

Deep Reinforcement Learning with Discrete Normalized Advantage Functions for Resource Management in Network Slicing

Chen Qi, Yuxiu Hua, Rongpeng Li, Zhifeng Zhao, and Honggang Zhang

Abstract—Network slicing promises to provision diversified services with distinct requirements in one infrastructure. Deep reinforcement learning (e.g., deep Q -learning, DQL) is assumed to be an appropriate algorithm to solve the demand-aware inter-slice resource management issue in network slicing by regarding the varying demands and the allocated bandwidth as the environment *state* and the *action*, respectively. However, allocating bandwidth in a finer resolution usually implies larger action space, and unfortunately DQL fails to quickly converge in this case. In this paper, we introduce discrete normalized advantage functions (DNAF) into DQL, by separating the Q -value function as a state-value function term and an advantage term and exploiting a deterministic policy gradient descent (DPGD) algorithm to avoid the unnecessary calculation of Q -value for every state-action pair. Furthermore, as DPGD only works in continuous action space, we embed a k -nearest neighbor algorithm into DQL to quickly find a valid action in the discrete space nearest to the DPGD output. Finally, we verify the faster convergence of the DNAF-based DQL through extensive simulations.

I. INTRODUCTION

Networks are becoming increasingly agile and flexible to provision diversified services with distinct requirements on latency and rate. Specifically, network slicing, which belongs to one of cutting-edge technologies in the 5G era, allows infrastructure providers to offer “slices” of resources (computational, storage and networking) with specified service license agreements (SLAs) [1]–[4]. However, in order to fully reap the desired merits like slice-level protection, envyfreeness, and load-driven elasticity [3], [4], end-to-end network slicing still faces a lot of technical challenges. For example, taking account of the limited spectrum, the slice-level protection could guarantee superior quality of experience (QoE) but also incurs degradation in spectrum efficiency (SE). Therefore, one typical question naturally arises like that how to intelligently allocate the spectrum to slices according to the dynamics of service request from mobile users in a coherent manner [5], so as to obtain satisfactory QoE in each slice at the cost of acceptable SE.

In order to address the aforementioned problem, one potential solution is to consider the (deep) reinforcement learning (RL). As a non-nascent algorithm, RL has been widely applied

in the field of cognitive radio [6] and green communications [7]. The recent well-known application success in Go [8] further proves the feasibility to utilize neural networks (NN) to approximate the value functions in classical RL with case-testified convergence stability, and triggers tremendous research attention in the communications and networking area to solve resource allocation issues in some specific fields like power control, green communications, cloud radio access networks, mobile edge computing and caching [3]. But, a common problem in these works is that researchers usually assume a rather limited small discrete action space to ensure the necessary convergence rate. For example, [3] realized the spectrum allocation per slice on the unit of MegaHertz and accordingly design a DRL framework with tens of possible actions. But, such a coarse-grained spectrum allocation solution inevitably decreases the SE when some slice has very few service activities. In a word, it urgently needs a rethink on DRL to better avoid the curse of dimensionality and quickly converge in larger action space.

Overall speaking, this paper aims to answer how to allocate the limited spectrum on a finer-grained resolution across slices based on an improved DRL. In particular, we revolutionize the calculation and approximation of the Q -value function in the deep Q -learning (DQL) as follows:

- Inspired by the normalized advantage functions (NAF) model [9], [10], we design a discrete NAF (DNAF) NN to separately approximate a state-value function term $V(s)$ and an advantage term $A(s, a)$, where s and a denote a state and an action respectively. Moreover, we have $Q(s, a) = V(s) + A(s, a)$. Hence, the common part of the Q -value function could be learnt across all actions.
- We utilize a deterministic policy gradient descent (DPGD)-based Q -learning [16] to replace the classical statistical policy gradient descent-based DQL, so as to directly yield the most suitable action for a specific state.
- In order to solve the issue that a DPGD method ignores the discreteness of the action space, we firstly output a proto action with the largest Q -value in the virtual continuous action space and then scramble it with an extra noise term. Finally, we embed a k -nearest neighbor (k -nn) algorithm to quickly select the closest discrete action.

