

Alternate Strategies for Retrieval in State-Spaces

Boris Kerkez and Michael T. Cox

Department of Computer Science and Engineering
Wright State University, Dayton, OH
{bkerkez,mcox}@cs.wright.edu

Abstract

In our previous research, we investigated the properties of case-based plan recognition with incomplete plan libraries. Incremental construction of plan libraries along with retrieval based on similarities among planning situations (rather than on similarities among planning actions) enables recognition in light of novel planning actions. In this paper we investigate the recognition behavior in situations where the recognizer fails to find past situations that match the currently observed situation at any level of abstraction. Such recognition behavior is especially common in early recognition stages when the rate of new bin observations is large. To cope with newly observed situations, we employ a retrieval scheme that utilizes a similarity measure among the states in the abstract state-space, based on the k-nearest neighbor similarity metric. Such a retrieval scheme may enable recognition in light of newly observed abstract situations. Properties of the retrieval in abstract state-spaces are investigated in two different planning domains. Experimental results show that improvements in the recognition process depend on the characteristics of a given planning domain.

Introduction

An important trait of intelligent agents in multi-agent environments is the ability to explain the behavior and motivation of other agents with which they interact. In the context of planning environments, the goal of inferring the intentions of planning agents is accomplished through plan recognition. Plan recognition is the task of inferring a planner's goals and the plans to achieve these goals from observations of the planner's behavior. Such intentional inferences may then be used to assist the planner by pointing out potential faulty plans (e.g., tutoring systems), or to obstruct the planner by changing the state of the planning environment (e.g., military counter-planning).

Plan recognition systems typically operate by matching observed partial plans with planning episodes in the plan library. Most systems operate with a complete plan library that contains all possible plans the planning agent may pursue. However, the completeness of plan libraries may introduce efficiency overhead (Lesh and Etzioni, 1996).

Furthermore, enumeration of all possible planning episodes may not be achievable in some complex planning domains.

In our previous work (Kerkez and Cox, 2002; Kerkez and Cox, 2001), we explored a case-based approach to plan recognition that enables reasoning despite incomplete domain knowledge presented in the form of incomplete plan libraries. Although case-based approaches (Bares *et al.*, 1994) and the use of abstraction (Bergman and Wilke, 1996) have been investigated in the context of plan recognition, our work incorporates these concepts in plan recognition in a novel manner (Kerkez and Cox, 2002). In the context of this work, indexing and retrieval are based on recalling the previously observed situations that are the most similar to the current situation at hand, and then using past similar situations to determine the planner's current course of action. This is in contrast to traditional recognition systems that reason in terms of planning actions with respect to the common initial planning state. As a consequence, our approach enables the recognizer to reason even with novel planning actions. After executing a novel action, the planner may find itself in a world state that is similar to some previously observed planning situation, and then refer to the planner's previous behavior in similar past situations in order to infer the planner's future intentions.

The main disadvantage of the retrieval based on the state similarity is that the currently observed situation may be novel at the highest level of abstraction. That is, the recognition system may fail to find past situations that are identical to the current one at the abstract level. This is because the plan libraries in this work are incomplete and therefore it cannot be guaranteed that the description of the current situation at any level of abstraction is contained in the storage structures. This is especially true in early stages of incremental recognition with incomplete libraries, when the rate of observation of new indexing structures (bins and equivalence classes) is rather large.

In this paper we describe an extension of the retrieval scheme that enables recognition in light of novel abstract situations. The system utilizes retrieval based on the similarity in the abstract state-space described in the next section. Because the abstract states are non-negative integer vectors, a feature vector similarity with a k-nearest neighbor similarity metric is appropriate for similarity assessment in the abstract state-space. The third section discusses situation similarity assessment at the abstract

level in more detail, while the last section illustrates experimental evaluations of this technique on two different planning domains.

Case-Based Plan Recognition

Although the planning situations, represented by the states of the planner's environment, are useful as indices for recognition with incomplete libraries, their practical use is limited due to potentially large state-spaces in complex planning domains. We cope with the state-space complexity issues by utilizing a multiple-level indexing scheme, which in turn enables efficient retrieval of past situations that are similar to the currently observed situation. Planner's world states, represented as collections of ground literals, can be abstracted into non-negative integer vectors by counting the number of occurrences of type-generalized predicates in a state, and placing the sum in the appropriate dimension. Figure 1 shows an example of state abstraction in the blocksworld planning domain.




