The Knowledge Engineering Review, Vol. 20:3, 315–320. © 2006, Cambridge University Press
doi:10.1017/S0269888906000567 Printed in the United Kingdom

Case-based recommender systems

DEREK BRIDGE[1], MEHMET H. GÖKER[2], LORRAINE McGINTY[3] and BARRY SMYTH[3]

[1]*Department of Computer Science, University College Cork, Ireland;*
e-mail: d.bridge@cs.ucc.ie
[2]*PricewaterhouseCoopers LLP, Center for Advanced Research, Ten Almaden Blvd., Suite 1600, San Jose, CA 95113, USA;*
e-mail: mehmet.goker@us.pwc.com
[3]*School of Computer Science and Informatics, University College Dublin, Belfield, Dublin 4, Ireland;*
e-mail: lorraine.mcginty@ucd.ie, barry.smyth@ucd.ie

Abstract

We describe recommender systems and especially case-based recommender systems. We define a framework in which these systems can be understood. The framework contrasts collaborative with case-based, reactive with proactive, single-shot with conversational, and asking with proposing. Within this framework, we review a selection of papers from the case-based recommender systems literature, covering the development of these systems over the last ten years.

1 Introduction

In our everyday lives we receive recommendations from many sources: from salespeople and movie critics to restaurant guides and acquaintances. Recommendations help us to decide which goods, services, or information to purchase or consume. In situations where choice is increasing, good recommendations are of increasing importance.

On-line *recommender systems* are a new source of recommendations. Such systems are becoming more commonplace, especially on the Internet. They can support us as we go about our on-line business, whether it be browsing our favorite on-line book store or researching next year's vacation. Recommender systems combine ideas from information retrieval and filtering, user modeling, machine learning, and human–computer interaction. Case-based reasoning has played a key role in the development of an important class of recommender system known as *content-based* or *case-based* recommenders.

This paper provides an overview of case-based recommenders. It presents a framework within which these and other recommender systems can be understood. For example, it contrasts collaborative with case-based, reactive with proactive, single-shot with conversational, and asking with proposing. Within this framework, it cites and describes work that is seminal or representative of the state-of-the-art.

2 Case-based and collaborative recommenders

There are two main classes of recommender system: those that employ *collaborative* approaches and those that employ *case-based* approaches. Collaborative approaches exploit user histories, usually in the form of ratings-based profiles. Recommendations come from the profiles of the active user's *recommendation partners*. The partners are users whose ratings correlate closely with the active user's ratings. A collaborative recommender will recommend items that are not already in the active user's profile but which her partners have rated highly.

https://doi.org/10.1017/S0269888906000567 Published online by Cambridge University Press

Collaborative recommender systems require user ratings for the items that are to be recommended. They do not require item descriptions, and this is what sets them apart from their content- or case-based cousins. Item descriptions (whether they be text-based or attribute-value based) are vital in case-based recommenders, which generate a set of recommendations for a target user by retrieving items whose descriptions best match the user's *query*.

A case-based reasoning (CBR) system will have a case base of cases (i.e. previously solved problems and their solutions). New problems are solved by transferring and adapting solutions that were used for similar problems in the past. CBR is a multi-step reasoning strategy, the details of which are covered admirably elsewhere (Aamodt & Plaza, 1994). For our purposes, we highlight one of the essential early steps: *retrieval*. In the retrieval step, the system receives a problem specification, searches through the case base, scores each case for similarity to the new problem specification, and selects the highest-scoring case(s), which are the subject of subsequent steps, such as *adaptation*.

There are obvious parallels between the CBR retrieval step and the way a recommender system should treat a user query. From a CBR viewpoint, the query serves as a problem specification, the item descriptions are cases, and similarity-based retrieval techniques select the best-matching items.

The use of similarity-based retrieval is a beneficial feature of case-based recommenders, giving advantages over more traditional exact matching techniques such as conventional database retrieval and classical constraint satisfaction techniques (Vollrath *et al*., 1998; Wilke *et al.,* 1998). For instance, if a user's query is *over-specified*, no item exactly matches the query. However, similarity-based retrieval techniques can nevertheless retrieve a set of useful *similar* items. Conversely, if the user's query is *underspecified*, too many items exactly match the query. Then, similarity allows us to rank items and even prune those with the lowest similarity scores.

