August 12, 2009

Review on Lukáš Chrpa's Dissertation Thesis Entitled Learning for Classical Planning

The main hypothesis elaborated by this thesis is that through algorithmic analysis of automatically constructed plans, one can induce formal knowledge that will help planners produce better plans, using smaller resources when solving future planning problems. The topic falls in the general research area that bootstraps machine learning techniques with optimization algorithms towards systems improving their performance by adapting themselves to particular problem domains. While these efforts have had, as the author correctly observes in the thesis, a rather long history, they have recently experienced a surprising bloom in fields such as propositional satisfiability or constraint satisfaction. Exploring novel techniques in the more special field of classical planning thus represents a hot topic and the timing for the thesis is perfect.

The way that best helped me understand the logical structure of the thesis was to see it as a composition of five building blocks. Two of them are foundational, providing the theoretical bases for the remaining three, that in turn involve heuristic ingredients. The first of the foundational stones is the *action dependency* theory that explores the degrees of freedom inside the structures representing plans. More specifically, the action dependency theory tells us on what conditions one can move individual actions back and forth with respect to other actions in the plan without invalidating it. The second theoretical body is the *entanglement theory* identifying operators that, informally put, are in an especially strong relation to the initial or goal state, a relation that cannot be altered by actions in the plans. I liked a lot both of these theories. While I did not grasp every line of all proofs, the presented algorithms rendering the theory operational made perfect sense to me.

The remaining three building stones serve in turn to achieve the grand goals of the thesis, exploiting the respective theories above. The first of them represents a rather obvious use of the action dependency framework. Specifically, it shows how the theory can reveal actions or action pairs that can be completely removed from the plan. Such tricks, however useful, are just post-processing of produced plans and have nothing to do with learning. This is likely the reason why they are only briefly stated in the thesis and supplemented with no experimental evaluation. The second heuristic component again uses the action dependency theory, this time to construct macro-operators. The compelling and to my knowledge innovative rationale is here as follows. Suppose that an action-grounding of a specific pair of operators is observed together in many plans constructed so far (the 'training plans'). Then it may indeed be a particularly useful combination of operators that likely would occur in future plans as well. There is thus good reason to promote this combination to the planner by defining a new operator combining the two. The innovative aspect of this strategy is that the respective actions do not have to be adjacent in the training plans, as long as the action dependency theory allows us to *make* them adjacent. The third method developed rests on the entanglement theory. By observing operators entangled to the initial or goal state in many plans, the method identifies action-groundings of these operators that will likely not make sense in the completion of future planning tasks. Of interest, the thesis shows a useful trick through which such groundings can be prevented by reformulating the lifted (non-ground) planning task specification. The latter two methods are extensively evaluated in experiments with tasks coming from the International Planning Competition and indicate clear superiority of planners equipped with the developed methods to their baseline variants.

There are two main things I miss in the thesis, which I detail on below. However, as the contributions above are of high quality and alone sufficient for a PhD thesis, my complaints hereafter are to be mostly taken as suggestions for future work.

Firstly, the thesis contains a formalization of the problem *background* and also a good formalization of the *solutions* to the problem. What is missing is a clear definition of the *problem* itself. That should clearly state the types and properties of the expected inputs and those of the desired outputs, commonly encompassing the various specific methods offered in the thesis. Such a problem statement that is now rather implicit would, in my view, be twofold. The simpler version, that does not involve learning, would have a plan on both the input and output part, such that the output plan is in some sense equivalent to the input plan yet shorter. The second problem, involving learning, would deal with operator-set reformulation and its precise statement would of course be more complex. A clear, up-front formulation of the problem would help the reader understand better what and why the author is doing. It would also chart a road map revealing territories that are undefined or ill-defined. For example, one such territory pertains to the expected character of the training set of plans. Training set properties represent a crucial aspect in any machine learning application. As of now, it is not clear whether the input plans are expected to follow any distributional properties (coming from the same planner or a defined mixture thereof?) nor e.g. whether the training plans should follow the same distribution as the testing plans (on which the degree of improvements is being measured). Lastly, the problem formalization would also entail criteria according to which one could decide which method is better in defined situations, either through measurements or possibly by providing theoretical bounds.

The second item on my wishlist pertains to the learning ingredient of the thesis, which is undoubtedly essential but which, unfortunately, sources very little from existing machine learning techniques. To remain brief, I will provide two examples where it would be beneficial to take a closer look at what machine learning theory has to offer. Firstly, both of the learning approaches (operator assembly and entanglement detection) need to set certain parameters such as frequency thresholds, and do so in an ad hoc manner. Here, the proverbial *lesson one* of statistical machine learning would rightly dictate to tune such parameters by cross-validation on the initial training set. Secondly, the thesis relies on the important machine learning technique of *generalization* but does so in a crude way and only implicitly, without proposing an explicit generality structure. For example, if certain action-pairs are found frequent, it is inductively inferred that their common generalization (the subsuming operator-pair) is characteristic for the plans typically produced in the domain. What I found crude about this is that this is only a two-step generalization whereas the generality spectrum between the fully grounded actions on one hand and the fully variabilized operators on the other hand obviously contains much more than just the two extremes and the empirically-best generalizer may indeed be somewhere in between of them. Generalization can obviously also go along axis other than variabilization of constants. For example, one could consider looking for frequent n-tuples of actions which could be iteratively generalized to $(n-1)$-tuples until a frequency threshold is met. In fact, the field of *inductive logic programming* provides detailed recipes on how the logical structures in question can be organized by generality into systematically searchable lattices.

Despite my two points of criticism above, which admittedly may be biased due to my own inclination to machine learning, I consider this a high-quality thesis with useful results, and fully fit for a defense.

Filip Železný (zelezny@fel.cvut.cz)
Czech Technical University
Faculty of Electrical Engineering
Department of Cybernetics