



Adaptive Information for Consumers of Healthcare

Alison Cawsey¹, Floriana Grasso², and Cécile Paris³

¹ School of Mathematical and Computer Sciences, Heriot Watt University, Edinburgh, UK,
alison@macs.hw.ac.uk

² Department of Computer Science, University of Liverpool, UK,
F.Grasso@csc.liv.ac.uk

³ CSIRO ICT Centre, Sydney, Australia, Cecile.Paris@csiro.au

Abstract. This chapter discusses the application of some of the technologies of the adaptive web to the problem of providing information for healthcare consumers. The particular issues relating to this application area are discussed, including the goals of the communication, typical content of a user model, and commonly used techniques. Two case studies are presented, and evaluation approaches considered.

15.1 Introduction

So far this book has looked at some of the techniques that have been developed for the adaptive web, focusing on how we model the user, and how we use that information in adapting the user's experience. In this chapter we show how some of these ideas apply to one particular application area: the provision of information to consumers of health care.

In recent years the way in which people are involved in their own health care has changed dramatically [47]. While, in the past, the almost exclusive source of information was the medical staff directly concerned with the provision of care, nowadays the Internet and the World Wide Web have provided new opportunities for a new generation of users, the "health information consumers". These have been defined by organisations like the American Medical Informatics Association as people who seek information on various aspects related to health and well being, like health promotion, disease prevention, management of long term conditions, and so on. Health information consumers are therefore not only patients, but also their family and friends, or simply people concerned about health.

An increasing number of people are now using the Internet to support their health-care [58], and the amount of information available on the Web continues to grow. The information needs of healthcare consumers are different from those of the members of the healthcare team (see [59, 53, 28] for some examples of research in health information systems aimed at health care providers). For example, patient-oriented health information systems may include providing information to promote patient choice, informed consent, self-care and shared patient-doctor decision-making (e.g., [46]). Providing such health information via adaptive web-based systems offers new possibilities

for pursuing public health objectives like providing knowledge and inducing behaviour change. Furthermore, recent studies have shown that web-based interventions (to provide knowledge and induce behavior change) can have more impact than non web-based interventions [73]. This includes increased knowledge about conditions and treatment, increased participation in health and more uptake of behaviour changes. In addition, sites that pointed readers to relevant, individually tailored material reported longer session times per web-visits and more visits.

There is also evidence that decontextualised, impersonal and generic health information, as typically found on the Internet, has less impact than health information tailored to the individual, at least in some situations (e.g., [3, 49, 18, 69, 70, 68]).

There has therefore been much interest in how we can design systems capable of tailoring information to the health care consumer, and exploiting the great potential to enhance health information and education through web delivery – applying ideas from adaptive web-based presentations and adaptive hypermedia to the problem of providing users with relevant, appropriate, understandable, and potentially persuasive information relating to their needs. There are particular issues in this area to be aware of, focusing now on patients as our main healthcare consumer.

First, we need to consider some of the goals of patient information and education. Patient information may be intended to *inform*, to *enable decision-making* or to *persuade*. We may, for example, want to: inform the user about their condition or about the side-effects of their treatment; give them enough information to enable them to take an active role in the decision-making concerning whether or not to have surgery; or persuade the patient to improve their diet. Persuading the user of a course of action may be part of encouraging patient *compliance* (or *adherence*) – we may want to encourage and motivate them to go along with the treatment regime proposed and take the necessary actions.

Whatever the objective of a healthcare communication, different patients have different individual needs. A good healthcare professional will recognise this and adjust the content and level of verbal information to the patient's perceived needs (both informational and emotional) and their level of understanding. He or she may also ensure that the language employed is both understandable and appropriate for a specific patient, remembering that, first, most patients are not medical experts, and, second, they might already be under considerable cognitive load and stress due to the situation.

This contrasts with current written sources of information (e.g., leaflets and websites) which are normally targeted at the typical patient, not at the individual. Yet written information is also of vital importance in healthcare communication. Verbal messages are often forgotten, while written information is there for reference, and potentially provides a shared information source for patient, family and friends. Recognising this, for example, a genetic counselor will always provide patients or carers a one- to two-page letter summarising the information that was given to them verbally during the consultation [4].

Given the need for personalised or tailored information and the benefits of written sources, many researchers have explored how we can automatically adapt the content of healthcare messages to the patient (or more generally, to the user). Information may be delivered through printed leaflets, online via adaptive websites, or through phone/text

messages. Similar methods of content adaptation can often be applied whatever the means of delivery. Conversely, while there are some peculiarities in each application area, general techniques for adaptive content presentation on the web apply whatever the domain. These general techniques are well described in, for example, Chapter 13 of this book [17]). In this chapter we will therefore concentrate on the issues that specifically arise in healthcare information.

First, we must take seriously issues of privacy, security and trust. Patients are unlikely to use a system where their personal medical details are potentially accessible by others. Furthermore, they need to trust the source of information. Second, in healthcare information, we are not just concerned with informing and educating, but also with the patient's emotional state and attitude (e.g., [3, 43, 35]). We have to take account of the patient's emotional needs and their willingness to accept and commit to change. An effective communication is not the one that is merely learned and remembered, but the one that enables the patient to talk about their problems and come to decisions or acceptance concerning their medical problems [55]. Finally, there may be an issue of control: i.e., patients may want to be able to control what information a system has about them and know how it is being used.

The rest of this chapter will look first at what we may be trying to achieve in personalised health communication, then at the user model (e.g., the attributes of the patient to whom the system is adapting the information) and at techniques that can be used to produce personalised healthcare information. Two case studies will be given, illustrating the range of applications and techniques. We will then look in detail at how personalised health care communication systems can be evaluated, and in particular whether evaluation methodologies from the medical domain can be usefully applied.

15.2 Health Education Goals

Before looking at *how* we personalise health materials, it is worth considering *why* in more detail. While different health professionals have different perspectives on this, two objectives are frequently discussed. The first objective is to support the patient in making decisions about their treatment (shared of course with the health professional team):

The overriding goal of patient education should be to support the patient's autonomous decision-making, not (as it has been conceptualized) to get patients to follow doctor's orders. [62]

The second objective often discussed is compliance (i.e., following the prescribed treatment and care plan). Compliance is a very important problem in health care, with many implications, both medical and socio-economical. It is estimated, for example, that in the European Union between 2% and 20% of the medical prescriptions never get to the pharmacy, and that about 125,000 deaths and 5-10% hospitalisations per year can be attributed to lack of compliance. Compliance might be achieved by a number of ways (e.g., [40]), and, conversely, non-compliance might be explained by a number of factors. For example, compliance has been shown to be correlated to the patient's understanding

of their condition and prescribed treatment (e.g., [30, 42]). Indeed, some patients need to understand the rationale for their treatment, and why it will work. For example, they may need to understand what a specific drug does. They might also need to understand why and how their own actions (e.g., exercise, taking medication) are necessary for success. With this understanding, they are more likely to follow the treatment regime recommended. But understanding alone is not enough. Patients also need to be committed to the treatment, and this may require convincing them of its necessity, by ensuring they both understand and truly believe the consequences of failing to follow a specific treatment (which might include a change of lifestyle). Finally, patients are more likely to follow a treatment or advice if they trust it and its prescriber.

