

A Survey on Deep Reinforcement Learning for Data Processing and Analytics

Qingpeng Cai¹, Can Cui¹, Yiyuan Xiong¹, Wei Wang¹,
Zhongle Xie², and Meihui Zhang², *Member, IEEE*

Abstract—Data processing and analytics are fundamental and pervasive. Algorithms play a vital role in data processing and analytics where many algorithm designs have incorporated heuristics and general rules from human knowledge and experience to improve their effectiveness. Recently, reinforcement learning, deep reinforcement learning (DRL) in particular, is increasingly explored and exploited in many areas because it can learn better strategies in complicated environments it is interacting with than statically designed algorithms. Motivated by this trend, we provide a comprehensive review of recent works focusing on utilizing DRL to improve data processing and analytics. First, we present an introduction to key concepts, theories, and methods in DRL. Next, we discuss DRL deployment on database systems, facilitating data processing and analytics in various aspects, including data organization, scheduling, tuning, and indexing. Then, we survey the application of DRL in data processing and analytics, ranging from data preparation, natural language processing to healthcare, fintech, etc. Finally, we discuss important open challenges and future research directions of using DRL in data processing and analytics.

Index Terms—Deep reinforcement learning, data processing and analytics, database, system optimization

1 INTRODUCTION

IN the age of big data, data processing and analytics are fundamental, ubiquitous, and crucial to many organizations which undertake a digitalization journey to improve and transform their businesses and operations. Data analytics typically entails other key operations such as data acquisition, data cleansing, data integration, modeling, etc., before insights could be extracted. Big data can unleash significant value creation across many sectors such as healthcare and retail [1]. However, the complexity of data (e.g., high volume, high velocity, and high variety) presents many challenges in data analytics and hence renders the difficulty in drawing meaningful insights. To tackle the challenge and facilitate the data processing and analytics efficiently and effectively, a large number of algorithms and techniques have been designed and numerous learning systems have

also been developed by researchers and practitioners such as Spark MLlib [2], and Rafiki [3].

To support fast data processing and accurate data analytics, a huge number of algorithms rely on rules that are developed based on human knowledge and experience. For example, shortest-job-first is a scheduling algorithm that chooses the job with the smallest execution time for the next execution. However, without fully exploiting characteristics of the workload, it can achieve inferior performance compared to a learning-based scheduling algorithm [4]. Another example is packet classification in computer networking which matches a packet to a rule from a set of rules. One solution is to construct the decision tree using hand-tuned heuristics for classification. Specifically, the heuristics are designed for a particular set of rules and thus may not work well for other workloads with different characteristics [5]. We observe three limitations of existing algorithms [6], [7]. First, the algorithms are suboptimal. Useful information such as data distribution could be overlooked or not fully exploited by the rules. Second, the algorithm lacks adaptivity. Algorithms designed for a specific workload cannot perform well in another different workload. Third, the algorithm design is a time-consuming process. Developers have to spend much time trying a lot of rules to find one that empirically works.

Learning-based algorithms have also been studied for data processing and analytics. Two types of learning methods are often used: supervised learning and reinforcement learning. They achieve better performance by direct optimization of the performance objective. Supervised learning typically requires a rich set of high-quality labeled training data, which could be hard and challenging to acquire. For example, configuration tuning is important to optimize the overall performance of a database management system (DBMS) [8]. There could be hundreds of tuning knobs that are correlated

- Qingpeng Cai, Can Cui, Yiyuan Xiong, and Wei Wang are with the National University of Singapore, Singapore 119077. E-mail: {qingpeng, cuican, yiyuan, wangwei}@comp.nus.edu.sg.
- Zhongle Xie is with the Zhejiang University and Institute of Computing Innovation, Zhejiang University, Hangzhou, Zhejiang 310027, China. E-mail: xiezl@zju.edu.cn.
- Meihui Zhang is with the Beijing Institute of Technology, Beijing 100811, China. E-mail: meihui_zhang@bit.edu.cn.

Manuscript received 5 Aug. 2021; revised 4 Feb. 2022; accepted 13 Feb. 2022. Date of publication 1 Mar. 2022; date of current version 3 Apr. 2023.

This work was supported by National Key Research and Development Program of China under Grant 2020YFB1708100. The work of Qingpeng Cai, Can Cui and Yiyuan Xiong's was supported by the Singapore Ministry of Education Academic Research Fund Tier 3 under MOE's official Grant MOE2017-T3-1-007.

(Corresponding author: Meihui Zhang.)

Recommended for acceptance by L. Chen.

Digital Object Identifier no. 10.1109/TKDE.2022.3155196

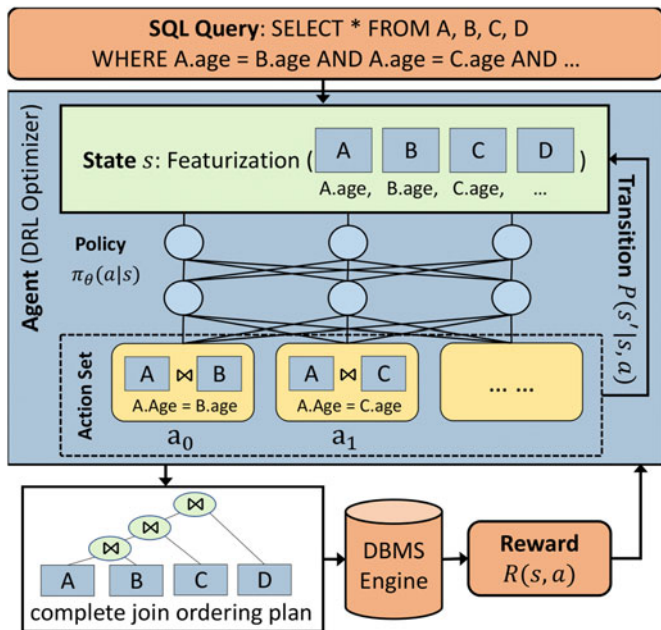


Fig. 1. The Workflow of DRL for Query Optimization. A, B, C and D are four tables.

in discrete and continuous space. Furthermore, diverse database instances, query workloads, hardware characteristics render data collection infeasible, especially in the cloud environment. Compared to supervised learning, reinforcement learning shows good performance because it adopts a trial-and-error search and requires fewer training samples to find good configuration for cloud databases [9]. Another specific example would be query optimization in query processing. Database system optimizers are tasked to find the best execution plan for a query to reduce its query cost. Traditional optimizers typically enumerate many candidate plans and use a cost model to find the plan with minimal cost. The optimization process could be slow and inaccurate [10]. Without relying on an inaccurate cost model, deep reinforcement learning (DRL) methods improve the execution plan (e.g., changing the table join orders) by interacting with the database [11], [12]. Fig. 1 provides a typical workflow for query optimization using DRL. When the query is sent to the agent (i.e., DRL optimizer), it produces a state vector via conducting featurization on essential information, such as the accessed relations and tables. Taking the state as the input, the agent employs neural networks to produce the probability distribution of an action set, where the action set could contain all possible join operations as potential actions. Each action denotes a partial join plan on a pair of tables, and the state will be updated once an action is taken. After taking possible actions, a complete plan is generated, which is then executed by a DBMS to get the reward. In this query optimization problem, the reward can be calculated by the real latency. During the training process with the reward signal, the agent can improve the policy and produce a better join ordering with a higher reward (i.e., less latency).

Reinforcement learning (RL) [13] focuses on learning to make intelligent actions in an environment. The RL algorithm works on the basis of exploration and exploitation to improve itself with feedback from the environment. In the past decades, RL has achieved tremendous improvements

in both theoretical and technical aspects [13], [14]. Notably, DRL incorporates deep learning (DL) techniques to handle complex unstructured data and has been designed to learn from historical data and self-exploration to solve notoriously hard and large-scale problems (e.g., AlphaGo [15]). In recent years, researchers from different communities have proposed DRL solutions to address issues in data processing and analytics [4], [16], [17]. We categorize existing works using DRL from two perspectives: system and application. From the system's perspective, we focus on fundamental research topics ranging from general ones, such as scheduling, to system-specific ones, such as query optimization in databases. We shall also emphasize how it is formulated in the Markov Decision Process and discuss how the problem can be solved by DRL more effectively compared to traditional methods. Many techniques such as sampling and simulation are adopted to improve DRL training efficiency because workload execution and data collection in the real system could be time-consuming [18]. From the application's perspective, we shall cover various key applications in both data processing and data analytics to provide a comprehensive understanding of the DRL's usability and adaptivity. Many domains are transformed by the adoption of DRL, which helps to learn domain-specific knowledge about the applications.

In this survey, we aim at providing a broad and systematic review of recent advancements in employing DRL in solving data systems, data processing and analytics issues. In Section 2, we introduce the key concepts, theories, and techniques in RL to lay the foundations. To gain a deeper understanding of DRL, readers could refer to the recently published book [19], which covers selected DRL research topics and applications with detailed illustrations. In Section 3, we review the latest important research works on using DRL for system optimization to support data processing and analytics. We cover fundamental topics such as data organization, scheduling, system tuning, index, query optimization, and cache management. In Section 4, we discuss using DRL for applications in data processing and analytics ranging from data preparation, natural language interaction to various real-world applications such as healthcare, fintech, E-commerce, etc. In Section 5, we highlight various open challenges and potential research problems. We conclude in Section 6. This survey focuses on recent advancements in exploring RL for data processing and analytics that spurs great interest, especially in the database and data mining community. There are survey papers discussing DRL for other domains. We refer readers to the survey of DRL for healthcare in [20], communications and networking in [21], and RL explainability in [22]. Another work [23] discusses how deep learning can be used to optimize database system design, and vice versa. In this paper, we use "DRL" and "RL" interchangeably.

2 THEORETICAL FOUNDATION AND ALGORITHMS OF REINFORCEMENT LEARNING

RL is targeted to solve the sequential decision making problem and the goal is to take actions with maximum expected rewards. In detail, the agent follows a policy to make a series of decisions (i.e., taking actions) in different states of

the environment, and the sequence of the states and the actions form a trajectory. To estimate whether the policy is good or not, each decision under the policy will be evaluated by the accumulated rewards through the trajectory. After evaluating the policy from the trajectories, the agent next improves the policy by increasing the probabilities of making decisions with greater expected rewards. By repeating these steps, the agent can improve the policy through trial-and-error until the policy reaches the optimal, and such a sequential decision-making process is modeled via Markov Decision Process (MDP).

2.1 Markov Decision Process

Mathematically, MDP, shown in Fig. 1, is a stochastic control process \mathcal{M} defined by a tuple with 5 elements, $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma)$, which are explained as follows.

- State \mathcal{S} : \mathcal{S} is the space for states that denote different situations in the environment and $s_t \in \mathcal{S}$ denotes the state of the situation at the time t .
- Action \mathcal{A} : \mathcal{A} is the space for actions that the agent can take; the actions can either be discrete or continuous, and $a_t \in \mathcal{A}$ denotes the action taken at the time t .
- Reward function $\mathcal{R}(s_t, a_t)$: It denotes the immediate reward of the action a_t taken under the state s_t .
- Transition function $\mathcal{P}(s_{t+1} = s' | s_t = s, a_t = a)$: It denotes the probability of transition to the state s' at the time $t + 1$ given the current state s and the taken action a at the time t .
- Discount factor $\gamma \in [0, 1]$: The total rewards of a certain action consist of both immediate rewards and future rewards, and the γ quantifies how much importance we give for future rewards.

We take the query optimization problem demonstrated in Fig. 1 to help explain the five components of the MDP. In this example, the state is expressed as a state vector, which summarizes the information of relations and tables that are assessed by the query q . In each state, the RL agent produces a probability distribution over all potential actions where each action denotes a partial join plan on a pair of tables. After repeating these two processes, it reaches a terminal state where the final join ordering is generated for an agent to execute, and all actions' target rewards are measured by the actual performance (i.e., latency) or a cost model. As for the transition function, the transitions of the states are always deterministic in both this problem and most of the other DB problems.

In RL, we aim to train the agent with a good policy π that is a mapping function from state to action. Through the policy, the agent can take a series of actions that will result in continuous changes in the states, and the sequence of the states and the actions following the policy π form a trajectory $\tau = (s_0, a_0, s_1, a_1, \dots)$. From each τ , we can evaluate the effect of each action by the accumulated rewards \mathcal{G} , and it consists of the immediate reward of this action and the discounted rewards of its following actions in the trajectory. The total result \mathcal{G} for the action a_t is as follows: $\mathcal{G}(\tau) = \sum_{t=0}^{\infty} \gamma^t r_t$, where γ quantifies how much importance we give for future rewards. With a bigger γ , the RL agent will be more likely to take any action that may have a less immediate reward at the current time but has a greater future reward in expectation.

