

When Two Case Bases Are Better Than One: Exploiting Multiple Case Bases

David B. Leake¹ and Raja Sooriamurthi²

¹ Computer Science Department, Indiana University, Lindley Hall 215
150 S. Woodlawn Avenue, Bloomington, IN 47405, U.S.A.
`leake@cs.indiana.edu`

² Department of Computer Science, University of West Florida
11000 University Parkway, Pensacola, FL 32514, U.S.A
`sraja@cs.uwf.edu`

Abstract. Much current CBR research focuses on how to compact, refine, and augment the contents of individual case bases, in order to distill needed information into a single concise and authoritative source. However, as deployed case-based reasoning systems become increasingly prevalent, opportunities will arise for supplementing local case bases on demand, by drawing on the case bases of other CBR systems addressing related tasks. Taking full advantage of these case bases will require *multi-case-base* reasoning: Reasoning not only about how to apply cases, but also about when and how to draw on particular case bases. This paper begins by considering tradeoffs of attempting to merge individual case bases into a single source, versus retaining them individually, and argues that retaining multiple case bases can benefit both performance and maintenance. However, achieving the benefits requires methods for *case dispatching*—deciding when to retrieve from external case bases, and which case bases to select—and for *cross-case-base adaptation* to revise suggested solutions from one context to apply in another. The paper presents initial experiments illustrating how these procedures may affect the benefits of using multiple case bases, and closes by delineating key research issues for multi-case-base reasoning.

1 Introduction

One of the early inspirations for case-based reasoning research was the desire to model how experiences affect individual human reasoning: how individual memories are organized, retrieved, and re-applied (Schank, 1982). Likewise, most CBR systems in research and applications capture knowledge as independent, individual systems, and studies of how to improve systems' access to good cases focus on single, unified case bases. For example, all the papers in the forthcoming *Computational Intelligence* special issue *Maintaining Case-Based Reasoning Systems* (Leake *et al.*, 2001) address issues in maintaining independent systems that each access a single case base. Drawing on a single well-maintained case base can significantly facilitate knowledge access. However, when numerous independent systems collect individual case bases, merging them into a single case

source may not be practical, and case base sharing raises its own issues: Case bases may reflect differences in their tasks, task environments, and even case representations, complicating the knowledge sharing process. This paper argues for the importance of developing methods to exploit the information available in multiple case libraries that may have been collected under differing circumstances. It surveys the issues and opportunities involved, and demonstrates that under some circumstances, two case bases can be better than one: Retaining multiple case bases can improve performance and provide useful information to aid case base maintenance.

Case-based reasoning is a natural method to support knowledge capture and reuse. Riesbeck (1996), for example, argues for the promise of CBR for intelligent components, integrated in other systems, and numerous stand-alone CBR applications exist for tasks such as diagnosis and help-desk support (e.g., (Auriol *et al.*, 2000)). The accumulation of individual case bases in these systems provides an opportunity for future CBR systems to draw not only on their own stored cases, but also on external case bases for related tasks. Just as thousands of topic-specific information sources are now available on the Web, multiple case bases may eventually provide a large-scale distributed, sharable information resource.

Previous research (e.g., (Hayes *et al.*, 1998; Martin *et al.*, 1999)) has addressed fundamental issues for distributed case bases, when those case bases are standardized. *Multi-case-base reasoning* strategically accesses and applies case bases that may have been accumulated in other contexts, for somewhat different tasks, but that still can be useful to augment the CBR system’s own competence. Multi-case-base reasoning requires supplementing the “eager” case base generation and refinement methods from case base maintenance research with a “lazy” approach to case base building: augmenting the local case base as needed by retrieving new cases from other case bases, and adjusting their solutions in light of overall differences between case bases.

Making effective use of cases from multiple idiosyncratic case bases depends on reasoning not only about the cases, but also about case properties that can be inferred from knowledge of their sources. It depends on developing methods for *case dispatching*—deciding when to retrieve from external case bases, and which case bases to select. It also depends on developing methods to address a new type of adaptation problem, *cross-case-base adaptation*, to adapt suggested solutions from one case base to apply to the needs of another. This paper examines the issues, opportunities, and tradeoffs of drawing on multiple case bases developed under differing circumstances. It first describes the motivations for this approach and the possible benefits of augmenting the information provided by individual cases with contextual information about the case bases from which they were drawn. It next presents experimental results demonstrating that when a local case base is sufficiently sparse, accessing a more competent external case base can improve performance, even if that case base reflects a different task and only simple cross-case-base adaptation is available. It closes with a proposal for key areas for future research in managing multiple case bases.

