Demand-Driven Discovery of Adaptation Knowledge

David McSherry
School of Information and Software Engineering
University of Ulster
Coleraine BT52 ISA
Northern Ireland

Abstract

A case-based approach to adaptation for estimation tasks is presented in which there is no requirement for explicit adaptation knowledge. Instead, a target case is estimated from the values of three existing cases, one retrieved for its similarity to the target case and the others to provide the knowledge required to adapt the similar case. With recursive application of the adaptation process, any problem space can be fully covered by fewer than *nk* selected cases, where *n* is the number of case attributes and *k* is the number of values of each attribute. Moreover, a *k* x *k* problem space is fully covered by any set of *2k* - 1 known cases provided there is no redundancy in the case library. Circumstances in which the approach is appropriate are identified by theoretical analysis and confirmed by experimental results.

1 Introduction

Addressing the problems of automatic adaptation has been identified as a central problem for the future of CBR [Leake, 1996]. Recent advances promise to reduce the overhead associated with the acquisition of adaptation knowledge by automating its discovery from case data (Hanney and Keane, 1996; Smyth and Keane, 1998], or by a case-based approach to adaptation in which *adaptation cases* are acquired from user adaptations or learned from adaptation experience [Leake *et al* 1996; Leaked a/., 1997].

This paper presents a case-based approach to adaptation for estimation tasks that differs from other approaches to adaptation in that there is no requirement for explicit adaptation knowledge. Instead, a target case is estimated from the values of three existing cases, one retrieved for its similarity to the target case and the others to provide the knowledge required to adapt the similar case. The transformation applied to the similar case is based on a simple adaptation heuristic called the *difference heuristic* [McSherry, 1998] or an alternative called the *ratio heuristic.*

The demand-driven discovery of adaptation knowledge that characterizes the approach is consistent with the demand-driven approach to problem solving that characterizes CBR itself [Aha, 1998]. The techniques presented have been implemented in a case-based reasoner for estimation tasks called CREST (Case-based Reasoning for ESTimation). Estimation tasks to which CBR has been applied include cost estimation of software projects [Bisio and Malabocchia, 1995], property valuation [Gonzalez and Laureano-Ortiz, 1992], and estimation of sentence duration in JUDGE [Riesbeck and Schank, 1989].

The retrieval criteria and adaptation heuristics used in CREST are presented in the following section. In Section 3, coverage of the problem space is shown to be significantly increased by recursive application of the adaptation process. In Section 4, experimental results are presented which show that certain departures from the assumptions on which the difference heuristic is based can be tolerated at the expense of some loss of coverage.

2 Case Retrieval and Adaptation

The value of a case is assumed to be determined by an unknown function /of its attributes A,, X_2....A. all of which are assumed to be discrete. The task of the case-based reasoner is to estimate the value of the case function for a target case from the values of existing cases in a case library. A problem space P of dimensions $K_1 \times K_2 \times ... \times k_n$ is defined by the cartesian product of the attribute domains D_1 D_2, D_n. A case is represented in CREST as a tuple $C = (a_1, a_2, ..., a)$, where for $1 \le i \le n$, a_i is the value of the attribute x_i for *C*. The known (or estimated) value of each case is also stored in the case library.

Table 1. Example library of 7 cases in the cylinder volume domain.

r	*h*	*Volume*
1	1	3.14
1	2	6.28
2	4	50.24
3	2	56.52
3	3	84.78
4	1	50.24
4	4	200.96

Case-based estimation of the volume of a cylinder provides an example estimation task in which the case function is known. In this domain, the case attributes are the radius (r) and height *(h)* of the cylinder. Seven cases in a 4 x 4 problem space are shown in Table 1. As illustrated in Figure 1, the problem space can be visualised as an array of points in the plane consisting of 4 rows and 4 columns. The filled circles represent the 7 known cases in the case library. The remaining points in the problem space are unknown cases which the case-based reasoner may be called upon to estimate.

The steps involved in the use of the difference heuristic to estimate the value of a target case C_1 are outlined below.

Step 1. Retrieve three existing cases C_2, C_3, C_4 such that C_1 and C_2 differ only in the value of a single attribute, C_3 and C_4 also differ only in the value of the same attribute, and the values of the distinguishing attribute for C_3 and C_4 are the same as its values for C_1 and C_2 respectively. The triple C_2, C_3, C_4 of retrieved cases is called an *estimation triple* for C_1.

