Ambient Personal Environment Experiment (APEX): A Cyber-Human Prosthesis for Mental, Physical and Age-Related Disabilities

David J. Atkinson, Bonnie J. Dorr, Micah H. Clark, William J. Clancey, Yorick Wilks
Institute for Human and Machine Cognition, 15 SE Osceola Avenue, Ocala, FL
{datkinson, bdorr, mclark, wclancey, ywilks}@ihmc.us

Abstract

We present an emerging research project in our laboratory to extend ambient intelligence (AmI) by what we refer to as "extreme personalization" meaning that an *instance* of ambient intelligence is focused on one or at most a few individuals over a very long period of time. Over a lifetime of co-activity, it senses and adapts to a person's preferences and experiences, and crucially, his or her (changing) special needs; needs that differ significantly from the normal baseline. We refer to our agent-based cyber-physical system as Ambient Personal Environment eXperiment (APEX). It aims to serve as *a Companion*, *a Coach*, and *a Caregiver*: crucial support for individuals with mental, physical, and age-related disabilities and those other people who help them. We propose that an instance of APEX, interacting socially with each of these people, is both a *social actor* as well as a *cyber-human prosthetic device*. APEX is an ambitious integration of multiple technologies from Artificial Intelligence (AI) and other disciplines. Its successful development can be viewed as a grand challenge for AI. We discuss in this paper three research thrusts that lead toward our vision: robust intelligent agents, semantically rich human-machine interaction, and reasoning from comprehensive multi-modal behavior data.

Introduction

Ambient intelligence (AmI) extends and combines earlier paradigms of pervasive computing with sensor networks, human-centered interfaces, mobility, Artificial Intelligence (AI), robotics, intelligent agents and the "Internet of Things." The concept is compelling as it promises to deliver an integrated computing, device and networking infrastructure that provides services while remaining largely hidden from the view of users.

The research themes in our laboratory seek to extend AmI by what we refer to as "extreme personalization", meaning that a given *instance* of APEX is focused on at most a few individuals, e.g., client, caregiver, and physician, over a very long period of time. Over a lifetime of co-activity, it senses and adapts to each person's preferences and experiences, and crucially, his or her (changing) special needs. We refer to this agent-based cyber-physical system as Ambient Personal Environment eXperiment (APEX). It serves as *a Companion*, *a Coach*, and a surrogate *Caregiver* to the client: crucial support roles for individuals with mental, physical, and age-related disabilities. It supports the primary caregiver, automating some tasks, advising, and facilitating interaction with the client. It provides the client's remote physician with an "extra set of eyes" for monitoring the client's progress and by being alert for anomalies requiring medical attention. We propose that an instance of APEX, interacting socially with each of these individuals, is both a *social actor* as well as a *cyber-human prosthetic* device (Hamilton 2001).

APEX must learn, adapt and perform in a natural environment that is rich with features, many of which are usually irrelevant or at least uncertain. To achieve this, APEX depends upon successful integration of multiple technologies from artificial intelligence, human-centered design, cognitive science, computational linguistics, human-machine interaction, and robotics. As such, it represents a major stretch goal that is likely to drive each area in some new directions. While incremental results will certainly be useful, the achievement of APEX surely represents an interesting grand challenge with high social value.

In the sections below, we present a brief overview of selected functional requirements and three research thrusts that lead toward our vision for APEX: (1) robust intelligent agents; (2) semantically rich human-machine interaction, and; (3) reasoning from comprehensive multi-modal behavior data. Each poses unique challenges, yet the technologies that support each area form a synergistic combination that we contend will lead us towards generalizable solutions.

We conclude with a discussion of related work, the broader impact of this type of system for AI and related research, and the potential social benefits as well as emerging risks to privacy and security.

Ambient Personal Environment

Our vision of APEX as a cyber-human prosthesis for persons with mental, physical and age-related disabilities is driven by two overarching functional requirements:

- The ability to <u>sense, interpret and change a person's environment</u> (e.g., physical objects, enclosed spaces, ambient attributes) in the context of delivering specific health-related support services.
- The ability, over potentially a lifetime of co-activity, to <u>learn and adapt</u> to an individual's changing special needs that arise from mental, physical, and age-related disabilities.

