

Deep Reinforcement Learning Applications

Subjects: Computer Science, Interdisciplinary Applications

Contributor: Xuanchen Xiang

Deep Reinforcement Learning (DRL) combines Reinforcement Learning and Deep Learning. It is more capable of learning from raw sensors or images as input, enabling end-to-end learning, which opens up more applications in robotics, video games, NLP, computer vision, healthcare, and more. A milestone in value-based DRL is employing Deep Q-Networks (DQN) to play Atari games by Google DeepMind in 2013.

Keywords: reinforcement learning ; deep reinforcement learning ; Markov decision process ; partially observable Markov decision process

1. Introduction

Where $V(s_t)$ is the estimated state value, initialized using a certain strategy; α is a rate that influences the convergence; G_t is the return from the time that the actor first visited the state-action pair (or sum of returns from each time it visited the pair), can be calculated as below: (12) $G_t = r_t + \gamma V(s_{t+1}) + \gamma^2 r_{t+2} + \dots + \gamma^T r_T$

Q-learning: (14) $Q(s, a) \leftarrow Q(s, a) + \alpha [r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$ ($s \leftarrow s'$)

SARSA: (15) $Q(s, a) \leftarrow Q(s, a) + \alpha [r + \gamma Q(s', a) - Q(s, a)]$ ($s \leftarrow s', a \leftarrow a'$)

A table of representative DRL techniques is shown in **Table 1**.

Table 1. Representative DRL Techniques.

	DRL Algorithms	Main Techniques
Value-based	DQN [1]	Experience Replay, Target Network, Clipping Rewards, Clipping Rewards, and Skipping Frames
	Double DQN [2]	Double Q-learning
	Dueling DQN [3]	Dueling Neural Network architecture
	Prioritized DQN [4]	Prioritized experience replay
	Bootstrapped DQN [5]	Deep exploration with DNNs
	Distributional DQN [6]	Distributional Bellman equation
	Noisy DQN [7]	Parametric noise added to weights
	Rainbow DQN [8]	Combine 6 extensions to the DQN
	Hierarchical DQN [9]	Hierarchical value functions
	Gorila [10]	Asynchronous training for multi agents
Policy-based	TRPO [11]	KL divergence constraint
	PPO [12]	Specialized clipping in the objective function
Actor-Critic	Deep DPG [13]	DNN and DPG
	TD3 [14]	Twin Delayed DDPG
	PGQ [15]	Policy gradient and Q-learning
	Soft Actor Critic (SAC) [16]	Maximum entropy RL framework

DRL Algorithms	Main Techniques
A3C ^[17]	Asynchronous Gradient Descent

2. Applications

Deep Reinforcement Learning has been widely utilized in many domains, as shown in **Table 2**, which is summarized from ^[18]. This chapter will discuss the recent advances in applications in games, robotics, Natural Language Processing (NLP), transportation, industrial applications, communication and networking, and more topics. Usually, different techniques may be employed in an application besides DRL, e.g., AlphaGo is trained using supervised learning and reinforcement learning. In this report, POMDP problems are taken into more consideration, but in reality, it is hard to define an application as a firm POMDP problem or not. We focus applications generally based on Deep Reinforcement Learning, not only on POMDP problems.

Table 2. Various DRL Applications, summarized from ^[18].

Domains	Applications
Healthcare	DTRs, HER/EMR, diagnosis
Education	Educational games, recommendation, proficiency estimation
Transportation	Traffic control
Energy	Decision control
Finance	Trading, risk management
Science, Engineering, Art	Math, Physics, music, animation
Business management	Recommendation, customer management
Computer systems	Resources management, security
Games	Board games, card games, video games
Robotics	Sim-to-real, control
Computer vision	Recognition, detection, perception
NLP	Sequence generation, translation, dialogue

2.1. Games

Researchers have continuously been studying RL algorithms with games as testbeds since long ago. This section discusses this research in board games, card games, and video games. This section refers to Li's overview ^[19].

2.1.1. Board Games

In the board games like Go, chess, checkers, and Backgammon, players have perfect information about each other, and the environment is deterministic (not MDPs). The AI research in two-person games can trace back to the 1950s, Claude Shannon and Alan Turing proposed computer programs that could challenge humans.

RL in Board Games

There have been significant steps in AI playing games towards the DRL direction. In 1994, TD-Gammon ^[20] was a Neural Network to estimate evaluation function with TD learning for Backgammon. Buro ^[20] also wrote papers proposing methods of playing Othello programs in the 1990s, for further references, see: <https://skatgame.net/mburo/publications.html>, accessed on 11 January 2021. Schaeffer et al. ^[21] developed Chinook, playing checkers, between 1989 and 2007. Ginsberg ^[22] is the inventor of GIB for playing the computer bridge. For the Hex games, Huang employed MCTS in MoHex ^[23], and Gao et al. ^[24] trained the neural networks by canonical maximum likelihood to enhance MoHex 2.0.

Mannen's thesis ^[25] implemented experiments in the training of neural networks as evaluation functions for a chess program. Block et al. ^[26] used reinforcement learning in chess engines, extending KinghtCap's ^[27] learning algorithm with a larger database. Later, Lai ^[28] presented Giraffe, using Deep Reinforcement Learning with automatic feature extraction in chess, outperformed other chess engines during the time.

AlphaGo

In 2016, AlphaGo [29] became the first computer Go program that won the professional Go player Lee Sedol. This is the historical moment that AI proves to be able to surpass human learning in board games. AlphaGo was trained by a combination of supervised learning and reinforcement learning of self-play. The networks also combined Monte Carlo simulation with value and policy networks. Later, Silver et al. [30] introduced ALphaGo Zero, based solely on RL without human data or guidance beyond game rules. A neural network is trained to predict AlphaGo or other winner's action selections. This resulted in higher quality and more robust self-play features. Later, Silver et al. [31] also came up with a general algorithm called AlphaZero. It realized superhuman performance in many other games and defeated champion programs in chess, shogi, and Go. Furthermore, Kimura et al. [32] applied AlphaZero to POMDP maze problems, and it turned out effective.

AlphaGo is a turning point in the application of Deep Reinforcement Learning. It has made tremendous progress in recent years.

2.1.2. Card Games

Card games, like many variants of Poker, UNO, and mahjong, have incomplete information. The agents and the environment are stochastic and uncertain. These games are usually POMDP problems.

Texas Hold'em Poker

One of the classic poker games is Texas Hold'em Poker. Davidson [33], Davidson et al. [34], Felix and Reis [35], and Bayes' Bluff presented by Southey et al. [36] used Neural Networks to model opponents in Texas Hold'em Poker, thus the network was able to predict the next actions of players correctly.

Zinkevich et al. [37] presented Counterfactual Regret Minimization(CFR). CFR is to compute an approximate Nash equilibrium for the abstract game by iteratively traversing the game tree. Based on the CFR technique, Bowling et al. [38] weakly solved Heads-up Limit Hold'em Poker. Furthermore, Deep CFR was discussed by Brown et al. [39].

DeepStack computer program presented by Moravčík et al. [40], uses deep neural networks as the value function for its depth-limited look-ahead, defeated professional poker players for the first time, in heads-up no-limit Texas Hold'em. More information about DeepStack can be found at <https://www.deepstack.ai/>, accessed on 11 January 2021.

Prior breakthroughs have been limited to settings with two players. Brown and Sandholm [41] proposed Pluribus, a superhuman AI capable of defeating elite human professionals in six-player no-limit Texas hold'em poker, with self-play through improved Monte Carlo CFR.

Hearts and Wizard

Hearts is a multi-player non-cooperative finite-state zero-sum imperfect-information game that can formulate a POMDP [42]. Fujita and Ishii [43] developed model-based RL for large-scale multi-player games with partial observability and applied it to Hearts, using Markov Chain Monte Carlo (MCMC) technique. Sturtevant and White [44] developed a player for the game of hearts of 4 players, based on stochastic linear regression and TD learning.

In 2017, Wagenaar's thesis [45] used a multi-layer perceptron (MLP) and Monte Carlo learning to play Game Hearts and realized DRL. However, more work is required to realize mature performance in the Hearts game.

Similar to Hearts, Wizard is also a multi-player imperfect-information poker game. RL has turned out effective for these POMDP games. Backhus et al. [46] applied RL in the Wizard game, separately dealing with forecasting and trick playing.

RLCard

In recent years, the full-width extensive-form fictitious play (XFP) was introduced by Heinrich and Heinrich and Lanctot [47], Neural Fictitious Self-Play (NFSP) was presented by Heinrich and Heinrich and Silver [48], and Unified Game-Theoretic Approach was proposed by Lanctot et al. [49]. We can see that RL strategies can perform well in betting games. These results give the inspiration for researchers to explore DRL in more card games.

Zha et al. [50] presented RLCard, an open-source toolkit for reinforcement learning research in card games. It provides various card game environments, including Blackjack, Leduc Hold'em, Texas Hold'em, UNO, Dou Dizhu, and Mahjong, with easy-to-use interfaces (see **Figure 1**).




Figure 1. RLCard: It supports various styles of card games and board games [50].

2.1.3. Video Games

Video games are valuable resources with video frames as inputs to RL/AI agents, in which computer games are the most popular testbeds for DRL algorithms. Similar to the MDP process as shown in **Figure 1**, DRL for video games employ CNNs to extract features from the video frames (environment) to be recognized by the agents. Recently, DRL has been applied in Sega, Nintendo, Xbox, etc. Shao et al. [51] made a comprehensive survey of DRL in video games. Based on this, **Figure 2** presents popular video-game platforms for RL research. In this section, we particularly dig into details about games with POMDP characteristics.




Figure 2. Popular video-game platforms for Reinforcement Learning research, summarized from [51].

Atari

A traditional platform, the Arcade Learning Environment (ALE) [52], provides an interface to lots of games in Atari 2600 and is a pioneer evaluation platform in RL research. As described in Chapter One, many value-based methods are tested in Atari 2600 and the un-mentioned techniques, such as DRQN [53] and ADRQN [54]. Ayta et al. [55] presented an exciting method of playing Atari exploration games by watching Youtube. They map unaligned videos from several sources to a common representation, then embed a single YouTube video in this representation to construct a reward function. Hence, the agent learns to imitate human gameplay. This method exceeds human-level performance on MONTEZUMA'S REVENGE, PITFALL! and PRIVATE EYE, with no environment rewards in games.

ViZDoom

Doom is a classical first-person shooter (FPS) game with a 3D environment, where agents observe perspective 2D projections from their positions as pixel matrices. The partial observability of game states is the main challenge. ViZDoom [56] is a Doom-based platform that offers API.

