

Cases, Context, and Comfort: Opportunities for Case-Based Reasoning in Smart Environments

David Leake, Ana Maguitman, and Thomas Reichherzer

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

Abstract. Artificial intelligence (AI) methods have the potential for broad impact in smart homes. Different AI methods offer different contributions for this domain, with different design goals, tasks, and circumstances dictating where each type of method best applies. In this chapter, we describe motivations and opportunities for applying *case-based reasoning* (CBR) to a human-centered approach to the capture, sharing, and revision of knowledge for smart homes. Starting from the CBR *cognitive model* of reasoning and learning, we illustrate how CBR could provide useful capabilities for problem detection and response, provide a basis for personalization and learning, and provide a paradigm for home-human communication to cooperatively guide performance improvement. The episodic memory developed by case-based reasoning systems within the home can both guide behaviors by the smart home and help its inhabitants to augment their own memories. After sketching how these capabilities could be served by case-based reasoning, the chapter discusses some design issues for applying CBR within smart homes and case-based reasoning research challenges for realizing the vision.

1 Introduction

Smart home environments have the potential for extensive impact on occupant comfort, convenience, and safety. Serious accidents can occur in the home because individuals are distracted or overwhelmed by ongoing events, and natural age-related declines in mobility and cognitive abilities may exacerbate problems, with occupants requiring monitoring and assistance (Craik & Salthouse 2000). For eldercare, smart home technology offers the prospect of a more independent life style in a private home compared to high-cost assisted living alternatives, or increased support and efficiency in group settings. Consequently, smart home technology may help alleviate the looming health-care crisis resulting from steady increases in health care cost and a rapidly aging population (Mann & Helal 2004; Mynatt, Essa, & Rogers 2000; Pollack 2005). Even for those who do not need special assistance, smart homes can facilitate daily tasks and provide an added measure of security by detecting and responding to emergencies.

Artificial intelligence (AI) methods for smart homes may provide both the flexibility to adapt to changing circumstances and the reasoning capabilities required to interpret events within the home and to make the right choices at the right times. In the introductory chapter to this volume, Augusto and Nugent (2005) argue for the suitability of

smart homes as an AI domain, both because of the potential payoff for AI methods and because their constrained task environment facilitates application of AI solutions. Because intelligent technology may be applied in many areas of a home, different types of support may suggest applying different types of AI methodologies. This chapter focuses on developing methods to enable smart homes to learn from observation of and direct interaction with their inhabitants, to be *adaptable* both to new tasks and to particular individuals' needs. The proposed approach is inspired by *human-centered computing* (HCC), in focusing on the home-inhabitant team rather than solely on the house as an autonomous agent.

The HCC approach dictates developing systems which can provide varying levels of autonomy and interaction, based on different individuals' needs and desires. In addition, it goes beyond simply requiring that the home attempt to "do the right thing" to requiring that the home can interact effectively with its inhabitants, sharing the burden with them in whatever ways will maximize overall performance. When the home can provide capabilities that the human lacks, or can increase efficiency, it should do so; when the human has unique capabilities or prefers to act independently, the home should defer and support the human's preference. This suggests that smart homes should support the *explanation* of their decisions, (1) to help an informed and capable occupant trust the home enough to accept its judgments and determine when to relinquish control, and (2) as background for an informed and capable inhabitant to directly adjust the home—reducing the learning burden on the system and making it feasible to refine system behavior more rapidly and reliably. It also suggests that the home should support a *simple and transparent learning process*, making it easy to revise system behaviors. These desiderata suggest exploring methods which learn in a way that people find understandable and offer simple processes for capturing new information.

Case-based reasoning (CBR) is a methodology for reasoning and learning based on the capture of specific experiences and their retrieval and adaptation to fit new needs. One of the inspirations for case-based reasoning comes from cognitive modeling of human reasoning and learning (e.g., (Kolodner 1994; Leake 1998)). Experience with fielded CBR systems suggests that humans are comfortable receiving information in the form of cases, and experiments show the usefulness of cases for explaining system reasoning (Cunningham, Doyle, & Loughrey 2003). Likewise, CBR provides a simple, easily comprehensible learning process. This chapter sketches how CBR could be applied to smart homes tasks such as responding to and learning from expectation failures, generalizing patterns of events to refine system expectations and to help caregivers and relatives to identify problems, as well as potentially interacting with inhabitants/caregivers through examples, to explain and adjust the homes' behaviors.

The chapter begins with a synopsis of case-based reasoning, including both the role of cases in human understanding, reasoning, and learning and the pragmatic motivations for applying case-based reasoning as an AI technology. It then illustrates how CBR can be applied to a number of distinct tasks desirable for smart homes, through multiple case-based processes. It then discusses special challenges for CBR to realizing the vision.

2 A Synopsis of CBR: Cases for Understanding, Learning, and Problem- Solving

One of the early foundations for case-based reasoning was Schank's study of human *re-minding* during understanding, which gave rise to Schank's Dynamic Memory Theory ((Schank 1982); summarized in (Schank & Leake 2002)). Dynamic Memory Theory addresses the relationship of human understanding, learning, and memory, with a central focus on how knowledge is structured, organized and revised based on experience. This approach is of interest to smart home development for three reasons. First, a fundamental need for smart homes is recognizing events and predicting following steps. Second, no smart home can be perfect: Smart homes will need to adapt to experience. Third, the use of human-like learning methods may help inhabitants to understand and interact with the smart home, increasing its acceptance.

In Dynamic Memory Theory, standard event sequences (such as the events involved in hosting a party—sending invitations, preparing food, welcoming guests, offering them food, etc.,) are characterized by sequences of basic components, called *scenes*, which normally describe events which take place in a single location, with a single purpose, and in a single time interval. These sequences are captured in knowledge structures called Memory Organization Packages, or “MOPs”. They are hierarchical, with shared structure enabling lower-level components to be shared by a number of MOPs.

A central focus of Dynamic Memory is the process by which an understander refines its knowledge structures, to improve future predictive power—a process of great importance to smart homes as well—and the role of reminders in that process. Each MOP provides certain expectations, which provide top-down guidance for understanding. If these expectations are not borne out, an expectation failure occurs. For example, with RFID tags, a smart house might detect an impaired occupant picking up a toothbrush, leading to the expectation for toothpaste to be picked up next. If, instead, the patient picked up a tube of shoe polish, an expectation failure would occur—and would require a warning.

A tenet of Dynamic Memory Theory is that failure episodes are stored under the processing structures in effect at the time of the failure, making them accessible as reminders when another similar failure occurs. Dynamic Memory Theory and later work in its tradition (e.g., (Domeshek 1992; Leake 1992b; Schank *et al.* 1990)) developed indexing vocabularies aimed at capturing the features needed for computer models to generate the types of reminders observed in humans. If we accept the general ability of people to generate useful reminders, we can expect that analogous reminders may be pragmatically useful for a smart home to generate, as a surrogate memory for the home's inhabitants. These reminders provide a starting point for case-based reasoning by either the human or the home.

