CHAPTER 4

Virtual Affective Agents and Therapeutic Games

Eva Hudlicka
Psychometrix Associates, Inc. and School of Computer Science, University of Massachusetts, Amherst, MA, USA

INTRODUCTION

The past decade has witnessed tremendous progress in the development of affective intelligent virtual agents and serious games (games used for training, coaching, therapy, and other non-entertainment purposes). These developments, in conjunction with an expanding interest in telemental health, have led to a growing recognition that these technologies can significantly enhance clinical practice and improve patient outcomes. This chapter provides a brief history of intelligent virtual affective agents and serious gaming technologies. The state-of-the-art in these areas and recent applications in health care, specifically behavioral health, are presented. Ethical and privacy considerations, as well as future prospects, are also discussed.

Intelligent virtual agents (IVAs) are synthetic virtual, computer-controlled characters that can interact with humans (Prendinger & Ishizuka, 2004). Virtual affective agents are IVAs capable of affective interaction, often emphasizing nonverbal interaction; that is, recognition of human emotions, and expression of "emotions" by the agent (Hudlicka et al., 2008; Hudlicka et al., 2009). Affective agents may also include affective user modeling and may explicitly model synthetic emotions within their architectures. These capabilities enable the agents to display degrees of emotional and social intelligence (e.g., awareness of social cues and user goals), to have some ability to form and manage relationships, and to adapt in real-time to the changing states and needs of the human user.

Both the visual appearance and the multi-modal interaction capabilities of these agents vary greatly, and can be customized to match the human user's preferences and needs. The embodiments range from cartoonish, animated characters or faces, to fully articulated full bodies, again, customizable to the user's preferences and needs. The interaction capabilities range from text-based, multiple-choice user input and

Artificial Intelligence in Behavioral and Mental Health Care.
DOI: http://dx.doi.org/10.1016/B978-0-12-420248-1.00004-0

Figure 4.1 Examples of virtual agents illustrating varying degrees of visual realism and a variety of embodiment types. From left to right: Woggles; coach Chris (top), coach Harmony (bottom); character from FearNot!; coach Laura; ECA Greta; agent Max.

text-based character output, to speech recognition and synthetic speech. The virtual agents also vary in the degree to which they can adapt to the user's knowledge, affective and motivational states, and their display of affective and social realism. Figure 4.1 shows examples of several virtual agents, illustrating the range of embodiments and degrees of visual realism. Increasingly, affective IVAs are being integrated in computer games, both in games for entertainment and in serious games.

Serious games are computer games developed for training and learning purposes, in contrast to games developed solely for entertainment. Games have a unique ability to engage the players, and to provide highly immersive learning, training and therapeutic environments that can be customized to the user's specific learning needs or therapeutic goals. Their potential in education and training, coaching, rehabilitation, and even psychotherapy has been increasingly recognized, and in spite of their relatively recent emergence within the past 2 decades, serious games represent the fastest-growing segment of the gaming market.

As is the case with games for entertainment, serious games typically provide a game "storyline", which evolves across distinct physical contexts within the simulated gameworld, and typically involves multiple non-playing characters (NPCs) and distinct tasks which the player aims to achieve as s/he progresses through the game levels. The skills to be learned or practiced are embedded within the game tasks, and the game levels provide progressively more challenging tasks. Depending on the type of task to be learned or the type of training or coaching, as well as on the age and abilities of the players, the gameplay may focus strictly on the "serious" task, or it may interleave these tasks with segments of gameplay designed only for entertainment. The latter is more typical for games aimed at children and younger users. Games have the potential to create highly customized learning, training and therapeutic environments and protocols. The gameplay levels, the reward structure, the NPCs, as well as the overall game storyline, can all be customized to provide an optimal learning and training experience for the user. As is the case with IVAs, the NPCs' appearance and behavior can be defined to match players' individual and cultural preferences as well as specific learning and training needs.

Both of these aforementioned technologies take advantage of two innate human needs and capabilities: the desire and ability to emotionally "connect" (i.e., to attach), and the desire and need to "play." To illustrate the potential of these technologies, consider the following questions. What if . . .

- a patient recovering from a major depressive episode could practice cognitive restructuring skills with the help of a virtual intelligent affective character or in the context of a customized serious game?
- a child on the autism spectrum could play with a virtual affective agent to learn social skills or interact with multiple such agents, in a customized learning environment embedded within a serious game?
- a patient with social phobia could attend "virtual parties" and practice approaching strangers and initiating and maintaining conversations?

These scenarios are rapidly approaching reality. Researchers are exploring the effectiveness of IVAs and social robots in helping children on the autism spectrum learn affective and social skills (Dautenhahn, 2007). Serious games are being used to help children and teens on the autism spectrum learn social skills and manage anxiety (Beaumont & Sofronoff, 2008), and games exist to help treat obsessive-compulsive disorder (OCD) (Brezinka, 2008; Brezinka & Hovestadt, 2007).

It must be emphasized that the discussion in this chapter in no way means to imply that virtual affective agents or serious games should be replacing human therapists or face-to-face therapy. These technologies cannot function at the level of an experienced, empathic human therapist, and likely never will. Rather, they have a unique role in the delivery of behavioral health care, both supportive of, and distinct from, the roles of human therapists. These include the following capabilities:

- Enhance dissemination of evidence-based treatment
- Make treatment more accessible (anytime/anywhere availability)
- Support treatment between sessions (facilitate homework and skills practice)
- Adapt to individual needs and cultural preferences
- Promote engagement and support motivation.

In some cases, technology may even be the preferred method of delivering services (e.g., for children on the autism spectrum or patients with social anxiety or agoraphobia). In addition, these technologies can also support enhanced diagnosis and assessment and play a role in research, where they can contribute to more mechanism-based diagnosis and treatment planning (e.g., Hudlicka, 2008a).

There are a number of unique benefits of both virtual affective agents and serious games, which make these technologies particularly well-suited for applications in health care in general, and behavioral health care in particular – both in the delivery of behavioral interventions and in training. This chapter introduces these technologies and discusses their applications in behavioral health.

BRIEF HISTORY OF VIRTUAL AFFECTIVE AGENTS AND SERIOUS GAMES

Both virtual affective agents and serious games are highly sophisticated technologies, which represent a culmination of decades of research and development across a number of sub-disciplines of computer science, engineering, and psychology. The two sections below provide a brief overview of the development of some of the core enabling technologies.