2 Problems for Case Base Combination

Most case-based reasoning systems are envisioned as reasoning based on a single case base; the goal of case-base maintenance is to assure the competence (e.g., (Smyth and McKenna, 1999a; Zhu and Yang, 1999)) and performance (e.g., (Leake and Wilson, 2000; Portinale *et al.*, 1999)) of this case base. This suggests that the natural approach to multiple case bases is to merge them into a single authoritative resource. Unfortunately, however, when multiple agents generate case bases, the practicality and benefits of merging may be decreased by problems of availability, efficiency, standardization, and maintenance.

Availability problems may arise if case bases are proprietary, requiring permission to access an entire case base to combine it. For example, in e-commerce, suppliers such as Amazon.com are willing to provide individual cases, such as records of books, to users, but not their entire case base.

Efficiency problems may arise if storing all cases locally results in the swamping utility problem for case retrieval, or if excessive case size causes space efficiency problems. If individual systems keep their own local case bases, tailored to their frequent needs, and access other case sources for individual supplementary cases as needed, these problems are avoided.

Standardization problems may arise if case features that are unspecified in the individual case bases become crucial when cases are shared. Even for closely related tasks, similar cases in different case bases may have different relevance to particular problems, because each one may implicitly reflect its different task circumstances. The AI-CBR travel case base, one of the standard benchmarks used by the CBR community, provides an example. That case base contains records of travel packages, with the information needed to match customer preferences (means of travel, destination, purpose, hotel, price of the package, etc.). It is possible to imagine world-wide travel agencies collecting sets of these types of cases individually, and combining them into a single centralized case base of recommendations. However, the cases omit a key contextual feature, without which the cases in the combined case base would be useless: the origin of the trip. Thus a client asking for the price of a trip to Paris would receive the same estimate, regardless of whether flying to Paris from London, or from Japan.

Even if two travel agencies are located within a single town, so that their locations may be similar enough not to affect the applicability of their cases, implicit aspects of their case collection process may be important to the appropriateness of their solutions. For example, if one agency has a wealthy clientele, its travel cases may tend to suggest luxurious options, so that combining its case base with the cases from an economy agency could increase the chance of a mis-match between client needs and retrieved cases. Such differences in task environments arise in many domains. For example, the advice provided by a system to diagnose engine problems and guide repairs must depend not only on the problem, but on factors such as the availability of tools, resources, and expertise to conduct the repairs. If case bases are combined, useful retrievals will depend on representing these factors, but it may be difficult to identify all the

factors relevant to case applicability in order to add explicit annotations to the case base.

Maintenance problems may also arise from combining and standardizing individual case bases. First, the combined case base may lose access to updates of the case bases from which it is drawn. In e-commerce, for example, if a collection of product case bases from different suppliers is combined into a standardized, centralized case base, it may rapidly become obsolete as the combined case base misses subsequent additions and revisions (e.g., price changes). Second, even if the original case bases remain static, a standardized version of the combined case base may become obsolete as the relationships between individual case bases change. In international e-commerce, one case base might quote prices in euros, and another in dollars. If these are combined into a case base with standardized prices in a single currency, cases become more comprehensible to the public for which they were standardized, but currency fluctuations would introduce errors after the fact.

All of the previous problems suggest difficulties that may be avoided by retaining multiple distinct case bases. In the following section, we consider the benefits that multiple case bases can provide beyond their cases alone.

3 Benefits of Multiple Idiosyncratic Case Bases

In case-based problem-solvers, the case base provides one type of information: its cases. The most basic way to exploit additional case bases is to retrieve and apply their cases, supplementing the competence of the local case base and providing a source of cases to be stored locally. If all case bases were developed in a standardized form, for standardized problems, this would be their only benefit. However, when different case bases reflect systematic differences, such as different tasks, domains, problem environments, case collection and validation procedures, or maintenance procedures, knowledge of a case's sources can enable a reasoner to make useful additional inferences that cannot be generated from the case in isolation. Because of the potential value of these inferences, the value of access to multiple idiosyncratic case bases may exceed the value of the union of the cases that they contain.

The knowledge that may be available from an external case base falls into three general categories. The first is simply *individual cases*, which supplement local competence and may be added to the local case base. The second is case base *descriptions and histories*, which provide information on the generation, previous use, and maintenance of the case base: how cases were collected, the types of problems and environment they were collected for, performance statistics, and how the case base was maintained. The third is *data for comparative analysis*, providing a source for identifying systematic differences in the contents of local and external case bases. The following sections discuss in more detail how each of these may be used.