Compliance with the doctor's treatment plan and autonomous decision making by the patient are sometimes presented as opposing points of view. However it seems more likely that both perspectives should be supported. It is not always appropriate to leave the patient to make the decisions, and they will often not want that role, while they may want to participate in the decision-making. Note also that patients might be more likely to comply to a treatment if they were involved in its choice.

Depending on the objective, different types of information might be provided to a patient. Where treatment choice is an issue, patients may receive background information about their conditions (e.g., what causes it if it is known, its symptoms, its consequences, what can be done about it), and specific information about the alternative treatments and why a particular one is more appropriate for them. For patients with chronic disorders (e.g., asthma, diabetes), appropriate information might include information that helps them manage their own care effectively, and that provides advice as to when to call out a health professional. In addition, there are today broader time-independent health promotion objectives, addressed to groups or the population at large, as opposed to an individual at a particular point in time. For a healthier society we want to promote a good diet, exercise, stopping smoking, avoiding direct exposure to the sun, and so on. While these are almost universally recognised goals, they may be more effectively achieved by addressing the individual – by personalising the advice and the information, e.g. [18, 70].

While supporting choice and promoting a particular course of action are perhaps the easiest health education goals to characterise, much of the information giving in healthcare has a less explicit objective. With more appropriate and understandable information, patients will usually feel more in control. If they know what will happen next, which health professionals will be managing their care, and how they should prepare for any treatment, then their anxiety is likely to be reduced. Anxiety and stress reduction is therefore another important objective in health education, but a difficult one to get right. Where patients have a poor prognosis, it is particularly difficult to get a balance between sensitivity and openness, and one that a machine is unlikely to achieve.

Currently most patient information is provided verbally or through leaflets, with an increasing number of patients turning to the web for further information [58] and an increasing numbers of health information websites (e.g., [13, 14, 12]). A typical website or leaflet will focus on a particular condition, and give general information, information about diagnosis, and information about treatment, including any options and alternatives and any actions that the patient can take to help themselves. These existing

resources are very much disease centred. They are not tailored to the patient's specific needs and knowledge. As a result they can sometimes be confusing or overwhelming to a patient. Adaptive and personalisation techniques open the way to more patient-centred sources of information and potentially more effective means of achieving the health education goals described above.

Effective health education is not just about making life better for the healthcare consumer, but it is also about making the process more efficient, using the available money and resources as effectively as possible and potentially saving our governments' money (e.g., [45]). By providing means for the patient (and their carers) to obtain information outside a doctor or hospital visit, there is a possibility to move some of the health care services to the home or the community. If care is to be shared between health professionals, community and patients, then each must have an appropriate level of understanding of the medical issues, as well as who to call when. If this is done effectively there is the potential to make better use of specialist expertise, and save on unnecessary hospital visits. Being able to automatically create personalised communications appropriate to context and need may prove to be a vital part of this process.

15.3 The User Model

Having briefly reviewed some of the objectives of patient education, we can turn back to how we can adapt health information to a patient, taking into account the particular goals that health professionals recognise as being important for that patient. In this section, we briefly discuss what needs to be captured in a user model in order to provide tailored information that achieves the objectives discussed above.¹

First, it will usually be necessary to acquire and capture factual information about the patient, their condition, current treatments, and so on. This information may be available in the patient record (e.g., [20]). It is thus possible (and relatively easy) to produce patient-centred information by starting with the information available from their record. Just this amount of tailoring is likely to be an improvement over a general health education leaflet that is typically disease-centric and does not take into account a patient's particular characteristics. For example, instead of including information related to all possible treatments for a condition, a patient-centric information system (or leaflet) may only contain information about treatments relevant to the patient. Similarly, if we know that the patient is being treated by a particular consultant in a specific hospital, the information might include whom to contact where, how to get to the hospital, where to park, information about visitor's hours, etc. In other words, it is possible to produce one coherent, concise and practical information source containing all the information that is important and relevant to the patient, and that he or she is likely to seek.

It is worth mentioning practical issues in using the patient record. There are of course major security and privacy issues when accessing this confidential information,

¹ We address in this chapter issues specific to patient oriented health information systems. For general overviews of user models for educational systems and personalised information access, see Chapters 1 [15] and 2 [31] of this book.

and using the patient record for web-based systems is still problematic. Typically web-based systems use more limited information on the patient's health obtained through an online questionnaire. Or, when they use more extensive data, they rely on a password based authentication. This is however likely to be insufficient in many cases, and it is expected that smart cards, private keys, or encryption will be increasingly used [6]. (See Chapter 21 of this book [44] for a discussion of issues related to privacy and security.)

Health education also shares the characteristics of traditional education in that it must be delivered at a level that will be understood by the individual concerned, taking into account at least the patient's literacy, medical and otherwise. (See Chapter 1 of this book [15]).

Most crucially, however, as well as information about the user's medical conditions and treatment, health education systems may need to take into account more complex factors, such as the patient's current mental and emotional state, their ability to make decisions and perform complex actions, or their acceptance of their disease. This is the case, for example, if the patient has just received news about a life-threatening disease, and his or her ability to absorb information may be impaired. In other cases, health education may be about changing attitudes and behavioural change – for example, a reason for providing information may be to convince patients to change their diet, to stop smoking or to start exercising. In these cases, then, the patients' motivation level, their willingness to accept treatment or make changes, as well as their desires and intentions all become important. It thus seems at least plausible that adapting materials to some of these factors will make written and online materials more effective. This in fact has already been shown, as discussed earlier, e.g., [18, 69, 70, 68].