Virtual Affective Agents

Implementing believable and engaging virtual agents requires a broad range of computational resources and capabilities, including: natural language processing (understanding and generation), speech processing (speech understanding and speech synthesis), dialog management, fundamental artificial intelligence (knowledge representation, automated reasoning, machine learning), affective computing (emotion recognition, expression of emotions by virtual characters, affective user modeling, emotion modeling, and cognitive-affective architectures), human–computer interaction, and computer graphics (3D modeling and animation). Clearly, it is beyond the scope of this chapter to discuss the numerous developments and achievements in these sub-disciplines over the past 50 years that have made the development of contemporary virtual agents possible. The brief historical overview below therefore highlights only a subset of these areas: embodiment, agent architectures, and affective computing.

Embodiment, in the context of virtual characters, refers to the type of visible appearance of an agent, depicted graphically on a display device (display devices range from small screens on mobile devices, to larger screens on iPads and laptops, large external displays, and finally to wall-sized virtual reality displays). Note that while most contemporary virtual agents are anthropomorphic, that is, they aim to resemble humans, this is not a necessary "feature." A number of engaging virtual agents have non-human, and even non-animal, form. In fact, one of the very first attempts at embodiment were the simple, non-anthropomorphic "Woggles" (see Figure 4.1).

Woggles were developed in the context of the OZ interactive drama project at Carnegie Mellon University (CMU) (Bates, Loyall, & Reilly, 1992). Woggles' agents were represented by oval shapes with eyes, and were capable of several simple movements (jumping, turning, sliding), squashing, puffing up, changing color, as well as eye movement, including tracking a moving object (Loyall & Bates, 1993). Woggles had some perceptual capabilities and could sense each other's presence, whether or not another Woggle was "looking" at them, and the presence of a human via a "sonar" sensor. In spite of their simplicity, and no attempt at visual, anthropomorphic realism, Woggles were able to convey some degree of affective realism and the users experienced a high-degree of engagement. Woggles' behavior was controlled by a relatively sophisticated agent architecture (the Tok architecture) (Reilly & Bates, 1992), that included both reactive and goal-directed behavior control, as well as an emotion modeling component that was based on the Ortony Clore Collins (OCC) model (discussed later) (Ortony, Clore & Collins, 1988).

As interest in anthropomorphic user interface agents grew, researchers began to explore various forms of agent appearance, ranging from highly simplified, cartoonish agents to fully articulated, fully embodied (head, upper torso, lower torso) 3D agents, aiming for a high-degree of visual realism (refer to Figures 4.1—4.3). Examples of the former include nutrition and exercise coach Harmony (Hayes-Roth, 2009; Hayes-Roth & Saker, 2010), virtual mindfulness coach Chris (Hudlicka, 2013), information presenter agent PPP persona (André, Rist, & Muller, 1998), medical education tutor Adele (Rickel & Johnson, 1999), and cartoonish characters in a system developed to help children cope with bullying (see character from FearNot! in Figure 4.1). Examples of more visually realistic agents include information presenter agent Jack (Norman & Badler, 1997) and instructional agent Steve, that was designed to provide naval personnel with system diagnostics training (Rickel & Johnson, 1999). An example of an agent aiming for a high degree of visual realism is the fully embodied agent Max (Becker, Nakasone, Prendinger, Ishizuka, & Wachsmuth, 2005), which has been evaluated in multiple contexts and roles, including a museum guide and a gaming partner (see Figure 4.1).

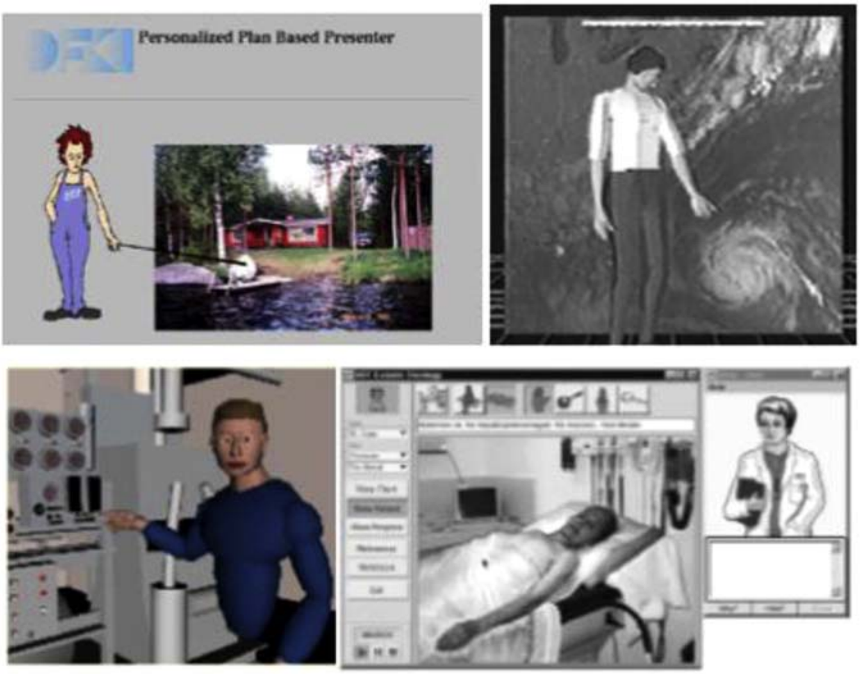

Figure 4.2 Examples of virtual agents illustrating different degrees of embodiment complexity (cartoonish vs full-body; simple vs anatomically based animation) and different interaction contexts: Presenter agents PPP Persona (upper left; André et al., 1998) and Jack (upper right; Norman & Badler, 1997); instructional agents Steve (lower left) and Adele (lower right).

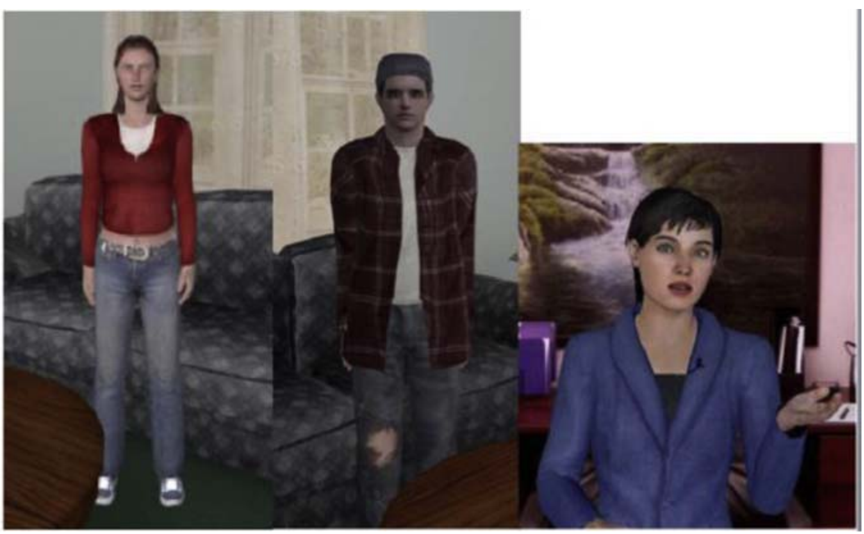

Figure 4.3 Examples of anthropomorphic, fully embodied virtual agents aiming for a high degree of visual realism. From left to right — virtual patients Justina and Justin (Rizzo et al., 2012) and counselor Amy (Lisetti et al., 2013).

