

Understanding Considerate Systems - UCS (pronounced: You See Us)

Ted Selker
Carnegie Mellon Silicon Valley
Ted.selker@sv.cmu.edu

INVITED PAPER

ABSTRACT

Interactions between people occur in a social realm. On the other hand, “things”, including devices for communication and computation, are generally socially deficient. Imagine socially aware systems moving from an interruption model of communication to an introduction model. To create considerate systems, there is a need to model social context, social behavior, and communication goals.

This paper describes early systems that work to understand and eliminate the socially disruptive qualities of the ubiquitous systems people increasingly use and rely on in all aspects of their personal, educational, social and business lives. We show performance improving systems: an instant message arrives after you have finished typing a sentence, not while you are forming it; a car waits for you to complete a difficult maneuver before giving you distracting feedback

This work relies on dynamic task, user, system, and communication models. The goal is to stimulate more work to understand and create considerate systems. Such systems will improve people’s experience and performance..

Social responsiveness can become the norm for the technology that pervades our lives.

KEYWORDS: Ambient Computing, Context Aware Computing, Considerate, User Model, Task Model, System Model,

1. INTRODUCTION

Peoples’ collaborations are increasingly intertwined with ubiquitous computing and communicating systems. The difficulties of operating these systems cause people to open windows instead of adjusting digital thermostats, or crash their cars because they are talking on their cell

phone while negotiating traffic. While research on how to build devices that are easy to understand and to use has been somewhat successful in influencing current ubiquitous computing systems, such research has not taught us how to make them collaborate with us in a *considerate* way. Computing systems and devices are self-centered and not aware of the context surrounding our interaction with them. This results in inappropriate and insensitive responses such as self-important beeps, pop-up windows, cryptic references to technical actions and requirements, etc. For example, a system/device discovers it has a problem and immediately notifies the user at an annoying time, with a “battery low” or an “update now” message while the user might be giving a presentation. Typical response paradigms for systems ignore the fact that the human’s availability to be disrupted is determined by the current context, and therefore a socially appropriate interruption is desirable or required. In view of the technological understanding of being, and in turn allowing users to be effective, systems must treat their communications with people as social communication [1, 2]. Systems must be adaptive and recognize and learn appropriate times and approaches to communicate a request or provide other feedback. In effect, to be ubiquitously helpful, our devices need a social sense based on behavioral norms, the user’s style, and the current context. This paper advocates extending beyond context-aware and ubiquitous computing with explicit models of social response.

As a consequence of their ubiquity, mobile systems have introduced an epidemic of dangerous situations, both physical and cognitive, for their users. We no longer sequester ourselves in an office to communicate or to concentrate on work. Losing situational awareness while walking across a busy road or driving is dangerous; I have even witnessed someone being lulled into inappropriately answering a text message while playing an instrument on stage. Today, our ubiquitous systems can encourage distractions, often impairing our ability to attend appropriately to learning or other cognitive or social tasks.

Providing our devices with socially appropriate response capabilities has been a long felt need. The quintessential example of a blinking "12:00" on old recorders has not been completely eradicated, even as VCRs disappear from our homes. There are finally beginning to be some calls for considerate computing [3], which we are here to answer and amplify. This paper discusses systems that improve educational, work, and entertainment tasks by supporting people with considerate communication.

We need more work to create an understanding of how to improve productivity by giving social skills to the computing and communication systems people rely on. Several examples of intelligent user interface scenarios are included in this paper. These scenarios demonstrate the value of dynamic social models to understand and communicate. These systems demonstrate how "considerate systems" might integrate dynamic models of social behavior with models of the user's abilities and goals, of what the system can do, and the tasks the user is engaged in, to provide productive feedback and engagement. Inspired by cognitive science and experiments, the success of considerate systems to improve collaboration between people and computers will be judged by new and improved scenarios.

1.1 Early Work

Early computer models of social interaction were apparently successful at reducing human responses to a few simple reactions - it was a mirage. Eliza reduced Rogerian therapy to ~40 rules. It successfully poked fun at people and therapy and provided an eerily realistic façade of therapeutic intelligence, but Eliza was not useful or extensible to other uses. More recently, Microsoft's Bob had a matrix of personalities that it professed could be tuned to help the variety of users that might encounter it [4]. These personality models were completely underwhelming inside an eye-grabbing, distracting and inconsiderate system that seemed to rarely solve a problem. Our past and continuing research contains several demonstrated human performance-enhancing social interaction tools.

General models of considerateness and social behavior will allow us to create systems with a true ability to utilize social skills across a broad range of situations. They will avoid being grossly inappropriate by speaking at the wrong level or time, being condescending, self-important, redundant, or distracting. The focus of considerate systems is to operationalize an understanding of the complexity of communication in social systems. Since we

have entered an era where it seems that more communication happens via mobile devices than at a desk, we need to understand how to deal with complex communication requests in a variety of physical/social settings. The social mediation to communications with humans is generally tacit, but is not required for some communications, such as a fire alarm. However, without having a sense of appropriateness, devices cannot treat communications differentially. There are simple tools that allow a cell phone user to assign different ring tones to different callers, allowing the user to decide whether to interrupt his activity to take a call. Aside from their obvious limitations in scope and lack of situational flexibility, current approaches put the programming and management burden on the user. Armed with technology for considerate systems, we can demonstrate ways of reducing the complexity and difficulty of imbuing devices with the capability for appropriate communication in social settings.

This paper promotes a paradigm shift from *user-centric* to *social-centric* interaction models; the social situation will become the most important consideration in how a device delivers interruptions. Considerate systems will be useful for improving communication, education, activity management and self improvement. This work draws on perceptual and cognitive psychology, sociology, human-computer interaction, artificial intelligence, and machine learning.

