Just-In-Time Context-Sensitive Questioning for Preventative Health Care

Stephen S. Intille **Kent Larson** **Chuck Kukla**

Massachusetts Institute of Technology
1 Cambridge Center, 4FL
Cambridge, MA 02142
intille@mit.edu

Abstract

This paper argues that there is an emerging opportunity for artificial intelligence researchers: the development of technologies that enable a new type of "proactive health care system" that continuously monitors "healthy patients" in their homes and motivates lifelong healthy behavior and health self-awareness. The components of the system, which would augment but not replace traditional health care, could be funded entirely by consumers without reliance on the financially-burdened health care industry. A research prototype of one component of such a system is described: a preventative monitoring system for congestive heart failure that uses "just-in-time," context-sensitive questioning. The system illustrates how sensors in the home used for entertainment and communication purposes might also be employed – at a negligible additional cost – for preventative health care applications.

The problem

The U.S. medical system and those in other countries face an impending crisis: how to pay for the care of an aging population. By 2020, the U.S. Census Bureau predicts that over 30% of the U.S. population will be over the age of 50 (Staff 2000). This growth combined with governmental health care shortfalls and the rising cost of medical procedures will place an enormous strain on the U.S. health care industry in the next 10 years. In 2000 alone, U.S. health care spending rose 7%, the largest increase in more than a decade (Kowalczyk 2002). More alarming, the Congressional Budget Office reports that "financing problems in the near term will be dwarfed by the crisis that could occur as the baby-boom generation reaches age 65" (Antos 1997). The current system provides excellent acute care for disease and injury, but current models of care and payment will not effectively handle the rising number of people needing chronic care.

Although there is a consensus that the present system is not fiscally sustainable, there is no agreement on how to fix it. Meanwhile, costs continue to rise for consumers 8-15% per year, and they feel increasingly frustrated by barriers to access and lack of care options. Corporations as well are strained by rapidly increasing health plan costs.

The health care industry does adopt new technologies but typically to improve existing procedures for acute care such as surgical procedures, not for preventative care. Further, many care givers are not reimbursed for preventative health care services. A transition from the current system of *sick* care that primarily reacts to medical crisis to one of *health* care that focuses on keeping healthy people healthy and early detection of emerging problems would seem to require strong government intervention that refocuses financial incentives on preventative care.

An emerging opportunity

The question we ask is, what can the engineering and design communities do to help alleviate the coming crisis? Looking at consumer spending behavior provides a guide.

Despite their rising health care premiums, individual consumers are actually increasing their voluntary expenditures on health care – only it is *outside* of the traditional medical system. Spending on products and services such as herbal supplements, fitness, dieting, and wellness is soaring. For instance, in 1997 "out-of-pocket monies spent on alternative medicine were comparable to the projected out-of pocket expenditures for all US physician services" (Eisenberg *et al.* 1998). Some of this spending is prompted by studies showing the benefits of exercise and good eating habits on long-term health. However, the growing public distrust of the long-term viability of the medical system is prompting individuals to make personal efforts to ensure that their own health is not being shortchanged.

Therefore, we see an opportunity for technologists: improve the average quality of care by developing personalized and preventative health care solutions based in the home using new technologies that can be purchased by the consumer *without* dependence on funding from the financially-strapped U.S. health care system. In short, we advocate the development of products and services that (1) use computational infrastructure (e.g. sensors) purchased by the consumer for entertainment and communication and not specifically for health care, (2) that lower the cost of providing expert knowledge to the consumer using context-aware information monitoring and data acquisition in the home, and (3) that can be marketed as devices that help the healthy stay healthy. The last point is an important one. Waiting until people get sick to provide care is more costly than early prevention. However, people who perceive themselves as healthy are less likely to invest in products and services that convey a sense of sickness instead of wellness.

