Agent-based support for behavior change

van Wissen, A.

2014

document version
Publisher's PDF, also known as Version of record

Link to publication in VU Research Portal

citation for published version (APA)
van Wissen, A. (2014). *Agent-based support for behavior change: Models and applications in health and safety domains*.

1

Introduction

This dissertation investigates how agent-based models can be used to (i) gain insight into processes of human behavior and behavior change, and (ii) develop systems that can effectively support behavior change. The research carried out to address this cuts across many disciplines, as there are many determinants for human behavior, and many aspects of human-computer interaction need to be considered when developing support systems. Rigorous study of all of these facets of behavior is even beyond the scope of 4 years of committed PhD research. This chapter is dedicated to introduce those particular determinants of behavior (change) and aspects of support systems that are the focus of this dissertation. In Section 1.1 it is explained how they relate to each other and why they are important for understanding human behavior and designing effective interventions. In Section 1.2 the research questions addressed in this dissertation are presented. A description of the methodology used to answer those questions can be found in Section 1.3. This section is followed by an outline of the dissertation in Section 1.5, where it is also described how the chapters correspond to the research questions. This introduction concludes with Section 1.4, in which the contributions of this dissertation to the research fields of intelligent support systems, crisis management, and eHealth are discussed.

1.1 Motivation

Human decision making is one of the most thoroughly researched subjects in the history of science. Even so, this has not diminished the number of challenges and unanswered questions that surround it. Decision making is notoriously complex as it involves many different facets, of which uncertainty, priming, option generation, norms, and evaluation are only a few. The reason why it has been a popular subject of many research papers, is because people are driven to improve their *understanding* of why people make particular choices and perform certain behaviors, and why they change.

There is no denying that as humans our decision making sometimes could use a little help. Human decision making relies on heuristics, which allows us to act quickly and proactively, to predict, adapt and act appropriately in different contexts. Such heuristics can help us when we have to deal with incomplete information, rapidly changing environments, and uncertainty about the effects of our behaviors. However, heuristics are fallible and we often lack good coping strategies to respond to unexpected situations. Furthermore, biases influence our choices. Even when we do rationally know what is the right or appropriate choice, that does not automatically mean that we choose it. Sometimes, we just don't feel like it or don't feel up to it. And change is hard.

“‘Most of the time what we do is what we do most of the time. Sometimes we do something new.” (Townsend and Bever, 2001, p.2) Our everyday behavior is determined for a significant part by our habits: sequences of behavior performed at particular times or places. Frequently, the action sequences are triggered when particular temporal or spacial conditions (cues) in context are met. Habits structure the way in which we interpret our experiences and influence when our attention is activated. They are characterized by a high degree of automaticity, which involves the direct association between context and response, but they can also interface with goals during learning and performance (Wood and Neal, 2007).

Once you become aware of such automatic responses, you can consciously try to learn new responses in accordance with your goal to replace the old ones. This can then result in behavior change. Imagine someone trying to lose weight. Many people automatically start eating when you put a bowl of peanuts in front of them at a party, regardless of whether they are hungry or whether they are particularly fond of peanuts. Being aware that you are one of those people automatically reaching for the peanuts can be a first step towards changing that behavior. Yet change often happens unconsciously, for example by the influence of others. If others are refraining themselves from eating, or choose the bite-size tomatoes over the peanuts each time, you may be triggered by that behavior to do the same.

Adopting new behavior comes with its own challenges. As put by Baumeister: “The ability to alter one’s own responses is one of the most important features of the human psyche and is substantially responsible for the immense range and diversity of human behavior as well as for the adaptive success of our species.” Yet coping with stress, regulating negative affect, and resisting temptations require *self-control*: the capacity to resist temptations (Baumeister, 2002). Constantly and consciously monitoring your behavior, as is required to stop eating the peanuts, is difficult. Furthermore, one can have conflicting goals (trying to lose weight versus satisfying a craving) or diminished resources to resist temptations. *Self-efficacy*, the measure of belief that one is able to act in accordance with one’s goals or tasks, can also greatly influence behavior change (Bandura, 1997). Adopting a new behavior and persisting in performing it requires confidence in one’s ability to do so.

Thus far, several prominent processes in behavior change have been brought up. Understanding such processes, knowing the determinants for human behavior, and ultimately being able to accurately predict that behavior yields great power. For with accurate prediction also comes the possibility of *creating interventions*: influencing attitudes or the decision process such that the outcome will be a different one, effectively establishing a *behavioral change*. Ideally, interventions that aim at behavior change will result in changes for the best. That is, after an intervention, the person chooses a ‘better’ option, an ‘optimized’ one, or the ‘right’ decision over a ‘wrong’ one. Creating interventions is essentially a manipulation of choice: someone is steered away from the option of executing one behavior in favor of another behavior. Yet people don’t always agree about the extent to which an outcome is good. Intervention designers often have their own agenda. Think of a doctor who uses behavioral insights in presenting medical treatment options to a patient. The same argument can be used to steer the patients towards two different outcomes. The doctor can choose to say that a procedure has an 80 percent chance that 90 of 100 patients will survive. On the other hand, the doctor can phrase it differently by saying that it comes with an 80 percent risk that 10 out of 100 patients will die. Although the options — and the probabilities of survival — remain the same, the way in which the choice is presented influences the outcome (see also Thaler and Sunstein (2008)). Consequently, when it comes to behavior change support, it is important to have a supporter who has your interest at heart. Someone who knows your motivations and your goals and can help you to stay on the right track. In short, a supporter who is on your team.

1.1.1 Agent-Based Support Systems

Intelligent agents are great candidates to be the supporters on your team. The term *intelligent agent* refers to a functional system that has at least some degree of autonomy, uses sensors to perceive the environment, interacts with that environment (including other agents and humans in it), can adapt to change and can act pro-actively (e.g., Wooldridge and Jennings (1995)). Importantly, agents are also capable of “taking on another’s goals” and can direct their activity towards achieving those goals (Russell and Norvig, 2003). Some intelligent agents are specifically designed to be part of larger electronic systems able to observe human activities and situational contexts, and to interact with them in a nonintrusive manner. Such agents are called *ambient agents*, and the environments in which they are embedded contain *ambient intelligence* (Aarts and de Ruyter, 2009).

The capabilities of intelligent agents enable them to be much more than mere tools. When introducing intelligent agents, Russell and Norvig originally defined them as being *rational*: for each possible percept sequence, the agent should select an action that is expected to maximize its performance measure, given the evidence provided by the history of everything the agent has perceived and whatever built-in knowledge the agent has. Such an agent thus always chooses to perform the action with the optimal expected outcome for itself from among all feasible actions. In many cases however, rational agents behave in manners that are counter-intuitive to people. In order to create more realistic representations of human behavior, there has been a growing interest in creating agents that are able to display more forms of human behaviors, for example showing behavior that is caused by inadequate use of reason or emotional distress. The study and development of such systems is called *affective computing* and reflects the the importance of emotions in human decision making (Picard, 2000). *Affective agents* are capable of expressing and acting on emotions, and to identify emotions of others. Such agents are for example deployed in virtual training environments, to simulate (hostile) soldiers and the emergence of riots or other collective behavior (van Diggelen, Muller, and van den Bosch, 2010; Silverman, Pietrocola, Nye, Weyer, Osin, Johnson, and Weaver, 2012).