3.1 Individual Cases

The most obvious potential benefit of accessing external case bases is to provide additional cases to augment those of the local case base, to solve problems outside its competence and possibly store those cases for future use. Another potential advantage, however, is to bring to bear additional cases even when solutions can be generated locally, in order to improve performance. In the machine learning community, research on ensemble learning has shown that combining ensembles of classifiers by weighted votes can often result in substantially better performance than the individual classifiers (Dietterich, 2000). The ramifications of ensemble methods for CBR are comparatively unexplored, but are potentially promising (for example, see (Cunningham and Zenobi, 2001) for a recent study of insights that they may provide for case representation). A necessary condition for the success of ensemble methods is that the individual classifiers be diverse, making their errors on different examples. Here not only the access to external cases, but also their diversity arising from case base differences, may be advantageous.

3.2 Case Base Descriptions and Histories

Pre-compiled descriptive information about a case base can provide a valuable guide to when its cases may apply. For example, knowledge about the source of a case base (e.g., the organization that generated it) can be used to assess its trustworthiness. If more detailed information is available, such as detailed task descriptions or information about how cases are collected, that information can also be used to determine the likely applicability of cases. Likewise, information on maintenance policies and the maintenance schedule can help in assessing the timeliness of the case base and can provide additional information about its likely reliability. All this information provides value to a system developer or maintainer beyond the value of the cases out of context. Ideally, this information could be made available to an automated case base selection process, to support large-scale case base sharing. To do so, however, will require methods for standardized branding and descriptions of case base contents and characteristics.

3.3 Data for Comparative Analysis

When two distinct case bases are available, analysis of their differences can provide additional useful information for guiding case base maintenance and standardization. When the two case bases are each internally consistent, and each is reasonably homogeneous (e.g., reflecting the solution preferences of a single user), their systematic differences may reflect important general case base characteristics. Multiple case bases may have three main types of differences:

1. Differences in their indices and case representations (here we focus on differing domain content, rather than differing forms).

2. Differences in their competence or problem distributions.
3. Differences in the solutions they suggest for problems in the intersection of their coverage, given consistent indexing schemes and fixed adaptation knowledge.

Detecting these types of differences can provide important information during the maintenance process. Although the differences are not guaranteed to be significant, they are a useful focusing device for determining whether case base maintenance is needed and how it should be applied, as well as for guiding cross-case-base knowledge application. For example, case base comparison may be useful for:

1. **Assessing case base reliability:** When a new case base overlaps with one whose trustworthiness is already known, comparison of the solutions for the overlapping cases can provide a reliability estimate for the new case base, which can then be extrapolated to non-overlapping case regions.
2. **Guiding standardization:** When two case bases use different representations and feature sets, but produce similar solutions, differences in their representations may reflect alternative representations of equivalent features. This suggests the potential for re-representation in standard form, or for the development of translation criteria for rerepresenting cases for transfer from one case base to another.
3. **Suggesting case base applicability conditions:** Differences in competence or problem distributions can provide information about the types of problems that tend to occur in different task environments, in turn suggesting problem distributions for which a particular case base is likely to be a useful resource. For example, if a particular travel agent has an extensive case base applicable to one region, that information may be useful for characterizing the case base, in order to facilitate its selection for problems relevant to that area.
4. **Guiding feature discovery:** When two case bases are both believed reliable, but provide divergent solutions on similar problems, their divergence suggests a possible gap in case representation: that relevant features of the task or environment in which the case bases were applied were left unstated. When both case bases use different feature sets, those differences provide a first suggestion of additional features to examine. For example, if two case bases have differing problem representations, and one case base provides greater accuracy, a useful heuristic is to consider whether the problem descriptions used in that case base are more appropriate.
5. **Guiding case discovery:** The existence of problems in one case base but not in another provides information about real problems that are not covered, suggesting possible competence gaps to fill.
6. **Guiding cross-case-base adaptation:** Once systematic differences have been identified between case bases, they may enable automatic conversion of cases from one case base to apply in another context. For example, if one product case base states prices in euros, and another in dollars, comparing

average prices (or prices for similar items) may provide an approximate euro-dollar conversion factor, enabling the system to use cases from one case base to predict the prices for missing items in the other. We demonstrate a simple application of this type of derived adaptation knowledge in the next section.