RL continuously evaluates the policy π and improves it until it reaches the optimal policy $\pi^* = \arg \max_{(\tau \sim \pi)} \mathcal{G}(\tau)$ where the agent always takes actions that maximize the expected return. To evaluate the policy π , RL algorithms estimate how good or bad it is for a state and a state-action pair by the function \mathcal{V} and function \mathcal{Q} respectively. Both of these two value functions are calculated according to the discounted return \mathcal{G} in expectation which can be written as

$$\mathcal{V}^\pi(s) = E_{\tau \sim \pi}[\mathcal{G}(\tau) | s_0 = s] \quad (1)$$

$$\mathcal{Q}^\pi(s, a) = E_{\tau \sim \pi}[\mathcal{G}(\tau) | s_0 = s, a_0 = a] \quad (2)$$

These two value functions have a close association where the $\mathcal{V}^\pi(s_t)$ is the expectation of the function \mathcal{Q} of all possible actions under the state s_t according to the policy π , and the $\mathcal{Q}^\pi(s_t, a_t)$ is the combination of the immediate reward of the action a_t and the expectation of all possible states' values after taking the action a_t under the state s_t . Hence, we have

$$\mathcal{V}^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) \mathcal{Q}^\pi(s, a) \quad (3)$$

$$\mathcal{Q}^\pi(s, a) = R(s, a) + \gamma \sum_{s' \in \mathcal{S}} \mathcal{P}(s'|s, a) \mathcal{V}^\pi(s') \quad (4)$$

Given a policy π , we can evaluate its value functions by Bellman equations [13] which utilize the recursive relationships of these value functions. Formally, Bellman equations deduce the relationships between a given state (i.e., function \mathcal{V}) or a given state-action pair (i.e., function \mathcal{Q}) and its successors which can be written as

$$\mathcal{V}^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) [R(s, a) + \gamma \sum_{s' \in \mathcal{S}} \mathcal{P}(s'|s, a) \mathcal{V}^\pi(s')] \quad (5)$$

$$\mathcal{Q}^\pi(s, a) = \sum_{s' \in \mathcal{S}} \mathcal{P}(s'|s, a) [R(s, a) + \gamma \sum_{a' \in \mathcal{A}} \pi(a'|s') \mathcal{Q}^\pi(s', a')] \quad (6)$$

By iterating the Bellman equations, we can easily obtain the value functions for a policy, and to compare policies, we define that the policy π is better than π' if the function \mathcal{V} according to the π is no less than the function \mathcal{V} according to the π' for all states, that is $\mathcal{V}^\pi(s) \geq \mathcal{V}^{\pi'}(s), \forall s$. It has been proven in [13] that the existence of the optimal policy π^* is guaranteed in the MDP problem, where $\mathcal{V}^*(s) = \max_{\pi} \mathcal{V}^\pi(s)$ and $\mathcal{Q}^*(s) = \max_{\pi} \mathcal{Q}^\pi(s)$. These two functions are defined as the optimal function \mathcal{V} and the optimal function \mathcal{Q} . We can obtain the optimal policy π^* by maximizing over the $\mathcal{Q}^*(\pi)$ which can be written as:

$$\pi^*(a|s) = \arg \max \mathcal{Q}^*(s, a) \quad (7)$$

To improve the policy, we apply the Bellman optimality equations [13] to update value functions by taking the action with maximum value instead of trying all possible actions. To facilitate the optimization of the policy, many RL techniques are proposed from different perspectives, and Fig. 2 provides a diagram outlining the broad categorization of these techniques, illustrating how these techniques can be applied.

2.2 Basic Techniques

Based on the representation of MDP elements, basic techniques can be categorized into two classes: *model-based* method

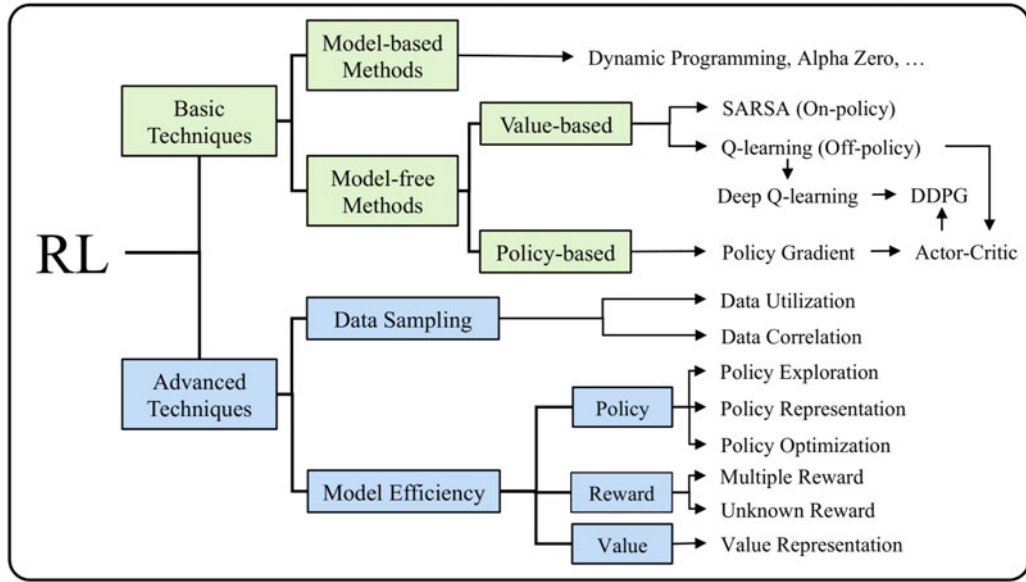


Fig. 2. Broad categorization of RL techniques.

and *model-free* method. The main difference is whether the agent has access to model the environment, i.e., whether the agent knows the transition function and the reward function. These two functions are already known in the model-based method where *Dynamic Programming* (DP) [24] and *Alpha-Zero* [14] are the classical methods which have achieved significant results in numerous applications. In these methods, agents are allowed to think ahead and plan future actions with known effects on the environment. Besides, an agent can learn the optimal policy from the planned experience which results in high sample efficiency.

In many RL problems, the reward and the transition function are typically unknown due to the complicated environment and its intricate inherent mechanism. For example, as illustrated in Fig. 1, we are unable to obtain the actual latency as the reward in the joint query optimization example. Besides, in the stochastic job scheduling problem [25], it is also impossible to directly model the transition function because of the randomness of the job arrivals in the practical scenarios. Hence, in these problems, agents usually employ model-free methods that can purely learn the policy from the experience gained during the interaction with the environment. Model-free methods can mainly be classified into two categories, namely the *value-based* method and the *policy-based* method. In the value-based method, the RL algorithm learns the optimal policy by maximizing the value functions. There are two main approaches in estimating the value functions that are *Mento-Carlo* (MC) methods and *Temporal difference* (TD) methods. MC methods calculate the $\mathcal{V}(s)$ by directly applying its definition, that is Equation (1). MC methods can directly update the value functions once they get a new trajectory τ as follows:

$$\mathcal{V}^\pi(s) \leftarrow \mathcal{V}^\pi(s) + \alpha(\mathcal{G}_{\tau \sim \pi}(\tau|s_0 = s) - \mathcal{V}^\pi(s)) \quad (8)$$

where $\alpha \in [0, 1)$ denotes the learning rate which controls the rate of updating the policy with new experiences. However, it has an obvious drawback that a complete trajectory requires the agent to reach a terminal state, while it is not

practical in some applications, such as online systems. Different from MC methods, the TD method builds on the recursive relationship of value functions, and hence, can learn from the incomplete trajectory. Mathematically, the update of TD methods can be written as

$$\mathcal{V}^\pi(s) \leftarrow \mathcal{V}^\pi(s) + \alpha(R(s, a) + \gamma\mathcal{V}^\pi(s') - \mathcal{V}^\pi(s)) \quad (9)$$

However, there is bias when estimating the function \mathcal{V} with TD methods because they learn from the recursive relationship. To reduce the bias, TD methods can extend the length of the incomplete trajectories and update the function \mathcal{V} by thinking more steps ahead, which is called *n-steps* TD methods. As n grows to the length of whole trajectories, MC methods can be regarded as a special case of TD methods where function \mathcal{V} is an unbiased estimate. On the other side of the coin, as the length n increases, the variance of the trajectory also increases. In addition to the above consideration, TD-based methods are more efficient and require less storage and computation, thus they are more popular among RL algorithms.

In value-based methods, we can obtain the optimal policy by acting greedily via Equation (7). The update of the function Q with TD methods is similar to the update of the function \mathcal{V} , and is as follows: $Q^\pi(s, a) \leftarrow Q^\pi(s, a) + \alpha(R(s, a) + \gamma Q^\pi(s', a') - Q^\pi(s, a))$ where the agent follows the policy π to take actions and follows the policy π' to maximize the function Q . If the two policies are the same, that is $\pi' = \pi$, we call such RL algorithms the *on-policy* methods where the *SARSA* [26] is the representative method. In addition, other policies can also be used in π' . For example, in *Q-learning* [27], the agent applies the greedy policy and updates the function Q with the maximum value in its successor. Its update formula can be written as: $Q^\pi(s, a) \leftarrow Q^\pi(s, a) + \alpha(R(s, a) + \gamma \max_{a'} Q^\pi(s', a') - Q^\pi(s, a))$. Both value-based methods can work well without the model of the environment, and Q-learning directly learns the optimal policy, whilst SARSA learns a near-optimal policy during exploring. Theoretically, Q-learning should converge quicker

than SARSA, but its generated samples have a high variance which may suffer from the problems of converging.

In RL, storage and computation costs are very high when there is a huge number of states or actions. To overcome this problem, DRL, as a branch of RL, adopts Deep Neural Network (DNN) to replace tabular representations with neural networks. For function \mathcal{V} , DNN takes the state s as input and outputs its state value $\mathcal{V}_\theta(s) \approx \mathcal{V}^\pi(s)$ where the θ denotes the parameter in the DNN. When comes to function \mathcal{Q} , It takes the combination of the state s and the action a as input and outputs the value of the state-action pair $\mathcal{Q}_\theta(s, a) \approx \mathcal{Q}^\pi(s, a)$. As for the neural networks, we can optimize them by applying the techniques that are widely used in deep learning (e.g., gradient descent). *Deep Q-learning network* (DQN) [28], as a representative method in DRL, combines the DNN with Q-learning and its loss function is as follows:

$$\mathcal{L}_w = \mathbf{E}_{\mathcal{D}}[(R(s, a) + \gamma \max_{a' \in \mathcal{A}} \mathcal{Q}_w(s', a') - \mathcal{Q}_w(s, a))^2] \quad (10)$$

where \mathcal{D} denotes the *experience replay* which accumulates the generated samples and can stabilize the training process.

Policy-based methods are another branch of the model-free RL algorithm that have a clear representation of the policy $\pi(a|s)$, and they can tackle several challenges that are encountered in value-based methods. For example, when the action space is continuous, value-based methods need to discretize the action which could increase the dimensionality of the problem, and memory and computation consumption. Value-based methods learn a deterministic policy that generates the action given a state through an optimal function \mathcal{Q} (i.e., $\pi(s) = a$). However, for policy-based methods, they can learn a stochastic policy (i.e., $\pi_\theta(a_i|s) = p_i, \sum_i p_i = 1$) as the optimal policy, where p_i denotes the probability of taking the action a_i given a state s , and θ denotes the parameters where neural networks can be used to approximate the policy. *Policy Gradient* [29] method is one of the main policy-based methods which can tackle the aforementioned challenges. Its goal is to optimize the parameters θ by using the gradient ascent method, and the target can be denoted in a generalized expression

$$\nabla_\theta J(\theta) = \mathbf{E}_{\tau \sim \pi_\theta} [R(\tau) \nabla_{\pi_\theta} \log \pi_\theta(a|s)] \quad (11)$$

The specific proof process can refer to [13]. Sampling via the MC methods, we will get the entire trajectories to improve the policy for the policy-based methods.

After training, the action with higher rewards in expectation will have a higher probability to be chosen and vice versa. As for the continuous action, The optimal policy learned from the Policy Gradient is stochastic which still needs to be sampled to get the action. However, the stochastic policy still requires lots of samples to train the model when the search space is huge. *Deterministic Policy Gradient* (DPG) [30], as an extension of the Policy Gradient, overcomes this problem by using a stochastic policy to perform sampling while applying deterministic policy to output the action which demands relatively fewer samples.

Both value-based methods and policy-based methods have their strengths and weaknesses, but they are not contradictory to each other. *Actor-Critic* (AC) method, as the integration of both methods, divides the model into two parts:

actor and critic. The actor part selects the action based on the parameterized policy and the critic part concentrates on evaluating the value functions. Different from previous approaches, AC evaluates the advantage function $\mathcal{A}^\pi(s, a) = \mathcal{Q}^\pi(s, a) - \mathcal{V}^\pi(s)$ which reflects the relative advantage of a certain action a to the average value of all actions. The introduction of the value functions also allows AC to update by step through the TD method, and the incorporation of the policy-based methods makes AC be suitable for continuous actions. However, the combination of the two methods also makes the AC method more difficult to converge. Moreover, *Deep Deterministic Policy Gradient* (DDPG) [31], as an extension of the AC, absorbs the advanced techniques from the DQN and the DPG which enables DDPG to learn the policy more efficiently.

In all the above-mentioned methods, there always exists a trade-off between exploring the unknown situation and exploiting with learned knowledge. On the one hand, exploiting the learned knowledge can help the model converge quicker, but it always leads the model into a local optimal rather than a globally optimal. On the other hand, exploring unknown situations can find some new and better solutions, but always being in the exploring process causes the model hard to converge. To balance these two processes, researchers have been devoting much energy to finding a good heuristics strategy, such as $\epsilon - greedy$ strategy, Boltzmann exploration (Softmax exploration), upper confidence bound (UCB) algorithm [32], Thompson sampling [33], and so on. Here, we consider the $\epsilon - greedy$, a widely used exploration strategy, as an example. $\epsilon - greedy$ typically selects the action with the maximal Q value to exploit the learned experience while occasionally selecting an action evenly at random to explore unknown cases. $\epsilon - greedy$ exploration strategy with m actions can be denoted as follow:

$$\pi(a|s) = \begin{cases} \epsilon/m + (1 - \epsilon) & a^* = \operatorname{argmax}_{a \in \mathcal{A}} \mathcal{Q}(s, a), \\ \epsilon/m & a \neq a^*. \end{cases} \quad (12)$$

$\epsilon \in [0, 1)$ is an exploration factor. The agent is more likely to select the action at random when the ϵ is closer to 1, and the ϵ will be continuously reduced during the training process.