Lample and Chaplot [57] used a divide and conquered method in Doom to divide the action space. They trained an action agent and a navigation agent. Different agents will act [58] based on the detection of an enemy. Akimov and Makarov [59] proposed Categorical 51 with Multi-step (C51M) architecture, outperformed baseline methods in ViZDoom environment.

In the survey [51], they mentioned the following important techniques that evaluated in ViZDoom: Wu and Tian [60] proposed a framework combining A3C with curriculum learning [61]; Parisotto and Salakhutdinov [62] developed Neural Map, a memory system with an adaptable write operator, using a spatially structured 2D memory image to store arbitrary environment information; Shao et al. [63] also introduced ACKTR, which teaches agents to battle in the ViZDoom environment.

TORCS

The Open Racing Car Simulator (TORCS) is a racing game with acceleration, braking, and steering as actions. It has been used as testbeds for visual-input games. It is also an essential platform for testing autonomous driving methods with its realistic feature.

In 2014, Koutník et al. [64][65] had Max-Pooling Convolutional Neural Network (MPCNN) compressor and recurrent neural network (RNN) controllers evolved successfully to control a car in TORCS based on visual input; DDPG [13], as mentioned, was assessed in TORCS; Wang et al. [66] adopted DDPG and selected sensor information from TORCS as inputs, evaluated on different modes in TORCS; Sallab et al. [67] formulated two categories of algorithms for lane-keeping assist: DQN for discrete actions and Deep Deterministic Actor-Critic Algorithm (DDAC) for continuous actions; Furthermore, Liu et al. [68] proposed DDPG method with expert demonstrations and supervised loss, simulated in TORCS, which improves efficiency and stability significantly; Verma et al. [69] presented Programmatically Interpretable Reinforcement Learning (PIRL), to generate interpretable agent policies and a new method called Neurally Directed Program Search (NDPS), to find a programmatic policy with maximal reward;

In addition to the works shown above, some techniques discussed by Shao et al. [51]: Sharma et al. [70] proposed Fine-Grained Action Repetition (FiGAR) to improve DDPG; Gao et al. [71] used TORCS to evaluate Normalized Actor-Critic (NAC); Mazumder et al. [72] developed a property called state-action permissibility (SAP), with which the agent learns to solve the lane-keeping problem in TORCS, with faster training process; Li et al. [73] break the vision-based lateral control system down into a perception module and a control module; Zhu and Zhao [74] uses DRL to train a CNN to perceive data.

Minecraft

Minecraft is a well-known sandbox video game where players can explore a blocky, procedurally-generated 3D world and build structures and earthworks with freedom in different game modes. It supports scenarios from navigation and survival to collaboration and problem-solving tasks, making it a rich environment to carry out RL.

Project Malmö [75] is an AI experimentation platform on top of Minecraft, created by Microsoft. Marlo, based on Malmö, provides a higher-level API and a more standardized RL-friendly environment for RL study. In 2019, Guss et al. [76] presented a comprehensive, large-scale, simulator-paired dataset of human demonstrations: MineRL, which provides an RL environment and competitions.

Some RL tasks created in Minecraft were effectively achieved. Liu et al. [77] constructed tasks with shared characteristics in Minecraft for multi-task DQN (MDQN), which turned out effective. Oh et al. [78] also designed tasks to evaluate DRL architectures: Memory Q-Network (MQN), Recurrent Memory Q-Network (RMQN), and Feedback Recurrent Memory Q-Network (FRMQN). Frazier [79] conducted experiments with two RL algorithms: Feedback Arbitration and Newtonian Action Advice, enabling human teachers to give action advice. Furthermore, HogRider [80] is the champion agent of MCAC in 2017, using a Q-learning approach with state-action abstraction and warm start; Tessler et al. [81] proposed a lifelong learning system that is able to transfer knowledge from one task to another, incorporated with Hierarchical-DRLN with a deep skill array and skill distillation; Jin et al. [82] utilized counterfactual regret minimization to update an approximation to an advantage-like function iteratively.

DeepMind Lab

DeepMind Lab [83] is a 3D customizable first-person game platform for agent-based AI research, a modified version of Quake III Arena. It is usually used to study how RL agents deal with complex tasks in large, partially observed, and visually diverse worlds.

Many researchers have created RL tasks in DeepMind Lab and tested algorithms. Johnston [84] introduced an unsupervised auxiliary task—"surprise pursuit". It quantifies the "surprise" the agent encounters when exploring the environment and trains a policy to maximize this "surprise".

There also are researchers who used DM-Lab to propose new architectures. As mentioned by [51], Mankowitz et al. [85] proposed Unicorn, which demonstrates robust continual learning; Schmitt et al. [86] had a teacher agent to kickstart the training of a new student agent. The method mentioned by Jaderberg et al. [87] trains plenty of agents concurrently from lots of parallel matches, where agents play cooperatively in teams, competing in randomly generated environments.

Later on, DeepMind also invented Psychlab [88], a psychophysics testing room embedded inside the world of DM-Lab. This framework allows DRL agents to be directly compared with humans, lifted from cognitive psychology and visual psychophysics.

Real-Time Strategy Games

Real-time strategy games (RTS) allow players to play the game in “real-time” simultaneously rather than to take turns to play. RTS like StarCraft, Dota, and LOL are popular platforms in AI research. The real-time games usually consist of thousands of time steps and hundreds of actions selected in real-time through the gameplay. StarCraft is a game in which players balance high-level economic decisions in the control of tons of units. In each time step, the agent receives imperfect information from observations. In 2019, DeepMind proposed AlphaStar ^[89], which defeats professional players for the first time. **Table 3** presents the techniques carried out in StarCraft.

MOBA (Multiplayer Online Battle Arena) games have two teams, and each team consists of five players. OpenAI Five ^[90] used Dota2 as the research platform and achieved expert-level performance, and is the first AI system that could defeat the world champions in an esports game.

Table 3. Techniques used in StarCraft, summarized based on ^[51].

StarCraft Algorithms	Main Techniques
GMEZO ^[91]	Micromanagement; greedy MDP; episodic zero-order exploration
BiCNet ^[92]	Multiagent bidirectionally coordinated network
Stabilising Experience Replay ^[93]	Importance sampling; fingerprints
PS-MAGDS ^[94]	Parameter sharing multi-agent gradient descent SARSA(
)	
MS-MARL ^[95]	Master-slave architecture; multi-agent
QMIX ^[96]	Decentralised policies in a centralised end-to-end training
NNFQ CNNFQ ^[97]	Neural network fitted Q-Learning
SC2LE ^[98]	FullyConv-A2C; baseline results
Relational DRL ^[99]	Self-attention to iteratively reason about the relations
TStarBots ^[100]	Flat action structure; hierarchical action structure
Modular architecture ^[101]	Splits responsibilities between multiple modules
Two-level hierarchical RL ^[102]	Macro-actions; a two-layer hierarchical architecture

2.2. Robotics

Robotics is widely used in science, industries, healthcare, education, entertainment, and many more domains. Robotics has many challenging problems involved, including perceptron, control system, operation system, etc. The development of Deep Reinforcement Learning paves the way in robotics. Many applications mentioned in the gaming section, including AlphaGo, are also involved in Robotics. While in this section, we focus on manipulation and locomotion.

For more surveys in Robotics, see ^[103] in 2009, a study of robot learning from demonstration (LfD), with which a policy is learned from demonstrations, provided by a teacher; Deisenroth ^[104] made a survey on policy search for robotics in 2013; In 2014, Kober and Peters ^[105] provided a general survey on RL in robotics; Tai et al. ^[106] presented a comprehensive survey for learning control in robotics from reinforce to imitation in 2018.

A 3D physics simulator MuJoCo physics engine ^[107], which stands for Multi-Joint dynamics with Contact, was developed by OpenAI for continuous control. Mainly based on which series algorithms are evaluated.

2.2.1. Manipulation

Dexterous manipulation is one of the most challenging control problems in robotics due to the barriers of high dimensionality, intermittent contact dynamics, and under-actuation in dynamic object manipulation; Dexterous manipulation can be widely used in different domains. These make it a hot topic in Robotics. Nguyen and La ^[108] provided a review of DRL for Robot Manipulation with the challenges, the existing problems, and future research directions.

For reinforcement learning, researchers made significant progress, making DRL in manipulation tasks accessible in recent years. For example, in 2008, Peters and Schaal ^[109] evaluated policy gradient methods on parameterized motor primitives, and Episodic Natural Actor-Critic ^[110] outperforms the others; In 2010, Theodorou et al. ^[111] presented Policy Improvement with Path Integrals (PI2); In 2010, Peters et al. ^[112] suggested Relative Entropy Policy Search (REPS) method; In 2011, Kalakrishnan et al. ^[113] used PI2 with force control policies. Nair et al. ^[114] demonstrated overcoming the exploration problem and solved the block stacking task using RL.

Multi-Fingered Hands

Dexterous multi-fingered hands can perform a wide range of manipulation skills. They are essential in research while it requires high dimensional observation and ample action space. Kumar et al. ^[115] demonstrated online planning based on an ADROIT platform, and Kumar et al. ^[116] described a model-based method combining optimizing Linear–Gaussian controllers models with KL-constrained optimization. It was applied on a pneumatically-actuated tendon-driven 24-DoF hand, which learned hand manipulation involving clockwise rotation of the object.

Rajeswaran et al. ^[117] then demonstrated that model-free DRL could scale up to complex manipulation tasks. It does not require a model of the dynamics and can optimize the policy directly. They used a combination of RL and imitation learning Demo Augmented Policy Gradient (DAPG), achieved in object relocation, in-hand manipulation, door opening, and tool use tasks.

However, in the real world, the interaction time with the environment is limited, the model-free methods perform poorly in such cases. Haarnoja et al. ^[118] studied how maximum entropy policies were trained using soft Q-learning. Policies learned using soft Q-learning can be composed to create new policies, the optimality of which can be enclosed knowing the divergence between the composing policies. By composing existing skills can efficiently improve the training process. This method is sample efficient than prior model-free DRL methods and can be performed for simulated and real-world tasks.

Zhu et al. ^[119] surveyed multi-fingered manipulation and used model-free DRL algorithms with low-cost hardware, succeeded in learning complex contact-rich behaviors.

Considering interacting with fragile objects, Huang et al. ^[120] realized gentle object manipulation, minimizing impact forces, with dynamics-based surprise and penalty-based surprise.