Additionally, since APEX must interact with multiple people, it must fill a niche that is complementary to the other actors in a disabled person's life. This requires:

- The ability to <u>function as an intelligent actor-agent</u> in the role of a *Companion*, a *Caregiver,* and *a Coach*, aiding and easing the burden to people in these traditional roles.

Therefore, it is critically important for the success of APEX not only that we address user needs that are common to the community of people with cognitive, physical, and age-related disabilities, but that we also consider the needs of the people who are part of the the APEX clients' lives. We envision a long-term relationship with differing levels of disability and sickness, though here we will refer to periods where such help is most needed.. It is probably more correct to say that APEX becomes adapted to the entire system—the people, activities, and environment as a whole, centered *around* the person with the disability.

While space precludes an extended discussion of the needs particular to each of these roles, there are several requirements that should be highlighted as they are broadly common to research on ambient assisted living.

Individual Client Needs

The practice of nursing observation of patients in a hospital setting provides some important topics to consider in AmI research. In current practice, nursing observation is necessarily intrusive. Many patients have negative reactions associated with a high level of intrusiveness. With lower levels of intrusion, patients report positive effects of observation including a sense of support. However, these benefits are negated if patients feel observers lack empathy or seem remote. They also react negatively if they feel they are not given sufficient information about the purpose of observation or a medical process to be provided by the nurse. Page (2006) provides an excellent review of relevant studies as well as citations to a rich set of original sources.

Therefore, new studies are required in order to help parameterize the idea of "an optimal sense of intrusiveness," that is, intrusiveness that elicits the *positive affect* of support without the *negative affect* associated with observation. Furthermore, APEX must simulate empathy based on modeling the client and use that to guide informative interactions (Bee et. al. 2010; Kearns et. al. 2014).

Caregiver Needs

Many clients have a human caregiver and it is important to recognize the importance of *augmenting* rather than *replacing* them, even as APEX provides essential help to the client that might otherwise fall to the caregiver. APEX must adapt and be functional with respect to the caregiver's role, preferences, and intentions, serving as the caregiver's agent even when the caregiver is not present. Augmentation, not replacement, is a common requirement for many applications of automation and this case is no different. Our approach is to position APEX as an aide to the caregiver, for example, by assisting in communication, reminding, observation, and so forth. Our recent studies, discussed below, with veterans who have suffered Traumatic Brain Injuries (TBI), have highlighted the potential benefits of mediating communication between client and caregiver using *Companion* agents (Wilks et. al. 2014).

Primary Physician Needs

Physicians must reconcile the dual need to provide excellent health care while avoiding excessive office or in-patient visits. Once a patient is discharged, e.g., from a Veterans Administration hospital brain injury rehabilitation unit, the primary physician's focus is on maintenance of stable day-to-day health, rehabilitative progress to the extent it is possible, and remaining alert to any signal that a client's condition requires immediate or near-term attention. These are three areas where APEX may usefully assist physicians. The early detection of clinically meaningful anomalies in client behavior would be very valuable to physicians and this is one area we are pursuing, building on previous work by our collaborators at the

Veterans Administration Polytrauma Unit in Tampa, discussed below (Kearns, Nams and Fozard 2010).

Technical Approach

Our general approach is to explore computational models that integrate core AI components with intelligent agent architectures to enable robust, trustworthy adaptive autonomy in the context of long-duration human-machine joint activity. In so doing, we will necessarily push the limits of core AI algorithms for natural language, multi-modal social interaction, theory of mind, and more. As we explore these models, it is apparent that the work will benefit greatly from multi-agent model-based discrete event simulation for guiding the design and tuning of complex cyber-physical intelligent agents, experimental design and the assimilation of vast quantities of sensor data to models for analysis and theory development (Clancey et. al. 1998, 2007). Our formulation of specific studies will be guided by analysis of parameterized exploratory simulations of use-case scenarios that we are now developing.

APEX includes a physical laboratory, built inside to resemble a small home. Previous work in this shared facility has focused on learning-by-observation for cooking tasks in a high quality kitchen. We are building out additional mock rooms where the walls and ceiling contain our multiple sensors and interactive devices including touch screens, structured lighting, motion capture, and more. Completion of the lab build depends on the timing and availability of appropriate funding; at the time of this writing we can only report, "in progress."