1.2 Implications

This work strives to create a field that will further the value of social understanding in system design. Lack of social grace is a huge barrier in life. Many of us are familiar with stories about autistics who struggle with this, even if they are not cognitively impaired. The current generation of computer and communication system interfaces is similar to an autistic person in being competent at their tasks and incompetent at their social awareness. By improving their social skills, however, we can provide systems with the capability to interact more successfully. Considerate systems research strives to create the principles of designing socially aware systems, to create future generations of devices that will not be "autistic". Today people interact with hundreds of millions of portable computers and billions of cell phones. The result will be future devices that improve collaborations with people and systems.

2. CONSIDERATENESS

2.1 Issues of Interruption and Disruption

People can accomplish more when they are faced with fewer disruptions in their work [5, 6, 7]. Interrupting events might be off-topic or on-topic, from machines or from people. When an interrupting signal is received and changes the topic a person is attending to, we call it a disruption. In collaborations with machines, as in a human conversation, interruptions might in fact be relevant; an active listener often brings up examples that change analysis, improve communication, or simplify what must be said. Communication in which an active listener responds with relevant interruptions, should, in fact, increase flow instead of disrupt the activity.

Mediating interruption can be done at a surface level, with a cognitive model of user, task and system or with a deep understanding of conversation. Keystroke and mouse-movement models [8, 9] have been successful in recognizing when a person is active. Such a surface-level model of activity allows a system to recognize certain elements and times of physical unavailability. In Engagement Tracker, for example, the size of the image of a baseball video on a screen was reduced when a user was active in their web browsing activity; loud cheering on the baseball feed made the image larger. A surface-interaction subsystem could help to mediate some reactions on cell phones or other devices as well.

A “conversation” might be immersive enough that all information relates to it. In a car, for example, the primary activity being performed must be driving. The CarCoach system therefore made all interruptions from the audio or vibrating feedback systems relevant to driving. Such immersive situations simplify analysis of system, tasks and user priorities. Understanding the contextual constraints of communication will be central to creating considerate systems and we plan to investigate the automobile setting in detail for this work.

2.2 Elements of Social Response

People are social animals. When using electronic devices, people treat user interfaces as though they are other people [2]. In this way, all user interfaces include a model of social interaction. Mostly, user interfaces assume that they are the most important social entity (i.e. are selfish) and show it by making utterances that cannot be ignored by the person they are interacting with. They will put up a window or dialogue box that must be dealt with; they will make noise regardless of who is talking; they will type in the place a person is typing, etc.

Many user interfaces expect that all communication should be on their terms, requiring their users to learn or look up every term and reference used to communicate with them. They assume that the user understands all of

their utterances, and has a deep enough commitment to the machine to look up or remember out what each utterance means. These interfaces (or rather their designers) are wrong. Users abandon devices that are too difficult or annoying to use. New products sit on shelves, avoided by consumers who are skeptical of more devices which might require more effort to use than the utility they provide.

Many user interfaces are manipulative, requiring their users to perform actions to increase their control and use of resources. Printers, for example, that will work with no extra operating system drivers installed tend to cajole and threaten their owners to take time to install many new programs, some adding function, some replacing function of other manufacturers, some promoting the company and its products, and some requesting payment for service, and so on. The cynicism and graspingness of such interfaces slows down our computers and makes user experience an obstacle course of traps, interruptions, and hijackers.

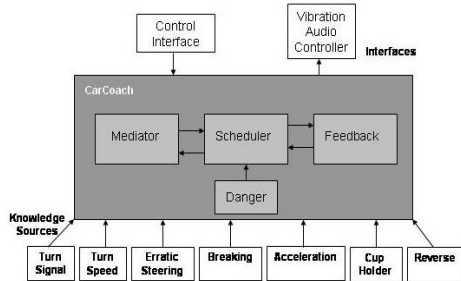
Many interfaces are arrogant, expecting that they need not know anything about their user and needn't change their understanding of the user. Such interfaces might require the same long process to be followed repeatedly; they require a user to only use the features supplied and can't be extended, customized or augmented. We view the value of dynamic social models as one part of the route towards creating considerate systems that are not self-centered, manipulative, or arrogant.

Without feedback, people often don't know if they are being understood. Considerate systems must integrate models of direct manipulation feedback, conversational feedback and background feedback. Direct manipulation feedback requirement varies depending on the modality and goal. The eyelid for example might react in 35 milliseconds [10], the eye/hand loop might need over 200 milliseconds [11], while some responses to warmth, smell etc. can be dramatically slower. It is also a fact that recalling a piece of information takes longer than recognizing it. What is a reasonable time for a considerate system to respond? Feedback is further complicated by other issues; we found that delaying feedback ½ second or more relative to turning performance in a car maneuver could improve learning [12] Figure 1. Was this cognitive load reduction, social turn taking, or something else? This and similar questions deserve more detailed exploration.

Social responses must be modeled in other ways as well. When someone is talking, it's okay to nod your head or say “um” or cough. It is less okay to state “WRONG!” or start whistling while another person is talking. Identifying social feedback that confirms communication rather than disrupts it is a key research area.

A shared language's job is to reduce communication cost. Considerate systems work focuses on explicit shared social language as a component of human system collaboration.

Affect has been widely studied [13]. Affective response might even have a deeper impact on communication than the feedback discussed above. When, how much, and what kind of affect will support the communication is important to the considerate system stance. Work such as CarCOACH car driving feedback system [12] has shown that affirmations make a huge difference to the impact of feedback. Excellent affective work is being pursued many places, but it needs to be more focused on scenarios and paradigms of response.



Imagine moving from a *speech recognition* scenario to a *listening* scenario. It would include feedback at a conversational pace such as the “uh huh”, “come again” and other indicators of people give to assure speakers they are being understood and should continue. A system implementing such a scenario would help communication with its operational model of social recognition and social behavior. The underlying social dynamics and their

Figure 1. CarCOACH architecture

A blackboard architecture for deciding how to give a variable schedule of feedback to a driver

and their associated surface language elements used to communicate should be further modeled; such work can be used to create reliable models of conversational turn taking and interruption in models of social response.