Step 2a. According to the difference heuristic, an estimate for the value of the target case is:

$$est_d(C_1) = est(C_2) + est(C_3) - est(C_4)$$

where *est*(C_2), *est*(C_3) and *est*(C_4) are the known (or estimated) values of C_2, C_3 and C_4 in the case library.

Adaptation based on the ratio heuristic differs only in the transformation applied to the value of the similar case.

Step 2b. According to the ratio heuristic, an estimate for the value of the target case is:

$$est_r(C_1) = \frac{est(C_2)\, est(C_3)}{est(C_4)}$$

To avoid division by zero, a necessary condition for its application is that *est*(C_4) > 0. The rationale for the difference heuristic is that provided the case function is an additive function of the form:

$$f(x_1, x_2, \ldots, x_n) = \sum_{i=1}^{n} g_i(x_i)$$

then for any case C_1 and estimation triple C_2, C_3 C_4 for C_1, $f(C_1) = f(C_2) + f(C_3) - f(C_4)$. It follows that the difference heuristic always gives the correct value for a target case provided the case function is an additive function and the values of existing cases are known without error [McSherry, 1998]. Similar justification of the analogous ratio heuristic follows an example of its use to estimate the volume of a target case C_1 = (4,2) in the cylinder volume domain.

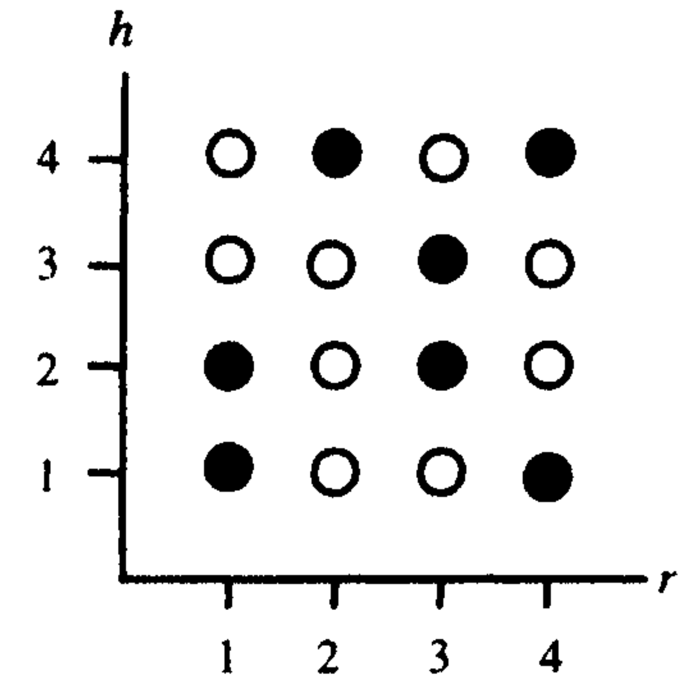

Figure 1. A 4 x 4 problem space in the cylinder volume domain.

An existing case which differs from the target case only in its radius is C_2 = (1,2), with an estimated volume of 6.28. Two existing cases are now retrieved from the case library which also differ only in their radii and which have the same radii as C_1 and C_2 respectively. Two such cases are C_3 = (4,1) and C_4 = (1,1). According to the ratio heuristic, an estimate for the value of the target case is therefore:

$$est_r(4,2) = \frac{est(1,2)\, est(4,1)}{est(1,1)} = \frac{6.28 \times 50.24}{3.14} = 100.48.$$

Since the volume of a cylinder is $\pi r^2 h$, the correct volume of the target case, based on the same approximate value of as the values of the known cases, is 3.14 x 4 x 2 = 100.48. As the following theorem shows, the ratio heuristic always gives the correct value for a target case provided the case function is multiplicative and the values of existing cases are known without error.

Theorem 1 *Provided the case function is a multiplicative function of the form:* $f(x_1, x_2, \ldots, x_n) = \prod_{i=1}^{n} g_i(x_i)$, *then for any target case* C *and estimation triple* C_2, C_3 C_4 for C_1,

$$f(C_1) = \frac{f(C_2) f(C_3)}{f(C_4)}.$$

The proof is similar to that of the analogous result for additive case functions [McSherry, 1998]. In addition to the target case (4,2) in the above example, four other cases can be estimated by adaptation based on the values of known cases. However, there remain 4 cases that cannot be estimated. For example, there is no estimation triple in the case library for the unknown case (2,3). Coverage of the problem space by the 7 known cases, defined as the percentage of cases that are known or can be estimated, is therefore only 75%. As shown in the following section, problem-space coverage is significantly increased by relaxing the condition that all 3 cases in an estimation triple for a target case must be known cases.

