

Case-based Acquisition of User Preferences for Solution Improvement in Ill-Structured Domains

Katia Sycara
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213, U.S.A.
katia@cs.cmu.edu

Kazuo Miyashita
Production Engineering Division
Matsushita Electric Industrial Co.
Kadoma, Osaka 571, Japan
miyasita@mcec.ped.mei.co.jp

Abstract

¹ We have developed an approach to acquire complicated user optimization criteria and use them to guide iterative solution improvement. The effectiveness of the approach was tested on job shop scheduling problems. The ill-structuredness of the domain and the desired optimization objectives in real-life problems, such as factory scheduling, makes the problems difficult to formalize and costly to solve. Current optimization technology requires explicit global optimization criteria in order to control its search for the optimal solution. But often, a user's optimization preferences are state-dependent and cannot be expressed in terms of a single global optimization criterion. In our approach, the optimization preferences are represented implicitly and extensionally in a case base. Experimental results in job shop scheduling problems support the hypotheses that our approach (1) is capable of capturing diverse user optimization preferences and re-using them to guide solution quality improvement, (2) is robust in the sense that it improves solution quality independent of the method of initial solution generation, and (3) produces high quality solutions, which are comparable with solutions generated by traditional iterative optimization techniques, such as simulated annealing, at much lower computational cost.

Introduction

We present an approach, implemented in the CABINS system, to demonstrate the capability of acquiring user context-dependent optimization preferences and reusing them to guide iterative solution optimization in ill-structured domains. This capability is very important for two main reasons. First, traditional search methods, both Operations Research-based and AI-based, that are used in combinatorial optimization, need explicit representation of objectives in terms of a cost function to be optimized

¹This research was partially supported by the Defense Advance Research Projects Agency under contract #F30602-88-C-0001. Most of the work was performed when the second author was a visiting scientist at the Robotics Institute at Carnegie Mellon University under the support of Matsushita Electric Industrial Co.

(Reeves 1993). In many practical problems, such as scheduling and design, optimization criteria often involve context- and user-dependent tradeoffs which are impossible to realistically consolidate in a cost function. Second, expert system approaches, while having the potential to capture context-dependent tradeoffs in rules, require considerable knowledge acquisition effort (Prerau 1990). Our approach uses case-based reasoning (CBR) which has been successful in dealing with exceptional data (Golding & Rosenbloom 1991; Ruby & Kibler 1992), acquiring user knowledge in complex domains (Chaturvedi 1992; McKay, Buzacott, & Safayeni 1988), and expending less effort in knowledge acquisition compared with knowledge acquisition for rule-based systems (Lewis, Minior, & Brown 1991). CABINS acquires, stores and reuses two categories of concepts that reflect user preferences (1) what heuristic local optimization action to choose in a particular context, and (2) what combinations of effects of application of a particular local optimization action constitutes an acceptable or unacceptable outcome. These are recorded in the case base and are used by CABINS to guide iterative optimization and induce optimization tradeoffs to evaluate the current solution. The optimization criteria are not explicitly represented as case features or in terms of a cost function but are implicitly and extensionally represented in the case base.

Previous case-based systems for incremental solution revision (e.g. (Hammond 1989; Veloso 1992)) have been motivated only by concerns of computational efficiency, preserving plan correctness rather than improving plan quality, and have assumed the existence of a strong domain model that provides feedback as to plan correctness. Case-based knowledge acquisition systems, (e.g. (Bareiss 1989)) require causal explanations from an expert teacher to acquire domain knowledge. In our approach neither the user nor the program are assumed to possess causal domain knowledge. The user's expertise lies in his/her ability to perform consistent evaluation of the results of problem solving and impart to the program cases of problem solving experiences and histories of evaluation tradeoffs.

In this paper, we present initial experimental re-

sults to test three hypotheses. First, our CBR-based incremental revision methodology shows good potential for capturing user optimization preferences in ill-structured domains, such as job shop scheduling, and re-using them to guide optimization. Second, the method is robust in the sense that it improves solution quality independent of the method of initial solution generation. Third, CABINS produces high quality solutions. To test this, we compared the solutions produced by CABINS with explicit optimization criteria, with solutions produced by simulated annealing (a well known iterative optimization technique (Johnson *et al.* 1991; Zweben, Deale, & Gargan 1990; Laarhoven, Aarts, & Lenstra 1992)) for the same criteria. Our investigation was conducted in the domain of job shop schedule optimization and the experimental results, shown in section confirmed these hypotheses.

