Lifting Techniques for Sequential Decision Making and Probabilistic Inference (Extended Abstract)

Ankit Anand
Department of Computer Science & Engineering
Indian Institute of Technology, Delhi
ankit.anand@cse.iitd.ac.in

Abstract

Many traditional AI algorithms fail to scale as the size of state space increases exponentially with the number of features. One way to reduce computation in such scenarios is to reduce the problem size by grouping symmetric states together and then running the algorithm on the reduced problem. The focus of this work is to exploit symmetry in problems of sequential decision making and probabilistic inference. Our recent work- ASAP-UCT defines new State-Action Pair (SAP) symmetries in Markov Decision Processes. We also apply these SAP symmetries in Monte Carlo Tree Search (MCTS) framework. In probabilistic inference, we expand the notion of unconditional symmetries to contextual symmetries and apply them in Markov Chain Monte Carlo (MCMC) methods. In future, we plan to explore interesting links in symmetry exploitation in different problems and aim to develop a generic symmetry based framework.

1 Introduction

Real world Artificial Intelligence(AI) problems like decision making, learning and probabilistic inference tasks must deal with a high degree of complexity in structure and uncertainty. Numerous compact representations like relational, factored and first order representations have been proposed where problem is normally solved by first grounding the objects in formulas and then running the appropriate algorithm on these grounded networks. Since grounded networks become exponential in size of representation, there has been a lot of interest recently in the field of lifting of these grounded networks [Mihalkova and Getoor, 2011].

Lifting a network essentially means identifying a group of nodes in the network which behave identically through out the algorithm and then combining them into a single node. A modified version of the algorithm is then run on this smaller lifted network. We explore lifting techniques in sequential decision making under uncertainty and probabilistic inference.

The problem of sequential decision making, often modeled as a Markov Decision Process (MDP), is a fundamental problem in the design of autonomous agents. Domain abstraction techniques have been used in past to compute symmetries to reduce the original MDP [Li *et al.*, 2006]. This can lead to significant savings in computation, but these have been predominantly implemented for offline planning. In a parallel thread, Monte-Carlo Tree Search (MCTS) algorithms such as UCT are an attractive online framework for large state spaces. However, MCTS search trees are constructed in flat state and action spaces, which can lead to poor policies for very large problems.

Recently, there has been lot of interest in combining these two orthogonal approaches of MCTS and exploiting symmetry. Our current work in this direction [Anand *et al.*, 2015a],[Anand *et al.*, 2015b] and [Anand *et al.*, 2016b] demonstrate the applicability of these approaches. We define the ASAP (Abstraction of State-Action Pairs) framework, which extends and unifies past work on domain abstractions by holistically aggregating *both* states and state-action pairs and uncovers a much larger number of symmetries. Further, we apply these symmetries in UCT algorithm and call the resulting MDP solver ASAP-UCT. ASAP-UCT is a batch algorithm where abstraction(symmetry) computation and UCT expansions operate in batches. This is further improved in OGA-UCT:On-the-Go Abstractions in UCT where abstractions are computed as we build the tree.

On the other hand, there has been a lot of interest in lifted inference where symmetries are exploited for efficient inference in graphical models. Symmetric variables (states) are collapsed into meta-variables (meta-states) and inference algorithms are run over the lifted graphical model instead of the flat one. In most of recent works, the lifting technique is tied to the specific algorithm being considered. Also, there is significant work on understanding the definition of symmetry.

In this light, our recent work [Anand *et al.*, 2016a] extends existing definitions of symmetry by introducing the novel notion of *contextual symmetry*. Two states that are not globally symmetric, can be contextually symmetric under some specific assignment to a subset of variables, referred to as the *context variables*. Contextual symmetry subsumes previous symmetry definitions and can represent a large class of symmetries not representable earlier.We further exploit contextual symmetries in Markov Chain Monte Carlo(MCMC) framework leading to significant gains.

Lastly, there is a strong co-relation between MDP solving methods and probabilistic inference as both of these algorithms depend upon local interactions between neighboring nodes. Under this theme, we would like to explore the possibility of unifying these different problems under a common *symmetry exploitation* framework so that a generic approach for solving these can be developed.

2 Current Progress and Focus

In our current work, we expand the existing notion of symmetries to incorporate many more symmetries present in a problem domain. We also illustrate the use of these new symmetries in state of art solvers. For example in Planning, we introduce Abstraction of State-Action Pairs(ASAP) and focus on MCTS Algorithms for exploiting these abstractions. In probabilistic inference, we introduce the idea of contextual symmetries and apply those in MCMC algorithms.

