DATA-DRIVEN EVALUATION OF TRAINING ACTION SPACE FOR REINFORCEMENT LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Training action space selection for reinforcement learning (RL) is conflict-prone due to complex state-action relationships. To address this challenge, this paper proposes a Shapely-inspired methodology for training action space categorization and ranking. To reduce exponential-time Shapely computations, the methodology includes a Monte Carlo simulation to avoid unnecessary explorations. The effectiveness of the methodology is illustrated using a cloud infrastructure resource tuning case study. It reduces the search space by 80% and categorizes the training action sets into dispensable and indispensable groups. Additionally, it ranks different training actions to facilitate superior RL model performance and lower cost. The proposed data-driven methodology is extensible to different domains, use cases, and machine learning algorithms.

1 INTRODUCTION

A reinforcement learning (RL) agent learns how to map situations to actions in order to maximize a long-term cumulative award signal in a given environment. Figure 1 shows the various artifacts of an RL algorithm. An RL problem is defined by a quartet of (S, A, P_a, R_a), where S is a set of states or the state space; A is a set of actions or the action space available to influence S; $P_a(s, s') = Pr(s_{t+1} = s'|s_t = s, a_t = a)$ is the transition probability which is the probability that action a in state s at time t will lead to state s' at time $t + 1$; and finally, $R_a(s, s')$ is the immediate reward signal received after transitioning from state s to state s', due to action a. An RL agent training involves recognizing an optimal policy function from a corpus of $\{(s_i, a_i)\}$. The reward function, R_a, is defined ab initio for an efficient goal accomplishment. The transition probability or state-action mapping is defined by the environment dynamics. In many real-life use cases, the agent cannot directly sense the effect of its actions on the environment. This is particularly true when the state-action relationship cannot be modeled by either closed-form analytical expressions [23, 25, 51] or explicit rules as in the popular games such as Chess [24], Go [41, 40], and Atari [30]. This challenge is well documented in the literature [44, 50]. To address this challenge for RL model training [45, 29], simulation-based action models [43] play a pivotal role. The optimal choice of simulation parameters such as the training action space is a non-trivial challenge because of the curse of dimensionality [15] and non-linearity [11].

Training data valuation [20, 10] and associated artisanal software engineering efforts constitute a large part of the machine learning (ML) life cycle (or MLOps). Yet, most research and development efforts [32] focus on algorithms and infrastructure. Production-grade MLOps needs to handle data lifecycle management (DLM) [36] challenges including: fairness and bias in labeled datasets [14], data quality [9], limitations of benchmarks [38], and reproducibility concerns [35]. For RL, the DLM challenges are further compounded due to complexities arising from non-linear state-action interactions, partially-observable processes, non-isometric action spaces [28], and strong domain specificity [19] of action models. With the recent emphasis on ML explainability, traditional supervised learning and deep learning research communities are actively working on systematic data-driven frameworks [13, 46] for training data valuation. We need equivalent frameworks for RL to streamline conflict-prone training action selection [8, 33, 16]. Depending on the agent-environment interactions, different training actions have different relative contributions to the RL agent performance. Some training actions are indispensable because of their unique positions in the parametric space. Other dispensable actions have different relative contributions to the reward function. Remarkably, a high-cardinality training action space in many cases leads to lower cumulative reward than a well-designed training

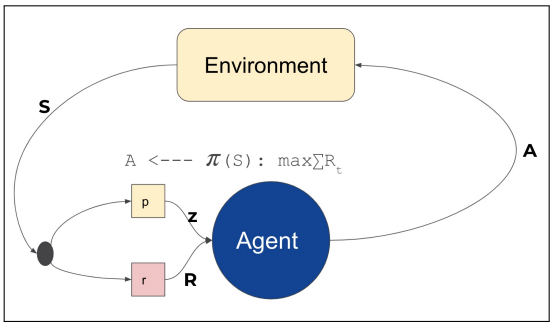

Figure 1: The schematic diagram for a typical reinforcement learning algorithm

action space. To the best of authors' knowledge, there is hardly any data-driven tool for the evaluation of training action space for RL. Such a tool leads to superior RL agent performance, lower model training and maintenance cost, and strong multi-disciplinary collaboration [31].