There are now many examples of fielded case-based recommenders. The electronics component manufacturer Analog Devices, for example, continues to use a case-based recommender to provide customers with more intuitive and flexible access to its catalog of operational amplifiers (OpAmps) (Vollrath *et al*., 1998; Wilke *et al*., 1998). Cases are composed of attribute–value pairs describing the technical features of each OpAmp. Because customers are seeking very specialized components for custom-built circuits, it is unlikely that any OpAmp will exactly match their requirements. However, the customers are usually satisfied by inexact matches that are sufficiently close to their needs.

The Analog Devices system is a good example of a *reactive* recommender system: the user provides an explicit query and the recommender system reacts with a recommendation response. It is possible for recommender systems to play a more *proactive* role, making recommendations without the need for an explicit query. For example, PTVPlus recommends television programs to users in the form of a personalized viewing guide, and bases its recommendations on the users' learned preferences rather than an explicit query (Smyth & Cotter, 1999). Incidentally, PTVPlus is also an example of a *hybrid* recommender system that combines case-based and collaborative recommendation techniques in order to maximize the accuracy of its recommendations. Relationships between CBR and collaborative recommendation continue to be investigated (Aguzzoli *et al*., 2002; O'Sullivan *et al*., 2002).

3 Recommendation dialogs

Many recommender systems adopt a *single-shot* recommendation strategy, returning a single set of suggestions to a user in a given session. In real-life, recommendation scenarios are rarely so shortlived, mostly because we are seldom able to fully specify our requirements up-front and we are rarely satisfied with the initial recommendations. *Conversational recommenders* adopt an iterative approach to recommendation. Users can elaborate their requirements, as part of an extended *recommendation dialog*.

https://doi.org/10.1017/S0269888906000567 Published online by Cambridge University Press

Different forms of conversational recommender systems can be distinguished by the way they elicit user requirements. For example, some conversational recommenders ask users a series of questions regarding their requirements (e.g., 'How much memory do you want?' in a recommender for personal computers). This is called *navigation-by-asking* (Shimazu, 2001, 2002). Systems may also or alternatively show the users particular products and elicit requirements in the form of *feedback* on the proposed products. This is called *navigation-by-proposing* (Shimazu, 2001, 2002). The issue of when to switch between the two types of navigation has been addressed in McGinty & Smyth (2003).

Systems engaging in navigation-by-asking face the problem of deciding the set of questions to ask in a session, and the ordering of those questions. Doyle & Cunningham (2000) were the first to report results on this topic within case-based recommender systems. They evaluate different question-selection criteria, including an entropy-based method inspired by work on inducing decision trees (Quinlan, 1986). Schmitt and his colleagues propose an alternative approach, called *simVar*, based on the variance in the similarity values (Kohlmaier *et al*., 2001; Schmitt *et al*., 2002; Schmitt, 2002). One advantage of *simVar* is that the knowledge it uses to choose the next question is the same knowledge that is used to make the next retrieval (i.e. the similarity values). The *simVar* approach can also accommodate the 'cost' to the user of answering each question so that the system can select questions to minimize dialog cost, rather than dialog length. See also Bergmann & Cunningham (2002) for other useful dynamic question-selection criteria. McSherry (2003c) has investigated the crucial issue of when the dialogue can be terminated without loss of solution quality.

Providing attribute values in response to explicit questions can place a significant burden on users. Indeed, sometimes users may not be able to answer a given question; for instance, their domain knowledge may be insufficient. This is what has motivated research into navigation-by-proposing. For example, Hammond *et al*. (1996) introduce the class of FIND-ME systems, the best known of which is the Entree restaurant recommender system. Entree presents restaurant recommendations to users; the user can then select one of the recommended restaurants and offer a *critique* or *tweak* of the selected restaurant. For example, the user might want a restaurant that is like the one she selects but cheaper, or like the one she selects but with French cuisine.

The FIND-ME approach, in which critiques furnish feedback within navigation-by-proposing, has enjoyed considerable success. It forms the basis, for example, of the Wasabi Personal Shopper, which is a domain-independent system whose applications include wine recommendation (Burke, 1999). The FIND-ME approach is further described in Burke (2002) and Burke *et al*. (1997). In a recent extension, compound critiques (ones that involve more than one attribute) are computed dynamically and offered to the user (Reilly *et al*., 2004). This has the potential to significantly reduce dialog length.