Grasping

Levine et al. ^[121] trained using guided policy search (GPS), which can perform policy search by supervised learning, where supervision is granted by a simple trajectory-centric RL method. It was evaluated in various tasks, like learning to insert, screw, fit, and place the objects. DRL has also been applied to learn grasping objects ^{[122][123]}, Popov et al. ^[124] introduced two extensions to the Deep Deterministic Policy Gradient (DDPG) algorithm to improve data efficiency with linear speedup for 16 parallel workers. The method succeeded in grasping objects and precisely stacking them on another in simulation. Kalashnikov et al. ^[125] introduced QT-Opt, a scalable self-supervised RL framework, enables closed-loop vision-based control. They achieved real-world grasping tasks with a high success rate based only on RGB vision perception.

Opening a Door

Gu et al. ^[126] presented Asynchronous NAF with practical extensions. The experiment shows success in several 3D manipulation skills in simulation and a door opening skill on real robots.

Cloth Manipulation

Regarding cloth manipulation, the state of the cloth can vary, unlike rigid objects. Tsurumine et al. ^[127] proposed Deep P-Network (DPN) and Dueling Deep P-Network (DDPN) to combine smooth policy update with feature extraction, evaluated in flipping a handkerchief and folding a t-shirt by a dual-arm robot. Jangir et al. ^[128] investigated the effectiveness of different textile state representations and the importance of speed and trajectory, solved dynamic cloth manipulation tasks by restricting the manipulator workspace.

Magnetic Manipulation

Magnets can be utilized in magnetic resonance imaging (MRI) applications, lifting machinery, levitation devices, jewelry, bearings, audio equipment, and hard disc drives. de Bruin et al. ^[129] deployed two experience replay databases in the method, one filled with experiences in the First In First Out (FIFO) manner. The experiences are overwritten with new

experiences in the other. This method limits the detrimental effects of exploration, tested on a simulated horizontal magnetic manipulation task.

2.2.2. Locomotion

Over the recent years, many RL methods have been proposed to control locomotion, locomotion examples are shown in **Figure 3**. On humanoid platforms such as Nao, Pepper, Asimo, and Robotis-op2 include the cameras ^[130], many algorithms were evaluated. The main goal for robot locomotion is to imitate human motions dynamically. Several main techniques such as DeepDPG ^[13], A3C ^[17], TRPO ^[11], etc., that were previously covered, were tested in robotics locomotion and turned out effective. They are the key to the vast development of DRL algorithms in locomotion ^{[131][132]}.

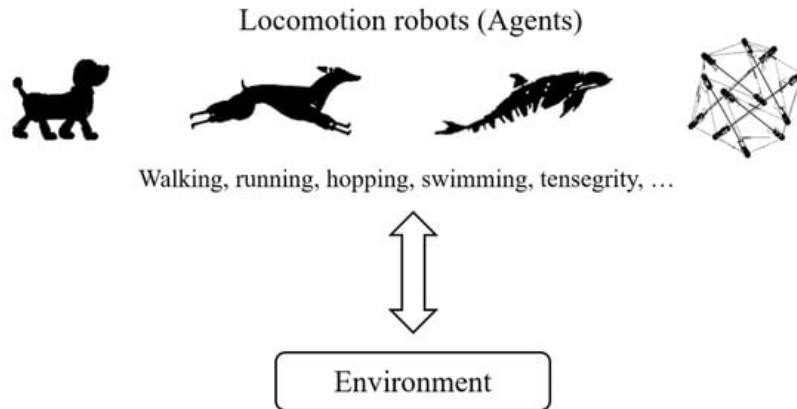


Figure 3. Locomotion examples.

In Robotics, many tasks are complex and poorly specified. Christiano et al. ^[133] had humans compare possible trajectories, used the data to learn a reward function, and optimized the learned reward function with RL and scaled up this method to DRL. In particular, the robotics tasks including walker, hopper, swimmer, etc., were implemented in MuJoCo and OpenAI Gym ^[134].

DeepLoco ^[135] adopted a two-level hierarchical control framework. First, low-level controllers (LLC) learn to operate at a timescale and achieve robust walking gaits. Second, high-level controllers (HLC) learn the policy at the timescale of steps by invoking desired step targets for the low-level controller. Both levels of the hierarchy use an identical style of the actor-critic algorithm.

To develop sound strategies for integrating or merging policies for multiple skills, Berseth et al. ^[136] proposed Progressive Learning and Integration via Distillation (PLAiD). Distillation refers to the combination of the policies of one or more experts, thus creating one controller to perform the tasks of a set of experts. The main blocks are policy transfer and multi-task policy distillation, evaluated on five different terrains with a 2D PD-biped. Compared with MultiTasker, PLAiD scales better.

Generally, model-free methods can learn a wide range of robotic skills but require a large number of samples to achieve good performance. Model-based algorithms can provide more efficient learning but are challenging to scale up to high-capacity models. Md-Mf ^[137] uses deep neural network dynamics models to initialize a model-free learner, then uses the policy trained by a model-based controller as the initial policy for a TRPO (model-free) algorithm. Testing locomotion skills on a simulated quadrupedal robot, this method surpasses the individual model-free and model-based approach.

Duan et al. ^[138] studied Truncated Natural Policy Gradient (TNPG) and compared it with TRPO and DDPG. They also provided a benchmark of continuous control problems for RL.

Following are some examples of different locomotion skills.

Walking/Running

Xie et al. ^[139] demonstrated DRL on a bipedal robot Cassie, developed by Agility Robotics. The framework used PPO and PD control on a simulated model of Cassie in the MuJoCo ^[140] simulation and is robust to disturbances.

Designing agile locomotion for robots usually requires human expertise and manual tuning. Tan et al. ^[140] presented a system to learn quadruped locomotion from scratch or allow users to guide the learning process. In addition, it narrows the gap between simulation and the real world, successfully developed two gaits, trotting, and galloping.

Haarnoja et al. [132] proposed an extension of the soft actor-critic (SAC) algorithm with dynamic temperature, based on maximum entropy RL. It was applied to a real-world Minitaur robot without any model or simulation. This sample-efficient and stable algorithm outperforms standard SAC, DDPG, TD3, and PPO.

Özalp et al. [130] evaluated Double Dueling Q Networks (D3QN) [141] and DQN [142] for locomotion using a Robotis-op2 humanoid robot.

Hopping

Peng et al. [143] introduced a mixture of actor-critic experts (MACE), constructed of multiple actor-critic pairs where each specializes in particular aspects of the motion. It can work with high-dimensional state descriptions for terrain-adaptive dynamic locomotion skills. With leaps or steps as output actions, the method was tested on seven classes of terrain and gives excellent results.

Swimming

Heess et al. [144] surveyed a hierarchical motor control architecture that encourages the low-level controller to focus on reactive motor control. The high-level controller directs behavior towards the task goal by transmitting a modulatory signal. Evaluated on a swimming snake, a quadruped, and a humanoid, this method failed in end-to-end learning but turned out effective in transferring tasks with sparse reward.

To learn stochastic energy-based policies for continuous states and actions, not only in tabular domains, Haarnoja et al. [145] introduced Soft Q-learning, which expresses the optimal policy via a Boltzmann distribution. The simulated swimming snake and quadrupedal robot results show that the method can effectively capture complex multi-modal behaviors and learn general-purpose stochastic policies.

To narrow the gap between simulation and the real world and generalize from training to testing, Pinto et al. [146] introduced robust adversarial reinforcement learning (RARL). The adversary learns to apply destabilizing forces on specific points on the system, encouraging the protagonist to retain a strong control policy. This method improves training stability, and is robust to training/testing conditions, and outperforms the baselines even without an adversary.

Verma et al. [147] studied that DRL can learn efficient collective swimming in unsteady and vortical flow fields.

Tensegrity

Tensegrity robots are composed of rigid rods connected by elastic cables, which display high tolerance to physical damage. With their unique properties, tensegrity robots are appealing for planetary exploration rovers. The primary locomotion of tensegrity robots is rolling.

Zhang et al. [148] applied an extension of mirror descent guided policy search (MDGPS) to periodic locomotion movements on a SUPERball tensegrity robot. The learned locomotion policies performed well, are more effective, and can generalize better than open-loop policies. Luo et al. [149] showed that MDGPS could efficiently learn even with limited sensory inputs.

2.2.3. Robotics Simulators

Robotics simulators are used to simulate and model the structures or actions of a physical robot without depending on the real-world machine. The use of simulators to develop robots allows for programs to be written and debugged offline, and the final versions can then be transferred to the real robot, thus saving cost and time. **Table 4** presents the popular robotics simulators nowadays.

Table 4. Robotics simulators, adapted from Wikipedia (2021) [150].

2.3. Natural Language Processing

Natural language processing (NLP) focuses on the communication between computers and humans—how to program computers to understand, analyze, and process natural language, including generation, language grounding, information extraction, dialogue, etc. Due to the unique features of NLP, DRL has been playing an important role. For more surveys, see [18][151][152][153][154][155].

2.3.1. Neural Machine Translation

Neural Machine Translation (NMT) refers to using machine learning software to translate content into another language. RL has been shown as an effective technique for NMT systems recently. Examples of machine translation providers are

Google Translate, Yandex Translate, Watson Language Translator, etc. Wu et al. [156] provided a literature review on RL for NMT, in which they compared several vital factors, including baseline reward and reward shaping, and studied how to train better NMT systems. They also proposed training with source-side and (or) target-side monolingual data and succeeded in the Chinese-English translation task.

He et al. [157] proposed dual-NMT for English-French dual translation. The English-to-French translation (primal) agent and French-to-English translation (dual) agent can form a closed loop and generate feedback signals to train the models by teaching each other through the RL process. This system works well on bilingual translation, especially by learning from monolingual data.

In the traditional NMT setting, the translator has to wait until the user completes the sentence, then translates. Satija and Pineau [158] presented an approach to interpret in real-time.

The open-source toolkit OpenNMT was described by Klein et al. [159], which prioritizes efficiency and modularity; Zhang et al. [160] introduced THUMT that supports minimum risk training and semi-supervised training.

For generative adversarial net (GAN) in NMT, BR-CSGAN [161] leverages the BLEU reinforced GAN to improve the NMT, and Wu et al. [162] introduced Adversarial-NMT with a CNN based adversary.

2.3.2. Dialogue

Dialogue systems generally have two categories—task-oriented dialog agents and chatbots. Task-oriented dialog agents usually have short Q-and-A conversations to help complete specific tasks or provide information, while chatbots are usually designed human-likely and used for entertainment [163].

Cuayáhuatl has done significant research on dialogue systems with RL [164][165][166][167]. In 2017, he presented SimpleDS [168], as shown in **Figure 4**. It's a publicly available dialogue system trained with DQN, which extends the traditional methods [167] and avoids manual feature engineering by directly selecting actions from raw text and user responses.