2.3 Advanced Techniques

This section mainly discusses some advanced techniques in RL which focus on efficiently using the limited samples and building sophisticated model structures for better representation and optimization. According to the different improvements, they can be broadly classified into two parts: data sampling and model efficiency.

2.3.1 Data Sampling

Data sampling is one of the most important concerns in training the DRL in data processing and analytics. In most applications, the sample generation process costs a great amount of time and computation resources. For example, a sample may refer to an execution run for workload and repartitioning for the database, which can take about 40 minutes [18]. Hence, to train the model with limited samples, we need to increase data utilization and reduce data correlation.

Data Utilization. Most DRL algorithms train the optimal policy and sample data at the same time. Instead of dropping samples after being trained, *experience replay* [34] accumulates the samples in a big table where samples are randomly selected during the learning phase. With this mechanism, samples will have a higher utilization rate and a lower variance, and hence, it can stabilize the training process and accelerate the training convergence. Samples after several iterations may differ from the current policy, and hence, *Growing-batch* [35] can continuously refresh the table and replace these outdated samples. In addition, samples that are far away from the current policy should be paid more attention and *Prioritized Experience Replay* [36] uses TD error as the priority to measure the sample importance, and hence, focus more on learning the samples with high errors. In a nutshell, with the experience replay, DRL cannot only stable the learning phase but also efficiently optimize the policy with fewer samples.

Data Correlation. Strong correlation of training data is another concern that may lead the agent to learn a sub-optimal solution instead of the globally optimal one. Apart from the experience replay, the mechanism of the distributed environments is another research direction to alleviate this problem. For example, the *asynchronous advantage actor-critic* (A3C) [37] and *Distributed PPO* (DPPO) [38] apply multi-threads to build multiple individual environments where multiple agents take actions in parallel, and the update is calculated periodically and separately which can accelerate the sampling process and reduce the data correlation.

2.3.2 Model Efficiency

RL model with better efficiency is the major driving force of the development of RL, and there are many researchers improving it from three major aspects, namely policy, reward function, and value function.

Policy. The policy-related techniques focus on stably and effectively learning a comprehensive policy, and the advanced techniques to efficiently learn the policy can be classified into three parts in detail, which are policy exploration, policy representation, and policy optimization.

a) *Policy exploration:* Its target is to explore as many actions as possible during the training process in case the policy will be trapped into the local optimal. For example, *entropy regularisation* [37] adds the entropy of the actions' probabilities into the loss item which can sufficiently explore the actions. Besides, adding noise to the action is another research direction to increase the randomness into policy exploration. The DDPG applies an Ornstein–Uhlenbeck process [39] to generate temporal noise \mathcal{N} which are directly injected into policy. *Noisy-Net* [40] incorporates the noise into the parameters of neural networks which is easy to implement, and it shows a better performance than the ϵ -greedy and entropy regularisation methods. Further, Plappert *et al.* [41] investigate an effective way to combine the parameter space noise to enrich the exploratory behaviors which can benefit both on-policy methods and off-policy methods.

b) *Policy representation:* The states in some RL problems are in a huge dimension which causes challenges during the training. To approximate a better policy, a branch of DRL models improve the policy representation by absorbing

convolutional neural networks (CNN) into DQN to analyze the data, such as *Dueling DQN* [42], *DRQN* [43], and so on. In addition, DRQN also incorporates the LSTM structure to increase the capacity of the policy which is able to capture the temporal information, such as speed, direction.

c) *Policy optimization:* The update of the value functions following the Equation (5) and 6 tends to overestimate the value functions and introduce a bias because they learn estimates from the estimates. Mnih *et al.* [44] separate the two estimation process by using two same Q-networks which can reduce the correlation of two estimation processes and hence, stabilize the course of training. However, the action with the maximum Q-value may differ between two Q-networks which will be hard for convergence. and *Double DQN* (DDQN) [45] alleviate the issue by disaggregating the step of selecting the action and calculating the max Q-value.

When we apply the policy-based RL methods, the learning rate of the policy plays an essential role in achieving superior performance. A higher learning rate can always maximize the improvement on a policy by step, but it also causes the instability of the learning phase. Hence, The *Trust Region Policy Optimization* (TRPO) [46] builds constraints on the old policy and new policy via KL divergence to control the change of the policy in an acceptable range. With this constraint, TRPO can iteratively optimize the policy via a surrogate objective function which can monotonically improve policies. However, the design of the KL constraint makes it hard to be trained, and *Proximal Policy Optimization* (PPO) [47] simplifies the constraint through two ways: adding it into the objective function, designing a clipping function to control the update rate. Empirically, PPO methods are much simpler to implement and are able to perform at least as well as TRPO.

Reward. Reward function as one of the key components in the MDP plays an essential role in the RL. In some specific problems, the agent has to achieve multiple goals which may have some relationships. For example, the robot can only get out through the door only if it has already found the key. To tackle this challenge, *Hierarchical DQN* [48] proposes two levels of hierarchical RL (HRL) models to repeatedly select a new goal and achieve the chosen goal. However, there is a limitation that the goal needs to be manually predefined which may be unknown or unmeasurable in some environments, such as the market and the effect of a drug. To overcome it, *Inverse RL* (IRL) [49] learns the rewards function from the given experts' demonstrations (i.e., the handcraft trajectories), but the agent in IRL can only prioritize the entire trajectories over others. It will cause a shift when the agent comes to a state that never appears before, and *Generative Adversarial Imitation Learning* (GAIL) [50], as an imitation learning algorithm, applies adversarial training methods to generate fake samples and is able to learn the expert's policy explicitly and directly.

Value. As we have mentioned earlier, the tabular representation of the value functions has several limitations which can be alleviated via DRL. Different from directly taking the state-action pair as the input to calculate the Q-function, *Dueling DQN* [42] estimates its value by approximating two separate parts that are the state-values and the advantage values, and hence, can distinguish whether the value is brought by the state or the action.

TABLE 1
Representative Works Using DRL for Data System Optimizations

Domain	Work	Algorithm	D(State)	D(Action)	DRL-based Approach	Open Source
Data organization	Analytical system data partition[16]	PPO	10 - 100	100 - 1000	Exploit workload patterns and Generate the tree	NO
	Database horizontal partition [18]	DQN	100	10	Navigate the partition search efficiently	NO
	UDF-centric workload data partition [54]	A3C	10	1-10	Exploit the features of partition and search	YES
	Time series data compression [55]	PG	100	10	Search parameters interactively	NO
Scheduling	Distributed job processing [4]	PG	100	10	Exploit the job dependencies and learn schedule decision	YES
	Distributed stream data [56]	DDPG	100	10-100	Learn schedule decision	NO
Tuning	Database configuration [9] [8]	DDPG	100	10	Search configuration parameters interactively	YES
Index	Index Selection [57]	CEM	100	10	Search the index interactively	NO
	R-tree construction [58]	DQN	10-100	10	Learn to generate the tree	NO
Query Optimization	Join order selection [11], [12], [59], [60]	PG, DQN, ...	10-100	1-10	Learn to decide the join order	Only [60]
Cache Management	View Materialization [61]	DQN	100	10	Model the problem as IIP and solve	NO

$D(X)$ denotes the approximate dimension of X space.

The aforementioned advanced algorithms and techniques improve and enhance the DRL from different perspectives, which makes DRL-based algorithms be a promising way to improve data processing and analytics. We observe that problems with the following characteristics may be amenable to DRL-based optimization. First, problems are incredibly complex and difficult. The system and application involve a complicated operational environment (e.g., large-scale, high-dimensional states) and internal implementation mechanisms, which is hard to construct a white-box model accurately. DRL can process complex data and learn from experience generated from interacting, which is naturally suitable for data processing and analytics where many kinds of data exist and are processed frequently. Second, the optimization objectives can be represented and calculated easily as the reward because the RL agent improves itself towards maximizing the rewards and rewards could be computed a lot of times during training. Third, the environment can be well described as MDP. DRL has been shown to solve MDP with theoretical guarantees and empirical results. Thus, problems involving sequential decision making such as planning, scheduling, structure generation (e.g., tree, graph), and searching could be expressed as MDP and a good fit for DRL. Fourth, collecting required labels of data massively is hard. Compared to supervised learning, DRL can utilize data efficiently to gain good performance.

3 DATA SYSTEM OPTIMIZATIONS

DRL learns knowledge about the system by interacting with it and optimizes the system. In this section, we focus on several fundamental aspects with regards to system optimization in data processing and analytics including data organization, scheduling, tuning, indexing, query optimization, and cache management. We discuss how each problem is formulated in MDP by defining three key elements (action,

state, and reward) in the system and solved by DRL. Generally, the states are defined by some key characteristics of the system. The actions are possible decisions (e.g., system configuration), that affect the system performance and the reward is calculated based on the performance metrics (e.g., throughput, latency). Table 1 presents a summary of representative works and the estimated dimension ranges on the state and action space of each work are added as signals on the DRL training difficulty. As a comparison, OpenAI Five [51], a Dota-playing AI, observes the state as 20,000 numbers representing useful game information and about 1,000 valid actions (like ordering a hero to move to a location) for per hero. Dota is a real-time strategy game between two teams of five players where each player controls a character called a “hero”.

3.1 Data Organization

3.1.1 Data Partitioning

Effective data partitioning strategy is essential to accelerate data processing and analytics by skipping irrelevant data for a given query. It is challenging as many factors need to be considered, including the workload and data characteristics, hardware profiles, and system implementation.

In data analytics systems, data is split into blocks in main memory or secondary storage, which are accessed by relevant queries. A query may fetch many blocks redundantly and, therefore, an effective block layout avoids reading unnecessary data and reduces the number of block accesses, thereby improving the system performance. Yang *et al.* [16] propose a framework called the qd-tree that partitions data into blocks using DRL over the analytical workload. The qd-tree resembles the classic k-d tree and describes the partition of multi-dimensional data space where each internal node splits data using a particular predicate and represents a subspace. The data in the leaf node is assigned to the same block. In the MDP, each state is a node representing the

subspace of the whole data and featured as the concatenation of range and category predicates. After the agent takes an action to generate two child nodes, two new states will be produced and explored later. The available action set is the predicates parsed from workload queries. The reward is computed by the normalized number of skipped blocks over all queries. They do not execute queries and a sampling technique is used to estimate the reward efficiently. The formulation of using DRL to learn a tree is similar to NeuroCuts [5] that learns a tree for packet classification. However, the qd-tree may not support a complex workload containing user-defined functions (UDFs) queries.

Horizontal partitioning in the database chooses attributes of large tables and splits them across multiple machines to improve the performance of analytical workloads. The design relies on either the experience of database administrators (DBAs) or cost models that are often inaccurate [10] to predict the runtime for different partitions. Data collection is too challenging and costly to train the accurate supervised learning model in the cloud environment. Hilprecht *et al.* [18] learn to partition using DRL on analytical workloads in cloud databases, on the fact that DRL is able to efficiently navigate the partition search and requires less training data. In the MDP, the state consists of two parts. The database part encodes whether a table is replicated, an attribute is used for partitioning, and which tables are co-partitioned. The workload part incorporates normalized frequencies of representative queries. Supported actions are: partitioning a table using an attribute, replicating a table, and changing tables co-partition. The reward is the negative of the runtime for the workload. One challenge is that the cost of database partitioning is high during training. To alleviate the problem, the agent is trained in the simulation environment and is further refined in the real environment by estimating the rewards using sampling. One limitation is that it may not support new queries well because only the frequency features of queries are considered. Durand *et al.* in [52], [53] utilize DRL to improve vertical partitioning that optimizes the physical table layout. They show that the DQN algorithm can easily work for a single workload with one table but is hard to generalize to random workloads.

For UDFs analytics workloads on unstructured data, partitioning is more challenging where UDFs could express complex computations and functional dependency is unavailable in the unstructured data. Zou *et al.* [54] propose the Lachesis system to provide automatic partitioning for non-relational data analytics. Lachesis translates UDFs to graph-based intermediate representations (IR) and identifies partition candidates based on the subgraph of IR as a two-terminal graph. Lachesis adopts DRL to learn to choose the optimal candidate. The state incorporates features for each partition extracted from historical workflows: frequency, the execution interval, time of the most recent run, complexity, selectivity, key distribution, number, and size of co-partition. In addition, the state also incorporates other features such as hardware configurations. The action is to select one partition candidate. The reward is the throughput speedup compared to the average throughput of the historical executions of applications. To reduce the training time, the reward is derived from historical latency statistics without partitioning the data when running the applications. One limitation is

that Lachesis largely depends on historical statistics to design the state and calculate the reward, which could lead to poor performance when the statistics are inadequate.

3.1.2 Data Compression

Data compression is widely employed to save storage space. The effectiveness of a compression scheme however relies on the data types and patterns. In time-series data, the pattern can change over time and a fixed compression scheme may not work well for the entire duration. Yu *et al.* [55] propose a two-level compression framework, where a scheme space is constructed by extracting global features at the top level and a compression schema is selected for each point at the bottom level. The proposed AMMMO framework incorporates compression primitives and the control parameters, which define the compression scheme space. Due to the fact that the enumeration is computationally infeasible, the framework proposes to adopt DRL to find the compression scheme. The agent takes a block that consists of 32 data points with the compressed header and data segment, timestamps, and metrics value as the state. The action is to select a scheme from compression scheme space and then the compression ratio is computed as the reward. The limitation is that the method may not work for other data types like images and videos.