Beyond the immediate usefulness of reminders to aid current tasks (e.g., by warning of a possible mistake or suggesting a useful past solution), such reminders can provide data for future generalization to generate new schemas. For example, the first time a patient confuses shoe polish with toothpaste, a warning may be sufficient, but if this happens often, the smart house may determine that the patient will reach for anything available, and might need to start guiding a visiting nurse to remove any potentially dangerous objects when preparing the house for the night. Likewise, case data

may be useful to refine responses such as warnings (e.g., whether a nurse should be called or an audible warning generated with a speech synthesizer is sufficient, and what volume warning is needed).

For a system trying to explain subject behavior, having access to cases for prior experiences provides a number of potential benefits. First, the cases focus attention on scenarios which actually *have happened* for the particular subject. Second, cases augment the system's built-in knowledge. Third, reminders can carry solutions or predictions which applied in the past, and which may be useful again. If the last time a subject overslept, he skipped breakfast, the home might not turn on the coffee maker, but instead could start warming up the car. If the heat is on, but the home does not become warm, a reminding that the same problem occurred when a door was left ajar, and was fixed by closing it, might provide a solution. Third, the case provided by the reminding may contain lessons from both successes and failures.

This process provides a first-pass suggestion of a framework for how smart homes may combine general schemas with cases to understand actions, predict the actions to follow—in order to support them—and note deviations in order to explain, react, and learn. For example, the schema for a daily routine might include getting up, dressing, eating breakfast, and so on. If a home's inhabitant fails to get up when the alarm goes off one morning, the anomaly might prompt reminders of other similar instances, for example, a case of prior illness.

In general, a case-based reasoning system exploits reminders to help interpret new situations or to solve new problems. When faced with a new problem, a CBR system retrieves prior cases for similar problems and adapts their solutions to fit new circumstances. When faced with a new situation to interpret, a CBR system retrieves prior cases for similar situations and compares and contrasts them to the current situation to form a new interpretation. The steps of the CBR process include situation assessment, to describe a situation in a vocabulary commensurate with the indices used in memory, then retrieval and similarity assessment of candidate cases, then adaptation to fit the case to the new situation and evaluation of the results (first by internal reasoning and then by application of the case to the new problem, providing real-world feedback). If problems are detected, the adaptation/evaluation cycle may be continued, until the final result is stored in the case base. Figure 1 sketches this process. Additional variations and extensions are possible—for example, drawing on multiple cases to provide parts of a solution. For a more in-depth overview, see, for example, (Aamodt & Plaza 1994; Kolodner & Leake 1996a; Leake 1996).

Flexible processes for situation assessment, similarity assessment, and case adaptation enable cases to be applied in new contexts and despite imperfect matches to past events. In smart homes, it might seldom be true that two tasks were performed exactly the same way and in the same context. To realize the full potential of CBR, a CBR system must generalize its indices—even as it keeps the cases themselves specific—to enable potentially relevant cases to be retrieved, must be able to assess similarity of differing situations, and must be able to adapt solutions to new needs. In practice, CBR applications will address these to different levels. Depending on the granularity of sensor information, for example, different amounts of indexing information will be needed.




Fig. 1. The CBR cycle

The CBR literature presents a number of practical motivations for using CBR (see (Leake 1996) for one more detailed treatment):

1. **Knowledge acquisition by example:** Providing knowledge in the form of cases can seem more natural to subjects than attempting to distill knowledge into general rules. Cases only need to be captured for situations that actually arise in practice, and the system can be fielded with a small set of “seed cases,” acquiring more cases as needed. However, more traditional knowledge acquisition problems may arise in case engineering or capturing supporting knowledge such as similarity or adaptation knowledge.
2. **Learning from both successes and failures:** Case acquisition not only accumulates successful cases as future guidance, but also failure cases which can serve as warnings of problems.
3. **Inertia-free learning:** Cases provided to the system may immediately be applied, without any need for re-training of the system or learning period during which the system responds incorrectly until more evidence is acquired.
4. **Simple and lazy learning:** Case acquisition to cover new problems can improve system coverage with a very simple learning process, and small initial cost. Effort to generalize cases, etc., need only be applied if similar circumstances arise in the future, only need to be done to the extent required by the new situation, and can immediately be tested with the feedback of that situation.
5. **Ability to draw on multiple types of knowledge:** What domain knowledge is available can be integrated into many parts of the CBR process, to improve indexing or ease adaptation.
6. **Problem-solving efficiency:** Case-based reasoning can be combined with other methods, such as rule-based or model-based reasoning, to provide speedup learning by reusing results of past reasoning rather than requiring them to be derived from scratch.
7. **Ability to function in poorly-understood domains:** When a domain is poorly understood, cases that have worked in similar situations may implicitly capture important factors that would not be captured in rules reflecting current understanding.

8. **Ability to provide operational and contextualized guidance:** Because cases are kept at an operational level rather than generalized, they may be easier to apply (e.g., a cooking plan could be stored including the specific places to look for an ingredient in a household, etc.)
9. **User acceptance:** Cases capturing similar prior situations can provide effective explanations for why a system performed as it did (Cunningham, Doyle, & Loughrey 2003).

3 Opportunities for CBR in Smart Homes

Case-based reasoning has been widely applied to advisory tasks, including the selection of assistive technology for smart houses (Wiratunga *et al.* 2004). However, the direct application of CBR to the internal function of smart homes remains unexplored. As many stands of CBR research have potential applications to smart homes, this section will outline the reasons for applying CBR to this domain and avenues for its application.

The characteristics of CBR suggest four motivations for applying CBR as an AI methodology within smart homes. First, case learning provides a simple method for refining and personalizing a smart home over time by storing cases for new event sequences, to be generalized into new schemas, and for learning appropriate responses. Second, its ability to reuse prior reasoning—regardless of the original reasoning process—can increase efficiency for expensive reasoning tasks with substantial potential payoffs for smart homes, such as abductive explanation of observed anomalies (Leake 1993) and use of these explanations to support multiple reasoning tasks explanations may be tailored to addressing individual needs (Leake 1992a). Third, the use of CBR can facilitate human-system interactions by enabling humans to provide unanalyzed examples (or provide feedback on examples) as a means of training the smart house. Fourth, references to cases provide an intuitive way for smart homes to explain their behaviors to people (Kolodner 1991; Cunningham, Doyle, & Loughrey 2003; Leake & McSherry 2005).

In the remainder of this section, we illustrate the potential usefulness of CBR to support selected smart home tasks. In this discussion, we will assume that the smart house is equipped with the array of sensors described in the introduction to this volume, providing capabilities to monitor properties such as the locations of inhabitants, on/off states of appliances, and vital signs of inhabitants. Later sections highlight some issues that must be addressed to develop CBR systems for smart homes.

