The Knowledge Engineering Review, Vol. 00:0, 1–4. c 2005, Cambridge University Press
DOI: 10.1017/S0000000000000000 Printed in the United Kingdom.

Case-based planning

MICHAEL T. COX[1], HECTOR MUNOZ-AVILA[2] and RALPH BERGMANN[3]
[1]*BBN Technologies, Cambridge, MA 02138;*
e-mail: mcox@bbn.com
[2]*Department of Computer Science and Engineering, Lehigh University, Bethlehem, PA 18015;*
e-mail: munoz@cse.lehigh.edu
[3]*University of Trier, Business Information Systems II, 54286 Trier, Germany;*
e-mail: bergmann@uni-trier.de

Abstract

We briefly examine case-based planning starting with the seminal work of Hammond. Derivational analogy represents an important shift of technical emphasis that helped mature the techniques. The choice of abstraction level is equally important. We conclude by discussing theoretical underpinnings and by providing some pointers to current directions.

1 Introduction

From the beginning (e.g., Fikes & Nilson, 1971; Korf, 1987; Allen, Hendler, & Tate, 1990), planning has largely been considered a search problem for finding a sequence of actions that can transform an initial state of the world into a given goal state. Researchers have proposed different approaches to planning that differ with respect to the composition of the search space. In state-space planning, for example, the search space is a graph where nodes are states and actions connect nodes (e.g., applying an action transforms the state of the world into an alternative state). In plan-space planning, nodes represent partially-ordered plans, and links represent planning actions that transform one partially-ordered plan into another. Independent of the particular kind of approach, planning can be seen as a process by which a partial plan is refined (Kambhampati & Srivastava, 1996). In state-space planning, a totally ordered plan is refined; in plan-space planning, a partial-ordered plan is refined.

Case-based planning (CBP) advocates a different view. Instead of searching to refine the current plan, it adapts cases to solve new problems. As discussed in the commentary on representation in this volume, sophisticated case structures in CBP need not be represented; just the problem description (in terms of initial and goal state) and the actual plan solution are needed. Cases may include annotations that describe how the plan was derived (Veloso, 1994; 1996) or that anticipate when adapting a particular plan may fail (Hammond, 1990; Ihrig & Kambhampati, 1997). Furthermore, alternate case representations exist. Cases may contain flat plan representations, hierarchical plan representations (Muñoz-Avila, Paulokat & Wess, 1994) or plan abstractions (Bergmann & Wilke, 1996). Regardless, retrieval and adaptation is the focus, rather than plan refinement from scratch, so the representations must support these processes.

2 CHEF

By developing the first case-based planner (CHEF), Hammond helped to define the case-based approach to problem solving and to explanation (Hammond, 1989; 1990). Given a set of goals and a current situation, the first task for CHEF is to find an old plan that solved a past problem that is similar to the current problem. The next tasks are to adapt the old solution to fit the new circumstances and to store the new solution so that it can be found and reused in the future. However in addition to old plans, Hammond illustrates the use of memory for plan adaptation, plan repair, and failure anticipation.

Because plan knowledge may be incomplete and inconsistent, a system such as CHEF can generate a plan that fails when executed. Hammond (1990) describes how to repair such failed plans in the domain of Chinese cooking. To perform the repair, a planner must explain how the failure occurred, use the explanation to find a set of repair strategies in memory, and then choose and execute the best repair. Once repaired the plan is then saved to memory so that it can be used to anticipate similar failures in the future. Thus CHEF uses a case-based approach to fix faulty plans as well as to solve new planning problems.

In his work, Hammond sets the stage for many subsequent avenues of research in the field. The theory introduces both case-based explanation (c.f., Schank, Kass, & Riesbeck, 1994) as well as case-based planning. It also emphasizes failure-driven learning, knowledge-based indexing, the role of cases for anticipation, and a cognitive-modeling approach toward artificial intelligence. These concepts typify early research into CBR and continue into the present.