This paper proposes a Shapely-inspired [47] algorithm to categorize and rank training action sets. It also assists in recognizing cut-off cardinality for the training action space to reduce unnecessary exploration and ensure polynomial time complexity. Additionally, a gradient-based action update rule is proposed for faster convergence for the RL loop. While Section 2 describes the algorithms for RL training action space evaluation, Section 3 illustrates the effectiveness of the algorithms in a specific case study. Finally, Section 4 present a succinct summary with possible future directions.

2 Algorithm

In RL agent design, training action selection is a conflict-prone process. This paper offers data-driven tools for streamlining RL training action selection. We present a Shapely-inspired algorithm [42], as shown in Algorithm 1, for optimal training action selection. The algorithm uncovers the indispensability of some training action sets for a given RL task. A training action set is deemed to be indispensable *iff* upon removal of that training action set the RL agent fails to accomplish the goal within a finite number of steps [22, 49, 21], $\mathcal{O}(1/\epsilon^2)$ (ϵ is error limit for the ML model for the environment, $\mathcal{F}$), and/or without escaping the parametric action boundary. On the other hand, the algorithm computes the cumulative reward for a given initial starting action, a_0, for each dispensable training action set. The cumulative reward can be used as a quantitative measure for the utility of a training action set for the given initial action, a_0 and could be used as a basis for ranking different training action sets.

As a Shapely-inspired algorithm, the cardinality of the training action set is a critical design consideration because the possible action space grows exponentially with the number of training actions. To address this computational issue, Algorithm 2 uses a truncated Monte Carlo method to compute a cut-off cardinality value and reduces unnecessary exploration.

Another important consideration is the step size for the action update. It is a conflict-prone design decision: while a small action update step leads to prohibitively slow convergence, a large action update step causes divergence. For a multi-dimensional non-isometric action space, the problem becomes more complicated because the movements in different action directions bring non-uniform cumulative rewards. Algorithm 3 proposes a computational procedure for optimal step-size for action update.

3 Case Study

In this section, we design a resource tuning example (in the cloud) to illustrate the effectiveness of the proposed algorithm for the training action space evaluation. We evaluate the performance of an MDP [18]-based RL agent, as shown in Figure 1. The state-action mapping and transition probability are modeled using time-series auto-regression [34] after a PCA-based dimensionality reduction [37] with time complexity $\mathcal{O}(klogk)$, where k is the number of principal components. The choice of the

Algorithm 1 Data-driven categorization and ranking for training action space evaluation

Input: Training state-action pairs: $(s_1, a_1), (s_2, a_2), .., (s_n, a_n)$, learning algorithm for the environment, $\mathcal{F}$, a reward function, threshold condition, step-size for action update, and an upper bound for the acceptable numbers of steps.

Output: Categorization vector, $\mathcal{V}$, for different training actions for a given initial action, a_0, and a parametric action boundary. It takes two different categorical values: {*dispensable*, *indispensable*}. Also, rank dispensable training actions based on their cumulative rewards.

```
S(A) ∈ a_i                                   ▷ This is the power set for the training action space A
S(S|A) ∈ s_i                                 ▷ The state space corresponding to each action set within the action power set
repeat
    for < a_i, s_i > do
        reward ← 0
        a_j ← a_o                            ▷ Initial action
        s_j ← F(a_j|a_i, s_i)                ▷ Machine learning model for the environment
        while Threshold condition is not satisfied and action is within the parametric boundary do
            a_j ← a_j + δ                    ▷ Action update
            s_j ← F(a_j|a_i, s_i)            ▷ State update
            reward ← reward − 1              ▷ More negative reward with time
            if Threshold condition is never satisfied within an acceptable number of steps then
                V(A − a_i) ← indispensable
            else
                V(a_i) ← dispensable
            end if
        end while
    end for
until all necessary training action sets are evaluated.
Ranking Vector, R ← dispensable actions sorted in order of the corresponding cumulative rewards
```

Algorithm 2 Monte Carlo method for cut-off cardinality computation for the training action space

Input: The power set for the training action space, $\mathcal{S}(A)$, the corresponding state space, $\mathcal{S}(S|A)$, learning algorithm for the environment, $\mathcal{F}$, an acceptable error margin, ϵ, and an acceptable time complexity, η.