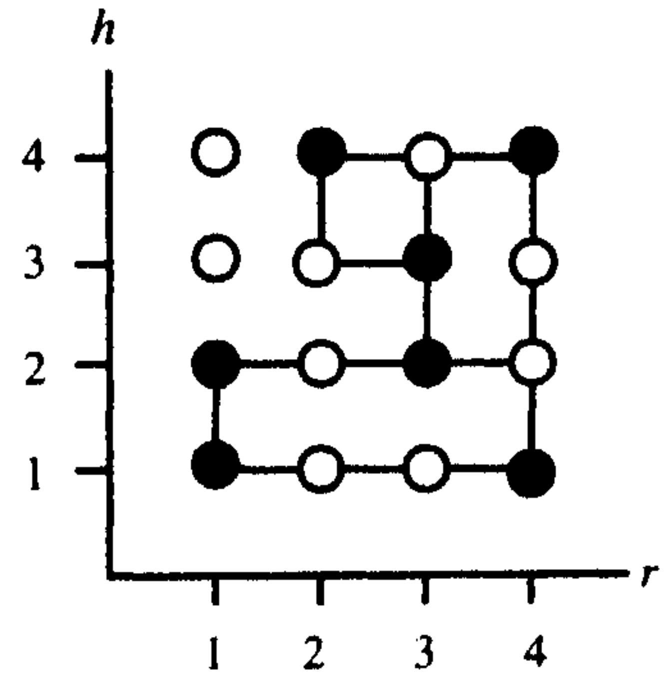

Figure 2. Recursive adaptation in the cylinder volume domain.

3 Recursive Adaptation

If the value of one of the cases in an estimation triple for a target case is unknown, it may be possible to estimate the value of the unknown case, which may in turn require the estimation of another unknown case. Recursive adaptation in the cylinder volume domain, with C_1 = (2,3) as the target case, is illustrated in Figure 2.

An estimation triple for the target case is provided by the known cases C_2 = (3,3) and C_3 = (2,4) and the unknown case C_4 = (3,4). An estimation triple for (3,4) is provided by the known cases (4,4) and (3,2) and the unknown case (4,2). Finally, an estimation triple for (4,2) is provided by the known cases (1,2), (4,1) and (1,1). According to the ratio heuristic, an estimate for the value of the target case is therefore:

$$est_r(2,3) = \frac{est(3,3)\ est(2,4)}{est_r(3,4)} = 4259.347 \times \frac{est_r(4,2)}{est(4,4)\ est(3,2)}$$

$$= 0.375 \times \frac{est(1,2)\ est(4,1)}{est(1,1)} = 37.68.$$

Again this is the correct value of the target case according to the known case function. In fact, all the unknown cases in the problem space can be estimated by a sequence of estimation triples in which at most one case is unknown. So even with recursive adaptation restricted to estimation triples in which at most one case is unknown, full coverage of the problem space is provided by the 7 known cases.

Recursive adaptation in a 3 x 3 x 5 problem space in the domain of property valuation is illustrated in Figure 3. The case attributes are location (1, 2 or 3), style (1, 2 or 3) and bedrooms (1, 2, 3, 4 or 5). Suppose that the value of a property in £1,000 is determined by the hypothetical case function:

$$f(loc, style, beds) = 20 + 10 \times 2^{loc-1} + 5 \times 2^{style-1} + 2 \times beds$$

and that the values of existing cases are known without error. An estimation triple for the target case C_1 = (1,3,3) is provided by the known cases C_2 = (2,3,3) and C_3 = (1,2,2)