Preventative medicine in the home

People spend their healthy days at home, but home is also where they get sick. It's where they feel comfortable, and it's where they communicate with their friends and family. Years ago the medical system was based out of the home. Doctors lived in the communities they served, saw their existing and future patients regularly, and traveled to homes to visit with people and treat them when they got sick. As medicine modernized, the home was relegated to a place that was abruptly left when one got sick for the presumed safety of the hospital and other institutional medical environments.

In January 2002, the authors organized a multi-disciplinary team of MIT researchers and collaborators and developed a proposal for technologies and systems to support personalized health care in the home for the year 2010.[1] Our goal was to generate discussion on how to use technology and design to improve health care *without any reliance on change from the existing health care system*. The "straw man" proposal was motivated by discussions between researchers in medicine, engineering, human factors, user interface design, and architecture and critiqued by participants in the medical system including patients, physicians, nurses, pharmacists, insurers, and administrators.

The system we propose is funded by the users themselves. For the most part, these are people who have not had any major medical issues; we call them "healthy patients" (HPs). HPs include children as well as adults, so that lifelong habits of preventative self-awareness can develop. HPs also include older individuals who may have chronic conditions that require monitoring and home treatment (e.g. medication compliance, glucose measurement, symptom logging).

HPs do not view themselves as sick or as prone to getting sick. They would purchase preventative health technologies primarily to help them feel good, healthy, and secure. It is for these reasons that people are already investing heavily in nutritional supplements, gym memberships, self-help books, dieting plans, and, in some exceptional cases, "boutique" medical care (Belluck 2002).

Many families spend a large percentage of their own income on technology. The applications of AI for preventative health care that we propose leverage off this investment. Our proposal touches on opportunities to impact care in the following ways:

- Preventative monitoring in the home for early detection of onset of medical conditions.
- Proactive human-computer interfaces that use context-aware sensing to motivate health-conscious behavior change over long periods of time, including medication compliance.
- Human-computer interface technologies that allow for new forms of information exchanges between healthy patients and preventative health advising teams.[2]
- Data collection from sensors in the home (e.g. biomonitors) that can aid medical professionals in diagnosis and treatment.
- Information-providing services that help someone navigate the existing health care infrastructure when they do get sick.
- Applications that increase self-awareness of health care in fun and entertaining ways, and applications and services that facilitate the creation of lifelong, personal medical records.

Each of these opportunities could possibly be funded entirely by consumers if they are developed to leverage off of existing (and emerging) technologies for the home. Traditional AI problems such as computational sensing, responsive user interface design, planning, and probabilistic inference will all play a role in reducing the cost barrier to levels most consumers can afford.

An example

To demonstrate how one component of the a future preventative health care system might work, we describe a prototype system for preventative monitoring that we have created in our laboratory. The goal is to monitor the healthy patient by asking the HP questions at the time and place when those questions will provide the most meaningful feedback to an onset diagnosis system and, ultimately, a medical professional.

This is one of several ongoing projects designed to demonstrate how preventative health care services can be created for the home environment. We are actively analyzing prior work on preventative medicine delivered to the home (e.g. via telephones (Friedman 1998)) and exploring how the successful systems can be extended using computational sensing and reasoning. In particular, we are designing algorithms that (1) passively and actively collect data from a user in the home setting without cognitively burdening the user, (2) use that data, aggregated over time, to establish the user's context, and (3) use the context to drive the human-computer interaction of preventative health care system components. Our systems rely on short bursts of interaction with the user. The algorithms must determine when and how (or *the right time and place*) – to acquire and convey preventative medicine information.

Based on interviews with medical professionals, we have developed a prototype system to automatically monitor an individual in the home for congestive heart failure (CHF) using input from sensors in the environment or placed on the person. In the United States alone, over $17 billion is spent annually on the treatment of congestive heart failure, and one-fifth of all hospitalizations for people over age 65 have a primary or secondary diagnosis of heart failure (National Heart, Lung, and Blood Institute 1996). If the onset of CHF

[1] The proposal is available at http://vismod.www.media.mit.edu/~intille/healthmtg/main.htm.