However, there is a simpler form of user feedback for use in navigation-by-proposing. The user might simply state a preference for one proposed item over the others that have been proposed. McGinty & Smyth (2002) call this *preference-based feedback*. It is akin to the *more like this* feature found in many Internet search engines. The simplicity of this form of feedback is particularly attractive in recommendation domains where users have very limited domain knowledge or where input modalities are limited, as they are with many Internet-enabled mobile communication devices. Importantly, preference-based feedback is *case-level* feedback; unlike critiquing, the feedback is not at the *feature-level*. This can limit its ability to guide retrieval. However, more sophisticated query revision (McGinty & Smyth, 2002) and/or case selection strategies (Smyth & McGinty, 2003) can improve the usefulness of preference-based feedback. For example, reductions in session length of up to 57% for preference-based feedback over the standard critiquing approach have been recorded (Smyth & McGinty, 2003).

Recent work in the area of *mixed-initiative* recommender systems promises to support more flexible models of interaction between users and recommender systems (Bridge, 2002; McSherry, 2002a). In mixed-initiative dialogue, there is an exchange of control between the two participants

https://doi.org/10.1017/S0269888906000567 Published online by Cambridge University Press

and a wider range of conversational moves may be open to each participant. In particular, the user should be able to volunteer information, as well as to provide information in response to system questions.

One crucial conversational move in CBR systems in general and recommender systems in particular is the provision of explanations to the user. The system may explain, for example, why it has asked a question; the explanation may be in terms of the effect the answer will have on its ability to discriminate between competing cases (McSherry, 2004). The system may also explain which of the user's requirements would result in a failure to retrieve any exact-matching cases (McSherry, 2003a). There remains an issue, however, of recognizing and explaining the system's confidence in its recommendations (see, for example, Reilly *et al.* (2005)).

4 The similarity assumption

Whether reactive or proactive, single-shot or conversational, navigation-by-asking or navigation-by-proposing, similarity plays a key role in case-based recommenders, just as it does in CBR in general. However, the pure similarity-based approach to retrieval is under attack. There is a growing acceptance that other factors have a role to play during item selection. For example, the CASPER system, which recommends job advertisements to users, combines query similarity and user relevance, selecting its recommendations because they are similar to the user's current query while at the same time being relevant to the user's known preferences (Bradley & Smyth, 2002). Similar motivations drive the Adaptive Place Advisor (Göker & Thompson, 2000), which attempts to learn the preferences of users during destination planning.

Equally, users can find that pure similarity-based approaches lead to recommendations that lack *diversity*. The likely success of a set of recommendations may be limited if the recommended items are too similar to each other. For example, recommenders designed to deliver a necessarily small set of recommendations to the screens of cellular phones must endeavor to recommend items that are similar to the user's query but different from each other. This increases the chance that at least one of the recommended items will satisfy the user.

The first in-depth investigation of the role of diversity in recommender systems appeared in Smyth & McClave (2001). The paper proposes a family of algorithms in which more items than are needed are retrieved based on similarity to the query. However, only a subset of these items are selected for display. The members of this subset are selected incrementally by a greedy algorithm that seeks to maximize their diversity. This relatively new idea has generated considerable interest and has prompted the development of other diversity-conscious retrieval algorithms (see, for example, McSherry (2002b)).

New approaches to retrieval are also under development. Compromise-driven retrieval, for example, selects cases using both similarity and compromise, the latter being based on comparing unsatisfied user requirements (McSherry, 2003c). Cases are *alike* if they involve the same compromises and, following McSherry (2002b), if cases are alike the system initially shows the user the one that is most similar to the query; this case acts as a representative of its group of alike cases.

Order-based retrieval is another new approach, with particular application to recommender systems (Bridge & Ferguson, 2002b). Rather than scoring the cases, order-based retrieval offers an expressive query language for defining and combining ordering relations; the result of query evaluation is to partially order the cases in the case base. The claims made for this new approach include: it is more expressive than similarity-based retrieval because it allows queries that naturally combine not just the user's preferred value (one for which similar values are sought) but also minimum values, maximum values, and ones the user wishes to avoid; it provides a natural semantics for critiques in navigation-by-proposing, and it returns inherently diverse result sets (Bridge & Ferguson, 2002a).

https://doi.org/10.1017/S0269888906000567 Published online by Cambridge University Press

5 Conclusions

Recommender systems in general and case-based recommenders in particular remain a vibrant research field. They have also been successfully deployed for approximately a decade. We have reviewed seminal papers and papers representative of the state-of-the-art and positioned them within a framework of concepts that contrasts collaborative with case-based, reactive with proactive, single-shot with conversational, and asking with proposing.