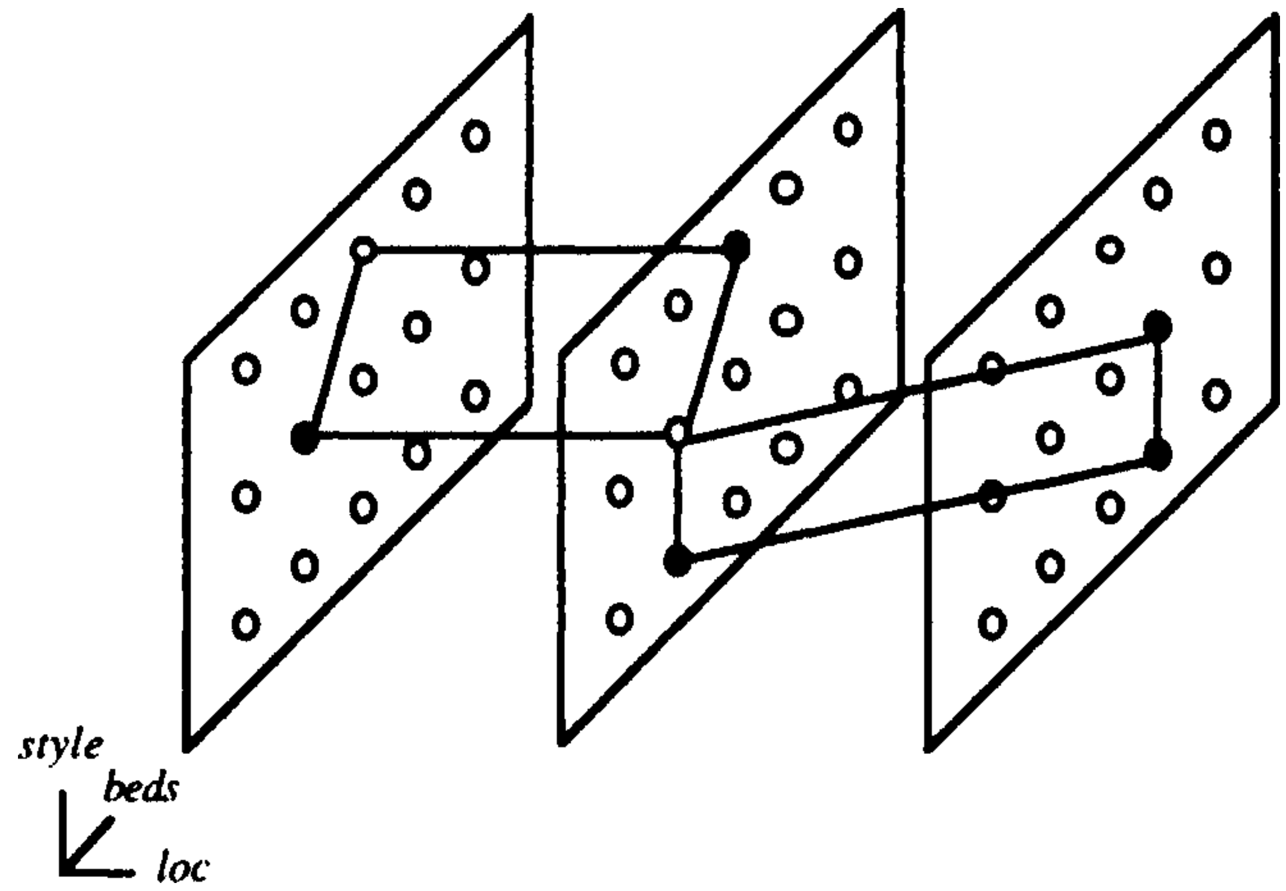

Figure 3. Recursive adaptation in a 3 X 3 x 5 problem space.

and the unknown case C_4 = (2,2,2). An estimation triple for (2,2,2) is provided by the known cases (2,1,2), (3,2,4) and (3,1,4). According to the difference heuristic, an estimate for the value of the target case is therefore:

$$est_d(1,3,3) = est(2,3,3) + est(1,2,2) - est_d(2,2,2)$$

$$= 66 + 44 - est(2,1,2) - est(3,2,4) + est(3,1,4) = 56$$

Since the case function is additive, this is the correct value of the target case.

3.1 Recursive adaptation in CREST

An algorithm for recursive adaptation based on the adaptation heuristic selected by the user has been implemented in CREST. The algorithm gives priority to estimation triples consisting only of known cases. If such an estimation triple for a target case does not exist, it searches for an estimation triple in which two cases are known and the third can be estimated by recursive application, similarly restricted, of the selected adaptation heuristic. For example, if the adaptation heuristic selected by the user is the difference heuristic and an estimation triple C_2, C_3, C_4 for a target case C_1 is found in which C_2 and C_3 are known cases and C_4 can be estimated by a sequence of estimation triples in which at most one case is unknown, then CREST's estimated value of the target case is:

$$est_d(C_1) = est(C_2) + est(C_3) - est_d(C_4)$$

Theorem 2 *In CREST, a problem space P of dimensions* $k_1 \times k_2 \times \ldots \times k_n$ *can be fully covered by* $\sum_{i=1}^{n} (k_i - 1) + 1$ *known cases. Moreover, the cases required to cover P can be selected in such a way that any target case can be estimated in at most n - 1 steps.*