Output: Cut-off cardinality, $|A|_{cutoff}$: all the other lower cadinality can be safely neglected.

```
S(A) ∈ a_i                                   ▷ This is the power set for the training action space A
S(S|A) ∈ s_i                                 ▷ The state space corresponding to each action set within the action power set
for k ← |A| to 1 do
    repeat
        s_j ← F(a_j|a_i, s_i)                ▷ a_j is the cross-validation point.
        if Error margin, ε, and computational complexity, η, are satisfied then
            status ← success
        else
            status ← failure
        end if
    until all necessary training action sets for a given cardinality is exhausted.
    if all statuses for a given cardinality, k, are failures then
        |A|_cutoff ← k + 1
    end if
end for
```

Algorithm 3 Computation of optimal step-size for action update

Input: A training action space, $\mathcal{A}$, the corresponding state space, $\mathcal{S}|\mathcal{A}$, the test statistic, t, and a threshold condition for the test statistic, c.

Output: Action update size, δ, for an optimal RL agent. Depending on the action space cardinality, δ can be multi-dimensional.

$\mathcal{S}(A) \in a_i$ ▷ This is the power set for the training action space A

$\mathcal{S}(S|A) \in s_i$ ▷ The state space corresponding to each action set within the action power set

while Threshold condition is not satisfied **do**

  $\delta_i \leftarrow \mathrm{a}_i(t-c) + b_i(t+c) + c_i(t)$ ▷ action update in i dimension

  $f \leftarrow \delta_i - a_i(t-c) - b_i(t+c) - c_i(t)$ ▷ functional representation

end while

while Convergence criteria is not satisfied **do**

  $(\mathbf{a}, \mathbf{b}, \mathbf{c})_{n+1} = (\mathbf{a}, \mathbf{b}, \mathbf{c})_n - \gamma \nabla f((\mathbf{a}, \mathbf{b}, \mathbf{c})_n)$

end while

Table 1: Relevant case study for cloud resource tuning

RL Artifact	Description
State	CPU utilization metrics
State Statistics	Median value of CPU utilization
Threshold	90% of CPU utilization
Action	Resource configuration set points: (# of vCPUs, Memory Size (GB))
Reward	Negative of total time steps required to satisfy the threshold condition
Parametric Boundary	Polygon defined by the parametric endpoints
Initial Action	$(6, 14)$: Arbitrarily assigned WLOG
Error Margin	5%
Acceptable Steps	400
Computational Complexity	Polynomial time

algorithms is purely driven by the nature of the training dataset. A more complex non-linear dataset warrants more complex sequential models such as long short-term memory (LSTM) [39] or attention based models [26]. For action updates, WLOG, an RL-based PID controller [1, 12] is used. Similar to time-series modeling, action update can be conducted by other policy learning algorithms such as SARSA [17].

As shown in Table 1, the objective of the RL agent for this case study is to quickly reduce high CPU utilization below a pre-assigned threshold for a given workload. In most infrastructure/cloud resource tuning technologies [48], CPU utilization represents a key metric. Therefore, the state space for this RL case study consists of virtual machine CPU utilization (%) metrics and the action space is defined by the VM resource set points, i.e., (# of vCPUs, Memory Size (GB)). The RL agent uses AWS boto3 SDK [2] to manipulate actions and AWS CloudWatch [3] for state space monitoring. The reward is defined as the number of time steps required by the agent to accomplish the objective multiplied by -1. The negative reward per time step was meant to push the agent to accomplish the task as fast as possible.

The training data for this case study was generated internally [4] with an open source library, *stress* (https://linux.die.net/man/1/stress). It uses a rectangular workload. The peak of the workload uses the stress command: *sudo stress –io 4 –vm 2 –vm-bytes 1024M –timeout 500s*. Essentially, the peak is running 4 *I/O* stressors and 2 VM workers spinning on malloc with 1024 MB per worker for 500 s. The simulated rectangular workload has a time period of 600 s: a high-stress phase of 500 s is followed by an inactive phase of 100 s. The RL training action space is spanned by the power set drawn from the five pairs of EC2 configuration set points as shown in Table 2. Using Algorithm 2, we notice that the power set below the cut-off cardinality of 4 produces trivial and unstable results. Therefore, WLOG, the analysis in this paper has been focused on six training action sets as shown in Table 4. That amounts to 81.25% reduction in the search space. For each EC2 instance in the training action space, the corresponding state metrics, i.e., CPU utilization (%) are shown in Figure 2. The