Fully exploiting solution relationships requires explaining whether observed patterns are actually significant—whether they can be explained in terms of the task or environment. In general, this may be difficult. For example, some case-based travel planners might accept plans with little margin for error if they conserve resources (e.g., accept travel plans with tight connections); others might favor more forgiving routings. When differences arise and cannot be explained internally, case pairs reflecting differences can be presented to a user or maintainer to explain (cf. (Shimazu and Takashima, 1996), which applies a similar approach to identify problems within a single case base).

However, it may sometimes be possible to reason from correlations, even without explanations. In the following experiments, we show how a simple calculation of the difference in ranges of numerical predictions can be used to improve performance by cross-case-base adaptation in the travel domain.

4 Experiments

To illustrate the issues involved in reasoning from multiple case bases, we performed initial experiments on the benefits of drawing on related but distinct case bases. Our goal was to explore how CBR system performance is affected by the interaction of (1) the competence of the local case base, (2) the case dispatching criteria used, (3) the availability of cross-case-base adaptation, and (4) the use of solution combination to exploit the availability of diverse case sources.

The system’s reasoning task was predicting the prices of travel packages. The data used were drawn from the AI-CBR travel case base at www.ai-cbr.org, which contains 1470 instances. Because the indicated prices vary widely for similar trips, after defining feature weightings for distance-weighted k-NN retrieval we selected a 681-case subset with reasonable problem-solution regularity. We defined a very simple case adaptation function: The prices of prior packages were adjusted proportionate to differences in their duration and number of travelers. More refined feature weights and adaptation criteria would have improved predictions, but because our goal was comparative, to study relative effects of drawing on an external case base, we did not tune the basic system.

To generate two case bases for related but distinct task contexts, we divided the travel cases according to hotel star ratings, simulating the division in travel packages that might arise for cases collected by two travel agencies, one catering to luxury and the other to economy travelers. The luxury (3–5 stars) case base, CB_1 , contained 352 cases. The economy (1–2 stars) case base, CB_2 , contained 329 cases. In the following runs the star weighting was suppressed by giving that feature a zero weighting.

4.1 Performance of Individual and Combined Case Bases

We first compared the predictive accuracy of CB_1 , CB_2 , and $CB_1 \cup CB_2$, each tested on itself by leave-one-out cross validation for distance-weighted 3-NN retrieval. This compares the performance of processing cases in the most appropriate individual case base, versus in a combined version. Performance was measured in two ways, (1) the prediction accuracy—defined as the percent of problems whose prices were predicted within 20% of the correct price—and (2) the average percent error of predicted prices. We expected that performance for each individual case base would be superior to the performance of $CB_1 \cup CB_2$, but that differences would be comparatively small, due to the availability of other features in the case representation (e.g., the hotel name) that should correlate with the luxury of the trip. In fact, the result was more marked than expected. When all cases were processed by $CB_1 \cup CB_2$, 61% of the cases were predicted with accuracy within 20% of the correct values, and the average percent error for predictions was 28%. When problems were dispatched to their corresponding case base, accuracy was 83% with an average percent error of 13% for the luxury case base, and 72% with 17% percent error for the economy case base. This illustrates how case base combination may impair performance by blurring distinctions in the task environment (for this example, the types of trips considered). Its main interest, however, is as a backdrop for the next experiment. That experiment will show that despite this performance drop from combining CB_1 and CB_2 , strategically drawing on CB_2 to augment the knowledge in CB_1 can actually *improve* performance compared to using CB_1 alone.

4.2 Augmenting a Local Case Base with External Retrievals

When multiple idiosyncratic case bases are available, a central question is whether those case bases can be used effectively. Because drawing on external case bases may be especially important when the local case base has limited competence (e.g., in the early phases of a CBR system or when case storage space is limited), we explored the effects of case dispatching starting from a set of local case bases of different sizes. Starting from CB_1 , we generated a series of case bases CB_1^* , of varying sizes, each consisting of a randomly-selected subset of CB_1 . CB_1^* simulates an incomplete “local” case base of a growing system.

In our tests, CB_2 (the economy travel case base) functions as an external case base that can be drawn on as a supplement to CB_1^* (a sparse version of the luxury travel case base). Problems from CB_1 (the full luxury travel case base) are used to test the predictive accuracy of CB_1^* by leave-one-out cross validation. Test problems from CB_1 are first directed to CB_1^* , which can either handle the query locally or dispatch it to CB_2 . The decision whether to dispatch a problem to CB_2 is made by calculating its average distance from the k closest cases in CB_1^* . If this distance exceeds a fixed threshold, the problem is dispatched to CB_2 . If CB_2 contains a closer case, the case from CB_2 is used, possibly with cross-case-base adaptation; otherwise, the system reverts to the solution from CB_1^* . Our tests used a very simple cross-case-base adaptation method: The range of

prices from CB_2 was linearly interpolated to map to the range of prices in CB_1 , and predictions from CB_2 were adjusted for case base differences by multiplying them by the corresponding scale factor.