2.3 Language Based Considerate Response

Vocalized human communication, might continue to be our most comfortable and pervasive sophisticated way of broadcasting our thoughts. The accuracy of Automatic Speech Recognition systems has been significantly improved, thanks to increasingly sophisticated statistical models, larger training data and increased computer power. Recognition systems can become a key component

in ubiquitous computing and communication. However, in real settings, performance often falls well below their assessed accuracy. A few of the reasons that they fall short are simple: the speaker enunciates words with a stressed way, human or physical noises come up around the speaker, the speaker stops to listen to another person speaking, the speaker responds to another person verbally, the speech is speeding up, the speaker or speech recognition system loses their place and shows hesitation or confusion, etc.

Improved models of transcription, stress, and noise should improve accuracy. In addition, considerate systems' research can look for new social modeling approaches, including the use of back channels of communication to improve accuracy. To address interpretation accuracy and social alignment, human listeners give speakers feedback utterances. The dynamic positive feedback that speakers give each other in a conversation is the core of the fluid adjustments that people make to keep a dialog on-track. Today, speech recognition systems do not adequately look for opportunities to give such lightweight feedback. Considerate systems should work to anticipate as well as sense problems. From across the room a host might watch conversationalists struggling over distractions and decide to intervene to help them. Simple examples of how modeling and interaction might address the above to improve communication:

- Emotion detection based on acoustic-prosodic and lexical features can detect if a user is stressed [14] or angry [15; 16]. The system might play encouraging sounds to calm the speaker's enunciation.
- Speaking rate is an important feature and it affects the speech recognition accuracy. The system might sigh lightly or otherwise suggest slower speaking when it detects high speaking rates.
- By explicitly evaluating background noise volume and qualities, the system can estimate the kind of place and the characteristics it has. These can be used to evaluate appropriate feedback.
- The system might initiate turn-taking for another conversation with an outside speaker.
- The system might provide feedback when the system can't tell if the speaker stopped speaking in order to listen to another person.

- The system might provide feedback regarding volume, helping speakers know to readjust the microphone position.

Considerate systems can evaluate how limited semantic analysis can be useful for feeding back key words to demonstrate to a person that the system is on track. As humans might hear “speech” and repeat back “beach”, so as long as it’s not disruptive, the speaker will correct them with a chuckle and with a sense of teamwork. So a considerate speech recognition system should try to help the speaker get a feeling that they are being interpreted correctly. In addition to standard language models that only calculate probabilities of a transcription hypothesis based on surface form n-grams, we could calculate the likelihood that a sentence is semantically sound. Such work should use semantic relationships, such as those demonstrated in the Open Mind [17] project, to estimate the semantic coherence and tone of the recognized speech. This will help develop considerate speech recognition systems that interact with a user when the recognized speech is semantically incoherent.

Non-verbal cues through an understanding of the local context, task, and user models are also important. For example, knowing someone is driving and merging onto a freeway should elicit a different interpretation of their lack of immediate response to a query than if a dialog is taking place within a coffee shop. A conversation taking place within a crowded setting with variable background noise might adapt to the higher and changing level of distraction. If the considerate system determines that the user is momentarily stressed, distracted, or focused on dealing with an interruption, then it should modulate its response appropriately, rather than always becoming more insistent or “loud”. Experiments should be conducted to understand how social cues that can be determined via recognition systems can augment other available cues to improve on overall considerateness.

Another aspect of language-based considerate response is *considerate conversation moderating*. Conversations are social creations. They are produced one step at a time as people carry out certain joint activities. As a conversation develops, words from previous utterances serve as verbal context for the following sentences. Conversations also reflect joint activities. Every joint activity has public goals, or mutually agreed-upon purposes for carrying them out. Because of this agreed-upon purpose, sentences from two speakers in a conversation should mostly relate to one topic, even though perhaps one is in a car and one is in a restaurant hoping to order for both people. The two peoples’ utterances will reveal the topic of the conversation and enable prediction of aspects of what the other’ response might be or should not be. These

characteristics of a conversation theoretically enable tracking of some of the social and verbal context of the conversation from both speakers. Such work will help to make language-based interactions more considerate.

Presupposing complete modeling of a domain has been problematic. We have productively used common sense approaches that add some semantic analysis opportunistically [18]. OpenMind [17], WordNet [19] and HowNet [20] also define semantic relationships which will be used to extend the Context Inference Network learned from data. Such considerate language modeling will require careful investigation of the tradeoffs between the use of data-rich and knowledge-rich techniques in the proposed research.

The goal is to move from speech recognition to systems that actively listen the way a person does in a conversation. Considerate systems work is intended to create a dialog instead of a monolog. Appropriate use of social cues generated by recognition systems and the overall context, task and user model can give a person interacting with a considerate system a belief that they are being listened to. While today’s voice systems might give some feedback, they are typically heavy-handed and disruptive. A system that can modulate how much feedback and acknowledgement is appropriate should greatly improve current speech recognition, enabling considerate speech recognition, improving speech translation, enabling considerate conversation, moderating and improving spoken dialogue systems, and enabling social context-aware dialog understanding and management. Systematically exploiting a variety of social, situational and linguistic cues in systems should help them have flexible, dynamic and graceful responsiveness which will support much more natural and helpful interactions with users.

2.4 Learning and Reasoning Infrastructure

Simple demonstrations of context-aware systems such as the Smart Threshold [21] have relied on rule-based systems. More complex systems such as Cognitive Adaptive Computer Help (COACH) which is described more below have relied on blackboard and/or probabilistic systems. Recent systems, such as CarCOACH, Figure 1., integrated machine learning into sensor evaluation (classifying steering activity as straight, turns, drunk, sleepy, etc.) to feed a blackboard system creating action. By creating systems that incorporate one or more machine learning components, we can integrate data-driven approaches into systems with prior knowledge (rule-based systems or systems using common sense and statistical learning and Bayesian reasoning). The ContextBuilder

Figure 2 is a graphical interactive system for assigning code, rules or machine learned test data to inputs, and conflict resolution rules between them to drive outputs. It was created by Shawn Sullivan and demonstrated as an engine for the CarCOACH and the Smart Spoon. It allows a user to define blackboard control for creating context aware systems.