So, the user model for a health education or health promotion system will very often include the information obtainable from the patient record, but may also include a whole range of cognitive factors, such as the ones mentioned above (e.g., current understanding, motivation and anxiety). The user model may capture factors related to different personality types (which might provide insight, for example, as to how a patient is likely to deal with change or bad news in different ways – e.g., [35]). This aspect is what makes the provision of healthcare information a challenge. In tackling this task, it is sensible to ground the user model and the information adaptation on well established behavioural theories. One example of such theories is the Stages of Change Model, or Trans-Theoretical Model [61]. The model assumes that people progress through very distinct stages of change on their way to improve health:

1. *precontemplation*: people at this stage see no problem with their behavior and have no intention of changing it. They mainly lack information, in the sense that they have not been presented yet with any convincing reason to change their behaviour. Often people are not very open to receiving advice.
2. *contemplation*: in this stage, people come to understand their problem, its causes, and start to think about taking action to solve it, but have no immediate plans. This is a delicate stage, as there is always the risk to miss the opportunity, and go back to precontemplation, because of laziness or old influences.
3. *preparation*: people are planning to take an action, and are putting together a plan, but have not taken any step yet. This is a sort of transition stage between the decision to act and the action itself. Often one of the causes of going back to a pre-

vious stage is that the plan is too ambitious, and the life style change planned is too drastic.

4. *action*: people are actually in the process of actively making behaviour changes. The concern here is to pay attention to negative emotions: anger, depression, anxiety, apathy, insecurity, etc., in order to prevent relapse.
5. *maintenance*: health behaviour continued on a regular basis. The state is more stable than the action one, but there is always the possibility of relapse.
6. (*termination*) at this stage, the former problem no longer presents any temptation or threat. Many people never reach this stage.

In addition to providing a classification of the user, the model suggests strategies for recognising and dealing with each stage of change, in terms of the information that should be presented at each stage.

For example, the precontemplator needs to identify the problem in the first place, so one may provide information on related problems. It is also likely that precontemplators have misconceptions about the consequences of their actions, so one should assess prior knowledge and clarify misunderstandings. On the other hand, those in the “action” stage mainly need to get things going, by means of tips and strategies to maintain and enhance their commitment. They need reinforcement too, and encouragement.

User models containing this kind of information have indeed already been used. For example, systems generating patient education with the goal of achieving behaviour change (e.g., diet, smoking) have captured patients’ attitude towards a specific change, exploiting the stages of change model [70, 64].

A user model in a health application can thus be quite complex. This leads to the question of how we obtain and update such a user model. Using the medical record is easy, and changes in the patient’s treatments is generally reflected in changes in their record. However, the patients’ record may not always be available to the health information system. In such cases, we need other ways to obtain the appropriate attributes of the patient. This of course provides a number of challenges. For some attributes it might be possible to let the patient fill in a simple questionnaire, but more thought is needed when considering how to capture some of the more subtle aspects that might support effective information provision (e.g., patients’ personality, mental and emotional state).

There are however various instruments that can be used here (mostly standardised questionnaires) which can be applied to ascertain personality type, stage of change, anxiety level, and so on. While these instruments may be seen as moderately intrusive and as potentially not always leading to accurate results, they are already used successfully in on-line health diagnostic and intervention applications, for example to treat depression using Cognitive Behavior Therapy (e.g., [21, 23]). One approach then is to use these existing tools to populate the user model of an accompanying health information system, which can then exploit this information to provide the patient with relevant information about their condition.

Increasingly, researchers are also investigating new, less intrusive, methods for capturing some characteristics of the user, in particular emotional state and stress/anxiety levels, such as by the use of physiological sensors (e.g., [60, 71, 19]). However, these are still at early stages of research, mostly applied to the domains other than healthcare,

and it is as yet unclear whether or how we can use these measures reliably outside of the experimental situation.

Finally, obtaining the attributes of the patient at one point in time is not enough. Having acquired details of the patient's current state and stored them in a user model, a system needs to be able to monitor the patient's state and update these details as their state changes. As health education is often about changing the patient's mental state (e.g., their beliefs, attitudes, anxiety, etc.), a system needs to be able to monitor these as well as the attributes related to their health problem. By monitoring the user and updating the model the system can both provide more appropriate and timely information, and also assess the effectiveness of its past interventions. For certain physiological attributes it is now possible to use small wearable monitoring devices to achieve this, providing a constantly updated model. However, in general these issues are still research challenges.

15.4 Techniques for Adaptation

As discussed previously, the goals of healthcare communication can be quite varied, including persuading the patient to take an action, enabling them to manage their care, supporting informed shared decision-making between patients and health professionals, and reducing stress and anxiety. From a broader perspective, the goal is to provide the patient with information that is *relevant* to their condition and to their situation, which enables them to understand and take control of their condition at a level appropriate to them. The specific goals will then depend on the patient and their situation, while their mental and medical state and the practical situation will influence how information can be best selected and presented to be most effective.

Having reviewed the goals, we also need to consider the nature of the communication itself. Many projects have simply generated personalised materials (leaflets or simple websites, e.g., [38, 64, 20]), where the content and style of the material is adapted to the user, but the interaction style is fixed and simple. Personalised email or text messages have also been used [48], but again with little dynamic interaction with the user. While recognising that the interaction or dialogue style may be important, in particular to acquire and maintain the user model, we focus here mostly on how content is selected, adapted and presented to the patient, given a user model. Indeed, techniques for adapting the information that is generated are similar whether a system is interactive or not. Dialogue issues are taken up in more detail when we discuss the HOMEY project [39] as a case study in the next section.

The most common techniques employed to produce tailored text-based material are based on Natural Language Generation (NLG)². In the health domain, several projects have used these techniques to generate adapted primarily text-based material (e.g., [54, 65, 20]). NLG techniques are concerned with the automatic production of coherent and appropriate textual documents from structured data [50, 51, 63]. They have also been applied in recent years to the generation of appropriate and coherent multimodal

² The reader is also referred to Chapter 13 of this book [17] for some techniques for adaptive content presentation on the web.

documents (e.g., [2, 29, 24] and hypertext presentations (e.g., [26, 56, 25]). Broadly, these techniques are divided into planning what to say (content) and deciding how to express it once there is a message to express. Planning what to say usually starts with communicative goals (e.g., persuade the hearer/reader to take some action, inform the hearer/reader of a fact or situation). The content planning process typically uses domain information from a database or knowledge base, and information about the user (from the user model/profile). The process thus selects information to present and organises it into a coherent whole. The output of this process is a sequence of primitive messages (e.g., informing the hearer/reader of a simple fact) which, given the user model, should achieve the communicative goal. Given a sequence of primitive messages, the question then arises as to how they should be expressed. For example, the following questions must be addressed: Should each fact constitute one sentence, or should they be conjoined? Should facts be announced bluntly or in a more indirect way? Should they be presented formally or less formally? What specific words and constructions should be used? Adapting this stage to the user may be as important as adapting the content and organisation of the information in healthcare communication, where the emotional state (and cultural status) of the user are important.