These early agents laid the foundations for contemporary *embodied conversational agents* (ECAs) (Cassell, Sullivan, Prevost, & Churchill, 2000), with increasingly complex and often highly visually realistic, anthropomorphic embodiments, multi-modal interaction capabilities including conversational capabilities, complex dialog management and nonverbal expressions (e.g., facial expressions, head movement and gaze, gestures, body language), and competence in a specific domain. Examples of ECAs are agents Greta (De Rosis, Pelachaud, Poggi, Carofiglio, & De Carolis, 2003) and Laura (Bickmore & Picard, 2005) (refer to Figure 4.1). Interaction with ECAs requires minimal or no training, due to their ability to engage in naturalistic interactions with humans through dialog and nonverbal expression. The ability to interact with human users without prior training is especially critical for populations such as the elderly, children, and individuals with disabilities, and therefore significantly extends their usefulness.

These complex, and increasingly visually and socially realistic, agents are also being integrated into both serious and entertainment games as NPCs and player avatars.

Serious Games

Serious games have been developed across diverse contexts and cover a wide range of domains, including: language and cultural skills, most academic subjects, sports and motor skills ("exergaming"), a variety of corporate and business skills, policy planning, and military and law enforcement training (Hudlicka, 2011a). In healthcare, serious games are used both for training and education purposes for healthcare professionals (e.g., emergency room triage, dental surgery, nursing care), and for therapeutic and educational purposes for patients and clients (e.g., education about diet and exercise, coaching to increase physical activity, pain management, social skills training) (Horne-Moyer, Moyer, Messer, & Messer, 2014).

It is beyond the scope of this chapter to review the rapid development of computer gaming technologies, which have progressed from simple 2D games such as Pong, PacMan, and Tetris to today's highly realistic and engaging, often multi-player, games. The realistic graphics, depicting complex, 3-D environments, visually and increasingly even affectively realistic NPCs, and real-time, multi-modal interaction contribute to levels of user engagement which are closer to immersive virtual reality and participatory film experiences than simple gameplays. The visual realism and dynamic interaction are supported by a range of enabling technologies that include

3D modeling and animation, simulation, artificial intelligence, and affective computing. Player engagement is further enhanced by real-time performance and complex gameplays and storylines, facilitated by advances in the underlying hardware, including minimally intrusive input devices (e.g., Microsoft Kinect), and networking technologies that support multiplayer games. Game engines have been developed to facilitate rapid and systematic game development (e.g., Unity) and specialized engines supporting the development of affective games are being explored (Popescu, Broekens, & van Someren, 2014; Hudlicka, 2009). Increasingly, adaptive gaming is being implemented, where the gameplay is modified in response to the players' changing state and goals. These adaptations include game level difficulty as well as actual game content. This development is particularly critical for serious games and is facilitated by user modeling technologies.

The agent technologies discussed above further contribute to the games' realism and users' engagement, as both the NPCs and player avatars can increasingly behave in an autonomous manner, and display socially and affectively realistic expressions and behavior. The former facilitated by a variety of artificial intelligence knowledge-based systems and planning techniques, the latter by affective computing technologies. Affective computing (Picard, 1997) is a sub-discipline of computer science addressing four core areas: emotion recognition by machines, emotion expression by agents and robots, affective user modeling, and emotion modeling and cognitive-affective agent architectures (Hudlicka, 2008b).

In fact, probably the most significant advances in gaming that are especially relevant for serious games in health care, particularly therapeutic games, are techniques from affective computing that facilitate the development of affective games and affect-adaptive games (Gilleade, Dix, & Allanson, 2005; Hudlicka, 2008a; Yannakakis & Paiva, 2014). This topic is discussed in more detail in the following section.

STATE OF THE ART

To the users, the most apparent aspects of virtual affective agents are their interactional capabilities and their appearance. Interactional capabilities include both verbal and nonverbal interaction. Appearance is defined by the agent's embodiment, which of course also determines the types of nonverbal interaction possible. To be effective, affective agents must also be believable. This global characteristic can be further deconstructed into several distinct capabilities, including verbal and nonverbal interaction and conversational abilities, affective and social realism, and the ability to adapt and personalize the interaction to the human user's needs (Ortony, 2001). This section discusses the state-of-the-art of these capabilities in more detail.

Nonverbal Interaction: Emotion Recognition, Emotion Expression, and Agent Embodiment

Nonverbal interaction is of course critical in any human relationship, and may be even more important than verbal interaction in "helping relationships" (e.g., coaching, teaching, support, therapy), as it is an essential component of conveying emotional understanding and empathy. It is also important for managing conversation flow and dialog, since many critical conversational cues are nonverbal, such as showing understanding or confusion, agreement or disagreement. Nonverbal interaction in humans is mediated via a number of expressive channels associated with the human body; primarily the face and the myriad of facial expressions that convey emotional and other mental states, but also eye gaze and head movement, hand gestures, body posture, and body movement. In human–agent interaction, nonverbal interaction requires monitoring and sensing of the user's nonverbal expressions and recognizing these in terms of states that are relevant for the interaction, primarily emotional states.

On the output side, nonverbal interaction is closely linked to agent embodiment, since some type of a physical form is necessary to convey nonverbal information (with the exception of speech, which allows nonverbal information to be conveyed via prosody). Again, the ability to express emotions by the agent is critical here, or, more precisely, to display visible manifestations that would be interpreted as emotions by the human user. (No claim is made that virtual agents can or should experience emotions.) Nonverbal interaction, emotion recognition and expression, and agent embodiment are therefore discussed together below.

Emotion Recognition and Expression

Recognition of human emotions by computers, and expression of "emotions" by agents and robots, represent two of the core areas of affective computing. The two processes can be viewed as inverses of each other. Both rely on the identification of semantic primitives (features) associated with each expressive channel, to facilitate the data analysis, the recognition process, and the rendering of the expressions (Hudlicka, 2005). Detailed discussion of these topics is beyond the scope of this chapter and the reader is referred to a number of recent reviews of emotion recognition and expression technologies (Gunes, Piccardi, & Pantic, 2008; Zeng, Pantic, Roisman, & Huang, 2009). A brief summary of the state-of-the-art is provided below.