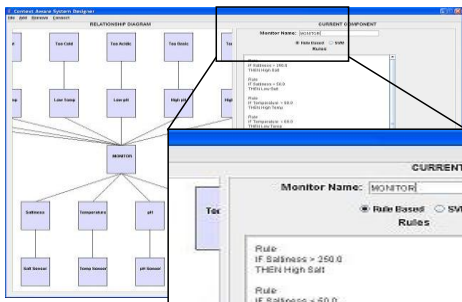


Figure 2. Context Builder

A graphical interface for creating context aware systems allows users to define inputs, a conflict resolving monitor and outputs. Each element can be conditioned with rules, results of a machine learning run on training data, or arbitrary code.

Considerate systems research must continue to create systems that use hybrid reasoning. We find that statistical machine learning is becoming especially helpful to evaluate alternative ways of representing knowledge sources in a blackboard system architecture.

Bayesian networks (BN) are also used in automated reasoning applications such as model-based diagnosis [22], medical diagnosis [23, 24], natural language understanding [25], probabilistic risk analysis [26], intelligent data analysis [27,28], and error correction coding [29,30]. Reflecting this wide range of application domains, there is also previous work on fusing heterogeneous digital data using BNs [31]. Research on Considerate Systems, too can utilize BNs to play an important role in modeling how context information (non-verbal social context) should be fused with verbal utterances, and for improving performance in view of the persistent noise, ambiguity and uncertainty in the real world.

2.5 Surface Social Models

A sense of social reaction that is recognized by the user of a system can be supplied by quite simple models. Created in 1966, Eliza [32] was the first known computer program to model conversational response. A second so-called “Chatterbot”, Parry was created by Colby in 1972 [33]. Eliza and Parry each modeled different roles in a therapeutic interaction. Parry, which modeled a paranoid schizophrenic, used a complex system of assumptions, attributions, and “emotional responses” triggered by shifting weights assigned to text input. Trained psychologists could not distinguish Parry’s transcripts from those of a paranoid human [34], making it the first computer program to pass the Turing test. Eliza, on the other hand, was based on a few dozen social reactions that a typical Rogerian therapist might utter [32]. Eliza was also very effective, and convincingly encouraged a feeling of therapeutic response in most users. The two systems actually had a conversation that brought forth both of their impressive response strengths. The impoverished but effective surface conversational response model of Eliza has been better remembered and celebrated than the much deeper personality model of Parry. It is remarkable for how a few encouraging phrases, used appropriately, can instill a sense of social responsiveness in a “patient”. People are strongly affected by simple surface social responses, which they interpret on a spectrum from rude to polite. A dog knows little of the semantics of human intellectual discourse, but understands how to use and respond to social surface cues that create positive affirmation or imply danger.

Simple speech acts such as “thank you”, “please”, “you’re welcome”, or “this is embarrassing” (the Firefox web browser response when it loses its way), can, when used appropriately, give a user interface a sense of social decorum which is effective although not backed up by a dynamic model of social response. Considerate systems can learn from Eliza-style surface social techniques based on static models of social dynamics and dynamic models utilizing an understanding of the task and user state to create enhanced system interactions.

Avatars also present a surface interaction to elicit social reactions that have been celebrated by some and reviled by others. The value of an embodied social agent in an avatar, when correctly done, can communicate emotions and even engender a sense of engagement. Indeed, an attentive looking avatar has been shown to help elicit a story from a child [35]. Microsoft’s Bob avatar, on the other hand, annoys users by attracting attention away from what a user is doing to tell them about their “problem” [4]. The surface nature of avatars or robots such as Valerie at CMU has often resulted in inadequate social models to support the affective response they try to portray [36]. Appropriately using them to present emotional reactions

will depend on reducing task requirements to simple acts such as eliciting more story as Cassel did, or be backed up by deep models of conversation such as Breazeal's Kismet, and Lockerd's social play systems [37, 38, 39].

The emotional language of facial communication is well studied [40]. Empathy Buddy is a system which used cartoon expressions to give an author of email feedback as to the tone of their messages [41], using Common Sense to evaluate the emotion of text in an email. When Empathy Buddy responded with words like "surprise", "sad", "happy", and "excited", people didn't notice; when, on the other hand these reactions were replaced with caricatures of facial expressions, people responded well, often re-editing their message. This example shows how choosing an appropriate surface social response (facial expressions in this case) can improve the communication ability of a user interface.

The emotional import of empathy, sadness, earnestness, etc., can overlay words and sounds. A variety of sonification researchers [42, 43] have sought to show the way sounds can connote information. Other work on pacing, prosody, intonation and tone has provided relevant results. Even simply listening to the tone of a conversation has been shown to be enough to evaluate the course and dynamics of the conversation; eg., who is in a power position, etc. These audio surface social cues are impressive and evident across different languages; dogs are excellent interpreters of voice tone. When and how audio feedback can improve considerate systems continues to be an area that can benefit from more work.

Simple movement of robotic toys, puppets and other animated things elicit social reaction. Robot body language is a developing field [37] - work of this kind and its impact on communication is well understood. Such social gestures as the other surface social cues, when used appropriately should improve social communication.

Imagery on a screen may fade, animate, move; words or sonification can indicate many things; devices might move as well. A taxonomy describing when and where such surface techniques will communicate well must be made for considerate systems to routinely use surface feedback techniques and cues that can be used to connote social response characteristics appropriate to the current state of an interaction.

2.6 Motivational and Captological Models

□ We meet people at a bar, a coffee shop or go for a walk to set the stage for a social experience. People use settings, rituals and other activities to encourage

themselves to reduce social distance and to find calm collaborative experiences. Such efforts can be important for fostering our productive work [44]. I created a digital cigarette project to demonstrate these ideas and show how interactive rituals can improve collaboration; the non-drug-delivery functions of a cigarette are used as a icebreaker, social lubricant and personal motivator .

