Towards Structural-Patterns Admissible Heuristics

Michael Katz*
Faculty of Industrial Engineering
Technion - Israel Institute of Technology
Haifa, Israel

Abstract

Considering admissible heuristics for sequentially-optimal planning, we suggest and study a generalization of the pattern-database homomorphism abstractions, called structural-patterns abstractions. The basic idea of structural patterns boils down to projecting the problem in hand to provably tractable fragments of optimal planning. The key motivation behind this generalization of PDBs is to alleviate the requirement for the projections to be of a low dimensionality.

Introduction

The difference between various algorithms for planning as heuristic search is mainly in the heuristic functions they define and use. Most typically, an (admissible) heuristic function for domain-independent planning is defined as the (optimal) cost of achieving the goals in an *over-approximating abstraction* of the planning problem in hand.Such an abstraction is obtained by relaxing certain constraints that are present in the specification of the real problem, and the desire is to obtain a *tractable* (that is, solvable in polynomial time), and, at the same time, *informative* abstract problem. The main question is thus: What constraints should we relax to obtain such an effective over-approximating abstraction?

Conceptually, one can distinguish between homomorphism and embedding abstractions, and the former is of our focus in this work. An *homomorphism abstraction* systematically contracts several states to create a single abstract state. Most typically, such a state-gluing is obtained by projecting the original problem onto a subset of its parameters, as if ignoring the constraints that fall outside the projection. Homomorphisms has been successfully explored in the scope of domain-independent pattern database (PDB) heuristics (Edelkamp 2001; Haslum, Bonet, & Geffner 2005), inspired by the (similarly named) problem-specific heuristics for search problems such as (k^2-1)-puzzles, Rubik's Cube, etc. (Culberson & Schaeffer 1998). A core property of the PDB heuristics is that the problem is projected onto a space of small (up to logarithmic) dimensionality so that reachability analysis in that space could be done by exhaustive search. Note that this constraint implies an inherent scalability limitation of the PDB heuristics—as the problems of interest grow, limiting patterns to logarithmic dimensionality will unavoidably make them less and less informative with respect to the original problems.

In this paper we suggest a generalization of the PDB abstractions to what we call *structural patterns*. In itself, the idea of structural patterns is simple, and it corresponds to projecting the original problem to provably tractable fragments of optimal planning. At least theoretically, moving to structural patterns alleviates the requirement for the projections to be of a low dimensionality. To materialize the idea of structural-patterns heuristics, we have started investigating the computational tractability of sequentially-optimal planning, and have already discovered numerous new problem classes for whose such optimization is tractable (Katz & Domshlak 2007b; 2007a). The results are based on exploiting numerous structural and syntactic characteristics of planning problems such as the structure of their causal graphs. Moreover, we have already shown that the idea of structural-patterns heuristics is not of a theoretical interest only—in (Katz & Domshlak 2007a) we suggest a concrete structural patterns abstraction based on decomposing the problem in hand along its causal graph, and show that the induced admissible heuristic can provide more informative estimates than its state-of-the-art alternatives.

From PDBs to Structural Patterns?

Given a problem $\Pi = \langle \mathcal{V}, A, I, G \rangle$, each subset of variables $\mathcal{V}' \subseteq \mathcal{V}$ defines a *pattern abstraction* $\Pi^{[\mathcal{V}']} = \langle \mathcal{V}', A^{[\mathcal{V}']}, I^{[\mathcal{V}']}, G^{[\mathcal{V}']} \rangle$ by intersecting the initial state, the goal, and all the actions' preconditions and effects with $\mathcal{V}'$ (Edelkamp 2001). The idea behind the PDB heuristics is elegantly simple. First, we select a (relatively small) set of subsets $\mathcal{V}_1, \dots, \mathcal{V}_m$ of $\mathcal{V}$ such that,

*Advisor: **Carmel Domshlak**, Technion.

for $1 \leq i \leq m$,

(a) $\Pi^{[\mathcal{V}_i]}$ is an over-approximating abstraction of Π,

(b) the size of $\mathcal{V}_i$ is sufficiently small to perform reachability analysis in $\Pi^{[\mathcal{V}_i]}$ by an (either explicit or symbolic) exhaustive search.