[2] These teams could consist of physicians, but also nurses, therapists, personal trainers, optometrists, dentists, nutritionists, and other professionals who could provide proactive health care advice.

Figure 1: (a) The room where our prototype was constructed and the visual tracking system estimating the position of two people. (b) Answering a question on the "wrist computer" at the table in the prototype room.

is identified prior to heart failure, the condition can often be treated with medication and diet changes, not expensive hospital stays.

Sensing

Our application uses the following data, acquired automatically in real time: (1) position in the home relative to fixed objects such as furniture, (2) weight, and (3) blood pressure (BP).[3] The primary contribution of this work is in the use of position data to infer context that is used to determine when to ask a user health-related questions. Some health conditions (e.g. CHF) cannot be readily identified through current biometric sensing alone. Additional information provided by answers to questions, however, similar to what a good doctor might ask during a check up, could expand the number of problems that could be detected or managed in the home.

We acquire position data using visual sensing (i.e. a camera placed in the ceiling). Position data could also be obtained using an indoor tagging system (e.g. (Priyantha, Chakraborty, & Balakrishnan 2000)). However, the use of visual sensors allows for unencumbering position detection, and the sensors are useful for a variety of other applications, including new forms of entertainment (e.g. see (Bobick *et al.* 1999)).

The fixed camera observes a room from above, as shown in Figure 1a. Real-time, closed-world tracking (Intille, Davis, & Bobick 1997) is used to estimate the position of people relative to fixed objects in the room such as the table, the pull-down bed, the counter, and the door. The fixed-object positions are provided manually to the system. Figure 1a shows two people being tracked, where the centroid of the box indicates the person's estimated position.[4]

The position data is fed into a set of heuristic "activity detectors" that look for events such as "at the table," "sleeping," and "recently sleeping." The sleeping detector, for example, simply outputs a confidence value that the person being tracked is sleeping based on position relative to a bed, time on the bed, and time of day. An active area of our research is how to detect such everyday activities without manual knowledge engineering.[5]

Interface

We assume that the person in the home of the near future will have a small, thin wrist computer with an approximately 2.5 by 2.5 inch curved color touch screen that can be worn comfortably like a watch.[6] We (somewhat awkwardly) approximate the existence of such a device using a PDA with a wireless modem that can be strapped to the arm, as shown in Figure 1b.

Our prototype uses the sensor data to determine an appropriate time to ask the user a simple question about how he or she is feeling. This question is conveyed via the wrist display. In most cases, the wrist device does not actively alert the user. The system displays questions appropriate for the user's context, and the questions will be removed if the user changes contexts and has not responded. If the user happens to look at the wrist display (e.g. to check the time) and sees the question, the user can answer the question with a click of a button. Alternatively, the user can simply tap the screen and the question disappears. Either way, the interruption is minimal. This is a passive interaction model. We assume that this wrist device will be used as a phone, personal planner, watch, and wireless Internet portal. The user, therefore, will attend to the device many times throughout the day. At some of these times, a preventative health message might be displayed that is context-sensitive to a person's activity.

The software is designed to continuously run in the background of the mobile computing device. It must not be perceived as disruptive, otherwise the user will simply turn it off in frustration. Our goal is to get the user in the habit of periodically answering simple questions about how he or she is feeling to the extent that it becomes second nature and the user stops making an effort to second guess the reasons for the questions. The passive interaction model minimizes the risk of annoying interruptions. However, relying entirely on passive inputs offers new challenges for the UI designer: interactions must be fast but can take place at unpredictable times and places. In this prototype, the information provided to the user is simply a text question with a multiple choice response. In ongoing work, we are experimenting with other methods of user interaction as well as designs that are appropriate for elderly users.