In making the "digital cigarette" with no drug delivery, I worked with my student, Winslow Burleson, to encourage Philip Morris Corporation to study our drug-free digital motivational social lubricant concept product. This digital artifact used song, vibration and visual feedback to provide more satisfying breaks, to encourage people to keep working, to allow them to give each other an affirmation gift and to define themselves inside or outside a group with the same or different "digital cigarette" behavior. Many things operate similarly to such a motivational cigarette, an object that is evaluated by a person and people around that person for its social, and captological opportunities [45]. Considerate systems research will develop tools for supporting social contexts for motivation collaboration.

3. INTELLIGENT SYSTEMS

3.1 Adaptive Interfaces

The first interface that used an adaptive user model that quantifiably demonstrated improved user experience may have been the Cognitive Adaptive Computer Help (COACH) system [46]. COACH improved user performance by progressively presenting help information based on demonstrated user experience and expertise. In teaching LISP, users completed five times as many exercises when the same help as the control group could select was presented proactively by a learning and reasoning system. A small user expert rule set collected a user model. Another small teaching expert rule set paced the introduction of topics. Another small presentation expert rule set chose how and where to present information to the user. A blackboard conflict resolution approach chose which things to present and how to present them.

Other knowledge-driven tutoring research systems of the time integrated smart interaction into the content of the response or evaluated structural issues in LISP code [47,48]. In contrast, COACH separated content from explicit models of user, teaching, and presentation. These same models were shown to be useful in teaching SGML and UNIX commands and eventually found their way into the OS/2 operating system to help drive its Smart Guides OS help system.

Formatted: Bullets and Numbering

COACH abstracted models of teaching and presentation. Considerate systems will also utilize abstracted socially aware models that can be shown to create considerateness across a variety of systems. The use of multiple knowledge sources can be seen as a collaborative society of mind [49, 50] approach that can improve the resilience of considerate systems. We must work to understand a variety of knowledge sources and to show how systems can choose them appropriately for considerate reactions. These systems must use conflict resolution and knowledge to decide when to use various surface feedback schemes of communication, how to recognize activity and topicality, and how to know when and how a person should be interrupted. Considerate systems evaluate a person's goals so as to only interrupt at an appropriate time without distracting the person or others from their focus.

3.2 Ubiquitous and Ambient Computing

Calls for research on ubiquitous and ambient computing postulated that sensor and effector networks would be available in most natural settings [51,52], and they now are. Many large research, product, and infrastructure deployments are underway that all work on this convergence of our surroundings and technology [53]. The implications of sensors and effectors everywhere, all connected to the computing infrastructure, is an exciting backdrop which considerate systems can take advantage of and extend.

3.3 Affective Computing

How affect and stress are communicated via tone and other forms of expression has been a focus of work in psychology for some time. Research on affective computing has resulted in the ability to create new sensors that can recognize affective reaction on everything from a computer to a vending machine [54]. These affective sensor models recognize new kinds of useful information about people [13]. Research on the understanding of emotional status is making great strides. Some products are now emerging that take advantage of facial expressions and tone. Some voice response systems are already reputed to automatically transfer a caller to an operator if the caller yells and swears at the system. Considerate systems must work to take advantage of and extend affective computing as important in all social feedback.

3.4 Context-Aware Computing

Context-aware computing was envisioned as using a convergence of sensors to recognize and respond to specific situations [55, 56, 57]. Products such as the Onstar system introduced in 1995 that works with an automotive network, GPS and a communication network to aid drivers, show the beginning of the convergence of heterogeneous sensor-driven services for human needs [58]. Demonstration systems that use the convergence of data from heterogeneous sources [59] are leading the way. I have previously shown that sensors with model-driven "virtual sensors" can increase our likelihood of understanding what is valuable at specific points in space and time in this stream of information. For example, our work on the CarCOACH [60, 61], demonstrated that a system can recognize and encourage a person to reduce driving errors by simply watching the speed, steering wheel, gas, brake, cup holder and blinkers. By combining driver-control information with social information such as how often and recently the person has been chided, the system was able to demonstrate that more considerate feedback, in the form of a variable schedule of reinforcement, reduced the distraction of feedback comments over simply immediately telling a driver how they were doing as they drove.

Since the early 1990s many new examples of using dynamic models in interfaces have been produced. My Context-Aware Computing group at MIT, for example, produced dozens of research platforms to show that reasoning, representation and learning (AI) capabilities could robustly improve human performance in natural, sometimes even dangerous, settings. These demonstrations range from instant messaging to beds, to cars and to kitchens, and even to electronic cigarettes [62]. Considerate systems will require continuous work on scenario-based demonstrations using dynamic models of user, task and system, with explicit dynamic social models.

4. BEYOND USABLE SYSTEMS

As described above, the technologies of ubiquitous and ambient computing, affective computing, and context-aware computing are valuable for knowing the state of a situation and a person. Much of this work turned hinge on designing appropriate uses of techniques to create "user friendly" scenarios and then integrating these techniques in working technologies. This led us to realize that the most valuable parts of these systems often had to do with arranging its context to naturally focus a person on the things that would allow the person and the system to communicate simply. As an example, the Smart Bed [63] holds a person's head stable on a pillow as they focus on a

ceiling; if the music is annoying it will “cause” them to blink “nervously” which changes the music program. To drive the system in this scenario, we tell people to simply exaggerate what they normally do: staring when interested, gazing around when not, winking when they really like things and blinking when they don’t. The eye is a social communication tool; by playing to the visual statements a person naturally makes and creating a language that characterizes it, people were able to lie down in the bed and immediately begin using it competently. Our Smart Door [21] capitalized on the fact that the threshold is a natural social demarcation; it used knocking, speaking your name, and touching schedules as the interface. We focused on the way people expect to get hold of another person and designed the interface to take advantage of this [64]. Considerate systems strive to become systematic about creating an understanding of which surface communication efforts (such as the door being opened or knocked on) can be used and how to create a scenario which lives in a social space and works in expected and reliable ways.