Figure 1. An example of a blocksworld domain state and its different representations.

Abstract states are employed as the first level of indexing, because many concrete states may share the identical abstract representation. Figure 2 illustrates the indexing structures used for storage and retrieval processes.




Figure 2. Indexing structures. Abstract states (as_i) point to bins with disjoint equivalence classes, containing concrete past situations (s_i) that point to the past cases in which they are contained through state pointers (dashed lines into the library).

Our abstract state-space contains abstracted states (as_i) and observed plans (P_i) in their abstract representations.

Abstract states point to structures called bins, which contain concrete states (s_i) sharing the identical abstract representation. A pseudo-isomorphic equivalence relation provides a natural partitioning of bin members into disjoint equivalence classes, serving as an additional level of indexing. Given the currently observed planning situation, the recognizer transforms it into an abstract representation, locates the matched bin and appropriate equivalence class, and retrieves all past situations from the matched equivalence class. The recognizer will then retrieve all previous cases containing the matched past situations via the state pointers into the plan library (see Figure 2). These past cases may then be utilized to infer the planner's current intent. Experimental results (Kerkez and Cox, 2002) show the effectiveness of the retrieval scheme based on situation reminding and multiple indexing levels in recognition with incomplete plan libraries.

The indexing scheme described above enables the recognizer to reason in light of novel planning actions. Upon execution of a novel action, the planner may reach a situation that is similar to some previously observed situations already present in the library. Past plans containing these similar past situations may then be used to predict the planner's intent although the current action observed is a completely new one. The indexing scheme will enable the recognition when the currently observed situation is new, provided that its abstract representation can be found in the abstract state-space. However, retrieval failure occurs when a completely new abstract state is observed.

Abstraction and Similarity

The plan recognition system described here recalls past situations that are similar to the currently observed situation and bases its predictions about the planner's current intent on the planner's intentions in similar situations observed in the past. Upon observation of an execution of some action by the planner, both the executed action and the current state reached after the action application are received by the recognizer. The current planning situation is then transformed into its abstract representation and the recognizer attempts to locate a bin that matches the observed situation by searching the abstract state-space.

In case of a successful match at the abstract state level, the retrieval process focuses on the found bin. The recognizer next attempts to find the equivalence class for the newly observed planning situation. If the current situation is a member of an existing equivalence class, past cases containing situations from the matched equivalence class are retrieved from the library via the state pointers (see Figure 2). Otherwise, a new equivalence class is created in the matched bin, and the newly observed situation is stored in the created equivalence class. In such cases, the recognizer retrieves past cases pointed to by the concrete situations in the matched bin. Although an equivalence class match is more focused than a match at

the bin level, the latter type of a match is very useful. Abstract indexing scheme guarantees that all states in a single bin are similar, because they all share the identical abstract representation.

The recognizer, however, is not able to find matches when the currently observed planning situation has a novel abstract representation. Match failure at the level of abstract states indicates presence of a newly observed abstract state and prevents the recognizer from finding similar past situations to guide the recognition process. Instead of attempting to find an exact match at the abstract state level, the recognizer may attempt to find abstract states that are similar to the abstract representation of the currently observed planning situation. Because abstract states are non-negative integer vectors, a k-nearest neighbor similarity metric can be used to locate similar states in the abstract state-space. Once such similar abstract states are found, the recognizer may be able to use past situations from similar bins to infer the planner's current intent.

The retrieval based on similarity among the abstract states is illustrated in Figure 3, which shows a portion of abstract state-space along with the indexing structures. Upon observing a new abstract state as_{new} , the recognizer creates a new bin (*Bin z*) with a single equivalence class in which the newly observed situation s_{new} is stored. By using




Figure 3. Indexing structures upon observation of a new state s_{new} with a novel abstract representation as_{new} .

the k-nearest neighbor similarity metric, the recognizer then locates all abstract states whose distance from the newly observed abstract state is less than or equal to the specified radius. Figure 3 shows an example in which two similar abstract states are found at distances d_{a1} and d_{a2} in the abstract state-space within the specified maximum distance. These abstract states point to *Bin c* and *Bin v*, that contain past situations at distances d_1 and d_2 . Although not a result of a direct match, these past situations may be similar to the current situation at hand and may provide guidance for inferring the planner's current intentions.