Table 2: Five different AWS EC2 resource pairs used in training action space

EC2 Type	of vCPUs	Memory Size (GB)
small t3a	2	2
medium t3a	2	4
large t3a	2	8
xlarge t3a	4	16
2xlarge t3a	8	32

Table 3: Training action categorization based on the valuation vector

Training Action	Category
small t3a	dispensable
medium t3a	dispensable
large t3a	indispensable
xlarge t3a	indispensable
2xlarge t3a	indispensable

training data was collected for a 24 hour period with 1 minute sampling interval. Using Algorithm 3, the optimal step-size for action update is identified to be 0.1 in both # of vCPUs and Memory Size (GB) dimensions.

With this set up, different RL models are developed with different training action sets and the corresponding rewards and ranks are noted in Table 4. Remarkably, a high cardinality training action set does not guarantee the best reward: the agent with all training actions,<*small t3a, medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>, does not yield the highest reward. In fact, the action set with <*medium t3a, large t3a, xlarge t3a, 2xlarge t3a*> yields the highest reward. This pattern could be attributed to the state-action interaction in this particular case study. First, the parametric distances between different EC2 instance pairs are not uniform. While the Euclidean distance between *small t3a* and *medium t3a* is equal to 2, that between *xlarge t3a* and *2xlarge t3a* is 16.5. The non-uniform spacing for training action space is a considerable deterrent [27] for RL adoption. Second, in this case study, the transient CPU utilization (%) patterns have undergone a material change from *xlarge t3a* (max 100%) to *2xlarge t3a* (max 73%). This indicates the strong influence of *2xlarge t3a* for the given RL task. Indeed, we noticed *2xlarge t3a* to be an indispensable training action. As shown in Table 3, a categorization of training actions can be inferred based on the valuation vector, $\mathcal{V}$, as described in Algorithm 1: two *dispensable* training actions are uncovered to be *small t3a, medium t3a* and three *indispensable* training actions to be *large t3a, xlarge t3a, 2xlarge t3a*.

Figure 3 shows the RL loop action with three training actions: <*small t3a, medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>, <*medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>, and <*small t3a, large t3a, xlarge t3a, 2xlarge t3a*>.

- For the first training action set of <*small t3a, medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>, the reward is -21. The left subplot in Figure 3(a) shows how the recommended action is evolving with time from an *arbitrary* initial point of (6, 14). The recommended points are superimposed on the training action parameter space to illustrate their relative position with respect to the parametric boundary which is defined by the trapezium with vertices: *{(2,2), (2,8), (8,32), (8,2)}* in the (# of vCPUs, Memory Size (GB)) space. The right subplot in Figure 3(a) shows how the median CPU utilization (%) comes below the *90%* threshold in 21 steps leading to -21 in reward. The blue dots represent the median CPU utilizations for different training set points.
- For the second training action set of <*medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>, the reward is -13 as shown in Figure 3(b).
- For the third training action set of <*small t3a, large t3a, xlarge t3a, 2xlarge t3a*>, the reward is -16 as shown in Figure 3(c).

Table 4: Rewards and ranks for different training actions. Some training action sets fail to satisfy the objective. Therefore, the rewards and ranks for them are noted as *none*
.

Training Action Set	Reward	Rank
<*small t3a, medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>	-21	3
<*medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>	-13	1
<*small t3a, large t3a, xlarge t3a, 2xlarge t3a*>	-16	2
<*small t3a, medium t3a, xlarge t3a, 2xlarge t3a*>	*none*	*none*
<*small t3a, medium t3a, large t3a, 2xlarge t3a*>	*none*	*none*
<*small t3a, medium t3a, large t3a, xlarge t3a*>	*none*	*none*