We expected that results would be superior when most problems were solved directly by CB_1^* , because both CB_1^* and the input test cases from CB_1 involve luxury travel. However, when luxury travel problems are solved using the economy case base, it may be possible to compensate somewhat by performing cross-case-base adaptation on estimates from CB_2 , to correct for the generally lower prices of the economy travel.

Our tests compared the predictive accuracy of:

1. **CB_1^*** : Predicting using only CB_1^* .
2. **$CB_1^* + CB_2$** : Predicting using CB_1^* if the input problem is within a fixed distance threshold of a case in CB_1^* ; else dispatching the problem to be solved by CB_2 if CB_2 contains a case closer to the current problem.
3. **$CB_1^* + CB_2 + \text{cross-CB adaptation}$** : Predicting as in (2), except that cross-case-base adaptation is applied to the solutions from CB_2 .
4. **Combined solution**: Predicting by averaging the prediction from method (2) with the prediction from method (3).

Because the results of methods 3–4 depend on the dispatching threshold, we tested 11 different distance thresholds defined so that the dispatch rates ranged from 100% (i.e., all cases sent to CB_2) to 0% (i.e., all problems solved locally by CB_1^*), in decrements of 10%. Testing was done with random CB_1^* s of sizes 60%, 20%, 5%, 2%, and 0.5% the size of CB_1 . The full experimental setup consisted of 30 random CB_1^* s (in 3 groups of 10 each) for each of the 5 case sizes and 11 distance thresholds, for a total of 1650 runs. The results across each group of 10 case bases were averaged.

As mentioned in Section 3.1, a potential advantage of access to multiple case bases is providing a diverse set of predictions that may be suitable for ensemble methods. Method (4) above, the combined method, tests a very simple form of combining results from two case bases with divergent characteristics.

Figures 1 and 2 illustrate the averaged results of 10 runs for different sparsity levels for random CB_1^* . When CB_1 (the luxury case base) is tested on its own problems by leave-one-out cross validation, its prediction accuracy is 75%. In the left-hand graph of Figure 1, CB_1^* has 60% the size of CB_1 and on average can correctly solve 73% of the cases in CB_1 —Drops in average accuracy levels as case bases become smaller reflect the expected decrease in competence when fewer cases are available. For this case base size, none of the methods match the performance of CB_1^* at any dispatching rate. However, even the simple cross-case adaptation strategy improves average performance compared to simply drawing on the supplementary case base. Despite the simplicity of the cross-case-base adaptation method used, its benefit compared to dispatching alone was highly consistent across all trials. The right-hand graph shows the average performance for CB_1^* with 20% of the cases, for which CB_1 alone achieves 64% predictive accuracy. Here the proportionate benefit of cross-case-base adaptation is greater. Also, the combined method slightly outperforms CB_1^* .

The left-hand side of Figure 2 shows average performance for a set of 10 CB_1^* s selected to contain 2% of the cases in CB_1 . Though the average accuracy of CB_1^* drops to 47%, due to the sparser case base, cross-case-base adaptation has a more pronounced effect, sometimes resulting in performance slightly superior to CB_1^* , and the combined prediction method markedly outperforms CB_1^* . In the right-hand side, for an even sparser case base that solves an average of 32% of the test problems, all alternative methods surpass the original case base for some dispatching rates, and both dispatching with cross-case-base adaptation and the combined method do consistently better than CB_1^* for high dispatching rates. The small peak at 70% of cases dispatched may illustrate the balance between problem similarity and cross-case-base differences. When larger percentages are dispatched, cases in CB_1^* that would give better predictions than cross-case-base adaptation are being bypassed; when smaller percentages are dispatched, the error in predicting from distant cases in CB_1^* is greater than the error introduced by case base differences. We have observed similar but more marked peaks for $CB_1^*+CB_2$ +cross-CB adaptation in individual runs.

Figure 3 illustrates two specific examples for the 5% CB_1^* , which is able to solve 53% of the problems. In the left-hand graph, the combined prediction method is again best, followed by $CB_1^*+CB_2$ +cross-CB-adaptation, both of which noticeably outperform CB_1^* and $CB_1^*+CB_2$. In the right-hand graph, both dispatching with cross-case-base adaptation and the combined method outperform CB_1^* alone, with slightly superior performance from $CB_1^*+CB_2$ +cross-CB-adaptation.