Let $h^{[\mathcal{V}_i]}(s)$ be the optimal cost of achieving the abstract goal $G^{[\mathcal{V}_i]}$ from the abstract state $s^{[\mathcal{V}_i]}$. To obtain an admissible heuristic, if the set of abstract problems $\Pi^{[\mathcal{V}_1]}, \ldots, \Pi^{[\mathcal{V}_k]}$ satisfy certain requirements of disjointness (Felner, Korf, & Hanan 2004; Edelkamp 2001), the PDB heuristic can be set to $h(s) = \sum_{i=1}^{m} h^{[\mathcal{V}_i]}(s)$. Otherwise, one can set $h(s) = \max_{i=1}^{m} h^{[\mathcal{V}_i]}(s)$.

The Achilles heel of the PDB heuristics is that each pattern (that is, each selected subset of variables $\mathcal{V}_i$) is required to be *small* so that reachability analysis in $\Pi^{[\mathcal{V}_i]}$ could be done by exhaustive search. In short, computing $h^{[\mathcal{V}_i]}(s)$ in polynomial time requires satisfying $|\mathcal{V}_i| = O(\log |\mathcal{V}|)$. Note that this constraint implies an inherent scalability limitation of the PDB heuristics—as the problems of interest grow, limiting patterns to logarithmic dimensionality will unavoidably make them less and less informative with respect to the original problems.

However, this is not necessarily the only way to proceed. In principle, given a SAS$^+$ problem $\Pi = \langle \mathcal{V}, A, I, G \rangle$, one can select a (relatively small) set of subsets $\mathcal{V}_1, \ldots, \mathcal{V}_m$ of $\mathcal{V}$ such that, for $1 \leq i \leq m$,

(a) $\Pi^{[\mathcal{V}_i]}$ is an over-approximating abstraction of Π,

(b) *the reachability analysis in* $\Pi^{[\mathcal{V}_i]}$ *is tractable (not necessarily due to the size of but) due to the specific structure of* $\Pi^{[\mathcal{V}_i]}$.

What is important here is that the second requirement can be satisfied even if the size of each selected pattern $\mathcal{V}_i$ is $\Theta(|\mathcal{V}|)$.

A priori, this generalization of the PDB idea to structural patterns is appealing as it allows using patterns of unlimited dimensionality. The pitfall, however, is that such structural patterns correspond to tractable fragments of sequentially-optimal planning, and the palette of such known fragments is extremely limited (Bäckström & Nebel 1995; Bylander 1994; Jonsson & Bäckström 1998; Jonsson 2007). In our work we aim at showing that this palette can still be extended, and such extensions may allow us materializing the idea of structural patterns heuristics.

Research Objectives

Our main objective is to extend and finalize the conceptual framework of structural-patterns admissible heuristics for domain-independent planning, to characterize and study possible instantiations of this framework, their effectiveness and computational efficiency, and to extend the developed heuristic estimates to richer formalisms of domain-independent planning.

Considering the essential ingredients of the structural-patterns framework, currently our primary goal is to *extend the pool of tractable subclasses of optimal classical planning.* It is now apparent that revealing the complexity hierarchy of subclasses of optimal classical planning is still very much an open problem. Our results so far show that progress in this direction is possible. In particular, currently we investigate various special topologies of the causal graphs, along with (practically interesting) local and global restrictions on the problem actions. Specifically, we have started investigating a few special cases of the causal graph structure, namely:

(1) Directed polytrees (poly-forests),

(2) Directed-path singly-connected DAGs, and

(3) Directed-path δ-connected DAGs.

Optimal planning for problems inducing such forms of causal graphs is considered both in general settings, as well as under additional limiting properties such as

(i) actions restricted to only unary effects,

(ii) (prevail) $O(1)$-dependence of the actions,

(iii) $O(1)$-bounded in-degree and/or out-degree of the causal graph, and

(iv) $O(1)$-bounded domains of the state variables.

For unary-effect problems with binary-valued variables, we have recently shown (Katz & Domshlak 2007b) that optimal planning is tractable for the problem fragments characterized by

(1) Directed polytree with $O(1)$-bounded in-degree,

(2) Directed polytree when $O(1)$-dependence of the actions.

In addition, for some minor extensions of these fragments we have shown that optimal planning is NP-hard. The latter results give us a better understanding of the boundaries of the optimal planning tractability.

