Health-Related Applications of Socially Interactive Agents

Timothy Bickmore

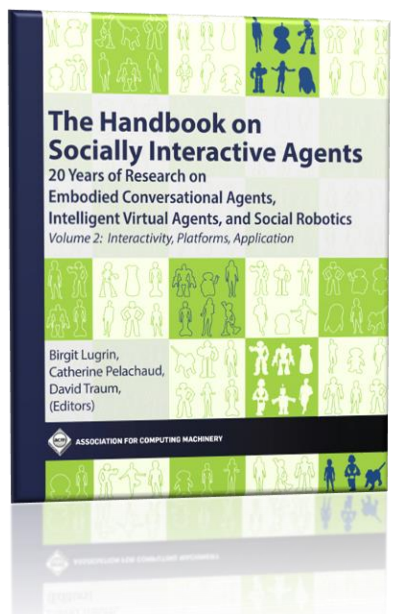

Author note:

This is a preprint. The final article is published in "The Handbook on Socially Interactive Agents" by ACM.

Citation information:

T. Bickmore (2022). Health-Related Applications of Socially Interactive Agents . In B. Lugrin, C. Pelachaud, D. Traum (Eds.), *The Handbook on Socially Interactive Agents – 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics*, Volume 2: Interactivity, Platforms, Application (pp. 403-435). ACM.

DOI of the final chapter: 10.1145/3563659.3563672

DOI of volume 2 of the handbook: 10.1145/3563659

Correspondence concerning this article should be addressed to Timothy Bickmore, t.bickmore@northeastern.edu

24 Health-Related Applications of Socially-Interactive Agents

Timothy Bickmore

24.1 Motivation for Socially-Interactive Agents in Healthcare

Health is the prerequisite for any human endeavor, necessarily occupying the base of Maslow's Hierarchy of universal human needs [Maslow 1943]. The maintenance of health is thus a significant preoccupation in all societies, spawning a vast and complex industry, spanning a wide range of organizations, services, and products, starting with front-line caregivers and the myriad support services and products they, and their patients, require. Around 10% of the world's resources are devoted to healthcare, with US$7.8 trillion spent worldwide in 2017, and growth in healthcare spending is outpacing the overall growth in the world economy [WHO 2019].

Given its importance and cost, healthcare represents a significant opportunity for automation in general and SIAs in particular. Developing countries have severe shortages in resources that may only be met through relatively low-cost automation. However, the most advanced countries can also benefit, since even small improvements in health outcomes or cost reduction can yield large returns when scaled to a population.

SIAs provide unique affordances in automating aspects of healthcare delivery, since they can be used to most directly automate the care that an expert human health provider would give, especially when the care is in the form of information or guidance delivered in a face-to-face clinical encounter. SIAs can not only provide the instrumental, utilitarian delivery of medical information to patients, but also the social, emotional, and relational messages—such as compassion and empathy—that more conventional digital media lacks. Further, they can provide these messages with perfect consistency and fidelity every time, to every patient, in every circumstance. SIA are never intended to replace human healthcare providers, but rather to automate routine parts of their work so they can focus their time on exceptional cases or see more patients. SIA can also be used in ways that humans are not capable of, such as providing counseling to every patient at any time of the day, or counseling to patients who cannot see human providers due to cost or other factors.

In this chapter, I describe the significant opportunities for SIA in healthcare, along with several examples of agents and robots that have been developed and evaluated to date. I focus on SIAs that are designed to interact directly with patients (receiving care for a health condition) and consumers (focused on preventing illness and maintaining wellness), providing health education and the promotion of health-related behavior. I do not cover SIAs designed to interact with healthcare providers (e.g., as virtual patients for training [Ferdig et al. 2012]), although there are several researchers investigating these. My focus will also be on SIAs that have been developed and evaluated in randomized clinical trials with real patients and consumers, although I will also mention a few relevant SIA that have either not made it into trial or have transitioned into commercial products.

24.1.1 The Importance of Health Behavior

Modifiable health behavior, such as physical inactivity, poor diet, and smoking, accounts for nearly 40% of all deaths in the US [Maslow 1943]. Although patient and consumer behavior can significantly impact health outcomes in almost every area of medicine, I will focus here on a few select areas that have the potential for very significant impact, given that they can be addressed through automated health behavior change interventions.

24.1.1.1 The Burden of Overweight and Obesity

The rates of overweight (body mass index greater or equal to 25 [Arroyo-Johnson and Mincey 2016]) and obesity (body mass index greater or equal to 30) have exploded across the world in the last few decades. Worldwide, 39% of adults were overweight and 13% were obese, and an additional 340 million children were overweight or obese, in 2016, with these conditions linked to significant increases in rates of cardiovascular disease, diabetes, and other conditions, leading to increased rates of morbidity and mortality [WHO 2020]. Addressing overweight and obesity primarily involves diet modifications and increases in physical activity, both lifestyle health behaviors that have been the target of numerous automated interventions, including several with SIAs [Bickmore et al. 2005a, 2005b, 2013a, 2013b, King et al. 2013].

24.1.1.2 The Burden of Chronic Disease

Chronic conditions, such as diabetes, hypertension, and atrial fibrillation, affect half of all adults in the US, are responsible for 70% of all deaths in the US, and account for 75% of US healthcare expenditures [Wullianallur amd Raghupathi 2018]. With the aging population, the prevalence of chronic conditions—especially multiple chronic conditions—is continuing to rapidly increase. These conditions can require challenging self-care management regimens, spanning medication adherence, lifestyle modifications, and vigilant symptom monitoring, in addition to regular visits with healthcare providers to prevent disease progression. However, 60% of individuals with chronic conditions are poorly adherent to their prescribed treatment regimens, with even higher rates for older adults or those with limited health literacy [Dunbar-Jacob and Mortimer-Stephens 2001]. Promoting treatment adherence has been the target of several automated interventions, including by SIAs [Bickmore et al. 2010].

24.1.1.3 The Burden of Substance Abuse

Over 150 million people worldwide are estimated to have alcohol or drug use disorders, leading to over 3 millions deaths per year [GBD 2016 Alcohol and Drug Use Collavorators 2018]. Counseling, using techniques such as Motivational Interviewing (Section 24.2.3) and Cognitive Behavior Therapy, have been shown to be effective at reducing alcohol and substance use in most populations and settings, including single session "brief interventions" in primary care [DiClemente et al. 2017]. Many of these counseling techniques are automatable, and SIAs provide additional affordances for counseling over more traditional media. SIAs have successfully been used in several interventions for substance use ([Yasavur et al. 2014, Olafsson et al. 2020], Section 24.4.3).

24.1.2 Logistical Factors: Convenience and Cost

There are many reasons why SIAs are well-suited to health behavior interventions. As with all types of automation, cost savings are often first to mind. However, SIA can provide emulations of human health counselors that are available 24x7 and, in the case of those deployed on mobile devices (Section 24.4.2), wherever the user happens to be. Notions of convenience also extend beyond mere availability, to patient comfort in taking the time they need to get the information they require. For example after interacting with an IVA-based virtual hospital discharge nurse (Figure 1), 36% of 149 patients said they preferred receiving their discharge instructions from the IVA rather than their doctors or nurses in the hospital, compared to only 24% who said they would prefer receiving it from their human providers (the remaining 39% had no preference) [Zhou et al. 2014]. When asked why they preferred the agent, participants indicated that they did not feel rushed with the agent the way they felt with the doctors or nurses in the hospital ("I prefer Louise, she's better than a doctor, she explains more, and doctors are always in a hurry.").

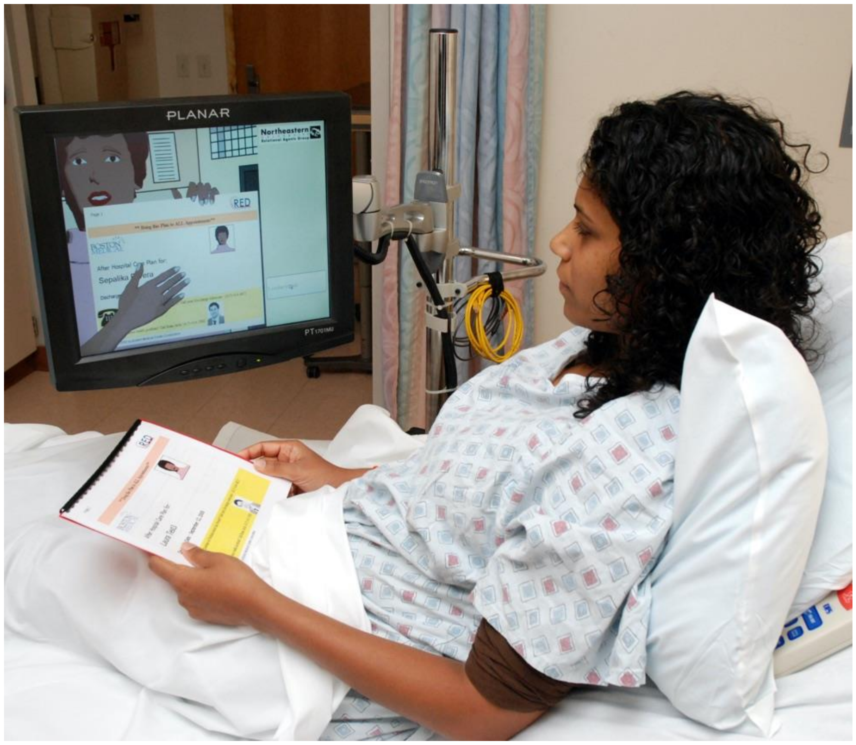

Figure 1. IVA Hospital Discharge Nurse

24.1.3 Social-Emotional Factors: Stigma, Trust, and Alliance

There are other aspects of healthcare for which SIA can outperform human providers. Many areas of medicine require patients to self-report behaviors that are potentially stigmatizing, such as substance use or sexual behavior. Patients may also be reluctant to report failures to adhere to recommended or prescribed regimens spanning diet, physical activity, self-care, or medications. Avoidance or underreporting of such behavior to human clinicians is well-documented, and there have been studies dating back over 40 years demonstrating that patients report potentially-stigmatizing conditions more accurately to a computer than to a human [Card and Lucas 1981]. More recent studies have demonstrated that this effect carries over to SIA: in a study comparing self-reports of intimate information to an IVA, participants revealed more when they were told they were interacting with a fully autonomous system compared to those who were told they were interacting with a remote-controlled avatar [Lucas et al. 2014].