The remainder of the paper is organized as follows: Section II talks about some necessary algorithmic background and formulates the system model. Section III gives the details of the DNAF based Q -Learning, while Section IV presents the

The authors are with Zhejiang University, Zheda Rd. 38, Hangzhou 310027, China (mail: {qichen7c, 21631087, lirongpeng, zhaozf, honggangzhang}@zju.edu.cn). Corresponding author: R. Li.

This work was supported by National Natural Science Foundation of China (No. 61701439, 61731002), Zhejiang Key R&D Plan (No. 2019C01002), the Fundamental Research Funds for the Central Universities.

related simulation results. Finally, Section V concludes the paper with a summary.

II. MATHEMATICAL BACKGROUND AND SYSTEM MODEL

A. Mathematical Background

RL tries to find a strategy π , which maps a state $s \in \mathbb{S}$ (i.e., the varying traffic per slice) to an action $\pi(s)$, (i.e., $\mathbf{a} \in \mathbb{A}$, allocated bandwidth per slice) to maximize the discounted accumulative reward starting from the state $s^{(0)} = s$. Formally, this accumulative reward is called as a state-value function, which can be calculated by [11]

$$V^\pi(s) = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k R(s, \pi(s) | s^{(0)} = s) \right] \quad (1)$$

where the positive parameter γ is the discount factor that maps the future reward to the current state. Given the diminishing importance of future cost than the current one, γ is smaller than 1.

Q-learning is an RL technique to obtain the strategy π . Specifically, a Q-learning agent attempts to learn the value of taking a specific action for a given state, i.e. Q-value, by constantly updating Q value in a temporal difference (TD) manner as

$$Q(s, \mathbf{a}) \leftarrow Q(s, \mathbf{a}) + \alpha (R(s, \mathbf{a}) + \gamma \max_{\mathbf{a}' \in \mathbb{A}} Q(s', \mathbf{a}') - Q(s, \mathbf{a})) \quad (2)$$

where α is the learning rate and s' denotes the state of the environment after taking action \mathbf{a} at state s .

In recent years, [8] proposes to use NNs to approximate Q-value function [12] so as to solve an RL problem with a tremendous state dimension. Mathematically, DQL trains an NN with parameters θ by minimizing the loss function between the real Q-value function $Q(s, \mathbf{a})$ and an NN-approximated one $Q^+(s, \mathbf{a} | \theta)$, which can be formulated as

$$\theta = \arg \min_{\theta'} L(\theta') = \arg \min_{\theta'} (Q(s, \mathbf{a}) - Q^+(s, \mathbf{a} | \theta'))^2 \quad (3)$$

Commonly, θ could be achieved by a gradient-based approach as

$$\theta \leftarrow \theta - \alpha \nabla L(\theta) \quad (4)$$

In addition, there are some tricks that can improve the performance of DQL, such as replay buffer [13], target network [14], prioritized replay [15], etc.

B. System Model

We consider an access network scenario in Fig. 1 consisting of multiple base stations (BSs), where there exists a list of existing slices $1, \dots, N$ sharing the aggregated bandwidth W and having fluctuating demands $\mathbf{d} = (d_1, \dots, d_N)$. We aim to maximize the expectation of the utility function $\mathbb{E}\{R(\mathbf{d}, \mathbf{w})\}$ by dynamically adjusting the allocated bandwidth $\mathbf{w} = (w_1, \dots, w_N)$ to each slice, where the notation $\mathbb{E}(\cdot)$ denotes to take the expectation of the argument. Moreover, the utility function is defined as the weighted sum of SE and QoE satisfaction ratio. Mathematically,

$$R = \zeta \cdot \text{SE} + \beta \cdot \text{QoE} \quad (5)$$

where ζ and β denotes the importance of SE and QoE. Our goal is to allocate the bandwidth to slices according to the traffic variations within each slice, that is,