In *emotion recognition*, human user data, typically from multiple channels (face, speech prosody, head movement, body movement, physiology) are analyzed to identify the user's emotional state. The computational task consists of mapping the filtered data onto a set of emotional states, via some pattern recognition classification algorithms. Currently, the recognized emotions are typically limited to the basic emotions (e.g., joy, sadness, anger, fear), but progress is being made in recognizing more complex, typically social emotions, such as shame, guilt, and pride. Success rates in forced-choice conditions for the basic emotions are beginning to approach those of humans (Zeng et al., 2009). Recognition of naturalistic emotions, that is, emotions displayed in non-laboratory settings and unconstrained interactions, is more difficult, and represents one of the major challenges in emotion recognition research.

Recognition rates thus depend greatly on the modality and channels used, on the quality of the sensed data, the feature sets used as input to the classification algorithm, and the classification algorithms themselves. Increasingly, multi-modal approaches are being used, combining primarily the visual (facial expression) and audio (speech content and prosody) channels. Physiological signals are also being incorporated to increase recognition accuracy, primarily signals reflecting the activity of the autonomic nervous systems (arousal). Contextual variables (e.g., recent history of the interaction, the tasks the user is attempting to accomplish) can also help increase recognition rates (e.g., Kapoor, Burleson, & Picard, 2007). In health care, recognition of user states ("smart sensing") is particularly important, and can augment and even replace often unreliable self-report data.

In *emotion expression*, the agent's available expressive channels are used to convey an "emotional state" to the human user, to promote engagement. These channels are similar to the human expressive channels, with the exception of the physiological channel. The computational task consists of mapping the emotion the agent should express onto the set of expressive channels available in the agent's embodiment.

The *expression of emotions* in agents requires two distinct components. The agent must not only depict the emotion in an appropriate manner, along its available expressive channels (e.g., face, speech, movement), appropriately synchronized, and depicting realistic affective dynamics, but it must also display an appropriate emotion for the specific context. Achieving these goals involves very different sets of methods and technologies, with the former closely linked to computer graphics (rendering and animation technologies), and the latter involving emotion models and agent architectures (discussed later).

Whereas in emotion recognition multiple channels provide a redundant source of data and enhance recognition accuracy, in emotion expression multiple channels present a challenge, by requiring that expression be coordinated and synchronized across the channels, to ensure character realism. For example, expressions of anger must involve consistent signals in speech, movement and gesture quality, facial expression, body posture, and specific action selection, with the manifestations evolving and decaying at appropriate rates across the distinct channels.

A number of markup languages have been developed to facilitate emotion expression, by providing vocabularies of channel-specific semantic primitives (e.g., the Facial Action Coding System (Ekman & Friesen, 1978)) and an intermediate, implementation-independent representation of the expressions, and by facilitating coordination among multiple expressive channels. While no standard language has yet emerged, the EmotionML markup language represents a good candidate (Schröder et al., 2011).

Embodiment

As outlined above, virtual affective agents take on a wide variety of embodiments, ranging from a "talking head", through upper torso and arms, to fully articulated 3D bodies. Agent developers must make careful choices regarding the agent's embodiment, to balance the need for expressiveness, and the complexity of the associated processing, with the needs of the users within the given interaction context.

In terms of their appearance, most affective agents are rendered to resemble humans, and take on anthropomorphic, often highly visually realistic, embodiments. It is important to point out that while increasingly visually realistic anthropomorphic agents are being developed, visual realism is not necessary for affective realism, and may in fact interfere with affective realism and believability, as evidenced by the phenomenon of uncanny valley (MacDorman, 2005; Mori, 1970). Uncanny valley refers to a drop in perceived believability as the agent's visual realism increases and its appearance increasingly resembles a human. The reason for this apparently counterintuitive finding is that the unconscious criteria used to evaluate believability and affective realism shift as the appearance of the agent becomes more human-like. In effect, for more cartoonish characters, our expectations are lower. Once the synthetic agent begins to resemble human appearance, our evaluation criteria become more stringent, we begin to expect human-like realism and effectiveness in verbal and nonverbal interaction, and when these expectations are not met, we feel that the character is not believable, or, worse, even disturbing.

To address this problem, researchers are exploring agent embodiments that are anthropomorphic but not aiming to be visually realistic, and several studies suggest that simpler, 2D, often cartoonish characters can be very effective in conveying nonverbal information and promoting engagement, particularly with children (Dautenhahn & Werry, 2004; Paiva et al., 2005). Finding the optimal balance between expressiveness afforded by highly visually realistic embodiments and the risk of reaching the uncanny valley represents one of the core challenges in emotion expression research.

In summary, significant progress has been made in nonverbal interaction capabilities of IVAs. Facilitated by developments in affective computing, virtual affective agents are increasingly able to recognize basic human emotions and display affective expressions (again, typically limited to a few of the basic emotions) to their human users, to enhance engagement and effectiveness. Agent researchers are exploring the specific elements of nonverbal interaction that contribute to more realistic, engaging interaction. This research is not only contributing to more effective agents, but also contributing to our understanding of complex phenomena such as empathy. For example, Kang et al. (2012) found that specific nonverbal behavior correlates with disclosure of intimate information in humans (more head tilts (vs nods), less eye contact and more gaze aversion, and more pauses) and that similar behavior in a virtual agent is perceived by a human as indicating the disclosure of intimate information, which can enhance a sense of intimacy and improve the therapeutic alliance.

Believability, Affective Realism, and Emotional Intelligence

The notion of believability comes from the narrative and dramatic arts (film and animation, drama, and literature), and is related to two other notions: "illusion of life" and "suspension of disbelief." The term *believability* was first used in agent research by Bates, who introduced the notion of a "believable software agent" (Bates, 1994). Believability includes a number of agent attributes, including domain competence, appearance, and affective and social realism. Two of the most critical components of believability are *affective realism* and *emotional intelligence*. *Affective realism* refers to the agent's ability to generate and display emotions that are appropriate for the interaction context, and significantly contributes to the effectiveness of user—agent relationships, particularly coaching relationships (Bickmore, Gruber, & Picard, 2005; Hayes-Roth, 2009). As discussed above, this requires that the agent react affectively in real-time, and display affective expression along all available expressive channels, in a coordinated fashion (e.g., speech content and prosody matches facial expression and gestures), and display realistic affective dynamics over time (i.e., onset and decay rates of emotions, or transitions between emotional states, should be consistent with the agent's personality).