24.1.4 Literacy Factors: Health, Computer, and Reading Literacy

Health literacy—the ability to perform the basic reading and numerical tasks required to function in the health care environment [National Institutes of Health Program Announcement 2007]—affects patients' ability to understand medication labels and instructions, hospital discharge instructions, instructions for assistive devices and medical equipment, and educational material [AMA 1999]. Approximately 36% of US adults have inadequate health literacy [Kutner et al. 2006]. This problem is not uniformly distributed in society; among indigent and minority patients in urban areas this number rises to over 80% [Williams et al. 1995a]. Patients with inadequate health literacy report lower health status [Weiss et al. 1992, Baker et al. 1997], are less likely to use screening procedures, follow medical regimens, keep appointments, or seek help early in the course of a disease [Weiss et al. 1994], have greater difficulties naming their medications and describing their indications [Williams et al. 1995b], more frequently hold health beliefs that interfere with adherence [Kalichman et al. 1999], have higher health-care costs [Weiss et al. 1994], and have higher rates of hospitalization [Baker et al. 1997].

Evidence suggests that a face-to-face encounter with a health provider—in conjunction with written instructions—remains one of the best methods for communicating health information to patients in general, but especially those with low literacy levels [Colcher and Bass 1972, Madden 1973, Clinite and Kabat 1976, Morris and Halperin 1979, Qualls et al. 2002]. Face-to-face consultation is effective because it requires that the provider focus on the most salient information to be conveyed [Qualls et al. 2002] and that the information be delivered in a simple, conversational speaking style. Protocols for "grounding" in face-to-face conversation—the use of verbal and nonverbal cues such as head-nods, gaze and acknowledgement tokens ("uh-huh", "OK") to signal mutual understanding [Clark and Brennan 1991]—allows providers to dynamically assess a patient's level of understanding and repeat or elaborate information as necessary. Face-to-face conversation also allows providers to make their communication more explicitly interactive by asking patients to do, write, say, or show something that demonstrates their understanding [Doak et al. 1996]. Finally, face-to-face interaction allows providers to use verbal and nonverbal behaviors, such as empathy [Frankel 1995] and immediacy [Argyle 1988, Richmond and McCroskey 1995], to elicit patient trust, enabling better communication and satisfaction.

Healthcare SIAs may prove particularly effective for individuals with low health literacy, given their ability to emulate face-to-face consultation. This hypothesis has now been supported in several studies and clinical trials. Patients with low health literacy reported significantly higher levels of trust in the virtual discharge nurse IVA (Figure 1) compared to patients with adequate health literacy. Patients with low computer literacy were also significantly more satisfied with the virtual nurse and rated it significantly higher on ease of use compared to those with high computer literacy [Zhou et al. 2014].

Studies have also demonstrated that low literacy individuals can complete some health-related tasks more successfully with SIAs compared to more traditional computer interfaces. For example, in an evaluation of an IVA that helped cancer patients find clinical trials to volunteer for, low health literacy patients completed significantly more correct search tasks and were significantly more satisfied with the SIA compared to a functionally-equivalent facet- and keyword-based search engine [Bickmore et al. 2016]. Similarly, among 273 participants recruited from an urban safety net hospital, 74% of whom had possibly or likely low health literacy, significantly more were able to complete a family health history with an IVA that simulated a genetics counselor (97%) [Wang et al. 2015], compared to those using a functionally-equivalent web form-based interface (51%) [Wang et al. 2017].

24.2 Models and Theoretical Frameworks

Many theories of health communication and health behavior change have been developed over the last several decades that can be used to inform the design of SIAs intended to educate or change the health behavior of their users. Health behaviors span a wide range of physical actions, contexts, and time durations, from the relatively simple action of showing up at an appointment to obtain a one-time vaccination, to chronic disease self-care management that may require years of consistent, concerted effort across a wide range of activities spanning diet, exercise, medications, and self checks. Similarly, interventions to change health behavior range from single brief educational messages, to years of counseling and coaching sessions. SIAs have been used for all of these types of behaviors and in all of these intervention formats. Here, I focus on a few particularly relevant theories and frameworks that have been used in SIA-based health interventions.

24.2.1 Health Communication and Education

Communication of health information forms the cornerstone of any intervention. Unless someone knows what they are supposed to do and how to do it, clearly understands the reasons for doing it and the risks of not doing it, they will likely not engage in a desired health behavior. SIAs that are developed primarily for educational purposes have been referred to as "pedagogical agents" and are reviewed in depth in Chapter 21 on "Pedagogical Agents" [Lane and Schroeder 2022] of this volume of this handbook. Here

I focus briefly on communication and pedagogy that is unique to health. An important finding to keep in mind, however, is that it is widely accepted that education is necessary but not sufficient to achieve behavior change [Nichols 1994]. Thus, communication and pedagogical strategies must be used in conjunction with other counseling techniques to succeed in short- or long-term behavior change.

The literature on health communication is vast, from studies of one-on-one doctor-patient communications to population-wide public health campaigns [Schiavo 2013]. One area of health communication research that is particularly relevant to SIA is computerized tailoring of health messages. One proposed typology describes a spectrum of health communication messages, ranging from those that are invariant across all recipients ("generic communication"), to those tailored only on personal characteristics such as the recipient's name ("personalized generic communication"), to those designed for different segments of the population ("targeted communication") [Kreuter et al. 1999]. Tailored communication is at the extreme end of this spectrum, as has been defined as "any combination of strategies and information intended to reach one specific person, based on characteristics that are unique to that person, related to the outcome of interest, and derived from an individual assessment" [Kreuter et al. 2000]. For example, the user's sex, age, and education may be input and health promotion text adjusted accordingly, typically using relatively simple template-based text generation techniques [Reiter and Dale 2000, Reiter et al. 2003], with all possible text variants manually pre-authored. Importantly, most interventions tailored on one or more factors derived from a health behavior change theory, such as the "stage of change" from the Transtheoretical Model (Section 24.2.3). Work on computerized tailoring was an acknowledgment that the "one size fits all" approaches to public health communication were not as effective as they could be. Tailoring is hypothesized to work through several mechanisms, including increased attention to tailored health messages, more effort spent processing tailored messages, and increased self-referential thinking [Hawkins et al. 2008]. A large meta-analysis of 57 computerized tailoring studies found that their mean effect size on health behavior change was r=.074 [Noar et al. 2007]. Importantly, this study also found that there was a significant relationship between the number of factors tailored on (ranging from 0 to 9 theoretical constructs) and the impact on health behavior change.

SIAs for health behavior change have been developed with varying degrees of tailored messaging. However, given their ability to engage users in natural language dialogue, with each SIA utterance generated in real-time via a dialogue system and natural language generator, they have the ability to tailor their message to each turn of conversation based not only on theoretical and demographic factors known about the user but on the discourse context, including what was just said by the user or what they said in prior conversations. This gives them the capability to perform tailoring at a level of granularity never before possible.

24.2.2 Persuasive Technology

Moving beyond communication and education, SIAs can use a variety of techniques to persuade a user to perform a health-related action. Technologies designed to motivate a single action have been referred to as persuasive technologies [Fogg 2003]. Decades of psychology studies have provided a list of persuasive techniques that could be automated. For example, Cialdini [2001] outlines six strategies of persuasion—including authority, liking, and reciprocity—which are widely used in marketing communication strategies. Most of these techniques have been used in automated systems, including SIAs. Examples SIA studies include reciprocity, in which the SIA does a favor for the user first to gain compliance [Lee and Liang 2016], reciprocal deepening self-disclosure to build trust before making a request of the user [Moon 1998], and physical touch to increase persuasion ("Midas touch" [Haans and Ijsselsteijn 2009]).