Job Shop Schedule Optimization

The job shop scheduling problem is one of the most difficult NP-hard combinatorial optimization problems (French 1982). Job shop scheduling deals with allocation of a limited set of resources to a number of activities (operations) associated with a set of jobs so as to respect given temporal relations (c.g. precedence relations among activities), temporal constraints (e.g. job release and due dates) and resource capacity restrictions in order to optimize a set of objectives, such as minimize tardiness, minimize work in process inventory (WIP), maximize resource utilization etc. Due to the tight interactions among scheduling constraints and the often conflicting nature of optimization criteria, it is impossible to assess with any precision the extent of schedule revision or the impact of a scheduling decision on the global satisfaction of optimization criteria. For example, in figure 1 moving forward the last activity of ORDER3 creates downstream cascading constraint violations. Therefore, a repair action must be applied and its repair outcome must be evaluated in terms of the resulting effects on scheduling objectives. In addition, the evaluation itself of what is a "high quality" schedule is difficult because of the need to balance conflicting objectives and trade-off among them. Such tradeoffs typically reflect user preferences, which are difficult to express as a cost function. For example, WIP and weighted tardiness are not always compatible with each other. As shown in figure 2, there are situations where a repair action can reduce weighted tardiness, but WIP increases. Which is a better schedule depends on user preferences.

CABINS *incrementally revises a complete but sub-optimal schedule* to improve its quality, based on flexible optimization tradeoffs. Revision-based approaches to scheduling have also been investigated by (Minton *et al.* 1990; Zweben, Deale, & Gargan 1990; Biefeld & Cooper 1991; Laarhoven, Aarts, & Lenstra 1992). In those systems, the initial schedule is repaired by several techniques, such as the min-conflict heuris-

tic or simulated annealing, to minimize the number of constraint violations or optimize a simple cost function (e.g. make-span) of the schedule. The value of incorporating context-dependent user preferences in operational scheduling environments is becoming increasingly recognized (e.g. (Mckay, Buzacott, & Safayeni 1988)) but adequate techniques are lacking.

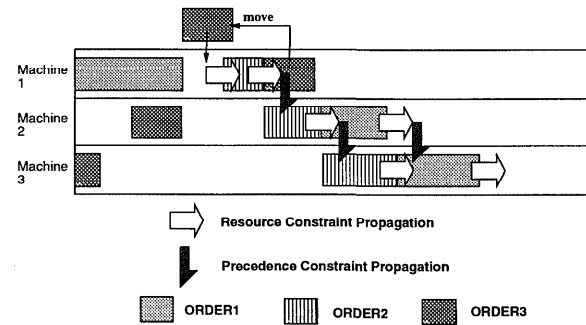


Figure 1: Example of Tight Constraint Interactions

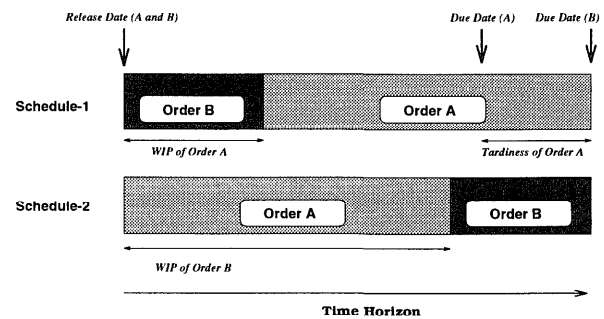


Figure 2: Example of Conflicting Objectives

CABINS Overview

CABINS is composed of three modules: (1) an initial schedule builder, (2) an interactive schedule repair (case acquisition) module and (3) an automated schedule repair (case re-use) module. To generate an initial schedule, CABINS can use any of several scheduling methods (e.g. traditional dispatching rules or a constraint-based scheduler).

Case representation

In each repair iteration, CABINS focuses on one activity at a time, the *focal activity*, and tries to repair it. A case in CABINS describes the application of a particular modification to a focal activity. Figure 3 shows the information content of a case. Our assumption, borne out by the experimental results, is that despite the ill-structuredness of the domain, the global, local and repair history features express (in an approximate

manner) domain regularities. The global features reflect an abstract characterization of potential repair flexibility for the whole schedule. High 'Resource Utilization Average', for example, often indicates a tight schedule without much repair flexibility. Associated with a focal_activity are local features that we have identified, based on those reported in (Ow, Smith, & Thiriez 1988), and which potentially are predictive of estimating the effects of applying a particular repair tactic to the schedule. For example, 'Predictive Shift Gain' predicts how much overall gain will be achieved by moving the current focal_activity earlier in its time horizon. In particular, it predicts the likely reduction of the focal_activity's waiting time when moved to the left within the repair time horizon.