3 Derivational analogy and abstraction

Since first introduced by Carbonell (1986) and operationalized into an implemented system (Carbonell & Veloso, 1988; Veloso & Carbonell, 1994), *derivational analogy* has been the subject of numerous studies. These include its application to partial-order planning (Ihrig & Kambhampati, 1997; Muñoz-Avila & Weberskirch, 1996), its integration with mixed-initiative planning (Cox & Veloso, 1997; Veloso, Mulvehill, & Cox, 1997), and its relevance to knowledge requirements (Cunningham, Slattery, & Finn, 1994). Rather than adapt an old solution as a function of the difference between the current and past planning contexts, this approach adapts the derivation of the past solution to derive a new solution or plan.

In case-based planning it is very important to properly consider the levels of *abstraction* on which cases are represented and on which the reuse of the solution or the replay of derivational traces take place. Bergmann and Wilke (1996) systematically analyzes the use of multiple case representations at different levels of abstraction. This paper provides a comparative survey of abstraction approaches in CBR with a focus on case-based planners. It introduces a general framework for analyzing and designing hierarchical CBR applications.

The basic idea behind these approaches is that a case base stores cases represented at several levels of abstraction. When a system must solve a new problem, it retrieves one or more cases at the 'appropriate' abstraction level(s) and reuses the solutions that the case(s) contain to derive a solution for the current problem. For these kinds of approaches, the literature uses the terms *hierarchical case-based reasoning* (Smyth & Cunningham, 1992), *stratified case-based reasoning* (Branting & Aha, 1995), and *reasoning with abstract cases* (Bergmann & Wilke, 1996). A significant advantage of introducing hierarchical representations is the greater flexibility of adaptation, thus increasing the coverage of a single case. Bergmann and Wilke (1996) demonstrate this in an empirical study done with a system called PARIS. They also show clear improvements in the overall planning performance. This result is in the line of similar research by Branting and Aha (1995) and Kambhampati and Hendler (1992).

The importance of the research of Au, Muñoz-Avila, and Nau (2002)[1] is two-fold: first, it provides a unifying framework for derivational analogy that covers state-space, plan-space planners and combinations of these. Second, it clarifies an apparent contradiction between various empirical studies about the performance of derivational analogy as a plan adaptation method and complexity results from Nebel and Koehler (1995). In the empirical studies (Veloso, 1994; Ihrig & Kambhampati, 1997; Muñoz-Avila & Weberskirch, 1996), adaptation with derivational analogy resulted in better problem-solving times than the underlying first-principles planner. However Nebel and Koehler proved that plan adaptation is harder than planning by first principles. The underlying condition for this theoretical result is what the authors called a conservative adaptation strategy (i.e., to reuse as much of the case as possible). Using the unifying framework for derivational analogy, Au and colleagues show that plan adaptation with

[1] This paper won the "Best Paper Award" at the 2002 International Conference on CBR.

derivational analogy does not meet the conservative condition, and, therefore, the complexity results of Nebel and Koehler do not apply.

4 Transformational analogy

Transformational analogy is a plan adaptation approach that advocates reusing an old solution (i.e., plan) as a function of the difference between the current and past planning contexts (Carbonell, 1983). In contrast to derivational analogy, cases contain the solution plans themselves rather than the sequence of derivations that led to the plan. This makes transformational analogy particularly suitable for situations in which only final plans exist without information concerning their generation. Most of the research on transformational analogy concentrates on using domain-specific rules to transform the old plans. However, Hanks & Weld (1995) introduced Systematic Plan Adaptor (SPA), a provable correct transformational approach that transforms a given partial-order plan to solve a new problem. SPA is domain-independent and makes use of partial-order planning techniques. Ram & Francis (1996) introduced Multi-Plan Adaptor (MPA), which extends SPA by allowing the combination of multiple plans using least-commitment strategies from partial-order planning.

