloydHub (FloydHub, 2019). Floydhub provides a near-zero setup cloud service that allows researchers to run deep learning simulations remotely on powerful machines and monitor learning process online. By integrating Floydhub into the Deep Learning Energy Market Simulation, powerful computing resources can be leveraged by non-experts for the purpose of exploring electricity market mechanism design.

10. Current State of Experiments

OAEMS has to date been used to simulate environments with hundreds of participants. Among these participants, up to six learning agents have been deployed on a single Floydhub instance. The limiting factor appears to be the RAM allocation for each agent, but this can be overcome by provisioning additional instances and using a centralised server to run the simulator.

11. 'Deep Learning Logbook' User Interface

OAEMS also provides a user interface for researchers to explore experimental results. After every model is trained, data about agent performance and bidding behaviour is uploaded to a remote web server, and a link to the corresponding web interface is propagated via email or Slack. This link presents the user with an editable record of the details of the simulation, the deep learning hyperparameters, simulated demand and price data, overall performance metrics and the AI agent's bidding behaviour via a bid stack explorer.

The output data can be interrogated from within the user interface to better understand the bidding strategies employed by AI market participants. Without the availability of a user interface, this data can be extremely difficult to interpret using standard charting tools. Screenshots of the Deep Learning Logbook software are shown below in Figure 3.

Figure 3: Deep Learning Logbook Screenshots

12. Conclusion

Electricity market mechanisms are extremely complex and require the coordination of hundreds of market participants with varying goals and facing different incentives. These mechanisms have historically been difficult to model due to both computational limitations and the difficulty of writing artificial agent-based bidders that can act on behalf of simulated participants. The lack of sophisticated behaviour modelling in electricity markets may become more pertinent as electricity systems transition to high penetrations of renewable energy, because renewable market participants have different sets of behavioural incentives that could shape the function of electricity markets.

Recent advances in deep learning have produced extremely competent general-purpose game playing agents that can autonomously develop winning strategies that out-compete human players. By developing an electricity market simulation 'game' that is compatible with the OpenAI gym environment API, OAEMS may provide researchers with an opportunity to test market theory hypotheses with intelligent, autonomous bidders competing to maximise profit.

13. Access

OAEMS can be found on github: https://github.com/UNSW-CEEM/energy-market-deep-learning

The Deep Learning Logbook can be found on github: https://github.com/UNSW-CEEM/market-control-room

References

Atienza, R. (2018). *Advanced Deep Learning with Keras*: Packt Publishing.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). 'Openai gym'. *arXiv preprint arXiv:1606.01540*.

Chen, J. X. (2016). 'The Evolution of Computing: AlphaGo'. *18*(4), 4-7. doi:10.1109/mcse.2016.74

Chen, Y., & Ledyard, J. O. (2010). Mechanism design experiments. In *Behavioural and Experimental Economics* (pp. 191-205): Springer.

Chrabaszcz, P., Loshchilov, I., & Hutter, F. (2018). 'Back to basics: Benchmarking canonical evolution strategies for playing atari'. *arXiv preprint arXiv:1802.08842*.

Elliston, B., Riesz, J., & MacGill, I. (2016). 'What cost for more renewables? The incremental cost of renewable generation – An Australian National Electricity Market case study'. *Renewable Energy, 95*, 127-139. doi:https://doi.org/10.1016/j.renene.2016.03.080

FloydHub. (2019). Floydhub - Deep Learning Platform - Cloud GPU. Retrieved from https://www.floydhub.com

Gibney, E. (2016). 'Google AI algorithm masters ancient game of Go'. *Nature News, 529*(7587), 445.

Gorman, N., Haghdadi, N., Bruce, A., & MacGill, I. (2018). *NEMOSIS–NEM Open Source Information Service; open-source access to Australian National Electricity Market Data*.

Hintjens, P. (2013). *ZeroMQ: messaging for many applications*: " O'Reilly Media, Inc.".

Marshall, L. (2019). NEM Jenga Bid Stack Explorer. Sydney, Australia: Github. Retrieved from https://github.com/UNSW-CEEM/jenga

McLean, B., & Elkind, P. (2013). *The smartest guys in the room: The amazing rise and scandalous fall of Enron*: Penguin.

Metz, C. (2016). 'Inside OpenAI, Elon Musk's wild plan to set artificial intelligence free'. *Wired. Retrieved from https://www. wired. com/2016/04/openai-elon-musk-sam-altman-plan-to-set-artificial-intelligence-free*.

National Electricity (NSW) Law, Australiasian Legal Information Institute § 11 (1997 1997).

Rahimiyan, M., & Mashhadi, H. R. (2010). 'An Adaptive Q-Learning Algorithm Developed for Agent-Based Computational Modeling of Electricity Market'. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 40*(5), 547-556. doi:10.1109/TSMCC.2010.2044174

Roughgarden, T. (2016). *Twenty lectures on algorithmic game theory*: Cambridge University Press.

Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*: Malaysia; Pearson Education Limited.

Sessa, P. G., Walton, N., & Kamgarpour, M. (2016). 'Exploring Vickery-Clarke-Groves Mechanism for Electricity Markets'.

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., . . . Hassabis, D. (2018). 'A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play'. *Science, 362*(6419), 1140-1144. doi:10.1126/science.aar6404

Sweeney, J. L. (2002). 'The California electricity crisis: Lessons for the future'. *BRIDGE-WASHINGTON-, 32*(2), 23-31.

Wang, F.-Y., Zhang, J. J., Zheng, X., Wang, X., Yuan, Y., Dai, X., . . . Yang, L. (2016). 'Where does AlphaGo go: From church-turing thesis to AlphaGo thesis and beyond'. *IEEE/CAA Journal of Automatica Sinica, 3*(2), 113-120.

Weare, C. (2003). *The California electricity crisis: causes and policy options*: Public Policy Instit. of CA.