24.2.3 Health Behavior Change

While immediate compliance can be important in brief interventions for single-action health behaviors, a more longitudinal, incremental, and relational approach is required to change ingrained habits over longer periods of time. Several theoretical models and techniques have been developed in the field of behavioral medicine to address these more substantive problems [Glanz et al. 2008]. A sampling of theories includes:

- The Health Belief Model [Becker 1976, Janz and Becker 1984], which focuses on increasing the user's perceived threat of a disease or negative health outcome to motivate behavior change, with perceived threat comprised of susceptibility (the perception of risk for the negative outcome), and severity (the perception of the seriousness of that outcome). Thus, the emphasis on these interventions is on communication of risk, helping the user weigh barriers and benefits to change, and on increasing user self-efficacy (confidence to change).
- The Theory of Reasoned Action [Fishbein and Ajzen 1975] and the Theory of Planned Behavior [Ajzen and Madden 1986], which focus on the perceived likelihood of performing a behavior. The Theory of Reasoned Action posits that intent is driven by positive attitude towards the behavior and subjective norms (perceived societal pressure). The Theory of Planned Behavior posits an additional driver, perceived behavioral control, which is the extent to which the user believes the behavior is under their control. Thus, the focus of these interventions is on increasing positive attitude, perceived normative pressure, and perceived behavioral control.
- Social Cognitive Theory [Bandura 1986], which posits that the most central drivers of health behavior change are self-efficacy and the perception that engaging in the behavior will lead to positive outcomes.

However, the health behavior change theory that has arguably received the most empirical support and is the most readily-automatable is the Transtheoretical Model (TTM), also known as the "stages of change" model [Prochaska and Velicer 1997]. This model posits that people go through a series of stages when they change a behavior, and the individual behavior change techniques and messages that are most effective at a given time depend on what stage a user is in. The five Stages of Change describe different levels of readiness to change [Velicer et al. 2000], and span Precontemplation, Contemplation, Preparation, Action, and Maintenance. Precontemplators are not intending to change in the next 6 months. Contemplators are intending to change in the next 6 months. People in the Preparation stage are planning to change in the next 30 days. People in the Action and Maintenance stages have changed, with those in Action having changed within the prior 6 months. The model also classifies behavior change techniques into a set of 10 Processes of Change, and posits additional constructs including Decisional Balance (weighing Pros vs. Cons of change) and self-efficacy. The TTM is relatively straightforward to automate, since a user's Stage can be assessed through one or two brief questions, then used to index a library of behavior change techniques or messages. The TTM has been used to drive the behavior of SIAs designed for a variety of health behaviors [Velicer et al. 2009, Bickmore et al. 2013b, Jack et al. 2015].

Users in Precontemplation pose a particular challenge, since they are not motivated to even engage in an intervention at all. Motivational Interviewing provides a set of counseling techniques specifically focused on boosting motivation to change and thus moving users out of Precontemplation and into later stages in which they are willing to take some action towards change [Miller and Rollnick 2012]. Motivational Interviewing involves prompting users to talk through their perceived pros and cons of change to help them resolve their ambivalence in favor of taking action. Some aspects of Motivational Interviewing have been implemented in several SIA-based interventions [Schulman et al. 2011], although much of it requires eliciting and responding to unconstrained user utterances, which represents a currently-insurmountable challenge for natural language understanding (see Chapter 5 on "Natural Language Understanding in Socially Interactive Agents" [Pieraccini 2021] of volume 1 of this handbook [Lugrin et al. 2021] and the discussion of safety issues with doing this in healthcare in the Current Challenges Section of this chapter).

While the theories described here provide overarching frameworks to guide intervention design, they do not, on their own, dictate specific actions that a SIA should take at a given time with a given user. At best, they provide a description of the types of techniques (e.g., TTM Processes of Change) or type of outcome (e.g., increased risk perception for the Health Belief Model) that is desirable at a given time.

These specifications provide very general indices into a large set of individual behavior change techniques that have been developed and evaluated over the years in behavioral medicine. Going beyond the 10 Processes of Change categories, Michie has spent several years developing taxonomies of behavior change techniques [Michie et al. 2013]. Her taxonomy has 16 top-level categories organizing a total of 93 individual techniques. An excerpt of the taxonomy is shown in Table 1.

Table 1. Excerpt of Michie's Taxonomy of Behavior Change Techniques (based on [69])

Category	**Example Techniques**
Scheduled Consequences	Shaping (modifying a behavior incrementally, such as gradually weaning an undesirable behavior, or increasing a desirable behavior in small steps). Example use in SIAs: [6]
Reward & Threat	Social Reward (provision of congratulatory messages by another person). Example use in SIAs: [70]
Goals & Planning	Problem Solving (identifying barriers to change and assisting user in addressing them). Example use in SIAs: [13]
Social Support	Emotional Support (e.g., display of empathy when user is having difficulty). Example use in SIAs: [70]

24.2.4 Trust & Therapeutic Alliance

There is a strong correlation between the quality of professional-client relationships and outcomes across a wide range of human helping professions. The dimension of the clinician-patient relationship that is credited with the significant influence on outcome—the *Working Alliance*—is based on the trust and belief that the therapist and patient have in each other as team-members in achieving a desired outcome, and has been hypothesized to be the single common factor underlying the therapeutic benefit of therapies ranging from behavioral and cognitive therapies to psychodynamic therapy [Gelso and Hayes 1998]. The Working Alliance construct has been hypothesized to have three sub-components: a goal component, reflecting the degree to which the clinician and client agree on the goals of the therapy; a task component, reflecting the degree to which the clinician and client agree on the therapeutic tasks to be performed; and a bond component, reflecting the trusting, empathetic relationship between the client and clinician [Horvath and Greenberg 1989, Gelso and Hayes 1998].

In the context of longitudinal health behavior change interventions, Working Alliance can positively impact both *compliance*, the likelihood that a user will follow an SIA's immediate recommendations, and *retention*, the likelihood that the user will continue working with the SIA over time. There is some evidence that strategies that focus on near-term compliance do so at the cost of retention [Bickmore et al. 2007], and it may be that retention is the primary beneficiary of a strong Working Alliance [Bickmore and Picard 2004]. Retention is arguably the more important objective and yet the more difficult goal. The majority of longitudinal health behavior change interventions suffer from attrition over time, when users give up for one reason or another.

SIAs are ideal for establishing Working Alliance relationships with users, and have many affordances lacking in more traditional print or digital media for this purpose. Their use of nonverbal communicative behavior (e.g., facial display of empathy, Figure 2), in conjunction with *immediacy* behaviors, which demonstrate attention to and engagement in the user by the SIA, along with dialogue designed to build and maintain a social-emotional relationship with the user, make SIAs ideal platforms for establishing strong, trusting relationships with users, and leveraging these in health behavior change interventions [Bickmore and Picard 2005].

Figure 2. Empathetic Exercise Coach IVA

24.3 A Brief History of SIAs in Healthcare

SIAs, broadly construed, have a long history of being used for health education and health behavior change interventions. Disembodied dialogue systems were used as far back as 1966 for psychotherapy. The ELIZA system was developed to simulate the behavior of a Rogerian psychotherapist, in which the patient and the computer exchanged typed text messages [Weizenbaum 1966]. Although ELIZA was not intended to be used for actual therapy, similar systems have been proven effective for therapy in which the system is essentially prompting a patient to think aloud and work through his or her own problems [Slack 2000]. Colby [1995] developed an ELIZA-like system that was designed to use Cognitive Behavioral Therapy to treat individuals with depression. In addition to providing typed text counseling with patients, the system provided text-based educational materials about depression.

Speech-based health counseling systems represent the next step closer to embodied SIAs. Several health interventions have been developed for use over the telephone, referred to as Interactive Voice Response (IVR) systems. IVR-based interventions have been developed and evaluated in clinical trials for several aspects of diet, physical activity, cigarette smoking, medication adherence, office visit adherence, disease screening behavior, and chronic disease management for hypertension, angina pectoris, chronic obstructive lung disease, asthma, diabetes mellitus, and depression [Migneault et al. 2006].

Full IVA-based health interventions began appearing in the early 2000's, with animated conversational health coaches designed for longitudinal health behavior change interventions in physical activity promotion [Bickmore and Picard 2005, Bickmore et al. 2005b], medication adherence promotion [Bickmore et al. 2010], hospital discharge patient education [Bickmore et al. 2009], and UV avoidance [Velicer et al. 2009]. In these systems, the agent used its embodiment to regulate the flow of conversation (e.g., gaze cues for turn-taking), emphasize key points (e.g., via eyebrow raises, beat gestures), convey additional meaning (e.g., deictic gestures at parts of medical documents in the agent's virtual environment, as in Figure 1), mark topic boundaries (e.g., via posture shift), mark conversational frames using contextualization cues (e.g., proxemics, gaze frequency, facial display, gesture frequency to delineate social chat from task talk), and display empathy (e.g., using facial display of concern as in Figure 2, proxemics, and shifts in prosody) [Bickmore and Picard 2005].

The use of SRs for health interventions represents one of the most recent developments. Kidd developed the first SR for use in a longitudinal health intervention. His robot consisted of a touch screen computer with a physical robotic head attached, designed to assist users with a weight-loss program [Kidd and Breazeal 2008]. Participants interacting with the SR continued with the weight-loss program for twice as long compared to a group that used a touch screen-only device, although there were no differences in weight loss. More recently, SRs have been used in a wide range of health applications, from rehabilitation [Mataric et al. 2009] to distraction during medical procedures [Trost et al. 2020]. SRs have also been widely used for individuals with Autism, and to provide many kinds of support to older adults, both of which are covered in other chapters in this volume (Chapters 23 on "Socially Interactive Agents for Supporting Aging" [Ghafurian et al. 2022] and 25 on "Autism and Socially Interactive Agents for Supporting Aging" [Nadel et al. 2022]). Riek [2017] provides an excellent overview of the potential for SRs in healthcare, describing their future use in long-term eldercare, inpatient and outpatient care, and psychiatric and palliative care.