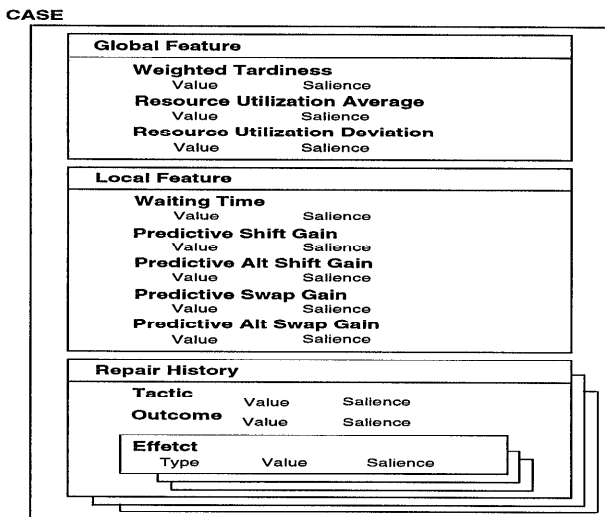


Figure 3: CABINS Case Representation

The repair history records the sequence of applications of successive repair tactics, the repair outcome and the effects. Repair effect values describe the impact of the application of a repair action on scheduling objectives (e.g. weighted tardiness, WIP). A repair outcome is the evaluation assigned to the set of effects of a repair action and takes values in the set ['acceptable', 'unacceptable']. Typically the outcome reflects tradeoffs among different objectives. If the application of a repair tactic results in a feasible schedule, the result is judged as either acceptable or unacceptable with respect to the repair objectives. An outcome is 'acceptable' if the user accepts the tradeoffs involved in the set of effects for the current application of a repair action. Otherwise, it is 'unacceptable'. The effect saliency is assigned when the outcome is 'unacceptable', and it indicates the significance of the effect to the repair outcome. This value is decided subjectively and interactively. The user's judgment as to balancing favorable and unfavorable effects related to a particular objective constitutes the explanation of the repair

outcome.

Case acquisition

To gather cases, sample scheduling problems are solved by a scheduler. CABINS identifies jobs that must be repaired in the initial sub-optimal schedule. Those jobs are sorted according to the significance of defect, and repaired manually by a user according to this sorting. For example, if the user's optimization criterion is to minimize order tardiness, the most tardy order is repaired first. The user selects a repair tactic to be applied. Tactic application consists of two parts: (a) identify the activities, resources and time intervals that will be involved in the repair, and (b) execute the repair by applying constraint-based scheduling to reschedule the activities identified in (a). Currently CABINS has 11 tactics and a flexible interface through which the user can define more.

After tactic selection and application, the repair effects are calculated and shown to the user who is asked to evaluate the outcome of the repair. If the user evaluates the repair outcome as 'acceptable', CABINS proceeds to repair another focal_activity and the process is repeated. If the user evaluates the repair outcome as 'unacceptable', s/he is asked to supply an explanation in terms of rating the saliency/importance of each of the effects. The repair is undone and the user is asked to select another repair tactic for the same focal_activity. The process continues until an acceptable outcome for the current focal_activity is reached, or the repair is given up. Repair is given up when there are no more tactics to be applied to the current focal_activity; in this situation, CABINS carries on repair of another activity. The sequence of applications of successive repair actions, the effects, the repair outcome, and the user's explanation for failed application of a repair tactic are recorded in the repair history of the case. In this way, a number of cases are accumulated in the case base.

Case re-use

Once cases have been gathered, CABINS repairs sub-optimal schedules without user interaction. CABINS repairs the schedules by (1) recognizing schedule sub-optimality, (2) focusing on a focal_activity to be repaired in each repair cycle, (3) invoking CBR with the set of global and local features as indices to decide the most appropriate repair tactic to be used for each focal_activity, (4) invoking CBR using the repair effect features (type, value and saliency) as indices to evaluate the repair result, and (5) when the repair result is unacceptable, deciding which repair tactic to use next. Note that in contrast to traditional local iterative optimization approaches, (e.g. tabu search, simulated annealing) where the schedule generated in the current iteration as a result of local revision is directly compared (in terms of its associated cost function) with the current schedule, in CABINS, evaluation of the re-

vision is provided by the case base, thus obviating the need for the presence of an explicit cost function.