$$\begin{aligned} & \arg \max_{\mathbf{w}} \mathbb{E}_t \{R(\mathbf{d}, \mathbf{w})\} \\ & = \arg \max_{\mathbf{w}} \mathbb{E} \{ \zeta \cdot \text{SE}(\mathbf{d}, \mathbf{w}) + \beta \cdot \text{QoE}(\mathbf{d}, \mathbf{w}) \} \\ \text{s.t.: } & \mathbf{w} = (w_1, \dots, w_N) \\ & w_1 + \dots + w_N = W \\ & w_i = k \cdot \Delta, \forall i \in [1, \dots, N] \\ & \mathbf{d} = (d_1, \dots, d_N) \\ & d_i \sim \text{Certain Traffic Model}, \forall i \in [1, \dots, N] \end{aligned} \quad (6)$$

where t denotes the temporal index, k is an integer and Δ is the minimum allocated bandwidth per slice. Notably, $\mathbf{d}(t)$ depends on both $\mathbf{d}(t-1)$ and $\mathbf{w}(t-1)$, since the maximum sending capacity of servers belonging to one service is tangled with the provisioning capabilities for this service. For example, the TCP sending window size is influenced by the estimated channel throughput.

The key challenge to solve (6) lies in the volatile demand variations without having known a priori due to the traffic model. Hence, DQL promises to be an appropriate solution to solve this problem. But, DQL converges slowly for large action space, since DQL needs to predict Q-values for each state-action pair. Unfortunately, the size of action space $|\mathbb{A}|$ increases exponentially along with the decrease of Δ and the increase of N , since we have

$$|\mathbb{A}| = \binom{\lfloor \frac{W}{\Delta} \rfloor - 1}{N - 1} \quad (7)$$

Therefore, it is inevitable to revolutionize the classical DQL.

III. THE DNAF BASED Q-LEARNING

Researchers in [9], [10] has suggested NAF as a potential solution to DQL with continuous action space by decomposing Q-value into a state-value function V and an advantage function A . Since the discreteness of action space for resource management in network slicing is different from the continuity in [9], [10], it is quite meaningful to re-investigate its effectiveness here.

Inspired by [9], [10], we build two separate NNs for the state-value function $V(s | \theta^V)$ parameterized by θ^V and the advantage function $A(s, \mathbf{a} | \theta^A)$ parameterized by θ^A , on top of a common NN collecting some general information, that is,

$$Q(s, \mathbf{a} | \theta^Q) = V(s | \theta^V) + A(s, \mathbf{a} | \theta^A) \quad (8)$$

Moreover, the advantage function approximation leverages the DPGD algorithm [16] to obtain the proto action $\mu(s | \theta^\mu)$ by an parameterized by θ^μ by

$$\begin{aligned} & A(s, \mathbf{a} | \theta^A) \\ & = -\frac{1}{2} (\mathbf{a} - \mu(s | \theta^\mu))^T \Lambda(s | \theta^L) \Lambda(s | \theta^L)^T (\mathbf{a} - \mu(s | \theta^\mu)) \end{aligned} \quad (9)$$

where $\Lambda(s | \theta^L)$ is a low-triangular matrix whose entries come from a linear output layer of an NN parameterized by θ^L [9]. Hence, θ^A is a concatenation of θ^μ and θ^L .