3.2 Scheduling

Scheduling is a critical component in data processing and analytics systems to ensure that resources are well utilized. Job scheduling in a distributed computing cluster faces many challenging factors such as workload (e.g., job dependencies, sizes, priority), data locality, and hardware characteristics. Existing algorithms using general heuristics such as shortest-job-first do not utilize these factors well and fail to yield top performance. To this end, Mao *et al.* [4] propose Decima to learn to schedule jobs with dependent stages using DRL for data processing clusters and improve the job completion time. In the data processing systems such as Hive [62], Pig [63], Spark-SQL [64], jobs could have up to hundreds of stages and many stages run in parallel, which are represented as directed acyclic graphs (DAGs) where the nodes are the execution stages and each edge represents the dependency. To handle parallelism and dependencies in job DAGs, Decima first applies graph neural network (GNN) to extract features as the state instead of manually designing them while achieving scalability. Three types of feature embeddings are generated. Node embedding captures information about the node and its children including the number of remaining tasks, busy and available executors, duration, and locality of executors. Job embedding aggregates all node embeddings in the job and cluster embedding combines job embeddings. To balance possible large action space and long action sequences, The action determines the job stage to be scheduled next and the parallelism limit of executors. The reward is based on the average job completion time. To train effectively in a job streaming environment, Decima gradually increases the length of training jobs to conduct curriculum learning [65]. The variance reduction technique [25] is applied to handle stochastic job arrivals for robustness. However, we note that Decima is

non-preemptive and does not re-schedule for higher priority jobs.

In distributed stream data processing, streams of continuous data are processed at scale in a real-time manner. The scheduling algorithm assigns workers to process data where each worker uses many threads to process data tuples and aims to minimize average data tuple processing time. Li *et al.* [56] design a scheduling algorithm using DRL for distributed stream data processing, which learns to assign tuples to work threads. The state consists of the scheduling plan (e.g., the current assignment of workers) and the workload information (e.g., tuple arrival rate). The action is to assign threads to machines. The reward is the negative tuple processing time on average. The work shows that DQN does not work well because the action space is large and applies DDPG to train the *actor-critic* based agent instead. To find a good action, the proposed method looks for k nearest neighbors of the action that the *actor* network outputs and selects the neighbor with the highest value that the *critic* network outputs. The algorithm is implemented on Apache Storm and evaluated with representative applications: log stream processing, continuous queries, and word count.

Many works have been recently proposed to improve scheduling using DRL [66], [67]. Query scheduling determines the execution order of queries, which has a great influence on query performance and resource utilization in the database system. SmartQueue [66] improves query scheduling by leveraging overlapping data access among queries and learns to improve cache hits using DRL. In addition, Tim *et al.* [67] design a scheduling system in SageDB using RL techniques. Other works using RL for scheduling include Bayesian RL for scheduling in heterogeneous clusters [68], operation scheduling in devices [69], application container scheduling in clusters [70], etc.

3.3 Tuning

Tuning the configuration of data processing and analytic systems plays a key role to improve system performance. The task is challenging because up to hundreds of parameters and complex relations between them could exist. Furthermore, other factors such as hardware and workload also impact the performance. Existing works often employ search-based or supervised learning methods. The former takes much time to get an acceptable configuration and the latter such as OtterTune [71] needs large high-quality data that is non-trivial to obtain in practice. Zhang *et al.* [9] design a cloud database tuning system CDBTune using DRL to find the best parameter in high-dimensional configuration space. The CDBTune formulates MDP as follows. The state is represented by the internal metrics (e.g., buffer size, pages read). The action is to increase or decrease the knob values. The reward is the performance difference between two states, which is calculated using throughput and latency. CDBTune takes several hours on offline training in simulation and online training in the real environment. Compared to OtterTune, CDBTune eases the burden of collecting large training data sets. In the experiments, CDBTune is shown to outperform DBA experts and OtterTune and improve tuning efficiency under 6 different workloads on four databases. One limitation of the approach is that the workload information is ignored and thus it may not perform well when the query workload is changed.

To address the issue, Li *et al.* [8] propose QTune that considers query information to tune the database using DRL. First, Qtune extracts features from SQL query including types (e.g., insert, delete), tables, and operation (e.g., scan, hash join) costs estimated by the database engine. The columns attributes and operations like selection conditions in the query are ignored. Subsequently, Qtune trains a DNN model to predict the difference of statistics (e.g., updated tuples, the number of committed transactions) in the state after executing the queries in the workload and updates the state using it. The action and reward design are similar to CDBTune. Additionally, QTune supports three levels of tuning granularity for balancing throughput and latency. For query-level, QTune inputs query vector and tries to find good knobs for each query. For workload-level, vectors for all queries are merged and used. For cluster-level, QTune employs a clustering method based on deep learning to classify queries and merge queries into clusters. One drawback of QTune is that the query featurization could lose key information such as query attributes (i.e., columns) and hurt the performance especially when the cost estimation is inaccurate. The prediction model for state changes is trained alone and needs accurate training data. An end-to-end training framework is therefore essential and a good direction to undertake.

3.4 Indexing

3.4.1 Database Index Selection

Database index selection considers which attributes to create an index to maximize query performance. Sharma *et al.* [57] show how DRL can be used to recommend an index based on a given workload. The state encodes selectivity values for workload queries and columns in the database schema and current column indexes. The action is to create an index on a column. The reward is the improvement compared to the baseline without indexes. The experiments show that the approach can perform as well or better as having indexes on all columns. Sadri *et al.* [72] utilize DRL to select the index for a cluster database where both query processing and load balancing are considered. Welborn *et al.* [73] optimize the action space design by introducing task-specific knowledge for index selection tasks in the database. However, these works only consider the situation where single-column indexes are built. Lan *et al.* [74] propose both single-attribute and multi-attribute indexes selection using DRL. Five rules are proposed to reduce the action and state space, which help the agent learn effective strategy easier. The method uses what-if caller [75] to get the cost of queries under specific index configurations without building indexes physically. These works conduct basic experiments with small and simple datasets. Extensive and large-scale experiments using real datasets are therefore needed to benchmark these methods to ensure that they can scale well.

3.4.2 Index Structure Construction

The learned index is proposed recently as an alternative index to replace the B⁺-Tree and bloom filter by viewing indexes as models and using deep learning models to act as indexes [76]. DRL can enhance the traditional indexes instead of replacing them.

Hierarchical structures such as the B⁺-tree and R-tree are important indexing mechanisms to locate data of interest efficiently without scanning a large portion of the database. Compared to the single dimensional counterpart, the R-tree is more complex to optimize due to bounding box efficiency and multi-path traversals. Earlier conventional approaches use heuristics to determine these two operations (i.e., choosing the insertion subtree and splitting an overflowing node) during the construction of the R-tree. Gu *et al.* [58] propose to use DRL to replace heuristics to construct the R-tree and propose the RLR-tree. The approach models two operations ChooseSubtree and Split as two MDPs respectively and combines them to generate an R-Tree. For ChooseSubtree, the state is represented as the concatenation of the four features (i.e., area, perimeter, overlap, occupancy rate) of each selected child node. More features are evaluated but do not improve the performance in the reported experiments. The action is to select a node to insert from top-k child nodes in terms of the increase of area. The reward is the performance improvement from the RLR-tree. For Split MDP, the state is the areas and perimeters of the two nodes created by all top-k splits in the ascending order of total area. The action is to choose one split rule from k rules and the reward is similar to that of ChooseSubtree. The two agents are trained alternately. As expected, the optimizations render the RLR-tree improved performance in range and KNN queries.

Graphs can be used as effective indexes to accelerate nearest neighbors search [77], [78]. Existing graph construction methods generally propose different rules to generate graphs, which cannot provide adaptivity for different workloads [79]. Baranchuk *et al.* [80] employ DRL to optimize the graph for nearest neighbors search. The approach learns the probabilities of edges in the graph and tries to maximize the search efficiency. It considers the initial graph and the search algorithm as the state. The action is to keep an edge or not. The reward is the performance for search. It chooses the TRPO [46] algorithm to train. The reported experimental results show that the agent can refine state-of-the-art graphs and achieve better performance. However, this approach does not learn to explore and add new edges to the initial graph that may affect the performance.

Searching and constructing a new index structure is another line of interesting research [81]. Inspired by Neural Architecture Search (NAS) [82], Wu *et al.* [83] propose an RNN-based neural index search (NIS) framework that employs DRL to search the index structures and parameters given the workload. NIS can generate tree-like index structures layer by layer via formalizing abstract ordered blocks and unordered blocks, which can provide a well-designed search space. The keys in the ordered block are sorted in ascending order, and the skip list or B⁺-Tree can be used. The keys in the unordered block are partitioned using customized functions and the hash bucket can be used. Overall, the whole learning process is similar to that of NAS.

3.5 Query Optimization

Query optimization aims to find the most efficient way to execute queries in database management systems. There are many different plans to access the query data that can have a large processing time variance from seconds to hours. The performance of a query plan is determined mostly by the

TABLE 2
Methods of Query Optimization

Method	Techniques	Training	Workload Adaptivity
Rejoin[11], DQ[12]	learn from execution experience	High	Low
SkinnerDB [84]	learn from current execution status	Medium	Medium
Bao[87]	learn to choose existing optimizers	Low	High

table join orders. Traditionally, query optimizers use certain heuristics combined with dynamic programming to enumerate possible efficient execution plans and evaluate them using cost models that could produce large errors [10]. Marcus *et al.* [11] propose Rejoin that applies DRL to learn to select better join orders utilizing past experience. The state encodes join tree structure and join predicates. The action is to combine two subtrees, where each subtree represents an input relation to join. The reward is assigned based on the cost model in the optimizer. The experiments show that ReJOIN can match or outperform the optimizer in PostgreSQL. Compared to ReJoin, DQ [12] presents an extensible featurization scheme for state representation and improves the training efficiency using the DQN [28] algorithm. Heitz *et al.* [60] compare different RL algorithms including DQN [28], DDQN [31], and PPO [47] for join order optimization and use a symmetric matrix to represent the state instead of vector. Yu *et al.* [59] introduce a graph neural network (GNN) with DRL for join order selection that replaces fixed-length hand-tuned vector in Rejoin [11] and DQ [12] with learned scalable GNN representation and better captures and distinguishes the join tree structure information. These works mainly differ in encoding what information and how to encode them.

Instead of learning from past query executions, Trummer *et al.* [84] propose SkinnerDB to learn from the current query execution status to optimize the remaining execution of a query using RL. Specifically, SkinnerDB breaks the query execution into many small time intervals (e.g., tens to thousands of slices per second) and processes the query adaptively. At the beginning of each time interval, the RL agent chooses the join order and measures the execution progress. SkinnerDB adopts a similar adaptive query processing strategy in Eddies [85] and uses the UCT algorithm [86], which provides formal guarantees that the difference is bounded between the rewards obtained by the agent and those by optimal choices. The reward is calculated by the progress for the current interval. A tailored execution engine is designed to fully exploit the learning strategy with tuple representations and specialized multi-way join algorithms. SkinnerDB offers several advantages. First, it is inherently robust to query distribution changes because its execution only depends on the current query. Second, it relies on less assumption and information (e.g., cardinality models) than traditional optimizers and thus is more suitable for the complicated environment where cardinality is hard to estimate. Third, it predicts the optimal join order based on real performance. However, it may introduce overhead caused by join order switching.

Learning-based methods that have been proposed to replace traditional query optimizers often incur a great deal of training overhead because they have to learn from scratch. To mitigate the problem, Bao [87] (the Bandit optimizer) is designed to take advantage of the existing query optimizers. Specifically, Bao learns to choose the best plan from the query plan candidates provided by available optimizers by passing different flags or hints to them. Bao transforms query plan trees into vectors and adopts a tree convolutional neural network to identify patterns in the tree. Then it formulates the choosing task as a contextual multi-armed bandit problem and uses Thompson sampling [33] to solve it. Bao is a hybrid solution for query optimization. It achieves good training time and is robust to changes in workload [87].

3.6 Cache Management

3.6.1 View Materialization

View materialization is the process of deciding which view, i.e., results of query or subquery, to cache. In database systems, a view is represented as a table and other queries could be accelerated by reading this table instead of accessing the original tables. There is an overhead of materializing and maintaining the view when the original table is updated. Existing methods are based on heuristics, which either rely on simple Least-Recently-Used rule or cost-model based approaches [88]. The performance of these approaches is limited because feedback from the historical performance of view materialization is not incorporated. Liang *et al.* [89] implement Deep Q-Materialization (DQM) system that leverages DRL to improve the view materialization process in the OLAP system. First, DQM analyzes SQL queries to find candidate views for the current query. Second, it trains a DRL agent to select from the set of candidates. Third, it uses an eviction policy to delete the materialized views. In the MDP, the state encodes view state and workload information. The action is to create the view or do nothing. The reward is calculated by the query time improvement minus amortized creation cost. Additionally, the eviction policy is based on credit and it evicts the materialized view with the lowest score.

Yuan *et al.* [61] present a different way that use DRL to automate view generation and select the most beneficial subqueries to materialize. First, the approach uses a DNN to estimate the benefits of a materialized view where features from tables, queries, and view plans are extracted. Then the approach models selection as an Integer Linear Programming (ILP) problem and introduce an iterative optimization method to figure it out. However, the method cannot guarantee convergence. To address the issue, the problem is formulated as the MDP. The state encodes the subqueries that are selected to materialize and status if queries use these materialized views. The action is to choose the subquery to materialize or not. The reward is the difference between benefit changes of two states. Both cost estimation and view selection models are trained offline using the actual cost of queries and benefits. Then the cost estimation model is used for the online recommendation for view materialization. Performance study shows its good performance; However, it lacks a comparison with DQM.