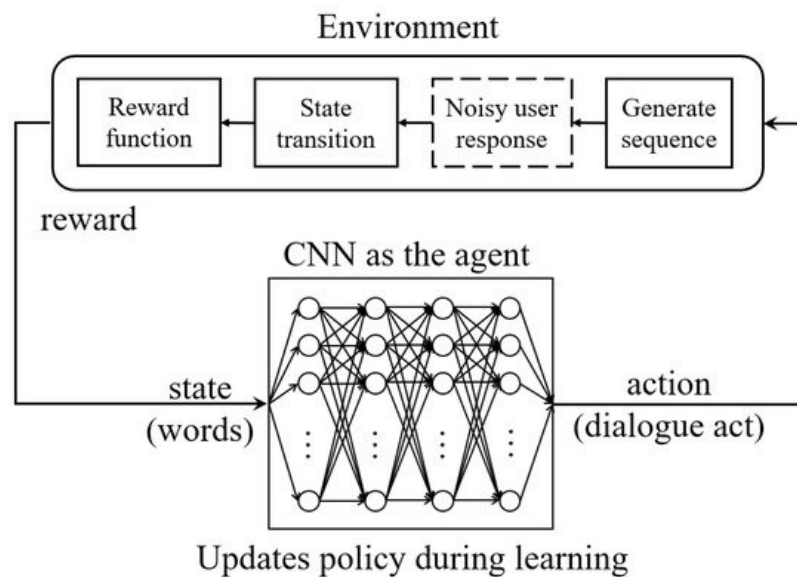


Figure 4. SimpleDS architecture [168].

Cuayáhuatl et al. [169] proposed a two-stage method—the first stage reduces the times of weight updates, and the second uses limited mini-batches sampled from experience replay for accelerating the induction of single or multi-domain dialogue policies. The training turns out five times faster than a baseline. Furthermore, Cuayáhuatl et al. [170] described an ensemble-based approach applied to value-based DRL chatbots. Actions are derived from sentence clustering, the training datasets are derived from dialogue clustering, and the reward function is derived from human–human dialogues without manual annotations. The ensemble of DRL agents turns out more promising than a single DRL agent or SEQ2SEQ model.

Most DRL systems perform poorly in exploring online in the environment or learn effectively from off-policy data. Jaques et al. [171] leveraged models pre-trained on data with KL-control to penalize divergence. They used uncertainty estimates with dropout to lower bound the target Q-values. Tested in the real world with humans, the Way Off-policy algorithm can

effectively learn from different reward functions from human interaction data and outperforms prior methods in off-policy batch RL.

Text Generation

Text generation is the basis of dialogue systems. Ranzato et al. [172] proposed mixed incremental cross-entropy reinforce (MIXER) with a loss function combining REINFORCE [173] and cross-entropy. For sequence prediction, Bahdanau et al. [174] proposed an actor-critic algorithm, inspired by Ranzato et al. [172]. Yu et al. [175] introduced sequence generative adversarial nets with policy gradient (SeqGAN). Li et al. [176] from Huawei Technologies proposed a new DL framework for paraphrase generation, consisting of a generator and an evaluator. The generator is modeled as the SEQ2SEQ learning model for paraphrase generation. As a deep matching model, the evaluator is trained via supervised learning or inverse RL in different settings for paraphrase identification.

Sentence Simplification

Sentence simplification is also crucial in dialogue systems, which aims to make sentences easier to read and understand for the agents. Zhang and Lapata [177] addressed this simplification issue with an encoder-decoder model with a DRL framework, known as Deep Reinforcement Sentence Simplification (DRESS). Their work succeeded in simplicity, grammaticality, and semantic fidelity to the input.

Coherent Dialogues

Some neural models of dialogue generation tend to be short-sighted, and modeling the future direction of dialogue is crucial to generating coherent and exciting conversations. Li et al. [178] applied DRL with policy search and curriculum learning on the SEQ2SEQ model [179], which can generate utterances that optimize future reward.

Goal-Oriented Dialogues

Dhingra et al. [180] proposed KB-InfoBot, an end-to-end dialogue agent to help people search Knowledge Bases (KBs). They offered a differentiable probabilistic framework to compute the posterior distribution (a Soft-KB lookup) to avoid symbolic queries used in previous works.

To evaluate persuasion strategies and negotiation dialogue, Keizer et al. [181] presented a comparative evaluation of various negotiation strategies on the game “Settlers of Catan” online. It outperforms the original rule-based strategy and human players.

Ilievski et al. [182] presented goal-oriented chatbot dialog management bootstrapping with transfer learning, which helps users achieve a predefined goal. Due to the domain specificity, the available data is limited to obtain. The transfer learning method is to boost the performances and mitigate the limitation. They describe that transfer learning improves performance significantly with a low number of goals or whole target domain data.

For social robot navigation, Ciou et al. [183] presented composite reinforcement learning (CRL) framework, with which agents learn social navigation with sensor input and reward update based on human feedback.

2.3.3. Visual Dialogue

Das et al. [184] introduced a goal-driven training for visual question answering (VQA) and dialogue agents. The experiment contains two cooperative robots, with one giving descriptions and the other giving answers. The results show that it outperforms supervised learning agents.

Thanks to the encoder-decoder architectures, End-to-end dialogue systems have developed rapidly. However, there are some drawbacks in encoder-decoder models when engaging in task-oriented dialogue. Strub et al. [185] introduced the DRL method to optimize vision-based task-oriented dialogues with REINFORCE algorithm. It resulted in significant improvement over a supervised baseline model, which was capable of generating coherent dialogue.

Visual Captioning

Evaluation metrics are created to evaluate image captioning. Standard syntactic evaluation metrics, such as BLEU, METEOR, and ROUGE, are usually not well correlated. SPICE and CIDEr are better correlated but are hard to optimize. Liu et al. [186] used a policy gradient method to optimize a linear combination of SPICE and CIDEr (SPIDER).

In reference to visual problems, image captioning is challenging due to the complexity of understanding the image content and describing it in natural language. Xu et al. [187] introduced an attention-based model for image caption generation.

Adaptive encoder-decoder attention model ^[188] with an LSTM extension as a visual sentinel decides whether to attend to the image and where to generate information. Ren et al. ^[189] introduced a decision-making framework for image captioning, which has a policy network and a value network. Both networks are learned using an actor-critic approach with a visual-semantic embedding reward.

Wang et al. ^[190] proposed an Adversarial Reward Learning (AREL) framework for visual storytelling. The model learns a reward function from human demonstrations, then optimizes policy search with the learned reward function.

Video captioning is to generate a description of the actions in the video automatically. Wang et al. ^[191] proposed a hierarchical RL framework for video captioning. A high-level module learns to create sub-goals, and a low-level module recognizes the actions to fulfill the sub-goal. The HRL model turns out a good performance on MSR-VTT and Charades Captions datasets. For video captioning, Pasunuru and Bansal ^[192] presented a mixed-loss policy gradient approach that allows for metric-based optimization, then improved with an entailment-corrected CIDE_r reward, successfully achieved on MSR-VTT.

Visual Relationship and Attribute Detection

Apart from generating labels and captioning, understanding relationships among objects and their attributes in the scene is challenging for computers. Liang et al. ^[193] proposed a deep Variation-structured Reinforcement Learning (VRL) framework for visual relationship and attribute detection. A directed action graph is built to represent semantic correlations between object categories, predicates, and attributes. Then based on the state information, a variation-structured traversal over the chart is used to generate an action set for each step. Predictions are made using DRL and the extracted phrases in the state vector. This method can significantly achieve better detection results on datasets involving relationships and attributes.

3. Conclusions

In this paper, we have presented the fundamental concepts of Markov Decision Processes(MDP), Partially Observable Markov Decision Processes (POMDPs), Reinforcement Learning (RL), and Deep Reinforcement Learning (DRL). Generally, we have discussed the representative DRL techniques and algorithms. The core elements of an RL algorithm are value function, policy, reward, model and planning, exploration, and knowledge. In DRL applications that were discussed, most of the algorithms applied are extensions of the representative algorithms, for example, DQN, DPG, A3C, etc., with some optimizations, respectively. We then discussed the applications of DRL in some popular domains. As DRL provides optimal policies and strategies, it has been widely utilized in gaming, including almost all kinds of games. As the leading entertainment among people, a better gaming experience is always desired. Another popular application of DRL is in Robotics. We discussed manipulation, locomotion, and other aspects of robotics where DRL can be applied. Robotics simulators are also discussed, where DRL is used to select the appropriate simulator when needed. Finally, the application of DRL in natural language processing (NLP) is presented.

As we witnessed, DRL has been developing fast and on enormous scales. Some groups pushed forward the development of DRL, like OpenAI, DeepMind, AI research office in Alberta, and research center led by Rich Sutton, and more. From the research, we have acknowledged the algorithms; variations, and the development of all the optimizations. Based on the review, we are more familiar with the usability of each method. For example, model-free methods work better in scaled-up tasks. The value function is key to RL, in DQN and its many extensions; policy gradient approaches have also been utilized in lots of applications, such as robotics, dialogue systems, machine translation, etc. New learning techniques have also been applied, such as transfer, curriculum learning, adversarial networks, etc. When we select a method or combinations of methods to apply in our models, the characteristics we need to consider are stability, convergence, generalizing power, accuracy, efficiency, scalability, robustness, safety, cost, speed, simplicity, etc. For example, in healthcare, we focus more on stability and security rather than simplicity or cost. This topic, along with applications in transportation, communications and networking, and industries, will be discussed in a Part 2 follow-up paper.

References

1. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; Riedmiller, M. Playing Atari with Deep Reinforcement Learning. arXiv 2013, arXiv:1312.5602.
2. van Hasselt, H.; Guez, A.; Silver, D. Deep Reinforcement Learning with Double Q-Learning. arXiv 2015, arXiv:1509.06461.