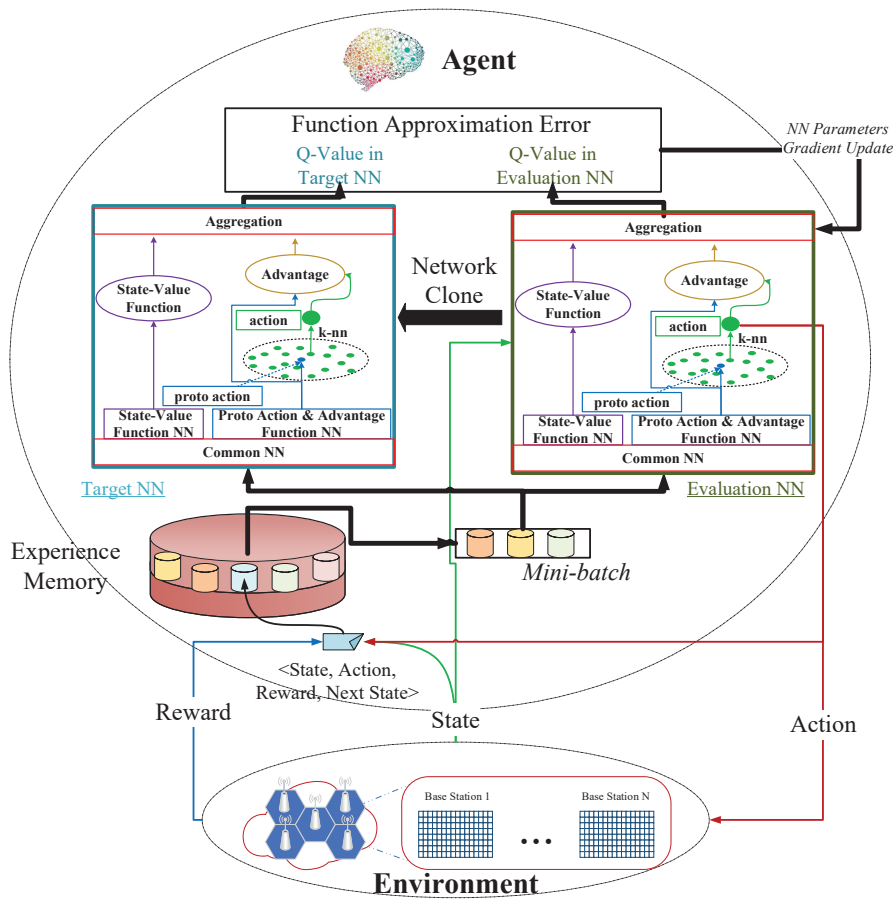


Fig. 1: An illustration of DNAF-based Q -learning for resource management in network slicing.

Since the action $\mu(s|\theta^\mu)$ might be an invalid action in the discrete action space \mathbb{A} , we first calculate the proto action \hat{a}^1 by adding $\mu(s|\theta^\mu)$ with a noise term. Notably, the added noise in selecting action procedure could be regarded as an alternative means to the ϵ -greedy strategy in classical RL for state exploration. Afterwards, k-nn is applied to find the closest valid action, that is,

$$g_k(\mathbf{a}) = \arg \min_{\mathbf{a}' \in \mathbb{A}} \|\mathbf{a}' - \hat{\mathbf{a}}\|_2 \quad (10)$$

In other words, the function $g_k(\cdot)$ is a k-nn mapping from a continuous space to a discrete set and returns k valid actions closest to the proto action. For $k > 1$, it also means to obtain the k nearest actions that maximize Q -value [17]. The case $k = 1$, which belongs to the key focus in this paper, is equivalent to simply select the nearest action.

We incorporate the aforementioned methods into the DQL and have the DNAF-based Q -learning in Algorithm 1.

IV. SIMULATION RESULTS AND NUMERICAL ANALYSES

In this part, we compare the convergence rate of the DNAF-based DQL and classical DQL. We simulate in one single BS scenario with three types of services (i.e., VoLTE, video, ultra-reliable low-latency communications (URLLC)) as in [3],

¹As depicted in Fig. 1, the proto action is plotted in the blue circle while the valid actions are presented in green ones.

[18] and correspondingly have three slices. Moreover, we attempt to allocate 10-MegaHertz bandwidth to these three slices and set the minimal bandwidth allocation resolution is 0.2-MegaHertz, thus leading to 1176 valid actions. Meanwhile, the network slice stops sending new packets to one user if 5 packets in the caching buffer for this user have not been successfully delivered or expired (e.g., exceeding the tolerant delay for that slice). Otherwise, each network slice sends traffic to its user following the settings in Table I.