[3]Miniature BP biosensors are in development in laboratories. Here we simulate weight and BP data. We assume that several pads are installed in doorways in the environment that periodically take weight measurements. The system we developed could run without this biometric data.

[4]In our current prototype we assume that there is only one individual within the environment, although in future versions the tracker may also use identity information obtained using tags or visual face recognition to allow the system to work for multiple people in the same space.

[5]Some of our current work is focussed on the use of extensions of naive Bayesian networks called tree-augmented networks (Friedman & Goldszmidt 1996). We are not, however, advocating the automatic construction of the diagnosis portion of the networks Bayesian networks described in following sections. These we envision being developed with medical experts and knowledge engineers.

[6]This device might also have an indoor positioning system that would replace the visual sensor used in the prototype system.

Using context

The key innovation of the CHF onset detection system is that it queries the user for preventative health data at an appropriate, meaningful time using computed context information. It works as follows.

A Bayesian network (Pearl 1988) is designed by a knowledge engineer to estimate the probability of CHF onset given evidence of symptoms associated with CHF as it develops in the home.[7] Bayesian networks have proven valuable for medical diagnosis systems (e.g. see (Heckerman & Shortliffe 1992)). The CHF network represents the probabilistic relationships between symptoms such as swelling, change in sleeping patterns, and difficulty breathing that result from CHF's root cause: inefficient heart pumping. The network used by our prototype has less than 20 nodes.[8] The Bayesian network is used to both estimate the probability that the home occupant is beginning to experience CHF given some evidence and to determine what question to ask the occupant in particular contexts that will improve the CHF onset estimate.

Orthopnea – one if the symptoms of CHF – is the condition of having difficulty sleeping while lying down. Some symptoms of orthopnea include general difficulty sleeping, sleeping with extra pillows (to prop oneself up), and falling asleep in chairs rather than beds. Each symptom node in the network, such as $S_{Orthopnea}$, is linked to an evidence node, such as $E_{Orthopnea}$. Evidence is supplied to the network via the evidence nodes, which encode $P(E|S)$. Each evidence node has a function associated with it, E_f that aggregates question data obtained from the user in the previous seven days.

For example, $P(E|S_{Orthopnia})$ encodes the probability that the user will answer a set of context-sensitive questions in a particular way given that the user being monitored has symptoms of Orthopnea. The evidence is acquired directly from the user via a set of context-specific multiple choice questions. In our prototype, each question is worded to yield either positive or negative evidence for a particular symptom. Each symptom has a set of $< Q, < A_P, A_N >, < P >>$ triples associated with it, where Q is a carefully worded question that could be posed to the user, A_P and A_N are the multiple choice answers representing positive support for Q or negative support, and $< P >$ is a set of boolean context predicates, $p_1, p_2, ..., p_n$. For $S_{Orthopnea}$, three of the triples are:

1. $Q =$ "Did you sleep with any extra pillows last night?", $P_P =$ "yes",$P_N =$ "no", $p_1 =$ nextToBed, $p_2 =$ recentlySleeping
2. $Q =$ "Did you sleep in that chair last night?", $P_P =$ "yes", $P_N =$ "no", $p_1 =$ nextToChair, $p_2 =$ recentlySleeping
3. $Q =$ "Did you get up unusually early because of trouble sleeping?", $P_P =$ "yes", $P_N =$ "no", $p_1 =$ nextToBed, $p_2 =$ recentlySleeping, $p_3 =$ earlyMorning

[7] Our network does not actually diagnose CHF; it only identifies warning signs that are likely to lead to a positive diagnosis.

[8] To deploy the system described here would require substantial amount of knowledge engineering with medical professionals in order to create an effective onset diagnosis network. This network might have many more nodes. For demonstration of the idea of "just-in-time" questioning for preventative medicine, our small network is sufficient.