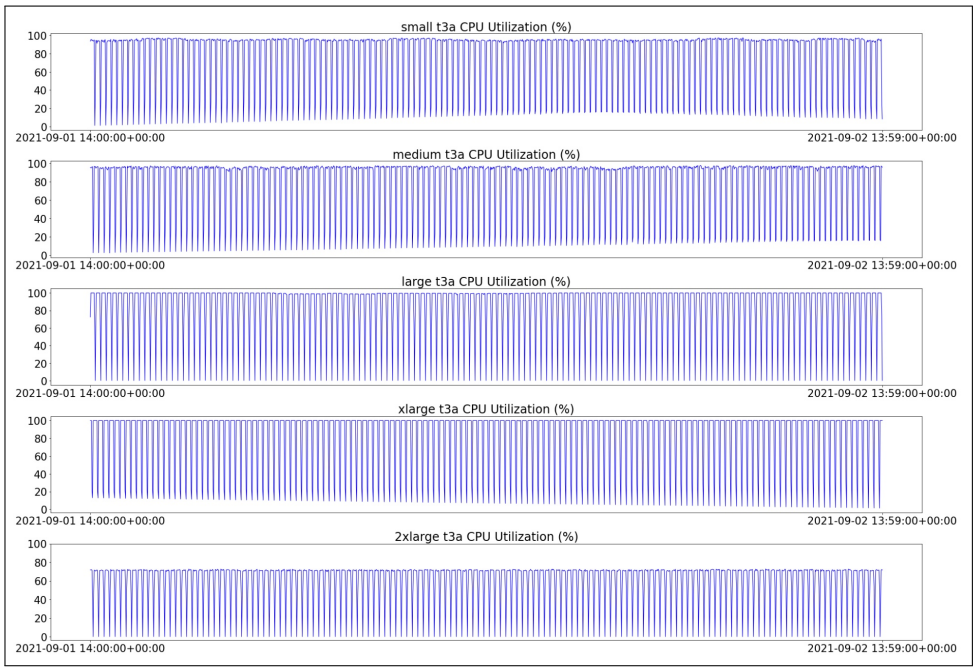

Figure 2: CPU utilization (%) responses on five different EC2 instances (Table 2) from 2PM-UTC 9/1/2021-2PM-UTC 9/2/2021 at 1 minute granularity. This training data was generated internally [4]. The median CPU utilization values (%) are noted to be equal to {95%, 95.5%, 99.5%, 100%, 72.5%}
.

As shown in Table 4, the reward scores can indeed be used for ranking different training action sets leading to a data-driven approach for training action selection. As shown in Figure 4, *large t3a* is an indispensable element in the training action space for the given RL agent. Without this training action, the RL agent fails to accomplish the goal of bringing the CPU utilization below the critical threshold. Similar observations can be made about *xlarge t3a* and *2xlarge t3a*.

4 Conclusion

This paper proposes a data-driven methodology for training action space evaluation for RL. The methodology offers a principled framework for training action space categorization and ranking within a finite computational time. It unleashes a strategy for superior model performance and lower modeling cost. Additionally, the proposed methodology is completely agnostic of use cases and machine learning algorithms. Therefore, it is a general-purpose methodology extensible to different domains including distributed computing, network traffic control, healthcare, automatic locomotion, building management system, and industrial controls, and different machine learning algorithms such as PCA, Autoencoder, ARIMA, LSTM, Transformer, PID, SARSA, DQN, and many others. For the