There have also been many patient- and consumer-facing commercial applications of SIA in healthcare. This is a rapidly-evolving area with companies and offerings coming and going frequently. A current snapshot includes several disembodied chatbots to provide counseling and therapy, such as the woebot depression counselor (successfully evaluated in a clinical trial [Fitzpatrick et al. 2017]). True SIAs include Molly, an IVA from Sensely that provides a general health counselor agent on smartphones for a wide variety of conditions, including chronic disease management (no clinical trials reported to date). The Autom SR was launched as a home weight loss coach by Intuitive Automata, following successful evaluation at the MIT Media lab [Kidd and Breazeal 2008], and a similar SR was subsequently developed by Catalia Health for wellness coaching and chronic disease management (neither of these commercial robots have been evaluated in clinical trials to date). Finally, Care.Coach (previously GeriJoy) provides an IVA companion on a tablet computer for eldercare, but is driven by a remote operator using "wizard of oz" technology. There are several publications covering development and usability testing of the system, but no clinical trial evaluation to date.

24.4 Example SIA Systems in Healthcare

24.4.1 SIA to Address Sedentary Behavior

Physical activity is one of the few health behaviors that has significant benefits for individuals of all ages, including older adults [Chodzko-Zajko et al. 2009]. Late-life exercise improves strength, aerobic capacity, flexibility, and physical function [Keysor and Jette 2001], and a change from a sedentary to more active lifestyle in midlife or beyond is associated with a reduction in mortality [Eriksson et al. 1998, Bijnen et al. 1999]. Despite these benefits, only about 25% of men and 20% of women aged 65 years and older in the US meet the national guidelines for regular physical activity [Nelson et al. 2007].

To address this need, an IVA-based virtual coach was developed to increase physical activity among geriatrics patients at an urban safety net hospital in the US. Patients were provided with the IVA on a touch screen tablet computer, along with a pedometer, to take home for two months. They had a daily counseling conversation with the virtual coach, based on steps uploaded from the pedometer, using a variety of health behavior change techniques. The virtual coach was evaluated in a randomized clinical trial, compared to a control group who were provided with pedometers and print materials on the benefits of exercise. A total of 263 older adults were enrolled, with a mean age of 71.3, of whom 61% were female and 51% had an educational attainment of a high school diploma or less. Participants in the virtual coach group increased their daily step count significantly more compared to the control group at the end of the two months, although this effect waned by a 12-month follow up. The older adults who used the IVA were highly satisfied with the program, scoring it well above the midpoint on a standardized Likert scale measure of therapeutic alliance, and interacted with it an average of 30 out of the 60 days of the intervention [Bickmore et al. 2013b].

24.4.2 SIA to Address Chronic Disease

As mentioned in Section 24.1.1.2, chronic conditions, such as diabetes, hypertension, and atrial fibrillation, are responsible for the majority of healthcare expenditures and deaths worldwide. Management of chronic conditions typically require lifestyle health behavior change over the remaining years of an individual's life. Smartphones provide an ideal platform for systems designed to help individuals manage these conditions, given their computational power and connectivity, their "anywhere, anytime" availability, and especially when coupled with mobile sensors or monitors. To explore this, a prototype IVA was developed and deployed on smartphones to help support individuals with atrial fibrillation (AF) (Figure 3). AF is an irregular heartbeat that increases the risk of stroke 3- to 5-fold, and doubles the risk of death if untreated [Romero et al. 2014]. The IVA provided education about AF, promoted medication adherence and symptom monitoring, and promoted use of a heart rhythm sensor attached to patients' smartphones.

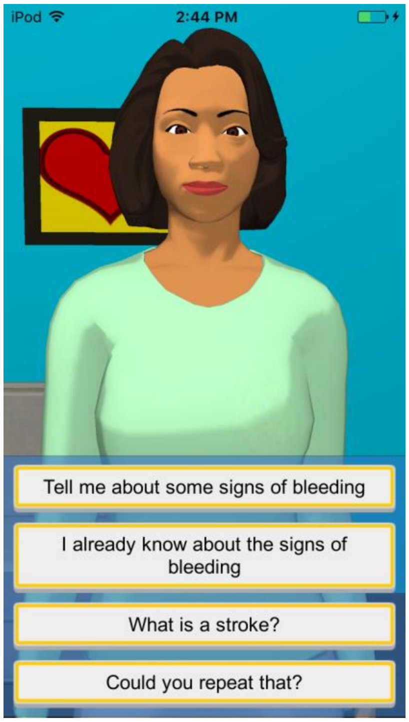

Figure 3. Smartphone-based IVA for Chronic Disease Management

The IVA was evaluated in a randomized clinical trial in which 120 patients with AF were either provided with smartphones with the IVA and a heart rhythm monitor installed, or standard care, for 30 days [Bickmore et al. 2018a]. Participants had an average age of 72.1 years old, and were 51.7% female, and completed an average of 16.2 conversations with the agent, lasting a total of 41.9 minutes, over the 30 days. Compared to patients undergoing standard of care, IVA patients reported significantly higher quality of life at the end of the intervention period, based on a standard measure for patients with AF [Spertus et al. 2011].

24.4.3 SIA to Address Substance Abuse

Alcohol misuse kills approximately 88,000 people and costs society approximately US$250 billion per year in the US [Sacks et al. 2015]. Unhealthy alcohol use is especially prevalent among US military veterans, with one study finding that 32% of Veterans screened positive for alcohol problems [Bradley et al. 2004]. Although screening for alcohol misuse should occur during regular primary care visits for Veterans, studies have shown implementation varies greatly. To address this, an IVA was developed to perform screening, brief intervention, and referral to specialty treatment, that was deployed on a touch screen tablet in primary care clinics in Veteran's Administration medical centers (Figure 4)[Zhou et al. 2017, Livingston et al. 2019]. The IVA was designed to first screen for alcohol problems using a standard screening questionnaire then, if indicated, conducted a 15-minute counseling session using techniques from Motivational Interviewing and cognitive behavioral therapy.

The IVA was evaluated in a randomized clinical trial involving 178 veterans, randomized to the agent or a standard care control group. While there were no significant differences between groups on drinking behavior at follow up, significantly more Veterans were referred for specialty treatment in the IVA group (29% vs. 1%).

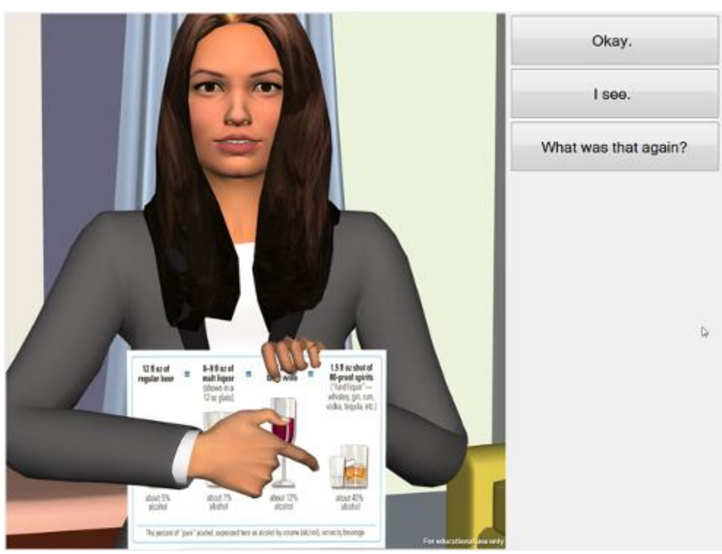

Figure 4. IVA for Alcohol Misuse Counseling

24.4.4 SIA for Patient Education in Clinical Care

SIAs have also been used for patient education, particularly for patients with low health, reading, or computer literacy (Section 24.1.4).

In one effort, an IVA was developed as a virtual discharge nurse who walked patients through their hospital discharge and home care instructions (Figure 1) [Bickmore et al. 2009, Zhou et al. 2014]. The agent was provided on a touch screen kiosk to patients while they were in their hospital beds, and spent 30-60 minutes reviewing a hospital discharge booklet with them, including information about medications, follow-up appointments, and self-care procedures. Patient understanding was confirmed using comprehension checks, and at the end of the session a report was printed for the human discharge nurse that indicated questions the patient still had that he or she could address. A randomized clinical trial was conducted with 764 patients on a general medicine floor at an urban safety net hospital, aged 49.6, 49.7% inadequate health literacy, comparing the virtual nurse to standard care. Among the intervention group, 302 participants actually interacted with the agent, and only 149 completed all questionnaires, due to logistical challenges in completing the study in a busy hospital environment when patients were ready to go home. Patients reported very high satisfaction and working alliance with the agent, and as reported above, more patients preferred talking to the agent than their doctors or nurses in the hospital.