3.1 Monitoring, problem detection and response

An essential capability for smart homes is the ability to monitor household events and inhabitant state. This monitoring may serve two purposes. One is to recognize instances of routine patterns, to enable the home to anticipate and satisfy standard needs (e.g., to turn on the coffee maker after the occupant gets up, or to guide an impaired inhabitant through normal steps of daily life). Another is to recognize unusual and potentially hazardous situations and respond. For both types of tasks, case-based methods can be useful.

Cases for monitoring and explanation Comparison to past cases can help to detect situations which have proven interesting in the past, with a case base directly capturing events of interest and responses to them. For example, a case could be indexed by a patient in a nursing home falling to the floor (as determined by the change in location of a sensor worn by the patient), and could directly suggest a response (e.g., making an urgent call to a nursing station, broadcasting to the inhabitant that help has been requested, and turning on house lights). By storing cases captured in similar circumstances (e.g., at the same nursing home), local details (e.g., telephone numbers) can be captured. Similar prior cases may also be used as a reference for comparing events, in order to detect anomalies. For example, if a patient's heart rate after climbing stairs is much higher than it was the previous time, that increase should be noticed.

Likewise, cases can play a role in monitoring more routine activities. Although these may in principle be as diverse as the number of ways humans interact with the domestic environment, particular individuals' tasks and strategies tend to recur, at least in similar forms. Cases provide a first vehicle for recognizing events, as well as data for generalizing the cases into schemas which can be used to track actions and form expectations for inhabitant activities (e.g., to turn on the bedroom light if the inhabitant is preparing for bed), as well as to detect, explain, and respond to expectation failures. Thus it provides an alternative to alternatives such as reinforcement learning (Mozer 2005), with the ability to explain system actions in terms of desired actions in similar situations.

CBR can also be useful as a way to explain anomalies detected by other AI methods. For example, (Leake 1992b) shows how a pattern-based method for anomaly detection can generate indices when generating explanations for the anomalies by case-based reasoning. The resulting explanations can then be made available to CBR or other AI methods to determine appropriate responses, once the cause of the unexpected event has been determined by CBR.

CBR for anytime response After detection of a situation requiring a response, CBR may again be useful, both for proposing responses and for tailoring problem response speed to problem urgency. When emergency conditions require very rapid response (e.g., when a heat sensor detects a fire), it may be crucial to begin actions without extensive reasoning, suggesting the use of anytime algorithms (Dean & Boddy 1988), which can be terminated at any time while still providing a meaningful answer, and which can return answers whose quality improves with increased processing time. Riesbeck (Riesbeck 1996) points out that CBR has "anytime" properties, as a system may immediately begin acting based on a similar case, while continuing to search for possibly more on-point cases and possible adaptations as time permits.

For example, given a fire, the system might immediately find a case for another fire involving turning on sprinklers, turn off room power and call for emergency services, all of which can be applied directly from a prior case. Other aspects of the case, such as calls to the previous patient's relatives, would require adaptation, and, perhaps additional effort (e.g., searching for the relative's number, or taking additional steps if the relative is hard to reach). Different types and levels of alarms could suggest different points at which to harvest the intermediate results of CBR. Applying this approach re-

quires methods for estimating problem severity, applicability of the current case and usefulness of its components, and expected adaptation cost for the cases in order to guide the response process.

3.2 Personalizing case application

A smart home guided by standard schemas can provide stereotyped responses, but not the differing needs of different users. Building up personalized sets of cases for different users addresses that problem. For example, suppose that a window was left open on a cool summer night, with the air conditioner on but not running. If, next morning, the house becomes warm enough for the air conditioner to start, one person might close the window (favoring temperature) and another might turn off the air conditioner (favoring fresh air). Case-based reasoning has already been applied to recommendation systems, with good results (e.g., (Smyth & Cotter 2005)).

For any new user, there will be a transition period until sufficient personalized cases are built up. During that period, the system must still rely on knowledge developed for other users. If that knowledge is drawn from cases, there are opportunities for adapting the CBR process to provide more personalized results, even when the cases are drawn from other users.

As proposed in (Leake 2002), developing personalized CBR systems requires replacing the task-centric view—that there is a single solution for each problem—with a user-centric view that supports multiple solutions, based on the user as well as the problem situation. This can be addressed by modeling the user as well as the task, and using both types of information to guide CBR at many levels.

Again, different CBR processes can play different roles. A secondary CBR process can be used to directly support personalization process (Blanzieri 2002), to classify a new individual based on individuals already classified; this can be an index provided to the main CBR process to reflect different users during retrieval, adaptation, evaluation and storage. Augmenting cases descriptions with user properties can improve the system ability to identify useful solutions.

3.3 Personalized smart home autonomy

However, not all advice is useful, and not all aid is desired; different individuals have varying needs and tolerances. Likewise, it may be desirable for a system to balance user preferences against general policies (e.g., comfort vs. energy costs (Mozer 2005)). Consequently, a smart home should provide *adjustable autonomy* (Musliner & Pell 1999), furnishing different levels of support depending on its own capabilities, user preferences, user needs, and general policies. Thus another useful type of personalization would consider potential actions and decide which should actually be performed for a particular inhabitant, or how to adapt them to make them more acceptable.

Decisions about whether to act depend on many factors, such as the user's physical and mental condition, the user's tolerance of system intervention, system's confidence in its assessment of a situation, the cost of action, the potential risk of failure to act, and so on. Because of the complexity of weighing these conflicting factors, it may be desirable either to rely on experience, using CBR to asking what this user or similar

users have favored in similar situations, or to exploit CBR to retrieve and adapt the results of prior calculations performed by other AI methods for similar prior problems.

3.4 Personalization of information presentation

In addition to deciding when to present information, a smart home must decide how to present it from a plethora of interface options. For example, one direct presentation approach is implemented in the REA system (Cassell 2001), a fully, virtually embodied interface agent that advises customers on real-estate purchases, uses a variety of verbal and non-verbal input modalities when interacting with its customers; the ambient display technologies that exist in the periphery of our perception and only appear if we choose to interact, thereby contributing to a calm environment (Weiser 1991). Examples of such technologies include Hello.Wall, which facilitates communication by presenting unobtrusive digital light patterns to individuals as they walk by, to alert them of new messages which can be accessed through a PDA if desired (Streitz *et al.* 2005). Ambient display technologies demonstrate the potential for interactions going far beyond traditional keyboard-mouse-based human-computer interaction, which require users to be located in a single spot and to explicitly initiate each interaction. In contrast, spoken-dialogue systems with a variety of input and output modalities allow for spontaneous interaction between humans and the system.