The success of the abstract similarity measure depends on the applicability of the planner's past intentions to guide the current recognition process. Past planning actions may not even be applicable in the current planning situation, which comes into play when dealing with local predictions concerning possible choices for the next planning action. Consequently, predictions based on the abstract similarity metric are less likely to be correct than predictions based on an exact match at the level of abstract states. This is especially the case when the radius for the abstract similarity measure is large and the k-nearest neighbor scheme retrieves a large number of bins that fall within the specified maximum distance. Nevertheless, some of these predictions may be correct and therefore may improve the overall prediction accuracy of the recognizer.

Experimental Results

In order to investigate the properties of the retrieval strategy based on partial matches in the abstract state-space, we investigated the performance of the Java-based implementation of plan recognition system utilizing similarities among situations as the basis for its reminding process. We focused our evaluations on plans in the logistics (Veloso, 1994) and the extended-STRIPS (Veloso and Carbonell, 1994) planning domains, executed by the PRODIGY (Carbonell *et al.*, 1992) state-space planner. The execution cycle of the planner was modified to monitor the intermediate planning states along with the executed planning actions. The recognition process was executed on a large scale, with about 60,000 and 13,000 observed and stored planning steps in the logistics and the extended-STRIPS planning domains, respectively. The results are averages of two different problem sets in each planning domain. These tests were performed to investigate the pattern of the increase in the overall prediction accuracy as we vary radii in our abstract retrieval scheme.

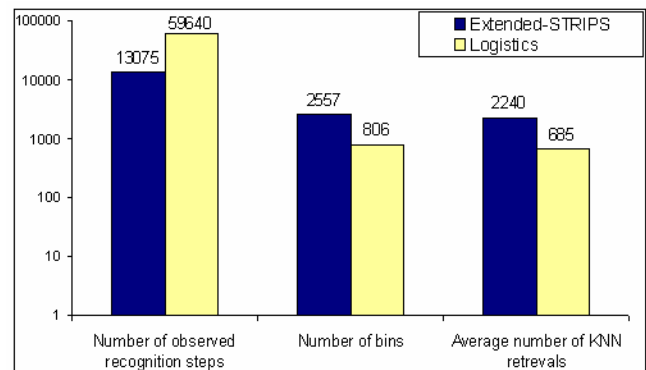


Figure 4. Logarithmic plot of several characteristic measurements for the extended-STRIPS (darker bars on left) and the logistics (lighter bars on right) domains.

The two domains we used for evaluation purposes have quite different characteristics. The logistics domain is characterized by a smaller abstract state-space, as can be

seen from the total number of bins for the two domains shown in the middle column of Figure 4. This results in a smaller number of retrievals based on the k-nearest neighbor similarity metric in the logistics domain, as can be seen in the right column of Figure 4. While the recognition process reaches a steady-state behavior in the logistics domain, it fails to do so in the extended-STRIPS domain, due to a high rate of observation of new abstract states even after 13,000 observed planning steps. A larger number of observed planning steps in the logistics domain, shown in the left column of Figure 4, is also an important factor for achieving the steady state behavior in this domain.




Figure 5. Percentages of increase in overall prediction accuracies with respect to retrieval without partial matches in the abstract state-space, for the extended-STRIPS domain.