The similarity between i-th case and the current problem is calculated as follows :

$$\exp\left(-\sqrt{\sum_{j=1}^N (SL_j^i \times \frac{CF_j^i - PF_j}{E-D_j})^2}\right)$$

where SL_j^i is the salience of j-th feature of i-th case in the case-base, and its value has been heuristically defined by the user. CF_j^i is the value of j-th feature of i-th case, PF_j is the value of j-th feature in the current problem, $E-D_j$ is the standard deviation of j-th feature value of all cases in the case-base. Feature values are normalized by division by a standard deviation of the feature value so that features of equal salience have equal weight in the similarity function.

An Example

We briefly illustrate the repair process with a very simple example schedule to be repaired shown in figure 4. The example has ten jobs (J_1, \dots, J_{10}) and each job has five activities with linear precedence constraints. (e.g. O_1^n BEFORE O_2^n, \dots, O_4^n BEFORE O_5^n). Resources R_1 and R_2, R_3 and R_5 are substitutable; resource R_4 is a bottleneck. Suppose that the job under repair is J_8 . This job has a weight of 2, a due date of 1250 and the scheduled end-time of its last activity is 1390. Hence it has a weighted tardiness of $2 \times (1390 - 1250) = 280$. Suppose the current focal_activity is O_4^8 . CBR is invoked with global features (weighted tardiness= 280, resource utilization average=0.544, resource utilization deviation=0.032) plus the set of local features as indices and selects swap as a repair tactic. One can see from the figure that this is a good choice since the focal_activity is scheduled on machine R_4 , which doesn't have any substitutable machine and any idle time in the repair time horizon (time between the end of O_3^8 and the end of O_4^8).

To apply swap, CABINS calculates the activity with which O_4^8 will be swapped. To do this, CABINS selects the activity which, if swapped with O_4^8 , will result in least amount of precedence constraint violations. In the example, activity O_4^4 is selected as the activity to be swapped with the current focal_activity O_4^8 . Job J_4 has weight 3 and weighted tardiness $3 \times (1370 - 1320) = 150$. The effect of applying the swap tactic is that O_4^8 and O_4^4 are unscheduled on R_4 and O_4^8 is re-scheduled to start at time 1090 (the start time of activity O_4^4 prior to the swap). The repair process resolves occurring constraint violations. The repaired schedule is shown in figure 5.

The effects of repairing O_4^8 are calculated. CABINS calculates the effects on J_8 and J_4 , the jobs affected by the application of the swap on O_4^8 . Machine utilization did not change but J_8 had an estimated decrease in weighted-tardiness of 180 time units and an estimated decrease in WIP of 200 units, J_4 had an increase in




Figure 4: Original Schedule Results




Figure 5: Schedule Results after Repair on O_4^8

weighted-tardiness of 150 units and an increase in WIP of 750 units. CBR is invoked using these effect values, weighted tardiness, WIP, as indices to determine whether this repair outcome is acceptable. The acceptability or unacceptability of the repair will depend on the biases reflected in the case base.

Evaluation of the Approach

We conducted a set of experiments to test the hypothesis that (1) our CBR-based incremental modification and re-use methodology could be effective in capturing user schedule optimization preferences and re-using them to control schedule optimization, (2) the approach is robust in that the schedules produced by CABINS consistently improve a schedule independent of the method used for initial schedule generation and (3) as an iterative optimization method, the approach produces schedules of high quality. These hypotheses are difficult to test since, due to the subjective and ill-defined nature of user preferences, it is not obvious how to correlate scheduling results with the captured preferences or how to define quality of a schedule whose evaluation is subjective.

To address these issues, we had to devise a method to test the hypotheses in a consistent manner. To do that, it is necessary to know the optimization criterion that would be implicit in the case base, so that the experimental results can be evaluated. In the experiments reported here, we used two different explicit criteria (weighted tardiness; WIP+weighted tardiness) to reflect the user's optimization criterion and built a

rule-based reasoner (RBR) that goes through a trial-and-error repair process to optimize a schedule. For each repair, the repair effects were calculated and, on this basis, since RBR had a predefined evaluation objective, it could evaluate the repair outcome in a consistent manner. Thus, we used RBR with different rules each time to generate different case bases (each with 1,000 cases)² for different explicit optimization objectives. Naturally, an objective, though known to us, is not known to CABINS and is only implicitly and indirectly reflected in an extensional way in each case base. By designing an objective into the RBR so it could be reflected in the corresponding case base we got an experimental baseline against which to evaluate the schedules generated by CABINS.