3.5 Explanation of system decisions

It has long been recognized that confidence in and acceptance of intelligent system decisions may benefit from increased user understanding of system reasoning (e.g., (Buchanan & Shortliffe 1984)). For non-impaired inhabitants, the ability to explain may be helpful as well, both to instill confidence and to enable systems to guide repair of problems (e.g., if a system sets an early alarm for Friday, believing that it is a work day, when the inhabitant works only Monday to Thursday).

Case-based reasoning provides a natural vehicle for explanations. Unlike, for example, neural networks, case-based reasoners can account for their decisions by presenting specific prior cases, which inhabitants can examine to assess their applicability. Early CBR work sometimes considered the case alone to be a sufficient explanation; recent work is developing a richer view (for a sampling of recent work, see (Leake & McSherry 2005)). In addition to providing arguments for the relevance of a case and a conclusion (e.g., in the form of comparisons and contrasts (Ashley & Risland 1987)), current research is examining how to explain other facets affecting system conclusions, such as how features of a case contribute to similarity calculations (Massie, Craw, & Wiratunga 2004). Each of these explanations in turn provides a point to which a user might provide feedback accounting for an erroneous conclusion and enabling system refinement.

3.6 Providing task-relevant reminders from captured cases

Elderly or infirm patients may have difficulty following normal task sequences. Consequently, the ability to capture and provide task guidance could be valuable. More

generally, memory augmentation, in the form of environment-aware systems which can use context to disambiguate requests for information, could provide valuable services in smart homes—from guiding tasks to providing recipes to more general questions.

Much work has already been done in the area of intelligent reminder systems, and many tools have been developed to augment human memory through context-based proactive assistance in homogeneous and regulated environments. Much of this work centers on the desktop paradigm, where tasks are often limited to editing, searching, entertainment and communication (Rhodes & Starner 1996; Budzik & Hammond 1999; Rhodes 2000; Budzik, Hammond, & Birnbaum 2001). Within this paradigm, case-based systems have already been applied to tutoring systems which monitor an individual's progress in a simulated environment, tracking behavior to detect potential problems and present video clips of cases with “reminders” warning about potential pitfalls (Burke & Kass 1996). While smart home environments would present many challenges beyond the simulated one, this work provides a general sketch for this approach. An additional issue will be the need to aid more general tasks and provide smooth integration, by being delivered using the right modality, at the right time, and at the right location. In general, making this determination may be a daunting task, making it appealing to gather examples, which may be combined with general rules for deciding where and when to provide guidance.

Systems that offer guidance as the individual completes daily activities have been developed primarily to assist the cognitively impaired. Some of these systems, such as Autominder (Pollack *et al.* 2003), are general purpose tools. They operate by monitoring the individual's task and perform schedule management functions that otherwise would require human assistance. Others have as a goal to assist more specific tasks. An example of such a system is COACH (LoPresti, Mihailidis, & Kirsch 2004), a tool developed to assist in the hand washing process by providing needed reminders. A primary motivation for exploring CBR for these tasks would be their capability for simple and lazy learning, and their ability to function in new domains by providing them with additional examples, making them especially suitable to adapt to diverse situations, without the need for a predefined plan of daily activities.

3.7 Enabling user instruction of smart homes

No smart home will be perfect; the home must learn. One form of learning is simply for the system to observe the user's actions. If the system mis-sets the morning temperature and the user adjusts it, the system can adapt the outcome of its stored case, replacing its initial choice with the user's corrected one. Observing and learning from user settings alone is a promising vision, but also has limitations in that few cues may be directly available to the system (Mozer 2005). In addition, if the smart home learns by methods requiring much training data, the system may continue in its old behavior for some time after the user's behavior changes, prompting user frustration.

The more sophisticated inferences a home requires, the more it may be difficult for a system to autonomously select appropriate behaviors and learn from the limited information that may be available by observation. Consequently, it may be desirable to enable users to choose to share some of the burden in exchange for faster learning or better performance. CBR provides a potential avenue for this. If cases are captured

in a comprehensible form and made available to users for interactive case adaptation, users can refine cases to directly reflect the conclusions they find appropriate. Numerous avenues for interaction with CBR systems have been explored (Aha & Munoz 2001), and this continues to be an active research area.

An interesting extension of these approaches would be to provide new paradigms for interacting with cases, such as enabling user actions in the home to automatically connect to refinements of case information. For example, if an individual following a recipe replaces ingredients, such changes might be easy to detect through the use of RFID tags, enabling the system to generate a new recipe without bothering the user, and to save that for later use. Because cases encode episodes, which may be presentable in an intuitive way, case presentation and editing interfaces might provide general-purpose methods to enable sophisticated users to adjust system behavior themselves through case editing. For example, some CBR research has pursued using concept maps (Novak & Gowin 1984) as the basis for browsable and editable cases (Cañas, Leake, & Wilson 1999; Leake & Wilson 2001).

4 Putting it into practice: Case generation and access for smart homes

Building any CBR application requires developing methods for the core CBR tasks of Section 2, as described in sources such as (Kolodner 1993; Watson 1997). This section focuses on two key issues, generating the system's initial case base and accessing the cases based on information available from the home.

4.1 Generating Case Bases

CBR systems are only as good as their cases. Consequently, case capture and engineering are central issues in fielding CBR applications, and a first step is determining the case representation. As discussed in (Kolodner & Leake 1996a), a case is a *contextualized piece of knowledge representing an experience that teaches a lesson fundamental to achieving the goals of the reasoner*. Cases must contain a *problem/situation description*, describing the state of the world when the episode recorded in the case occurred, a *solution* or *interpretation*; and the *outcome* of applying the solution. The solutions/interpretations that cases provide may take different forms, ranging from simple textual advice to an inhabitant (e.g., "Are you forgetting to take out the garbage?") to rich structured representations (e.g., the plan for a recipe). The preferred representation will depend on the task, in light of a tradeoff between easy generation (e.g., for textual cases) and the ability for cases to support further reasoning (e.g., for a recipe, for which the home might suggest how to revise the recipe when substituting for an unavailable ingredient).

In order for a CBR-based system to be useful as soon as it is fielded, developers must "jump start" the system with initial knowledge capturing the events, actions, and behaviors of individuals in the home relevant to system tasks. For a CBR system which monitors and predicts behaviors, one option would be to base initial processing on generalized schemas; another would be for developers to manually identify prototypical or

critical events and activities for the system to handle, such as common, regular activities and behaviors (e.g. when people get up in the morning, when they leave the house, etc.), or to analyze automatically generated cases to determine which to store. Additional cases may be added automatically by varying the values of case features within a predefined range.