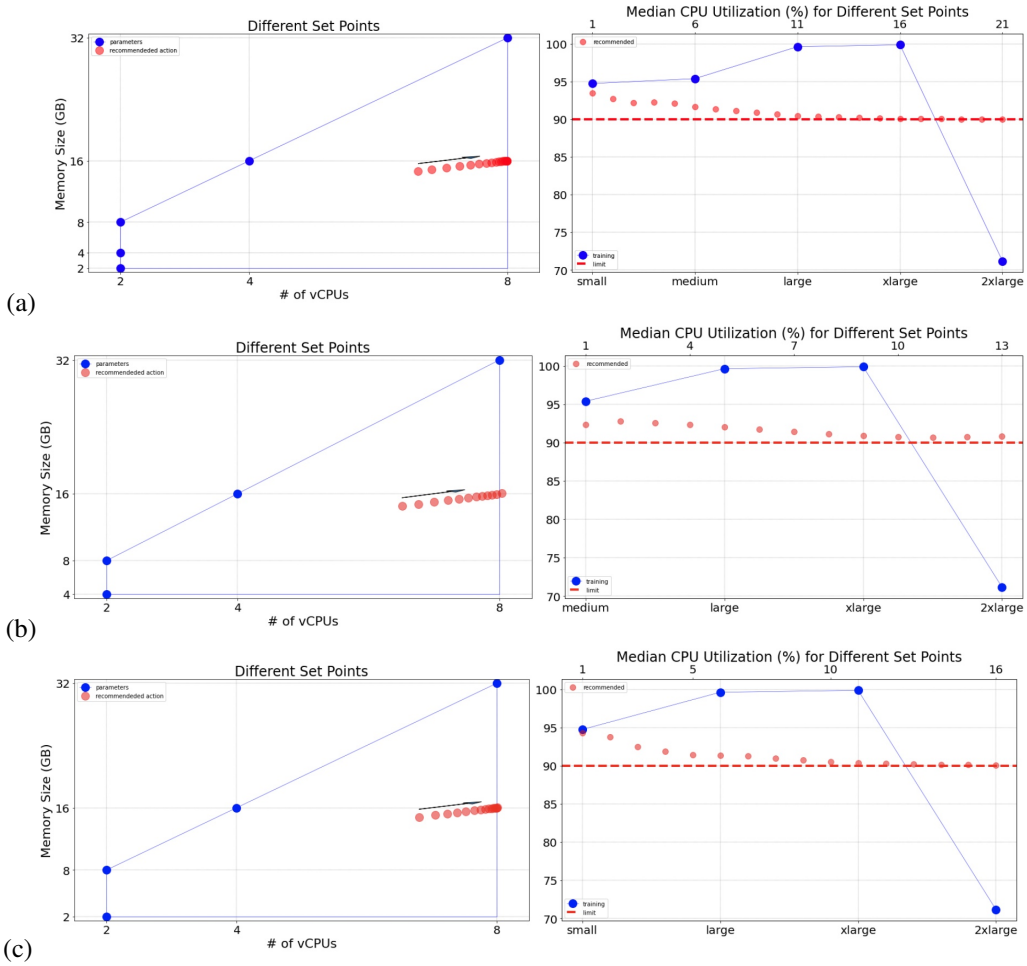

Figure 3: Examples of dispensable actions. (a) RL loop with all training actions, <*small t3a, medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>. The reward is noted to be -21. (b) RL loop with a training action set of <*medium t3a, large t3a, xlarge t3a, 2xlarge t3a*>. The reward is noted to be -13. (c) RL loop with a training action set of <*small t3a, large t3a, xlarge t3a, 2xlarge t3a*>. The reward is noted to be -16. Both *small t3a* and *medium t3a* are noted to be dispensable actions.

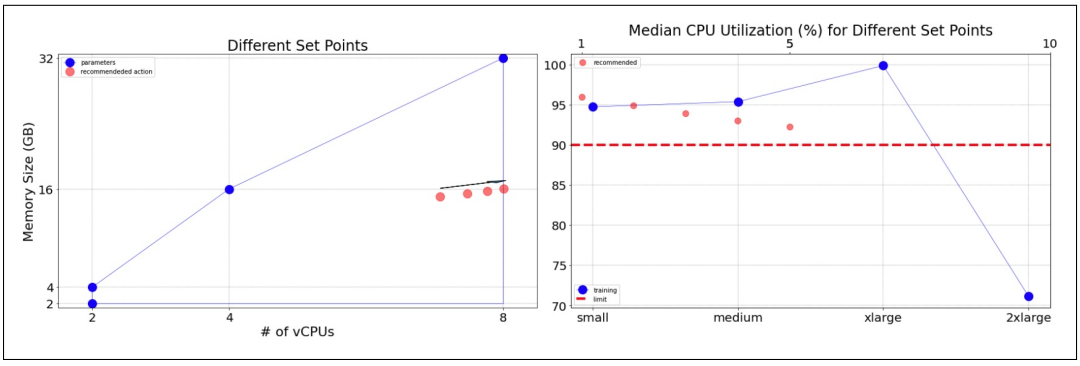

Figure 4: RL loop with <*small t3a, medium t3a, xlarge t3a, 2xlarge t3a*> action space. The agent could never accomplish the goal, therefore, *large t3a* is an indispensable element in the training action space. Similarly, *xlarge t3a* and *2xlarge t3a* are two indispensable elements in the training action space.

next phase of the development for this data-agnostic methodology, we are planning to contribute a general RL design library to relevant open source projects [5, 6, 7].