We believe that the variation in results of individual runs is accounted for by our very simple linear cross-case-base adaptation function. When the range of values in CB_1^* corresponds naturally to the range in CB_2 , performance is maximized. When it does not, performance suffers. We expect that case dispatching would be more helpful for denser case bases if this function were refined. We are also investigating methods for assessing the quality of cross-case-base adaptation during processing, to predict when (and whether) dispatching will be useful, in order to adjust dispatching criteria.

Even with the present cross-case-base adaptation function, the results suggest that in the early phases of a CBR system, when its case base is sparse, dispatching selected cases to a denser external case base can improve performance, even if the external case base is suboptimal for the task. Because the combined method consistently outperformed CB_1^* , they also suggest the potential value of using multiple case bases with cross-case-base adaptation to provide diverse data sources for ensemble predictions.

5 Towards “Case Boutiques:” A Research Agenda

The previous discussion argues that comparing multiple case bases can provide valuable information for case base maintenance, and the experiments show that drawing on even a sub-optimal external case base can help supplement a system’s own case base. Exploiting multiple case bases, however, depends first on their




Fig. 1. Predictive accuracy for CB_1^* containing 60% (left) and 20% (right) of the cases in CB_1 , averaged over 10 runs.

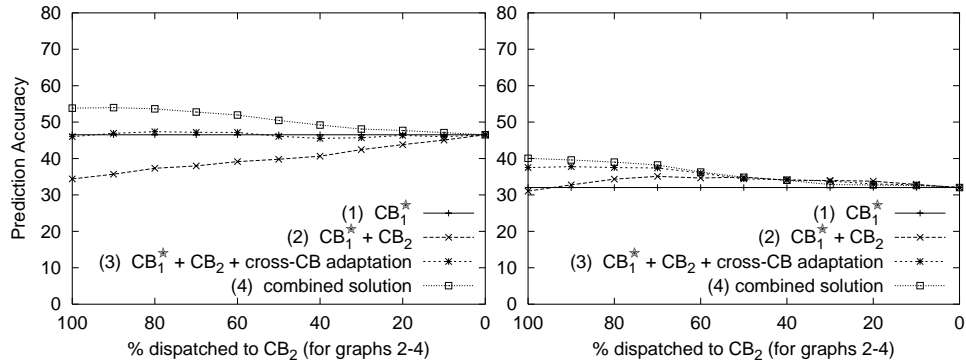


Fig. 2. Predictive accuracy for CB_1^* containing 2% (left) and .5% (right) of the cases in CB_1 , averaged over 10 runs.




Fig. 3. Sample individual results for CB_1^* containing 5% of the cases in CB_1 .

availability. Thus a central research issue is how to build up sharable task-based case libraries and to describe their contents in a way that will permit the right case bases to be identified efficiently. Data warehousing research and applications suggest a useful parallel. With terabyte storage now available, storage of available data is not a major issue, but effective access is. This has led to the development of “data boutiques” that provide specialized applications for particular tasks. Analogously, “case boutiques” can be developed to provide similar access.

The idea of case base sharing is not new; as early as 1995, Inference Corporation formed a knowledge publishing division to sell case bases for particular tasks. However, this approach was viewed as a means for a one-time “jump start” to building an individual case base, rather than part of a supplemental resource to augment a local case base on demand. Exploiting multiple case bases on demand requires addressing new multi-case-base issues for steps that parallel the basic steps of CBR, but that apply to case bases as a whole, rather than to individual cases:

- **Situation assessment:** Determining the general task context in which a particular problem is being solved, to express it in a vocabulary compatible with the case base description vocabulary.
- **Case base indexing:** Characterizing the types of task contexts and problems for which a particular case base may be useful. This requires a vocabulary to describe overall task types and solution characteristics, competence characteristics (e.g., areas of high density), representations used, etc.
- **Multi-source retrieval:** Determining when to dispatch cases to external case bases, which sources are most appropriate, and how to convert indices to apply to the external case base’s own retrieval mechanisms.
- **Cross-case-base adaptation:** Revising retrieved cases’ representations and contents based on general characteristics of the cases in the external case base.
- **Multi-case-base maintenance:** Determining how to distribute cases—when to split, merge, or standardize collections of cases. Case deletion in standard case base maintenance research may be replaceable by export of cases to new case bases; and case discovery by importation of new cases from other case bases. Multi-case-base maintenance issues also include determining how to apply comparative information to focus the maintenance process for individual cases.