Cases may be useful if their history suggests that they would be frequently used in the home, if they would help solve critical problems (e.g., for emergencies), if they would fill gaps in the system problem coverage, or if they would be easily adaptable, increasing system flexibility. The result is an initial case base with a representative set of cases capable of adapting to the specific requirements. Given the range of CBR processes that may apply within a smart home, multiple case bases may be needed, e.g., for monitoring inhabitant vital signs, monitoring home conditions, etc.

4.2 Information sources and situation assessment

Accessing the right cases requires methods for retrieving relevant cases, based on information from the smart home. Work on pervasive computing for smart homes has developed sensor technologies and smart appliances to monitor individuals and their physical environment, and has built middleware to provide services and facilitate communication among the different components. A number of research projects have focused on developing an infrastructure to deal with context in smart home environments (e.g., the Aware Home Research Initiative (Dey, Abowd, & Salber 1999; Meyer & Rakotonirainy 2003)), providing a foundation on which case-based smart home applications could build.

Generally, a pervasive computing architecture for smart homes can be broken down into a physical layer, a middleware service layer, and an application layer, which together provide a platform delivering high-level information abstracted from the sensors, detectors, actuators, smart appliances and wired or wireless communication devices to integrate new technologies and services in smart homes (examples include the Gator Tech Smart House (Helal *et al.* 2005) and MavHome (Cook *et al.* 2000)). The CBR layer, like other AI methods for smart homes, can be integrated into the pervasive computing architecture, to process a comparatively high-level information stream.

Existing projects address many component issues for deriving a number of types of information which might be useful for a CBR system to generate indices for retrieval. For example, the Pfinder project studies tracking individuals in a smart room (Pentland 1996), and the Sociometer project studies how to track people's interaction with others (Choudhury & Pentland 2003), potentially providing a basis for generating more abstract indices.

To facilitate humans interactively providing the system with supplementary information (e.g., by an inhabitant simply telling the system that he or she wants to bake a cake), information provided to situation assessment could also be based on human-like input and output modalities such as speech, body postures, eye movements and other non-verbal gestures. All such interactions must be able to feed into the CBR system, and their interpretation may itself require considerable AI processing to provide the information used in indexing. For example, the Trips system (Allen *et al.* 2001), a spoken-dialogue computer system for planning, parses human language into practical

dialogues that capture what the user meant by an utterance, applies domain-independent problem-solving models, and domain-specific task models and is capable of recognizing user intentions, which could then become part of the context used by CBR retrieval.

4.3 Capturing context and context-based indexing

In order to access the right cases, cases must be organized; the *indexing problem* is the problem of assuring that a case is accessed whenever appropriate (Kolodner & Leake 1996b). The CBR system's situation assessment process must be able to generate suitable indices. A number of general indexing vocabularies have been developed for tasks such as indexing explanations of anomalies (Leake 1992b), and could be exploited in the smart homes setting. However, for smart homes, contextual factors will play a key role in indexing as well, and the range of potential information is extensive. For example, Dey proposes that "Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves" (Dey 2001). Capturing context for reasoning in smart environments is challenging because of the need to determine relevant classes of features and how they can be recognized based on numerous heterogeneous sources. In the proposed scenario for this volume, which assumes simple on/off device information, the informativeness of each individual sensor's information may be fairly small—determining useful information may depend on the context of other sensor readings and sensor changes over time.

It is difficult to determine in advance which aspects of a situation are important to include in context-based indices, because feature relevance will change from situation to situation. However, general vocabularies can be developed to capture fundamental properties. For example, (Gross & Specht 2001) proposes including dimensions for *identity* (information about user's short and long term needs, interests, preferences, knowledge, etc.) as also discussed in Section 3.2, *location, time* (while tracking time is simple, using temporal information for planning and decision-making poses several challenges; see the discussion of temporal reasoning in smart homes in this volume (?)) (CITATION TO BE INSERTED), and *environment and activity*, describing the artifacts and the physical location of the current situation. Awareness of the environment beyond the house could be relevant to issue recommendations or alerts. For example, during a rainy day a smart home could remind the inhabitant to carry an umbrella. Efforts to develop suitable vocabularies can draw on considerable active research on context modeling (for a sampling, see (Schulz, Roth-Berghofer, & Leake 2005)).

A smart home system makes predictions from its observations based on current system knowledge and the current state of the environment. A discrepancy between a prediction and the actual outcome may suggest the need for updates to system knowledge, to learn from a new situation. A difficult problem for both prediction and learning is that effects may depend on hidden context, causing smart home systems to make incorrect recommendations or to generate false alerts due to incomplete information. Thus an important goal for smart home systems is to detect context changes without being explicitly informed about them and to quickly recover from context changes, adjusting their hypotheses to fit the new context. Approaches to this problem as it relates

to concept learning may also be valuable for smart home systems (Widmer & Kubat 1996).

The situation assessment and indexing problems are open-ended. Developing vocabularies for describing fairly specific features, such as patient vital signs, will be straightforward. Developing principled methods to derive highly abstract features is still an open challenge for long-term research (e.g., to recognize that an occupant is having a bad day).

5 Some Research and Practical Issues for Realizing the Vision

While CBR provides an appealing methodology for smart homes, many specific issues must be addressed to fully realize its promise. This section highlights a few which may be especially salient for the smart homes domain.

Case Engineering: In broad outline, CBR simplifies knowledge acquisition by alleviating the need to develop rules (for a rule-based system) or to acquire extensive training data (e.g., for neural network approaches). However, CBR requires its own knowledge sources, such as indexing and similarity knowledge and knowledge to adapt cases to fit new needs. While it may be possible to develop, e.g., standardized types of indices for the smart homes domain (following the example of standardized indexing vocabularies of (Domeshek 1992; Leake 1992b; Schank *et al.* 1990)), enabling multiple projects to leverage this research, a major effort would be required to develop indices or similarity criteria covering a wide range of scenarios.

Segmentation and distillation of cases from the information stream: When a CBR system captures cases based on observing actions in a “messy” real-world environment, events may need to be recognized from sensor changes over time, and new methods will be needed to delineate case-relevant information and determine case boundaries in an information stream.

Integrating similarity judgments with additional information sources: In the basic CBR model, each problem is addressed by identifying the most similar case and adapting its solution to the new situation. However, in problem detection tasks, it may not be desirable to predict problems whenever the most similar case was a problem case, due to the risk of false alarms. Addressing this issue may require combining case-based prediction with other reasoning methods. For example, CBR could provide candidate predictions which could then be further scrutinized by Bayesian methods (e.g., based on patient history, probabilities of particular events, etc.) as part of the case evaluation process.

