Tractable relaxations for Limited Memory Influence Diagrams

Victor Cohen[1], Axel Parmentier, Guillaume Obozinski, Vincent Leclere, and Joseph Salmon
[1]CERMICS, École Nationale des Ponts et Chaussées, Marne-la -Vallée
victor.cohen@enpc.fr

Limited Memory Influence Diagrams (LIMIDs) are a flexible probabilistic graphical model framework to express stochastic optimization problems, such as Markov Decision Processes (MDP) and partially observable MDP. The random variables are represented by a probabilistic graphical model whose vertices are partitioned into three types : chance, decision, and utility vertices. The user chooses a policy, i.e., the distribution of the decision vertices conditionally to their parents, in order to maximize the expected utility. The standard algorithm to find an optimal policy, Single Policy Update (SPU), is a minibatch coordinate descent. It iteratively updates the conditional distribution of one decision vertex, the conditional distribution of all other decision vertices being fixed. It is a heuristic and may converge to a local optimum. LIMIDs for which SPU is guaranteed to converge to a global optimum are called soluble, and are the ones for which finding an optimal policy is easy. Only sufficient conditions for a LIMID to be soluble were previously known. We prove a necessary and sufficient condition for a LIMID to be soluble that can be checked in linear time in the size of the graph. We then leverage this condition to propose algorithms building soluble relaxations of non-soluble LIMIDs. These relaxations provide upper-bounds on the value of an optimal solution, as well as good heuristic solutions. Our algorithms enable to strike a balance between the quality of the bound obtained and the computational difficulty of the relaxation we build.