With the development of affective agents that can express, for example, empathy, the view of agents as *social actors* instead of tools has been given more attention. According to Fogg, Cuellar, and Danielson, agents that function as social actors can persuade people to change their attitudes and behaviors by providing social support, reacting to attitudes or behaviors, and leveraging social rules and dynamics (Fogg et al., 2009). Such agents can be used to create a Behavior Change Support System (BCSS), which is “an information system designed to form, alter or reinforce attitudes, behaviors or an act of complying without using deception, coercion or inducements” (Oinas-Kukkonen, 2010, p.6). These computerized systems essentially function as a coach, and are also referred to as *e-coaches* (Warner, 2012). Such persuasive systems aim to motivate, facilitate and maintain change. This aim to maintain the change is an important aspect of behavior change support systems, which often focus on long-term behavior rather than single instances. Although the lists of required capabilities for such support systems vary, there are three criteria that are commonly accepted with respect to their design. First, the agent must be able to deal with context and timing. Second, the agent must be able to adapt to the user and act proactively. And third, the agent should use forms of personalization and tailoring in communication with the user (Aarts and de Ruyter, 2009; Aarts, Harwig, and Schuurmans, 2001). The following subsections describe these three criteria in more detail.

1.1.1.1 Adaptive and proactive strategies

To support behavior change, an agent has to be able to anticipate actions of a human, by predicting his or her behavior and the effects thereof. Behavior change is a long-term process that involves attitude change, learning new behaviors and repetition. If a support agent is able to predict relapse or a change of attitude, motivation or commitment, this information can help it to pick the most appropriate type of intervention. For example, focusing on the dangers of skipping medication might be a good strategy if the agent suspects a relapse in adherence (for example because the patient is on holiday). On the other hand, if the agent predicts that the person will be able to keep up his or her behavior change, then it might not be necessary to provide any reminders or suggestions. In fact, it could even prove counterproductive to intervene at this point, as the message might trigger re-evaluation or plain annoyance. In short, the agent strategy should adapt to observed and predicted behavior, and if neccessary let the agent act proactively.

1.1.1.2 Personalization and tailoring

For successful interventions, an assessment of the person's character and current condition (e.g., is he/she hungry, angry, sad, tired?) is also essential. Although character or personality can be assumed more or less stable over a period of time (Costa and McCrae, 1994), a person's condition can fluctuate heavily and depends on cognitive, physical and emotional dynamics. Mood, beliefs, desires, stress and physical discomfort are a few of the factors that can affect behavior and they can change rapidly over time. Much work has been done on trying to enrich ambient agents with the capabilities to sense human states (one can think of biophysical measurements, face and voice recognition, heart rate monitors), yet it remains a complex challenge. The purpose of such measurements is to be able to deliver personalized and tailored information that targets those factors that enable the person to change. *Tailored* materials are "materials that are intended to reach one specific person, are based on characteristics that are unique to that person, are related to the outcome of interest, and have been derived from an individual assessment" (Kreuter, Stretcher, and Glassman, 1999, p.276).[1] Many psychological studies have studied the effect of tailored information in comparison to more general information. The evidence suggests that tailored messages are likely to produce better results in terms of behavior change (e.g., Brug, Campbell, and van Assema (1999); Noar, Benac, and Harris (2007); Revere and Dunbar (2001); Skinner, Campbell, Rimer, Curry, and Prochaska (1999)).

1.1.1.3 Context and timing

Related to personalization and tailoring is the ability to be aware of and be able to respond to the current context, including location, time and the social environment. With respect to stimulating behavior change, timing is an essential aspect. The suggestion to eat a salad for lunch when you just finished eating at your favorite burger joint is less effective than receiving that same suggestion at 11 AM that morning, or having received it yesterday when you were grocery shopping. If someone is meeting up with friends who smoke, it might be more important

[1] As pointed out by Kreuter et al., "[t]he distinctions between tailored, targeted, personalized, andother forms of health communication are important ones, yet the terms have sometimes been used interchangeably in the research literature." (Kreuter et al., 1999, p.276). In contrast to Kreuter et al., Oinas-Kukkonen and Harjumaa argue that tailored material is material that "fits the potential needs, interests, personality, usage context, or other factors relevant to a user group" (Oinas-Kukkonen and Harjumaa, 2008, p.170). As such, he argues that tailored material should be regarded as Kreuter et al. would *targeted* material: "[material] intended to reach some specific subgroup of the general population, usually based on one or more demographic characteristics shared by its members" (Kreuter et al., 1999, p.276). In this work, the definition of tailoring suggested by Kreuter et al. is used.

to remind him of his intentions not to smoke than when he is visiting his non-smoking parents. Giving appropriate suggestions is about *kairos*: “providing the right information at the best time” (Andrew, Borriello, and Fogarty, 2007, p. 259). For agents to be able to initiate interventions, give relevant information and suggest appropriate actions, they have to get the timing right.

1.1.2 Social Context

Focusing on individual processes of decision making is often not sufficient to change behavior, as many cognitive processes and behavioral actions are the result of complex and repeated interactions between multiple individuals and their environment over time. Humans are easily influenced by actions and statements of others (Barsade, 2002; Loewenstein, 1989). Collective decisions can deviate much from those that would be made individually, as our actions and cognitive states are influenced by what other people do, wear or say. As such, social context is an important determinant for individual decision making. Interaction with others is crucial for learning and adopting new behaviors. Several social principles can be distinguished that play a role in behavior change, among them principles of social learning, social comparison, social facilitation and competition. Unfortunately, this also means that we can learn and adopt the mistakes others are making. Think for example of peer pressure of friends to take up smoking. Another example comes from studies that have suggested that obesity may spread in social networks through social ties, where social distance appears to be more important than geographic distance (Christakis and Fowler, 2007).

Diffusion models seek to explain how and when flows of information, dispositions and actions travel through groups. They can be used to analyze the origins and development of *contagious* and *emergent* behavior in crowds. The structure of the network is important for such considerations. One’s position in a network for instance, can identify one as a leader or a follower. But network structure alone is not enough to explain how different people with different preferences, beliefs, intentions and goals, are able to quickly – and often unconsciously – adopt the same view or behavior. In social neuroscience neural mechanisms have been discovered that can account for the (unconscious) process of recognizing mental states of others. It has been found that some neurons are *mirror neurons*: they are active for certain actions or bodily changes of a person as well as when the person observes someone else intending or performing the action or body change (Iacoboni, 2009a; Rizzolatti and Fabbri-Destro, 2010). In his book *Mirroring people*, Iacoboni states that studies “strongly support the hypothesis that we understand the mental states of others by simulating them in our brain, and we achieve this end by way of mirror neurons” (Iacoboni, 2009b, p.34). Indeed, when states of others are mirrored by one’s own states, which at the same time are connected via neural circuits to states crucial to one’s feelings and actions, then this provides an effective basic mechanism for how in a social context people fundamentally affect each other’s actions and feelings.

These principles, along with network awareness and social context information, can be used to aid persuasive technology (Fogg, 2003; Oinas-Kukkonen and Harjumaa, 2009). For instance, they can be deployed to achieve persuasion through implicit, unconscious processes. Moreover, support systems could be the mediators that uses these principles to establish who in the user’s social environment would be the one most likely to succeed in persuading the user to change.