Emotional intelligence refers to a broader set of affective skills that includes perceiving and recognizing emotions in self and others, understanding the causes and consequences of emotions (in self and others), regulating emotions in the self, and managing emotions in others (Mayer & Salovey, 1995). Emotionally intelligent agents are expected to not only recognize the users' emotions, but understand their meaning within the given context and determine the best response, affective and nonaffective, that will help the users achieve their goals. Emotional intelligence is of course particularly critical for agents used in behavioral care, where the user's emotional state is often the focus of the interaction.

Personalization and Adaptation Capabilities

Much as we become more effective and engaging communicators as our knowledge of our interaction partner increases, so can affective agents become more effective as they "get to know" the user's preferences and idiosyncratic characteristics. To be engaging, an affective agent's interaction should be personalized to the user's preferences, needs, and goals, and adapt to his/her changing cognitive, emotional, and possibly physical state (e.g., fatigue, disability). A key benefit of virtual agents is that both their appearance (ethnicity, gender, age, style of dress) and their behavior (personality, emotional expressiveness, manner of addressing the user, small talk topics, degree of self-disclosure) can be personalized. For example, an agent interacting with a teen might address the user informally ("Hey, are you ok?"), whereas an agent interacting with an elderly user would be more formal ("How are you doing today, Mrs. Smith?").

Personalization and adaptation (user-adapted interaction) refer both to the ability to customize the agent's appearance and behavior to the user's preferences, and to the agent's ability to adapt its behavior to the user's long-term goals and needs, and the user's changing state and short-term goals. User-adapted interaction is mediated by knowledge about the user stored in an *affective user model*, which is dynamically constructed and updated during the interaction. Affective user models contain information about the types of emotions the user is likely to experience in a given context, an affective user profile (Hudlicka, 2011a) and how those emotions are manifested by the user, within the modalities the agent is able to sense (Hudlicka & McNeese, 2002). Since affective behavior can be highly idiosyncratic, affective models typically involve a learning component, which enables the identification of characteristic affective patterns by tracking the user's behavior over time. For example, frustration may be manifested in User A by agitated speech, but in User B by increasingly long delays between utterances. Significant existing research in intelligent tutoring systems provides the knowledge and methods supporting affective user modeling (e.g., D'Mello, Craig, Witherspoon, Mcdaniel, & Graesser, 2008; Forbes-Riley, Rotaru, & Litman, 2008; Mcquiggan, Mott, & Lester, 2008;Yannakakis, Hallam, & Lund, 2008).

Alternatively, affective user models can attempt to reconstruct the user's mental apparatus, including internal constructs such as beliefs, expectations, goals, and plans, and attempt to simulate the user's actual cognitive-affective processing, within the context of the specific gameplay. Such models then provide a "deeper" representation of the user's mental architecture. In other words, the affective user model may try to simulate the user's own appraisal of the on-going situation, and infer his/her goals and beliefs. These models then begin to resemble cognitive-affective architectures, and may be capable of simulating some of the user's own cognitive/affective processing; e.g., the cognitive appraisal processes resulting in a particular emotion. Clearly, the construction of these types of models is much more challenging than simple recognition of patterns in user behavior and research in this area represents one of the current challenges in affective computing and affect-adaptive interaction.

Putting It All Together: Agent Architectures

The broad range of capabilities outlined above requires a number of distinct functionalities and their coordination, as well as a significant amount of domain knowledge and user-specific data. Agents therefore require an associated system architecture, referred to as *cognitive agent architecture* or,

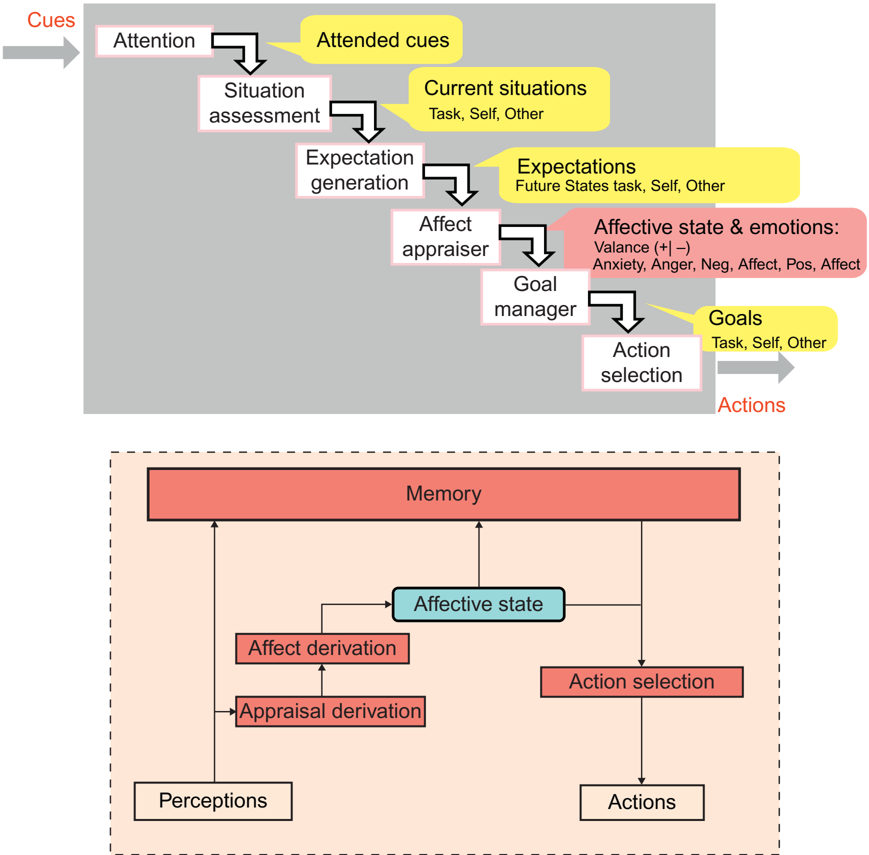

Figure 4.4 Examples of cognitive-affective agent architectures (top) MAMID architecture, implementing a see—think/feel—do sequence and focusing on modeling of emotion effects on cognition and personality-induced affective variability (Hudlicka, 2002, 2003); (bottom) FaTIMA architecture (Paiva et al., 2005).

when including affective processing, *cognitive-affective agent architecture*. An agent architecture consists of several modules, each responsible for a particular aspect of the agent's processing or behavior. Figure 4.4 shows two examples of agent architectures, illustrating the types of modules typically included and their interconnections.

