DOWNLOAD PDF REINFORCEMENT LEARNING IN ROBOTICS A SURVEY FILETYPE

Chapter 1 : [] Emotion in Reinforcement Learning Agents and Robots: A Survey

in reinforcement learning the designer of a control task â^— Bielefeld University, CoR-Lab Research Institute for Cognition and Robotics, UniversitÃ¤tsstr. 25, Bielefeld, Ger-.

Besides, there seems to be very little resources detailing how RL is applied in different industries. Reinforcement Learning is a very general framework for learning sequential decision making tasks. And Deep Learning, on the other hand, is of course the best set of algorithms we have to learn representations. And combinations of these two different models is the best answer so far we have in terms of learning very good state representations of very challenging tasks that are not just for solving toy domains but actually to solve challenging real world problems. The rest of the article is organized as follows. Section I is a general introduction. Section II presents the applications of RL in different domains and a brief description of how it was applied. Section VI is conclusion. Introduction to Reinforcement Learning RL, known as a semi-supervised learning model in machine learning, is a technique to allow an agent to take actions and interact with an environment so as to maximize the total rewards. An Introduction Imagine a baby is given a TV remote control at your home environment. Then the curious baby will take certain actions like hitting the remote control action and observe how would the TV response next state. As a non-responding TV is dull, the baby dislike it receiving a negative reward and will take less actions that will lead to such a result updating the policy and vice versa. The study of RL is to construct a mathematical framework to solve the problems. For example, to find a good policy we could use valued-based methods like Q-learning to measure how good an action is in a particular state or policy-based methods to directly find out what actions to take under different states without knowing how good the actions are. However, the problems we face in the real world can be extremely complicated in many different ways and therefore a typical RL algorithm has no clue to solve. For example, the state space is very large in the game of GO, environment cannot be fully observed in Poker game and there are lots of agents interact with each other in the real world. Researchers have invented methods to solve some of the problems by using deep neural network to model the desired policies, value functions or even the transition models, which therefore is called Deep Reinforcement Learning. There are lots of good stuffs about RL online and interested readers can visit awesome-rl , argmin and dennybritz. Applications This part is written for general readers. At the same time, it will be of greater value for readers with some knowledge about RL. Resources management in computer clusters Designing algorithms to allocate limited resources to different tasks is challenging and requires human-generated heuristics. State space was formulated as the current resources allocation and the resources profile of jobs. For action space, they used a trick to allow the agent to choose more than one action at each time step. Then they combined REINFORCE algorithm and baseline value to calculate the policy gradients and find the best policy parameters that give the probability distribution of actions to minimize the objective. Click here to view the code on Github. Tested only on simulated environment though, their methods showed superior results than traditional methods and shed a light on the potential uses of multi-agent RL in designing traffic system. Five agents were put in the five-intersection traffic network, with a RL agent at the central intersection to control traffic signaling. The state was defined as eight-dimensional vector with each element representing the relative traffic flow of each lane. Eight choices were available to the agent, each representing a phase combination, and the reward function was defined as reduction in delay compared with previous time step. Robotics There are tremendous work on applying RL in Robotics. Readers are referred to [10] for a survey of RL in Robotics. The RL component was the guided policy search to generate training data that came from its own state distribution. Demo of the paper. Web System Configuration There are more than configurable parameters in a web system and the process of tuning the parameters requires a skilled operator and numerous trail-and-error tests. The reconfiguration process can be formulated as a finite MDP. The authors used the model-free Q-learning algorithm to do the task. Although the authors used some other technique like policy initialization to remedy