In this thesis processes of contagion are also considered a particular type of persuasion. Contagion often happens unconsciously, both from view of the source (who may not be aware that his/her actions, intentions or emotions influence others) and the receiver (who may not be aware that his/her actions, intentions and emotions are influenced by those of others). As such, it does not fall under the traditional view of persuasion that focuses on a conscious intentionality from the source. However, unconscious processes such as contagion often contribute to and enhance persuasive effects. Most of the factors that operate on the peripheral route of persuasion consist of these unconscious processes. Think for example of the effects of a source's nervousness, attractiveness, or easy smiles on the willingness to accept information. Given their considerable impact on persuasion, such unconscious and unintentional processes should have a place in the definition of persuasion. This is consistent with the view of Gass and Seiter that "many influence attempts take place without any conscious awareness on the part of the persuader" (Gass and Seiter, 2007, p.24). Influence processes such as contagion that contribute to the attitude or behavior change of others can be considered part of 'borderline' persuasion (see also Gass and Seiter (2007)).

1.1.3 Persuading to Change: Human vs Computer

Tailored information is often used to *persuade* people to change their attitude or behavior or both. Being persuaded to change your behavior or attitude can come about in different ways. Sometimes it happens because arguments and information are presented to which leads one to (re)consider one's current behavior. Yet just as often, persuasion occurs without any deliberation, but because one is influenced by other characteristics than content, such as source expertise or attractiveness. The *Elaboration Likelihood Model* (ELM) of persuasion is a dual process theory that distinguishes two routes of persuasion for attitude formation and change: the central and the peripheral route (Petty and Cacioppo, 1986). In general, the two different routes involve different types of persuasion processes. The central route can be considered to involve explicit (controlled), conscious processes, while the peripheral route involves implicit (automatic), unconscious processes. Central-route persuasion requires someone to think and deliberate in order to determine the validity of presented information or an arguments' merits. This means that one's cognitive response to the message determines its persuasive outcome. For persuasion to succeed in this way, a person must have the ability and motivation to cognitively process it. Peripheral persuasion on the other hand, does not involve elaboration of the message through cognitive processing. Instead, it relies on the impact of environmental characteristics on one's feelings, for example the perceived credibility of the source, geographical or social distance to the source, or the source's authority. The peripheral route is a mental shortcut which accepts (or rejects) information based on external cues, rather than thought.

It seems that no other means could be more effective in persuading humans to change than interpersonal communication. First, as addressed in Section 1.1.1.2, literature suggests that more specific, personalized information produces the best results for behavior change, and humans can deliver just that type of information, as they are generally good at sensing data from the environment (such as recognizing facial expressions or locations), adapting to new information, and learning from past interactions (for example, whether or not a friend responds well to negative motivation). Furthermore, copied behaviors of others makes up a large part of human behavior (see Section 1.1.2). To date, human intervention is used on a large scale, think for example of health counselors, personal coaches, and AA sponsors.

Recently, however, computerized interventions have shown their own potential in the domain of behavior change. With the growing developments in the field of ambient intelligence, technologies are increasingly embedded and integrated into everyday objects and environments. Ambient technologies, because of their wide availability and rapid distribution create new opportunities for reaching many people in a short time. The use of display screens, mobile sensor devices, and smart phones holds a lot of promise for cost-effective spread of information, while at the same time using context-driven personal information. These factors make ambient technologies optimal candidates for delivering persuasive communication (Kaptein, 2012; Noar et al., 2007).

One can imagine however, that taking advice from a computer agent is different from taking advice from one of your colleagues, and that you may be more likely to copy the behavior of your dearest friend than that of a graphical representation on your computer. Over the past decade, researchers have investigated whether people respond differently to computers or humans trying to persuade them. They found that people have similar social responses towards computers as they do towards humans. *The Media Equation* is a communication theory that follows from these results, stating that people treat computers and other technologies as social actors (Reeves and Nass, 1996). For instance, several studies have found that people apply gender and ethnicity stereotypes to computers (Nass and Moon, 2000; Nass, Moon, and Green, 1997), and social factors such as politeness, reciprocity, authority and trust have been shown to influence human-computer relationships (see e.g., Bickmore and Picard (2005); Fogg and Nass (1997); Nass and Moon (2000); van Wissen, van Diggelen, and Dignum (2009)).

Recent work indicates that there is a distinction between those social considerations that are the same in both human-computer interaction and human-human interaction, and those that seem to differ. For example, Blount showed that in an ultimatum game people are more likely to accept lower offers from computers than from human proposers (Blount, 1995). Also, there seems to be an asymmetry with respect to how people attribute trust to themselves and to decision support systems (van Dongen and van Maanen, 2006). Although the evidence does indicate that people form a social relationship with their technological devices, much remains unclear about which social factors are involved and under what circumstances.

1.1.4 Models and Simulations of Behavior and Change

In order to support humans in making better decisions and to help them to change, it is important to have an understanding of why they behave the way they do. As argued by Lehto amongst others, designing systems for behavior change requires a thorough understanding of the problem domain and underpinning theories (Lehto, 2012). Supporting behavior change is about diagnosing why a person performs ‘undesirable’ behavior and providing feedback such that he or she can improve. As such, it is important to have a grasp of the key determinants of individual states and social interactions, and the interactions between the two. Such insights can be a first step towards successful prediction and influence on the adjustment of behavior, by giving direction to expected outcomes of, for example, changes in context.

Creating a *model*, that is, a (simplified) representation of the phenomenon under investigation, helps to focus on its most important elements and relations. A challenge when creating models for human decision making is that the model needs to account for:

- a heterogeneous population
- nonlinear behavior
- adaptive behavior
- dynamics over time
- changing environments
- emergent behavior

This list of requirements relates closely to the necessary capabilities of agents that function as Behavior Change Support Systems (being adaptive and proactive, use personalization and tailoring, and deal with context and timing, see Section 1.1.1). It should therefore come as no surprise that *Agent-Based Models* (ABMs) are very suitable to meet the list of requirements (Bonabeau, 2002; Smith and Conrey, 2007). First, they provide a means to define individual processes that include for example cognitive states, reasoning and action selection. This means that different agents can have different traits, goals and preferences, resembling heterogeneity between people. Second, as mentioned in Section 1.1.1, intelligent agents do not necessarily exhibit rational behavior. Human behavior is from time to time irrational and nonlinear and as such, realistic models of human behavior have to include emotions and biases. Third, in an ABM, each agent is affected by its environment, which changes over time and consists for a large part of other agents. The agents can learn from the interactions in that environment, be adaptive, and as such the agents can express behavior change. Fourth, models that operate at an aggregated level — e.g., models of a crowd of many people — are often difficult to create due to the complexity of coordination and information exchange. Using multi-agent systems or team-based agents offers a way to address this challenge (Fan and Yen, 2004). Existing agent-based systems that help children deal with cyberbullying (van der Zwaan, Geraerts, Dignum, and Jonker, 2012), that support humans to make appropriate bids in negotiations (Lin, Oshrat, and Kraus, 2009), or that pro-actively advice inexperienced incident commanders (Bergmann, 2007) show that agents are indeed very suitable to be the supporters on your team.

In agent-based modeling, two sources of information are important (see Russell and Norvig (2003)). The first is how the world evolves outside of the agent: how the human behaves and how the context changes. The second is how the actions of the agent affect that world. Especially for agents that aim to change human behavior (e.g., function as a coach), the predicted effects of its actions are relevant. Such agents will be most effective if they choose actions that have the highest chance of achieving successful behavior change. Also, if the agents are to be successful, they have to adapt when the results of their chosen actions are undesirable. The models can be used to perform *simulations*, which display the dynamics of the agent (population) and the environment over time. Simulating cause-and-effect scenarios and the influence of context can lead to a better understanding of the theory and its implications. At the same time, simulations provide a means to examine possible effects of interventions.