Robust Intelligent Agents

An instantiation of APEX is an autonomous, intelligent, and social agent (at times personified and/or embodied as discussed elsewhere in this paper) that takes the role of a life-long companion, providing highly personalized and a dynamic, ever-changing degree of assistance and support for healthy assisted living. In this way, APEX differs substantial from other agent technology applied to healthcare (Isern, Sánchez, and Moreno 2010). The capability to predict, plan, and manage physical effects along with attention to individual behavior along psychological and social dimensions of the disabled individual, caregiver and primary physician requires a high degree of *shared awareness*. This forms a basis for human-machine trust, a foundation of teamwork and the adaptive autonomy for effective human-machine joint control that the APEX problem domain requires (Atkinson, Clancey and Clark 2014).

Successful sensing and planning along behavioral and psycho-social dimensions may be achieved by using predictive cognitive models to underlie the system's "theory of mind" regarding others (Premack and Woodruff 1978). Though it is already a challenge to model typical humans, for APEX the matter is complicated by the fact that many clients will be cognitively impaired. Some progress has been made toward modeling various cognitive impairments such as Autism and Alzheimer's disease using existing cognitive architectures (Matessa 2008; Serna, Pigot, and Rialle 2007), but the development of atypical cognitive models remains difficult. For clients with TBI, the applicability of existing cognitive architectures is unclear particularly because of the range of function that people with TBI can have from very low (with no initiative or memory) to relatively high. The nature of TBI is such that the individual's impairments are at once both profound and highly unique, which serves to undermine common modeling assumptions and architectural commitments regarding cognitive processes and capabilities.

We envision APEX to have access to real-time and historical observation data; therefore, our approach to client modeling is to exploit machine learning (e.g., statistical, inductive techniques) where possible. This includes using the methodologies of behavior analysis (Cooper, Heron, Heward 2007) to develop predictive models of client behavior that is contingent on objects and events in the environment and the client's history therewith (observables as opposed to invisible cognitive processes).[1]

Our data-driven approach to the creation and maintenance of client models naturally allows accommodation, integration, and adaptation to learned user preferences, observed long-term trends (such as recovery or disease progression), and event-triggered short-term phase changes (such as the temporary effects of a recently taken medication). Finally, while the purpose of APEX is not therapeutic, at times assistance and joint action may require APEX to motivate or gain compliance from the client. Behavior analysis provides an appropriate methodology and ethical framework for such manipulations.

As observers and aggregators of various forms of personal information (e.g., behavioral, medical), there are numerous privacy and security concerns with computer systems like APEX (e.g., safeguarding against accidental, illegal, or malicious compromise of data; means for individuals to exercise control over their personal data). Moreover, the autonomy imputed to intelligent agents brings with it issues of ethics and whether such agents are or ought to be ethically bound. For example, imagine that a client confides in APEX that he or she is contemplating

[1] This is not to say that we are abandoning main stream cognitive modeling; we are trying to forge a happy marriage of cognitive and behaviorist methods. Sustained discussion of the relative merits of each is beyond the scope of this paper.

suicide. Does APEX have a duty, ethically, legally, and/or morally (Wilks and Ballim 1990), to report the client's statement to others – and if so, is APEX liable if, in fact, it was only a case of gallows humor? Do we want APEX to have privileged confidentiality like an attorney or do we want it to be a mandatory reporter like other medical professionals? And how should privacy, data ownership, and ethical duties be weighed against equality in a convalescent or group home setting? Will APEX behave as a loyal friend, and be dedicated first and foremost to the client? Or will it appear as an agent representing the caregiver and physician? How are such dual purposes reconciled to establish the trust of all the players? Innovative home-based services like APEX provide a new impetus for academic and social discussion of such ethical concerns and risks, which are far from resolved.

Semantically Rich Human-Machine Interaction

The disabled and/or aging users whom we are targeting with APEX pose a variety of unique challenges for human-machine interaction. Cognitive disabilities may impair both interpretation and generation of language; physical disabilities may impair one or more signal channels, e.g., vision, speech, gesture; general aging may affect communication tempo and other attributes of interaction.