While natural language techniques remain important, health education and information provision has recently become of wide interest (see, for example, a number of workshops and symposia related to this topic – e.g., [9, 34, 1]), and other techniques have been investigated in healthcare communication for both health care providers and consumers. These include speech (for example the generation of voice messages over the telephone, e.g., [33], as is briefly described below), search and summarisation (e.g., [52, 28]), hypermedia and virtual reality (c.f. [34]) and Embodied Conversational Agents (ECA), e.g., [8].

15.5 Case Studies

Having introduced some of the issues and techniques in personalised patient information, we can turn to two specific projects that have made use of these techniques in practical applications. The first project illustrates how fairly simple patient-centred materials can be generated given a patient record. The second project is more ambitious, providing adaptive advice in the context of a multi-modal dialogue.

15.5.1 Personalised Information for Patients with Cancer

The first case study (Piglit) is a project concerned with creating personalised materials (online and written) for patients undergoing treatment for diabetes [11] and cancer [20]. The main goal of the project has been to provide materials that are patient centred, and which allow the patient to quickly access additional materials of interest. The techniques used have been generally simple, but the systems and approaches produced have been thoroughly evaluated with many patients.

In this project the patient's medical record is used as the main source of information about the patient. While there are many formats for computerised records, it

will hold information on, at least, the patient's medical conditions and treatments. The Piglit project provided patients with online access to this record, with hyperlinks allowing access to explanatory information about their conditions and treatment. These explanatory pages were generated dynamically from a simple knowledge base of medical information, used in combination with information from the patient record, allowing the explanations to be geared to the patient's likely information needs.

Figure 1 illustrates an example page of information for a patient with prostate cancer. Italicised terms were hyperlinks taking the patient to more general, but hopefully relevant information about their condition.

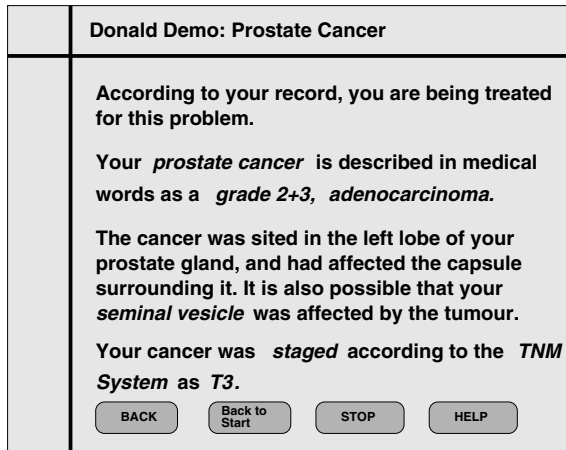


Fig. 15.1. Personalised Health Information linked to Medical Record

This system used simple planning methods from natural language generation to determine content. These techniques were used to plan what to say given a topic (e.g., prostate cancer) and a particular patient record. A simple knowledge base contained the facts about conditions, treatments, and other relevant medical concepts, and this could be exploited to generate pages of information.

The system was evaluated in a large randomised controlled trial (see section on evaluation) [41], which compared the personalised version with a general information system providing very similar information, and with standard leaflets. Patients using the computer used a touch screen system located in a room in the cancer centre where they were receiving treatment. They also received printouts of the information presented in their session. 525 patients participated in the study. Questionnaires were used to gauge the patients views of the system, and also to assess their anxiety levels (before and after the intervention). Statistical analysis of the results was done to compare personalised *vs* non personalised and computer *vs* leaflet.

The results showed that the patients receiving the personalised book were (at a statistically significant level) more likely to think they had learned something new ($p = 0.02$) and that the information was relevant ($p = 0.03$), and were more likely to show the information to family and friends ($p = 0.035$). However, this is perhaps not

surprising as they received information about their own specific conditions/treatment not available to the group receiving general information. Unexpectedly, we found that patients with personalised information showed better improvement in anxiety over three months than those with more general information, despite receiving information on their condition that might be worrying. Three months after the intervention 37% of patients in the general computer information group were still anxious compared with only 19% in the personal information group, with the intervention a significant predictor of anxiety level ($p = 0.001$).

15.5.2 Personalised Home Monitoring to Support Continuity of Care

Our second case study is concerned with the needs of chronic patients, such as Hypertension or Diabetes patients. For these patients, the main objective of a health care system is to ensure compliance to the therapeutic and lifestyle regime over long term periods. In hypertensive patients, for example, it is crucial to maintain a healthy lifestyle (in terms of nutrition habits, doing exercise, stopping smoking and so on), and also to carefully monitor blood pressure, heart rate and weight. Frequent visits to the cardiovascular unit would be highly beneficial, both to keep the doctors updated with the patient's situation and to reinforce the health promotion message to the patient. However, this is difficult to implement, both because of the lack of resources from the hospital, and because of the tendency for patients, especially from long term patients, to relax their attendance to the meetings.

We describe one solution to this problem, as proposed by the HOMEY project [7, 33, 39]³, that developed a system able to efficiently communicate with the patient and to improve the information flow between patients and medical staff. The aim is to allow patients to use the system to communicate as frequently as possible their test results to their care team, while the system has the opportunity to enquire about lifestyle changes, and update the patient's record. The medical staff read the updates, and possibly make new recommendations, which in turn are stored in the system and passed to the patient at the next contact. In order to achieve this, the system makes use of natural language dialogue technology.

Many dialogue systems are based on a "scripted conversation" approach. This means that the main structure of the dialogue is fixed once and for all by the dialogue designer, in order to have control of what can happen in the dialogue. While this is simple and effective for very focussed applications (like telephone banking), it becomes rapidly expensive and too inflexible for complex situations, where there are many objectives to take into account and a very large domain knowledge, like the medical domain. In these situations, sophisticated "intelligent" dialogue systems are more appropriate. The HOMEY system is based on intelligent dialogue technology, and is able to manage a conversation with a patient, adapted to the patient's needs, preferences, and clinical history, but also taking into account the physician's goals. The system supports multimodal input and output, combining the generation of dynamic HTML pages and

³ This is a project funded by the European Community in the programme "Health Care for Citizens" (5th Framework, Project No. IST-2001-32434). The project started in 2001 and was completed in 2004.

VoiceXML [72] sentence fragments. This allows the user to contact the system either with a simple phone call, or with a traditional computer connection. In the latter case, the user is free to choose a speech input/output, a keyboard input with a visual output, or a combination of both.