Continuing considerate systems work will create a body of work which investigates the specific social issues of feedback in systems. All interactions with people occur in a social realm [2]. The right way of presenting feedback makes a difference. We have demonstrated an augmented reality kitchen in which text projected on appliances distracted people. However, people’s performance was improved with ecological feedback such as waves projected above the sink, fire above a burner and the sound of a cold wind when a freezer opened [65]. The surface characteristics of feedback presentation deeply affected its usefulness.

Systems have been made more usable with fewer steps per action, more understandable interfaces, better presentations of feedback and actions, etc. However, many of our frustrations come from the systems not changing their behavior when we are around others, frightened, angry, in a hurry or distracted. Considerate systems research focuses on creating systems that will work better with peoples’ social reactions. Considerate systems must embody feedback in operational models for teaching, comfort and captology to create a social response; they must choose how to give this feedback to match peoples abilities, expectations and to minimize disruption of human goals.

4.1 Evaluation

The value of social responses in improving human interactions with context-aware systems requires evaluation. Objective experiments comparing relevant

control systems can be implemented to demonstrate success of considerate systems.

Controlled experiments must be interpreted to understand the results in the appropriate context. Our Demonstrations by [8] show difficulties with using heat as a feedback. Still, most people enjoy the warmth of the touch of another’s hand. This example typifies the need to evaluate even such surface characteristics in light of new scenarios in order to understand the task and contextual constraints that determine their usefulness. Creating and testing scenarios is an important part of understanding the impact of appropriate and timely social feedback - and even for testing considerate system theories.

Disruption is a measure of how much interruptions affect focus. While the disruption work done for desktops computing is a good beginning [9, 66] new studies must continue to be performed for new scenarios such as language translation.

Electronic pets and video games have already touched the public’s fancy by including social elements as the focus of their appeal; the Tamagotchi, the Furby, Animal Crossing and the Sims, etc., offer social interactives that are part of the way many people collaborate, relax or reduce stress. Considerate systems will aim more at the personal support people can give themselves to be productive than the aspirational regimes of personal change that captology might work towards. Some of the most productive work will follow how using the notions of feedback and social interaction to give a person small rewards can help keep them in their preferred emotional state [44].

5.4 CONCLUSIONS

Considerate systems must hold up their side of an interaction with appropriate behavior. A considerate system must be able to present itself with the adequate social aspects for the stage of an interaction it finds itself in. Part of its response must layer the appropriate social reaction or coloring onto its responses so as to engage and not offend the person or people it is communicating with. We might give such systems the appropriate social behavior characteristics someone who knows you intimately, has your interests at heart and is selfless. We must start simply with use models of social interaction and research on social cues, where possible, to imbue considerate systems with a capability to socially interface with the user in a contextually-appropriate and useful fashion.

Considerate computing focuses on creating socially responsive systems. This starts from generalizing work

Formatted: Bullets and Numbering

such as context-aware considerate scenarios described above in this paper. Such work benefits from building on work in ambient intelligence, ubiquitous computing, context-aware computing, affective computing and captology.

This paper is a call for creating considerate systems as a field related to human collaboration. Operational social models for collaboration with systems will require integrating work in human computer interaction, social sciences, cognitive science, perceptual sciences, using scenario-based design, machine learning, and language-based understanding. Considerate theory, techniques and working scenarios can help build a world with fewer frustrations that is simpler to interpret and react to.

ACKNOWLEDGEMENTS

This paper is based on a proposal which Steven Rosenberg worked tirelessly to get me to write. His many edits eliminated hundreds of typos and think-os. The work was also edited by Ole Mengshoel, who brought ideas and edits to the learning aspect. It was also edited by Joy Zhang, who brought depth to the language considerations. Martin Griss encouraged the writing and made valued comments. Finally Ellen Shay looks at everything.