Case adaptation: Automated case adaptation has long been recognized as a key challenge for case-based reasoning (Barletta 1994; Leake 1996; Mark, Simoudis, & Hinkle 1996). Although many methods exist (see (Mantaras *et al.* 2006) for a recent survey)), it can still be difficult to capture the needed case adaptation knowledge. A number of systems have explored interactive adaptation methods (e.g., (Smith, Lottaz, & Faltings

1995)), including methods in which system adaptation is augmented by the user as needed (Gervasio *et al.* 1998)). With appropriate interfaces, such methods might provide a means for competent inhabitants and care-givers to initially guide the application of prior cases as needed, in conjunction with methods to allow traces of their adaptations to themselves be captured and reused by CBR, as in (e.g., (Leake, Kinley, & Wilson 1995)).

Case-base maintenance: As CBR systems have received extended use, maintenance of CBR systems has become an active CBR research area (e.g., (Leake *et al.* 2001)). In smart homes, changes in the home and inhabitants (e.g., during the course of a terminal disease) would require methods to address both the volume of new cases and the changing circumstances which might render old cases or similarity metrics obsolete (Leake & Wilson 1999; Zhang & Yang 1998).

Event mining for case generation: The monitoring processes of smart homes will provide a rich stream of information to mine for cases. However, segmenting these into meaningful cases—both indices and responses—may be difficult to automate. Some CBR research has considered mining cases from sources such as database records (e.g., (Yang & Cheng 2003)) and has considered issues in capturing and controlling continuous phenomena through CBR (Ram & Santamaria 1993). However, this area remains largely unexplored.

Exploiting multiple case sources: As already mentioned, a question for CBR in smart homes concerns how to enable the system to function during early use, when little experience is available. If the smart home system begins with an initial set of standard schemas, those schemas will provide a basis for initial processing, but will not reflect special individual needs.

An interesting alternative is case-based experience sharing. Many opportunities to support human decision making can come from other people's experiences. When an individual encounters a situation that is novel to him or her, but potentially successful solutions can be collected from the experiences of other individuals, CBR can empower the individual with the ability to make more informed decisions. Collective case memories provide the knowledge that a single individual may lack.

Approaches such as *multi-case-base reasoning* (MCBR) (Leake & Sooriamurthi 2004) are designed to address gaps in a single reasoner's case base by drawing on external case bases and adapting externally provided cases to reflect differences in task and user characteristics. Applying MCBR requires developing adaptation methods not only for particular tasks, but also for differences in user characteristics. In the smart homes domain, MCBR approaches might also need to address instances of cases which simply conflict, by choosing between a range of alternatives for handling a particular situation.

It may also require addressing potential differences in the vocabularies in which cases are captured. Otherwise, cases and events may fail to match simply due to differences in terminology between case descriptions and event descriptions. For example, an event *stove-turned-on* cannot trigger the selection of a case describing the same event as *range-turned-on*. This vocabulary issue is well known in AI, and can be overcome by

using ontologies (Gruber 1993) to provide standardized representations (e.g., an ontology could simply specify that *stove-turned-on* and *range-turned-on* are equivalent). In addition, ontologies may capture the information required to facilitate further reasoning about cases or to enable their adaptation to new circumstances.

Internal system confidence: The smart homes domain has a marked difference from many CBR applications. Many CBR systems are advisory to an expert, who makes the final decisions; autonomous systems making judgments about health and safety issues must address substantial risk issues. Recent research has begun to examine how CBR systems can confidence in their own solutions (Cheetham & Price 2004). However, this is only a first step towards assuring the needed system reliability.

Trust: Related to internal confidence is external trust in the quality of system decisions. Regardless of the objective performance of a CBR systems, its practical use depends on those with authority over the smart home having sufficient trust to place certain tasks under system control. As mentioned previously, presentation of cases may itself constitutes a useful form of explanation, helping to build user trust in a recommendation. Current CBR research is augmenting this type of explanation with new methods (see (Leake & McSherry 2005) for a collection of this work). Special challenges for building trust may arise if CBR systems are fielded with a limited set of cases, to be augmented as needed in response to failures. In that case, a certain level of failure will be expected during the learning process.

Privacy, security, and control: For smart home technology to gain wide acceptance, issues of privacy, security, and control are of great importance. Occupants must feel in control of the information collected—they must be assured that their privacy is protected, that information is secure from illegal access, and that it is shared and disseminated only to appropriate parties, for the occupants benefit. For example, for eldercare, any abnormalities in a person's behavior that may indicate a serious change in health or well being may need to be shared with relatives and caretakers, while daily routines and more personal choices need to be private. Thus, the smart home environment must provide a flexible security system to adapt to the privacy and security needs of its residents. One strategy for alleviating privacy issues as cases are shared might be to aggregate cases from a number of users before distribution (Smyth *et al.* 2005).

In agent-based systems, issues of privacy and security control have been addressed through the use of policies (Schreckenghost, Martin, & Thronesbery 2002; Barrett 2004). These may restrict both what an agent, or a smart home can and also what they must do, allowing system developers to specify a "legal" framework within agents operate. Applying a CBR approach to control privacy and security has the advantage of not requiring a full "legal" framework to fully function. For example, an initial case base may include the most important restrictions and obligations that the system must follow for privacy and security protection. Subsequently, the CBR system may acquire additional cases derived from interactions with residents that model exceptions to an initially more restrictive and generic framework. However, the use of such cases raises additional issues for, for example, security of case sources and reliability of case selection procedures.

6 Conclusion

The previous sections have discussed three main areas in which the use of case-based reasoning may provide benefits for smart homes: supporting personalization, supporting interactive adjustment of the system by the user, and facilitating customization and knowledge acquisition by the developer. The connections of CBR to human reasoning make cases a promising way to communicate knowledge and to explain system actions, potentially increasing trust. This is important because the success of smart homes will rest not only on their ability to fulfill human needs, but on the willingness of inhabitants or caregivers to entrust themselves or their charges to the home's care.

Case-based methods promise to be applicable to many different processes within smart homes, and case-based methods can be combined with other approaches to exploit their strengths and capture their results for efficient reuse. Likewise, case-based smart home installations may be designed to share experiences with other smart homes, to assure their competence not only for common cases, but for rarely-occurring emergencies. Thus smart homes offer a rich potential area for the application of CBR.

References

- Aamodt, A., and Plaza, E. 1994. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications* 7(1):39–52. <http://www.iiaa.csic.es/People/enric/AICom.pdf>.
- Aha, D., and Munoz, H. 2001. *Interactive Case-Based Reasoning*, volume 14. Kluwer. Special issue of *Applied Intelligence*.
- Allen, J. F.; Byron, D. K.; Dzikovska, M.; Ferguson, G.; Galescu, L.; and Stent, A. 2001. Toward conversational human-computer interaction. *AI Magazine* 22(4):27–37.
- Ashley, K., and Rissland, E. 1987. Compare and contrast, a test of expertise. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, 273–284. San Mateo, CA: AAAI.
- Augusto, J., and Nugent, C. 2005. Smart homes can be smarter. In Augusto, J., and Nugent, C., eds., *Smart Homes Can be Smarter*. Berlin: Springer. In press.
- Barett, R. 2004. People and policies: Transforming the human-computer partnership. In *Proceedings of the 5th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2004)*, 111–116. Yorktown Heights, NY: IEEE Computer Society.
- Barletta, R. 1994. A hybrid indexing and retrieval strategy for advisory CBR systems built with ReMind. In *Proceedings of the Second European Workshop on Case-Based Reasoning*, 49–58.
- Blanzieri, E. 2002. A cognitive framework for personalization of the cbr cycle. In *ECCBR Workshop on Case Based Reasoning and Personalization*.
- Buchanan, B., and Shortliffe, E. 1984. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley.
- Budzik, J., and Hammond, K. 1999. Watson: Anticipating and contextualizing information needs. In *62nd Annual Meeting of the American Society for Information Science*.

