Abstract

Ever since the dawn of computing devices, making machines intelligent has been an ambitious dream of humans. While the primary focus in the last decade of research has been to develop data-driven Machine Learning (ML) solutions, the design decisions while modelling both the problem definition and the solution approach have been critical in developing effective ML solutions. One such design choice is to add structure to the problem definition and/or in the solution approach by introducing symbols and their relations, like that in predicate and first-order logic and similarly in object-centric Deep Learning. Various problem definition frameworks and solution approaches in ML leverage structured representations to define and study a class of problems where individual problems are distinct from each other but are closely related to each other. For example, in Markov Logic Networks in Statistical Relational Learning, we can write a single first-order domain (called theory) that can be grounded using constants to create unique domain instances. Similarly, in the case of the task of planning, the structure is leveraged, for example, in Relational Markov Decision Processes, to define a set of (possibly infinite) domain instances such that these are related to each other by sharing the same underlying first-order transition function. In this work, we study the use of structure for learning in three crucial ML tasks: inference, planning and reinforcement learning. Specifically, we focus on the task of learning solutions that can be transferred from one instance of domain to another.

In chapter 3, we study the task of transfer in Lifted inference that reduces the complexity of inference in relational probabilistic models by identifying groups of constants (or atoms) which behave symmetric to each other. We present the first application of lifting rules for marginal-MAP (MMAP). We define a new equivalence class of (logical) variables, called Single Occurrence for MAX (SOM), and show that the solution lies at the extreme with respect to the SOM variables, i.e., predicate groundings differing only in the instantiation of the SOM variables take the same truth value. Further, we define a sub-class *SOM-R* (SOM Reduce) and exploit properties of extreme assignments to show that MMAP inference can be performed by reducing the domain of SOM-R variables to a single constant. We refer to our lifting technique as the *SOM-R* rule for lifted MMAP.

In the chapters 4 and 5, we study how to learn generalized neural policies in Relational Markov Decision Processes (RMDP) such that they can be transferred from one domain instance (variation) to an unseen one in a zero-shot manner. First, in chapter 4, we carefully identify problems in the existing state-of-the-art approach, called *Symbolic Network* (SymNet), that first converts a given instance of RDMP into a graph (called *instance-graph*) and then uses a Graph Neural Network (GNN) based architecture to learn a policy network. We then propose SymNet2.0 that systematically handles these problems in SymNet by introducing a new instance graph and augments the GNN-based policy network to learn better policies than SymNet. Next, in chapter 5, we identify that due to using a GNN of fixed depth, SymNet2.0 struggles to learn policies that exploit long-range dependencies. As a remedy, we propose SymNet3.0, where we first construct a novel *influence graph* characterized by edges capturing one-step influence (dependence) between nodes based on the transition model. We then define *influence distance* between two nodes as the shortest path between them in this graph – a feature we exploit to represent long-range dependencies in the instance-graph of SymNet3.0.

In the final chapter 6, we focus on the task of learning generalized policies in Reinforcement learning that are capable of transferring the neural policy learned on a set of training environments to a set of unseen but related environments. We present *Object Centric Reinforcement Learning Agent (ORLA)*, an object-centric approach for model-free RL in perceptual domains. Given a perceptual input state, ORLA first identifies the set of objects in the scene; it then constructs a symbolic graph constituting various identified objects and their relations. Finally, a Graph Attention Network (GAT) based architecture is employed over the extracted object positions to learn a dense state embedding, which is then decoded to get the final policy that generalizes to unseen environments.

Finally, we conclude the thesis by outlining our contributions and discussing various directions of future work to extend the work presented in the thesis.