We believe that addressing these issues, thus developing a foundation for “case boutiques,” is a promising way to leverage the independent knowledge of individual CBR systems.

6 Related Work

6.1 Hierarchical Retrieval and Web Source Selection

The idea of dispatching cases to particular case bases is related to research on hierarchical retrieval (Watson and Perera, 1997) and footprint similarity (Smyth

and McKenna, 1999b). Both of these methods can be seen as determining a region to which to dispatch an input case, within a single case base. Likewise, the potential growth of sharable independent case bases has parallels to the increasing availability of specialized information sources on the web, and the issues involved in developing methods to determine which web sources to access for a particular query (e.g., (Leake and Scherle, 2001; Sugiura and Etzioni, 2000)). Just as Apple's Sherlock XML-style plug-ins encode wrapper information about search engines, analogous methods could be used to facilitate access to external case bases.

Some existing web data, such as FAQ files, already provide a resource that has been used for on-demand exploitation by textual CBR techniques (Burke *et al.*, 1997). Issues in how to convert between different representations for information are being addressed in work on ontologies on the *wrapper generation problem* (e.g., (Ashish and Knoblock, 1997)).

6.2 Relationship to Multi-Agent and Distributed CBR

(Prasad *et al.*, 1996) describe an approach in which multiple agents each cooperatively access their individual case bases to contribute subparts of a solution. (Martin *et al.*, 1999) describes an approach to knowledge reuse in which peer agents each maintain independent cases and share them as needed. Both these situations differ from the current task in assuming that all cases have a consistent representation and consistent solutions, removing the need for cross-case-base adaptation and the need for comparative inferences. (Hayes *et al.*, 1998) present methods for efficient distributed CBR when using a single standardized case format, and propose a potential method to facilitate case communication: CBML, an XML application to serve as a standard for large-scale case distribution. Of work on distributed CBR, the most closely related is (McGinty and Smyth, 2001), in which retrievals from multiple case bases are used to compensate for experience gaps when recommending travel routes reflecting individual preferences.

6.3 Relationship to Case Base Maintenance

Current case base maintenance research focuses on the issues of generating high-quality individual case bases, with the aim of producing a unified, consistent body of cases for a given problem class. These methods provide a valuable means to improve the quality of the case bases to be exploited by multiple-case-base methods. However, they are based on the fundamental assumption that all cases are available for combination, that all are collected for a single task and environment, and that the goal is to eagerly process cases. Multi-case-base reasoning must instead address the issues of facilitating lazy case access on demand. Thus it is crucially concerned not only with assuring the quality of cases, but also the quality of information about case bases as a whole. As described in section 3.3, however, case base comparison may prove a valuable tool for focusing maintenance of individual case bases.

7 Conclusion

The increasing use of CBR systems and web-based communication provides an opportunity to improve the performance of case-based reasoners by developing methods for strategically combining their local case information with cases drawn from external sources containing similar cases, even if those case bases may be designed for different task environments. This paper has considered the benefits and difficulties of using multiple case bases that may reflect different tasks and environments. It has illustrated that the ability to dispatch cases to an alternative case base—even when that case base reflects systematic task differences—can help improve system performance, provided that solution transfer is supported by cross-case-base adaptation.

Making effective use of external case bases requires developing methods for *case dispatching*—deciding when a problem should be handled by retrieval from an external case base, and which external case base to select—and for performing *cross-case-base adaptation* to transform solutions to fit new contexts. By combining these operations with the standard case-based reasoning process, a CBR system may supplement its competence through just-in-time knowledge access from external knowledge sources.

Acknowledgments

David Leake’s research is supported in part by NASA under awards NCC 2-1035 and NCC 2-1216. Raja Sooriumurthi’s research is supported in part by a professional development leave from the University of West Florida. The authors would like to thank the anonymous reviewers for their helpful comments.