Our focus is on the use of multiple modalities for human-APEX interaction. In any given interactions, modalities will dynamically adjust in composition and manner of use (e.g., signals and protocols) based on context, client capability, and other communication exigencies of the moment (e.g., urgency to take medication on time).

Interaction must address the complexity of human-machine trust, especially when APEX must behave in a dominant manner to coach and guide behavior. Many veterans with TBI, for example, have a strong distrust of authority. Other users could simply fail to comply because they are skeptical of APEX's competence, or feel it is "hiding something." Compliance of APEX with the constraints and demands of human social interaction is paramount (Atkinson and Clark 2013).

To address the trust-related concerns as well as the possibilities of providing (1) a unique modality for non-verbal interaction and,(2) active physical assistance, we are investigating the use of humanoid robots as an in-home "avatar" for APEX. An embodied avatar, much like Embodied Conversational Agents (ECA), will evoke human social expectations and interaction very effectively (Schaefer, Billings and Hancock 2012) with quantifiable risks and benefits (Dorneich 2012). For a cogent review of research, application, and evaluation of embodied conversational agents, see Cassell (2000). Recent research projects in this domain include SEMAINE (Schroder, 2010), VHTookit (Hartholt et al. 2013), and Companions (Wilks et al. 2011).

In the near-term we are planning to investigate the proposition that a humanoid avatar, compared to a disembodied visual interface, will perform better in guiding client navigation in the home (e.g., to the medicine cabinet) and ensuring compliance with a pre-determined schedule of activities. A very small mobile robot would likely be sufficient for this purpose but would be incapable of performing physical labor.

A more robust robotic system would be capable of providing direct physical assistance to the client, such as helping in the kitchen, finding and providing the television remote control, assisting in standing, or dispensing medication (Figure 1). Mobility could be provided on the ground (e.g., wheels or legs) or via an overhead rail.

Figure 1: A Robot Assisting with Medication

A significant component of our rich human-machine interaction is that of automatic speech processing, with an emphasis on understanding of, and adaptation to, speech that is impaired. Borrowing from the field of machine translation (Dorr 1993), we adopt a paradigm in which the notion of divergence is central.

To illustrate the concept of divergence across languages, consider three properties (vocabulary, pronunciation, and syntactic structure) coupled with the differences across these for four languages: Spanish, Portuguese, English, and Chinese. We may consider a language to be similar to another language in "vocabulary" if there is a shared orthography, in "syntax" if the grammars are the same, and in "pronunciation" if they contain similar phonological forms. The most similar language pair of these four (aside from the language to itself) is Spanish-Portuguese, which shares all three features. The most radically divergence pair in is Spanish-Chinese, where there no similarities are associated with any of these three features.

We apply this same notion of divergence to the problem of "speech functioning," constraining our language pair to asymptomatic English speech compared to impaired English speech. In this case, the divergence properties to be studied are articulatory and disfluency patterns. We develop and apply techniques for detecting such divergences and leverage these to enable adaptive automatic speech recognition. The goal is to adapt to both deterioration and improvement in speech, within the same person, over time. For example, in Amyotrophic Lateral Scleroris, speech is likely to become more impaired, whereas with Traumatic Brain Injury, the speech is likely to become less impaired.

The closest speech processing study to the divergence approach described above is by Biadsy et al. (2011), who investigated the variation of speech properties under intoxicated and sober conditions. However, this earlier work was applied to the detection of intoxication (vs. sobriety), not the *degree* of intoxication. Rudzicz et al. (2014) employed another approach for recognizing impaired speech to answer a similar yes/no question (Alzheimer's vs. no Alzheimer's). Although the notion of "degree" was not the focus of these earlier studies, we leverage the incidental but significant discovery that pronunciation varies systematically within categories of speech impairment. This discovery is critical for correlating the divergence from a baseline English and provides a foundation for adapting speech recognition technology to varying degrees of impairment.