REFERENCES

[1] https://www.ni.com/en-us/innovations/white-papers/06/pid-theory-explained.html. [Online; accessed 30-September-2021].

[2] https://instances.vantage.sh/. [Online; accessed 20-September-2021].

[3] https://aws.amazon.com/cloudwatch/. [Online; accessed 30-September-2021].

[4] https://github.com/rghosh8/ec2-rl-dataset. [Online; accessed 30-September-2021].

[5] https://github.com/ray-project/ray/tree/master/python/ray. [Online; accessed 30-September-2021].

[6] https://gym.openai.com/. [Online; accessed 30-September-2021].

[7] https://github.com/kubeflow/katib. [Online; accessed 30-September-2021].

[8] James F Cavanagh et al. "Conflict acts as an implicit cost in reinforcement learning". In: *Nature communications* 5.1 (2014), pp. 1–10.

[9] Kate Crawford and Trevor Paglen. *Excavating ai*. 2019.

[10] Daniel Fryer, Inga Strümke, and Hien Nguyen. "Shapley values for feature selection: the good, the bad, and the axioms". In: *arXiv preprint arXiv:2102.10936* (2021).

[11] Paul Garnier et al. "A review on deep reinforcement learning for fluid mechanics". In: *Computers & Fluids* 225 (2021), p. 104973.

[12] Meysam Gheisarnejad and Mohammad Hassan Khooban. "An intelligent non-integer PID controller-based deep reinforcement learning: Implementation and experimental results". In: *IEEE Transactions on Industrial Electronics* 68.4 (2020), pp. 3609–3618.

[13] Amirata Ghorbani and James Zou. "Data shapley: Equitable valuation of data for machine learning". In: *International Conference on Machine Learning*. PMLR. 2019, pp. 2242–2251.

[14] Naman Goel and Boi Faltings. "Deep bayesian trust: A dominant and fair incentive mechanism for crowd". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 01. 2019, pp. 1996–2003.

[15] Jiequn Han, Arnulf Jentzen, and E Weinan. "Overcoming the curse of dimensionality: Solving high-dimensional partial differential equations using deep learning". In: *arXiv preprint arXiv:1707.02568* (2017), pp. 1–13.

[16] Yousuf Hashmy et al. "Wide-area measurement system-based low frequency oscillation damping control through reinforcement learning". In: *IEEE Transactions on Smart Grid* 11.6 (2020), pp. 5072–5083.

[17] J Fernando Hernandez-Garcia and Richard S Sutton. "Understanding multi-step deep reinforcement learning: a systematic study of the DQN target". In: *arXiv preprint arXiv:1901.07510* (2019).

[18] Mark N Howell and Matt C Best. "On-line PID tuning for engine idle-speed control using continuous action reinforcement learning automata". In: *Control Engineering Practice* 8.2 (2000), pp. 147–154.

[19] Vladimir Ilievski et al. "Goal-oriented chatbot dialog management bootstrapping with transfer learning". In: *arXiv preprint arXiv:1802.00500* (2018).

[20] Ruoxi Jia et al. "Towards efficient data valuation based on the shapley value". In: *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR. 2019, pp. 1167–1176.

[21] Nan Jiang and Alekh Agarwal. "Open problem: The dependence of sample complexity lower bounds on planning horizon". In: *Conference On Learning Theory*. PMLR. 2018, pp. 3395–3398.

[22] Nan Jiang et al. "Contextual decision processes with low Bellman rank are PAC-learnable". In: *International Conference on Machine Learning*. PMLR. 2017, pp. 1704–1713.

[23] Nate Kohl and Peter Stone. "Policy gradient reinforcement learning for fast quadrupedal locomotion". In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*. Vol. 3. IEEE. 2004, pp. 2619–2624.

[24] Matthew Lai. "Giraffe: Using deep reinforcement learning to play chess". In: *arXiv preprint arXiv:1509.01549* (2015).

[25] Jacky Liang et al. "GPU-accelerated robotic simulation for distributed reinforcement learning". In: *Conference on Robot Learning*. PMLR. 2018, pp. 270–282.