Evidence nodes average 10 context-sensitive questions each in our system; a commercial implementation would have many more potential questions for each type of evidence. These questions would be created by medical professionals based on their experience interviewing patients; however, the types of questions that must be asked are often different than those asked in the clinic where a physician can engage in a dialogue with followup questions. The CHF prevention system questions must stand alone and be worded so that in the given context, P, they rule out other confounding explanations about why the user's behavior has recently changed.

Consider the situation where the user has just woken from an afternoon nap on the pull-out bed in our laboratory space. The user casually looks at the wrist display to check the time and sees Question 1 above. In this situation, only Question 1 of the three questions is valid: all the predicates p_i must match the current context. Each time a question is answered for a particular symptom, E_f aggregates all the responses received for that symptom in the last 7 days and produces certainty value that is entered as virtual evidence (Pearl 1988) in the Bayesian network. In our prototype, E_f simply estimates the support for a symptom computing the ratio of positive to negative responses in the last 7 days weighted by the number of responses.

A response of "yes," therefore, will increase the evidence for $S_{orthopnea}$ which may increase $P(CHFOnset = true|evidence)$, depending upon the structure of the network. If the user does not happen to look at the wrist display, the question simply disappears when the user's context changes.

Each time a new response is received, the evidence is propagated through the network. The software then rank orders the evidence nodes by the impact that changing that evidence would have on $P(CHFOnset|evidence)$. The question predicates are then compared with the current context; those with predicates that do not match are eliminated. Finally, a remaining question is selected from the top of the rank-order list. If this question has not been asked in the last week, it is the question displayed in the current context. Otherwise, the next question remaining in the rank order is selected. Once a question is answered there is a time delay in which no questions will be asked to prevent the user from being bombarded with questions.

Questions appropriate for many contexts are developed for each node by medical professionals. Asking a question at the wrong time can provide meaningless data. For instance, since the feet swell towards the end of the day and after a long walk, a system that tries to obtain evidence for swelling of the feet (a sign of CHF onset) by asking if shoes feel tight after the user has done substantial walking is more likely to acquire misleading evidence of CHF than a system that avoids asking this question at that time. A good time to

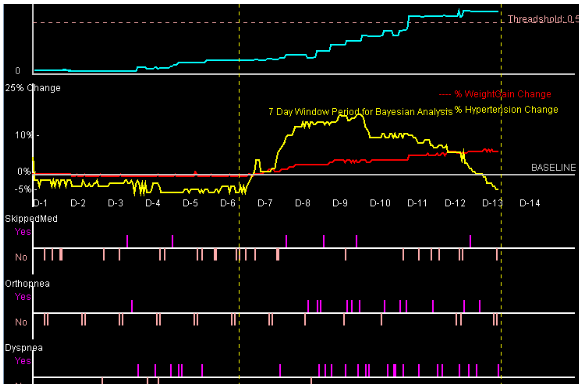

Figure 2: A partial screen shot of the control panel on the computer collecting and analyzing data using the system described in the paper. The top line indicates the system's confidence in CHF onset. The second line shows simulated weekly change in weight and blood pressure. The remaining lines show some of the symptoms that are represented in the Bayesian network and the binary responses received over a period of 13 days (up is positive evidence from a question, down is negative evidence). The dotted lines represent the week-long period of data being considered at one time.

ask this question is right after the person gets up. Although our current system is limited to detecting activities that are inferred entirely from location estimates, future extensions could use more specific activities such as "getting dressed." Future extension might also use less brittle mechanisms for matching question context to detected sensor context.

Figure 2 demonstrates a scenario where data has been entered for a simulated user who is experiencing the onset of CHF. Each small notch indicates that a question was answered for that symptom positively (above) or negatively (below). Our demonstration system starts with simulated data. A user can then enter our laboratory space with a PDA device. The tracking system sends detected contextual states to the diagnosis system that selects a question and sends it via a wireless link to the PDA. If the user answers the question the data is sent back to the diagnosis system, combined with the last week of data, and entered into the network. We estimate that in a real-life setting users would answer approximately 4-8 questions per day. As the user enters more evidence indicating CHF symptoms, the top line shows the P(CHF) eventually surpassing a user-defined threshold. At this time the user and/or a physician might be notified that a visit to a clinic is warranted.