the large state space and computational complexity of the problem instead of the potential combinations of RL and neural network, it is believed that the pioneering work has paved the way for future research in this area. Chemistry RL can also be applied in optimizing chemical reactions. The application is a great one to demonstrate how RL can reduce time-consuming and trial-and-error work in a relatively stable environment. Personalized Recommendations Previous work of news recommendations faced several challenges including the rapid changing dynamic of news, users get bored easily and Click Through Rate cannot reflect the retention rate of users. In practice, they constructed four categories of features, namely A user features and B context features as the state features of the environment, and C user-news features and D news features as the action features. The authors also employed other techniques to address other challenging problems, including memory replay, survival models, Dueling Bandit Gradient Descent and so on. Please refer to the paper for details. The details of the implementation are left to users to investigate. Generally speaking, Taobao ad platform is a place for merchants to place a bid in order to display ad to the customers. This could be a multi-agent problem because the merchants are bidding against each other and their actions are interrelated. In the paper, merchants and customers were clustered into different groups to reduce computational complexity. The state space of the agents indicated the cost-revenue status of the agents, action space was the bid continuous , and reward was the revenue caused by the customer cluster. Games RL is so well-known these days because it is the mainstream algorithm used to solve different games and sometimes achieve super-human performance. RL vs linear model vs Human. AlphaGo, trained with countless human games, already achieved super-human performance by using value network and Monte Carlo tree search MCTS in its policy network. Yet, the researchers later on thought back and tried a purer RL approach â€” train it from scratch. Deep Learning More and more attempts to combine RL and other deep learning architecture can be seen recently and they showed impressive results. RL and RNN is another combinations people used to try new idea. Deepmind showed [9] how to use generative models and RL to generate programs. In the model, the adversarially trained agent used the signal as rewards to improve the actions, instead of propagating the gradients to the input space as in the GAN training. Input vs Generated result. What you need to know before applying RL to your problem There are several things needed before RL can be applied: You do not necessarily need to use RL in your problem and sometimes you just cannot use RL. You may want to check if your problem has some of the following characteristics before deciding to use RL: Lots of iterations are needed before a RL algorithm to work. Therefore, a simulated environment that can correctly reflect the real world is needed. You world need to formulate your problem into a MDP. You need to design the state space, action space, reward function and so on. Your agent will do what it is rewarded to do under the constraints. You may not get the results you want if you design the things differently. There are different RL algorithms you can choose and questions to ask yourself. You want to directly find out the policy or you want to learn the value function? You want to go model-free or model-based? Do you need to combine other kinds of deep neural network or methods to solve your problems? To stay objective and fair, you are also warned about the shortcomings of RL and here is a great post about it. Intuitions from other disciplines RL has a very close relationship with psychology, biology and neuroscience. If you think about it, what a RL agent does is just trial-and-error: And this is exactly how human learns to make a decision. Besides, the exploration and exploitation problem, credit assignment problem, attempts to model the environment are also something we face in our everyday life. The Economics theory can also shed some light on RL. In particular, the analysis of multi-agent reinforcement learning MARL can be understood from the perspectives of game theory, which is a research area developed by John Nash to understand the interactions of agents in a system. What could RL possibly achieve in the future RL still has lots of problems and cannot be used easily. Yet, as long as more efforts are put in solving the problems, RL would be influential and impactful in the following ways: Maybe it is too much to say RL can one day evolve into artificial general intelligence AGI , but RL surely has the potential to assist and work with human. Just imagine a robot or a virtual assistant working with you and taking your actions into its considerations to take actions in order to achieve a common goal. Understanding

the consequences of different strategies: Life is so amazing because time will not go back and things just happen once. Yet, sometimes we would like to know how things could be different at least in the short term if I took a different action? Or would Croatia has a greater chance to win the World Cup if the coach used another strategy? Of course, to achieve this we would need to model the environment, transition functions and so on perfectly and also analyse the interactions between the agents, which seems to be impossible at the moment. Conclusion This article just showed some of the examples of RL applications in various industries. They should not limit your RL use case and as always, you should use first principle to understand the nature of RL and your problem. If you are a decision maker of a company, I hope this article is enough to persuade you to rethink about your business and see if RL can be potentially used. If you are a researcher, I hope you would agree with me that although RL still has different shortcomings, it also means it has lots of potentials to improve and lots of research opportunities. What are your thoughts? Can you think of any problem that RL could solve? Resource Management With deep Reinforcement Learning. A reinforcement learning approach to online web systems auto-configuration. In Distributed Computing Systems, Playing atari with deep reinforcement learning. Reinforcement Learning in Robotics: End-to-end Training of Deep Visuomotor Policies. Mastering the game of go with deep neuralnetworks and tree search.

Chapter 2 : Reinforcement learning in robotics: A survey

Reinforcement learning offers to robotics a framework and set of tools for the design of sophisticated and hard-to-engineer behaviors. Conversely, the challenges of robotic problems provide both.

Chapter 3 : Robotics â€“ Towards Data Science

Reinforcement learning offers to robotics a framework and set of tools for the design of sophisticated and hard-to-engineer behaviors. Conversely, the challenges of robotic problems provide both inspiration, impact, and validation for developments in reinforcement learning.

Chapter 4 : Reinforcement Learning in Robotics: A Survey - The Robotics Institute Carnegie Mellon Univer

Reinforcement Learning in the Context of Robotics Robotics as a reinforcement learning domain diï¬€ers considerably from most well-studied reinforcement learning benchmark problems.

Chapter 5 : [] Deep Reinforcement Learning for Robotic Manipulation-The state of the art

Reinforcement Learning Provides feedback in terms of a scalar objective function Measures and evaluates a step's performance Contains a state, action, and policy function.

Chapter 6 : Applications of Reinforcement Learning in Real World

Reinforcement Learning (RL) is a subfield of Machine Learning where an agent learns by interacting with its environment, observing the results of these interactions and receiving a reward (positive or negative) accordingly.

Chapter 7 : CiteSeerX â€” Reinforcement Learning in Robotics: A Survey

18 Reinforcement Learning in Robotics: A Survey (a) OBELIX robot (b) Zebra Zero robot (c) An autonomous blimp (d) Sarcos humanoid DB Fig. This ï¬•gure illustrates robots to which reinforcement learning has been applied.