1.1.5 Applications

The design and development of systems for behavior change is evidently a multifaceted and interdisciplinary issue. As such, it is a challenge to create systems that are specific enough to deal with the specifics of one person, but at the same time are general enough to be applied to a variety of domains. There are however two main types of domains in which humans can especially benefit from behavior change support systems.

The first type captures domains with *highly demanding tasks*, that is, in which decisions have to be made in uncertain and dynamic environments, with time constraints and many actors. These are scenarios in which decision making is stressful and a lot of risks are involved. One could think of team operations in combat settings, disaster management, or crowd control. In these domains, agent-based technology has proven to be a promising approach both to model and simulate large crowds (addressing for example emergency evacuations (Tsai, Bowring, Marsella, and Tambe, 2011) and riots (Chao and Li, 2011)), as well as to contribute to team efforts in demanding situations (such as in the battlefield (Yen, Fan, Sun, Hanratty, and Dumer, 2006) and in space (Smets, van Diggelen, Neerincx, Bradshaw, Jonker, de Rijk, Senster, ten Thije, and Sierhuis, 2010)).

The second type captures domains in which humans have particular *difficulties in functioning in accordance with their intentions and goals*. Such difficulties occur most notably when humans experience some psychological or physical disturbances, as is for example the case for patients with a chronic disease who have to adhere to a therapy. The health domain has been the subject of many studies focussing on behavior change (e.g., Kroeze, Werkman, and Brug (2006); Norman, Zabinski, Adams, Rosenberg, Yaroch, and Atienza (2007)). *eHealth* is the term used to describe this health care practice supported by computer-tailored interventions and interactive technologies. Mobile phones and websites are used to improve patients' insight into the development of their own situation and to stimulate their self-monitoring and therapy adherence. But also non-patients are known to struggle with adhering to their health goals. When one experiences low self-control or pressure from the social environment, one may be prevented from committing to actions that will achieve the intended change with respect to, for example, environmental responsibility or increase in physical activity. Intelligent agents are increasingly deployed to enhance supportive systems in these domains (e.g., Smith, Cavazza, Charlton, Zhang, Turunen, and Hakulinen (2008); Sokolova and Fernández-Caballero (2009)).

Several chapters in this work apply agent-based models to domains with these characteristics. Most of them are related to eHealth, but some focus on group decision making in emergency circumstances. See for more details the outline in Section 1.5.

1.2 Research questions

Following the motivation as set out in Section 1.1, the research challenges addressed in this dissertation focus on behavior changes in the domains of health and safety. The thesis deals with the following question:

How can agent-based systems effectively support behavior change using computational models?

In this thesis, two methods for answering this question are explored. One part (Part II) of the thesis is dedicated to the use of computational models and agent-based simulations of the phenomena that affect behavior change. The questions that are addressed in this part of the thesis contribute to a better understanding of the processes related to behavior change by creating theory-driven computational models based on theories in the field. These questions are:

1. How can insights into processes of learning, habit learning, mirroring, social contagion and behavior change be translated into computational models?
2. Can agent-based computational models adequately reproduce and simulate human behavior?
3. How can agent-based simulations based on computational models be used to explore effective interventions?

Another part (Part III) of the thesis is concerned with the application of the models and is dedicated to empirical studies on how computational models can be applied and evaluated in agent-based systems. Designing for behavior change presents many challenges and thus far there is only limited evidence of the kind and quality of the effects of behavior change support systems. With respect to evaluating and designing agent support systems for behavior change the following questions are addressed:

4. How do complex models of behavior that integrate individual and social processes perform compared to simpler models?
5. Do agent-based behavior change support systems produce behavior change?
6. Which social factors play a role when humans interact with agents and how do they affect these interactions?

1.3 Method and Approach

This dissertation focuses on models and analyses of decision making processes and behavior change from both an individual and a collective perspective. The models aim to shed light on the workings and implications of theories from the field at different aggregation levels, such as physiological, cognitive, behavioral and social levels. Furthermore, the purpose is to investigate how these models contribute to the creation of intelligent applications that are able to establish a certain behavior change of the user. The process of developing, verifying and validating models for this purpose consists of several phases, which will be addressed in the remainder of this section.

1.3.1 Model Development

1.3.1.1 Theory and Conceptual Models

Existing theories on decision making and behavior change can greatly contribute to the development of models that are able to realistically capture behavioral and mental processes. *Theory-driven* models use theories from the experts in a domain, which enriches them with a solid knowledge base of the subject. Fishbein and Cappella state that "[t]heories of behavioral prediction and behavior change are useful because they provide a framework to help identify the determinants of any given behavior, an essential first step in the development of successful interventions to change that behavior." (Fishbein and Cappella, 2006, p.S1). The theories used to create the models in this dissertation originate from the research areas of psychology, sociology, physiology, and occasionally, philosophy. A literature review can help to identify those theories relevant to the research question to create a *domain model*. Often however, the theories are informal, ambiguous and use different terms to describe the same concepts. They frequently provide conceptual definitions, but for translation to a model *operational* definitions are necessary. Formalizing theories to create operational definitions requires very strict and precise definitions of the components of a theory and the relations between them. It is during this process of formalization that the theory is scrutinized, and will be checked for inconsistencies, ambiguity and incompletion. As such, formalizing theories can in turn contribute to the development of those theories by:

- clarifying which assumptions are part of the theory
- comparing and integrating existing theories
- finding errors or ommitances in the theory
- strictly defining the components and their relations

1.3.1.2 Agent-Based Modeling

When creating a model, many choices have to be made. First and foremost the model type has to be determined. The research problem often influences the model type, as models can be defined at different levels of abstraction and designed to model different aspects of a system. Some models are more suitable to deal with very detailed process descriptions while other models are better suited to recognise patterns without any lower-level specifications. As addressed in Section 1.1.4, agent-based models are able to capture the dynamics of behavior change at both an individual and social level. As a result, "[t]his approach allows for the observation of the large-scale consequences of the theoretical assumptions about agent behavior when the behaviors are carried out in the context of many other agents and iterated dynamically over an extended period of time." (Smith and Conrey, 2007, p.88)

In particular, ABMs are a class of *computational models* that can consist of formalizations of different kinds, for example combining elements of game theory, computational sociology, and emergency and evolutionary programming, to name only a few. In this work the agents are required to use the domain knowledge to reason about human performance and how they can establish a change, if desired. Contrary to a number of researchers who have posed ABM as an alternative to variable-based modeling (e.g., Smith and Conrey (2007)), others have proposed agent-based models should consist of a set of differential equations describing the dynamics of the components of a system and their relations (Bonabeau, 2002). A hybrid approach incorporates both qualitative states and transitions, and numerical dynamic systems based on difference or differential equations. This approach results in models that can describe state dynamics over time for which certain qualitative properties holds. The computational models in this dissertation are based on these hybrid models in order to reflect how variables in dynamical systems change over time. As an illustration, consider an update rule for (a value of) a particular belief of an agent:

$$Belief_1(t+\Delta t) = Belief_1(t) + \eta \cdot (\omega_i \cdot Information_i + (1-\omega_j) * Information_j) \cdot \Delta t$$

In this equation, η can be considered the update factor (is the agent a fast or a slow learner?), and ω the connection strengths of the two pieces of infomation that contribute to the belief.