Intelligent dialogue systems need to keep track of many of the user's characteristics, in order to handle a real time dialogue. In addition to information on the evolution of the disease and the treatment, for instance, the user's goals and beliefs about the medical treatment will be needed for the system to be able to better promote, justify or reinforce the particular piece of advice that it gives. Also, the history of the past interactions with the system will give information on what to ask and what to talk about the next time round. For example, if the user had said he would try to stop smoking, the system may want to check whether the plan had been implemented, and if not, give some more motivation for the user to start doing so.

Generally speaking, when producing intelligent dialogue systems, many phenomena have to be taken into account. These can be broadly divided into *high level* and *low level* phenomena. The former include very general notions like the goals of the dialogue participants, or the strategies for producing persuasive messages, and so on, which are assumed to be independent both of the language and on the output medium. The latter include what it takes to actually produce the single message, like the grammar of the language to use, whether to use speech or text, and so on. These low level issues may be important in adaptation too. For instance the system can try to use the user's vocabulary as much as possible, in order to be better understood.

In the HOMEY project, both levels are taken into account and are dealt with in an architecture based on the concept of *abstract task specification* (see Fig. 2). This structure gives information on two important aspects of what the system should do next: the "plan" representing the high level task to be executed (such as take patient's measurements and make a decision on referral to the clinic), and the definitions of the objects involved in this plan (such as, "heart", "blood pressure", "measurement devices", etc.), together with the relationships among them, that come from a "domain ontology", that is a conceptual representation of the domain.

The abstract task specification is then transformed into the *high level dialogue specification*. The main purpose of this specification is to give some structure to the conversation. The initial dialogue structure depends on the task specification. For instance, if the plan says that a decision cannot be taken until all the patient's measurements are in, then the first part of the dialogue will involve asking the user to report his measurements. This initial structure is however flexible, and can adapt to the way in which the dialogue evolves. For instance, consider the following dialogue:

System: What is your heart rate?

Patient: What do you mean?

Here the user asks for clarifications before replying to the system's question, so the system will have a new "obligation" to fulfill, in addition to those coming from the task specification, and has to take a decision about what to do next (typically, the obligations coming from the user will be dealt with first).

Also, the user may take some initiative in reacting to the system's question. For example, consider the following dialogue:

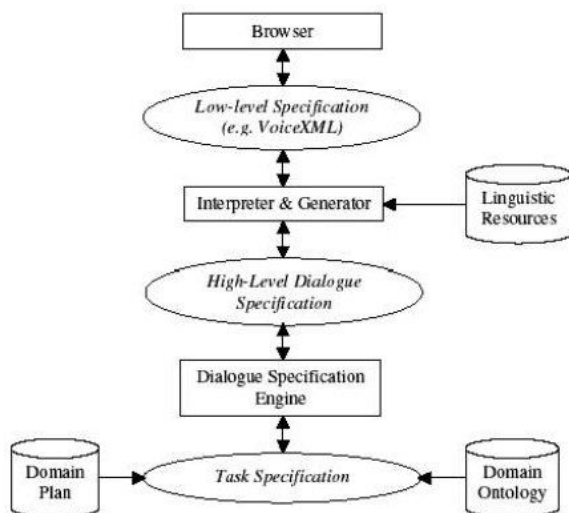


Fig. 15.2. Architecture of the HOMEY system

System: What is your heart rate?

Patient: My heart rate is 67,
and my blood pressure values are 90 and 120.

Here the user has anticipated a system's question by providing more information than requested, so the dialogue specification has to account for the fact that this sub-task has already been accomplished, and may move on. This knowledge comes again from the task specification, which says that two tasks (ask heart rate and ask pressure values) can be both part of a super-task (ask patient's measurements).

Another type of task that can be included in the plan concerns checking on the user's lifestyle, and perhaps reinforcing some of the recommendations coming from the physician's goals. This can lead to dialogues like the following:

System: Are you still swimming two times a week?

Patient: Yes.

System: Are you still smoking?

Patient: Yes, 5 cigarettes per day.

System: You should stop smoking.

The examples above show how dialogues might occur in the setting where the user contacts the system via telephone. In these cases the system will typically output a single new move at a time, e.g. a question. In the multi-modal context, the output could be one or more HTML forms, where several questions are presented, and where, if the user has asked for both voice and visual output, VoiceXML will utter the first question on the form, while the language model will enable the user to answer any question on the form in the preferred order.

The system has been evaluated in two Italian hospitals with a hypertension unit, in two studies. The first one involved fictitious patients, that is a number of volunteers who were assigned a disease profile. This study was mainly done to assess the usability of the system. In the second study, the aim was to assess whether using the system would actually make some difference to the health of the patients. A clinical trial was performed, where about 300 patients of the units were assigned randomly to two groups, only one of which using the (telephone based) HOMEY system. The average blood pressure of the two groups of patients was measured before and after the period under observation. While both groups had a significant decrease in blood pressure, the statistical results suggest a trend whereby the group of patients using the system had a greater systolic pressure decrease than the other group. From the point of view of the user's satisfaction, the evaluation was also successful, as it is testified by the fact that, even after the trial was completed, there is still a good number of patients that decided to continue to use the system to report their data.

An extension to the system is currently being investigated [57] in which the user model is enhanced based on three theories of behaviour change: the Social Cognitive Theory, the Health Action Process Approach [5, 67], and the already mentioned Stages of Change Model.

15.6 Evaluation and Uptake

We now return to the issue of how we evaluate adaptive or personalised healthcare information systems. Before we do so, however, it is worth mentioning that, typically, systems are first designed based on a requirements analysis: this is where the designers of the system spend time with both the expert and the intended users to elicit requirements for the system. This stage is an important aspect of usability engineering for any system. (See Chapter 24 of this book [32] for more information on Usability engineering.). For healthcare information systems researchers have used various techniques for this, including applying knowledge acquisition techniques to elicit expert knowledge of healthcare communication [66], and studying doctor-patient interaction in a natural setting [16].

Once a system designed and implemented, evaluation is crucial. We need to consider what it is that we are trying to claim for the personalised systems, and second, what techniques we can use to demonstrate that our claims or hypotheses are valid.

Some benefits of a personalised system lie in the subjective opinions of users. Perhaps a personalised system is *preferred* by users, and provides information they perceive as more *relevant*. The main way we can assess this is by questionnaire, asking users to rate or compare systems (e.g., [54]), or by simply monitoring actual uptake/use of the systems if freely available.

However, often we will be trying to influence things such as the patient's understanding of their condition, anxiety levels, level of compliance, willingness to take a test, or even their state of health (mental or physical). If a system is developed to affect these things then they naturally need to be measured in the evaluation. In the Piglit project, for example, the aim was for the patient to understand their condition better and so feel more in control. Their preferences and their state of anxiety were thus measured.