Proof The values of the case attributes $x_1, x_2, \ldots, x_n$ can be assumed without loss of generality to be coded as positive integers. Let L be the set of all cases in P in which the value of at most one of the case attributes is greater than one. To prove that P is fully covered by L, which can be seen to contain the required number of cases, let

$$C_1 = (a_1, \ldots, a_{m-1}, a_m, 1, \ldots, 1)$$

be an unknown target case, where m is the index of the rightmost attribute whose value for the target case is greater than one. An estimation triple for the target case is provided by the cases:

$$C_2 = (a_1, \ldots, a_{m-1}, 1, 1, \ldots, 1)$$
$$C_3 = (1, \ldots, 1, a_m, 1, \ldots, 1)$$
$$C_4 = (1, \ldots, 1, 1, 1, \ldots, 1)$$

Of the three cases in the estimation triple, only C_2 can be unknown, and only if $m > 2$. If C_2 is unknown, an estimation triple for it can similarly be constructed in which at most one case is unknown and the index of the rightmost attribute whose value for the unknown case is greater than one is less than m - 1.

In this way, a sequence of estimation triples can be constructed in each of which at most one case is unknown and in which the index of the rightmost attribute whose value for the unknown case is greater than one is decreasing. Since every case in P for which the value of only a single attribute is greater than one is a known case, an estimation triple containing no unknown cases must be reached in at most m - 1 steps. Since $m < n$, it follows as required that the target case can be estimated in at most n - 1 steps.

With the case library in CREST strategically populated as described in Theorem 2, the retrieval of cases required to estimate a target case can be guided in such a way that the complexity of the estimation task increases only linearly with the number of case attributes. However, an unrestricted choice of seed cases is seldom possible in practice. Often the case library originates from archival sources and the acquisition of additional cases is costly. In these circumstances, it is important to minimize redundancy in the case library.

3.2 Mutually Independent Cases

Mutual independence in a case library, as defined below, requires the absence of redundancy in the library. For example, the 7 known cases which fully cover the 4 x 4 problem space illustrated in Figure 2 are mutually independent.

Definition 1 *Given a problem space P and a library L of known cases, the cases in L are mutually independent if there exists no* $C \in L$ *such that C can be estimated by CREST from L - {C).*

Lemma 1 *If P is a* k_1x k_2 *problem space which is fully covered in CREST by a library L of known cases, then every row and column of P must include at least one of the cases in L.*

Proof It is easily verified that if any row or column in P includes none of the cases in L, then no case in that row or column can be estimated.

Lemma 2 *If P is a* k_1 k_2 *problem space which is fully covered in CREST by a library L of known cases, and P' is the extended problem space defined by adding a new value to the domain of one of the case attributes in P, then for any known* $C \in P' - P$, P' *is fully covered by* $L \cup \{C\}$

Proof Let $D_1 = \{a_1, a_2, \ldots, a_{k_1}\}$ and $D_2 = \{b_1, b_2, \ldots, b_{k2}\}$ be the domains of the case attributes x_1 and x_2 in P respectively, and consider for example the $(k_1 + 1)$ x k_2 problem space defined by adding a new value a_Q to D_1 Let C = (A_0, b_j) be a known case in P' - P, and let $C_1 = (A_0, B_1)$ be any target case in $P' - (P \cup \{C\})$. Since P is fully covered by L, the row of P for which $x_2 = b_1$ includes at least one of the cases in L by Lemma 1. If $C_2 = (a_i, b_t)$ is such a case, an estimation triple for C_1 is provided by C_2, C_3, C_4 where $C_3 = C$, and $C_4 = (a_i, b_j)$. Of the three cases in the estimation triple, only C_4 can be unknown, in which case it can be estimated from L. It follows as required that C_1 can be estimated from $L \cup \{C\}$

Lemma 3 *In CREST, a problem space P of dimensions k* x k, *where* $k \geq 2$, *is fully covered by any set of 2k -1 known cases provided they are mutually independent.*

Proof The proposition is clearly true for k = 2, since a 2x2 problem space is fully covered by any set of 3 cases. Assume the proposition is true for k-m, where m > 2, and let L be any set of $2m$ + 1 mutually-independent cases in a problem space P of dimensions $(m + 1)$ x (w +1). Of the m + 1 rows in P, at least one must contain less than 2 of the cases in L. Similarly, at least one of the columns in P must contain less than 2 of the cases in L. On the other hand, there cannot be a row or column in P which contains none of the cases in L. For example, if such a row were to exist, its deletion together with the deletion of a column containing less than 2 of the cases in L would leave an m x m problem space containing at least $2m$ mutually-independent cases. By assumption, however, such a problem space is fully covered by any set of $2m$ - 1 mutually-independent cases, and cannot therefore contain $2m$ such cases.