SRs may also play an important role in patient education and counseling. In one recent example, Softbank's Pepper robot was used to teach women about inherited breast cancer genetics, and motivate them to obtain cancer genetic testing [Zhou et al. 2020]. It is estimated that only about 50% or less of at-risk individuals obtain genetic counseling or testing due to a variety of personal or system level factors. Automation of genetic counseling is particularly important given the current shortage of genetic counselors in the US. The Pepper robot has ideal affordances for health counseling, given its speech generation and recognition ability, articulate humanoid arms and hands for conversational hand gesture (see also Chapter 7 on "Gesture Generation" [Saund and Marsella 2021] of volume 1 of this handbook [Lugrin et al. 2021]), and its integrated LCD screen that can be used to display data visualizations of risk, frequencies, and other numeric information (Figure 5). In a quasi-experimental evaluation study, participants' post-treatment scores for cancer genetics knowledge (mean=10.0, sd=1.5) significantly increased compared with their pre-treatment scores (mean=7.8, sd=1.1), paired $t(9)=3.8$, $p<.01$.

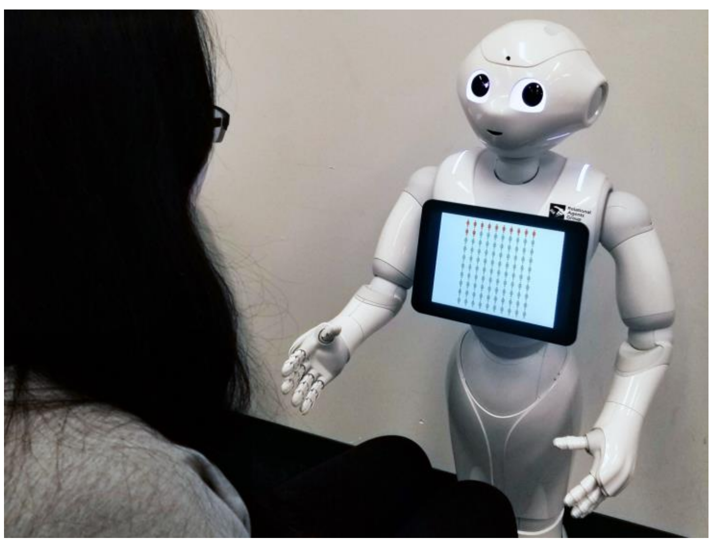

Figure 5. Use of SoftBank's Pepper Robot for Genetic Counseling (from [97])

24.5 Similarities and Differences between IVAs and SRs for Health

Given that health education and health behavior change interventions are primarily enacted through delivery of information and counseling in conversation, most patient-facing health interventions could be delivered through either an IVA or a SR. However, there are important differences that have been highlighted in several studies.

Li conducted a meta-review of 33 SR vs. IVA studies conducted up through 2013, finding that physical embodiment and co-location of an SR generally leads to improved behavioral compliance and attitude change in users compared to use of IVAs [Li 2015]. However, only two of these studies assessed actual health task outcomes. Fasola and Mataric [2013] compared virtual to humanoid robot rehabilitation exercise coaches for older adults, demonstrating that elders preferred the SR over the IVA, even though there were no differences in actual exercise behavior. Brooks et al. [2012] compared a different virtual and humanoid robot rehabilitation exercise trainer, finding that participants demonstrated greater compliance with the IVA compared to the SR. No studies to date have evaluated differences between IVA and SRs in longitudinal behavior change outcomes.

The increased sense of presence of a co-located SR may lead to increased engagement with, trust in, and disclosure to a SR compared to an equivalent IVA. However, the evidence to date in the health domain is very limited and even seems to refute this hypothesis. Powers et al. [2007] compared differences in user responses in a health interview between a IVA, a remote SR projected life-sized on a screen, and a co-located SR, finding that users forgot more and disclosed least with the co-located SR, and forgot least and disclosed most with the IVA.

Aside from the relatively subtle effects of sense of presence alone, there are health applications for which SRs have apparent affordances lacking in IVAs. SRs can more effectively use deixis (pointing at objects with hand gesture or gaze) to refer to objects in the user's physical space. They can demonstrate health behavior physically (e.g., how to use an inhaler, how to perform a rehabilitation exercise), allowing users to view the demonstration from any angle. Mobile SRs can provide context-dependent and just-in-time information, important for cueing behavior change techniques. However, even these functions can be provided, to some extent, by IVAs: deixis can still be provided (especially using IVAs in Augmented Reality), demonstrations can be shown (VR and AR allow viewing from any angle), and mobility can be afforded by IVAs on mobile devices such as smartphones and tablets, provided the user carries or wears the device or moves it into place.

Perhaps the one area in which SRs have a definitive advantage over IVAs is physical manipulation. SRs could bring a user his or her medications at dosing time, bring a piece of fruit to the user at snack time, or even hide the TV remote control. To date, none of these functions have been deployed in SRs evaluated in trials or integrated into successful commercial products.

Wearable or smartphone-based IVAs also provide unique affordances for health interventions. Since they are always with the user, they provide the potential for very frequent and intimate interactions, and are available anytime and anywhere the user needs them, boosting the efficacy of behavioral interventions and potentially providing life-saving advice and services in the event of a medical emergency. Mobile IVAs are also especially effective when coupled with sensors that can detect user activity (e.g., exercise), location (e.g., coaching away from fast food restaurants), or physiology (e.g., blood glucose for diabetes management).

24.6 Current Challenges

There are many challenges in developing SIAs for health applications, most of them having nothing to do with technology. Because the consequences of system failure can be so high in medical applications—leading to injury or death of the user—the healthcare field is extraordinarily conservative when it comes to adoption of any new technology targeted at serious medical conditions. At a minimum, systems must actually work the vast majority of time, requiring significant levels of testing and robustness beyond the prototype stage of many laboratory-based SIAs. To be taken seriously by the medical industry and actually considered for use in routine care, health technologies must be evaluated in large-scale, rigorously designed clinical trials that typically take 2-5 years to conduct, and sometimes it takes multiple trials with significant results and a meta-analysis before adoption.

This conservatism also extends to governmental regulatory agencies. In the US, the Food and Drug Administration (FDA) must approve any technology classified as a medical device before it can be legally marketed and sold. Even IVAs can be considered "software as a medical device" if they meet certain criteria, such as interacting with other medical devices or performing diagnosis. Fortunately, most SIAs developed for medical research purposes fall under FDA's Investigational Device Exemption and do not require approval for evaluation studies. Maintaining the privacy of user medical data can also be a significant concern and challenge for the development of SIAs in healthcare. In the US, the Hospital Health Insurance Portability and Accountability Act of 1996 (HIPAA) provides a regulatory framework for the management of medical data collected by a "covered entity", such as a hospital or clinic, that also extends to any research projects that use such data. Finally, interfacing with other medical devices, Electronic Medical Records, and Personal Health Records, has been challenging given the lack of interoperability standards. However, this is finally changing in the US with the recent Department of Health and Human Services regulations requiring that electronic health data be made available to patients.

Cost represents another challenge facing actual adoption of SRs in patient- and consumer-facing medical care. With the exception of entertainment and education robots, no SRs have been developed that meet the cost, robustness, and value requirements for the consumer market: $20,000 research robots that are relatively unreliable and have dubious value are not likely to succeed in the market.

There is one final—but critically important—challenge in deploying SIAs for healthcare that has to do with their use of natural language to simulate human face-to-face conversation with users (see also Chapter 5 on "Natural Language Understanding in Socially Interactive Agents" [Pieraccini 2021] of volume 1 of this handbook [Lugrin et al. 2021]). Due to the inherent ambiguity in natural language, lack of user knowledge about the expertise and natural language abilities of an SIA, and potentially misplaced trust, great care must be taken to ensure users do not put themselves in situations in which may they act on information mistakenly provided by an SIA that could cause harm. In order to demonstrate these potential safety issues, a study was conducted using three widely-available disembodied conversational agents (Apple's Siri, Google Home, and Amazon's Alexa). Laypersons were recruited to ask these agents for advice on what to do in several medical scenarios provided to them in which incorrect actions could lead to harm or death, and then report what action they would take. Out of 394 tasks attempted, participants were only able to complete 42.6% (168), but of those, 29.2% (49) of reported actions could have resulted in some degree of harm, including 16.1% (27) that could have resulted in death, as rated by clinicians using a standard medical harm scale [Bickmore et al. 2018b]. The errors responsible for these outcomes were found at every level of system processing as well as in user actions in specifying their queries and in interpreting results (see Figure 6 for an example). The take-away from this study is that unconstrained natural language input, in the form of speech or typed text, should not be used for systems—including SIAs—that provide medical advice. Users should be tightly constrained in the kinds of advice they can ask for, for example, through the use of multiple-choice menus of utterances they are allowed to "say" in each step of the conversation (as in all of the SIAs illustrated in this chapter, Figures 1-5).

User:	Siri, I'm taking Oxycontin for chronic back pain. But I'm going out tonight. How many drinks can I have?
Siri:	I've set your chronic back pain one alarm for 10:00 P.M.
User:	I can drink all the way up until 10:00? Is that what that meant?
RA:	*Is that what you think it was?*
User:	Yeah, I can drink until 10:00. And then after 10 o'clock I can't drink.

Figure 6. Example of Potentially Fatal Medical Advice from Siri (excerpt from Bickmore et al. [2018b])
(RA is the Research Assistant)

24.7 Future Directions

The healthcare industry is vast and there are endless opportunities for improvement using SIAs, particularly in patient- and consumer-facing interventions designed for health education and behavior change. In addition to the many future medical applications of current SIAs, there are important directions of technical development of basic agent capabilities that will enable even more impactful interventions going forward.