The HOMEY project was concerned more with improving the patient's ability to manage their own care, and hence the patient's blood pressure was measured, as a measure indicating good self-management for the relevant condition.

Both example projects used randomised clinical trials to measure the benefits of the system. Patients are randomly assigned to one of two (or more) interventions (e.g., personalised system versus non-personalised) and then appropriate measures are taken relating to the system goals. These measures, as we have seen, can be anything from the objective and concrete blood pressure readings to the patient's perception of information relevance. Differences in results between groups can be measured for statistical significance, and we can attempt to explain these differences as due to the different interventions.

While randomised trials are the *gold standard* for evaluation in the medical domain, we can question their utility from the perspective of the Computer Science researcher. Randomised trials are difficult to design, very expensive and time consuming to run, and, for information systems, interpreting the results is often difficult, mostly because there might be many factors outside the system itself that can affect the results. This may explain why results have often been negative or insignificant, and is not always possible to draw clear conclusions (see [65] for a discussion of negative results). Even for positive results, it may be hard to know where exactly to attribute the differences, as the differences between intervention groups is rarely reducible to one single factor. In Piglit, for example, the total *content* available to patients *and* the starting point for the navigation (medical record) was different. Yet, randomised trials are being used routinely in the medical domain, and researchers have been able to use them to draw conclusions on the benefit of various treatments, or to obtain information on the impact of various information methods (e.g., web *vs* non-web, tailored *vs* generic, etc.) [73, 68].

More such trials are in fact both planned and in progress (e.g., [27, 22]). As a community, adaptive hypermedia researchers may need to learn from medical researchers to be able to better evaluate the effectiveness of their systems. One issue will always remain, though: that of the cost involved and the need for large numbers of users. This is not always practical. At first glance it may appear that, for web based systems, the experimenter has access to a large pool of potential users at low cost – users could be recruited (e.g., by email) and allocated randomly between two or more systems. However, in most cases these users would not be truly representative of the target user group, and it may prove hard to maintain contact with the users over an extended period.

While randomised trials will always be needed in this area, and are necessary when it is the (long term) health of the user we are aiming to improve, alternative methods of evaluation should also be considered. Less expensive and time consuming methods can be used to measure usability, learning and memory, while user preferences can perhaps best be measured both through questionnaires and by looking at actual uptake and use (as in HOMEY). If we can make two different systems available to patients, the one that they choose and continue to use is clearly the preferred one.

This brings us to the question of how we get these systems accepted and used in practice. This is partly, but not just, a question of demonstrating their benefits. There is a huge range of obstacles to changing healthcare practice (see [37] for a brief discussion). Some of this relates to healthcare as a monolithic institution, somewhat resistant to

change, with many stakeholders. However, other obstacles appear when we consider how to change the way information is communicated. Healthcare is an area where trust and privacy are of key importance. Healthcare information must be trusted, come from and validated by reliable sources, and automatically generated adapted information may not meet that criteria, or may not be seen as meeting that criteria. Patient information must be protected and not accessible outside the healthcare team, so even the implicit information available in a seen-over-the-shoulder personalised page of information may result in patient trust being compromised and uptake of a personalised system being reduced. For practical uptake, it is often the apparently trivial issues, like how to find a quiet and private corner of the waiting room for an information point, that prove the hardest to satisfy.

15.7 Conclusions

In this chapter, we have looked at the issue of using adaptive techniques to provide health information and education. We have argued that these techniques show great promise and may open new horizons in this domain, with potential for significant improvements over non-tailored information material. We also described some of the challenges that arise in this domain, in particular issues of trust and privacy, problems of acquiring and constantly updating the user model required to provide sophisticated tailoring, and the cost and difficulty of evaluation. Yet, it remains an exciting domain, one in which, on the one hand, even simple techniques can already bring real benefits and impact, and, on the other hand, new challenges arise. We note, in fact, that there is a growing interest in health applications, in particular for health information systems (e.g., [9, 10, 36, 34, 1]), not only to educate patients but also to assist health professionals (e.g., [28, 53]), to promote better communication both amongst health professionals and between patients and their health care team, and, finally, to provide diagnostic tools and assist in health care provision [21, 23]. Further research is required to assist in these goals and provide systems capable of facilitating this communication and adapting appropriately to the context, at all the required levels.

References

1. Alm, N., Abe, S., Kuwahara, N., eds.: The First International Workshop on Cognitive Prostheses and Assisted Communication (CPAC), Sydney, Australia (January 2006) <http://www.irc.atr.jp/cpac2006/>.
2. André, E., Rist, T.: Generating coherent presentations employing textual and visual material. *Artificial Intelligence Review* **9**(2/3) (1995) 147–165 (Special volume on the Integration of Natural Language and Vision Processing).
3. Auerbach, S., Martelli, M., Mercuri, L.: Anxiety, information, interpersonal impacts, and adjustment to a stressful health care situation. *Journal of Personality And Social Psychology* **44**(6) (June 1983) 1284–96
4. Baker, D., Eash, T. J., S., Uhlmann, W.: Guidelines for writing letters to patients. *Journal of Genetic Counseling* **11**(5) (2002) 399–418