3.6.2 Storage

Cache management impacts the performance of computer systems with hierarchical hardware structures directly. Generally, a caching policy considers which objects to cache, to evict when the cache is full to maximize the object hit rate in the cache. In many systems, the optimal caching policy depends on workload characteristics. Phoebe [90] is the RL-based framework for cache management for storage models. The state encodes the information from a preceding fixed-length sequence of accesses where for each access, nine features are extracted including data block address, data block address delta, frequency, reuse distance, penultimate reuse distance, average reuse distance, frequency in the sliding window, the number of cache misses, and a priority value. The action is to set a priority value ranging within $[-1, 1]$ to the data. The reward is computed from if the cache is hit or missed and values are 1 and -1 respectively. It applies the DDPG algorithm to train the agent. Periodical training is employed to amortize training costs in online training. In network systems, one issue is that the reward delay is very long in systems with a large cache, i.e., CDN cache can host up to millions of objects. Wang *et al.* [91] propose a subsampling technique by hashing the objects to mitigate the issue when applying RL on caching systems.

4 DATA ANALYTICS APPLICATIONS

In this section, we shall discuss DRL applications from the perspective of data processing and data analytics. These two categories of DRL applications form indispensable parts of a pipeline, in which data processing provides a better basis for data analytics. In addition, these two categories share some overlapping topics, making these topics mutually motivating and stimulating. We have summarized the technical comparisons of different applications in Table 3. We shall first discuss DRL applications in data preparation and then in data analytics.

4.1 Data Preparation

4.1.1 Entity Matching

Entity matching is a data cleaning task that aligns different mentions of the same entity in the context. Clark *et al.* [92] identify the issue that the heuristic loss function cannot effectively optimize the evaluation metric B^3 , and propose using reinforcement learning to directly optimize the metric. The problem is formulated as a sequential decision problem where each action is performed on one mention of a document. The action maps the mention to an entity in the database at each step by a mention ranking model. Then the reward is calculated using the evaluation metric B^3 . This work originally proposes scaling each action's weight by measuring its impact on the final reward since each action is independent. However, this work does not consider the global relations between entities. Fang *et al.* [93] propose a reinforcement learning framework based on the fact that an easier entity will create a better context for the subsequent entity matching. Specifically, both local and global representations of entity mentions are modeled and a learned policy network is devised to choose from the next action (i.e., which entity to recognize). However, the selection of the easier entity to learn

TABLE 3
Representative Works for RL Applications

Domain	Work	Algorithm	D(State)	D(Action)	DRL-based Approach
Data processing application	Entity matching [93], [94]	PG	100 - 1000	100 - 1000	Select target entity from the candidate entities
	Database interaction with natural language [95], [96]	PG	100 - 1000	100 - 1000	Learn to generate the query
	Feature engineering [17]	DQN	100	1-10	Select features and model feature correlations in states
	Exploratory data analysis [97]	A3C	10-100	100000	Learn to query a dataset for key characteristics
	Abnormal detection [98]	IRL	1-10	1-10	Learn the reward function for normal sequences
	AutoML pipeline generation [99]	DQN	10	100	Learn to select modules of a pipeline
Healthcare	Treatment recommendation [100]	DDPG	10	100-1000	Select treatment from candidate treatments
	Diagnostic inference [101]	DQN	100-1000	1-10	Learn diagnostic decision
	Hospital resource allocation [102]	DDPG	100	1000-10000	Learn resource scheduling
Fintech	Portfolio optimization [103]	Q-Learning	100	100	Select the portfolio weights for stocks
	Trading [104], [105]	IRL	1-10	10	Learn the reward function of trading behaviors
	Fraud detection [105]	IRL	100	10-100	Learn the reward function of trading behaviors
E-Commerce	Online advertising [106]	DQN	1-10	1-10	Learn to schedule the advertisements
	Online recommendation [107]	DQN	100	10000	Learn to schedule recommendations
	Search results aggregation [108]	DQN	10-100	10-100	Learn to schedule search results
Others	User profiling [109]	DQN	100-1000	1000-10000	Select users' next activities by modeling spatial semantics
	Spammer detection [110]	PG	100	100	Search for the detector by interacting with spammers
	Transportation [111]	PG	1000-10000	1000	Learn to schedule transportation

$D(X)$ denotes the approximate dimension of X space.

the context could be less powerful than context modeling with more recent techniques in NLP such as the transformer.

4.1.2 Database Interaction With Natural Language

To facilitate query formulation for relational databases, there have been efforts in generating SQL queries from various other means that do not require knowledge of SQL and schema. Zhong *et al.* [94] propose to generate SQL from a natural language using RL. For queries formed by a natural language, the model Seq2SQL will learn a policy transforming the queries into SQL queries. The transformed queries will then be executed in the database system to get results. The results will be compared with the ground truth to generate RL rewards. Earlier work [95] using a generic autoencoder model for semantic parsing with Softmax as the final layer may generate unnecessarily large output spaces for SQL query generation tasks. Thus the structure of SQL is used to prune the output space of query generating and policy-based reinforcement learning to optimize the part which cannot be optimized by cross-entropy. However, RL is observed to have limited performance enhancement by [111] due to unnecessary modeling of query serialization.

Efficiently querying a database of documents is a promising data processing application. Karthik *et al.* [112] propose collecting evidence from external sources of documents to boost extraction accuracy to original sources where data might be scarce. The problem is formulated as an MDP problem, where each step the agent needs to decide if current extracted articles are accepted and stop querying, or these articles are rejected and more relevant articles are queried. Both data reconciliation (from original sources) and

data retrieval (from external sources) are represented as states. Extraction accuracy and penalties for extra retrieval actions are reflected in the reward function.

4.1.3 Feature Engineering

Feature engineering can be formulated as a single-agent reinforcement learning problem to search for an optimal subset of features in a large space: the agent selects one feature at each action step. The state is the current feature subset. A reward is assigned to the agent based on the predictive performance of the current features subset. Liu *et al.* [17] propose a method to reformulate feature engineering as a multi-agent reinforcement learning problem. The multi-agent RL formulation reduces the large action space of a single agent since now each of the agents has a smaller action space for one feature selection. However, this formulation also brings challenges: interactions between agents, representation of the environment, and selection of samples. Three technical methods in [17] have been proposed to tackle them respectively: adding inter-feature information to reward formulation, using meta statistics, and deep learning methods to learn the representation of the environment, and Gaussian mixture to independently determine samples. However, although this formulation reduces the action space, the trade-off is using more computing resources to support more agents' learning. Also, the method is difficult to scale to a large feature space.

4.1.4 Exploratory Data Analysis

Exploratory data analysis (EDA) is useful for users to understand the characteristics of a new dataset. In [96], the

problem is formulated as an MDP. The action space is the combination of a finite set of operators and their corresponding parameters to query a dataset. The result of a query shows the characteristics of the dataset. The characteristics are modeled as the state, which is represented by descriptive statistics and recent operators. The reward signal measures the interestingness, diversity, and coherency of the characteristics by an episode of EDA operations. DRL is applied to the non-differential signals and discrete states in MDP. However, challenges arise when applying deep reinforcement learning given a large number of possible actions as parameterized operations (i.e., for each type of operation, the corresponding possible action is the Cartesian product of all parameters' possible values). In [96], a two-fold layer architecture is proposed to replace a global softmax layer into two local layers, which effectively reduces the intractable large numbers of actions. However, the global interactions of operations and attributes are not considered.

4.1.5 Abnormal Detection

Abnormal detection is important for high-stake applications such as healthcare (e.g., predicting patients' status) and fintech (e.g., financial crime). Based on the assumptions, there are two approaches to this problem. One approach models the dynamics in the unlabeled datasets as a sequential decision process where the agent performs an action on each observation. Oh *et al.* [97] propose to use IRL to learn a reward function and a Bayesian network to estimate a confidence score for a potential abnormal observation. To achieve this, the prior distribution of the reward function is assumed. Then a reward function is sampled from the distribution to determine the sample generating policy, which generates sample background trajectories. As explained by the reward part of Section 2.3.2, experts' trajectories are observed. With these experts' trajectories and sample background trajectories, the parameters of the reward function are updated and thus the policy is improved. The sequence of actions is the input into the neural network. This network is trained to learn the normal pattern of a targeted agent and to predict if the next observation is abnormal or not. However, this approach relies too much on mining unlabeled datasets and ignores the labeled dataset. To address this issue, another approach also uses DRL but focus on the Exploit-Explore trade-off on both unlabeled and labeled dataset. Pang *et al.* [113] propose a DRL model with a sampling function to select data instances from both the unlabeled and labeled dataset. This sampling function helps the DRL model to exploit the scarce but useful labeled anomaly data instances and to explore the large unlabeled dataset for novel anomaly data instances. Thus, more anomaly data instances are selected to train the DRL model with better model capacity.

4.1.6 AutoML Pipeline Generation

Pipeline generation includes generating all data processing and analytics steps or modules to perform ML tasks. Heffetz *et al.* [98] propose a grid-world to represent all possible families of each step of a data pipeline as cells and connect all possible cells as a graph. Subsequently, a hierarchical method is used to reduce the space of all actions and represent all

actions by layers of clusters. Finally, the state representations are inputs to the value sub-network in a DQN network, and action representations are inputs to evaluate the advantage-to-average sub-network.

4.2 Healthcare

Healthcare analytics has gained increasing attention in tandem with the advancement of healthcare treatment and availability of medical data and computational capacity [114]. Naturally, a great amount of effort has been spent on applying DRL to healthcare. As before, implementing DRL-based models in healthcare requires the understanding of the application context and defining the key elements of MDP. However, differences occur in the approaches to learning better decisions: learning the motivation of expert decisions by IRL, learning better decisions without an expert by interacting with an environment or interacting with an environment with expert decisions as supervising signals.

4.2.1 Treatment Recommendation

Treatment recommendation systems are designed to assist doctors to make better decisions based on electronic health records. However, the doctors' prescriptions are not ground truth but valuable suggestions for high stake medical cases. The ground truth is the delayed condition of the patients. Thus model predictions must not deviate from the doctors' judgments too much, and not use those judgments as true labels. To tackle this challenge, Wang *et al.* [99] propose an architecture to combine supervised learning and reinforcement learning. This model reduces the inconsistency between indicator signals learned from doctor's prescriptions via supervised learning and evaluation signals learned from the long-term outcome of patients via reinforcement learning. In the formulated MDP, the domain expert makes a decision based on an unknown policy. The goal is to learn a policy that simultaneously reduces the difference between the chosen action of the agent and the expert's decision and to maximize the weighted sum of discounted rewards.

4.2.2 Diagnostic Inference

Using DRL to perform diagnosis can provide a second opinion in high-intensity diagnosis from historical medical records to reduce diagnostic errors. Ling *et al.* [100] propose modeling the integration of external evidence to capture diagnostic concept as a MDP. The objective is to find the optimal policy function. The inputs are case narratives and the outputs are improved concepts and inferred diagnoses. The states are a set of measures over the similarity of current concepts and externally extracted concepts. The actions are whether to accept (part of) the extracted concepts from external evidence. The environments are the top extracted case narratives from Wikipedia as the document pool for concepts extraction and a knowledge base for evaluating the intermediate results for current best concepts. The rewards are evaluated based on an external knowledge base mapping from the concepts to the diagnoses. The whole process is modeled by DQN. At each step, narrative cases and evidence are extracted, which provide the initial concepts and external concepts. The state representing the agent's confidence in the learned concept is duly calculated. Then the state is sent to

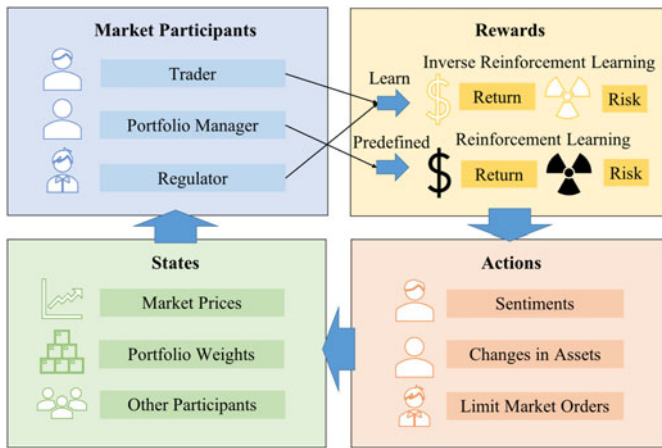


Fig. 3. DRL in fintech applications.

the DQN agent to estimate the reward to model the long-run accuracy of the learned concept by the agent. Iteratively, the model converges with better concepts and diagnoses.

4.2.3 Hospital Resource Allocation

Allocating limited hospital resources is the key to providing timely treatment for patients. In [101], the problem is formulated as a classification problem where the patients' features are given and the target is to predict the location of admissions. The RL framework uses a student network to solve the classification problem. The weights of the student network are used as states, which are fed into a teacher network to generate actions to select which batch of data to train the student network. The accuracy of the classification is used as the reward. This method provides a view on the resource allocation problem from a curriculum learning perspective. However, the temporal information of the data samples is not considered but it could affect resource allocation since some hours during a day could have fewer patients than the others.

4.3 Fintech

Reinforcement learning has wide applications in the finance domain. First, reinforcement learning has brought new perspectives to let the finance research community revisit many classic financial research topics. For example, traditional financial research topics such as option pricing that are typically solved by the classic Black–Scholes model can be steered through with a data-driven insight by reinforcement learning [115]. Second, portfolio optimization, typically formulated as a stochastic optimal control problem, can be addressed by reinforcement learning. Finally, the agents are financial market participants with different intentions. Reward functions can be learned to model these intentions, and hence, make better decisions as illustrated in Fig. 3. We refer readers with further interest in finance to [116].