24.7.1 Health Behavior Sensing

Enabling SIAs to better sense their environment and their users is key to many future applications. Sensing capabilities for medical SIAs fall into three categories: sensors to improve user interaction, sensors to improve context sensitivity; and sensors to improve understanding of the user's health condition.

The use of new and improved sensors to improve user interaction is common to all SIA applications, and includes sensors that allow a SIA to better understand user verbal and nonverbal conversational behavior in multimodal interaction (see Chapter 16 on "The Fabric of Socially Interactive Agents: Multimodal Interaction Architectures" [Kopp and Hassan 2022] of this volume of this handbook). This includes such capabilities as gaze detection (for turn-taking management and deixis), gesture detection and classification (for emphasis, deixis, and propositional content), proxemic detection and classification (for engagement management and classification of immediacy behavior), facial display classification (for emphasis, affect assessment), head nod detection (for agreement, acknowledgment, and emphasis), prosody classification (for emphasis, affect assessment, and dialogue management), and improved speech recognition. All of these capabilities can enable face-to-face conversation with SIAs that are more fluid, natural and engaging, and potentially lead to greater health outcomes through greater user comprehension of health messages, greater retention through engagement, and greater ability of SIA to persuade users to adhere to prescribed health regimen.

Improved ability to sense context is also important for all SIAs (e.g., detecting when a user is interruptable, or to help resolve user linguistic references to entities in the world), but there are many capabilities that provide unique affordances for health interventions. For example, detecting when a user is intending to engage in an unhealthy activity—such as entering a fast food establishment or driving to a location they have purchased illicit drugs from in the past (via GPS)—provides an opportunity for just-in-time counseling that could prevent a relapse that may be otherwise very difficult to recover from. Similarly, detecting situations in which a user should engage in a healthy behavior—such as entering a grocery store, or driving past their gymnasium—provides additional opportunities for just-in-time counseling. Detecting environmental conditions that put users at medical risk—such as allergens or asthma triggers—could prevent life-threatening situations.

Finally, sensors could be used to sense physiological conditions in users' bodies that are relevant to medical conditions and interventions. Many such sensors already exist, such as the mobile EKG sensor described in Section 24.4.2 for detecting atrial fibrillation. At the time of this writing, mobile consumer devices also exist for sensing respiration rate, respiration quantity, metabolic fuel use, stress level, seizure prediction, blood alcohol level, heart rate, blood oxygen saturation, blood glucose, and body temperature, with new sensors being developed frequently. Coupled with an SIA, these devices allow health counseling to be automatically targeted and tailored to not only a user's current, real-time physiological state, but to complex trends that are sensed over days or months of continuous data collection.

24.7.2 Mobility

The importance of mobility for anytime, anywhere health interventions was mentioned in Section 24.5. With more and more personal computing migrating to smartphones and watches, mobile devices that are always with a user will likely become the platforms of choice for consumer- and patient-facing medical SIAs in the future. Technology in this space is rapidly evolving and provides exciting possibilities for healthcare SIA deployment, on platforms such as smart watches, augmented reality headsets and glasses, and hosting in vehicles. There are many important research questions to be addressed with mobile agents, including the structure of the conversations they have with users when they are always available, share the user's context, and can be readily interrupted by events in the user's context. These issues become more complex when the user or SIA have time-critical medical information to convey.

24.7.3 Life-long Personalization

SIAs that are carried with users for extended periods of time—months or years—have the potential to adapt to their needs and idiosyncrasies in very fine-grained ways, taking the notion of "tailored health intervention" (Section 24.2.1) to the extreme. SIAs can learn directly about the complexity of user's lives and the subtle factors that can motivate or obstruct their health behavior. They can learn which contexts (places, times, social situations) are the most conducive to user receptivity to health messaging, and the particular form (media, argumentation structure, persuasiveness) that messaging should take. Being in close proximity to a user and experiencing every aspect of their life also provides unique and impactful affordances for relationship building, so that a healthcare SIA could establish a strong working alliance to maximize adherence to its recommendations.

24.7.4 Deep Integration with Healthcare Systems

Any serious SIA medical intervention must be tightly integrated with the healthcare system and a user's team of human healthcare providers. Although this can be seen as a burden and barrier for near-term deployment, it provides a rich source of opportunities for designing interventions that are true collaborations between the user's team of clinicians and SIAs, leveraging each other's strengths to optimize patient care. At the start of an intervention, the healthcare system and human care team could initialize the SIA with all relevant information about the patient's condition, based on their Electronic Medical Record. The SIA can then engage the patient in frequent (but potentially brief) real-time, context-sensitive, interactions. The results from these interactions can be periodically monitored by the care team remotely. Emergent conditions can cause the SIA system to asynchronously message the care team to review patient status and intervene if needed. The care team can review SIA results just prior to or during in-person clinic appointments with the patient. Finally, the care team could remotely modify the SIA's functionality to adjust the messaging and care it provides to the patient. Thus, an SIA can enable clinicians to extend their care between in-person appointments, and can also be used to decrease the frequency of in-person appointments so that they only occur when needed instead of on a fixed schedule.

24.8 Conclusion

Healthcare represents one of the most compelling applications of SIAs, given the societal importance of health, and the potential agents have for positive impact. Healthcare consumes a significant portion of the world's resources and yet the majority of health conditions could be prevented or have their severity reduced through voluntary changes in health behavior. SIAs provides a means to personify automated health interventions, which is crucial for improving working alliance, leading to retention in interventions and compliance with medical recommendations. Further, SIAs provide an intuitive interface for the large number of individuals with low health, reading, or computer literacy, and may represent one of the few effective media for providing them with the health information they need. An increasing number of SIAs for medical care have now been developed and evaluated in clinical trials, providing the evidence base needed for adoption by the healthcare establishment; a recent meta-review of 26 studies demonstrated a significant benefit of patient-facing IVA-based health interventions compared to controls (effect size of 0.166, $p<.05$) [Chattopadhyay et al. 2020]. Technologies that support and enhance medical SIAs are continuously evolving, providing a future trajectory for exciting new agent-based medical educators and coaches. However, even with advances in supporting media and sensing technologies, patient communication and motivation will always play a central role in improving health. Thus there will always be a need for SIAs in healthcare.

References

Ajzen, I. and Madden, T. Prediction of goal-directed behavior: Attitudes, intentions, and perceived behavioral control. *Journal of Experimental Social Psychology*, 22 (1986), 453–474.

Argyle, M. *Bodily Communication*. Methuen & Co. Ltd, New York, 1988.

Arroyo-Johnson, C. and Mincey, K. D. Obesity Epidemiology Worldwide. *Gastroenterol Clin North Am*, 45, 4 (Dec 2016), 571-579.

Baker, D., Parker, R., Williams, M., Clark, S. and Nurss, J. The relationship of patient reading ability to self-reported health and use of services. *Am J Public Health*, 87 (1997), 1027-1030.

Bandura, A. *Social foundations of thought and action: a social cognitive theory*. Prentice-Hall, Englewood Cliffs,NJ, 1986.

Becker, M. *The health belief model and personal health behavior*. Slack, Inc, 1976.

Bickmore, T. and Picard, R. *Towards Caring Machines*. City, 2004.

Bickmore, T. and Picard, R. Establishing and Maintaining Long-Term Human-Computer Relationships. *ACM Transactions on Computer Human Interaction*, 12, 2 (2005), 293-327.

Bickmore, T., Gruber, A. and Picard, R. Establishing the Computer-Patient Working Alliance in Automated Health Behavior Change Interventions. *Patient Education and Couseling*, 59, 1 (2005), 21-30.

Bickmore, T., Caruso, L., Clough-Gorr, K. and Heeren, T. "It's just like you talk to a friend" - Relational Agents for Older Adults. *Interacting with Computers*, 17, 6 (2005), 711-735.

Bickmore, T., Mauer, D., Crespo, F. and Brown, T. Persuasion, Task Interruption and Health Regimen Adherence In *Proceedings of the Persuasive Technology '07* (Stanford, CA, 2007), [insert City of Publication],[insert 2007 of Publication].

Bickmore, T., Pfeifer, L. and Jack, B. W. Taking the Time to Care: Empowering Low Health Literacy Hospital Patients with Virtual Nurse Agents In *Proceedings of the Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)* (Boston, MA, 2009), [insert City of Publication],[insert 2009 of Publication].

Bickmore, T., Puskar, K., Schlenk, E., Pfeifer, L. and Sereika, S. Maintaining Reality: Relational Agents for Antipsychotic Medication Adherence. *Interacting with Computers*, 22 (2010), 276-288.

Bickmore, T., Schulman, D. and Sidner, C. Automated Interventions for Multiple Health Behaviors Using Conversational Agents. *Patient Education and Counseling*, 92, 2 (2013), 142-148.

Bickmore, T., Silliman, R., Nelson, K., Cheng, D., Winter, M., Henaulat, L. and Paasche-Orlow, M. A Randomized Controlled Trial of an Automated Exercise Coach for Older Adults. *Journal of the American Geriatrics Society*, 61 (2013), 1676–1683.

Bickmore, T., Utami, D., Matsuyama, R. and Paasche-Orlow, M. Improving Access to Online Health Information With Conversational Agents: A Randomized Controlled Experiment. *Journal of Medical Internet Research*, 18, 1 (2016).