Fig. 2 gives an illustration of the reward variations with respect to the episode index. Here, the reward is defined in (5). As for the DNAF-based DQL, inspired by the noise settings in [13], we first assume the noise obeys normal distribution and is multiplied by an attenuating coefficient, which gradually decays with the number of iterations and ultimately fixes at zero after 3000 iterations. It can be observed from Fig. 2 that regardless of the values of α , the DNAF-based DQL could converge after 4000 - 6000 episodes, while the classical DQL still changes dramatically with no sign of convergence even after 10000 episodes. Therefore, it could safely come to the conclusion that the DNAF-based DQL could converge more rapidly than the classical DQL. On the other hand, Fig. 2 also provides the performance result of the equal-allocation strategy where we intuitively allocate the bandwidth according to the number of slices and verified that the DNAF-based DQL yield superior performance than the equal-allocation

Algorithm 1 The DNAF based Q -learning

- 1: Randomly initialize a normalized Q network $Q(s, a|\theta^Q)$, where θ^Q is a concatenation of θ^V and θ^A .
- 2: Initialize a target network Q' with weight $\theta^{Q'} \leftarrow \theta^Q$.
- 3: Initialize replay buffer $\mathbb{R} \leftarrow \emptyset$.
- 4: Initialize an episode index $t = 0$.
- 5: **repeat**
- 6: At episode t , the agent observes the state s_t .
- 7: The agent calculates a proto-action $\hat{a}_t = \mu(s_t|\theta^\mu) + \mathcal{N}_t$ and determines the closest action $a_t = g_1(\hat{a}_t)$ by (10).
- 8: The agent receives the reward $R(s_t, a_t)$ and observes a new state s_{t+1} for the environment.
- 9: The agent stores the episode experience (s_t, a_t, R_t, s_{t+1}) in \mathbb{R} .
- 10: The agent samples a mini-batch \mathbb{R}_{mbatch} of experiences from \mathbb{R} .
- 11: The agent sets $y_i = R_i + \gamma V'(s_{i+1}|\theta^{Q'})$ and gets estimated Q -value $Q_i = Q(s_i, a_i|\theta^Q)$ by (5) and (6), $\forall (s_i, a_i, R_i, s_{i+1}) \in \mathbb{R}_{mbatch}$.
- 12: The agent updates the weights θ^Q for the evaluation network by leveraging gradient descent algorithm to $\frac{1}{|\mathbb{R}_{mbatch}|} \sum_i ((y_i - Q_i)^2)$.
- 13: The agent clones the Q network to the target network Q' every C episodes by assigning the weights Q' as $\theta^{Q'} = \theta^Q$.
- 14: The episode index is updated by $t \leftarrow t + 1$.
- 15: **until** A predefined stopping condition (e.g., the gap between y_i and Q_i , the episode length, etc) is satisfied.

TABLE I: A Brief Summary of Key Settings for Traffic Generation Per Slice

	VoLTE	Video	URLLC
Bandwidth	10 MHz		
Scheduling	Round robin per slot (0.5 ms)		
Slice Band Adjustment (Q-Value Update)	1 second (2000 scheduling slots)		
Channel	Rayleigh fading		
User No. (100 in all)	46	46	8
Distribution of Inter-Arrival Time per User	Uniform [Min = 0, Max = 160ms]	Truncated Pareto [Exponential Para = 1.2, Mean = 6 ms, Max = 12.5 ms]	Exponential [Mean = 180 ms]
Distribution of Packet Size	Constant (40 Byte)	Truncated Pareto [Exponential Para = 1.2, Mean = 100 Byte, Max = 250 Byte]	Truncated Lognormal [Mean = 2 MB, Standard Deviation = 0.722 MB, Maximum = 5 MB]
SLA: Rate	51 kbps	5 Mbps	10 Mbps
SLA: Latency	10 ms	10 ms	5 ms