4.3.1 Dynamic Portfolio Optimization

The portfolio optimization problem is challenging because of the high scale of the dimensionality and the high noise-to-signal ratio nature of stock price data. The latter problem of noisy observation can cause uncertainty in a learned policy. Therefore, [102] proposes a novel model structure based

on the Q-learning to handle noisy data and to scale to high dimensionality. The quadratic form of reward function is shown to have a semi-analytic solution that is computationally efficient. In the problem formulation, the agent's actions are represented as the changes in the assets at each time step. The states are the concatenation of market signals and the agent's holding assets. This method enhances Q-learning by introducing an entropy term measuring the noise in the data. This term acts as a regularization term forcing the learned policy to be close to a reference policy that is modeled by a Gaussian distribution.

4.3.2 Algorithm Trading Strategy Identification

Identification of algorithm trading strategies from historical trades is important in fraud detection and maintaining a healthy financial environment. [104] proposes using IRL to learn the reward function behind the trading behaviors. The problem is formulated as an Inverse Markov Decision Process (IMDP). The states are the differences between the volumes of bid orders and ask orders, which are discretized into three intervals based on the values of the volumes. The actions are the limit and market order discretized into 10 intervals each by their values. The prior distribution of the reward function is a Gaussian Process parameterized by θ . Given θ , the approximation of the posterior distribution of reward is performed by maximum a posteriori (MAP). This step would give a MAP estimated value of the reward. θ is optimized by a log-likelihood function on the posterior of observations. The optimization process can be proved to be convex which guarantees the global minimum. The learned features are then used to identify and classify trading strategies in the financial markets.

4.3.3 Sentiment-Based Trading

One of the main predictors in stock trading is sentiment, which drives the demand of bid orders and asks orders. Sentiment scores are often represented by unstructured text data such as news or twitters. [103] proposes treating the sentiment as the aggregated action of all the market participants, which has the advantage of simplifying the modeling of the numerous market participants. Specifically, the sentiment scores are categorized into three intervals: high, medium, and low as the action spaces. Compared to previous works, the proposed method can model the dependency between the sentiment and the market state by the policy function. This method is based on Gaussian Inverse Reinforcement Learning [117] similar to [104] as discussed at the beginning of Section 4.3, which is effective at dealing with uncertainty in the stock environment. This method provides a method for modeling market sentiments. However, as IRL faces the challenge of non-uniqueness of reward [19] of one agent's actions, the method does not address how aggregated actions of multiple market participants can infer a unique reward function.

4.4 E-Commerce

4.4.1 Online Advertising

With the increasing digitalization of businesses, sales and competition for market shares have moved online in tandem.

As a result, online advertising has been increasing in its presence and importance and exploiting RL in various aspects. One of the topics in online advertising, bidding optimization, can be formulated as a sequential decision problem: the advertiser is required to have strategic proposals with bidding keywords sequentially to maximize the overall profit. In [105], the issue of using static transitional probability to model dynamic environments is identified and a new DRL model is proposed to exploit the pattern discovered from dynamic environments.

Including but not limited to advertising, Feng *et al.* [118] propose to consider the whole picture of multiple ranking tasks that occurred in the sequence of user's queries. A new multi-agent reinforcement learning model is proposed to enable multiple agents to partially observe inputs and choose actions through their own actor networks. The agents communicate through a centralized critic model to optimize a shared objective. This allows different ranking algorithms to reconcile with each other when taking their own actions and consider the contextual information.

4.4.2 Online Recommendation

The problem of an unstabilized reward function arises because of the dynamic environment in the online recommendation. For example, user preference is modeled as the reward in DRL and it changes unexpectedly when a special discount happens for some products. In [106], a random stratified sampling method is proposed to calculate the optimal way of stratifying by allocating more samples to the strata with more weighted variance. Then the replay sampling is improved to consider key attributes of customers (e.g., gender, age, etc.), which are less volatile in the dynamic environment. This allows the modeling of reward function based on sampling from a pool with a longer horizon, thus reducing the bias in the estimation of the reward function. Lastly, the dynamic environment poses a challenge in setting an optimal policy used in regretting. A new method in [106] is proposed to train an offline model to calculate a real-time reward for a subset of customers to approximate a reference policy, that is used as an offset in the reward recalibration to stabilize the performance of the DRL algorithm.

4.4.3 Search Results Aggregation

Aggregating useful search results in online shopping search is important to improve the shopping experience. However, the challenge of aggregating heterogeneous data sources is often encountered. The heterogeneous data sources in online shopping are different product categories such as a shoe brand group or a particular topic group, each of which is a ranking system. A new model in [107] is proposed to decompose the task into two sub-tasks. The first one is to select a data source for the current page of search results based on historical users' clicks on previous pages. Learning to select the correct data source for each page is a sequential decision-making problem. The second sub-task is to fill the sequence of a page by selecting the best source from the candidate sources. However, the items from different sources cannot be directly compared because of their heterogeneous nature. The problem is solved by formulating the sub-task

as an RL task to let an agent fill up the sequence. However, one limitation of this method is that lacking full annotations of item relevance scores may constrain the model's performance on various scenarios [107].

4.5 Other Applications

DRL has been applied to various other applications. These DRL methods are often used with a knowledge graph, confounders, or game theory to model application-specific dynamics. These methods are not only well motivated from their respective applications but also general enough to be applied in other applications. However, these methods often fail to be evaluated by experiments in other applications.

The problem of mobile user profiling aims to identify user profiles to provide personalized services. In [108], the action is the selection of a place of visit. The environment is comprised of all users and a knowledge graph learning the semantic connections between the spatial entities. The knowledge graph is updated once a user's new activity is performed and then affects the agent's prediction. The state is the embedding of a user and the knowledge graph for the current time step. The reward is determined by several metrics measuring the similarity between the predicted spatial entities and the ground truth. This method considers the spatial semantics of entities but does not consider how the change of a user's key attributes (e.g., career) will affect activity prediction and policy learning, which could cause instability in policy updating.

In the transportation system, drivers often get recommendations and provide feedback in return to improve the service. However, the recommendation often fails when drivers make decisions in a complex environment. To address this issue, in [110] a new method is proposed to model hidden causal factors, called confounders, in a complex environment. Specifically, the framework in [50] is extended to include the confounders. First, all three elements (i.e., policy agent, environment, confounder) are treated as agents. The effect of a confounder is modeled as the policy of the hidden agent, which takes the observation and action of the policy agent as inputs and performs an action. The environment in turn takes the action based on inputs of the hidden agent's action and the policy agent's action and observation.

The problem of spammer detection aims to detect spam generating strategies. The challenge is that the detectors only detect easier spams while missing spams with strategies. In [109], the problem is formulated as two agents countering each other. One agent is the spammer, whose policy is to maintain a distribution of spam strategies and the action is to sample from the distribution. Another agent is the detector, whose state is the detection results after a spam attack and the action is to identify the spam. The rewards of two agents are measured by winning or losing revenue manipulation, respectively. The limitation of this method is that there is no guarantee for equilibrium.

5 OPEN CHALLENGES AND FUTURE DIRECTIONS

RL approaches provide strong alternatives to traditional heuristics or supervised learning-based algorithms. However, many challenges remain to be addressed.

5.1 Open Challenges for System Optimization

5.1.1 MDP Formulation and Lack of Justification

The design of MDP impacts the performance and efficiency of the RL algorithm greatly. The state should satisfy Markov property that its representation contains enough relevant information for the RL agent to make the optimal decision. It should summarize the environment compactly because a complicated state design will cause more training and inference costs. The action space should be designed carefully to balance learning performance and computational complexity. The reward definition directly affects the optimization direction and the system performance. Additionally, the process of reward calculation can involve costly data collection and computation in the data systems optimization. Currently, many works rely on experimental exploration and experience to formulate MDP while some works exploit domain knowledge to improve the MDP formulation by injecting task-specific knowledge into action space [73]. Generally, MDP can influence computational complexity, data required, and algorithm performance. Unfortunately, many works lack ablation studies of their MDP formulations and do not justify the design in a convincing manner. Therefore, automation of MDP formulation remains an open problem.

5.1.2 RL Algorithm and Technique Selection

RL algorithms and techniques have different tradeoffs and assumptions. Value-based DRL algorithms like DQN are not stable and guaranteed convergence. Policy-based DRL algorithms like TRPO and PPO are often not efficient. Model-based DRL algorithms do not guarantee that a better model can result in a better policy. Value-based methods assume full observability while policy-based ones assume episodic learning. Off-policy algorithms are usually more efficient than on-policy algorithms in terms of sample efficiency. One example is that DQ [12] uses off-policy deep Q-learning to increase data efficiency and reduce the number of training queries needed. Training efficiency can be a big concern for DRL-based system optimization, especially when the workload of the system could change dramatically and the model needs to be retrained frequently. Generally, RL algorithms and techniques selection affect the training efficiency and effectiveness greatly.

5.1.3 Integration With Existing Systems

Integrating RL-based methods into the real system more naturally and seamlessly faces many challenges. The RL agent has to be evolved when the system environment changes (e.g., workload) and the performance is degraded. We need to design new model management mechanisms to monitor, maintain, and upgrade the models. Furthermore, we find that the RL-based solutions can be lightweight or intrusive. The lightweight approach in which the RL agent is not designed as a component of the system, e.g., using RL to generate the qd-tree [16], is easier to integrate into the system because it does not change the architecture of the system dramatically. In contrast, the intrusive approach such as using RL models for join order optimization [11] is deeply embedded in the system and hence may need a redesign and optimization of the original system architecture to

support model inference efficiently. SageDB [67] proposes to learn various database system components by integrating RL and other ML techniques. Nevertheless, the proposed model-driven database system is yet to be fully implemented and benchmarked. It is likely that the data system architecture needs to be overhauled or significantly amended in order to graft data-driven RL solutions into the data system seamlessly to yield an overall performance gain.

5.1.4 Reproducibility and Benchmark

In the data system optimization problem, RL algorithms are not easy to be reproduced due to many factors such as lacking open source codes, workload, historic statistics used, and the unstable performance of RL algorithms. The landscape of problems in system optimization is vast and diverse. It could prevent fair comparison and optimization for future research works and deployments in practice. Lacking benchmarks is another challenge to evaluate these RL approaches. The benchmarks are therefore to provide standardized environments and evaluation metrics to conduct experiments with different RL approaches. There are some efforts to mitigate the issue. For example, Park [119] is an open platform for researchers to conduct experiments with RL. However, it only provides a basic interface and lacks system specifications. There is much room to improve with regards to the reproducibility and benchmark in order to promote the development and adoption of RL-based methods [120].

5.2 Open Challenges for Applications

5.2.1 Lack of Adaptability

There is a lack of adaptability for methods on a single component of a data pipeline to the whole. For example, many works focus on data cleaning tasks such as entity matching. However, little works have shown their efficiency in deploying their model in an end-to-end data pipeline. These works treat the tasks isolatedly from other tasks in the pipeline, thereby limiting the pipeline's performance. In healthcare, each method is applied in different steps of the whole treatment process, without being integrated and evaluated as one pipeline. One possible direction could be considering DRL as a module in the data pipeline optimization. However, data pipeline optimization has been focusing on models simpler than DRL to enable fast pipeline evaluation [121]. How to efficiently incorporate DRL into the data pipeline optimization remains a challenge.

5.2.2 Difficulty in Comparison With Different Applications

To date, most works with generalized contributions are only evaluated domain-specifically. Research questions are often formulated in their own platform as in E-Commerce. This presents difficulty in evaluating the methods for different environments. For example, the confounders modeling hidden causal factors in [110] can also contribute to DRL modeling in E-commerce. This is because modeling customers' interests are always subject to changing environments and a new environment may contain hidden causal factors. For example, consumers are more willing to buy relevant products for certain situations such as Covid-19. Thus a general

DRL method is yet to show the robustness and effectiveness under the environment of different applications.

5.2.3 Lack of Prediction in Multi-Modality

In healthcare and finance, multiple sources of data bring different perspectives. For example in healthcare, electronic health records, image scans, and medical tests can provide different features for accurate prediction. In addition, these sources of data with different sample frequencies provide contextual information for modeling a patient's visits to the hospital or symptom development. However, most innovations in healthcare focus on one particular source of data. How to integrate the contextual information with multi-modality effectively remains an unsolved difficult problem.

5.2.4 Injecting Domain Knowledge in Experience Replay

In high-stake applications such as healthcare and finance, injecting domain knowledge can make decision making in RL more robust and explainable. One possible way is to inject the knowledge of human beings' experience into an agent's experience pool as a prior distribution for the policy. For example, in dynamic portfolio optimization, a portfolio manager could have a large source of experience for risk management and profit optimization. Such experience could be useful for warming up the agent's exploration in the search space. Some works have shown positive effects of domain knowledge injection on selecting important experiences (i.e., transition samples) [36]. Notwithstanding, it remains a big challenge to inject useful and relevant knowledge from the experience into the agent's experience pool.

5.3 Future Research Directions

5.3.1 Data Structure Design

DRL provides an alternative way to find good data structures through feedback instead of designing them based on human knowledge and experience, e.g., decision tree [5] and the qd-tree [16]. These trees are optimized better because they are learned by interacting with the environment. DRL has also been effective in graph designs (e.g., molecular graph [122]). However, large-scale graph generation using DRL is difficult and daunting because it involves a huge search space. Generating other important structures using DRL remains to be explored. Idreos *et al.* [81] propose a *Data Alchemist* that learns to synthesize data structures by DRL and other techniques including Genetic Algorithms and Bayesian Optimization. In summary, DRL has a role in the design of more efficient data structures by interacting and learning from the environment. These indexes have to be adaptive to different data distributions and workloads.