We evaluated the approach on a benchmark suite of 60 job shop scheduling problems where parameters, such as number of bottlenecks, range of due dates and activity durations were varied to cover a range of job shop scheduling problem instances with the following structure. Each problem class has 10 jobs of 5 operations each and 5 machines. Two parameters were used to cover different scheduling conditions: a range parameter controlled the distribution of job due dates and release dates, and a bottleneck parameter controlled the number of bottleneck resources. Six groups of 10 problems each were randomly generated by considering three different values of the range parameter, and two values of the bottleneck configuration (1 and 2 bottleneck problems). These problems are variations of the problems originally reported in (Sadeh 1991). Our problem sets are, however, different in two respects: (a) we allow substitutable resources for non-bottleneck resources whereas the original problems did not, and (b) the due dates of jobs in our problems are tighter by 20 percents than in the original problems. We also tested the approach on another set of 60 problems of 20 orders and 5 resources with similar results.

A cross-validation method was used to evaluate the learning capability of CABINS. Each problem set in each class was divided in half. The training sample was repaired by RBR to gather cases. These cases were then used for case-based repair of the validation problems. We repeated the above process by interchanging the training and test sets. Reported results are for the validation problem sets.

Experimental Results

Figures 6 show the performance of CABINS using “weighted tardiness” case base (labeled in the figures as CABINS(WT)) vs performance of CABINS using the “weighted tardiness and WIP” case base (labeled in the figures as CABINS(WT+WIP)). The cases constituted the only source of knowledge for CABINS. In

²Since a case represents the application of one repair tactic to an activity, if, for example, 5 repair tactics are utilized in an attempt to successfully repair an activity, then 5 cases would be created.

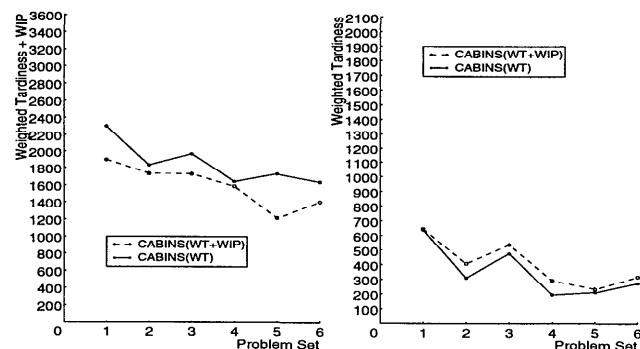


Figure 6: Scheduling Results with Different Case Bases

	Wei.Tar.	WIP	Total	CPU Sec.
EDD	956.0	1284.6	2240.6	0.1
CABINS	349.5	1311.2	1660.7	73.5
SA	340.5	1333.4	1673.9	388.2
WSPT	584.0	1241.0	1825.0	0.1
CABINS	321.0	1254.9	1575.9	72.1
SA	328.5	1320.4	1684.9	398.3
R&M	556.0	1242.0	1798.0	0.1
CABINS	305.3	1264.9	1570.2	84.9
SA	330.1	1290.8	1620.9	450.5
CBS	1173.0	1481.0	2654.0	17.4
CABINS	405.3	1195.0	1600.3	296.5
SA	395.5	1220.0	1615.5	1380.0

Table 1: Repair by CABINS and SA based on Different Methods of Initial Schedule Generation

other words, there was no objective given to CABINS explicitly. The case-bases were used both as a source of suitable repairs, and also as a source of advice regarding repair evaluation. From the results we observe that CABINS(WT) generated higher quality schedules with respect to minimizing weighted tardiness than CABINS(WT+WIP). Conversely, CABINS(WT+WIP) generated higher quality schedules with respect to WIP, and weighted tardiness plus WIP than CABINS(WT). These results indicate that CABINS can acquire different and subjective user preferences.

In order to test the hypothesis that CABINS consistently improves schedule quality independent of the method of initial schedule generation, we generated initial schedules for the benchmark suite of problems using three different state-of-the-art dispatch scheduling heuristics (EDD, WSPT, R&M) (Morton 1992) and a constraint-based scheduler (CBS). The optimization objective was WT+WIP. Table 1 presents the average of all 60 problems in the benchmark and shows that CABINS improved schedule quality independent of method to create the initial schedule. To test the

hypothesis that CABINS generates schedules of high quality, we compared the schedules generated by CABINS against schedules generated by simulated annealing with the explicit objective of WT+WIP. Table 1 shows that CABINS generated schedules of comparable quality but was on the average 4-5 times more efficient than simulated annealing.