3. Wang, Z.; Schaul, T.; Hessel, M.; van Hasselt, H.; Lanctot, M.; de Freitas, N. Dueling Network Architectures for Deep Reinforcement Learning. In Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 20–22 June 2016.
4. Schaul, T.; Quan, J.; Antonoglou, I.; Silver, D. Prioritized Experience Replay. arXiv 2016, arXiv:1511.05952.
5. Osband, I.; Blundell, C.; Pritzel, A.; Van Roy, B. Deep Exploration via Bootstrapped DQN. arXiv 2016, arXiv:1602.04621.
6. Bellemare, M.G.; Dabney, W.; Munos, R. A Distributional Perspective on Reinforcement Learning. In Proceedings of the 34th International Conference on Machine Learning, Sydney, Australia, 6–11 August 2017.
7. Fortunato, M.; Azar, M.G.; Piot, B. Noisy Networks for Exploration. arXiv 2018, arXiv:1706.10295.
8. Hessel, M.; Modayil, J.; van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; Silver, D. Rainbow: Combining Improvements in Deep Reinforcement Learning. In Proceedings of the 32nd AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018.
9. Kulkarni, T.D.; Narasimhan, K.R.; Saeedi, A.; Tenenbaum, J.B. Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation. arXiv 2016, arXiv:1604.06057.
10. Nair, A.; Srinivasan, P.; Blackwell, S.; Alcicek, C.; Fearon, R.; Maria, A.D.; Panneershelvam, V.; Suleyman, M.; Beattie, C.; Petersen, S.; et al. Massively Parallel Methods for Deep Reinforcement Learning. arXiv 2015, arXiv:1507.04296.
11. Schulman, J.; Levine, S.; Moritz, P.; Jordan, M.; Abbeel, P. Trust Region Policy Optimization. In Proceedings of the 32nd International Conference on Machine Learning, Lille, France, 6–11 July 2015.
12. Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; Klimov, O. Proximal Policy Optimization Algorithms. arXiv 2017, arXiv:1707.06347.
13. Lillicrap, T.P.; Hunt, J.J.; Alexander Pritzel, N.H.; Erez, T.; Tassa, Y.; Silver, D.; Wierstra, D. Continuous Control with Deep Reinforcement Learning. arXiv 2016, arXiv:1509.02971.
14. Fujimoto, S.; van Hoof, H.; Meger, D. Addressing Function Approximation Error in Actor-Critic Methods. arXiv 2018, arXiv:1802.09477.
15. O'Donoghue, B.; Munos, R.; Kavukcuoglu, K.; Mnih, V. Combining policy gradient and Q-learning. arXiv 2017, arXiv:1611.01626.
16. Haarnoja, T.; Zhou, A.; Abbeel, P.; Levine, S. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018.
17. Mnih, V.; Badia, A.P.; Mirza, M.; Graves, A.; Harley, T.; Lillicrap, T.P.; Silver, D.; Kavukcuoglu, K. Asynchronous Methods for Deep Reinforcement Learning. In Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 20–22 June 2016.
18. Li, Y. Deep Reinforcement Learning: An Overview. arXiv 2018, arXiv:1701.07274.
19. Tesauro, G. TD-Gammon, a Self-Teaching Backgammon Program, Achieves Master-Level Play. *Neural Comput.* 1994, 6, 215–219.
20. Buro, M. How Machines have Learned to Play Othello. *IEEE Intell. Syst.* 1999, 14, 12–14.
21. Schaeffer, J.; Burch, N.; Björnsson, Y.; Kishimoto, A.; Muller, M.; Lake, R.; Lu, P.; Sutphen, S. Checkers Is Solved. *Science* 2007, 317, 1518–1522.
22. Ginsberg, M.L. GIB: Imperfect Information in a Computationally Challenging Game. *J. Artif. Intell. Res.* 2001, 14, 303–358.
23. Huang, S.C.; Arneson, B.; Hayward, R.B.; Müller, M.; Pawlewicz, J. MoHex 2.0: A Pattern-Based MCTS Hex Player. In *Computers and Games*; van den Herik, H.J., Iida, H., Plaat, A., Eds.; Series Title: Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2014; Volume 8427, pp. 60–71.
24. Gao, C.; Hayward, R.; Muller, M. Move Prediction Using Deep Convolutional Neural Networks in Hex. *IEEE Trans. Games* 2018, 10, 336–343.
25. Mannen, H.; Wiering, M. Learning To Play Chess Using Reinforcement Learning with Database Games. Master's Thesis, Cognitive Artificial Intelligence Utrecht University, Utrecht, The Netherlands, October 2003.
26. Block, M.; Bader, M.; Tapia, E.; Ramírez, M.; Gunnarsson, K.; Cuevas, E.; Zaldivar, D.; Rojas, R. Using Reinforcement Learning in Chess Engines. *Res. Comput. Sci.* 2008, 35, 31–40.
27. Baxter, J.; Tridgell, A.; Weaver, L. KnightCap: A chess program that learns by combining TD(λ) with game-tree search. arXiv 1999, arXiv:cs/9901002.

28. Lai, M. Giraffe: Using Deep Reinforcement Learning to Play Chess. arXiv 2008, arXiv:1509.01549.
29. Silver, D.; Huang, A.; Maddison, C.J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. Mastering the game of Go with deep neural networks and tree search. *Nature* 2016, 529, 484–489.
30. Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. Mastering the game of Go without human knowledge. *Nature* 2017, 550, 354–359.
31. Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 2018, 362, 1140–1144.
32. Kimura, T.; Sakamoto, K.; Sogabe, T. Development of AlphaZero-based Reinforcement Learning Algorithm for Solving Partially Observable Markov Decision Process (POMDP) Problem. *Bull. Netw. Comput. Syst. Softw.* 2020, 9, 69–73.
33. Davidson, A. Using Artificial Neural Networks to Model Opponents in Texas Hold'em. Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, November 1999. Available online: <http://www.spaz.ca/aaron/poker/nnpoker.pdf> (accessed on 1 July 2021).
34. Davidson, A.; Billings, D.; Schaeffer, J.; Szafron, D. Improved Opponent Modeling in Poker. Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, January 2002. Available online: <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.480.8603&rep=rep1&type=pdf> (accessed on 1 July 2021).
35. Felix, D.; Reis, L.P. An Experimental Approach to Online Opponent Modeling in Texas Hold'em Poker. In *Advances in Artificial Intelligence—SBIA 2008*; Zaverucha, G., da Costa, A.L., Eds.; Springer: Berlin/Heidelberg, Germany, 2008; pp. 83–92.
36. Southey, F.; Bowling, M.; Larson, B.; Piccione, C.; Burch, N.; Billings, D.; Rayner, D.C. Bayes' Bluff: Opponent Modelling in Poker. arXiv 2012, arXiv:1207.1411.
37. Zinkevich, M.; Johanson, M.; Bowling, M.; Piccione, C. Regret Minimization in Games with Incomplete Information. In *Advances in Neural Information Processing Systems 20*; Platt, J.C., Koller, D., Singer, Y., Roweis, S.T., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2008; pp. 1729–1736.
38. Bowling, M.; Burch, N.; Johanson, M.; Tammelin, O. Heads-up limit hold'em poker is solved. *Science* 2015, 347, 145–149.
39. Brown, N.; Lerer, A.; Gross, S.; Sandholm, T. Deep Counterfactual Regret Minimization. arXiv 2018, arXiv:1811.00164.
40. Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; Bowling, M. DeepStack: Expert-level artificial intelligence in heads-up no-limit poker. *Science* 2017, 356, 508–513.
41. Brown, N.; Sandholm, T. Superhuman AI for multiplayer poker. *Science* 2019, 365, 885–890.
42. Ishii, S.; Fujita, H.; Mitsutake, M.; Yamazaki, T.; Matsuda, J.; Matsuno, Y. A Reinforcement Learning Scheme for a Partially-Observable Multi-Agent Game. *Mach. Learn.* 2005, 59, 31–54.
43. Fujita, H.; Ishii, S. Model-based Reinforcement Learning for Partially Observable Games with Sampling-based State Estimation. *Neural Comput.* 2007, 19, 3051–3087.
44. Sturtevant, N.R.; White, A.M. Feature Construction for Reinforcement Learning in Hearts. In *International Conference on Computers and Games*; Springer: Berlin/Heidelberg, Germany, 2006; pp. 122–134.
45. Wagenaar, M.; Wiering, D.M. Learning to Play the Game of Hearts Using Reinforcement Learning and a Multi-Layer Perceptron. Bachelor's Thesis, University of Groningen, Groningen, The Netherlands, 2017. Available online: <http://fse.studenttheses.ub.rug.nl/id/eprint/15440> (accessed on 11 July 2021).
46. Backhus, J.C.; Nonaka, H.; Yoshikawa, T.; Sugimoto, M. Application of reinforcement learning to the card game Wizard. In *Proceedings of the 2013 IEEE 2nd Global Conference on Consumer Electronics (GCCE)*, Tokyo, Japan, 1–4 October 2013; pp. 329–333.
47. Heinrich, J.; Lanctot, M. Fictitious Self-Play in Extensive-Form Games. In *Proceedings of the 32nd International Conference on Machine Learning*, Lille, France, 6–11 July 2015; pp. 805–813.
48. Heinrich, J.; Silver, D. Deep Reinforcement Learning from Self-Play in Imperfect-Information Games. arXiv 2016, arXiv:1603.01121.
49. Lanctot, M.; Zambaldi, V.; Gruslys, A.; Lazaridou, A.; Tuyls, K.; Perolat, J.; Silver, D.; Graepel, T. A Unified Game-Theoretic Approach to Multiagent Reinforcement Learning. In *Advances in Neural Information Processing Systems 30*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; pp. 4190–4203.