1.3.1.3 Languages and Environments

The implementation of the models can be done using different languages, environments and tools. The **Temporal Trace Language** (TTL, Bosse, Jonker, van der Meij, Sharpanskykh, and Treur (2009)) is a hybrid modeling language that supports the formal specification and analysis of properties of a model. It is based on first-order logic and can capture both qualitative and quantitative aspects of dynamic properties. TTL can be used to formally specify (non-local) dynamic properties and to check these properties automatically against formal traces. Traces are time slices consisting of sequences of (values of) states. The checking is done to determine which properties are true in a model and which are false. For example, the following statements denote that the state property p holds and does not hold, respectively, in trace γ at time t.

$$state(\gamma, t) \models p$$
$$state(\gamma, t) \not\models p$$

LEADSTO (Language and ENvironment for Analysis of Dynamics by SimulaTiOn, (Bosse, Jonker, van der Meij, and Treur, 2005)) is an executable language that can be used to model direct dependencies between two state properties. It uses temporal-causal relations to describe state properties of the form:

$$\alpha \twoheadrightarrow_{e,f,g,h} \beta$$

which means that if state property α holds for a certain time interval with duration g, then after some delay (between e and f) state property β will hold for a certain time interval of length h. LEADSTO can use variables, which makes it easier to use than modeling languages that use only propositional or qualitative causal statements. Furthermore, it can graphically display simulations that were done with the model.

Although TTL and LEADSTO are useful for formal verification, creating properties of desired outcomes, and checking for temporal patterns, the majority of the models in this dissertation have been developed using **MATLAB** (MATLAB, 2010). Matlab is a numerical computing environment and programming language that allows for extensive data analysis, numerical computations, algorithm testing and parameter estimations. Its main advantage over LEADSTO and TTL is that is is computationally very efficient, even with large amounts of data. When modeling the interactions between many agents or groups of agents, this is especially appealing. Furthermore, MATLAB has extensive visualization options for plotting the performed simulations.

1.3.2 Verification and Validation

To check whether a developed model performs as intended, it needs to be verified and validated. *Verification* is the process of evaluating software to determine whether the products of a given development phase satisfy the conditions imposed at the start of that phase.[1] In short, it answers the question: does the model perform as intended? *Validation* is the process of evaluating software during or at the end of the development process to determine whether it satisfies specified requirements.[1] In other words, can it describe human behavior accurately? For this, model data has to be compered to observations from the real world.

1.3.2.1 Model Verification

Model checking, or *property checking*, is the exhaustive and automatic checking of whether a given model meets its specification. It is about checking hypotheses within the model about the expected behavior of the model. Several methods for property checking can be used, although they vary in their degree of exhaustiveness and automaticity. The first is to ask an expert for *face validity* of the model: whether it 'looks like' it will measure what it is supposed to measure. Secondly, a tool for automatic checking can be used. For example, TTL provides a TTL Checker which can automatically validate whether global properties identified from literature hold for all traces in the simulation (if it cannot, it provides a counter example). This ensures that the model adheres to the theories is was founded on. Contrary to model checking, this method is not exhaustive, as the found properties only hold for all analyzed traces and not all possible traces. A third method is to use mathematical analysis to examine whether the equilibria of the model match expected values based on the characteristics of the theory. Fourth, simulations can be performed to check that the model shows realistic behavior, for example comparing the simulation outcomes to known dynamics in the domain.

1.3.2.2 Model Validation

For agent-based models to scientifically contribute to theory and interventions, they must be subject to empirical validation. Of course, empirical data is not always easily available nor desirable to generate (think for example of models for dealing with stress, emergency evacuations or mass hysteria). Following Smith and Conrey, "[e]ven when empirically based validation of a model is difficult or impossible [...], ABM can be valuable [...] for the goal of understanding, comparing, integrating and ultimately improving theory" (Smith and Conrey, 2007, p.100). Yet validation with empirical data is crucial for the purpose of testing the accuracy of the model and its usability, and to answer the questions of whether it is able to accurately explain or predict the outcomes. Also, empirical data helps to fine-tune the model. By comparing data points from a data set to the model outputs, the error of the model is established. Parameter tuning techniques

[1] Carnegie Mellon University. Registered service mark. Capability Maturity Model (cmmi-sw v1.1) ieee-std-610.

can then be used to match the data output more closely to the empirical data. After parameter tuning the model is able to simulate the data more accurately.

The validation process is concerned with the question of whether the model can realistically reproduce or predict human behavior. For systems to succesfully interact and cooperate with humans, user experience also has to be addressed in the validation. Validated surveys or log data can provide information on how users experience the system and whether the agent is able to correctly reason about the preferences and cognitive or bodily states of the user.

1.4 Contributions

The aim of this dissertation is to contribute to the support of human behavior change using computational models. The dissertation advances both the current state of knowledge on human decision making, as well as the current state of the art with respect to developing support systems for behavior change.

1.4.1 Contributions to the Understanding and Modeling of Human Behavior and Change

The interactions of beliefs, intentions, and emotions are the foundation of human behavior are complex and are often underestimated or underexposed in models of behavior. For example, much research has been done on extracting sentiment and mood from tweets, but the measurement of emotional dynamics over time and their effect in social networks is fairly new (see Chapter 7). Similarly, while there has been much work on diffusion, little of it takes into account the complexity of emotional and cognitive states that affect each other not only at an intraperson level, but also at an interperson level, as is done in Chapters 4 and 5. Recent developments in social neuroscience have resulted in theories that greatly contribute to the explanation and prediction of behavior. These theories have identified dynamics that have a significant impact on diffusion characteristics. In contemporary research, only limited attention is given to these characteristics in agent-based modeling, despite their explanatory power. This thesis addresses that gap by proposing several computational models in a novel approach that link together the internal states of an agent (including beliefs, emotions and intentions), its behaviors, and its perceptions (i.e., the way it perceives and interprets the environment and behaviors of other agents in it). By formalizing the theories that are at the core of the models, the various parameters can be adjusted and examined, which consequently leads to a more thorough understanding of the theories and the complex interplay of their constructs. These insights are very relevant in supporting (collective) decision making.

Another way in which this dissertation contributes to the understanding and modeling of behavior and its adaptivity, is by using the models to create simulations and predictions of behavior. These can provide convincing arguments that the models and the theories on which they are based are an adequate representation of human behavior. The insights obtained from the computational model of habit formation (Chapters 3,5) can for example be an aid for policy-makers in making predictions of habits formation and spread in a population. In a world where the average life span of people keeps increasing, where obesity has nearly doubled since 1980[1], where health care costs are rising due to those factors, and where sustainable living and environmental impact are major concerns, the creation of policies that stimulate healthy habits and lifestyle change are crucial.

[1]World Health Organization, http://www.who.int/mediacentre/factsheets/fs311/en/, last visited: February 2nd 2014

Models that are shown to be well suited to capture the behavior dynamics are extremely significant and relevant for society. For instance, obtaining and utilizing data of real emergency behavior, especially for large bodies of people over stretches of time, is difficult to come by. There has been limited use in the literature of qualitative data to validate models of collective decision making in emergency scenarios. As such, the evaluation presented in Chapter 11 provides a rare view into the dynamics of panic situations and demonstrates the added value of a model that captures the contagion of mental states between people.

1.4.2 Contributions to Developing Support Systems

The development of computerized support systems for behavior change is a rapidly growing research field, though still in its infancy. The purpose of this thesis is not just to provide models of human behavior and decision making, but to go a step further and use these models to support behavior change. To achieve this, some preliminary and some more advanced steps are taken towards realizing behavior change support systems.