The existence of a row and column in P each containing exactly one of the cases in L has now been established. The point of the problem space at which this row and column meet cannot be one of the cases in L; otherwise deletion of the row and column would again leave an m x m problem space containing $2m$ mutually-independent cases. Deletion of the row and column must therefore leave an m x m problem space containing 2m - 1 mutually-independent cases, which by assumption must fully cover the reduced problem space.

Extending the reduced problem space by adding back the points in the deleted row except for the point in the deleted column is equivalent to adding a single value to the domain of the second attribute in the reduced problem space. The m x $(m + 1)$ problem space so defined contains $2m$ of the cases in L and by Lemma 2 is fully covered by these cases. Adding back the points in the deleted column now restores the problem space to its original size. The reconstructed problem space contains all of the $2m + 1$ cases in L and again by Lemma 2 is fully covered by these cases. The proposition is therefore true for $k = m + 1$ and so by induction for all k £ 2.

Theorem 3 *In CREST, a problem space P of dimensions* K_1 x ky. *where* $k_1, k_2 \geq 2$, *is fully covered by any set of* $k_1 + k_2 - 1$ *known cases provided they are mutually independent.*

Proof For a given value of $k_2 \geq 2$, the proposition is true for $k_1 = k_2$ by Lemma 3. Assume it is true for $K_1 = m$, where $m \geq k_2$, and let L be any set of $m + k_2$ mutually-independent cases in a problem space P of dimensions $(m + 1) \times k_2$. Of the $m + 1$ columns in P, at least one must contain less than 2 of the cases in L; otherwise the number of mutually-independent cases would be more than $m + k_2$. On the other hand, there cannot be a column in P which contains none of the cases in L; otherwise its deletion would leave an ,$m \times k_2$ problem space containing $m + k_2$ mutually-independent cases which by assumption is impossible. One of the columns in P can therefore be deleted to leave an m x k_2 problem space containing $m + k_2 - 1$ mutually-independent cases, which by assumption must fully cover the reduced problem space. When the problem space is restored to its original size by adding back the points in the deleted column, it follows by Lemma 2 that the original problem space is fully covered by the $m + k_2$ cases in L. The proposition is therefore true for $k_1 = m + 1$ and so by induction for all $k_1 > k_2$- By a similar argument it is also true for all K_2 $>K_1$ value of k, > 2.

As shown in the following section, the usefulness of Theorem 3 is not limited to two-dimensional problem spaces.

4 Interaction Between Case Attributes

While the difference heuristic is perhaps the more widely applicable of the two adaptation heuristics, its assumption of an additive case function with no interaction between case attributes may be unrealistic in many estimation tasks. For example, consider the following hypothetical function, according to which the contributions of style and bedrooms to the value (£) of a property depends on its location.

$$f\{loc, style, beds) = 20{,}000 + 10{,}000 \times 2^{loc-l} + 2{,}500 \times loc \times 2^{style{\sim}l} + 2{,}000 \times loc \times beds$$