50. Zha, D.; Lai, K.H.; Cao, Y.; Huang, S.; Wei, R.; Guo, J.; Hu, X. RLCARD: A Toolkit for Reinforcement Learning in Card Games. arXiv 2019, arXiv:1910.04376.
51. Shao, K.; Tang, Z.; Zhu, Y.; Li, N.; Zhao, D. A Survey of Deep Reinforcement Learning in Video Games. arXiv 2019, arXiv:1912.10944.
52. Bellemare, M.G.; Naddaf, Y.; Veness, J.; Bowling, M. The Arcade Learning Environment: An Evaluation Platform for General Agents. *J. Artif. Intell. Res.* 2013, 47, 253–279.
53. Hausknecht, M.; Stone, P. Deep Recurrent Q-Learning for Partially Observable MDPs. arXiv 2017, arXiv:1507.06527.
54. Zhu, P.; Li, X.; Poupart, P.; Miao, G. On Improving Deep Reinforcement Learning for POMDPs. arXiv 2018, arXiv:1704.07978.
55. Aytar, Y.; Pfaff, T.; Budden, D.; Paine, T.; Wang, Z.; de Freitas, N. Playing hard exploration games by watching YouTube. In *Advances in Neural Information Processing Systems 31*; Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2018; pp. 2930–2941.
56. Kempka, M.; Wydmuch, M.; Runc, G.; Toczek, J.; Jaskowski, W. ViZDoom: A Doom-based AI Research Platform for Visual Reinforcement Learning. In *Proceedings of the 2016 IEEE Conference on Computational Intelligence and Games (CIG)*, Santorini, Greece, 20–23 September 2016.
57. Lample, G.; Chaplot, D.S. Playing FPS Games with Deep Reinforcement Learning. arXiv 2016, arXiv:1609.05521.
58. Papoudakis, G.; Chatzidimitriou, K.C.; Mitkas, P.A. Deep Reinforcement Learning for Doom using Unsupervised Auxiliary Tasks. arXiv 2018, arXiv:1807.01960.
59. Akimov, D.; Makarov, I. Deep Reinforcement Learning with ViZDoom First-Person Shooter. *CEUR Workshop Proc.* 2019, 2479, 3–17.
60. Wu, Y.; Tian, Y. Training Agent for First-Person Shooter Game with Actor-Critic Curriculum Learning. In *Proceedings of the International Conference on Learning Representations*, Toulon, France, 24–26 April 2017.
61. Bengio, Y.; Louradour, J.; Collobert, R.; Weston, J. Curriculum Learning. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, Montreal, QC, Canada, 14–18 June 2009; Association for Computing Machinery: New York, NY, USA, 2009; pp. 41–48.
62. Parisotto, E.; Salakhutdinov, R. Neural Map: Structured Memory for Deep Reinforcement Learning. arXiv 2017, arXiv:1702.08360.
63. Shao, K.; Zhao, D.; Li, N.; Zhu, Y. Learning Battles in ViZDoom via Deep Reinforcement Learning. In *Proceedings of the 2018 IEEE Conference on Computational Intelligence and Games (CIG)*, Maastricht, The Netherlands, 14–17 August 2018; pp. 1–4.
64. Koutník, J.; Schmidhuber, J.; Gomez, F. Online Evolution of Deep Convolutional Network for Vision-Based Reinforcement Learning. In *Proceedings of the International Conference on Simulation of Adaptive Behavior*, Aberystwyth, UK, 23–26 August 2014; pp. 260–269.
65. Koutník, J.; Schmidhuber, J.; Gomez, F. Evolving deep unsupervised convolutional networks for vision-based reinforcement learning. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation—GECCO '14*, Vancouver, BC, Canada, 12–16 July 2014; ACM Press: Vancouver, BC, Canada, 2014; pp. 541–548.
66. Wang, S.; Jia, D.; Weng, X. Deep Reinforcement Learning for Autonomous Driving. arXiv 2019, arXiv:1811.11329.
67. Sallab, A.E.; Abdou, M.; Perot, E.; Yogamani, S. End-to-End Deep Reinforcement Learning for Lane Keeping Assist. arXiv 2016, arXiv:1612.04340.
68. Liu, K.; Wan, Q.; Li, Y. A Deep Reinforcement Learning Algorithm with Expert Demonstrations and Supervised Loss and its application in Autonomous Driving. In *Proceedings of the 2018 37th Chinese Control Conference (CCC)*, Wuhan, China, 25–27 July 2018; pp. 2944–2949.
69. Verma, A.; Murali, V.; Singh, R.; Kohli, P.; Chaudhuri, S. Programmatically Interpretable Reinforcement Learning. arXiv 2019, arXiv:1804.02477.
70. Sharma, S.; Lakshminarayanan, A.S.; Ravindran, B. Fine grained action repetition for deep reinforcement learning. arXiv 2017, arXiv:1702.06054.
71. Gao, Y.; Xu, H.; Lin, J.; Yu, F.; Levine, S.; Darrell, T. Reinforcement Learning from Imperfect Demonstrations. arXiv 2018, arXiv:1802.05313.
72. Mazumder, S.; Liu, B.; Wang, S.; Zhu, Y.; Liu, L.; Li, J. Action Permissibility in Deep Reinforcement Learning and Application to Autonomous Driving. In *KDD'18 Deep Learning Day*; ACM: London, UK, 2018.

73. Li, D.; Zhao, D.; Zhang, Q.; Chen, Y. Reinforcement Learning and Deep Learning based Lateral Control for Autonomous Driving. *arXiv* 2018, arXiv:1810.12778.
74. Zhu, Y.; Zhao, D. Driving Control with Deep and Reinforcement Learning in The Open Racing Car Simulator. In *Neural Information Processing*; Cheng, L., Leung, A.C.S., Ozawa, S., Eds.; Springer International Publishing: Cham, Switzerland, 2018; pp. 326–334.
75. Johnson, M.; Hofmann, K.; Hutton, T.; Bignell, D. The Malmo Platform for Artificial Intelligence Experimentation. *IJCAI*. 2016. pp. 4246–4247. Available online: <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1045.1281&rep=rep1&type=pdf> (accessed on 1 July 2021).
76. Guss, W.H.; Houghton, B.; Topin, N.; Wang, P.; Codel, C.; Veloso, M.; Salakhutdinov, R. MineRL: A Large-Scale Dataset of Minecraft Demonstrations. *arXiv* 2019, arXiv:1907.13440.
77. Liu, L.T.; Dogan, U.; Hofmann, K. Decoding multitask DQN in the world of Minecraft. In *Proceedings of the 13th European Workshop on Reinforcement Learning (EWRL)*, Barcelona, Spain, 3–4 December 2016; p. 11.
78. Oh, J.; Chockalingam, V.; Singh, S.P.; Lee, H. Control of Memory, Active Perception, and Action in Minecraft. *arXiv* 2016, arXiv:1605.09128.
79. Frazier, S. Improving Deep Reinforcement Learning in Minecraft with Action Advice. *arXiv* 2019, arXiv:1908.01007.
80. Xiong, Y.; Chen, H.; Zhao, M.; An, B. HogRider: Champion Agent of Microsoft Malmo Collaborative AI Challenge. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, New Orleans, LA, USA, 2–7 February 2018.
81. Tessler, C.; Givony, S.; Zahavy, T.; Mankowitz, D.J.; Mannor, S. A Deep Hierarchical Approach to Lifelong Learning in Minecraft. *arXiv* 2016, arXiv:1604.07255.
82. Jin, P.H.; Levine, S.; Keutzer, K. Regret Minimization for Partially Observable Deep Reinforcement Learning. *arXiv* 2017, arXiv:1710.11424.
83. Beattie, C.; Leibo, J.Z.; Teplyashin, D.; Ward, T.; Wainwright, M.; Küttler, H.; Lefrancq, A.; Green, S.; Valdés, V.; Sadik, A.; et al. DeepMind Lab. *arXiv* 2016, arXiv:1612.03801.
84. Johnston, L. Surprise Pursuit Auxiliary Task for Deepmind Lab Maze. n.d. p. 7. Available online: <http://cs231n.stanford.edu/reports/2017/pdfs/610.pdf> (accessed on 1 July 2021).
85. Mankowitz, D.J.; Zidek, A.; Barreto, A.; Horgan, D.; Hessel, M.; Quan, J.; Oh, J.; van Hasselt, H.; Silver, D.; Schaul, T. Unicorn: Continual Learning with a Universal, Off-policy Agent. *arXiv* 2018, arXiv:1802.08294.
86. Schmitt, S.; Hudson, J.J.; Zidek, A.; Osindero, S.; Doersch, C.; Czarnecki, W.M.; Leibo, J.Z.; Küttler, H.; Zisserman, A.; Simonyan, K.; et al. Kickstarting Deep Reinforcement Learning. *arXiv* 2018, arXiv:1803.03835.
87. Jaderberg, M.; Czarnecki, W.M.; Dunning, I.; Marris, L.; Lever, G.; Castañeda, A.G.; Beattie, C.; Rabinowitz, N.C.; Morcos, A.S.; Ruderman, A.; et al. Human-level performance in first-person multiplayer games with population-based deep reinforcement learning. *arXiv* 2018, arXiv:1807.01281.
88. Leibo, J.Z.; de Masson d'Autume, C.; Zoran, D.; Amos, D.; Beattie, C.; Anderson, K.; Castañeda, A.G.; Sanchez, M.; Green, S.; Gruslys, A.; et al. Psychlab: A Psychology Laboratory for Deep Reinforcement Learning Agents. *arXiv* 2018, arXiv:1801.08116.
89. Vinyals, O.; Babuschkin, I.; Czarnecki, W.M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D.H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 2019, 575, 350–354.
90. Open, A.I.; Berner, C.; Brockman, G.; Chan, B.; Cheung, V.; Dębiak, P.; Dennison, C.; Farhi, D.; Fischer, Q.; Hashme, S.; et al. Dota 2 with Large Scale Deep Reinforcement Learning. *arXiv* 2019, arXiv:1912.06680.
91. Usunier, N.; Synnaeve, G.; Lin, Z.; Chintala, S. Episodic Exploration for Deep Deterministic Policies: An Application to StarCraft Micromanagement Tasks. *arXiv* 2016, arXiv:1609.02993.
92. Peng, P.; Yuan, Q.; Wen, Y.; Yang, Y.; Tang, Z.; Long, H.; Wang, J. Multiagent Bidirectionally-Coordinated Nets for Learning to Play StarCraft Combat Games. *arXiv* 2017, arXiv:1703.10069.
93. Foerster, J.N.; Nardelli, N.; Farquhar, G.; Torr, P.H.S.; Kohli, P.; Whiteson, S. Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning. *arXiv* 2017, arXiv:1702.08887.
94. Shao, K.; Zhu, Y.; Zhao, D. StarCraft Micromanagement With Reinforcement Learning and Curriculum Transfer Learning. *IEEE Trans. Emerg. Top. Comput. Intell.* 2019, 3, 73–84.
95. Kong, X.; Xin, B.; Liu, F.; Wang, Y. Revisiting the Master-Slave Architecture in Multi-Agent Deep Reinforcement Learning. *arXiv* 2017, arXiv:1712.07305.