In the overarching APEX framework, the studies above are significantly enhanced beyond individual speech recognition experiments, in three ways:

- We benefit from the potential for embedding this technology into the three paradigms mentioned above (companions, humanoid avatar, and robotic systems) to enable conversations with a computer.
- We leverage the paradigms above to investigate interactive dialog that includes informal language understanding, in the face of disfluencies such as filled pauses (*uh*), repeated terms (*I-I-I know*), and repair terms (*she—I mean—he*).
- We are able to investigate pragmatic interpretation of language and action, thus undertaking intention recognition. Sensor input (visual, tactile, etc.) enables the understanding of utterances that are otherwise uninterpretable due to speech impairment, e.g., *Fill it with rockbee* may be understood with gesture toward a coffee cup may be understood as *Fill it with coffee*.

Reasoning from Comprehensive Multi-modal Behavioral Data

A major challenge for APEX, indeed for many adaptive intelligent agents that are focused on human-computer interaction, is collecting, aggregating and analyzing very large amounts of longitudinal data from heterogeneous sensors. APEX requires a symbolic, temporal representation of "now," that is, what happened previously that led to the present situation, and what is likely to happen in the future under various hypothetical conditions. From such a basis, APEX must maintain situational awareness, interpret behavior, and infer intent. This capability is a fundamental basis for real-time reasoning about human behavior in the context of environment dynamics. It provides essential support for decision-making and closed-loop physical automation. Finally, experiential knowledge is fodder for non-real time reflection that leads ultimately to the machine learning and adaptation we believe is required.

We will use automated coding of multi-modal behavioral data to achieve situational awareness of the client. The field of behavioral signal processing (BSP) has used automation to model abstract human behaviors in relevant, realistic scenarios, mitigating previous manual behavioral sciences coding schemes. An overview of automated methods that are maturing rapidly includes discussion of social cues, affect, and emotion (Black et. al. 2011).

An early application of BSP technology is currently fielded as a "Smart Home' by the Tampa VA Hospital Polytrauma Center (Jaziewicz et. al. 2011). This Smart Home continuously collects and analyzes client location and orientation data, as well as every interaction of clients with clinical and medical staff. Early data mining analyses using a BSP method called Fractal D have provided insight into gait and walking behavior (e.g., wandering) that would otherwise not have been detected or quantitatively documented if dependent on human observation alone (Kearns, Nams and Fozard 2010).

Our approach to level-one data processing of the multi-modal sensor data acquired by APEX will include an approach for well-structured behaviors (e.g., "sit down"). We will generate a probabilistic template built from training examples based on motion-capture data. Well-known algorithms such as stochastic context-free grammars can be used to make probabilistic matches to such templates using limited sensor data (Abowd et. al. 2002).

Level-two processing will assimilate these tokenized situational and behavioral data to dynamic world models that represent the evolution of situations, intentions, activities, and other elements of shared awareness (Atkinson, Clancey and Clark 2014). Previous studies have shown the viability of detecting clinically relevant changes in behavior using this type of longitudinal sensor data and activity recognition algorithms (Dawadi and Cook 2014).

Reasoning in APEX will be driven by goals that range from baseline policies that always constrain possibilities (e.g., keep the client safe) to goals that reflect physician general guidance extending over some period (e.g., take the medicine twice each day at mealtimes), to reactively generated goals that are a function of interaction with the client or care-giver, or other situational exigencies of the moment (e.g., the stove is left on after cooking is complete).

Conclusion

In the sections above, we have presented our research and vision for the Ambient Personal Environment eXperiment (APEX). This is an exciting new project that brings together disparate areas of artificial intelligence and promises to reveal new challenges for cyber-physical systems and robotics in the context of ambient intelligence applications.

Integrated system studies, such as APEX, require multidisciplinary contributions. If those studies are performed in a common experimental environment, such as the one we are constructing, it will facilitate both collaboration and technology integration, thereby increasing the chances for success in creating valuable system-level advances that address important individual and social needs.

Our system-level approach gives us the opportunity to investigate challenging use cases of interest to clients, clinicians, and caregivers. These include health monitoring, remote health care, support for rehabilitation and independent living, and systems that promote and help ensure health through ambient persuasive technologies.