REFERENCES

- [1] M. Heidegger, "The Question Concerning Technology," BASIC WRITINGS, Ed. David Krel, HarperCollins, New York, NY, 1993.
- [2] B. Reeves & C. Nass, THE MEDIA EQUATION: HOW PEOPLE TREAT COMPUTERS, TELEVISION, AND NEW MEDIA LIKE REAL PEOPLE AND PLACES, Cambridge University Press, Cambridge, UK, 1996.
- [3] W. Gibbs, "Considerate Computing," *Scientific American*, Vol. 55, pp. 55-61, 2004.
- [4] M. Newman, "Bob is dead; long live Bob- Failed Microsoft interface may have been well ahead of its time," *Pittsburgh Post-Gazette*, Sunday, May 23, 1999.
- [5] T. Gillie, and D. Broadbent. "What makes Interruptions Disruptive? A Study of Length, Similarity and Complexity," *Psychological Research*, Vol.50, pp. 43-250, 1989.
- [6] E. Cutrell, M. Czerwinski, and E. Horvitz. "Notification, Disruption and Memory: Effects of Messaging Interruptions on Memory and Performance," In Hirose, Michitaka (Ed.), *Human-Computer Interaction INTERACT '01*, (Tokyo, July 9-13), IOS Press (for IFIP), pp. 263-269, 2000.
- [7] E. Arroyo, T. Selker, and W. Wei, "Usability Tool for Analysis of Web Designs Using Mouse Tracks," *CHI '06 Extended Abstracts on Human Factors in Computing Systems*, Montreal, Canada, pp. 484-489, April 2006.
- [8] E. Arroyo, and T. Selker, "Self-adaptive multimodal-interruption interfaces," *Proceedings International Conference on Intelligent User Interfaces (IUI'03)*, 2003.
- [9] E. Horvitz, C. Kadie, T. Paek and D. Hovel, "Models of attention in computing and communication: From principles to applications," *Communications of ACM*, Vol.46, No. 3, pp. 52-59, 2003.
- [10] A. Glorig, "The Eyeblick Response As A Test For Hearing," *Journal of Speech and Hearing Disorders*, Vol.18, No. 4, pp. 373-378, 1953.
- [11] "The Human Benchmark," [website], March 3, 2010, Available:<http://www.humanbenchmark.com/tests/reactiontime/index.php>.
- [12] T. Sharon, "An Advanced Driver Warning Framework Incorporating Educational Warning. MS thesis, MIT Media Lab, Cambridge, MA, May 2003.
- [13] R. Picard, AFFECTIVE COMPUTING, MIT Press, Cambridge, MA, 1997.
- [14] S. Scherer, H. Hofmann, and G. Palm, "Emotion Recognition From Speech: Stress Experiment," *Proceedings of the Sixth International Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, <http://www.lrecconfor.org/proceedings/lrec2008/>, 2008.
- [15] D. J. Litman, and K. Forbes-Riley, "Predicting student emotions in computer-human tutoring dialogues," *ACL '04: Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, Association for Computational Linguistics, Morristown, NJ, USA, p. 351, 2004.
- [16] F. Burkhardt, M. Van Ballegooy, K-P. Engelbrecht, T. Polzehl, and J. Stegmann, "Emotion Detection in dialog Systems: Applications, Strategies and Challenges," *Proceedings of International Conference on Affective Computing & Intelligent Interaction*, 2009.
- [17] H. Liu and P. Singh, "ConceptNet: A Practical Commonsense Reasoning Toolkit," *BT Technology Journal*, Vol. 22, Kluwer Academic Publishers, 2004.
- [18] E. Arroyo, "Implicit Metrics of Activity and Attention," PhD thesis, MIT Media Lab, Cambridge, MA, May 2007.
- [19] C. Fellbaum, ed, WORDNET: AN ELECTRONIC LEXICAL DATABASE, The MIT Press, Cambridge, MA, 1998.

- [20] Dong, Z. and Q. Dong., HOWNET AND THE COMPUTATION OF MEANING, World Scientific Publishing Co. Inc., River Edge, NJ, 2006.
- [21] H. Yan, T. Selker, "Context-Aware Office Assistant," *Proceedings of the 5th International Conference on Intelligent User Interfaces*, New Orleans, LA, pp. 276-279, January 2000.
- [22] U. Lerner, R. Parr, D. Koller, and G. Biswas, "Bayesian fault detection and diagnosis in dynamic systems," *Proceedings of the 17th national Conference on Artificial Intelligence (AAAI-00)*, pp. 531-537, 2000.
- [23] S. Andreassen, M. Woldbye, B. Falck, and S.K. Andersen. "MUNIN - A causal probabilistic network for interpretation of electromyographic findings," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pp. 366-372, Milan, Italy, August 1987.
- [24] M.A. Shwe, B. Middleton, D.E. Heckerman, M. Henrion, E.J. Horvitz, H.P. Lehmann, and G.F. Cooper, "Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base: I. The probabilistic model and inference algorithms," *Methods of Information in Medicine*, Vol. 30, No. 4, pp. 241-255, 1991.
- [25] N. Chater and C. D. Manning. "Probabilistic models of language processing and acquisition," *Trends in Cognitive Sciences*, Vol. 10, No. 7, pp. 335-344, 2006.
- [26] H. Langseth and L. Portinale. "Bayesian networks in reliability," *Reliability Engineering and System Safety*, Vol. 92, No.1, pp. 92-108, 2007.
- [27] P. Jones, C. Hayes, D. Wilkins, R. Bargar, J. Sniezek, P. Asaro, O. J. Mengshoel, D. Kessler, M. Lucenti, I. Choi, N. Tu, and J. Schlabach, "CoRAVEN: Modeling and design of a multimedia intelligent infrastructure for collaborative intelligence analysis," *Proceedings of the International Conference on Systems, Man, and Cybernetics*, pp. 914-919, San Diego, CA, October, 1998.
- [28] C. C. Ruokangas and O. J. Mengshoel. "Information Filtering using Bayesian networks: Effective interfaces for aviation weather data," *Proceedings of the 2003 International Conference on Intelligent User Interfaces*, pp. 280-283, Miami, FL, 2003.
- [29] R. G. Gallager, "Low density parity check codes," *IRE Transactions on Information Theory*, Vol. 8, pp. 21-28, Jan 1962.
- [30] D. J. C. MacKay, "Information Theory, Inference and Learning Algorithms" Cambridge University Press, Cambridge, UK, 2002.
- [31] K. J.Ezawa and T. Schuermann. "Fraud/uncollectible debt detection using a Bayesian network based learning system: A rare binary outcome with mixed data structures," *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 157-166, Montreal, Canada, 1995.
- [32] J. Weizenbaum. "ELIZA - A Computer Program For the Study of Natural Language Communication Between Man And Machine," *Communications of the ACM*, Vol. 9, No.1, pp. 36-45, January 1966.
- [33] R. Schank and K. Colby. "Computer Models of Thought and Language," W. H. Freeman, New York, NY, 1973.
- [34] C. Teuscher, ALAN TURING: LIFE AND LEGACY OF A GREAT THINKER, Springer, Heidelberg, Germany, 2006.
- [35] J. Cassell, "More than Just Another Pretty Face: Embodied Conversational Interface Agents," *Communications of the ACM*, Vol.43, No. 4, pp. 70-78, 2000.
- [36] R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, Al. Schultz, J. Wang. "Designing Robots for Long-Term Social Interaction," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005.
- [37] C. L. Breazeal, DESIGNING SOCIABLE ROBOTS, The MIT Press, Cambridge, MA, 2002.
- [38] A. Brooks, J. Gray, G. Hoffman, A. Lockerd (Thomaz), H. Lee, and C. Breazeal, "Robot's play: interactive games with sociable machines," *Proceedings of the International Conference on Advances in Computer Entertainment (ACE)*, June 2004.
- [39] L. Thomaz, "Socially Guided Machine Learning," MIT Media Lab PhD dissertation, June 2006.
- [40] P. Ekman, "Basic Emotion," T. Dalgleish and M. Power (Eds.), HANDBOOK OF COGNITION AND EMOTION., John Wiley & Sons Ltd, Sussex, UK, 1999.
- [41] H. Liu, H. Lieberman, and T. Selker. "A Model of Textual Affect Sensing Using Real-World Knowledge," *Proceedings of the 8th International Conference on Intelligent User Interfaces*, Miami Beach, FL, pp. 125-132, January 2003.
- [42] W. Gaver. "The SonicFinder: An Interface that Uses Auditory Icons," *Human-Computer Interactions*, Vol. 4, No. 1, 1989.
- [43] M. Blattner, D. Sumikawa, and R. Greenberg. "Earcons and icons: Their structure and common design principles," *Human Computer Interaction*, Vol. 4, No. 1, pp. 11-44, 1989.
- [44] T. Selker, "Fostering Motivation and Creativity for Computer Users, Special Creativity and Computation Support.," *International Journal of Human-Computer Studies*, Vol. 63, No. 4-5, pp. 410-421, October 2005.