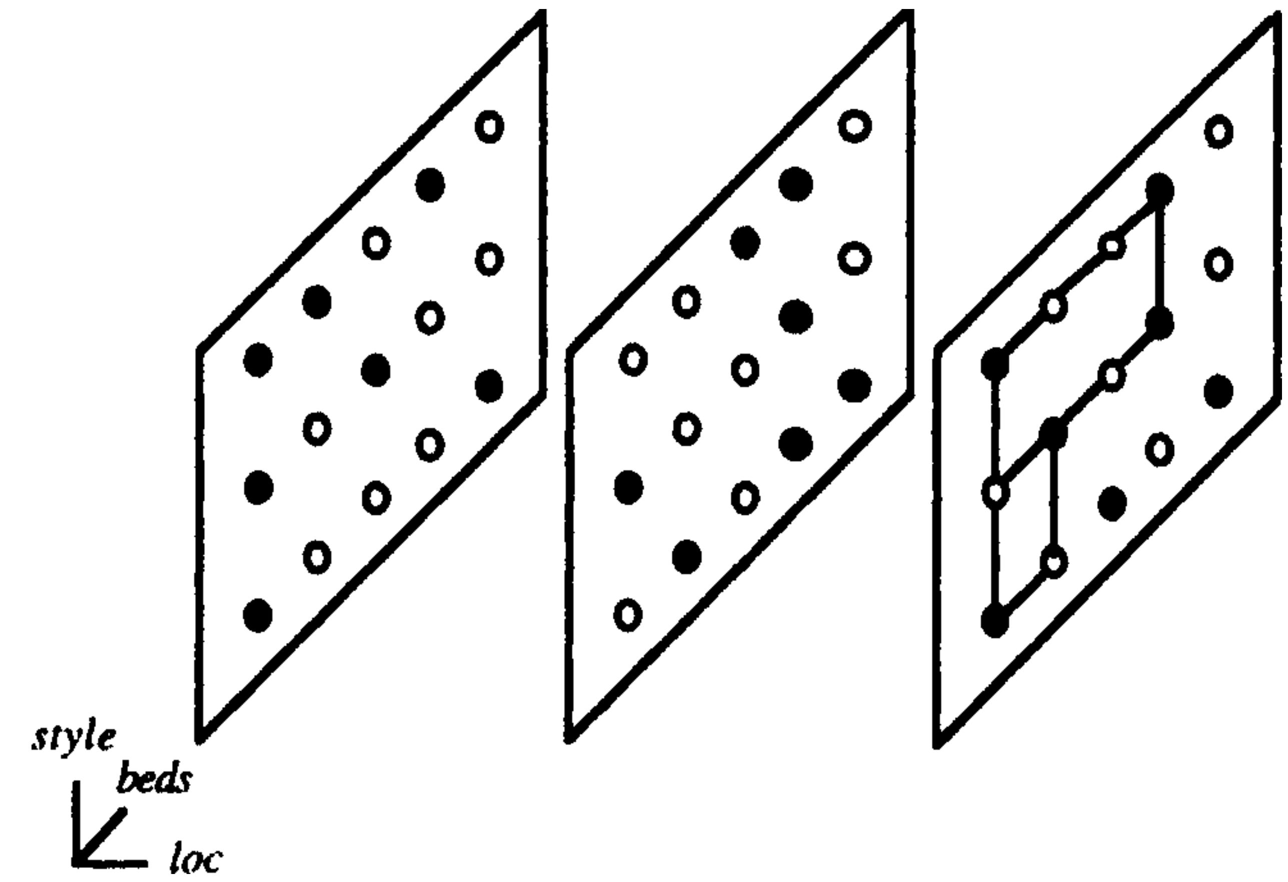

Figure 4. Estimation of the target case (3,1,2) in the property valuation domain.

The interaction caused by *loc* is likely to affect the accuracy of estimates based on the difference heuristic. However, with *loc* held constant, the function reduces to an additive function of *style* and *beds.* Provided *loc* is held constant, adaptation based on the difference heuristic must therefore give the correct value of a target case, subject to the absence of errors in the values of existing cases [McSherry, 1998]. The trade-off in this strategy is reduced coverage of the problem space. In the presence of interaction, full coverage of the 3 x 3 x 5 problem space illustrated in Figure 4 can no longer be provided by 9 cases as described in Theorem 2. However, as illustrated by the experiment described below, the problem space can be fully covered by 21 known cases even with the estimation process modified to eliminate the effects of interaction.

For each location, a set of 7 mutually-independent cases in the 1 x 3 x 5 problem space defined by that location was randomly generated. The case library is shown in Figure 4. An estimation error of e x £1,000, where e is a random integer between -3 and +3, was added to the value of each of the 21 generated cases according to the case function. CREST was then used to estimate each of the remaining 24 cases in the 3 x 3 x 5 problem space with *loc* held constant in the estimation of each case. For example, to estimate a target case with *loc* = 3, only estimation triples involving cases with *loc* = 3 were used. This is equivalent to separately applying CREST to the 1 x 3 x 5 problem space defined by each location. As such a problem space is equivalent to a 3 x 5 problem space, it follows from Theorem 3 that each of the 1 x 3 x 5 problem spaces is fully covered by the 7 mutually-independent cases it contains.

Recursive adaptation was necessary in only three cases: (3,2,3), (3,2,5) and (3,1,2). The sequence of estimation triples used by CREST to estimate the target case (3,1,2) is shown in Figure 4. Table 2 shows the value of each of the 8 unknown cases in location 3 according to the case function and its value as estimated by CREST. The average estimation error for the 8 cases, attributable only to the errors introduced in the values of existing cases, is £2,750.

Table 2. Estimated values (£) of 8 cases in the property valuation problem space.