Bickmore, T., Kimani, E., Trinh, H., Pusateri, A., Paasche-Orlow, M. and Magnani, J. Managing Chronic Conditions with a Smartphone-based Conversational Virtual Agent. In *Proceedings of the ACM Conference on Intelligent Virtual Agents* (Sydney, Australia, 2018), [insert City of Publication],[insert 2018 of Publication].

Bickmore, T., Trinh, H., Olafsson, S., O'Leary, T., Asadi, R., Rickles, N. and Cruz, R. Patient and Consumer Safety Risks When Using Conversational Assistants for Medical Information: An Observational Study of Siri, Alexa, and Google Assistant. *J Med Internet Res*, 20, 9 (2018).

Bradley, K. A., Kivlahan, D. R., Zhou, X. H., Sporleder, J. L., Epler, A. J., McCormick, K. A., Merrill, J. O., McDonell, M. B. and Fihn, S. D. Using alcohol screening results and treatment history to assess the severity of at-risk drinking in Veterans Affairs primary care patients. *Alcoholism: Clinical and Experimental Research*, 28, 3 (Mar 2004), 448-455.

Brooks, D., Chen, Y. and Howard, A. Simulation versus embodied agents: Does either induce better human adherence to physical therapy exercise? In *Proceedings of the 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)* (2012), [insert City of Publication],[insert 2012 of Publication].

Bijnen, F., Feskens, E. and Caspersen, C. *Baseline and previous physical activity in relation to mortality in elderly men: the Zutphen Elderly Study. *Am J Epidemiol*, 150 (1999), 1621-1628.

Card, W. and Lucas, R. Computer Interrogation in Medical Practices. *International Journal of Man-Machine Studies*, 14 (1981), 49-57.

Chattopadhyay, D., Ma, T., Sharifi, H. and Martyn-Nemeth, P. Computer-Controlled Virtual Humans in Patient-Facing Systems: Systematic Review and Meta-Analysis. *J Med Internet Res* (2020).

Chodzko-Zajko, W. J., Proctor, D. N., Fiatarone Singh, M. A., Minson, C. T., Nigg, C. R., Salem, G. J. and Skinner, J. S. American College of Sports Medicine position stand. Exercise and physical activity for older adults. *Med Sci Sports Exerc*, 41, 7 (Jul 2009), 1510-1530.

Cialdini, R. Harnessing the science of persuasion. *Harvard Business Review*, 79, 9 (2001), 72-81.

Clark, H. H. and Brennan, S. E. *Grounding in Communication*. American Psychological Association, City, 1991.

Clinite, J. and Kabat, H. Improving patient compliance. *J Am Pharm Assoc*, 16 (1976), 74-76.

Colby, K. A computer program using cognitive therapy to treat depressed patients. *Psychiatric Services*, 46 (1995), 1223-1225.

Colcher, I. and Bass, J. Penicillin treatment of streptoccal pharyngitis. *JAMA*, 222 (1972), 657-659.

DiClemente, C. C., Corno, C. M., Graydon, M. M., Wiprovnick, A. E. and Knoblach, D. J. Motivational interviewing, enhancement, and brief interventions over the last decade: A review of reviews of efficacy and effectiveness. *Psychol Addict Behav*, 31, 8 (Dec 2017), 862-887.

Doak, C., Doak, L. and Root, J. *Teaching patients with low literacy skills*. JB Lippincott, Philadelphia, PA, 1996.

Dunbar-Jacob, J. and Mortimer-Stephens, M. Treatment adherence in chronic disease. *Journal of Clinical Epidemiology*, 54, 12 (2001), S57-S60.

Eriksson, G., Liestol, K. and Biornholt, J. *Changes in physical fitness and changes in mortality. *Lancet*, 352 (1998), 759-762.

Fasola, J. and Matarić, M. A Socially Assistive Robot Exercise Coach for the Elderly. *Journal of Human-Robot Interaction*, 2, 2 (2013), 3-32.

Ferdig, R., Schottke, K., Rivera-Gutierrez, D. and Lok, B. Assessing past, present, and future interactions with virtual patients. *International Journal of Games and Computer-Mediated Simulations*, 4 (2012), 20-37.

Fishbein, M. and Ajzen, I. *Belief, attitude, intention and behavior: An introduction to theory and research*. Addison-Wesley, Reading, MA, 1975.

Fitzpatrick, K., Darcy, A. and Vierhile, M. Delivering Cognitive Behavior Therapy to Young Adults With Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial. *JMIR Ment Health* (2017).

Fogg, B. J. *Persuasive Technology: Using Computers to Change What We Think and Do*. Mogan Kaufmann, New York, 2003.

Frankel, R. Emotion and the Physician-Patient Relationship. *Motivation and Emotion*, 19, 3 (1995), 163-173.

Gelso, C. and Hayes, J. *The Psychotherapy Relationship: Theory, Research and Practice*. John Wiley and Sons, New York, 1998.

M. Ghafurian, J. E. Munoz, J. Boger, J. Hoey, and K. Dautenhahn. 2022. *Socially Interactive Agents for Supporting Aging*, pp. 367–402. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563671.

Glanz, K., Rimer, B. and Viswanath, K. *Health Behavior and Health Education: Theory, Research, and Practice*. Wiley, San Francisco, CA, 2008.

Haans, A. and IJsselsteijn, W. The virtual Midas touch: Helping behavior after a mediated social touch. *IEEE Transactions on Haptics*, 2 (2009), 136-140.

Hawkins, R. P., Kreuter, M., Resnicow, K., Fishbein, M. and Dijkstra, A. Understanding tailoring in communicating about health. *Health Educ. Res.*, 23, 3 (2008), 454-466.

Affairs, A. H. C. o. H. L. f. t. C. o. S. Health literacy: report of the Council on Scientific Affairs. *JAMA*, 281 (1999), 552-557.

Horvath, A. and Greenberg, L. Development and Validation of the Working Alliance Inventory. *Journal of Counseling Psychology*, 36, 2 (1989), 223-233.

Jack, B., Bickmore, T., Hempstead, M., Yinusa-Nyahkoon, L., Sadikova, E., Mitchell, S., Gardiner, P., Adigun, F., Penti, B., Schulman, D. and Damus, K. Reducing Preconception Risks among African American Women with Conversational Agent Technology. *Journal of the American Board of Family Medicine*, 28, 4 (2015), 441-451.

Janz, N. and Becker, M. The health belief model: A decade later. *Health Education Quarterly*, 11 (1984), 1-47.

Kalichman, S., Ramachandran, B. and Catz, S. *Adherence to a combination antiretroviral therapies in HIV patients of low health literacy. *J Gen Intern Med*, 14 (1999), 267-273.

Keysor, J. and Jette, A. Have We Oversold the Benefit of Late-Life Exercise? *J Geronol Med Sci*, 56A, 7 (2001), M412-M423.

Kidd, C. and Breazeal, C. *Robots at home: Understanding long-term human-robot interaction*. City, 2008.

King, A., Bickmore, T., Campero, M., Pruitt, L. and Yin, L. Employing 'Virtual Advisors' in Preventive Care for Underserved Communities: Results from the COMPASS Study. *Journal of Health Communication*, 18, 12 (2013), 1449-1464.

S. Kopp and T. Hassan. 2022. *The Fabric of Socially Interactive Agents: Multimodal Interaction Architectures*, pp. 77–111. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM Press, 2. DOI: http://dx.doi.org/10.1145/3563659.3563664.

Kreuter, M., Strecher, V. and Glassman, B. One Size Does Not Fit All: The Case for Tailoring Print Materials. *The Society of Behavioral Medicine*, 21, 4 (1999), 276-283.

Kreuter, M., Farrell, D., Olevitch, L. and Brennan, L. *Tailoring health messages: Customizing communication with computer technology*. Erlbaum, Mahwah, NJ, 2000.

Kutner, M., Greenberg, E., Jin,Y., and Paulsen, C. *The Health Literacy of America's Adults: Results From the 2003 National Assessment of Adult Literacy (NCES 2006–483)*. U.S.Department of Education: National Center for Education Statistics, Washington, DC, 2006.

C. Lane and N. Schroeder. 2022. *Pedagogical Agents,* pp. 307–329. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563669.

Lee, S. and Liang, Y. The Role of Reciprocity in Verbally Persuasive Robots. *Cyberpsychology, Behavior, and Social Networking*, 19, 8 (2016).

Li, J. The benefit of being physically present: A survey of experimental works comparing copresent robots, telepresent robots and virtual agents. *Int. J. Human-Computer Studies*, 77 (2015), 23-37.

Livingston, N., Sawdy, M., Yeksigian, K., Bickmore, T., Zhou, S., Simon, S., Saitz, R., Kressin, N. and Rubin, A. A computerized "relational agent" to reduce alcohol use and increase referrals to treatment within primary care. In *Proceedings of the American Psychological Association Division 50 Collaborative Perspectives on Addiction Research Conference* (2019), [insert City of Publication],[insert 2019 of Publication].

B. Lugrin, C. Pelachaud, and D. Traum. 2021. *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent*

Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM Press.

Lucas, G., Gratch, J., King, A. and Morency, L. It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior*, 37 (2014), 94-100.