5. Bandura, A.: *Social Foundations of Thought and Action: A Social Cognitive Theory*. Prentice-Hall, Englewood Cliffs, NJ (1986)
6. Bellazzi, R., Montani, S., Riva, A., Stefanelli, M.: Web-based telemedicine systems for home-care: technical issues and experiences. *Computer Methods and Programs in Biomedicine* **64** (2001) 175–187
7. Beveridge, M., Milward, D.: Definition of the high level task specification language. Deliverable D11, HOMEY Project, IST-2001-32434 (2003)
8. Bickmore, T.: *Relational Agents: Effecting Change Through Human-Computer Relationships*. PhD thesis, MIT Media Arts & Sciences (2003) Available at: <http://www.ccs.neu.edu/home/bickmore/bickmore-thesis.pdf>.
9. Bickmore, T., ed.: *AAAI 2004 Fall Symposium on Dialogue Systems for Health Communication*, Crystal City Hyatt, Washington DC (2004) <http://www.aaai.org/Library/Symposia/Fall/fs04-04.php>.
10. Bickmore, T., ed.: *The 2005 AAAI Fall Symposium on "Caring Machines: AI in Eldercare"*, Arlington, Virginia (November 2005) <http://www.aaai.org/Library/Symposia/Fall/fs05-02.php>.
11. Binsted, K., Cawsey, A., Jones, R.: Generating personalised patient information using the medical record. In Barahona, P., Stefanelli, M., Wyatt, J., eds.: *5th Conference on Artificial Intelligence in Medicine Europe (AIME'95)*. Volume 934 of *Lecture Notes in Artificial Intelligence*. (1995) 29–41
12. <http://www.bluepages.anu.edu.au> BluePages, A site providing information about depression.
13. <http://www.thebreastcancersite.com> Breast cancer information site (1).
14. <http://www.breastcancer.org/> Breast cancer information site (2).
15. Brusilovsky, P., Millan, E.: User models for adaptive hypermedia and adaptive educational systems. In Brusilovsky, P., Kobsa, A., Niejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Volume 4321 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin Heidelberg New York (2007)
16. Buchanan, B., Carenini, G., Mittal, V., Moore, J.: Designing computer-based frameworks that facilitate doctor-patient collaboration. *Artificial Intelligence in Medicine* **12** (1995) 171–193
17. Bunt, A., Carenini, G., Conati, C.: Adaptive content presentation for the web. In Brusilovsky, P., Kobsa, A., Niejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Volume 4321 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin Heidelberg New York (2007)
18. Campbell, M., DeVellis, B., Stretcher, V., Ammerman, A., DeVellis, R., Sandler, R.: Improving dietary behavior: the effectiveness of tailored messages in primary care settings. *American Journal of Public Health* **84**(5) (1994) 783–787
19. Carberry, S., de Rosis, F., eds.: *Proceedings of the User Modelling 2005 Workshop on Adapting the Interaction Style to Affective Factors*, Edinburgh, UK (2005) <http://www.di.uniba.it/intint/UM05/WS-UM05.html>.
20. Cawsey, A., Jones, R., Pearson, J.: The evaluation of a personalised health information system for patients with cancer. *User Modeling and User-Adapted Interaction* **10**(1) (2001) 47–72
21. Christensen, H., Griffiths, K.: The prevention of depression using the internet. *Medical Journal of Australia* **177** (suppl) (2002) S122–5
22. Christensen, H., Griffiths, K., Jorm, A.: Delivering interventions for depression by using the internet: randomised controlled trial. *British Medical Journal* (23 January 2004) doi:10.1136/bmj.37945.566632.EE Downloadable from bmj.com.
23. Christensen, H., Griffiths, K., Korten, A.: A web-based cognitive behavior therapy: Analysis of site usage and changes in depression and anxiety scores. *Journal of Medical Internet Research* **4**(1) (2002) e3

24. Colineau, N., Paris, C.: Task-driven information presentation. In: Proceedings of OZCHI'03, Brisbane, Australia (November 2003)
25. Colineau, N., Paris, C., Wu, M.: Actionable information delivery. *Revue d'Intelligence Artificielle (RSTI RIA)* **18**(4) (2004) 549–576 Special Issue on Tailored Information Delivery.
26. De Carolis, B., de Rosis, F., Andreoli, C., Cavallo, V., De Cicco, M.: The dynamic generation of hypertext presentations of medical guidelines. *The New Review of Hypermedia and Multimedia* **4** (1998) 67–88
27. Duszynski, A., Flight, I., Wilson, C., Turnbull, D., Cole, S., Young, G.: Intersecting electronic decision support with user modelling in a web-based consumer decision aid for colorectal cancer screening. [34] <http://www.csc.liv.ac.uk/floriana/UM05-eHealth/>.
28. Elhadad, N., Kan, M., Klavans, J., McKeown, K.: Customization in a unified framework for summarizing medical literature. *Artificial Intelligence in Medicine* **33**(2) (2005) 179–198
29. Feiner, S., McKeown, K.: Automating the generation of coordinated multimedia explanations. In Maybury, M., Wahlster, W., eds.: *Readings in Intelligent Multimedia Interfaces*. Morgan Kaufmann Publishers, Inc., San Francisco, Ca (1998) 89–98
30. Fisher, R.: Patient education and compliance: a pharmacist's perspective. *Patient Education and Counseling* **19**(3) (June 1992) 261–71
31. Gauch, S., Speretta, M., Chandramouli, A., Micarelli, A.: User profiles for personalized information access. In Brusilovsky, P., Kobsa, A., Niejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Volume 4321 of Lecture Notes in Computer Science. Springer-Verlag, Berlin Heidelberg New York (2007)
32. Gena, C., Weibelzahl, S.: Usability engineering for the adaptive web. In Brusilovsky, P., Kobsa, A., Niejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Volume 4321 of Lecture Notes in Computer Science. Springer-Verlag, Berlin Heidelberg New York (2007)
33. Giorgino, T., Quaglini, S., Rognoni, C., Baccheschi, J.: The homey project: a telemedicine service for hypertensive patients. [34] <http://www.csc.liv.ac.uk/floriana/UM05-eHealth/>.
34. Grasso, F., Cawsey, A., Paris, C., Quaglini, S., Wilkinson, R., eds.: Working notes of the UM-2005 workshop on Personalisation for e-Health, Edinburgh, UK (July 2005) <http://www.csc.liv.ac.uk/floriana/UM05-eHealth/>.
35. Green, N.: Affective factors in generation of tailored genomic information. [19] <http://www.di.uniba.it/intint/UM05/WS-UM05.html>.
36. Green, N., Bickmore, T., eds.: The AAAI 2006 Spring Symposium on "Argumentation for Consumers of Healthcare", Stanford University, California (March 2006) <http://www.aaai.org/Library/Symposia/Spring/ss06-01.php>.
37. Grol, R.: Personal paper: Beliefs and evidence in changing clinical practice. *British Medical Journal* **315** (1997) 418–421
38. Hirst, G., DiMarco, C., Hovy, E., Parsons, K.: Authoring and generating health-education documents that are tailored to the needs of the individual patient. In Jameson, A., Paris, C., Tasso, C., eds.: *Proceedings of the Sixth International Conference on User Modeling (UM'97)*, Sardinia, Springer Wien New York (1997) 107–119
39. <http://turing.eng.it/pls/homey/homey.home> The Homey Project website.
40. Jaret, P.: 10 ways to improve patient compliance. *Hippocrates* **15**(2) (2001) http://www.hippocrates.com/FebruaryMarch2001/02features/02feat_compliance.html.
41. Jones, R., Pearson, J., McGregor, S., Cawsey, A., Barret, A., Craig, N., Atkinson, J., Harper Gilmour, W., McEwen, J.: Randomised trial of personalised computer-based information for cancer patients. *British Medical Journal* **319** (1999) 1241–47
42. Kahn, G.: Computer-based patient education: A progress report. *M.D. Computing* **10**(2) (1993) 93–99
43. Kessler, S., ed.: *Genetic Counseling: psychological Dimensions*. Academic Press, NY (1979)