5 Additional research and conclusion

We summarized research on important directions in CBP, including derivational analogy, transformational analogy, and abstraction. Much work from the community is absent in our coverage, however. Significant applications in CBP include the work of Ricci and colleagues (Avesani, *et al.,* 2000; Ricci *et al.*, 1999) in the domain of forest fire management. Although not fielded, the HICAP system (Muñoz-Avila, *et al*., 1999) demonstrates practical CBP in the domain of noncombatant evacuation operations. Among the many new research directions, three warrant mention. A case-based approach to plan recognition stores old observations to predict future planning behavior (Kerkez & Cox, 2002; 2003); whereas, case-based mixed-initiative planning (e.g., Cox & Veloso, 1997; Muñoz-Avila, *et al*., 2001) integrates humans as active problem solvers with case-based planners that initiate independent actions and goals. Finally the area of plan merging (Muñoz-Avila & Weberskirch, 1997; Veloso, 1997) investigates new means for combining and adapting multiple cases for single solutions. The article here simply constitutes an introduction to a broad selection of important efforts in case-based planning.

References

Allen, J, Hendler, J, & Tate, A 1990. *Readings in Planning*, San Francisco: Morgan Kaufmann.

Au, T C, Muñoz-Avila, H & Nau, D S 2002. "On the complexity of plan adaptation by derivational analogy in universal classical planning framework" in S. Craw & A. Preece, eds *Advances in Case-Based Reasoning: 6th European Conference*, , Berlin: Springer, pp 13-27.

Avesani, P, Perini, A & Ricci, F. 2000. "Interactive case-based planning for forest fire management" *Applied Intelligence* 13(1) pp 41-57.

Bergmann, R & Wilke, W 1996. "On the role of abstraction in case-based reasoning" in Smith, I & Faltings, B, eds, *Proceedings of the European Workshop on Case-Based Reasoning*, Berlin: Springer, pp 28-43.

Branting, K & Aha, D W 1995. "Stratified case-based reasoning: Reusing hierarchical problem solving episodes" in *Proceedings of the International Joint Conference on Artificial Intelligence* pp 384-390.

Carbonell, J G 1983 "Learning by analogy: Formulating and generalizing plans from past experience" in *Machine Learning: An Artificial Intelligence Approach*, pp 137-161.

Carbonell, J G 1986. "Derivational analogy: A theory of reconstructive problem solving and expertise acquisition" in Michalski, R, Carbonell, J & Mitchell, T eds, *Machine Learning: An Artificial Intelligence Approach, Vol. 2.* Los Altos, California: Morgan Kaufmann, pp 371-392.

Carbonell, J G & Veloso, M M 1988. "Integrating derivational analogy into a general problem solving architecture" in Kolodner, J L, ed, *DARPA Workshop on Case-Based Reasoning*, pp 104-124.

Cox, M T & Veloso, M M 1997. "Supporting combined human and machine planning: An interface for planning by analogical reasoning" in Leake, D B & Plaza, E eds, *Case-Based Reasoning Research and Development, Proceedings of the 2nd International Conference on Case-Based Reasoning*, Berlin: Springer, pp 531-540.

Cunningham, P, Slattery, S & Finn, D 1994. “Knowledge engineering requirements in derivational analogy” in Wess, S, Althoff, K-D & Richter, M eds, *Proceedings of the First European Workshop on Case-Based Reasoning*, Berlin: Springer, pp 234-245.

Fikes, R E & Nilsson, N J 1971. “Strips: A new approach to the application of theorem proving to problem solving” *Artificial Intelligence,* 2 pp 189-208.

Ram, A & Francis, A 1996. “Multi-plan retrieval and adaptation in an experience-based agent” in Leake, DB ed, *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, Menlo Park, CA: AAAI Press.

Hammond, K J 1989. *Case-Based Planning: Viewing planning as a memory task*, Boston: Academic Press.

Hammond, K J 1990. “Explaining and repairing plans that fail” *Artificial Intelligence.* 45 pp 173-228.

Hanks, S and Weld, D 1995. “A domain-independent algorithm for plan adaptation” *Journal of Artificial Intelligence Research*, 2 pp 319-360.

Ihrig, L & Kambhampati, S 1997. “Storing and indexing plan derivations through explanation-based analysis of retrieval failures” *Journal of Artificial Intelligence Research* 7 pp 161-198.