- Budzick, J.; Hammond, K.; and Birnbaum, L. 2001. Information access in context. *Knowledge based systems* 14(1–2):37–53.
- Burke, R., and Kass, A. 1996. Retrieving stories for case-based teaching. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press. 93–109.
- Cañas, A.; Leake, D.; and Wilson, D. 1999. Managing, mapping, and manipulating conceptual knowledge. In *Proceedings of the AAAI-99 Workshop on Exploring Synergies of Knowledge Management and Case-Based Reasoning*, 10–14. Menlo Park: AAAI Press.
- Cassell, J. 2001. Embodied conversational agents. *AI Magazine* 22(4):67–83.
- Cheetham, W., and Price, J. 2004. Measures of solution accuracy in case-based reasoning systems. eccbr 2004. In Funk, P., and González, P., eds., *ECCBR-2004: Advances in Case-Based Reasoning*, 106–118. Berlin: Springer Verlag.
- Choudhury, T., and Pentland, A. 2003. Sensing and modeling human networks using the sociometer. In *Wearable Computers (ISWC 2003)*, 216–222. IEEE Press.
- Cook, D. J.; Youngblood, M.; Heierman, E. O.; Gopalratnam, K.; Rao, S.; Litvin, A.; and Khawaja, F. 2000. Mavhome: An agent-based smart home. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications*, 521–524. IEEE Computer.
- Craik, F. I. M., and Salthouse, T. A. 2000. *The Handbook of Aging and Cognition, 2nd Ed.* Mahwah, NJ: Lawrence Erlbaum.
- Cunningham, P.; Doyle, D.; and Loughrey, J. 2003. An evaluation of the usefulness of case-based explanation. In *Case-Based Reasoning Research and Development: Proceedings of the Fifth International Conference on Case-Based Reasoning, ICCBR-03*, 122–130. Berlin: Springer-Verlag.
- Dean, T., and Boddy, M. 1988. An analysis of time-dependent planning. In *Proceedings of the seventh national conference on artificial intelligence*, 49–54. San Mateo, CA: Morgan Kaufmann.
- Dey, A.; Abowd, G.; and Salber, D. 1999. A context-based infrastructure for smart environments. In *Proceedings of the 1st International Workshop on Managing Interactions in Smart Environments (MANSE '99)*, 114–128.
- Dey, A. 2001. Understanding and using context. *Personal Ubiquitous Computing* 5(1):4–7.
- Domeshek, E. 1992. *Do the Right Things: A Component Theory for Indexing Stories as Social Advice*. Ph.D. Dissertation, The Institute for the Learning Sciences, Northwestern University.
- Gervasio, M.; Iba, W.; Langley, P.; and Sage, S. 1998. Interactive adaptation for crisis response. In *Proceedings of the AIPS-98 Workshop on Interactive and Collaborative Planning*.
- Gross, T., and Specht, M. 2001. Awareness in context-aware information systems. In *Proceedings of Mensch & Computer*, 221–232. Wiesbaden: Teubner Verlag.
- Gruber, T. R. 1993. A translation approach to portable ontologies. *Knowledge Acquisition* 5(2):199–220.
- Helal, S.; Mann, W.; El-Zabadani, H.; King, J.; Kaddoura, Y.; and Jansen, E. 2005. The gator tech smart house: A programmable pervasive space. *Computer* 38(3).

- Kolodner, J., and Leake, D. 1996a. A tutorial introduction to case-based reasoning. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press. 31–65.
- Kolodner, J., and Leake, D. 1996b. A tutorial introduction to case-based reasoning. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press. 31–65.
- Kolodner, J. 1991. Improving human decision making through case-based decision aiding. *AI Magazine* 12(2):52–68.
- Kolodner, J. 1993. *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.
- Kolodner, J. 1994. From natural language understanding to case-based reasoning and beyond: A perspective on the cognitive model that ties it all together. In Schank, R., and Langer, E., eds., *Beliefs, Reasoning, and Decision Making: Psycho-Logic in Honor of Bob Abelson*. Hillsdale, NJ: Lawrence Erlbaum. 55–110.
- Leake, D., and McSherry, D. 2005. *Explanation in Case-Based Reasoning*. Kluwer. Special issue of *Artificial Intelligence Review*, In press.
- Leake, D., and Sooriamurthi, R. 2004. Case dispatching versus case-base merging: When MCBR matters. *International Journal of Artificial Intelligence Tools* 13(1):237–254.
- Leake, D., and Wilson, D. 1999. When experience is wrong: Examining CBR for changing tasks and environments. In *Proceedings of the Third International Conference on Case-Based Reasoning*, 218–232. Berlin: Springer Verlag.
- Leake, D., and Wilson, D. 2001. A case-based framework for interactive capture and reuse of design knowledge. *Applied Intelligence* 14:77–94.
- Leake, D.; Smyth, B.; Wilson, D.; and Yang, Q., eds. 2001. *Maintaining Case-Based Reasoning Systems*. Blackwell. Special issue of *Computational Intelligence*, 17(2), 2001.
- Leake, D.; Kinley, A.; and Wilson, D. 1995. Learning to improve case adaptation by introspective reasoning and CBR. In *Proceedings of the First International Conference on Case-Based Reasoning*, 229–240. Berlin: Springer Verlag.
- Leake, D. 1992a. Constructive similarity assessment: Using stored cases to define new situations. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 313–318. Hillsdale, NJ: Lawrence Erlbaum.
- Leake, D. 1992b. *Evaluating Explanations: A Content Theory*. Hillsdale, NJ: Lawrence Erlbaum.
- Leake, D. 1993. Focusing construction and selection of abductive hypotheses. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 24–29. San Francisco: Morgan Kaufmann.
- Leake, D. 1996. CBR in context: The present and future. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press. 3–30. <http://www.cs.indiana.edu/leake/papers/a-96-01.html>.
- Leake, D. 1998. Cognition as case-based reasoning. In Bechtel, W., and Graham, G., eds., *A Companion to Cognitive Science*. Oxford: Blackwell. 465–476.
- Leake, D. 2002. Personalized cbr: Challenges and illustrations. In *ECCBR: Workshop on Case Based Reasoning and Personalization*.