Any review of such an immense body of research is necessarily partial. Much more could be said about recommender systems in general, and about the role that ideas from CBR can play. For example, the focus of most research to date has been on 'off-the-shelf' products. Recommender systems for customizable or configurable products have not been widely investigated (see Stahl & Bergmann (2000)). The role of CBR in building user models and reusing previous shopping session experience has also received little attention (see Göker & Thompson (2000) and Ricci *et al.* (2002)).We look forward to the development of these and other areas over the next decade!

References

Aamodt, A and Plaza, E, 1994, Case-based reasoning: foundational issues, methodological variants, and system approaches. *Artificial Intelligence Communications* **7**(1), 39–59.

Aguzzoli, S, Avesani, P and Massa, P, 2002, Collaborative case-based recommender systems. In Craw, S and Preece, A. (eds.), *Proceedings of the 6th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 460–474.

Bergmann, R and Cunningham, P, 2002, Acquiring customer's requirements in electronic commerce. *Artificial Intelligence Review* **18**(3–4), 163–193.

Bradley, K and Smyth, B, 2002, Personalized information ordering: a case study in online recruitment. In Bramer, M (ed.) *Proceedings of Expert Systems 2003—Research and Development in Intelligent Systems XIX*, pp. 279–292.

Bridge, D and Ferguson, A, 2002a, Diverse product recommendations using an expressive language for case retrieval. In Craw, S and Preece, A (eds) *Proceedings of the 6th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 43–57.

Bridge, D and Ferguson, A, 2002b, An expressive query language for product recommender systems. *Artificial Intelligence Review* **18**(3–4), 269–307.

Bridge, D, 2002, Towards conversational recommender systems: a dialogue grammar approach. In Aha, DW (ed.) *Proceedings of the Workshop on Mixed-Initiative Case-Based Reasoning, Workshop Program at the 6th European Conference in Case-Based Reasoning*, pp. 9–22.

Burke, RD, Hammond, KJ and Young, BC, 1997, The FindMe approach to assisted browsing. *IEEE Expert* **12**(5), 32–40.

Burke, R, 1999, TheWasabi Personal Shopper: a case-based recommender system. In *Proceedings of the 11th National Conference on Innovative Applications of Artificial Intelligence*. Menlo Park, CA: AAAI Press, pp. 844–849.

Burke, R, 2002, Interactive critiquing for catalog navigation in e-commerce. *Artificial Intelligence Review* **18**(3–4), 245–267.

Doyle, M and Cunningham, P, 2000, A dynamic approach to reducing dialog in on-line decision guides. In Blanzieri, E and Portinale, L (eds) *Proceedings of the 5th European Workshop on Case-Based Reasoning*. Berlin: Springer, pp. 49–60.

Göker, MH and Thompson, CA, 2000, Personalized conversational case-based recommendation. In Blanzieri, E and Portinale, L (eds) *Proceedings of the 5th European Workshop on Case-Based Reasoning*. Berlin: Springer, pp. 99–111.

Hammond, KJ, Burke, R and Schmitt, K, 1996, A case-based approach to knowledge navigation. In Leake, DB (ed.) *Case-Based Reasoning: Experiences, Lessons, & Future Directions*. Menlo Park, CA: AAAI Press, pp. 125–136.

Kohlmaier, A, Schmitt, S and Bergmann, R, 2001, A similarity-based approach to attribute selection in user-adaptive sales dialogs. In Aha, DW and Watson, I (eds) *Proceedings of the 4th International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 306–320.

McGinty, L and Smyth, B, 2002, Comparison-based recommendation. In Craw, S and Preece, A (eds) *Proceedings of the 6th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 575–589.

https://doi.org/10.1017/S0269888906000567 Published online by Cambridge University Press

McGinty, L and Smyth, B, 2003, Tweaking critiquing. In Mobasher, B and Anand, S (eds) *Proceedings of the Workshop on Personalization and Web Techniques, Workshop Program at the 18th International Joint Conference on Artificial Intelligence*, pp. 20–27.

McSherry, D, 2002a, Diversity-conscious retrieval. In Craw, S and Preece, A (eds) *Proceedings of the 6th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 219–233.

McSherry, D, 2002b, Recommendation engineering. In van Harmelen, F (ed.) *Proceedings of the 15th European Conference on Artificial Intelligence*, pp. 86–90.