96. Rashid, T.; Samvelyan, M.; de Witt, C.S.; Farquhar, G.; Foerster, J.N.; Whiteson, S. QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. arXiv 2018, arXiv:1803.11485.
97. Tang, Z.; Zhao, D.; Zhu, Y.; Guo, P. Reinforcement Learning for Build-Order Production in StarCraft II. In Proceedings of the 2018 Eighth International Conference on Information Science and Technology (ICIST), Cordoba/Granada/Seville, Spain, 30 June–6 July 2018; pp. 153–158.
98. Vinyals, O.; Ewalds, T.; Bartunov, S.; Georgiev, P.; Vezhnevets, A.S.; Yeo, M.; Makhzani, A.; Küttler, H.; Agapiou, J.P.; Schrittwieser, J.; et al. StarCraft II: A New Challenge for Reinforcement Learning. arXiv 2017, arXiv:1708.04782.
99. Zambaldi, V.F.; Raposo, D.; Santoro, A.; Bapst, V.; Li, Y.; Babuschkin, I.; Tuyls, K.; Reichert, D.P.; Lillicrap, T.P.; Lockhart, E.; et al. Relational Deep Reinforcement Learning. arXiv 2018, arXiv:1806.01830.
100. Sun, P.; Sun, X.; Han, L.; Xiong, J.; Wang, Q.; Li, B.; Zheng, Y.; Liu, J.; Liu, Y.; Liu, H.; et al. TStarBots: Defeating the Cheating Level Builtin AI in StarCraft II in the Full Game. arXiv 2018, arXiv:1809.07193.
101. Lee, D.; Tang, H.; Zhang, J.O.; Xu, H.; Darrell, T.; Abbeel, P. Modular Architecture for StarCraft II with Deep Reinforcement Learning. arXiv 2018, arXiv:1811.03555.
102. Pang, Z.J.; Liu, R.Z.; Meng, Z.Y.; Zhang, Y.; Yu, Y.; Lu, T. On reinforcement learning for full-length game of starcraft. In Proceedings of the AAAI Conference on Artificial Intelligence, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 4691–4698.
103. Argall, B.D.; Chernova, S.; Veloso, M.; Browning, B. A survey of robot learning from demonstration. *Robot. Auton. Syst.* 2009, 57, 469–483.
104. Deisenroth, M.P. A Survey on Policy Search for Robotics. *Found. Trends Robot.* 2013, 2, 1–142.
105. Kober, J.; Peters, J. Reinforcement Learning in Robotics: A Survey. In *Learning Motor Skills: From Algorithms to Robot Experiments*; Springer International Publishing: Cham, Switzerland, 2014; pp. 9–67.
106. Tai, L.; Zhang, J.; Liu, M.; Boedecker, J.; Burgard, W. A Survey of Deep Network Solutions for Learning Control in Robotics: From Reinforcement to Imitation. arXiv 2015, arXiv:1612.07139.
107. Todorov, E.; Erez, T.; Tassa, Y. Mujoco: A physics engine for model-based control. In Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vilamoura/Algarve, Portugal, 7–12 October 2012; pp. 5026–5033.
108. Nguyen, H.; La, H. Review of Deep Reinforcement Learning for Robot Manipulation. In Proceedings of the 2019 Third IEEE International Conference on Robotic Computing (IRC), Naples, Italy, 25–27 February 2019; pp. 590–595.
109. Peters, J.; Schaal, S. Reinforcement learning of motor skills with policy gradients. *Neural Netw.* 2008, 21, 682–697.
110. Peters, J.; Schaal, S. Natural actor-critic. *Neurocomputing* 2008, 71, 1180–1190.
111. Theodorou, E.; Buchli, J.; Schaal, S. Reinforcement learning of motor skills in high dimensions: A path integral approach. In Proceedings of the 2010 IEEE International Conference on Robotics and Automation, Anchorage, AK, USA, 3–7 May 2010; pp. 2397–2403.
112. Peters, J.; Mülling, K.; Altun, Y. Relative entropy policy search. In Proceedings of the AAAI Conference on Artificial Intelligence, Atlanta, GA, USA, 11–15 July 2010; Volume 10, pp. 1607–1612.
113. Kalakrishnan, M.; Righetti, L.; Pastor, P.; Schaal, S. Learning force control policies for compliant manipulation. In Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, San Francisco, CA, USA, 25–30 September 2011; pp. 4639–4644.
114. Nair, A.; McGrew, B.; Andrychowicz, M.; Zaremba, W.; Abbeel, P. Overcoming Exploration in Reinforcement Learning with Demonstrations. In Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 21–25 May 2018; pp. 6292–6299.
115. Kumar, V.; Tassa, Y.; Erez, T.; Todorov, E. Real-time behaviour synthesis for dynamic hand-manipulation. In Proceedings of the 2014 IEEE International Conference on Robotics and Automation (ICRA), Hong Kong, China, 31 May–7 June 2014; pp. 6808–6815.
116. Kumar, V.; Todorov, E.; Levine, S. Optimal control with learned local models: Application to dexterous manipulation. In Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA), Stockholm, Sweden, 16–21 May 2016; pp. 378–383.
117. Rajeswaran, A.; Kumar, V.; Gupta, A.; Schulman, J.; Todorov, E.; Levine, S. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. arXiv 2017, arXiv:1709.10087.
118. Haarnoja, T.; Pong, V.; Zhou, A.; Dalal, M.; Abbeel, P.; Levine, S. Composable Deep Reinforcement Learning for Robotic Manipulation. In Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 21–25 May 2018; pp. 6244–6251.

119. Zhu, H.; Gupta, A.; Rajeswaran, A.; Levine, S.; Kumar, V. Dexterous Manipulation with Deep Reinforcement Learning: Efficient, General, and Low-Cost. In Proceedings of the 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019; pp. 3651–3657.
120. Huang, S.H.; Zambelli, M.; Kay, J.; Martins, M.F.; Tassa, Y.; Pilarski, P.M.; Hadsell, R. Learning Gentle Object Manipulation with Curiosity-Driven Deep Reinforcement Learning. arXiv 2019, arXiv:1903.08542.
121. Levine, S.; Finn, C.; Darrell, T.; Abbeel, P. End-to-End Training of Deep Visuomotor Policies. arXiv 2015, arXiv:1504.00702.
122. Pinto, L.; Gupta, A. Supersizing Self-supervision: Learning to Grasp from 50K Tries and 700 Robot Hours. arXiv 2015, arXiv:1509.06825.
123. Levine, S.; Pastor, P.; Krizhevsky, A.; Quillen, D. Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection. arXiv 2016, arXiv:1603.02199.
124. Popov, I.; Heess, N.; Lillicrap, T.P.; Hafner, R.; Barth-Maron, G.; Vecerík, M.; Lampe, T.; Tassa, Y.; Erez, T.; Riedmiller, M.A. Data-efficient Deep Reinforcement Learning for Dexterous Manipulation. arXiv 2017, arXiv:1704.03073.
125. Kalashnikov, D.; Irpan, A.; Pastor, P.; Ibarz, J.; Herzog, A.; Jang, E.; Quillen, D.; Holly, E.; Kalakrishnan, M.; Vanhoucke, V.; et al. QT-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation. arXiv 2018, arXiv:1806.10293.
126. Gu, S.; Holly, E.; Lillicrap, T.; Levine, S. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In Proceedings of the 2017 IEEE International Conference on Robotics and Automation (ICRA), Marina Bay Sands, Singapore, Singapore, 29 May–3 June 2017; pp. 3389–3396.
127. Tsurumine, Y.; Cui, Y.; Uchibe, E.; Matsubara, T. Deep reinforcement learning with smooth policy update: Application to robotic cloth manipulation. *Robot. Auton. Syst.* 2019, 112, 72–83.
128. Jangir, R.; Alenya, G.; Torras, C. Dynamic Cloth Manipulation with Deep Reinforcement Learning. arXiv 2019, arXiv:1910.14475.
129. de Bruin, T.; Kober, J.; Tuyls, K.; Babuška, R. Improved deep reinforcement learning for robotics through distribution-based experience retention. In Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Daejeon, Korea, 9–14 October 2016; pp. 3947–3952.
130. Özalp, R.; Kaymak, C.; Yildirim, Ö.; Ucar, A.; Demir, Y.; Güzeliş, C. An Implementation of Vision Based Deep Reinforcement Learning for Humanoid Robot Locomotion. In Proceedings of the 2019 IEEE International Symposium on INnovations in Intelligent SysTEms and Applications (INISTA), Sofia, Bulgaria, 3–5 July 2019; pp. 1–5.
131. Danel, M. Reinforcement learning for humanoid robot control. POSTER May 2017.
132. Haarnoja, T.; Zhou, A.; Ha, S.; Tan, J.; Tucker, G.; Levine, S. Learning to Walk via Deep Reinforcement Learning. arXiv 2018, arXiv:1812.11103.
133. Christiano, P.F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; Amodei, D. Deep Reinforcement Learning from Human Preferences. In Advances in Neural Information Processing Systems 30; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; pp. 4299–4307.
134. Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; Zaremba, W. Openai gym. arXiv 2016, arXiv:1606.01540.
135. Peng, X.B.; Berseth, G.; Yin, K.; Van De Panne, M. DeepLoco: Dynamic Locomotion Skills Using Hierarchical Deep Reinforcement Learning. *ACM Trans. Graph.* 2017, 36.
136. Berseth, G.; Xie, C.; Cernek, P.; Van de Panne, M. Progressive reinforcement learning with distillation for multi-skilled motion control. arXiv 2018, arXiv:1802.04765.
137. Nagabandi, A.; Kahn, G.; Fearing, R.S.; Levine, S. Neural Network Dynamics for Model-Based Deep Reinforcement Learning with Model-Free Fine-Tuning. In Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 21–25 May 2018; pp. 7559–7566.
138. Duan, Y.; Chen, X.; Houthoofd, R.; Schulman, J.; Abbeel, P. Benchmarking deep reinforcement learning for continuous control. In Proceedings of the International Conference on Machine Learning, New York, NY, USA, 20–22 June 2016; pp. 1329–1338.
139. Xie, Z.; Berseth, G.; Clary, P.; Hurst, J.; van de Panne, M. Feedback Control For Cassie With Deep Reinforcement Learning. In Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Madrid, Spain, 1–5 October 2018; pp. 1241–1246.