Conclusions

We have presented a case-based approach to acquire user optimization preferences and reuse them to guide iterative solution optimization in ill-structured domains. We demonstrated the effectiveness of the approach in capturing user preferences and creating efficiently high quality solutions on job shop scheduling problems. One crucial issue is how much effort should be spent to capture "enough" number of cases for "sufficient" solution quality improvement. This is an issue we are currently pursuing. Initial experiments to determine case base size versus quality improvement have shown that a case base of 800 cases gives on the average 20% higher quality improvement at 15% lower computational cost than a case base of 400 cases. It seems that the effort expended to capture a big number of cases can be amortized by future repeated use of the case base to get high quality schedules efficiently. More importantly, CABINS can acquire those cases from user's interaction during the process of solution improvement, thus imposing low additional effort on the user but enhancing solution improvement. We believe that CABINS has the potential for accommodating acquisition of user preferences that change over time. Future work will investigate this issue.

References

- Bareiss, R. 1989. *Exemplar-based knowledge acquisition : a unified approach to concept regression, classification, and learning*. New York, NY: Academic Press.
- Biefeld, E., and Cooper, L. 1991. Bottleneck identification using process chronologies. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*.
- Chaturvedi, A. 1992. Acquiring Implicit Knowledge in a Complex Domain. *Expert Systems with Applications*.
- French, S. 1982. *Sequencing and Scheduling: An Introduction to the Mathematics of the Job-Shop*. New York, NY: Ellis Horwood.
- Golding, A. R., and Rosenbloom, P. S. 1991. Improving Rule-Based Systems Through Case-Based Reasoning. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 22-27. AAAI.
- Hammond, K. J. 1989. *Case-Based Planning : Viewing Planning as a Memory Task*. New York, NY: Academic Press.
- Johnson, D.; Aragon, C.; McGeoch, L.; and Schevon, C. 1991. Optimization By Simulated Annealing: An Experimental Evaluation, Part II (Graph Coloring and Number Partitioning). *Operations Research*.
- Laarhoven, P. J. M. V.; Aarts, E. H. L.; and Lenstra, J. K. 1992. Job shop scheduling by simulated annealing. *Operations Research* 40(1):113-125.
- Lewis, L.; Minior, D.; and Brown, S. 1991. A Case-Based Reasoning Solution to the Problem of Redundant Engineering in Large Scale Manufacturing. *International Journal of Expert Systems* 4(2):189-201.
- Mckay, K.; Buzacott, J.; and Safayeni, F. 1988. The scheduler's knowledge of uncertainty: The missing link. In *Proceedings of IFIP Working Conference on Knowledge Based Production Management Systems*.
- Minton, S.; Johnston, M. D.; Philips, A. B.; and Laird, P. 1990. Solving large-scale constraint satisfaction and scheduling problems using a heuristic repair method. In *Proceedings, Eighth National Conference on Artificial Intelligence*, 17-24. Boston, MA.: AAAI.
- Morton, T. E. 1992. *HEURISTIC SCHEDULING SYSTEMS: With Application to Production Systems and Product Management*. Pittsburgh, PA.: GSIA, Carnegie Mellon University. Course Textbook.
- Ow, P. S.; Smith, S. F.; and Thiriez, A. 1988. Reactive plan revision. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 77-82. St-Paul, Minnesota: AAAI.
- Prerau, D. S. 1990. *Developing and Managing Expert Systems: Proven Techniques for Business and Industry*. Reading, MA: Addison-Wesley.
- Reeves, C., ed. 1993. *Modern Heuristic Techniques for Combinatorial Problems*. New York: Halsted Press.
- Ruby, D., and Kibler, D. 1992. Learning Episodes for Optimization. In *Machine Learning : proceedings of the Ninth International Workshop (ML92)*, 379-384.
- Sadeh, N. 1991. *Look-Ahead Techniques for Micro-Opportunistic Job Shop Scheduling*. Ph.D. Dissertation, School of Computer Science, Carnegie Mellon University.
- Veloso, M. M. 1992. *Learning by Analogical Reasoning in General Problem Solving*. Ph.D. Dissertation, School of Computer Science, Carnegie Mellon University.
- Zweben, M.; Deale, M.; and Gargan, M. 1990. Any-time rescheduling. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control*, 251-259. San Diego, CA.: DARPA.