A first step toward creating such systems is to investigate how a model can be used to explore effective interventions. Such analyses are for example provided in Chapter 6, where, although the interventions are not yet really applied, it is hypothesised how they may be used in a support system for learning. In Chapter 7 it is examined how sentiment analyses extracted from tweets can be used by an ambient intelligent system to explore effective interventions for adherence to exercise behavior. These particular explorations are unique in the sense that they do not only consider direct persuasion of the user, but also interventions that target (connections to) those in the social environment to undertake action and help the user to commit to long-term behavior change. That is, the agent persuades the social supporters, who in turn persuade the target.

Going beyond these explorations, this thesis also introduces a fully functioning agent-based support system for lifestyle change (see Chapter 2,8). The approach is original and contributes to the current state of the art in several ways. First, the system omits human supervision and automatically provides tailored feedback to the user. Second, this feedback is selected based on a computational model incorporating theories of behavior change. The design of health support systems rarely includes psychologically relevant theories, yet it is an important step in the development process. Third, a novel way to reason about and increase the impact of behavior change interventions by targeting 'bottlenecks' (likely causes of non-adherence) is presented. Fourth, the system integrates many technical components and provides support on three different lifestyle domains. As such, the chosen approach is innovative, of an interdisciplinary nature, and captures many of the challenges of designing behavior change support systems. Initial tests and evaluations have been performed (see Chapters 8,9,10), which have inspired observations and guidelines for developing BCSSs (see Chapter 13).

The final contribution of this thesis is to the field of human-agent interaction. It offers insights into the social relationship between humans and agents, which can greatly benefit the design of agents that are able to effectively support or coach humans. A *Wizard-of-Oz* approach, in which subjects interact with what they believe is an autonomous agent system, but which is actually (being operated by) an unseen human being, was used to identify how perception influences factors such as trust and fairness in interactions (see Chapters 12,10).

1.5 Dissertation Outline

Reflecting the different methodologies addressed in Section 1.3, this dissertation is divided by the method of analysis used in the chapters:

Figure 1.1: Structure of the thesis

<table>
<tr><td rowspan="7">How can agent-based systems effectively support behavior change using computational models?</td><td colspan="2">PART I:
Introduction</td><td>CHAPTER 1</td></tr>
<tr><td rowspan="2">PART II:
Models & Simulations</td><td>individual processes</td><td>CHAPTER 2-3</td></tr>
<tr><td>social processes</td><td>CHAPTER 4-7</td></tr>
<tr><td rowspan="2">PART III:
Empirical validations</td><td>individual processes</td><td>CHAPTER 8-10</td></tr>
<tr><td>social processes</td><td>CHAPTER 11-12</td></tr>
<tr><td colspan="2">PART IV:
Discussion & Conclusion</td><td>CHAPTER 13</td></tr>
</table>

PART I concernes the introduction of the subject of agent-based modeling and behavior change. This part presents the motivation for the performed research and covers the research questions. It also addresses related work in the domains of decision making, contagion and persuasion.

PART II embeds those chapters that cover the identification, delineation and explicit testing of theories in the field. Answers are sought to questions such as which theories and constructs are most pertinent for a specific behavior, under what circumstances, and for which individuals or groups. In some cases theories are combined, incorporating unique contributions of different theories to create a new one. Those theories are then formalized to create agent-based models, which are in turn verified using one of the methods discussed in 1.3.2.1.

PART III embeds those chapters that cover different types of empirical validations of agent-based models that aim to analyze behavior or establish change. To validate models that aim to analyze behavior, validations are done by matching existing data to the outcome of the models in order to determine how well they perform. For models that are used in systems to establish change, the effects of interventions are also taken into account. This part also contains chapters that address how the interaction with agent-based (support) systems affects the social relationship between humans and agents.

PART IV provides discussion of the research and the answers to the research questions. In this part the also the limitations of this work are addressed, and ideas for future work are presented.

The chapters in each part are ordered by the type of process they study: individual or social. The structure of the thesis can be found in Figure 1.1. Parts II and III form the core of this thesis. Their chapters will be discussed in a bit more detail below.

============= *PART II* =============

Chapter 2 introduces a model of behavior change called COMBI, which combines several prominent theories in the field of health behavior change. The model is designed to be used by an agent-based system for therapy adherence and is able use model-based reasoning to identify the user's stage of behavior change and the determinants of non-compliance. The system uses persuasive mobile interventions to support the user to move through the stages of change.

In this chapter the theories of health behavior change are discussed and it is demonstrated how the model and simulations following from it can contribute to support behavior change.

Chapter 3 continues to explore agent-based support for healthy living. In this chapter the focus is on habits. Habits often present a barrier to attempts to change behavior, since they are automatic behaviors that guide human decision making unconsciously. A computational model for habit formation is developed which incorporates the influence of cues in the environment and goal-directed behavior. The model was implemented in LEADSTO en several simulations were peformed. Also, TTL was used to check whether the identified patterns from the theories could be verified.

Chapter 4 approaches behavior from a collective point of view and proposes a neurologically inspired model to capture the dynamics and diffusion of behavior in groups. Neurological principles such as mirroring and the somatic marker hypothesis are discussed and incorporated in the model. Property verifications (using TTL) and mathematical analyses were performed to demonstrate patterns in the model's output. Several simulations show the dynamics between beliefs and emotions in groups.

Chapter 5 builds on the research of Chapters 3 and 4. The subject of this chapter is habits in a social context, and considers both the individual and the social mechanisms that influence habit formation. Modeling approaches from both Chapter 3 and 4 are incorporated in a computational model of habit contagion. Simulations show how habits can change when the agents interact in small or large groups (consisting of 2 and 100 agents, respectively). The results are compared with behavior patterns identified in literature.

Chapter 6 focuses on the core of all behavior change and habit formation: learning. In this chapter, insights from cognitive, affective and social neuroscience are used to capture the dynamics of emotions and social interactions in learning. In particular, it is examined how a computational model of such dynamics in learning can account for the differences experienced by different types of learners (e.g., active or reflective learners). Multiple simulations are performed to demonstrate the flexibility and scope of the model.

Chapter 7 creates a foundation for an agent-based system that puts the knowledge gained in the previous chapters into practice. The core of the system is a model of how emotions (both your own and those of others) affect behavior. Information from social interactions (such as text feeds from Twitter) is used to establish the user's state with respect to adhering to a lifestyle goal. Furthermore, several intervention techniques are proposed that can stimulate the support of members in the user's social network.

============= *PART III* =============

Chapter 8 provides validations for the model of behavior change proposed in Chapter 2. It is shown how this model is employed in an agent-based Behavior Change Support System called eMate, which is comprised of an app, a smart pill box and a website. eMate is designed to support patients who have diabetes type two, cardiovascular disease or HIV.

Chapter 9 describes another validation study of the COMBI model proposed in Chapter 2. Contrary to the studies performed in Chapter 8, this study was done with healthy participants. The purpose of the study was to use the eMate system to encourage students to take the stairs more often. The performance and reasoning strategies of the COMBI model are analyzed using the logs and survey data.

Chapter 10 is based on the study introduced in Chapter 9, but focuses on the Wizard-of-Oz setup of the experiment. In particular, it addresses how perceptions of the nature of the coach (human or computer) can influence the (perceived and actual) effectiveness of the eMate system. Using validated surveys as well as logs, the effect of these perceptions is measured with regard to improvements in stage of change, number of stairs taken, and self-reports of effectiveness and trust.

Chapter 11 presents a computational model for collective decision making, drawing from insights gained in Chapter 4. The model is used to analyze mutual influences between beliefs, intentions and emotions in groups. Several Matlab simulations are performed to demonstrate the model dynamics. Furthermore, the computational model is validated for emergent crowd behavior using empirical data from an incident in Amsterdam on the 4th of May in 2010. The model is calibrated using parameter tuning and subsequently the results are compared to those of similar decision making models.