140. Tan, J.; Zhang, T.; Coumans, E.; Iscen, A.; Bai, Y.; Hafner, D.; Bohez, S.; Vanhoucke, V. Sim-to-Real: Learning Agile Locomotion For Quadruped Robots. arXiv 2018, arXiv:1804.10332.
141. Huang, Y.; Wei, G.; Wang, Y. V-D D3QN: The Variant of Double Deep Q-Learning Network with Dueling Architecture. In Proceedings of the 2018 37th Chinese Control Conference (CCC), Wuhan, China, 25–27 July 2018; pp. 9130–9135.
142. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J. Human-level control through deep reinforcement learning. *Nature* 2015, 518, 529–533.
143. Peng, X.B.; Berseth, G.; Van de Panne, M. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Trans. Graph. (TOG)* 2016, 35, 1–12.
144. Heess, N.; Wayne, G.; Tassa, Y.; Lillicrap, T.P.; Riedmiller, M.A.; Silver, D. Learning and Transfer of Modulated Locomotor Controllers. arXiv 2016, arXiv:1610.05182.
145. Haarnoja, T.; Tang, H.; Abbeel, P.; Levine, S. Reinforcement Learning with Deep Energy-Based Policies. arXiv 2017, arXiv:1702.08165.
146. Pinto, L.; Davidson, J.; Sukthankar, R.; Gupta, A. Robust Adversarial Reinforcement Learning. arXiv 2017, arXiv:1703.02702.
147. Verma, S.; Novati, G.; Koumoutsakos, P. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proc. Natl. Acad. Sci. USA* 2018, 115, 5849–5854.
148. Zhang, M.; Geng, X.; Bruce, J.; Caluwaerts, K.; Vespignani, M.; SunSpiral, V.; Abbeel, P.; Levine, S. Deep reinforcement learning for tensegrity robot locomotion. In Proceedings of the 2017 IEEE International Conference on Robotics and Automation (ICRA), Singapore, Singapore, 29 May–3 June 2017; pp. 634–641.
149. Luo, J.; Edmunds, R.; Rice, F.; Agogino, A.M. Tensegrity Robot Locomotion Under Limited Sensory Inputs via Deep Reinforcement Learning. In Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 21–25 May 2018; pp. 6260–6267.
150. Wikipedia. Robotics Simulator. Available online: https://en.wikipedia.org/wiki/Robotics_simulator (accessed on 1 July 2021).
151. Wang, W.Y.; Li, J.; He, X. Deep reinforcement learning for NLP. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts, Melbourne, Australia, 15–20 July 2018; pp. 19–21.
152. Sharma, A.R.; Kaushik, P. Literature survey of statistical, deep and reinforcement learning in natural language processing. In Proceedings of the 2017 International Conference on Computing, Communication and Automation (ICCCA), Greater Noida, India, 5–6 May 2017; pp. 350–354.
153. Xiong, C. Recent Progress in Deep Reinforcement Learning for Computer Vision and NLP. In Proceedings of the 2017 Workshop on Recognizing Families in the Wild, RFIW '17, New York, NY, USA, 27 October 2017; Association for Computing Machinery. p. 1.
154. Lapan, M. Deep Reinforcement Learning Hands-On: Apply Modern RL Methods, with Deep Q-Networks, Value Iteration, Policy Gradients, TRPO, AlphaGo Zero and More; Packt Publishing Ltd.: Birmingham, UK, 2018.
155. Gao, J.; Galley, M.; Li, L. Neural approaches to conversational AI. In Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, Ann Arbor, MI, USA, 8–12 July 2018; pp. 1371–1374.
156. Wu, L.; Tian, F.; Qin, T.; Lai, J.; Liu, T. A Study of Reinforcement Learning for Neural Machine Translation. arXiv 2018, arXiv:1808.08866.
157. He, D.; Xia, Y.; Qin, T.; Wang, L.; Yu, N.; Liu, T.Y.; Ma, W.Y. Dual Learning for Machine Translation. In *Advances in Neural Information Processing Systems 29*; Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2016; pp. 820–828.
158. Satija, H.; Pineau, J. Simultaneous machine translation using deep reinforcement learning. In Proceedings of the ICML 2016 Workshop on Abstraction in Reinforcement Learning, International Conference on Machine Learning, New York, NY, USA, 19–24 June 2016.
159. Klein, G.; Kim, Y.; Deng, Y.; Senellart, J.; Rush, A.M. OpenNMT: Open-Source Toolkit for Neural Machine Translation. arXiv 2017, arXiv:1701.02810.
160. Zhang, J.; Ding, Y.; Shen, S.; Cheng, Y.; Sun, M.; Luan, H.; Liu, Y. THUMT: An Open Source Toolkit for Neural Machine Translation. arXiv 2017, arXiv:1706.06415.
161. Yang, Z.; Chen, W.; Wang, F.; Xu, B. Improving Neural Machine Translation with Conditional Sequence Generative Adversarial Nets. arXiv 2017, arXiv:1703.04887.

162. Wu, L.; Xia, Y.; Tian, F.; Zhao, L.; Qin, T.; Lai, J.; Liu, T.Y. Adversarial neural machine translation. In Proceedings of the 10th Asian Conference on Machine Learning, PMLR, Beijing, China, 14–16 November 2018; pp. 534–549.
163. Jurafsky, D.; Martin, J.H. Speech & Language Processing—Third Edition Draft; Pearson Education India: New Delhi, India, 2019.
164. Cuayáhuítl, H.; Renals, S.; Lemon, O.; Shimodaira, H. Evaluation of a hierarchical reinforcement learning spoken dialogue system. *Comput. Speech Lang.* 2010, 24, 395–429.
165. Cuayáhuítl, H.; Dethlefs, N. Spatially-aware dialogue control using hierarchical reinforcement learning. *Acm Trans. Speech Lang. Process. (TSLP)* 2011, 7, 1–26.
166. Cuayáhuítl, H.; Kruijff-Korbyová, I.; Dethlefs, N. Nonstrict hierarchical reinforcement learning for interactive systems and robots. *Acm Trans. Interact. Intell. Syst. (TiiS)* 2014, 4, 1–30.
167. Cuayáhuítl, H.; Keizer, S.; Lemon, O. Strategic Dialogue Management via Deep Reinforcement Learning. *arXiv* 2015, arXiv:1511.08099.
168. Cuayáhuítl, H. SimpleDS: A Simple Deep Reinforcement Learning Dialogue System. In *Dialogues with Social Robots: Enablements, Analyses, and Evaluation*; Jokinen, K., Wilcock, G., Eds.; Springer Singapore: Singapore, 2017; pp. 109–118.
169. Cuayáhuítl, H.; Yu, S. Deep reinforcement learning of dialogue policies with less weight updates. In *Proceedings of the Interspeech 2017*, Stockholm, Sweden, 20–24 August 2017.
170. Cuayáhuítl, H.; Lee, D.; Ryu, S.; Cho, Y.; Choi, S.; Indurthi, S.; Yu, S.; Choi, H.; Hwang, I.; Kim, J. Ensemble-based deep reinforcement learning for chatbots. *Neurocomputing* 2019, 366, 118–130.
171. Jaques, N.; Ghandeharioun, A.; Shen, J.H.; Ferguson, C.; Lapedriza, À.; Jones, N.; Gu, S.; Picard, R.W. Way Off-Policy Batch Deep Reinforcement Learning of Implicit Human Preferences in Dialog. *arXiv* 2019, arXiv:1907.00456.
172. Ranzato, M.; Chopra, S.; Auli, M.; Zaremba, W. Sequence level training with recurrent neural networks. *arXiv* 2015, arXiv:1511.06732.
173. Williams, R.J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.* 1992, 8, 229–256.
174. Bahdanau, D.; Brakel, P.; Xu, K.; Goyal, A.; Lowe, R.; Pineau, J.; Courville, A.; Bengio, Y. An actor-critic algorithm for sequence prediction. *arXiv* 2016, arXiv:1607.07086.
175. Yu, L.; Zhang, W.; Wang, J.; Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, San Francisco, CA, USA, 4–9 February 2017.
176. Li, Z.; Jiang, X.; Shang, L.; Li, H. Paraphrase Generation with Deep Reinforcement Learning. *arXiv* 2017, arXiv:1711.00279.
177. Zhang, X.; Lapata, M. Sentence simplification with deep reinforcement learning. *arXiv* 2017, arXiv:1703.10931.
178. Li, J.; Monroe, W.; Ritter, A.; Jurafsky, D.; Galley, M.; Gao, J. Deep Reinforcement Learning for Dialogue Generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Austin, TX, USA, 1–4 November 2016.
179. Sutskever, I.; Vinyals, O.; Le, Q.V. Sequence to sequence learning with neural networks. *arXiv* 2014, arXiv:1409.3215.
180. Dhingra, B.; Li, L.; Li, X.; Gao, J.; Chen, Y.; Ahmed, F.; Deng, L. End-to-End Reinforcement Learning of Dialogue Agents for Information Access. *arXiv* 2016, arXiv:1609.00777.
181. Keizer, S.; Guhe, M.; Cuayáhuítl, H.; Efstathiou, I.; Engelbrecht, K.P.; Dobre, M.; Lascarides, A.; Lemon, O. Evaluating Persuasion Strategies and Deep Reinforcement Learning Methods for Negotiation Dialogue Agents. *ACL*. In *Proceedings of the 15th Conference of the European chapter of the Association for Computational Linguistics*, Valencia, Spain, 3–7 April 2017.
182. Ilievski, V.; Musat, C.; Hossmann, A.; Baeriswyl, M. Goal-Oriented Chatbot Dialog Management Bootstrapping with Transfer Learning. *arXiv* 2018, arXiv:1802.00500.
183. Ciou, P.H.; Hsiao, Y.T.; Wu, Z.Z.; Tseng, S.H.; Fu, L.C. Composite reinforcement learning for social robot navigation. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, 1–5 October 2018; pp. 2553–2558.
184. Das, A.; Kottur, S.; Moura, J.M.; Lee, S.; Batra, D. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*, Venice, Italy, 22–29 October 2017; pp. 2951–2960.

185. Strub, F.; De Vries, H.; Mary, J.; Piot, B.; Courville, A.; Pietquin, O. End-to-end optimization of goal-driven and visually grounded dialogue systems. arXiv 2017, arXiv:1703.05423.
186. Liu, S.; Zhu, Z.; Ye, N.; Guadarrama, S.; Murphy, K. Improved Image Captioning via Policy Gradient Optimization of SPIDER. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017.
187. Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 2048–2057.
188. Lu, J.; Xiong, C.; Parikh, D.; Socher, R. Knowing When to Look: Adaptive Attention via A Visual Sentinel for Image Captioning. arXiv 2016, arXiv:1612.01887.
189. Ren, Z.; Wang, X.; Zhang, N.; Lv, X.; Li, L.J. Deep Reinforcement Learning-Based Image Captioning With Embedding Reward. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017.
190. Wang, X.; Chen, W.; Wang, Y.; Wang, W.Y. No Metrics Are Perfect: Adversarial Reward Learning for Visual Storytelling. arXiv 2018, arXiv:1804.09160.
191. Wang, X.; Chen, W.; Wu, J.; Wang, Y.F.; Wang, W.Y. Video Captioning via Hierarchical Reinforcement Learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18–22 June 2018.
192. Pasunuru, R.; Bansal, M. Reinforced video captioning with entailment rewards. arXiv 2017, arXiv:1708.02300.
193. Liang, X.; Lee, L.; Xing, E.P. Deep variation-structured reinforcement learning for visual relationship and attribute detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 848–857.

Retrieved from <https://encyclopedia.pub/entry/history/show/28763>