The latter is a topic area fraught with ethical concerns (Verbeek 2009). Persuasive technology applied in health care requires especially careful consideration and discussion of methodological and ethical factors with respect to informed consent and privacy. The notion of *informed consent* is very important in any care system embedded in a society with strong legal constraints and recourse such as the US. Thus, it would be of great interest if an APEX-like agent could also elicit informed consent from clients after a process of explanation and conversation based on deep knowledge of them (Wilks and Ballim 1990). That would not only economize on expensive professional time, but would be a genuine cognitive advance into an area where an automaton was able to make an informed judgment about a client's mental state: that of understanding and consequent consent to procedures.

It is our hope that this research, and those of our colleagues working on ambient assisted living technology, will eventually help meet the needs of an overburdened health system and an aging society. That burden is severe. Of the two million soldiers who have served in Iraq and Afghanistan, as many as eight hundred thousand have suffered traumatic brain injuries that resulted in some level of cognitive impairment[2]. We seek to provide essential technology that allows these wounded veterans and others with special needs to remain in their homes and participate in society with a high quality of life.

Key contributions and points to remember are:

- Ambient intelligence (AmI) for people with cognitive, physical or age-related disabilities requires *extreme personalization*, adaptation significantly beyond the baseline of AmI for general users.
- Extremely personalized systems must be able to learn and adapt to a person's changing needs over a life-time of co-activity.
- Unique individual cognitive or physical disabilities require an AmI to interact flexibly with a client through multiple modalities whose needs cannot always be foreseen without actual experience.
- An AmI that provides support to a person with disabilities is both an actor-agent and, by virtue of substituting for lost abilities and augmenting others, a cyber-human prosthesis.
- APEX is our AmI laboratory for bringing together the multiple disciplines and technologies required in order to achieve this vision.

Acknowledgements

The authors gratefully acknowledge Adam Dalton, Ian Perera and their other colleagues at IHMC who contribute to the APEX project. The contribution by David Atkinson to this work was supported in part by the Air Force Office of Scientific Research under grant FA9550-12-1-0097.

References

Abowd, G., Bobic, A.F., Essa, I.A., Mynatt, E.D. and Rogers, W. 2002. The Aware Home: A living laboratory for technologies for successful aging. In *AAAI Technical Report WS-02-02*. Menlo Park, CA: AAAI Press.

Atkinson, David J. 2009. Robust Human-Machine Problem Solving. Final Report Contract Number FA2386-09-4005, Air Force Office of Scientific Research. Pensacola, FL: Institute for Human and Machine Cognition.

Atkinson, D.; Friedland, P.; and Lyons, J. 2012. Human-Machine Trust for Robust Autonomous Systems. In *Proc. of the 4th IEEE Workshop on Human-Agent-Robot Teamwork, in conjunction with the 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI 2012).* Boston, USA: ACM Press.

Atkinson, D.J. and Clark, M.H. 2013. Autonomous Agents and Human Interpersonal Trust: Can We Engineer a Human- Machine Social Interface for Trust? *In Trust and Autonomous Systems: Papers from the 2013 AAAI Spring Symposium. Technical Report No. SS-13-07.* Menlo Park, CA: AAAI Press.

Atkinson, D.J., Clancey, W.J. and Clark, M.H. 2014. Shared Awareness, Autonomy and Trust in Human-Robot Teamwork. In *Artificial Intelligence and Human-Computer Interaction: Papers from the 2014 AAAI Fall Symposium. Technical Report No. FS-14-01.* Menlo Park, CA: AAAI Press.

[2] Personal communication to author by a senior physician in the Tampa VA Polytrauma unit who is involved in VA planning for these veterans.

Bee, N., Andre, E. Vogt, T. and Gebhard, P. 2010. The Use of Affective and Attentive Cues in an Empathic Computer-Based Companion. Yorick Wilks, ed.. *Close Engagements with Artificial Companions: Key, Social, Psychological, Ethical and Design Issues*. 131-142. Benjamins Publishing Company.

Black, M.P., et al. 2011. Toward automating a human behavioral coding system for married couples' interactions using speech acoustic features. *Speech Communication*. doi:10.1016/j.specom.2011.12.003

Cassell, J. 2000. Nudge Nudge Wink Wink: Elements of Fact-to-Face Conversation for Embodied Conversational Agents. In *Embodied Conversational Agents.* Cassell, J., Sullivan, J., and Prevost, S. eds. Vol. 1. Cambridge, MA: MIT Press.