- LoPresti, E. F.; Mihailidis, A.; and Kirsch, N. 2004. Assistive technology for cognitive rehabilitation: State of the art. *Neuropsychological Rehabilitation* 14(1-2):5–39.
- Mann, C. W., and Helal, S. 2004. Pervasive computing research on aging, disability and independence. In *Proceedings of the 2004 International Symposium on Applications and the Internet Workshops (SAINTW'04)*, 244–248. Tokyo, Japan: IEEE Computer Society.
- Mantaras, R.; McSherry, D.; Bridge, D.; Leake, D.; Smyth, B.; Craw, S.; Faltings, B.; Maher, M.; Cox, M.; Forbus, K.; Keane, M.; Aamodt, A.; and Watson, I. 2006. Retrieval, reuse, revise, and retention in cbr. *Knowledge based systems*. In press.
- Mark, W.; Simoudis, E.; and Hinkle, D. 1996. Case-based reasoning: Expectations and results. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press. 269–294.
- Massie, S.; Craw, S.; and Wiratunga, N. 2004. A visualisation tool to explain case-based reasoning solutions for tablet formulation. In *Proceedings of the 24th SGA International Conference on Innovative Techniques and Applications of Artificial Intelligence*. Berlin: Springer-Verlag.
- Meyer, S., and Rakotonirainy, A. 2003. A survey of research on context-aware homes. In *CRPITS '03: Proceedings of the Australasian information security workshop conference on ACSW frontiers 2003*, 159–168. Darlinghurst, Australia: Australian Computer Society.
- Mozer, M. 2005. Lessons from an adaptive house. In Cook, D., and Das, R., eds., *Smart environments: Technologies, protocols, and applications*. Hoboken, NJ: Wiley. chapter 12, 273–294.
- Musliner, D., and Pell, B., eds. 1999. *Agents with Adjustable Autonomy*. Menlo Park, CA: AAAI Press. Technical Report SS-99-06.
- Mynatt, E.; Essa, I.; and Rogers, W. 2000. Increasing the opportunities for aging-in-place. In *Proceedings of the ACM Conference on Universal Usability*, 65–71. Arlington, VA, USA: ACM.
- Novak, J., and Gowin, D. 1984. *Learning How to Learn*. New York: Cambridge University Press.
- Pentland, A. 1996. Smart rooms. *Scientific American* 274(4):68–76.
- Pollack, M. E.; Brown, L.; Colbry, D.; McCarthy, C. E.; Orosz, C.; Peintner, B.; Ramakrishnan, S.; and Tsamardinos, I. 2003. Autominder: An intelligent cognitive orthotic system for people with memory impairment. *Robotics and Autonomous Systems* 44(3-4):273–282.
- Pollack, M. 2005. Intelligent technology for an aging population: The use of AI to assist elders with cognitive impairment. *AI Magazine* 26(2):9–24.
- Ram, A., and Santamaria, J. 1993. Continuous case-based reasoning. In *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning*, 86–93. Washington, DC: AAAI. AAAI Press technical report WS-93-01.
- Rhodes, B., and Starner, T. 1996. The remembrance agent: A continuously running automated information retrieval system. In *Proceedings of The First International Conference on The Practical Application of Intelligent Agents and Multi Agent Technology (PAAM '96)*, 487–495.

- Rhodes, B. J. 2000. Margin notes: Building a contextually aware associative memory. In *The Proceedings of the International Conference on Intelligent User Interfaces (IUI '00)*.
- Riesbeck, C. 1996. What next? The future of CBR in postmodern AI. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press. 371–388.
- Schank, R., and Leake, D. 2002. Natural language understanding: Models of roger schank and his students. In *Encyclopedia of Cognitive Science*. London: Nature Publishing Group. 189–195.
- Schank, R.; Osgood, R.; Brand, M.; Burke, R.; Domeshek, E.; Edelson, D.; Ferguson, W.; Freed, M.; Jona, M.; Krulwich, B.; Ohmayo, E.; and Pryor, L. 1990. A content theory of memory indexing. Technical Report 1, Institute for the Learning Sciences, Northwestern University.
- Schank, R. 1982. *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge, England: Cambridge University Press.
- Schreckenghost, D.; Martin, C.; and Thronsbury, C. 2002. Specifying organizational policies and individual preferences for human-software interaction. In *Proceedings of AAAI Fall Symposium on Etiquette for Human-Computer Work, Technical Report FS-02-02*, 32–39. North Falmouth, Massachusetts: AAAI.
- Schulz, S.; Roth-Berghofer, T.; and Leake, D., eds. 2005. *Modelling and retrieval of context: Proceedings of the second international workshop MRC 2005*. CEUR. In press.
- Smith, I.; Lottaz, C.; and Faltings, B. 1995. Spatial composition using cases: IDIOM. In *Proceedings of First International Conference on Case-Based Reasoning*, 88–97. Berlin: Springer Verlag.
- Smyth, B., and Cotter, P. 2005. Personalized electronic program guides for digital TV. *AI Magazine* 22(2):89–98.
- Smyth, B.; Balfe, E.; Freyne, J.; Briggs, P.; Coyle, M.; and Boydell, O. 2005. Exploiting query repetition and regularity in an adaptive community-based web search engine. *User Modeling and User-Adapted Interaction* 14:383–423.
- Streitz, N.; Röcker, C.; Prante, T.; van Alphen, D.; Stenzel, R.; and Magerkurth, C. 2005. Designing smart artifacts for smart environments. *Computer* 38(3).
- Watson, I. 1997. *Applying Case-Based Reasoning: Techniques for Enterprise Systems*. San Mateo, CA: Morgan Kaufmann.
- Weiser, M. 1991. The computer for the 21st century. *Scientific American* 265(3):94–104.
- Widmer, G., and Kubat, M. 1996. Learning in the presence of concept drift and hidden contexts. *Machine Learning* 23(1):69–101.
- Wiratunga, N.; Craw, S.; Taylor, B.; and Davis, G. 2004. Case-based reasoning for matching SMARHOUSE technology to people's needs. *Knowledge based systems* 17:139–146.
- Yang, Q., and Cheng, S. 2003. Case mining from large databases. In *Case-Based Reasoning Research and Development: Proceedings of the Fifth International Conference on Case-Based Reasoning, ICCBR-03*, 691–702. Berlin: Springer-Verlag.

Zhang, Z., and Yang, Q. 1998. Towards lifetime maintenance of case base indexes for continual case based reasoning. In *Proceedings of the 1998 International Conference on AI Methodologies, Systems and Applications (AIMSA-98)*, 489–500. Berlin: Springer Verlag.