At an abstract level, an agent architecture performs a "see—think—do" sequence: it senses and processes external data ("see"), engages in situation assessment, problem-solving, decision-making, and planning ("think"), and executes its plans to accomplish its goals, including verbal and nonverbal interaction goals and behavior in a simulated environment ("do"). In affective agents, this sequence is augmented to "see—think/feel—do", and affective agents therefore also implement a subset of a range of emotion-specific functions, as outlined above: sensing and recognition of the user's emotional states, affective user modeling, generation of the agent's emotion in response to particular interactions or situations (emotion modeling), and expressing this emotion to the user. Emotion recognition, affective user modeling, and emotion expression were outlined above. Below is a brief description of emotion modeling in cognitive-affective architectures.

Emotion modeling enables the agents to generate emotions dynamically, in response to evolving situations, both interpersonal situations (e.g., an affective reaction by the user) and situations in the simulated environment the agent may be sharing with the human user. Once an emotion is generated within the agent, it can be expressed to the human user, and it can influence both internal processing within the agent architecture, much as human emotions influence cognitive processes, and exert an effect on behavior and action selection. Most agent emotion models represent a subset of the basic emotions (e.g., joy, fear, anger, sadness), but increasingly researchers are attempting to develop models of more complex, social emotions (e.g. Becker, Tausch, & Wagner, 2011). These are, of course, particularly important for affectively realistic agents in behavioral health applications.

From a computational perspective, emotion modeling requires two categories of processes: those mediating emotion generation, and those mediating emotion effects (Hudlicka, 2008b, 2011b). The former is required to dynamically generate affectively and socially realistic agent behavior. The latter is necessary for translating particular emotions into the associated patterns of expression, in terms of the multiple channels available (e.g., speech, facial expression, gestures), and, possibly, to implement the effects of emotions on internal processing (attention, perception, memory, and reasoning).

In the majority of existing agent architectures, the emphasis is on emotion generation, and the subsequent relatively simple mapping of the emotion into the available expressive channels (e.g., agent's facial expression, hand gestures) and the agent's behavior. Emotion generation is typically implemented via models of cognitive appraisal (André, Klesen, Gebhard, Allen & Rist, 2000; Bates et al., 1992; Broekens, DeGroot, & Kosters, 2008; Reilly, 2006). Most of these appraisal models are based on either the OCC model (Ortony, 2002), or the explicit appraisal dimension theories developed by the componential emotion theorists (Smith and Kirby, 2000; Scherer, Schorr, & Johnstone, 2001) (e.g., *novelty, valence, goal relevance and congruence, responsible agent, coping potential*).

Typically, symbolic AI methods are used to implement the stimulus-to-emotion mapping, whether this is done via an intervening set of appraisal dimensions, or directly from the domain stimuli to the emotions. In general, the complexity of this process lies in analyzing the domain stimuli (e.g., features of a game situation, behavior of game characters, player behavior) and extracting the appraisal dimension values. This may require representation of complex mental structures, including the game characters' and players' goals, plans, beliefs, and values, their current assessment of the evolving game situation, and expectations of future developments, as well as complex causal representation of the gameplay dynamics. Rules, semantic nets, and Bayesian belief nets are some of the frequently used formalisms implementing this mapping.

Much less frequently, emotion effects on cognition are included in agent architectures, and modeled via parametric representations of distinct emotion states and their effects on the processing within the architecture modules (Hudlicka, 1998, 2002, 2003, 2008c). Models of emotion effects on expressive behavior and action selection are typically implemented via direct mapping of a particular emotion onto specific behavior, or patterns of expression along one or more modalities; e.g., if "happy" then "jump" and "smile." Some models implement mappings onto components of expressive behaviors, rather than fully formed expressions; mapping values of individual appraisal dimensions onto elements of facial expressions, e.g., eyebrow position (Scherer, 1992).

It is beyond the scope of this chapter to discuss the extensive work in emotion modeling that supports increasingly sophisticated affective modeling in agent architectures, including models of emotion effects on cognitive processing, to implement affective biases (Hudlicka, 2008c), and models of emotion regulation and emotion contagion (Bosse, Broekens, Dias, & van der Zwaan, 2014). For a recent review of existing work in emotion modeling see Hudlicka (2011b, 2014a) and affective agent architectures see Hudlicka (2014b).

Nonplaying Characters and Player Avatars

Not every game has or needs characters. For example, the popular game Tetris does not have any characters, and a number of serious games focusing on skills training do not include NPCs. However, most entertainment games, and many serious games, require or benefit from the presence of game characters. Characters in games fall into two categories: NPCs and player avatars. NPCs are characters controlled by the game software, which interact with the player and play various roles within the gameplay, as determined by the game plot and narrative. NPCs vary in their degree of autonomy and sophistication, increasingly take advantage of the technologies developed in intelligent agents and affective computing research, and are becoming more autonomous and believable, thereby contributing to a more engaging player experience. Player avatar characters represent the player within the context of the game.

Affective and Affect-Adaptive Gaming

Recognizing the key role that emotion plays in learning, in the training of new cognitive and affective skills, and in the acquisition of new behavioral skills, and the elimination of undesirable behaviors (e.g., addictions), developers of serious games are increasingly recognizing the importance of affective gaming (Sykes, 2004; Gilleade et al., 2005). Gilleade and colleagues (2005) captured the objectives of affective gaming in a succinct statement, describing a progression of functionalities an affective game should support: "Assist me, Challenge me, Emote me."

Affective gaming focuses on the integration of emotion into game design and development, and includes the following areas: recognition of player emotions, adaptations of the gameplay to the players' affective states, and modeling and expression of emotions by NPCs (Hudlicka, 2008b). Affective gaming thus emphasizes the use of virtual affective agents as NPCs and the use of affective player modeling and affect-adaptive gameplay.

Within the broad area of affective gaming, Hudlicka has previously suggested that a distinction be made between affect-sensitive games and affect-centered games. The former being games that recognize and adapt to the player's emotional state, whereas the latter are games where emotions play a central role, and whose explicit purpose is to train affective and social skills, or to aid in psychotherapy (Hudlicka, 2011a). Continued progress in the development of affective game engines will facilitate the development of affect-centered therapeutic games.

Overview of Recent Applications in Health Care

Serious games are being explored in health care across a number of contexts, including patient education and psycho-education, medical decision-making, support for behavior change (e.g., in exercise, nutrition, and specific health behaviors such as dental health), rehabilitation and physical therapy, pain reduction, and psychotherapy. Games are also increasingly used in the training of healthcare professionals. A recent review of serious games in health care identified over 20 serious games used in training of medical professionals, and some intended for use by patients; e.g., for pain management, decision-making, patient education, and health behavior coaching (Ricciardi & De Paolis, 2014). Examples of games developed for patients and clients include "Escape from Diab" (archimage.com) and "Squire's Quest" (http://www.squiresquest.com), health behavior coaching games to help children adopt healthy eating and exercise habits, and a game for motor rehabilitation following stroke or brain trauma (Burke et al., 2009).