Overall benefits of the retrieval based on partial abstract similarity can be observed in Figure 5 and Figure 6, which show the percentages of overall increase in prediction accuracies (i.e., number of correct next action predictions) at both levels of abstraction, with respect to the retrieval without the abstract similarity. We evaluated the k-nearest neighbor scheme with the Euclidian distance measure. Evaluations span across several radii for a ball around a given feature vector representing the maximum allowable

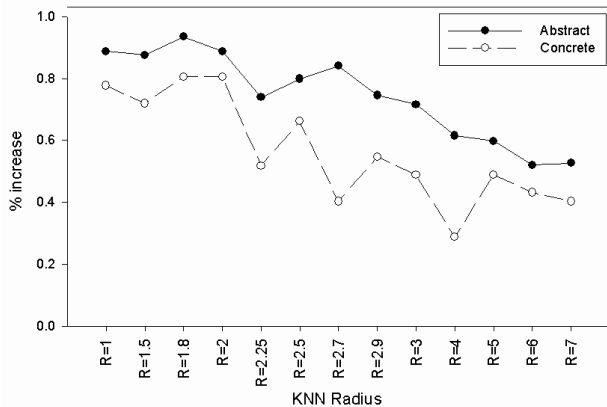


Figure 6. Percentages of increase in overall prediction accuracies with respect to retrieval without partial matches in the abstract state-space, for the logistics domain.

distance at which similar abstract states may be located. Notice that because our abstract feature vectors consist of non-negative integers, the figures show evaluations that consider all possible distance intervals between 1 and 3 inclusively, as well as several radii at larger distances, as indicated on the x-axis of evaluation figures.

The improvements in the overall prediction accuracies for the extended-STRIPS domain shown in Figure 5 are much greater than improvements in the logistics domain shown in Figure 6. As expected, accuracies are also greater at the abstract level of prediction, because appropriate adaptation of action arguments is needed for correct predictions at the concrete level. Notice that the largest increase in accuracies of the logistics domain occurs for small radii, shown by the best increases in prediction accuracy for the first four maximum distances in Figure 6. The extended-STRIPS domain shown in Figure 5 has a single best radius of 1.5, which includes all abstract states whose distance from the currently observed abstract state is larger than square root of 2 and less than square root of 3.

Evaluation results for both domains indicate that large radius values cause smaller improvements in prediction




Figure 7. Percentages of correctly predicted actions for attempts following retrieval in the abstract state-space of the extended-STRIPS domain.

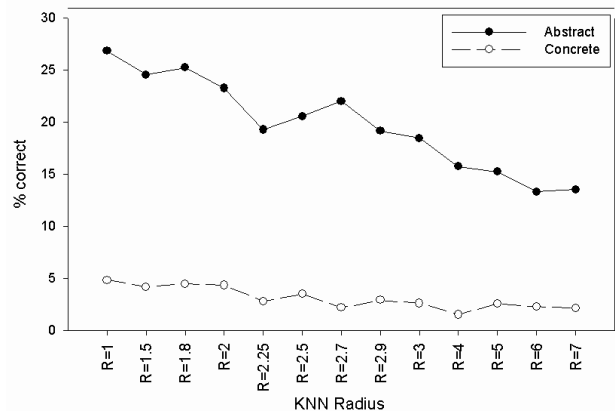


Figure 8. Percentages of correctly predicted actions for attempts following retrieval in the abstract state-space of the logistics domain.

accuracy. Figure 7 and Figure 8 represent percentages of correct predictions made after the retrieval with utilization of k-nearest neighbor similarity metric in the extended-STRIPS and the logistics domains, respectively. Note the different y-axis scales for two domains. The percentages shown concern only predictions made after k-nearest neighbor based retrieval. Essentially, the curves represent the rate of successful utilization of the abstract state-space based retrieval. Again, smaller radii seem to produce the best percentages of correct predictions, indicating that the largest overall prediction improvements coincide with the most successful prediction utilization rates.

Better recognition performance in the extended-STRIPS planning domain is likely due to the differences in state-space characteristics between these two domains. The experimental results we presented indicate that the domains where recognition process is not in its steady-state, such as the extended-STRIPS domain, benefit the most from the partial similarity matches in the abstract state-space. Furthermore, the size of the abstract state-space has a significant influence on the success of the retrieval in the abstract state-space. Due to a smaller abstract state-space, the logistics domain produces less retrieval attempts at the abstract state level, even though this domain has a larger number of overall recognition steps observed (see Figure 4) and contains about three times more observed concrete situations than the extended-STRIPS domain.