Loc	*Style*	*Beds*	*Value*	*Estimate*
3	2	1	81,000	84,000
3	1	2	79,500	77,500
3	3	2	102,000	97,000
3	2	3	93,000	92,000
3	3	3	108,000	102,000
3	1	4	91,500	91,500
3	2	5	105,000	106,000
3	3	5	120,000	116,000

At the expense of additional computational effort, more accurate results may be achievable by basing the estimate of the target case on the average of the estimates provided by all available estimation triples (or sequences of estimation triples) rather than a single estimate as in this experiment.

5 Conclusions

A case-based approach to adaptation for estimation tasks has been presented in which a target case is estimated from three existing cases, one retrieved for its similarity to the target case, and the others to provide the knowledge required to adapt the similar case. Recursive application of the adaptation process significantly increases problem-space coverage, with full coverage in ideal circumstances by a relatively small number of selected cases.

Even with recursive adaptation, the requirement for the similar case to differ from the target case only in the value of a single attribute may restrict coverage in realistic problem spaces. The effects on retrieval efficiency and coverage of allowing the similar case to differ from the target case in more than one attribute are being investigated. Another interesting possibility is the development of interactive tools to guide the acquisition of cases to maximize coverage and minimize redundancy in the case library. Support for interactive authoring of cases is a topic of increasing importance in CBR research [Smyth and McKenna, 1998].

As confirmed by experimental results, certain departures from the assumptions on which the difference heuristic are based can be tolerated at the expense of some loss of coverage. However, in many domains, the case function may be neither an additive function nor reducible to an additive function, while the ratio heuristic's assumption of a multiplicative case function may be equally untenable. In some potential applications, the nature of the relationship between the value of a case and its attributes may be such that different adaptation heuristics are appropriate in different parts of the problem space. Further research will investigate the use of introspective reasoning to automate the selection of an appropriate adaptation heuristic, or combination of heuristics, thereby increasing the autonomy and scope of the case-based reasoner.

References

[Aha, 1998] David Aha. The omnipresence of case-based reasoning in science and application. *Expert Update,* 1:29-45, 1998.

[Bisio and Malabocchia, 1995) Rossella Bisio and Fabio Malabocchia. Cost estimation of software projects through case base reasoning. In *Proceedings of the First International Conference on Case-Based Reasoning,* pages 11-22, Lausanne, 1995. Springer-Verlag.

[Gonzalez and Laureano-Ortiz, 1992] Avelino Gonzalez and Raymond Laureano-Ortiz. A case-based reasoning approach to real estate property appraisal. *Expert Systems with Applications,* 4:229-246, 1992.

[Hanney and Keane, 1996] Kathleen Hanney and Mark Keane. Learning adaptation rules from a case-base. In *Proceedings of the Third European Workshop on Case-Based Reasoning,* pages 178-192, Lausanne, 1996. Springer-Verlag.

[Leake, 1996] David Leake. CBR in context : the present and future. In David Leake, editor, *Case-Based Reasoning: Experiences, Lessons & Future Directions,* pages 3-30. AAAI Press/MIT Press, Cambridge, Massachusetts, 1996.

[Leake *et al,* 1996] David Leake, Andrew Kinley, and David Wilson. Learning to improve case adaptation by introspective reasoning and CBR. In David Leake, editor, *Case-Based Reasoning: Experiences, Lessons & Future Directions,* pages 185-197. AAAI Press/MIT Press, Cambridge, Massachusetts, 1996.

[Leake *et ai,* 1997] David Leake, Andrew Kinley, and David Wilson. Learning to integrate multiple knowledge sources for case-based reasoning. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence,* pages 246-251, Nagoya, Japan, August 1997. International Joint Conferences on Artificial Intelligence.

[McSherry, 1998) David McSherry. An adaptation heuristic for case-based estimation. In *Proceedings of the Fourth European Workshop on Case-Based Reasoning,* pages 184-195, Dublin, September 1998. Springer-Verlag.

[Riesbeck and Shank, 1989] Christopher Riesbeck and Roger Shank. *Inside Case-Based Reasoning.* Lawrence Erlbaum Associates, New Jersey, 1989.

[Smyth and Keane, 1998] Barry Smyth and Mark Keane. Adaptation-guided retrieval: questioning the similarity assumption in reasoning. *Artificial Intelligence,* 102:249-293, 1998.

[Smyth and McKenna, 1998] Barry Smyth and Elizabeth McKenna. Modelling the competence of case-bases. In *Proceedings of the Fourth European Workshop on Case-Based Reasoning,* pages 208-220, Dublin, September 1998. Springer-Verlag.