5.3.2 Interpretability

The underlying logic behind the DRL agent is still unknown. In high-risk application areas such as healthcare, the adoption of DRL will be a big issue in the case that these approaches make wrong decisions and people do not know why it happens due to lack of interpretability. Many techniques have been proposed to mitigate the issue and provide interpretability [22]. However, they neglect domain knowledge from

related fields and applications and the explanations are not effective to human users. To instill confidence in the deployment of DRL-based systems in practice, interpretability is an important component and we should avoid treating DRL solutions as black boxes especially in critical applications.

5.3.3 Robustness by Causal Reasoning

Modeling real-world applications by DRL inevitably suffers from the problem of distribution changes. The real world has independent physical mechanisms that can be seen as different modules. For example, an image is subjected to the light of the environment. Given the modular property, a structural type of modeling focusing on factorizing the causal mechanisms can extract the invariant causal mechanisms and show robustness cross distribution changes [123]. One research direction towards DRL robust decision making is to perform sampling from past actions from a causal perspective. Given the invariance property of causal mechanisms, past actions can be reused by capturing the invariant mechanisms in a changing environment.

5.3.4 Extension to Other Domains

Beyond existing works, many classic problems in the data system and analytics could potentially be solved by DRL. For example, Polyjuice [124] learns the concurrency control algorithm for a given workload by defining fine-grained actions and states in the context of concurrency control. Though they use an evolutionary algorithm to learn and outperform a simple DRL baseline, we believe that there are huge potentials to further improve DRL for niche applications. Hence, we expect that more problems will be explored and solved with DRL in various domains in the near future.

5.3.5 Towards Intelligent and Autonomous Databases

Although DRL algorithms could provide breakthrough performance on many tasks than traditional methods, many issues need to be addressed towards intelligent and autonomous databases. First, database schema could be updated and DRL models trained on the previous snapshots may not work. DRL algorithms need to tackle generalization [125]. Second, it would be so costly and infeasible to train models from scratch for each scenario and setting. Transfer learning from existing models could be a potential way to ease the workload greatly. Third, we have to choose appropriate DRL algorithms automatically, in the same spirit as AutoML. Fourth, current DBMS systems were designed without considering much about the learning mechanism. A radically new DBMS design may be proposed based on the learning-centric architecture. To support intelligent and autonomous database systems, DRL models intelligent behaviors and may provide a solid basis for achieving artificial general intelligence based on reward maximization and trial-and-error experience [126].

6 CONCLUSION

In this survey, we present a comprehensive review of recent advances in utilizing DRL in data processing and analytics. The DRL agent could learn to understand and solve various tasks with the right incentives. First, we introduce basic

foundations and practical techniques in DRL. Next, we survey and review DRL for data processing and analytics from two perspectives, systems and applications. We cover a large number of topics ranging from fundamental problems in system areas such as tuning and scheduling to important applications such as healthcare and fintech. Finally, we discuss key challenges and future directions for applying DRL in data processing and analytics. We hope the survey would serve as a basis for research and development in this emerging area, and better integration of DRL techniques into data processing pipelines and stacks.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive comments and suggestions. Qingpeng Cai, Can Cui, and Yiyuan Xiong contribute equally to this work.

REFERENCES

- [1] J. Manyika et al., *Big Data: The Next Frontier for Innovation, Competition, and Productivity*. Washington, DC, USA: McKinsey Global Institute, 2011.
- [2] X. Meng et al., "MLlib: Machine learning in apache spark," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 1235–1241, 2016.
- [3] W. Wang et al., "Rafiki: Machine learning as an analytics service system," *Proc. VLDB Endowment*, vol. 12, no. 2, pp. 128–140, 2018.
- [4] H. Mao, M. Schwarzkopf, S. B. Venkatakrisnan, Z. Meng, and M. Alizadeh, "Learning scheduling algorithms for data processing clusters," in *Proc. ACM Special Int. Group Data Commun.*, 2019, pp. 270–288.
- [5] E. Liang, H. Zhu, X. Jin, and I. Stoica, "Neural packet classification," in *Proc. ACM Special Int. Group Data Commun.*, 2019, pp. 256–269.
- [6] B. Vamanan, G. Voskuilen, and T. Vijaykumar, "EffiCuts: Optimizing packet classification for memory and throughput," *Proc. ACM Special Int. Group Data Commun.*, vol. 40, no. 4, pp. 207–218, 2010.
- [7] W. Li, X. Li, H. Li, and G. Xie, "CutSplit: A decision-tree combining cutting and splitting for scalable packet classification," in *Proc. IEEE Conf. Comput. Commun.*, 2018, pp. 2645–2653.
- [8] G. Li, X. Zhou, S. Li, and B. Gao, "QTune: A query-aware database tuning system with deep reinforcement learning," *Proc. VLDB Endowment*, vol. 12, no. 12, pp. 2118–2130, 2019.
- [9] J. Zhang et al., "An end-to-end automatic cloud database tuning system using deep reinforcement learning," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2019, pp. 415–432.
- [10] V. Leis, A. Gubichev, A. Mirchev, P. Boncz, A. Kemper, and T. Neumann, "How good are query optimizers, really?," *Proc. VLDB Endowment*, vol. 9, no. 3, pp. 204–215, 2015.
- [11] R. Marcus and O. Papaemmanouil, "Deep reinforcement learning for join order enumeration," in *Proc. 1st Int. Workshop Exploiting Artif. Intell. Techn. Data Manage.*, 2018, pp. 1–4.
- [12] S. Krishnan, Z. Yang, K. Goldberg, J. Hellerstein, and I. Stoica, "Learning to optimize join queries with deep reinforcement learning," 2018, *arXiv:1808.03196*.
- [13] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 2018.
- [14] D. Silver et al., "A general reinforcement learning algorithm that masters chess, shogi, and go through self-play," *Science*, vol. 362, no. 6419, pp. 1140–1144, 2018.
- [15] D. Silver et al., "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [16] Z. Yang et al., "Qd-tree: Learning data layouts for big data analytics," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2020, pp. 193–208.
- [17] K. Liu, Y. Fu, P. Wang, L. Wu, R. Bo, and X. Li, "Automating feature subspace exploration via multi-agent reinforcement learning," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 207–215.
- [18] B. Hilprecht, C. Binnig, and U. Röhms, "Learning a partitioning advisor for cloud databases," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2020, pp. 143–157.
- [19] H. Dong, H. Dong, Z. Ding, S. Zhang, and Chang, *Deep Reinforcement Learning*. Berlin, Germany: Springer, 2020.
- [20] C. Yu, J. Liu, S. Nemati, and G. Yin, "Reinforcement learning in healthcare: A survey," *ACM Comput. Surveys*, vol. 55, no. 1, pp. 1–36, 2021.
- [21] N. C. Luong et al., "Applications of deep reinforcement learning in communications and networking: A survey," *IEEE Commun. Surv. Tut.*, vol. 21, no. 4, pp. 3133–3174, Oct.–Dec. 2019.
- [22] E. Puiutta and E. M. Veith, "Explainable reinforcement learning: A survey," in *Proc. Int. Cross-Domain Conf. Mach. Learn. Knowl. Extraction*, 2020, pp. 77–95.
- [23] W. Wang, M. Zhang, G. Chen, H. Jagadish, B. C. Ooi, and K.-L. Tan, "Database meets deep learning: Challenges and opportunities," *ACM SIGMOD Rec.*, vol. 45, no. 2, pp. 17–22, 2016.
- [24] R. Bellman, "Dynamic programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.
- [25] H. Mao, S. B. Venkatakrisnan, M. Schwarzkopf, and M. Alizadeh, "Variance reduction for reinforcement learning in input-driven environments," in *Proc. Int. Conf. Learn. Representations*, 2019, pp. 1–20.
- [26] G. A. Rummery and M. Niranjan, *On-Line Q-Learning Using Connectionist Systems*, vol. 37. Cambridge, U.K.: Cambridge Univ. Press, 1994.
- [27] C. J. Watkins and P. Dayan, "Q-learning," *Mach. Learn.*, vol. 8, no. 3/4, pp. 279–292, 1992.
- [28] V. Mnih et al., "Playing atari with deep reinforcement learning," 2013, *arXiv:1312.5602*.
- [29] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Proc. Int. Conf. Neural Informat. Process. Syst.*, 2000, pp. 1057–1063.
- [30] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic policy gradient algorithms," in *Proc. Int. Conf. Mach. Learn.*, 2014, pp. 387–395.
- [31] T. P. Lillicrap et al., "Continuous control with deep reinforcement learning," 2015, *arXiv:1509.02971*.
- [32] P. Auer, "Using confidence bounds for exploitation-exploration trade-offs," *J. Mach. Learn. Res.*, vol. 3, no. Nov, pp. 397–422, 2002.
- [33] W. R. Thompson, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples," *Biometrika*, vol. 25, pp. 285–294, 1933.
- [34] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Mach. Learn.*, vol. 8, no. 3/4, pp. 293–321, 1992.
- [35] S. Lange, T. Gabel, and M. Riedmiller, "Batch reinforcement learning," in *Reinforcement Learning*, Berlin, Germany: Springer, 2012, pp. 45–73.
- [36] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," 2015, *arXiv:1511.05952*.
- [37] V. Mnih et al., "Asynchronous methods for deep reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1928–1937.
- [38] N. Heess et al., "Emergence of locomotion behaviours in rich environments," 2017, *arXiv:1707.02286*.
- [39] G. E. Uhlenbeck and L. S. Ornstein, "On the theory of the brownian motion," *Phys. Rev.*, vol. 36, 1930, Art. no. 823.
- [40] M. Fortunato et al., "Noisy networks for exploration," 2017, *arXiv:1706.10295*.
- [41] M. Plappert et al., "Parameter space noise for exploration," 2017, *arXiv:1706.01905*.
- [42] Z. Wang, T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas, "Dueling network architectures for deep reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1995–2003.
- [43] M. Hausknecht and P. Stone, "Deep recurrent Q-learning for partially observable MDPs," 2015, *arXiv:1507.06527*.
- [44] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [45] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *Proc. Conf. Assoc. Advance. Artif. Intell.*, vol. 30, no. 1, pp. 2094–2100, 2016.
- [46] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1889–1897.
- [47] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.
- [48] T. D. Kulkarni, K. R. Narasimhan, A. Saeedi, and J. B. Tenenbaum, "Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation," 2016, *arXiv:1604.06057*.
- [49] A. Y. Ng, S. J. Russell et al., "Algorithms for inverse reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2000, Art. no. 2.