2.1 ASAP and Abstractions in UCT

ASAP Framework: In [Anand *et al.*, 2015a], we introduce the idea of State-Action Pair (SAP) abstractions in addition to previously defined state abstractions. Any MDP can be represented in the form of AND-OR graph of states and SAP nodes. We define a mutually recursive formulation of abstract state and SAP nodes. We say that two SAP nodes belong to same abstract SAP node if and only if they transition to same abstract states with same probabilities and incur equal costs. Similarly, two states belong to same abstract state if and only if we can define a mapping of their applicable actions leading to same abstract SAP nodes. Here, two state-actions(SAPs) could be still symmetric even though their parent states may be not symmetric. This uncover many more symmetries in domains where there is no state symmetry.
ASAP in UCT: We provide algorithms to apply ASAP abstractions in UCT. ASAP-UCT, our first algorithm, is a batch algorithm which alternates in two phases. One phase consists of an abstraction computation routine that uses the existing UCT tree to induce groups of symmetric nodes called abstract nodes. The second phase is the (modified) UCT algorithm, which uses the abstractions computed in first phase to make use of roll outs more effectively. Further, in our recent work [Anand *et al.*, 2016b], we illustrate that this batch computation does not achieve the full potential of abstractions. Our solution is OGA-UCT in which abstraction computation is tightly integrated into the MCTS algorithm. It also has several desirable properties, including (1) rapid use of new information in modifying existing abstractions, (2) elimination of the expensive batch abstraction computation phase, and (3) focusing abstraction computation on important parts of the sampled search space. Our empirical investigations reveal that ASAP-UCT is 26% better than existing approaches and OGA-UCT further improves upon ASAP-UCT by 18%.

2.2 Contextual Symmetries in Graphical Models:

Most of the existing lifted inference algorithms exploit sets of variables (states) that are *unconditionally* symmetric. We extend the notion of symmetries to *contextual* symmetries [Anand *et al.*, 2016a]: set of states that are symmetric under a given context (variable-value assignment to a subset of variables). We also give an algorithm to compute contextual symmetries through graph isomorphism and show how they can be used for efficient inference in MCMC settings by proposing Contextual MCMC algorithm.

3 Scope and Future Proposals

Finally, we believe that exploiting symmetries could lead to significant improvements in many algorithms not limited to planning and probabilistic inference. Our initial investigation with it and development of ASAP-UCT and Contextual MCMC, clearly show the first step in this direction. Nevertheless, there are significant challenges with-in this, which we would like to address in future.
Advancing Contextual Symmetry Work: Adapting popular inference algorithms like Belief Propagation, MC-SAT to use contextual symmetries could be an important next step in area of probabilistic inference. Even approximate contextual symmetry could be applicable to many real world problems and could lead to huge gains in inference.
Learning Symmetries: With advances in machine learning techniques, an interesting approach to compute abstractions is by learning the symmetries in state space. A recent work by Srinivasan et. al.(2015) suggests the use of nearest neighbor approach to improve exploration in UCT. Learning abstractions could play a critical role in online real time algorithms where computing symmetries could be a bottleneck.
Development of Generic Abstraction Framework: Due to local nature of computation, there is significant overlap of techniques used to exploit symmetries in both planning and probabilistic inference. We intend to study this correlation in detail and wish to develop a generic abstraction framework for both these fields.

Acknowledgements

I am thankful to Mausam, Parag Singla, Aditya Grover and Ritesh Noothigattu for collaborating in this work. I am also grateful to Tata Consultancy Services for supporting me through TCS Research Scholarship program.

References

[Anand *et al.*, 2015a] Ankit Anand, Aditya Grover, Mausam, and Parag Singla. ASAP-UCT: Abstraction of State-Action Pairs in UCT. In *IJCAI*, 2015.

[Anand *et al.*, 2015b] Ankit Anand, Aditya Grover, Mausam, and Parag Singla. A Novel Abstraction Framework for Online Planning. In *AAMAS*, 2015.

[Anand *et al.*, 2016a] Ankit Anand, Aditya Grover, Mausam, and Parag Singla. Contextual Symmetries in Probabilistic Graphical Models. In *IJCAI*, 2016.

[Anand *et al.*, 2016b] Ankit Anand, Ritesh Noothigattu, Mausam, and Parag Singla. OGA-UCT: On-the-Go Abstractions in UCT. In *ICAPS*, 2016.

[Li *et al.*, 2006] Lihong Li, Thomas J Walsh, and Michael L Littman. Towards a Unified Theory of State Abstraction for MDPs. In *ISAIM*, 2006.

[Mihalkova and Getoor, 2011] Lilyana Mihalkova and Lise Getoor. Lifted graphical models: A survey. *CoRR*, abs/1107.4966, 2011.