Challenges

This prototype system raises many questions that we leave for future work. First, how should questions be worded so that they elicit the intended information while appearing harmless and not provoking user concern? Second, will people find the passive interface model acceptable or will they still get annoyed and disable the device? Third, what activities and contextual states need to be recognized to identify early onset of common medical conditions, and can these activities be identified using technology that can be easily retrofitted into the existing homes without jeopardizing the user's sense of privacy? Fourth, what are acceptable false positive and false negative rates for such a system if it were to be widely deployed? Finally, would any organization risk creating such a monitoring system if it could be held liable for any false negatives?

Conclusions

We describe a prototype of a context-aware system for preventative medical monitoring in the home using computational sensing and Bayesian reasoning to illustrate one example of the potential for applications of AI-related research to improve preventative health care for elderly. We are currently redesigning this system so that it can be easily installed in environments with standard 8 foot ceilings and so that it can use activity recognition sensors trained via supervised learning. This preventative monitoring tool is just one of a complimentary set of consumer-funded tools we foresee that could augment the traditional health care system using AI technologies and enable new forms of proactive health care.

Acknowledgments

We would like to acknowledge Joseph Su, who implemented the congestive heart failure demonstration system and Dan Carlin, Verlé Harrop, and Jennifer Beaudin for conversations that influenced the ideas presented here.

References

Antos, J. R. 1997. CBO testimony: 2010 and beyond: Preparing Medicare for the Baby Boomers. U.S. Congressional Budget Office Testimony.

Belluck, P. 2002. Doctors' new practices offer deluxe service for deluxe fee. *The New York Times*.

Bobick, A.; Intille, S.; Davis, J.; Baird, F.; Campbell, L.; Ivanov, Y.; Pinhanez, C.; Schütte, A.; and Wilson, A. 1999. The KidsRoom: A perceptually-based interactive and immersive story environment. *PRESENCE: Teleoperators and Virtual Environments* 8(4):367–391.

Eisenberg, D.; Davis, R.; Ettner, S.; Appel, S.; Wilkey, S.; Van Rompay, M.; and Kessler, R. 1998. Trends in alternative medicine use in the United States, 1990-1997. *Journal of the American Medical Association* 280(18):1569–1575.

Friedman, N., and Goldszmidt, M. 1996. Building classifers using Bayesian networks. In *Proc. Nat. Conf. on Artificial Intelligence*, 1277–1284. AAAI Press.

Friedman, R. 1998. Automated telephone conversations to assess health behavior and deliver behavioral interventions. *Journal of Medical Systems* 22(2):95–102.

Heckerman, D., and Shortliffe, E. 1992. From certainty factors to belief networks. *Artificial Intelligence in Medicine* 4:35–52.

Intille, S.; Davis, J.; and Bobick, A. 1997. Real-time closed-world tracking. In *Proc. of the IEEE Conf. on*

Computer Vision and Pattern Recognition, 697–703. Los Alamitos, CA: IEEE Computer Society Press.

Kowalczyk, L. 2002. Spending on health care rises 7 percent. *Boston Globe* Section D1.

National Heart, Lung, and Blood Institute. 1996. Data fact sheet: Congestive heart failure in the United States: A new epidemic.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems : Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann Publishers.

Priyantha, N.; Chakraborty, A.; and Balakrishnan, H. 2000. The cricket location-support system. In *Proc. of the Sixth Annual ACM International Conference on Mobile Computing and Networking (MOBICOM)*.

Staff, S. I. 2000. Projections of the total resident population by 5-year age groups, and sex with special age categories: Middle series, 2016 to 2020. Technical report, U.S. Census Bureau, Population Division.