Kambhampati, S & Hendler, J 1992. “A validation-structure-based theory of plan modifications” *Artificial Intelligence* 55(23) pp 193-258.

Kambhampati, S & Srivastava, B 1996. “Universal classical planner: An algorithm for unifying state-space and plan-space planning” in Ghallab, M & Milani, A eds, *New Directions in AI Planning,* Amsterdam: IOS Press, pp 61-75.

Kerkez, B & Cox, M T 2002. “Local predictions for case-based plan recognition” in Craw, S & Preece, A, eds, *Advances in Case-Based Reasoning: 6th European Conference,* Berlin: Springer, pp 189-203.

Kerkez, B & Cox, M T 2003. Incremental case-based plan recognition with local predictions. *International Journal on Artificial Intelligence Tools: Architectures, languages, algorithms* 12(4), pp 413-464.

Korf, R E 1987. “Planning as search: a quantitative approach” *Artificial Intelligence* 33(1) pp 65-88.

Muñoz-Avila, H, Aha, D W, Nau D S, Breslow, L A, Weber, R & Yamal, F 2001. “SiN: Integrating case-based reasoning with task decomposition” in *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-2001),* Menlo Park, California: AAAI Press, pp 999-104.

Muñoz-Avila, H, Paulokat J & Wess S 1994. “Controlling nonlinear hierarchical planning by case replay” in: *Proceedings of the European Workshop on Case-based reasoning*, Berlin: Springer, pp 266-279.

Muñoz-Avila, H & Weberskirch F 1996. “Planning for manufacturing workpieces by storing, indexing and replaying planning decisions” in *Proceedings of the 3rd International Conference on AI Planning Systems (AIPS-96),* Menlo Park, California: AAAI-Press, pp 150-157.

Muñoz-Avila, H & Weberskirch F 1997. “A case study on the mergeability of cases with a partial-order planner” in Steel, S & Alami, R, eds, *Recent Advances in AI Planning. Proceedings of ECP-97*, Berlin: Springer, pp 325-338.

Muñoz-Avila, H, McFarlane, D, Aha, D W, Ballas, J, Breslow, L A & Nau, D 1999. “Using guidelines to constrain interactive case-based HTN planning” in *Proceedings of the Third International Conference on Case-Based Reasoning*, Berlin: Springer, pp 288-302.

Nebel, B & Koehler, J 1995. Plan reuse versus plan generation: A theoretical and empirical analysis. *Artificial Intelligence* 76(1-2) pp 427-454.

Ricci, F, Avesani, P & Perini, A 1999. “Cases on fire: Applying CBR to emergency management” *New Review of Applied Expert Systems* 5 pp 175-190.

Schank, R C, Kass, A & Riesbeck, C 1994. *Inside Case-Based Explanation*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Smyth, B & Cunningham, P 1992. “Deja vu: A hierarchical case-based reasoning system for software design” in B. Neumann, ed, *10th European Conference on Artificial Intelligence*. Hoboken, New Jersey: John Wiley and Sons, pp 587-589.

Veloso, M M 1994. *Planning and Learning by Analogical Reasoning*. Berlin: Springer.

Veloso, M M 1996. “Flexible strategy learning using analogical replay of problem solving episodes” in Leake, D B, ed, *Case-Based Reasoning: Experiences, Lessons, and Future Directions,* Menlo Park, California: AAAI Press / The MIT Press, pp 137-150.

Veloso, M M 1997. “Merging strategies for multiple case plan replay” in Leake, D B & Plaza, E, eds, *Proceedings of the 2nd International Conference on Case-Based Reasoning* Berlin: Springer..

Veloso, M M & Carbonell, J G 1994. “Case-based reasoning in PRODIGY” in Michalski, R S & Tecuci, G, eds, *Machine Learning: A multistrategy approach IV,* San Mateo, California: Morgan Kaufmann.

Veloso, M M, Mulvehill, A M & Cox, M T 1997. “Rationale-supported mixed-initiative case-based planning” in *Proceedings of the Ninth Innovative Applications of Artificial Intelligence Conference,* Menlo Park, California: AAAI Press The MIT Press, pp 1072-1077.