Chapter 12 is also concerned with group interactions. It does not propose a model for decision making, but instead examines how social factors are influenced by human-computer interactions. The study has a Wizard-of-Oz setup with similarities to the one presented in Chapter 10. A web-based testbed is used to create a dynamic virtual environment where repeated interaction between participants is enforced. The participants are deceived about the nature of their team mates (human or agent). Using survey and log data, it is investigated which social factors affect human decision making in choosing teammates.

The chapters are based on the following publications (the authors are mentioned in alphabetical order and have made a comparable contribution to the work, unless the publication is marked by a *, in which case the authors are mentioned by degree of contribution):

Chapter 2 will appear as (the first part of):

Klein, M.C.A., Mogles, N., Wissen, A. van. Intelligent Mobile Support for Therapy Adherence and Behavior Change. *Journal of Biomedical Informatics*, 2014

Chapter 3 has been published as:

Klein, M.C.A., Mogles, N., Treur, J., Wissen, A. van. A Computational Model of Habit Learning to Enable Ambient Support for Lifestyle Change. In: K.G. Mehrota et al. (eds.): *Proceedings of the 24th International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems* (IEA-AIE'11), Part II. Lecture Notes in Artificial Intelligence, vol. 6704, pp. 130–142. Springer Verlag, 2011.

Chapter 4 has been published as:

Hoogendoorn, M., Treur, J., Wal, C.N. van der, Wissen, A. van. Agent-Based Modelling of the Emergence of Collective States Based on Contagion of Individual States in Groups. *LNCS Transactions on Computational Collective Intelligence*, vol. 3, 2011, pp. 152-179.

Chapter 5 has been published as:

Klein, M.C.A., Mogles, N., Treur, J., Wissen, A. van. Contagion of Habitual Behaviour in Social Networks: an Agent-Based Model. In: *Proceedings of the 4th International Conference on Social Computing* (SocialCom'12), IEEE Computer Society Press, 2012.

Chapter 6 is based on published papers:

Treur, J., Wissen, A. van. On the Impacts of Emotion on Learning in a Social Context: A Conceptual and

Computational Analysis. In: *Proceedings of the 12th International Conference on Intelligent Agent Technology* (IAT'12), IEEE Computer Society Press, 2012.
Treur, J. and Wissen, A. van, Conceptual and Computational Analysis of the Role of Emotions and Social Influence in Learning. In: *Proceedings of the Third World Conference on Learning, Teaching and Educational Leadership* (WCLTA'12), Procedia Social and Behavioral Sciences, Elsevier, 2012.

Chapter 7 has been published as:
Breda, W. van, Treur, J., and Wissen, A. van, Analysis and Support of Lifestyle via Emotions Using Social Media. In: *Proceedings of the 4th International Conference on Social Informatics* (SocInfo'12), Lecture Notes in Computer Science, vol. 7710, pp. 275-291, Springer Verlag, 2012.

Chapter 8 will appear as (the second part of):
Klein, M.C.A., Mogles, N., Wissen, A. van. Intelligent Mobile Support for Therapy Adherence and Behavior Change. *Journal of Biomedical Informatics*, 2014

Chapter 9 will appear as:
Kamphorst, B.A., Klein, M.C.A, and Wissen, A. van. Autonomous e-coaching in the wild: Empirical Validation of a Model-Based Reasoning System. *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems* (AAMAS 2014), Lomuscio, Scerri, Bazzan, Huhns (eds.), May, 5–9, 2014, Paris, France, 2014.

Chapter 10 is under review as:
Kamphorst, B.A., Klein, M.C.A, and Wissen, A. van. Human involvement in e-coaching: effects on effectiveness, perceived influence and trust, 2014.

Chapter 11 has been published as:
Bosse, T., Hoogendoorn, M., Klein, M.C.A., Treur, J., Wal, C.N. van der, and Wissen, A. van. Modelling Collective Decision Making in Groups and Crowds: Integrating Social Contagion and Interacting Emotions, Beliefs and Intentions. *Autonomous Agents and Multi-Agent Systems Journal* (JAAMAS), volume 27 (1), 2013, pp. 52-84.

Chapter 12 has been published as:
*Wissen, A. van, Gal, Y., Kamphorst, B.A., Dignum, V. Human-Agent Team Formation in Dynamic Environments. Computers in Human Behavior, 28(1), 2012, pp. 23-33

References

Aarts, E. and B. de Ruyter (2009). New research perspectives on ambient intelligence. *Journal of Ambient Intelligence and Smart Environments 1*(1), 5–14.

Aarts, E., R. Harwig, and M. Schuurmans (2001). Ambient intelligence. In P. Denning (Ed.), *Ambient Intelligence in The Invisible Future: The Seamless Integration Of Technology Into Everyday Life*, pp. 235–250. New York: McGraw-Hill Companies.

Andrew, A., G. Borriello, and J. Fogarty (2007). Toward a systematic understanding of suggestion tactics in persuasive technologies. In *Proceedings of the 2nd international conference on Persuasive technology*, PERSUASIVE'07, Berlin, Heidelberg, pp. 259–270. Springer-Verlag.

Bandura, A. (1997). *Self-efficacy: The exercise of control*. New York: Freeman.

Barsade, S. (2002). The ripple effect: Emotional contagion and its influence on group behavior. *Administrative Science Quarterly 47*(4), 644–675.

Baumeister, R. (2002). Yielding to temptation: Self-control failure, impulsive purchasing, and consumer behavior. *ournal of Consumer Research 28*(4), 670–676.

Bergmann, R. (2007). Ambient intelligence for decision making in fire service organizations. In *Ambient Intelligence*, pp. 73–90. Springer.

Bickmore, T. W. and R. W. Picard (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions of Computer-Human Interaction 12*(2), 293–327.

Blount, S. (1995). When social outcomes aren't fair. *Organizational Behavior and Human Decision Processes 63*(2), 131–144.

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America 99*(3), 7280–7287.

Bosse, T., C. M. Jonker, L. van der Meij, A. Sharpanskykh, and J. Treur (2009). Specification and verification of dynamics in agent models. *International Journal of Cooperative Information Systems 18*(01), 167–193.

Bosse, T., C. M. Jonker, L. van der Meij, and J. Treur (2005). Leadsto: a language and environment for analysis of dynamics by simulation. In *Multiagent System Technologies*, pp. 165–178. Springer.

Brug, J., M. Campbell, and P. van Assema (1999). The application and impact of computer-generated personalized nutrition education: A review of the literature. *Patient Education and Counseling 36*, 145–156.

Chao, W.-M. and T.-Y. Li (2011). Simulating riot for virtual crowds with a social communication model. In *Computational Collective Intelligence. Technologies and Applications*, pp. 419–427. Springer.

Christakis, N. and J. Fowler (2007). The spread of obesity in a large social network over 32 years. *The new england journal of medicine*, 370–379.

Costa, P. and R. McCrae (1994). Set like plaster? evidence for the stability of adult personality. In *Can Personality Change?*, pp. 21–40. Washington, DC: American Psychological Association.

van Diggelen, J., T. Muller, and K. van den Bosch (2010). Using artificial team members for team training in virtual environments. In *Proceedings of the 10th international conference on Intelligent virtual agents*, IVA'10, Berlin, Heidelberg, pp. 28–34. Springer-Verlag.

van Dongen, K. and P.-P. van Maanen (2006). Under-reliance on the decision aid: A difference in calibration and attribution between self and aid. In *Proceedings of the Human Factors and Ergonomics Society's 50th Annual Meeting*, Volume 50, San Francisco, USA, pp. 225–229.