- [45] B. J. Fogg. PERSUASIVE TECHNOLOGY: USING COMPUTERS TO CHANGE WHAT WE THINK AND DO, Morgan Kaufmann, San Francisco, CA, 2003.
- [46] T. Selker. "Cognitive Adaptive Computer Help (COACH), A Case Study," chapter in ADVANCES IN COMPUTERS, Vol. 47, M. Zelkowitz, Editor, Elsevier Press, Maryland Heights, MO, pp. 69-137, 1998.
- [47] J. Anderson, F. Conrad, A. Corbett, "Skill acquisition and the Lisp tutor," *Cognitive Science*, Vol.13, pp. 467-505, 1989.
- [48] G. Fischer, A. Lemke, T. Mastaglio, A. Morch, "Using critics to empower users," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Empowering People*, pp. 337-347, Seattle, Washington, 1990.
- [49] M. Minsky, THE SOCIETY OF MIND, Simon & Schuster, New York, NY, 1988.
- [50] L.D. Erman, and V. R. Lesser, "A multi-level organization for problem-solving using many diverse cooperating sources of knowledge," *International Joint Conferences of Artificial Intelligence*, Vol. 4, pp. 483-490, 1975.
- [51] M. Weiser, "The Computer for the 21st Century," *Scientific American Special Issue on Communication, Computers and Networks*, September, 1991.
- [52] P. P. Maglio, and C.S. Campbell, "Tradeoffs in displaying peripheral information," *Proceedings of the CHI 2000 Conference on Human Factors in Computing Systems*, April 1-6, The Hague, The Netherlands. ACM Press. pp. 241-248, 2000.
- [53] N. Streitz and R. Wichert, "Road-Mapping Research in Ambient Computing and Communication Environments," *Towards The Human City, EU Interlink International Cooperation Activities in Future and Emerging ICTs*, December 1, 2009.
- [54] C. J. Lee, C. I. Jang, T. D. Chen, J. Wetzel, Y. B. Shen, T. Selker, E. Arroyo, W. Wei, "Attention Meter: A Vision-Based Input Toolkit for Interaction Designers," *CHI '06 Extended Abstracts on Human Factors in Computing Systems, Montreal, Canada*, pp. 1007-1012, April 2006.
- [55] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, P. Steggles, "Towards a Better Understanding of Context and Context-Awareness," *HUC 1999*, pp. 304-307.
- [56] H. Lieberman, T. Selker. "Out of Context: Computer Systems That Adapt to, and Learn from Context," *IBM Systems Journal*, Vol.39, No.3-4, pp. 617-632, October 2000.
- [57] T. Selker, W. Bursleson. "Context-Aware Design and Interaction in Computer Systems," *IBM Systems Journal*, Vol. 39, No.3-4, pp. 880-891, October 2000.
- [58] "Wikipedia," Onstar, [website], March 3, 2010, Available: <http://en.wikipedia.org/wiki/OnStar>.
- [59] M. Yin, M. Griss, SCATEAGENT: CONTEXT AWARE SOFTWARE AGENTS FOR MULTI-MODAL TRAVEL; APPLICATIONS OF AGENT APPLICATIONS OF AGENT TECHNOLOGY IN TRAFFIC AND TRANSPORTATION, pp. 69-84, Birkhäuser, Basel, Switzerland, 2005.
- [60] T. Sharon, T. Selker, L. Wagner and A. J. Frank, . in *Proc. of the IEEE Int. Conf. on Software - Science, Technology & Engineering*, Washington, DC, 2005, pp. 13-22.
- [61] E. Arroyo, S. Sullivan, T. Selker, "CarCoach: a Polite and Effective Driving Coach," *CHI '06 Extended Abstracts on Human Factors in Computing Systems, Montreal, Canada*, pp. 357-362, April 2006.
- [62] T. Selker, "Context-Aware," [website], March 3, 2010, Available: <http://context.media.mit.edu/press/index.php/projects/>.
- [63] T. Selker, J. Scott, and W. Bursleson, "A Test-Bed for Intelligent Eye Research," *LREC 2002, Workshop on Multi-Modal Resources and Multi-Modal System Evaluation, Canary Islands, Spain*, pp. 78-83, June 2002.
- [64] Y. Hao, T. Selker, "Context-aware office assistant," *Proceedings of the 5th international Conference on Intelligent User Interfaces New Orleans, LA*, , pp. 276-279, January 09-12, 2000.
- [65] L. Bonanni, C. Lee, T. Selker, "Attention-Based Design of Augmented Reality Interfaces "(poster),. *CHI '05 Extended Abstracts on Human Factors in Computing Systems, Portland, OR*, pp. 1228-1231, April 2005.