Cassell, J. 2001. More Than Just Another Pretty Face: Embodied Conversational Agents. *Communications of the ACM*, 43(4).

Cervantes, L., Lee, Y-S., Yang, H., Ko, S-h, and Lee, J. 2007. Agent-Based Intelligent Decision Support for the Home Healthcare Environment. M.S. Szczuka, et. al (Eds.). In *Proceedings of the 1st international conference on Advances in hybrid information technology*. LNAI Vol 4413 pp. 414-424. Berlin: Springer-Verlag.

Clancey, W.J., Sachs, P., Sierhuis, M., van Hoof, R. 1998. Brahms: Simulating Practice for Work Systems Design. *International Journal on Human-Computer Studies* 49:831–865.

Clancey, W.J. 2002. Simulating Activities: Relating Motives, Deliberation, and Attentive Coordination. *Cognitive Systems Research* 3(3):471–499.

Clancey, W.J., Sierhuis, M., Alena, R., Berrios, D., Dowding, J., Graham, J.S., Tyree, K.S., Hirsh, R.L., Garry, W.B., Semple, A., Buckingham Shum, S.J., Shadbolt, N. and Rupert, S. 2007. *Automating CapCom Using Mobile Agents and Robotic Assistants*. NASA Technical Publication 2007-214554. Washington, D.C.

Cooper, J.O., Heron, T.E. and Heward, W. L. 2007. *Applied Behavior Analysis*. 2nd Ed. Pearson Education, Inc.

Dawadi, P. and Cook, D.J. 2014. Smart Home-Based Longitudinal Functional Assessment. In *Proceedings of UBICOMP 2014*. 1217-1224. Seattle, WA: ACM.

Dorr, B.J. 1993. *Machine Translation: A View from the Lexicon*. Cambridge, MA: MIT Press.

Dorneich, M.C., Ververs, P.M., Mathan, S., Whitlow, S., and Hayes, C.C. 2012. Considering Etiquette in the Design of an Adaptive System. *Journal of Cognitive Engineering and Decision Making*. 6(2):243–265.

Hamilton, S. 2001. Thinking outside the box at the IHMC. *IEEE Computer*. 34(1):61-71.

Hartholt, A., Traum, D., Marsella, S. C., Shapiro, A., Stratou, G., Leuski, A., Morency, L.-P., and Gratch, J. 2013. All Together Now: Introducing the Virtual Human Toolkit. In *International Conference on Intelligent Virtual Humans* Edinburgh.

Isern, D., Sánchez, D. and Moreno, A. 2010. Agents applied in health care: A review. *I J of Med Informatics* 79(3):145-166.

Jasiewiccz, J., Kearns, W., et. al. 2011. Smart rehabilitation for the 21st century: The Tampa Smart Home for veterans with traumatic brain injury. Guest Editorial. *JRRD* 48(8).

Kaluža, B., Mirchevska, V., Dovgan, E., Luštrek, M. and Gams, M. 2010. An Agent-based Approach to Care in Independent Living. B. de Ruyter et. al (Eds.): *Ambient Intelligence, Lecture Notes in Computer Science.* 6439:177-168. Berlin: Springer-Verlag.

Kearns W.D., Nams V.O, Fozard J.L. 2010. Tortuosity in movement paths is related to cognitive impairment. Wireless fractal estimation in assisted living facility residents. *Methods Inf Med.* 49(6):592–98. [PMID:20213038] DOI:10.3414/ME09-01-0079

Kearns, W., Fozard. J., Webster, P., and J.M. Jasiewicz. 2014. Location aware smart watch to support ageing in place. *Gerontechnology*, 13(2):223.

Kiesler, S. 2005. Fostering common ground in human-robot interaction. In *Proceedings of IEEE International Workshop on Robot and Human Interactive Communication.* 729-734. Pittsburgh, PA: IEEE Press.