References

- [Ashish and Knoblock, 1997] N. Ashish and C. Knoblock. Wrapper generation for semi-structured internet sources. In *SIGMOD Record*, volume 26, pages 8–15, 1997.
- [Auriol *et al.*, 2000] E. Auriol, P. Avesani, B. Bartsch-Spoerl, M. Goeker, M. Manago, K. Pulaski, and W. Wilke, editors. *Workshop Notes, Workshop on Innovative Customer Centered Applications (ICCA)*, Trento, Italy, 2000. EWCBR-2K.
- [Burke *et al.*, 1997] R. Burke, K. Hammond, V. Kulyukin, S. Lytinen, N. Tomuro, and S. Schoenberg. Question answering from frequently asked question files: Experiences with the FAQFinder system. *AI Magazine*, 18(2):57–66, Summer 1997.
- [Cunningham and Zenobi, 2001] P. Cunningham and G. Zenobi. Case representation issues for case-based reasoning from ensemble research. In *Case-Based Reasoning Research and Development: Proceedings of the Fourth International Conference on Case-Based Reasoning*, Berlin, 2001. Springer Verlag. In press.
- [Dietterich, 2000] T. Dietterich. Ensemble methods in machine learning. In *Proceedings of the First International Workshop on Multiple Classifier Systems*, pages 1–15, Berlin, 2000. Springer Verlag.
- [Hayes *et al.*, 1998] Conor Hayes, Pádraig Cunningham, and Michelle Doyle. Distributed CBR using XML. In *Proceedings of the KI-98 Workshop on Intelligent Systems and Electronic Commerce*, 1998.

- [Leake and Scherle, 2001] D. Leake and R. Scherle. Towards context-based search engine selection. In *Proceedings of the 2001 International Conference on Intelligent User Interfaces*, pages 109–112, 2001.
- [Leake and Wilson, 2000] D. Leake and D. Wilson. Remembering why to remember: Performance-guided case-base maintenance. In E. Blanzieri and L. Portinale, editors, *Proceedings of the Fifth European Workshop on Case-Based Reasoning*, pages 161–172, Berlin, 2000. Springer Verlag.
- [Leake et al., 2001] D. Leake, B. Smyth, Q. Yang, and D. Wilson, editors. *Maintaining Case-Based Reasoning Systems*. Blackwell, 2001. Special issue of *Computational Intelligence*. In press.
- [Martin et al., 1999] F. Martin, E. Plaza, and J.-L. Arcos. Knowledge and experience reuse through communications among competent (peer) agents. *International Journal of Software Engineering and Knowledge Engineering*, 9(3):319–341, 1999.
- [McGinty and Smyth, 2001] L. McGinty and B. Smyth. Collaborative case-based reasoning: Applications in personalised route planning. In *Case-Based Reasoning Research and Development: Proceedings of the Fourth International Conference on Case-Based Reasoning*, Berlin, 2001. Springer Verlag. In press.
- [Portinale et al., 1999] L. Portinale, P. Torasso, and P. Tavano. Speed-up, quality, and competence in multi-modal reasoning. In *Proceedings of the Third International Conference on Case-Based Reasoning*, pages 303–317, Berlin, 1999. Springer Verlag.
- [Prasad et al., 1996] M. V. Nagendra Prasad, V. Lesser, and S. Lander. Reasoning and retrieval in distributed case bases. *Journal of Visual Communication and Image Representation*, 7(1):74–87, 1996.
- [Riesbeck, 1996] C. Riesbeck. What next? The future of CBR in postmodern AI. In D. Leake, editor, *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, pages 371–388. AAAI Press, Menlo Park, CA, 1996.
- [Schank, 1982] R.C. Schank. *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge University Press, Cambridge, England, 1982.
- [Shimazu and Takashima, 1996] H. Shimazu and Y. Takashima. Detecting discontinuities in case-bases. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, volume 1, pages 690–695, Menlo Park, CA, 1996. AAAI Press.
- [Smyth and McKenna, 1999a] B. Smyth and E. McKenna. Building compact competent case-bases. In *Proceedings of the Third International Conference on Case-Based Reasoning*, pages 329–342, Berlin, 1999. Springer Verlag.
- [Smyth and McKenna, 1999b] B. Smyth and E. McKenna. Footprint-based retrieval. In *Proceedings of the Third International Conference on Case-Based Reasoning*, pages 343–357, Berlin, 1999. Springer Verlag.
- [Sugiura and Etzioni, 2000] A. Sugiura and O. Etzioni. Query routing for web search engines: Architecture and experiments. In *Proceedings of Proceedings of the Ninth World Wide Web Conference (WWW9)*, pages 417–429, 2000.
- [Watson and Perera, 1997] I. Watson and R. Perera. The evaluation of a hierarchical case representation using context guided retrieval. In *Proceedings of the Second International Conference on Case-Based Reasoning*, pages 255–266, Berlin, 1997. Springer Verlag.
- [Zhu and Yang, 1999] Jun Zhu and Qiang Yang. Remembering to add: Competence-preserving case-addition policies for case base maintenance. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 234–241. Morgan Kaufmann, 1999.