44. Kobsa, A.: Privacy-enhanced web personalization. In Brusilovsky, P., Kobsa, A., Nijedl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Volume 4321 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin Heidelberg New York (2007)
45. Kun, L.: Telehealth and the global health network in the 21st century. from homecare to public health informatics. *Computer Methods and Programs in Biomedicine* **64**(3) (2001) 155–167
46. Large, S., Arnold, K.: Evaluating how users interact with nhs direct online. [34] <http://www.csc.liv.ac.uk/floriana/UM05-eHealth/>.
47. Lewis, D., Eysenbach, G., Kukafka, R., Stavri, P., Jimison, H., eds.: *Consumer Health Informatics - Informing Consumers and Improving Health Care*. Springer-Verlag (2005)
48. Marsden, J.: Primary prevention of osteoporosis in young british women; a comparison of stage-based, tailored email versus non-tailored intervention. In: *Proceedings of Consumer Health Informatics Network for Scotland Conference*. (October 2001) 5 <http://www.phis.org.uk/pdf.pl?file=pdf/jenny%20marsden.pdf>.
49. Marshall, W., Rothenberger, L., Bunnell, S.: The efficacy of personalized audiovisual patient-education materials. *Journal of Family Practice* **19**(5) (November 1984) 659–63
50. McKeown, K.: The TEXT system for natural language generation: An overview. In: *Proceedings of the 20th Annual Meeting of the ACL (ACL'82)*. (1982) 113–120
51. McKeown, K.: Discourse strategies for generating natural-language text. *Artificial Intelligence* **27**(1) (1985) 1–42
52. McKeown, K., Chang, S., Cimino, J., Feiner, S., Friedman, C., Gravano, L., Hatzivassiloglou, V., Johnson, S., Jordan, D., Klavans, J., Kushniruk, A., Patel, V., Teufel, S.: PERSIVAL a system for personalized search and summarization over multimedia healthcare information. In: *Proceedings of the 1st ACM/IEEE Joint Conference on Digital Libraries, JCDL'01, Roanoke, Virginia (June 2001)* 331–340
53. McKeown, K., Pan, S., Shaw, J., Jordan, D., Allen, B.: Language generation for multimedia healthcare briefings. In: *Proceedings of the fifth conference on Applied natural language processing, Washington, DC (1997)* 277–282
54. Mittal, V., Carenini, G., Moore, J.: Generating patient specific explanation in migraine. In: *Proceedings of the 18th Annual Symposium on Computer Applications in Medical Care, Washington DC, McGraw-Hill Inc. (1994)* 5–9
55. Neumark, D.: Providing information about advance directives to patients in ambulatory care and their families. *Oncology Nursing Forum* **21**(4) (1994) 771–5
56. Paris, C., Wan, S., Wilkinson, R., Wu, M.: Generating personal travel guides - and who wants them? In: *Proceedings of the 8th International Conference on User Modeling (UM2001)*. Volume 2109 of *Lecture Notes in Computer Science*., Sonthofen, Germany, Springer-Verlag (July 2001) 251–253
57. Piazza, M., Giorgino, T., Azzini, I., Stefanelli, M., Luo, R.: Cognitive human factors for telemedicine systems. In Fieschi, M., ed.: *MEDINFO 2004*, Amsterdam, IOS Press (2004) 974–978
58. Powell, J., Clarke, A.: The www of the world wide web: Who, what and why? *Journal of Medical Internet Research* **4**(4) (February 2002) doi: 10.2196/jmir.4.1.e4
59. Pratt, W., Sim, I.: Physician's information customizer (PIC): using a shareable user model to filter the medical literature. In: *MEDINFO 1995, Vancouver, B.C., Canada (1995)* 1447–1451
60. Prendinger, H., Mori, J., Ishizuka, M.: Recognizing, modeling and responding to users' affective states. In: *Proceedings of the 10th International Conference on User Modeling (UM'05)*. Volume 3538 of *Lecture Notes in Computer Science*, Edinburgh, UK (2005) 60–69

61. Prochaska, J., Di Clemente, C.: Stages of change in the modification of problem behavior. In Hersen, M., Eisler, R., Miller, P., eds.: *Progress in Behavior Modification*. Volume 28. Sycamore Publishing Company, Sycamore, IL (1992) 183–218
62. Redman, B.: *Advances in Patient Education*. Springer (2004)
63. Reiter, E., Dale, R.: Building applied natural-language generation systems. *Journal of Natural-Language Engineering* **3** (1997) 57–87
64. Reiter, E., Osman, L.: Tailored patient information: some issues and questions. In: *Proceedings of the ACL-1997 Workshop on From Research to Commercial Applications: Making NLP Technology Work in Practice*. (1997) 29–34
65. Reiter, E., Robertson, R., Osman, L.: Lessons from a failure: Generating tailored smoking cessation letters. *Artificial Intelligence* **144** (2003) 41–58
66. Reiter, E., Sripada, S., Robertson, R.: Acquiring correct knowledge for natural language generation. *Journal of Artificial Intelligence Research* **18** (2003) 491–516
67. Schwarzer, R.: Social-cognitive factors in changing health-related behavior. *Current Direction in Psychological Science* **10** (2001) 47–51
68. Skinner, C., Campbell, M., Rimer, B., Curry, S., Prochaska, J.: How effective is tailored print communication? *Annals of Behavioral Medicine* **21** (1999) 290–298
69. Skinner, C., Strecher, V., Hospers, H.: Physicians' recommendation for mammography: do tailored messages make a difference? *American Journal of Public Health* **84**(1) (1994) 43–49
70. Stretcher, V., Kreuter, M., Boer, D., D., Kobrin, S., Hospers, H., Skinner, C.: The effects of computer-tailored smoking cessation messages in family practice settings. *Journal of Family Practice* **39**(3) (September 1994) 267–70
71. Von Wilamowitz-Moellendorff, M., Mueller, C., Jameson, A.: Recognition of time pressure via physiological sensors: is the user motion a help or a hindrance? [19] <http://www.di.uniba.it/intint/UM05/WS-UM05.html>.
72. W3C: Voice extensible markup language (voicexml) version 2.0. W3C Recommendation 16th March 2004 www.w3.org/TR/voicexml20/.
73. Wantland, D., Portillo, C., Holzemer, W., Slaughter, R., McGhee, E.: The effectiveness of web-based vs non-web-based interventions: A meta-analysis of behavioral change outcome. *Journal of Medical Internet Research* **6**(4) (2004) e40