Lee, J.J. Modeling the Dynamics of Nonverbal Behavior on Interpersonal Trust for Human-Robot Interactions. *In Trust and Autonomous Systems: Papers from the 2013 AAAI Spring Symposium. Technical Report No. SS-13-07.* Menlo Park, CA: AAAI Press.

Luštrek, M., Kaluža, B., Dovgan, E., Pogoreic, B. and Gams, M. 2009. Behavior Analysis Based on Coordinates of Body Tags. In *Proceedings of the European Conference on Ambient Intelligence.* Pp. 14-23. Berlin: Springer-Verlag.

Matessa, M. et. al. 2008. An ACT-R Representation of Information Processing in Autism. In *Proceedings of the 30th Annual Conference of the Cognitive Science Society.* Love, B.C., McRae, K. and Sloutsky, V.M., Eds. 2168-2173. Austin, TX: Cognitive Science Society.

Mihailidis, A., Fernie G.F., Eng, P., Barbenel J.C., and Eng, C. 2001. The Use of Artificial Intelligence in the Design of an Intelligent Cognitive Orthosis for People with Dementia, Assistive Technology. *The Official Journal of RESNA*, 13(1):23-39. DOI: 10.1080/10400435.2001.10132031

Page, M.J. 2006. Methods of observation in mental health inpatient units. *Nursing Times.* 102(22).

Pérez, J.G., op den Akker, R. and Evers, V. 2013. Acceptable roboticS COMPanions for AgeiNg Years. Final Report Project 287625. EU Seventh Framework Program. University of Twente. Twente, NL.

Premack, D. and Woodruff, G. December 1978. Does the Chimpanzee have a Theory of Mind? *Behavioral and Brain Sciences* 1(4):515-526.

Schaefer, K.; Billings, D.; and Hancock, P. 2012. Robots vs. Machines: Identifying User Perceptions and Classifications. In Proc. of the IEEE Int. Multi-Disciplinary Conf. on Cognitive Methods in Situation Awareness and Decision Support, 168–171.

Schroder, M. 2010. The SEMAINE API: Towards a Standards-Based Framework for Building Emotion-Oriented Systems. *Advances in Human Computer Interaction*, 2(2).

Serna, A. Pigot, H., and Rialle, V. 2007. Modeling the Progression of Alzheimer's Disease for Cognitive Assistance in Smart Homes. *User Modeling and User-Adapted Interaction* 17(4):415-438.

Verbeek, P-P. 2009. Ambient Intelligence and Persuasive Technology: The Blurring Boundaries Between Human and Technology. *Nanoethics*. 3:231–242

Wilks, Y. and Ballim, A. 1989. Shifting the belief engine into higher gear. *Proc. of the Int. Conf. on AI Methodology Systems Applications.* Elsevier. Pp. 11–20.

Wilks, Y., and Ballim, A. 1990. Liability and Consent. In N. Narayanan and T. Bennun, eds. *Law, Computers and Artificial Intelligence*. Norwood, NJ: Ablex.

Wilks, Y. 2011. Protocols for Reference Sharing in a Belief Ascription Model of Communication. *In Proc. AAAI Fall Symposium on Advances in Cognitive Systems.* AAAI Press. Pp. 337–344.

Wilks, Y., Catizone, R., Worgan, S., Dingli, A., Moore, R., Field, D., and Cheng, W. 2011. A Prototype for a Conversational Companion for Reminiscing about Images. *Computer Speech and Language*. 25(2):140-157.

Wilks, Y., Jasiewicz, J., Catizone, R., Galescu, L., Martinez, K. and Rugs, D. 2014. CALONIS: An Artificial Companion within a Smart Home for the Care of Cognitively Impaired Patients. *In Proceedings of the 12th International Conference on Smart Homes and Health Telematics*. Denver Colorado.

Wilks, Y., ed. 2010. *Artificial Companions in Society: scientific, economic, psychological and philosophical perspectiv*es. John Benjamins: Amsterdam.

Wilks, Y. 2010. On being a Victorian Companion. *In Artificial Companions in Society: scientific, economic, psychological and philosophical perspectives*. John Benjamins: Amsterdam.

Wilks, Y. 2010. Artificial Companions. *In Artificial Companions in Society: scientific, economic, psychological and philosophical perspectives.* John Benjamins: Amsterdam.