Fan, X. and J. Yen (2004). Modeling and simulating human teamwork behaviors using intelligent agents. *Physics of Life Reviews 1*(3), 173–201.

Fishbein, M. and J. N. Cappella (2006). The role of theory in developing effective health communications. *Journal of Communication 56*(s1), S1–S17.

Fogg, B., G. Cuellar, and D. Danielson (2009). Motivating, influencing, and persuading users: An introduction to captology. In J. J. A. Sears A. (Ed.), *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, pp. 113–146. New York: Taylor Francis.

Fogg, B. J. (2003). *Persuasive Technology: Using computers to change what we think and do*. San Fransisco: Morgan Kaufmann Publishers.

Fogg, B. J. and C. Nass (1997). Silicon sycophants: The effects of computers that flatter. *International Journal of Human-Computer Studies 46*(5), 551–561.

Gass, R. and J. Seiter (2007). *Persuasion, social influence and compliance gaining*. Pearson.

Iacoboni, M. (2009a). Imitation, empathy, and mirror neurons. *Annual Reviews in Psychology 60*, 653–70.

Iacoboni, M. (2009b). *Mirroring People: The Science of Empathy and How We Connect with Others*. New York: Picador.

Kaptein, M. C. (2012). Personalized persuasion in ambient intelligence. *Journal of Ambient Intelligence and Smart Environments 4*(3), 279–280.

Kreuter, M., V. Stretcher, and Glassman (1999). One size does not fit all: The case for tailoring print materials. *Annals of Behavioral Medicine 21*(4), 276–283.

Kroeze, W., A. Werkman, and J. Brug (2006). A systematic review of randomized trials on the effectiveness of computer-tailored education on physical activity and dietary behaviors. *Annals of Behavioral Medicine 31*(3), 205–223.

Lehto, T. (2012). Designing persuasive health behavior change interventions. In *Critical Issues for the Development of Sustainable E-health Solutions*, pp. 163–181. Springer.

Lin, R., Y. Oshrat, and S. Kraus (2009). Investigating the benefits of automated negotiations in enhancing people's negotiation skills. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pp. 345–352. International Foundation for Autonomous Agents and Multiagent Systems.

Loewenstein, G. (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology 57*(3), 426–441.

MATLAB (2010). *version 7.10.0 (R2010a)*. Natick, Massachusetts: The MathWorks Inc.

Nass, C. and Y. Moon (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues 56*(1), 81–103.

Nass, C., Y. Moon, and N. Green (1997). Are computers gender-neutral? gender stereotypic responses to computers. *Journal of Applied Social Psychology 27*(10), 864–876.

Noar, S., C. Benac, and M. Harris (2007). Does tailoring matter? meta-analytic review of tailored print health behavior change interventions. *Psychological Bulletin 133*(4), 673–693.

Norman, G. J., M. F. Zabinski, M. A. Adams, D. E. Rosenberg, A. L. Yaroch, and A. A. Atienza (2007). A review of ehealth interventions for physical activity and dietary behavior change. *American journal of preventive medicine 33*(4), 336–345.

Oinas-Kukkonen, H. (2010). Behavior change support systems: The next frontier for web science. *Proceedings of the WebSci10: Extending the Frontiers of Society On-Line*.

Oinas-Kukkonen, H. and M. Harjumaa (2008). A systematic framework for designing and evaluating persuasive systems. In H. Oinas-Kukkonen (Ed.), *Lecture Notes of Ccomputer Science*, Volume 5033 of *PERSUASIVE'08*, pp. 164–176. Springer-Verlag Berlin Heidelberg.

Oinas-Kukkonen, H. and M. Harjumaa (2009). Persuasive systems design: Key issues, process model, and system features. *Communications of the Association for Information Systems 24*(28).

Petty, R. E. and J. T. Cacioppo (1986). *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*. New York: Springer-Verlag.

Picard, R. W. (2000). *Affective Computing*. Cambridge, MA: MIT Press.

Reeves, B. and C. Nass (1996). *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places (CSLI Lecture Notes Series, no. 63)*. Center for the Study of Language and Informatics.

Revere, D. and P. Dunbar (2001). Review of computer-generated outpatient health behavior interventions: clinical encounters 'in absentia'. *Journal of the American Medical Informatics Association 8*, 62–79.

Rizzolatti, G. and M. Fabbri-Destro (2010). Mirror neurons: from discovery to autism. *Experimental Brain Research 200*(3-4), 223–237.

Russell, S. J. and P. Norvig (2003). *Artificial Intelligence: A Modern Approach* (2 ed.). Pearson Education.

Silverman, B., D. Pietrocola, B. Nye, N. Weyer, O. Osin, D. Johnson, and R. Weaver (2012). Rich socio-cognitive agents for immersive training environments: case of nonkin village. *Autonomous Agents and Multi-Agent Systems 24*(2), 312–343.

Skinner, C., M. Campbell, B. Rimer, S. Curry, and J. Prochaska (1999). How effective is tailored print communication. *Annals of Behavioral Medicine 21*, 290–298.

Smets, N. J., J. van Diggelen, M. A. Neerincx, J. M. Bradshaw, C. M. Jonker, L. J. de Rijk, P. A. Senster, O. ten Thije, and M. Sierhuis (2010). Assessing human-agent teams for future space missions. *IEEE Intelligent Systems*, 46–53.

Smith, C., M. Cavazza, D. Charlton, L. Zhang, M. Turunen, and J. Hakulinen (2008). Integrating planning and dialogue in a lifestyle agent. In *Intelligent Virtual Agents*, pp. 146–153. Springer.

Smith, E. R. and F. Conrey (2007). Agent-based modeling: A new approach for theory building in social psychology. *Personality and social psychology review 11*(1), 87–104.

Sokolova, M. V. and A. Fernández-Caballero (2009). Modeling and implementing an agent-based environmental health impact decision support system. *Expert Systems with Applications 36*(2), 2603–2614.

Thaler, R. and C. Sunstein (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press.

Townsend, D. and T. Bever (2001). *Sentence comprehension: The integration of habits and rules*. Cambridge, MA: MIT Press.

Tsai, J., E. Bowring, S. Marsella, and M. Tambe (2011). Empirical evaluation of computational emotional contagion models. In *Intelligent Virtual Agents*, pp. 384–397. Springer.

Warner, T. (2012). E-coaching systems: Convenient, anytime, anywhere, and nonhuman. *Performance Improvement 51*(9), 22–28.

van Wissen, A., J. van Diggelen, and V. Dignum (2009). The effects of cooperative agent behavior on human cooperativeness. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems*, pp. 1179–1180.

Wood, W. and D. Neal (2007). A new look at habits and the habit-goal interface. *Psychological Review 114*(4), 843–863.

Wooldridge, M. and N. R. Jennings (1995). Intelligent agents: Theory and practice. *Knowledge Engineering Review 10*(2).

Yen, J., X. Fan, S. Sun, T. Hanratty, and J. Dumer (2006). Agents with shared mental models for enhancing team decision makings. *Decision Support Systems 41*(3), 634–653.

van der Zwaan, J., E. Geraerts, V. Dignum, and C. Jonker (2012). User validation of an empathic virtual buddy against cyberbullying. *Annual Review of Cybertherapy and Telemedicine 2012: Advanced Technologies in the Behavioral, Social and Neurosciences 181*, 243.