- [50] J. Ho and S. Ermon, "Generative adversarial imitation learning," 2016, *arXiv:1606.03476*.
- [51] C. Berner *et al.*, "Dota 2 with large scale deep reinforcement learning," 2019, *arXiv:1912.06680*.
- [52] G. C. Durand *et al.*, "Gridformation: Towards self-driven online data partitioning using reinforcement learning," in *Proc. 1st Int. Workshop Exploiting Artif. Intell. Techn. Data Manage.*, 2018, pp. 1–7.
- [53] G. C. Durand, R. Piriyeve, M. Pinnecke, D. Broneske, B. Gurumurthy, and G. Saake, "Automated vertical partitioning with deep reinforcement learning," in *Proc. Eur. Conf. Adv. Databases Informat. Syst.*, 2019, pp. 126–134.
- [54] J. Zou *et al.*, "Lachesis: Automated generation of persistent partitionings for big data applications," *Proc. VLDB Endowment*, vol. 14, 2021, pp. 1262–1275.
- [55] X. Yu, Y. Peng, F. Li, S. Wang, X. Shen, H. Mai, and Y. Xie, "Two-level data compression using machine learning in time series database," in *Proc. IEEE 36th Int. Conf. Data Eng.*, 2020, pp. 1333–1344.
- [56] T. Li, Z. Xu, J. Tang, and Y. Wang, "Model-free control for distributed stream data processing using deep reinforcement learning," *Proc. VLDB Endowment*, vol. 11, pp. 705–718, 2018.
- [57] A. Sharma, F. M. Schuhknecht, and J. Dittrich, "The case for automatic database administration using deep reinforcement learning," 2018, *arXiv:1801.05643*.
- [58] T. Gu, K. Feng, G. Cong, C. Long, Z. Wang, and S. Wang, "The RLR-Tree: A reinforcement learning based R-Tree for spatial data," 2021, *arXiv:2103.04541*.
- [59] X. Yu, G. Li, C. Chai, and N. Tang, "Reinforcement learning with Tree-LSTM for join order selection," in *Proc. IEEE Int. Conf. Data Eng.*, 2020, pp. 1297–1308.
- [60] J. Heitz and K. Stockinger, "Join query optimization with deep reinforcement learning algorithms," 2019, *arXiv:1911.11689*.
- [61] H. Yuan, G. Li, L. Feng, J. Sun, and Y. Han, "Automatic view generation with deep learning and reinforcement learning," in *Proc. IEEE Int. Conf. Data Eng.*, 2020, pp. 1501–1512.
- [62] A. Thusoo *et al.*, "Hive: A warehousing solution over a map-reduce framework," *Proc. VLDB Endowment*, vol. 2, no. 2, pp. 1626–1629, 2009.
- [63] C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins, "Pig latin: A not-so-foreign language for data processing," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2008, pp. 1099–1110.
- [64] M. Armbrust *et al.*, "Spark SQL: Relational data processing in spark," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2015, pp. 1383–1394.
- [65] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proc. Int. Conf. Mach. Learn.*, 2009, pp. 41–48.
- [66] C. Zhang, R. Marcus, A. Kleiman, and O. Papaemmanouil, "Buffer pool aware query scheduling via deep reinforcement learning," 2020, *arXiv:2007.10568*.
- [67] T. Kraska *et al.*, "SageDB: A learned database system," in *Proc. Conf. Innov. Data Syst. Res.*, 2019, pp. 1–10.
- [68] S. Banerjee, S. Jha, Z. Kalbarczyk, and R. Iyer, "Inductive-bias-driven reinforcement learning for efficient schedules in heterogeneous clusters," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 629–641.
- [69] Y. Gao, L. Chen, and B. Li, "Spotlight: Optimizing device placement for training deep neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1676–1684.
- [70] L. Wang, Q. Weng, W. Wang, C. Chen, and B. Li, "Metis: Learning to schedule long-running applications in shared container clusters at scale," in *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.*, 2020, pp. 1–17.
- [71] D. Van Aken, A. Pavlo, G. J. Gordon, and B. Zhang, "Automatic database management system tuning through large-scale machine learning," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2017, pp. 1009–1024.
- [72] Z. Sadri, L. Gruenwald, and E. Lead, "DRLindex: Deep reinforcement learning index advisor for a cluster database," in *Proc. 24th Symp. Int. Database Eng. Appl.*, 2020, pp. 1–8.
- [73] J. Welborn, M. Schaarschmidt, and E. Yoneki, "Learning index selection with structured action spaces," 2019, *arXiv:1909.07440*.
- [74] H. Lan, Z. Bao, and Y. Peng, "An index advisor using deep reinforcement learning," in *Proc. ACM Int. Conf. Informat. Knowl. Manage.*, 2020, pp. 2105–2108.
- [75] S. Chaudhuri and V. Narasayya, "Autoadmin "what-if" index analysis utility," *ACM SIGMOD Rec.*, vol. 27, no. 2, pp. 367–378, 1998.
- [76] T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis, "The case for learned index structures," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2018, pp. 489–504.
- [77] Y. A. Malkov and D. A. Yashunin, "Efficient and robust approximate nearest neighbor search using hierarchical navigable small ICDE graphs," *Pattern Anal. Mach. Intell.*, vol. 42, no. 4, pp. 824–836, 2018.
- [78] W. Dong, C. Moses, and K. Li, "Efficient K-nearest neighbor graph construction for generic similarity measures," in *Proc. 20th Int. Conf. World Wide Web*, 2011, pp. 577–586.
- [79] M. Wang, X. Xu, Q. Yue, and Y. Wang, "A comprehensive survey and experimental comparison of graph-based approximate nearest neighbor search," 2021, *arXiv:2101.12631*.
- [80] D. Baranchuk and A. Babenko, "Towards similarity graphs constructed by deep reinforcement learning," 2019, *arXiv:1911.12122*.
- [81] S. Idreos *et al.*, "Learning data structure alchemy," *Bull. IEEE Comput. Soc. Tech. Committee Data Eng.*, vol. 42, no. 2, pp. 46–57, 2019.
- [82] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," 2016, *arXiv:1611.01578*.
- [83] S. Wu, X. Yu, X. Feng, F. Li, W. Cao, and G. Chen, "Progressive neural index search for database system," 2019, *arXiv:1912.07001*.
- [84] I. Trummer, J. Wang, D. Maram, S. Moseley, S. Jo, and J. Antonakakis, "SkinnerDB: Regret-bounded query evaluation via reinforcement learning," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2019, pp. 1153–1170.
- [85] K. Tzoumas, T. Sellis, and C. S. Jensen, "A reinforcement learning approach for adaptive query processing," in *DB Technical Report*, Aalborg Univ., 2008.
- [86] L. Kocsis and C. Szepesvári, "Bandit based monte-carlo planning," in *Proc. Eur. Conf. Mach. Learn.*, 2006, pp. 282–293.
- [87] R. Marcus, P. Negi, H. Mao, N. Tatbul, M. Alizadeh, and T. Kraska, "Bao: Making learned query optimization practical," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2021, pp. 1275–1288.
- [88] L. L. Perez and C. M. Jermaine, "History-aware query optimization with materialized intermediate views," in *Proc. IEEE Int. Conf. Data Eng.*, 2014, pp. 520–531.
- [89] X. Liang, A. J. Elmore, and S. Krishnan, "Opportunistic view materialization with deep reinforcement learning," 2019, *arXiv:1903.01363*.
- [90] N. Wu and P. Li, "Phoebe: Reuse-aware online caching with reinforcement learning for emerging storage models," 2020, *arXiv:2011.07160*.
- [91] H. Wang, H. He, M. Alizadeh, and H. Mao, "Learning caching policies with subsampling," in *Proc. NeurIPS Mach. Learn. Syst. Workshop*, 2019, pp. 1–5.
- [92] K. Clark and C. D. Manning, "Deep reinforcement learning for mention-ranking coreference models," 2016, *arXiv:1609.08667*.
- [93] Z. Fang, Y. Cao, Q. Li, D. Zhang, Z. Zhang, and Y. Liu, "Joint entity linking with deep reinforcement learning," in *Proc. World Wide Web Conf.*, 2019, pp. 438–447.
- [94] V. Zhong, C. Xiong, and R. Socher, "Seq2SQL: Generating structured queries from natural language using reinforcement learning," 2017, *arXiv:1709.00103*.
- [95] L. Dong and M. Lapata, "Language to logical form with neural attention," 2016, *arXiv:1601.01280*.
- [96] O. Bar El, T. Milo, and A. Somech, "Automatically generating data exploration sessions using deep reinforcement learning," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2020, pp. 1527–1537.
- [97] M.-H. Oh and G. Iyengar, "Sequential anomaly detection using inverse reinforcement learning," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2019, pp. 1480–1490.
- [98] Y. Heffetz, R. Vainshtein, G. Katz, and L. Rokach, "DeepLine: AutoML tool for pipelines generation using deep reinforcement learning and hierarchical actions filtering," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2020, pp. 2103–2113.
- [99] L. Wang, W. Zhang, X. He, and H. Zha, "Supervised reinforcement learning with recurrent neural network for dynamic treatment recommendation," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 2447–2456.
- [100] Y. Ling *et al.*, "Diagnostic inferring via improving clinical concept extraction with deep reinforcement learning: A preliminary study," in *Proc. 2nd Mach. Learn. Healthcare Conf.*, 2017, pp. 271–285.
- [101] R. El-Bourri, D. Eyre, P. Watkinson, T. Zhu, and D. Clifton, "Student-teacher curriculum learning via reinforcement learning: Predicting hospital inpatient admission location," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 2848–2857.
- [102] M. Dixon and I. Halperin, "G-Learner and GIRL: Goal based wealth management with reinforcement learning," 2020, *arXiv:2002.10990*.

- [103] S. Y. Yang, Y. Yu, and S. Almahdi, "An investor sentiment reward-based trading system using gaussian inverse reinforcement learning algorithm," *Expert Syst. Appl.*, vol. 114, pp. 388–401, 2018.
- [104] S. Y. Yang, Q. Qiao, P. A. Beling, W. T. Scherer, and A. A. Kirilenko, "Gaussian process-based algorithmic trading strategy identification," *Quantitative Finance*, vol. 15, no. 10, pp. 1683–1703, 2015.
- [105] J. Zhao, G. Qiu, Z. Guan, W. Zhao, and X. He, "Deep reinforcement learning for sponsored search real-time bidding," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 1021–1030.
- [106] S.-Y. Chen, Y. Yu, Q. Da, J. Tan, H.-K. Huang, and H.-H. Tang, "Stabilizing reinforcement learning in dynamic environment with application to online recommendation," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 1187–1196.
- [107] R. Takanobu, T. Zhuang, M. Huang, J. Feng, H. Tang, and B. Zheng, "Aggregating e-commerce search results from heterogeneous sources via hierarchical reinforcement learning," in *Proc. World Wide Web Conf.*, 2019, pp. 1771–1781.
- [108] P. Wang, K. Liu, L. Jiang, X. Li, and Y. Fu, "Incremental mobile user profiling: Reinforcement learning with spatial knowledge graph for modeling event streams," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 853–861.
- [109] Y. Dou, G. Ma, P. S. Yu, and S. Xie, "Robust spammer detection by nash reinforcement learning," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 924–933.
- [110] W. Shang, Y. Yu, Q. Li, Z. Qin, Y. Meng, and J. Ye, "Environment reconstruction with hidden confounders for reinforcement learning based recommendation," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 566–576.
- [111] X. Xu, C. Liu, and D. Song, "SQLNet: Generating structured queries from natural language without reinforcement learning," 2017, *arXiv:1711.04436*.
- [112] K. Narasimhan, A. Yala, and R. Barzilay, "Improving information extraction by acquiring external evidence with reinforcement learning," 2016, *arXiv:1603.07954*.
- [113] G. Pang, A. van den Hengel, C. Shen, and L. Cao, "Toward deep supervised anomaly detection: Reinforcement learning from partially labeled anomaly data," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2021, pp. 1298–1308.
- [114] C. Lee *et al.*, "Big healthcare data analytics: Challenges and applications," in *Handbook of Large-Scale Distributed Computing in Smart Healthcare*, Berlin, Germany: Springer, 2017.
- [115] I. Halperin, "The QLBS Q-learner goes NuQlear: Fitted Q iteration, inverse RL, and option portfolios," *Quantitative Finance*, vol. 19, no. 9, pp. 1543–1553, 2019.
- [116] F. D. Matthew, H. Igor, and B. Paul, "Machine learning in finance: From theory to practice," Berlin, Germany: Springer, 2021.
- [117] S. Levine, Z. Popovic, and V. Koltun, "Nonlinear inverse reinforcement learning with gaussian processes," in *Proc. Int. Conf. Neural Informat. Process. Syst.*, 2011, pp. 19–27.
- [118] J. Feng *et al.*, "Learning to collaborate: Multi-scenario ranking via multi-agent reinforcement learning," in *Proc. Conf. World Wide Web*, 2018, pp. 1939–1948.
- [119] H. Mao *et al.*, "Park: An open platform for learning-augmented computer systems," in *Proc. Int. Conf. Neural Informat. Process. Syst.*, 2019, pp. 2494–2506.
- [120] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, "Deep reinforcement learning that matters," in *Proc. Conf. Assoc. Advance. Artif. Intell.*, 2018, pp. 3207–3214.
- [121] Z. Luo *et al.*, "MLCask: Efficient management of component evolution in collaborative data analytics pipelines," in *Proc. IEEE Int. Conf. Data Eng.*, 2021, pp. 1655–1666.
- [122] J. You, B. Liu, R. Ying, V. Pande, and J. Leskovec, "Graph convolutional policy network for goal-directed molecular graph generation," in *Proc. Int. Conf. Neural Informat. Process. Syst.*, 2018, pp. 6412–6422.
- [123] B. Schölkopf *et al.*, "Toward causal representation learning," *Proc. IEEE*, vol. 109, no. 5, pp. 612–634, May 2021.
- [124] J. Wang *et al.*, "Polyjuice: High-performance transactions via learned concurrency control," in *Proc. 15th USENIX Symp. Oper. Syst. Des. Implementation*, 2021, pp. 198–216.
- [125] C. Packer, K. Gao, J. Kos, P. Krähenbühl, V. Koltun, and D. Song, "Assessing generalization in deep reinforcement learning," 2018, *arXiv:1810.12282*.
- [126] D. Silver, S. Singh, D. Precup, and R. S. Sutton, "Reward is enough," *Artif. Intell.*, vol. 299, 2021, Art. no. 103535.




Qingpeng Cai received the BSc degree from Zhejiang University, China, in 2019 and currently working toward the PhD degree with the National University of Singapore. He was a research intern with the Zhejiang University CCNT lab in 2017 and the National University of Singapore SSI lab in 2018. His research interests lie in healthcare, time-series representation learning, and reinforcement learning.



Can Cui received the MSc degree from Singapore Management University. He is currently working toward the PhD degree with the School of Computing, National University of Singapore. His research interests lie in the interdisciplinary area of machine learning and finance.



Yiyuan Xiong received the BSc degree from the School of Electronics Engineering and Computer Science, Peking University, China, in 2017. He is currently working toward the PhD degree with the School of Computing, National University of Singapore, Singapore. His research interests lie in the intersection between database and machine learning, especially reinforcement learning.



Wei Wang received the BSc degree from the Renmin University of China, Beijing, China, in 2011 and the PhD degree from the National University of Singapore, Singapore, in 2017. He is a researcher with ByteDance. His research interests include deep learning systems and applications for multimedia data. He has served as a PC member for the ACM Multimedia, VLDB, SIGMOD, ICDE, and DASFAA.



Zhongle Xie received the BSc degree from Shanghai Jiao Tong University in 2014 and the PhD degree from the National University of Singapore in 2020. He is currently a researcher with the Institute of Computing Innovation, Zhejiang University. His research interests include indexing, query processing and database systems. He has served as a member of the review board for VLDB'21.



Meihui Zhang (Member, IEEE) received the BEng in computer science from the Harbin Institute of Technology, China, in 2008, and the PhD in computer science from the National University of Singapore, Singapore, in 2013. She is currently a professor with the Beijing Institute of Technology, Beijing, China and was an assistant professor with the Singapore University of Technology and Design (SUTD), Singapore from 2014 to 2017. Her research interests include data analytics, massive data integration, distributed and blockchain systems. She has served as a vice PC chair/associate editor for VLDB'18, IEEE ICDE'18, VLDB'19, VLDB'20, SIGMOD'21, and ICDE'22.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.