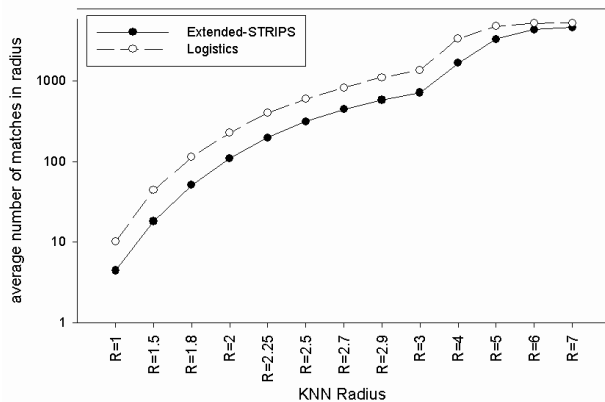


Figure 9. Logarithmic plot for average numbers of state matches within a specified radius of the k-nearest neighbor based retrieval.

Another factor that influences the effectiveness of abstract state-space based retrieval is the size of the retrieval set. In general, the size of the retrieval set is determined by the characteristics of a planning domain in question, with smaller radii resulting in a smaller number of matches produced. This can be seen in Figure 9, which shows a logarithmic plot of the average number of concrete past situations found within the radius specified during retrieval in the abstract state-space. Large retrieval sets produced by larger radii have a direct influence on the increase in prediction accuracy, as indicated by Figure 5 and Figure 6. When retrieval sets are large, the recognizer has a greater chance of choosing an incorrect prediction, resulting in a decrease in the overall prediction accuracies.

Conclusions

We presented an extension to the case-based plan recognition scheme with incomplete plan libraries, concerning the retrieval based on similarity among planning situations, represented by the states of the planner's environment. This extension utilizes a k-nearest neighbor similarity metric that enables the retrieval in the abstract state-space, which occurs when the current planning situation being observed has a novel abstract representation. Although evaluations show that the retrieval based on partial matches in the abstract state-space is effective, the degree of effectiveness depends on the state-space characteristics of a given domain. The results presented in this paper indicate that retrieval in the abstract state-space is more effective in planning domains with larger abstract state-spaces, where steady-state recognition behavior has not been reached. Our future research efforts will investigate abstract state-space based retrieval in planning domains with different state-space characteristics.

References

- Bares, M., Canamero, D., Delannoy, J. F., & Kodratoff, Y. (1994). XPlans: Case-based reasoning for plan recognition. *Applied Artificial Intelligence* 8, 617-643.
- Bergmann, R., & Wilke, W. (1996). On the Role of Abstractions in Case-Based Reasoning. In *EWCBR-96 European Conference on Case-Based Reasoning*. Springer.
- Carbonell, J. G., Blythe, J., Etzioni, O., Gil, Y., Joseph, R., Kahn, D., Knoblock, C., Minton, S., Perez, A., Reilly, S., Veloso, M., & Wang, X. (1992). *PRODIGY 4.0: The Manual and Tutorial* (Tech. Rep. No. CMU-CS-92-150). Carnegie Mellon University, Department of Computer Science, Pittsburgh, PA.
- Kerkez, B., & Cox, M. T. (2002). Local predictions for case-based plan recognition. In S. Craw & A. Preece (Eds.) *Advances in case-based reasoning: 6th European Conference, ECCBR 2002 Proceedings* (pp. 189-203). Berlin: Springer.
- Kerkez, B., & Cox, M. T. (2001). Case-based plan recognition using state indices, In D. W. Aha & I. Watson (Eds.), *Case-based Reasoning Research and Development: Proceedings of 4th international conference on case-based reasoning* (pp. 227-242). Berlin: Springer.
- Lesh, N., & Etzioni, O. (1996). Scaling up goal recognition. In *Proceedings of the Fifth Internat. Conference on Principles of Knowledge Representation and Reasoning* (pp 178-189).
- Veloso, M. (1994). Planning and learning by analogical reasoning. Berlin: Springer.
- Veloso, M., & Carbonell, J. G. (1994). Case-based reasoning in PRODIGY. In R. S. Michalski & G. Tecuci (Eds.), *Machine learning IV: A multistrategy approach* (pp.523-548). San Francisco: Morgan Kaufmann.