[26] Yuxuan Liang et al. "Geoman: Multi-level attention networks for geo-sensory time series prediction." In: *IJCAI*. Vol. 2018. 2018, pp. 3428–3434.

[27] Sandeep Manjanna, Herke van Hoof, and Gregory Dudek. "Reinforcement learning with non-uniform state representations for adaptive search". In: *2018 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE. 2018, pp. 1–7.

[28] Riccardo Marin et al. "Correspondence learning via linearly-invariant embedding". In: (2020).

[29] Volodymyr Mnih et al. "Human-level control through deep reinforcement learning". In: *nature* 518.7540 (2015), pp. 529–533.

[30] Volodymyr Mnih et al. "Playing atari with deep reinforcement learning". In: *arXiv preprint arXiv:1312.5602* (2013).

[31] Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh. "Agnostic federated learning". In: *International Conference on Machine Learning*. PMLR. 2019, pp. 4615–4625.

[32] Christoph Molnar, Giuseppe Casalicchio, and Bernd Bischl. "Interpretable machine learning–a brief history, state-of-the-art and challenges". In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer. 2020, pp. 417–431.

[33] Błażej Osiński et al. "Simulation-based reinforcement learning for real-world autonomous driving". In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 6411–6418.

[34] Charles W Ostrom. *Time series analysis: Regression techniques*. 9. Sage, 1990.

[35] Joelle Pineau et al. "Improving Reproducibility in Machine Learning Research". In: *Journal of Machine Learning Research* 22 (2021), pp. 1–20.

[36] Neoklis Polyzotis et al. "Data lifecycle challenges in production machine learning: a survey". In: *ACM SIGMOD Record* 47.2 (2018), pp. 17–28.

[37] Vikas Raunak, Vivek Gupta, and Florian Metze. "Effective dimensionality reduction for word embeddings". In: *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*. 2019, pp. 235–243.

[38] Roberta Rocca and Tal Yarkoni. "Putting psychology to the test: Rethinking model evaluation through benchmarking and prediction". In: (2020).

[39] Alaa Sagheer and Mostafa Kotb. "Time series forecasting of petroleum production using deep LSTM recurrent networks". In: *Neurocomputing* 323 (2019), pp. 203–213.

[40] David Silver et al. "Mastering the game of Go with deep neural networks and tree search". In: *nature* 529.7587 (2016), pp. 484–489.

[41] Satinder Singh, Andy Okun, and Andrew Jackson. "Learning to play Go from scratch". In: *Nature* 550.7676 (2017), pp. 336–337.

[42] Mukund Sundararajan and Amir Najmi. "The many Shapley values for model explanation". In: *International Conference on Machine Learning*. PMLR. 2020, pp. 9269–9278.

[43] Richard S Sutton. "Dyna, an integrated architecture for learning, planning, and reacting". In: *ACM Sigart Bulletin* 2.4 (1991), pp. 160–163.

[44] Sebastian Thrun and Anton Schwartz. "Issues in using function approximation for reinforcement learning". In: *Proceedings of the Fourth Connectionist Models Summer School*. Hillsdale, NJ. 1993, pp. 255–263.

[45] Hado Van Hasselt. "Reinforcement learning in continuous state and action spaces". In: *Reinforcement learning*. Springer, 2012, pp. 207–251.

[46] Hilde JP Weerts, Werner van Ipenburg, and Mykola Pechenizkiy. "A human-grounded evaluation of shap for alert processing". In: *arXiv preprint arXiv:1907.03324* (2019).

[47] Eyal Winter. "The shapley value". In: *Handbook of game theory with economic applications* 3 (2002), pp. 2025–2054.

[48] Yu Zhang, Jianguo Yao, and Haibing Guan. "Intelligent cloud resource management with deep reinforcement learning". In: *IEEE Cloud Computing* 4.6 (2017), pp. 60–69.

[49] Dongruo Zhou, Jiafan He, and Quanquan Gu. "Provably efficient reinforcement learning for discounted mdps with feature mapping". In: *International Conference on Machine Learning*. PMLR. 2021, pp. 12793–12802.

[50] Hao Zhu et al. "Deep reinforcement learning for mobile edge caching: Review, new features, and open issues". In: *IEEE Network* 32.6 (2018), pp. 50–57.

[51] Henry Zhu et al. "Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost". In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE. 2019, pp. 3651–3657.