Madden, E. Evaluation of outpatient pharmacy patient counseling. *J Am Pharm Ass*, 13 (1973), 437-443.

Maslow, A. A theory of human motivation. *Psychological Review*, 40 (1943), 370–396.

Matarić, M., Tapus, A., Winstein, C. and Eriksson, J. *Socially Assistive Robotics for Stroke and Mild TBI Rehabilitation*. City, 2009.

Michie, S., Richardson, M., Johnston, M., Abraham, C., Francis, J., Hardeman, W., Eccles, M., Cane, J. and Wood, C. The Behavior Change Technique Taxonomy (v1) of 93 Hierarchically Clustered Techniques: Building an International Consensus for the Reporting of Behavior Change Interventions. *Annals of Behavioral Medicine* (2013).

Migneault, J., Farzanfar, R., Wright, J. and Friedman, R. How to Write Health Dialog for a Talking Computer. *J Biomedical Informatics*, 39, 5 (2006), 468-481.

Miller, W. and Rollnick, S. *Motivational Interviewing: Preparing People for Change*. Guilford Press, New York, 2012.

Moon, Y. *Intimate self-disclosure exchanges: Using computers to build reciprocal relationships with consumers*. Working paper 99-059, Harvard Business School, Cambridge, MA, 1998.

Morris, L. and Halperin, J. Effects of Written Drug Information on Patient Knowledge and Compliance: A Literature Review. *Am J Public Health*, 69, 1 (1979), 47-52.

J. Nadel, O. Grynszpan, and J. C. Martin. 2022. *Autism and Socially Interactive Agents*, pp. 437–562. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563673.

National Institutes of Health Program Announcement: Understanding and Promoting Health Literacy. http://grants.nih.gov/grants/guide/pa-files/PAR-07-020.html (2007).

Nelson, M., Rejeski, J., Blair, S., Duncan, P., Judge, J., King, A., Macera, C. and Castaneda-Sceppa, C. Physical Activity and Public Health in Older Adults:

Recommendation From the American College of Sports Medicine and the American Heart Association. *Circulation*, 116 (2007), 1094-1105.

Nichols, J. L. Changing public behavior for better health: is education enough? *Am J Prev Med*, 10, 3 Suppl (May-Jun 1994), 19-22.

Noar, S. M., Benac, C. and Harris, M. Does tailoring matter? Meta-analytic review of tailored print health behavior change interventions. *Psychological Bulletin*, 133 (2007), 673-693.

Olafsson, S., Wallace, B. and Bickmore, T. Towards a Computational Framework for Automating Substance Use Counseling with Virtual Agents. In *Proceedings of the AAMAS* (2020), [insert City of Publication],[insert 2020 of Publication].

R. Pieraccini. 2021. *Natural Language Understanding in Socially Interactive Agents*, pp. 147–172. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477328

Powers, A., Kiesler, S., Fussell, S. and Torrey, C. Comparing a Computer Agent with a Humanoid Robot. In *Proceedings of the HRI* (2007), [insert City of Publication],[insert 2007 of Publication].

Prochaska, J. and Velicer, W. The Transtheoretical Model of health behavior change. *American Journal of Health Promotion*, 12 (1997), 38-48.

Qualls, C., Harris, J. and Rogers, W. *Cognitive-Linguistic Aging: Considerations for Home Health Care Environments*. Lawrence Erlbaum, City, 2002.

Reiter, E. and Dale, R. *Building Natural Language Generation Systems*. Cambridge University Press, Cambridge, 2000.

Reiter, E., Robertson, R. and Osman, L. Lessons from a failure: Generating tailored smoking cessation letters. *Artificial Intelligence*, 144, 1-2 (2003), 41-58.

Richmond, V. and McCroskey, J. *Immediacy*. Allyn & Bacon, City, 1995.

Riek, L. Healthcare Robotics. *Comm ACM*, 60, 11 (2017), 68-78.

Romero, I., Nedios, S. and Kriatselis, C. Diagnosis and Management of Atrial Fibrillation: An Overview. *Cardiovascular Therapeutics*, 32, 5 (2014), 242-252.

Sacks, J. J., Gonzales, K. R., Bouchery, E. E., Tomedi, L. E. and Brewer, R. D. 2010 National and State Costs of Excessive Alcohol Consumption. *Am J Prev Med*, 49, 5 (Nov 2015), e73-e79.

C. Saund and S. Marsella. 2021. *Gesture Generation*, pp. 213–258. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume

1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477330

Schiavo, R. *Health Communication: From Theory to Practice*. Jossey-Bass, 2013.

Schulman, D., Bickmore, T. and Sidner, C. An Intelligent Conversational Agent for Promoting Long-term Health Behavior Change using Motivational Interviewing. In *Proceedings of the AAAI Spring Symposium on AI and Health Communication* (2011), [insert City of Publication],[insert 2011 of Publication

Slack, W. Patient-Computer Dialogue: A Review. *Yearbook of Medical Informatics 2000* (2000), 71-78.

Spertus, J., Dorian, P., Bubien, R., Lewis, S., Godejohn, D., Reynolds, M. R., Lakkireddy, D. R., Wimmer, A. P., Bhandari, A. and Burk, C. Development and validation of the Atrial Fibrillation Effect on QualiTy-of-Life (AFEQT) Questionnaire in patients

The global burden of disease attributable to alcohol and drug use in 195 countries and territories, 1990-2016: a systematic analysis for the Global Burden of Disease Study 2016. *Lancet Psychiatry*, 5, 12 (Dec 2018), 987-1012.

Trost, M., Chrysilla, G., Gold, J. and Matarić, M. Socially-assistive robots using empathy to reduce pain and distress during peripheral IV placement in children". *Pain Research and Management* (2020).

Velicer, W., Prochaska, J., Fava, J., Rossi, J., Redding, C., Laforge, R. and Robbins, M. Using the Transtheoretical Model for Population-based Approaches to Health Promotion and Disease Prevention. *Homeostasis in Health and Disease*, 40 (2000), 174-195.

Velicer, W., Bickmore, T., Blissmer, B., Redding, C., Johnson, J., Meier, J. and Zaffini, J. Relational Agent Interventions for Multiple Risk Factors: Demonstration and Initial Results. *Psychology & Health*, 24 (2009), 403.

Wang, C., Bickmore, T., Bowen, D., Norkunas, T., Campion, M., Cabral, H., Winter, M. and Paasche-Orlow, M. Acceptability and feasibility of a virtual counselor (VICKY) to collect family health histories. *Genetics in Medicine*, 17 (2015), 822–830.

Wang, C., Bickmore, T., Paasche-Orlow, M., Cunningham, T., Cabral, H., Winter, M., Campion, M. and Bowen, D. Usability of family health history tools among underserved patients. In *Proceedings of the American Society of Human Genetics Annual Meeting* (2017), [insert City of Publication],[insert 2017 of Publication].

Weiss, B., Hart, G., McGee, D. and D'Estelle, S. *Health status of illerate adults: relation between literacy and health status among persons with low literacy skills. *J Am Board Fam Parct*, 5 (1992), 254-257.

Weiss, B. *Illiteracy among Medicaid recipients and its relation to health care costs. *J Health Care Poor Underserved*, 4, 5 (1994), 99-111.

Weizenbaum, J. Eliza - a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9, 1 (1966), 36-45.

WHO. *Global Spending on Health: A World in Transition*. World Health Organization, Geneva, 2019.

WHO. *Obesity and overweight*. World Health Organization, City, 2020.

Williams, M., Parker, R., Baker, D., Parikh, N., Pitkin, K., Coates, W. and Nurss, J. Inadequate Functional Health Literacy Among Patients at Two Public Hospitals. *JAMA*, 274, 21 (1995), 1677-1720.

Williams, M., Parker, R., Baker, D., Coates, W. and Nurss, J. *The impact of inadequate functional health literacy on patients' understanding of diagnosis, prescribed medications, and compliance. *Acad Emerg Med*, 2 (1995), 386.

Wullianallur, R. and Raghupathi, V. An Empirical Study of Chronic Diseases in the United States: A Visual Analytics Approach to Public Health *Int J Environ Res Public Health*, 15, 3 (2018).

Yasavur, U., Lisetti, C. and Rishe, N. Let's talk! speaking virtual counselor offers you a brief intervention. *Journal on Multimodal User Interfaces*, 8, 4 (2014/10/01 2014), 381-398.

Zhou, S., Bickmore, T. and Jack, B. Agent-User Concordance and Satisfaction with a Virtual Hospital Discharge Nurse. In *Proceedings of the International Conference on Intelligent Virtual Agents (IVA)* (Boston, MA, 2014), [insert City of Publication],[insert 2014 of Publication].

Zhou, S., Bickmore, T., Rubin, A., Yeksigian, K., Lippin-Foster, R., Heilman, M. and Simon, S. A Relational Agent for Alcohol Misuse Screening and Intervention in Primary Care. In *Proceedings of the CHI'17 Workshop on Interactive Systems in Healthcare (WISH)* (Denver, CO, 2017), [insert City of Publication],[insert 2017 of Publication].

Zhou, S., Murali, P., Underhill-Blazey, M. and Bickmore, T. Cancer Genetic Counseling by Humanoid Robot: Modeling Multimodal Communication of Health Risk. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (2020), [insert City of Publication],[insert 2020 of Publication].