Automated transfer in reinforcement learning

Summary

Originally, artificial intelligence (AI) was created to comprehend human intellect. Since then, different fields and sub-domains of AI have been established. Among these fields, maybe the closest to humans, is that of reinforcement learning (RL). In RL, agents try to act in potentially unknown environments using only trial and error, a technique that has proven liable in a variety of human learning settings. In human learning however, neutral information and objective thinking is impossible. Almost always, humans make use directly or indirectly of already acquired knowledge to aid learning in a new task. Albeit, most RL agents are tabula rasa learners. If RL is to mimic human learning, then knowledge reuse is essential. Transfer Learning (TL) is that sub-field of AI dedicated to targeting these challenges. Although successful in different applications, the strive to create autonomous transfer agents is still far-fetched. Taylor and Stone [107], define an autonomous transfer agent as having the following three capabilities: (1) appropriate selection of source task(s), (2) successful learning of the relation between the source and target task(s), and (3) effective knowledge transfer between the tasks.

In this dissertation the previous three problems are targeted and different contributions are made. Firstly, two novel and automated approaches to learn the relationships between RL tasks are proposed. On a high level, these novel methods work by discovering a unifying, rich, and descriptive space that is capable of potentially representing both tasks. These spaces are then used as the basis for learning the intertask mapping between different tasks. Secondly, using this learnt relation, effective and efficient knowledge transfer is executed by using state-of-the-art RL algorithms, such as, least squares policy iteration, and fitted Q-iteration. Here, the transferred knowledge is in form of transitions that are used to bias sample-efficient RL algorithms in their action-selection scheme, thus, leading to an increase in the learning performance. Finally, two data driven similarity measures between different source and target RL tasks are presented. Such measures, allow for the automatic determination of appropriate source task(s) to a given target. The measures adopt and extend Deep Belief Networks (DBNs) for issues of effectiveness and robustness. It is shown that these measure are not only capable of discovering different dynamical phases in the same family of systems, but are also able of relating cross domain similarities between different tasks.

Samenvatting

Reinforcement learning is een methode waarin agenten proberen om doormiddel van trial-and-error te handelen in een potentieel onbekende omgeving, een techniek die haar waarde heeft bewezen in een verscheidenheid aan menselijke leertaken. Mensen leren echter zelden aan de hand van puur neutrale informatie en objectief denken. In tegendeel, mensen maken bijna altijd direct of indirect gebruik van bestaande kennis wanneer een nieuwe taak geleerd wordt. Als reinforcement learning tot doel heeft om menselijk leren te benaderen, dan is hergebruik van kennis essentieel. Transfer Learning is het gebied binnen kunstmatige intelligentie dat zich met deze kwestie bezighoudt. Hoewel Transfer Learning reeds succesvol is in verschillende toepassingen, is het streven naar volledig autonome lerende agenten die deze techniek gebruiken nog een open probleem.

Dit proefschrift draagt op meerdere manieren bij aan een oplossing die het automatisch hergebruiken van kennis mogelijk maakt. Ten eerste worden er twee technieken voorgesteld om automatisch de relatie (inter-task mapping) tussen verschillende reinforcement learning taken te leren. Ten tweede wordt laten zien hoe effectieve en efficinte kennisoverdracht kan worden bereikt in geavanceerde reinforcement learning technieken. Tot slot worden er twee maatstaven gepresenteerd om verschillende leertaken met elkaar te vergelijken, waarmee automatisch een geschikte brontaak gevonden kan worden voor een gegeven doeltaak.