McSherry, D, 2003a, Explanation of retrieval mismatches in recommender system dialogues. In Aha, DW (ed.) *Proceedings of the Workshop on Mixed-Initiative Case-Based Reasoning, Workshop Program at the 5th International Conference on Case-Based Reasoning*. Trondheim, Norway: Norwegian University of Science and Technology, pp. 191–199.

McSherry, D, 2003b, Increasing dialogue efficiency in case-based reasoning without loss of solution quality. In Gottlob, G and Walsh, T (eds) *Proceedings of the 18th International Joint Conference on Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann, pp. 121–126.

McSherry, D, 2003c, Similarity and compromise. In Ashley, KD and Bridge, DG (eds) *Proceedings of the 5th International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 291–305.

McSherry, D, 2004, Explanation in recommender systems. In Cunningham, P and McSherry, D (eds) *Proceedings of the Workshop on Explanation in Case-Based Reasoning, Workshop Program at the 7th European Conference on Case-Based Reasoning*. Madrid, Spain: Universidad Complutense Madrid, pp. 125–134.

O'Sullivan, D, Wilson, D and Smyth, B, 2002, Improving case-based recommendation: a collaborative filtering approach. In Craw, S and Preece, A (eds) *Proceedings of the 6th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 278–291.

Quinlan, JR, 1986, Induction of decision trees. *Machine Learning* **1**, 81–106.

Reilly, J, McCarthy, K, McGinty, L and Smyth, B, 2004, Dynamic critiquing In Funk, P and González-Calero, PA (eds) *Proceedings of the 7th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 763–777.

Reilly, J, Smyth, B, McGinty, L and McCarthy, K, 2005, Critiquing with confidence. In Muñoz-Avila, H and Ricci, F (eds) *Proceedings of the 6th International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 436–450.

Ricci, F, Arslan, B, Mirzadeh, N and Venturini, A, 2002, ITR: a case-based travel advisory system. In Craw, S and Preece, A (eds) *Proceedings of the 6th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 613–641.

Schmitt, S, Dopichaj, P and Domínguez-Marín, P, 2002, Entropy-based vs. similarity-influenced: attribute selection methods for dialogs tested on different electronic commerce domains. In Craw, S and Preece, A (eds) *Proceedings of the 6th European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 380–394.

Schmitt, S, 2002, *simVar*; a similarity-influenced question selection criterion for e-sales dialogs. *Artificial Intelligence Review* **18**(3–4), 195–221.

Shimazu, H, 2001, ExpertClerk: navigating shoppers' buying process with the combination of asking and proposing. In Nebel, B (ed.) *Proceedings of the 17th International Joint Conference on Artificial Intelligence*. Berlin: Springer, pp. 1443–1448.

Shimazu, H, 2002, ExpertClerk: a conversational case-based reasoning tool for developing salesclerk agents in e-commerce webshops. *Artificial Intelligence Review* **18**(3–4), 223–244.

Smyth, B and Cotter, P, 1999, Surfing the digital wave: generating personalised TV listings using collaborative, case-based recommendation. In Althoff, KD., Bergmann, R and Branting, LK (eds) *Proceedings of the 3rd International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 561–571.

Smyth, B and McClave, P, 2001, Similarity vs. diversity. In Aha, DW and Watson, I (eds) *Proceedings of the 4th International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 347–361.

Smyth, B and McGinty, L, 2003, The power of suggestion. In Gottlob, G and Walsh, T (eds) *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, San Francisco, CA: Morgan Kaufmann, pp. 127–132.

Stahl, A and Bergmann, R, 2000, Applying recursive CBR for the customization of structured products in an electronic shop. In Blanzieri, E and Portinale, L (eds) *Proceedings of the 5th European Workshop on Case-Based Reasoning*. Berlin: Springer, pp. 297–308.

Vollrath, I, Wilke, W and Bergmann, R, 1998, Case-based reasoning support for online catalog sales. *IEEE Internet Computing* **2**(4), 45–54.

Wilke, W, Lenz, M and Wess, S, 1998, Intelligent sales support with CBR. In Lenz, M, Bartsch-Spörl, B, Burkhard, HD and Wess, S (eds) *Case-Based Reasoning Technology: From Foundations to Applications*. Berlin: Springer, pp. 91–113.

https://doi.org/10.1017/S0269888906000567 Published online by Cambridge University Press