While the majority of existing serious games in health care, whether for professional training or for patients, focus on physical health (e.g., surgery, emergency medicine, pain management), games focusing on behavioral care are beginning to emerge (Brezinka & Hovestadt, 2007). These include social skills training (Beaumont & Sofronoff, 2008); support for children experiencing divorce, based on family therapy (www.ziplandinteractive.com); cognitive-behavioral treatment of OCD (Brezinka, 2008); game to help veterans overcome posttraumatic stress disorder (PTSD) (Rizzo et al., 2010); and a game designed to motivate adolescents for solution-focused therapy ("Personal Motivator" (Coyle, Matthews, Sharry, Nisbet, & Doherty, 2005)).

An important distinction must be made here between games designed for entertainment only that are used in health care and behavioral care contexts as adjuncts to traditional therapy vs serious games developed expressly for the purpose of direct administration of healthcare interventions and education, including interventions for support, symptom reduction, and behavior change in psychotherapy.

Examples of the former are any entertainment games where the gameplay is incorporated into traditional treatment. For example, games developed solely for entertainment purposes (e.g., Tetris) can be used to help train children with ADHD in attention focusing strategies and first-person shooter (FPS) games have been used to help patients visualize their immune system attacking tumors (Sajjad, Hanan Abdullah, Sharif, & Mohsin, 2014). A number of entertainment games are also used for physical rehabilitation purposes (e.g., WiiFit, Dance Dance Revolution). (The effectiveness of these rehabilitation games is greatly enhanced by input devices, which facilitate gross motor movement exercise and tuning (e.g., Dance Dance Revolution dance mat, WiiMote, and Microsoft Xbox Kinect).)

Examples of the latter are games designed to address a particular symptom (e.g., obsessive-compulsive behavior) or teach a particular emotional or social skill (e.g., emotion regulation, handling difficult social situations such as bullying). These are the games that are emphasized in the discussion below.

The majority of these games are in the exploratory stages, with only a few pilot studies having been conducted to date to evaluate the games' effectiveness. This is especially the case for psychotherapeutic games, which represent both the most challenging applications of the serious gaming technologies (vs psycho-education, training), and also the most promising. Ricciardi and De Paolis (2014) report that only nine of the reviewed games were subjected to an evaluation study, with six of the studies reporting improved training results with game-based training compared with traditional training. Examples of serious games in behavioral care are discussed below.

Pediatric Pain Management — "Free Dive"

Free Dive is a game designed to help children control pain during medical procedures (http://www.breakawaygames.com/serious-games/solutions/healthcare/; Figure 4.5). Free Dive presents a visually rich, underwater environment to the player, whose objective is to swim around with the fish and sea turtles in search of a hidden treasure. The game concept is based on the underlying assumption that distracting attention away from painful stimuli reduces the subjective experience of pain. Data from a pilot study suggest that children are able to tolerate painful procedures more easily

Figure 4.5 Examples of serious games for children: "Free Dive" (game to control pain by serving as a distraction activity during medical procedures) (left) and "Ricky and the Spider" (game to treat OCD) (right).

while playing the game. Another pain management game was developed for adults, based on EEG technologies and biofeedback principles (Sourina, Wang, & Nguyen, 2011).

OCD in Children – "Ricky and the Spider"

Ricky and the Spider is a game designed to treat OCD in children (Brezinka, 2008; Figure 4.5). The game is based on a cognitive-behavioral approach to OCD treatment. The players are provided psycho-education about the condition, and are supported in creating a hierarchy of symptoms and provided with opportunities for simulated exposure and response prevention: these are established, evidence-based approaches to treating OCD. The players are also taught techniques for externalizing their symptoms, as a means of reducing anxiety. The game characters include "Dr. Owl", who provides advice and guidance (a stand-in for a therapist), "Spider", who represents the OCD condition and issues commands for the other game characters to engage in OCD behavior, and two characters with OCD symptons: "Ricky the grasshopper", who must hop in a specific pattern, and "Lisa the Ladybug", who must count her polka dots before falling asleep. The Spider threatens Ricky and Lisa with terrible consequences if they do not follow his orders. The game is played by a child under the supervision of a therapist.

The game is being used by several thousand users across more than 40 countries. Data from an initial evaluation (Brezinka, 2013), conducted with 18 children and 13 therapists, were positive, with the children reporting satisfaction with use of the game during treatment, and the therapist reporting that the children enjoyed playing the game and experienced increased motivation for treatment. The symptoms of OCD were significantly reduced in 15 children, and unchanged in one child.

PTSD in War Veterans – "Virtual Iraq"

Virtual Iraq is, strictly speaking, a virtual reality (VR) environment, rather than a game. However, due to the increasing overlap between these two types of technologies, and the influential role that Rizzo's work has played in computerized technologies for behavioral health, Virtual Iraq is included here.

Virtual Iraq implements a VR-mediated exposure therapy approach to the treatment of PTSD in veterans returning from the Iraq and Afghanistan wars. Exposure therapy is a behavioral therapy treatment of

PTSD, based on systematic desensitization and habituation. PTSD patients are provided with head-mounted displays and exposed to realistic wartime scenarios in Iraq. Sound and motion are included, to simulate explosions and shaking vehicles. Some scenarios include smells. The scenarios include multiple locations (desert, city), and ranges of stimuli intensity and potential (simulated) danger. The patients are exposed to increasingly intense scenes, under the supervision of a therapist. An early evaluation study showed promising results, with 16 out of 20 patients not meeting the criteria for PTSD following treatment (Rizzo et al., 2010).

These initial results are promising and some existing reviews suggest that VR-based exposure therapy is as effective as traditional exposure therapy (Goncalves, Pedrozo, Coutinho, Figueira, & Ventura, 2012) and may be even more effective in some situations (Malbos, Boyer, & Lançon, 2013). This is a significant finding, since computer-delivered therapy has a number of advantages, most importantly cost and accessibility, and may even be a preferred mode of treatment for some patients. However, as is the case with most of the technologies reviewed in this chapter, systematic, randomized controlled trials are still lacking in this area and there is need for more standardization in the therapeutic environments and treatment manuals (Motraghi, Seim, Meyer, & Morissette, 2014).