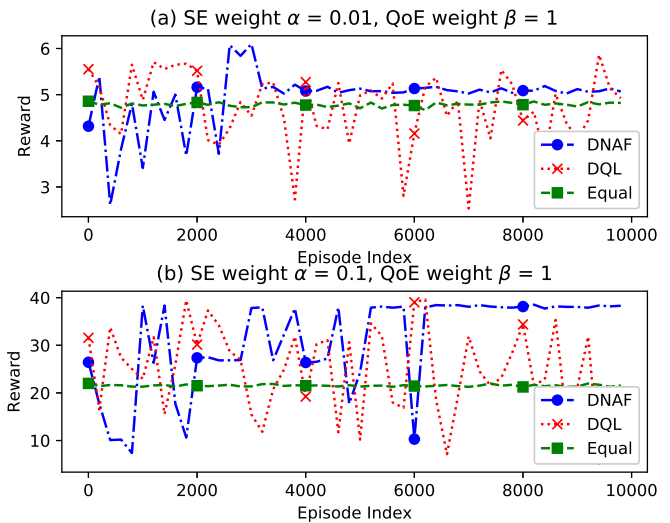


Fig. 2: The convergence rate of the DQF-based DQL and classical DQL.




Fig. 3: The cumulative reward of the DQF-based DQL, classical DQL, equal-allocation strategy.

strategy. Furthermore, Fig. 3 gives the cumulative reward of the DQF-based DQL, classical DQL, equal-allocation strategy. It can be found in Fig. 3 that the DQF-based DQL could obtain superior performance than the equal-allocation strategy while DQL yield some inferior performance than the equal-allocation strategy due to its slow convergence rate.

Fig. 4 presents the SE and QoE satisfaction ratio applying the learnt policy after 10000 episodes. It can be observed that when $\alpha = 0.01$, which implies that SE is on a par with QoE satisfaction ratio, the DQF-based DQL could yield superior performance on SE and QoE satisfaction ratio simultaneously than the classical DQL. On the other hand, when α takes a larger value (i.e., 0.1) to put more focus on SE, compared than the classical DQL, the DQF-based DQL learns a policy giving significantly higher SE but degrades the QoE satisfaction ratio for some slices. Notably, some evaluation metrics produced by the DQF-based DQL policy are not always superior (or even inferior) to the classical




Fig. 4: The snapshot of SE and QoE satisfaction ratio applying the learnt policy after 10000 episodes.




Fig. 5: The reward of the DQAF-based DQL with the action noise obeying normal distribution and uniform distribution.

DQL, since it is still a very challenging research topic to design RL with multiple conflicting rewarding metrics. In this paper, we simply choose the weighted sum of two conflicting metrics (i.e., SE and QoE) as the reward in RL. Despite the intuitiveness of this direct summation, our simulation results have demonstrated that we cannot guarantee to simultaneously obtain superior performance for both SE and QoE. Therefore, we have left this inspiring and interesting topic as our future works.

Fig. 5 shows the comparison between normal distributed noise and uniform distributed ones, respectively. It can be observed that both cases could lead to convergent learning policy and exhibit trivial performance difference.

V. CONCLUSION

In this paper, we have discussed how to accelerate the convergence rate of the classical DQL in large action space,

so as to satisfy the requirements for finer-resolution resource management in network slicing. In particular, we have applied the DQAF into DQL, by separating the Q -value function as a state-value function term and an advantage term and exploiting a DPGD algorithm to avoid the unnecessary calculation of Q -value for every state-action pair. Furthermore, we have embedded a k-nn algorithm into DQL to quickly find a valid action in the discrete action space. We have also verified that compared than the classical DQL, the DQAF-based DQL exhibits faster convergence and superior performance. Hence, we believe our works could contribute to enhancing the applicability of DQL in network slicing. However, there still exist some research issues to be solved, in particular the need to further improve DQL to guarantee minimal slice SLAs, capably adapt non-stationary traffic demands, and smartly design the reward for multi-conflicting metrics.