We have also started investigated optimal planning for problems with multi-valued variables. For such unary-effect problems we have show in (Katz & Domshlak 2007a) that optimal planning is tractable if

(1) the causal graph induces a directed fork with a root r, and either (i) for all $v \in \mathcal{V}$, we have $|Dom(v)| = O(1)$, or (ii) $|Dom(r)| = 2$,

(2) the causal graph induces a directed inverted fork with a root r, and $|Dom(r)| = O(1)$.

Using these results, in (Katz & Domshlak 2007a) we have introduced and looked into a concrete structural patterns abstraction based on decomposing the problem into a set of fork and inverted fork components of its causal graph, combined with abstracting the domains of certain variables within these individual components. Our preliminary analysis shows that the induced admissible heuristic can provide more informative estimates than its state-of-the-art alternatives.

The basic principles of the structural patterns framework motivate further research in numerous directions, and in particular, in (1) discovering new islands of tractability of optimal planning, and (2) translating and/or abstracting the general planning problems into such islands. In our ongoing work we aim at pursuing both these directions by "mining" the tractable fragments of optimal planning, and by performing formal and empirical analysis of alternative schemes for abstracting general planning problems to meet the specification of such islands of tractability. In addition, we plan to start investigating numerous additional issues in using tractable subclasses of optimal planning in homomorphism abstractions for planning as heuristic search. In particular, we plan to devote our efforts to the following research questions.

1. *Optimization of structural patterns selection.* Having established a set of structural classes of planning problems for which optimal planning is tractable, the next step is to formalize the criteria for selecting concrete structural patterns for a given planning problem. The pitfall here is that the number of alternative structural patterns (e.g., the number of different subforests of a given causal graph) can be exponential in the size of the problem description. It is also apparent that some choices of structural patterns will be more informative than the other. First, we plan to provide a concrete formal model for optimizing the outcome of structural patterns selection. Second, we will aim at suggesting some tractable approximations for this optimization problem (as the latter is naturally expected to be NP-hard in itself.)

2. *Optimization of variable domains abstraction.* As we show in (Katz & Domshlak 2007a), selecting a set of structural patterns based on the causal graph decomposition alone might be insufficient (or, at least,informativeness-wise sub-optimal.) In particular, while selecting sub-graphs of the problem's causal graph to form our structural patterns, we might need to further abstract the domains of the multi-valued variables underlying the nodes of the causal graph. For instance, such a domain abstraction will be essential if the problem is described over general multi-valued variables, while the structural patterns are required to be defined over $O(1)$-valued variables. For good and for bad, here as well we have a substantial degree of freedom, and thus ideally we should provide a concrete formal model for optimizing the process of variable domain abstraction. We believe that a substantial progress in this direction can be achieved by bridging between the relevant principles of structural patterns and PDHs, as well as by exploring the structure and interplay between the variables' domain-transition graphs.

3. *Extensions to richer formalisms.* One of the challenging directions we would like to pursue at the later stages of our research is this of extending structural-pattern heuristics to some richer planning formalisms. Extending the classical planning, such formalisms allow problem specification to use numerical variables, temporal actions, delayed effects, etc., and for now it is not clear to what degree the homomorphism abstractions can be useful in dealing with such rich formalisms.

References

Bäckström, C., and Nebel, B. 1995. Complexity results for SAS^+ planning. *Comp. Intell.* 11(4):625–655.

Bylander, T. 1994. The computational complexity of propositional STRIPS planning. *AIJ* 69(1-2):165–204.

Culberson, J., and Schaeffer, J. 1998. Pattern databases. *Comp. Intell.* 14(4):318–334.

Edelkamp, S. 2001. Planning with pattern databases. In *ECP*, 13–34.

Felner, A.; Korf, R. E.; and Hanan, S. 2004. Additive pattern database heuristics. *JAIR* 22:279–318.

Haslum, P.; Bonet, B.; and Geffner, H. 2005. New admissible heuristics for domain-independent planning. In *AAAI*, 1163–1168.

Jonsson, P., and Bäckström, C. 1998. State-variable planning under structural restrictions: Algorithms and complexity. *AIJ* 100(1–2):125–176.

Jonsson, A. 2007. The role of macros in tractable planning over causal graphs. In *IJCAI*, 1936–1941.

Katz, M., and Domshlak, C. 2007a. Structural patterns heuristics. In *submitted to the ICAPS-07 Workshop on Heuristics for Domain-independent Planning: Progress, Ideas, Limitations, Challenges*.

Katz, M., and Domshlak, C. 2007b. Structural patterns of tractable sequentially-optimal planning. In *ICAPS*.