REFERENCES

- [1] X. Zhou, *et al.*, "Network slicing as a service: Enable industries own software-defined cellular networks," *IEEE Commun. Mag.*, vol. 54, no. 7, pp. 146–153, Jul. 2016.
- [2] R. Li, *et al.*, "Intelligent 5G: When cellular networks meet artificial intelligence," *IEEE Wireless Commun.*, vol. 5, no. 24, pp. 175–183, Oct. 2017.
- [3] R. Li, *et al.*, "Deep reinforcement learning for resource management in network slicing," *IEEE Access*, vol. 6, pp. 74 429–74 441, Nov. 2018.
- [4] J. Zheng and G. de Veciana, "Elastic multi-resource network slicing: can protection lead to improved performance?" *arXiv:1901.07497 [cs]*, Jan. 2019. [Online]. Available: <http://arxiv.org/abs/1901.07497>
- [5] S. Vassilaras, *et al.*, "The algorithmic aspects of network slicing," *IEEE Commun. Mag.*, vol. 55, no. 8, pp. 112–119, 2017.
- [6] X. Chen, *et al.*, "Stochastic power adaptation with multiagent reinforcement learning for cognitive wireless mesh networks," *IEEE Trans. Mob. Comput.*, vol. 12, no. 11, pp. 2155–2166, Nov. 2013.
- [7] R. Li, *et al.*, "TACT: A transfer actor-critic learning framework for energy saving in cellular radio access networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 4, pp. 2000–2011, Apr. 2014.
- [8] V. Mnih, *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015. [Online]. Available: <http://www.nature.com/nature/journal/v518/n7540/full/nature14236.html>
- [9] S. Gu, *et al.*, "Continuous deep Q-learning with model-based acceleration," *arXiv:1603.00748 [cs]*, Mar. 2016. [Online]. Available: <http://arxiv.org/abs/1603.00748>
- [10] Z. Wang, *et al.*, "Dueling network architectures for deep reinforcement learning," *arXiv:1511.06581 [cs]*, Nov. 2015. [Online]. Available: <http://arxiv.org/abs/1511.06581>
- [11] R. Sutton and A. Barto, *Reinforcement learning: An introduction*. Cambridge University Press, 1998. [Online]. Available: <http://webdocs.cs.ualberta.ca/~sutton/book/ebook/>
- [12] K. Hornik, *et al.*, "Multilayer feedforward networks are universal approximators," *Neural Networks*, vol. 2, no. 5, pp. 359–366, Jan. 1989.
- [13] T. P. Lillicrap, *et al.*, "Continuous control with deep reinforcement learning," *arXiv:1509.02971 [cs, stat]*, Sep. 2015. [Online]. Available: <http://arxiv.org/abs/1509.02971>
- [14] H. v. Hasselt, *et al.*, "Deep reinforcement learning with double Q-learning," in *Proc. AAAI 2016*, ser. AAAI'16. Phoenix, Arizona: AAAI Press, 2016, pp. 2094–2100. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3016100.3016191>
- [15] T. Schaul, *et al.*, "Prioritized Experience Replay," *arXiv:1511.05952 [cs]*, Nov. 2015. [Online]. Available: <http://arxiv.org/abs/1511.05952>
- [16] D. Silver, *et al.*, "Deterministic policy gradient algorithms," in *Proc. ACM ICML 2014*, Beijing, China, Jun. 2014. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3044805.3044850>
- [17] G. Dulac-Arnold, *et al.*, "Deep reinforcement learning in large discrete action spaces," *arXiv:1512.07679 [cs, stat]*, Dec. 2015. [Online]. Available: <http://arxiv.org/abs/1512.07679>
- [18] NGMN, "NGMN radio access performance evaluation methodology," Oct. 2007. [Online]. Available: https://www.ngmn.org/fileadmin/user_upload/NGMN_Radio_Access_Performance_Evaluation_Methodology.pdf