Social and Emotion Regulation Skills for Children on the Autism Spectrum – "Secret Agent Society"

Secret Agent Society is a game designed to teach children and adolescents on the autism spectrum social skills and emotion regulation skills (Beaumont & Sofronoff, 2008; Figure 4.6). The specific skills include the fundamental emotion intelligence skills (recognition of emotions in self and others) and emotion regulation strategies, as well as a variety of social problem-solving skills (e.g., how to deal with teasing, bullying, losing a game).

The game was developed as a component of an integrated social skills training program for children with Asperger syndrome: the "Junior Detective Training Program." (The other components of the program include group training in social skills and parent and teacher training.) Within the game, the child acts as a junior detective, whose task is to recognize the suspects' emotions. The game progresses through several levels of difficulty, with the lower levels involving recognition of emotions in game characters from their facial expressions, body posture, and speech prosody. In higher levels, the children are taught how to recognize more complex social emotions (guilt, embarrassment). In the

Figure 4.6 Secret Agent Society – game designed to teach children on the autism spectrum social and emotion regulation skills.

highest level of the game, children are asked to complete missions where they learn and practice strategies for dealing with bullying and appropriately playing with others. For example, the player might feel frustrated at losing a game and must identify the emotion (frustration) and then select a strategy for dealing with the emotion in a socially acceptable manner. The player is provided with multiple-choice options and asked to select the best response for the situation; e.g., congratulate other player, yell at other player, talk to a mentor agent in the game, jump on a trampoline, or punch a wall. The children are also taught how to recognize their own emotions via physiological self-assessment scales. To help children learn emotion regulation skills, the game presents a series of anxiety-inducing situations (e.g., flying an airplane) and possible coping strategies for anxiety control. The player chooses a coping strategy and is provided with feedback regarding its appropriateness for the situation.

An earlier version of the game was evaluated via a randomized controlled trial ($N = 49$; 7.5–11-year-olds) over a 2-month period, and yielded positive results: 76% of the children not only improved their social functioning and reached the level of neurotypical children, but these gains were maintained over a 5-month period (Beaumont & Sofronoff, 2008). Maintaining social skills gains following training, and generalizing to non-training contexts, is a major challenge for children on the autism spectrum.

The data regarding the maintenance of gains following game-based training are therefore particularly encouraging, and reflect the promise that gaming technologies provide for this population.

APPLICABLE ETHICAL AND PRIVACY CONSIDERATIONS

Any technology that has the potential to track human behavior, and collect personal data, must address privacy and data security considerations. These considerations have been codified into the legal system via the HIPAA regulations, which are actively enforced. Of course, there is the very serious issue of unwarranted and potentially illegal government spying, which makes any HIPAA regulations meaningless. In this regard, it is quite surprising that no professional organization in behavioral health has yet raised this issue, as it has a direct impact on health care, and particularly on telemedicine and telemental health.

However, the issues of data privacy, data sharing, and security are addressed elsewhere in this book, and are not specific to virtual affective agents and therapeutic games. This section therefore focuses on ethical and privacy considerations that are specific to these technologies. These contexts, in conjunction with the significant emphasis on the user's emotion, raise an entirely new set of ethical considerations: *affective privacy, the ethics of affect induction in users, and the ethics of virtual relationships*, that is, relationships with virtual affective agents.

Affective Privacy

Our emotions, perhaps even more so than our thoughts, are likely the most personal and private aspects of our lives. The development and use of applications that sense, infer, or track our emotions therefore presents considerable and as yet unexplored ethical challenges. This is especially the case in any applications in behavioral health, where the users may be addressing a particularly painful experience, or reveal emotions and thoughts that could have negative repercussions on their lives if they were made public, or revealed to other parties (e.g., employers, insurance companies). User modeling, and especially affective user modeling, thus presents an ethical challenge since these models may contain the most guarded personal information about the users: the emotions they feel, including "undesirable" emotions; the events that trigger those emotions, including triggers that may be considered inappropriate, etc.

Technological advances are increasingly facilitating the collection of affective data, often without our knowledge: facial expressions can be recognized from data collected by the ubiquitous video cameras that watch most of our activities in public spaces and in the workplace; our posts on social networks can be, and are being, analyzed, to "mine" affective data; our telephone conversations are recorded. We have no control over these activities and no option to opt out. Even in situations where we willingly agree to have our data collected, there are privacy concerns: data are lost, compromised, stolen. A number of technologists are predicting an end to privacy. However, we interpret this based on historical possibilities, and thus our worst privacy breach nightmare might be that our financial data will be compromised. We do not generally expect that data about our most intimate mental and affective states are being collected, analyzed, stored, and perhaps used. And yet these are precisely the types of data that will be collected as behavioral healthcare applications begin to proliferate.

The existing concerns about data security and privacy protection are thus multiplied when we consider the personal nature of data now being collected. As remote sensing and emotion recognition technologies continue to advance, it will be increasingly possible to intrude into our private lives. It is therefore essential that healthcare professionals actively address the ethical challenges associated with the privacy of affective data, and become active and proactive, both within their specific professional organizations and more broadly at the political level, as government surveillance increases.

Emotion Induction

Both virtual affective agents and serious games not only have the potential to induce strong emotions, but may often be designed with the express purpose of inducing specific emotions, including negative emotions. Virtual agents acting as coaches may be designed to induce affection, so that they are viewed as empathic, that their message is trusted, and they can be more persuasive (e.g., induce behavior change such as more exercise, less smoking, better nutrition). In some behavioral healthcare applications, agents and games may be designed to induce a negative or unpleasant emotion; for example, when the technology aims to implement exposure-based treatments that require some degree of anxiety to be effective, such as PTSD or OCD treatments based on exposure.

Of course, computers already induce a range of negative emotions on a daily basis, without explicitly attempting to, and computer games designed for entertainment also induce a range of emotions, including frustration, disappointment, even rage. However, the use of these technologies in behavioral health, where emotion manipulation may be desirable and necessary to achieve the therapeutic goals, presents a unique set of ethical concerns.

How can we ensure that the induced emotions will not overwhelm the user, or have a deleterious effect in the future? What if systematic desensitization occurs to the wrong stimulus (e.g., simulated violence or abusive behavior)? What emotions can ethically be induced in the user? In therapeutic contexts, it is generally accepted that the induction of some negative emotions is not only acceptable or unavoidable, but even necessary for treatment (e.g., anxiety induction during exposure therapies; shame experience in a safe, supportive setting to help individuals reorganize some trauma-induced cognitive-affective schema).

Most existing active therapeutic systems are designed to be used under the supervision of a clinician (e.g., "Ricky and the Spider" OCD treatment game, The "Secret Agent Society" social skills and emotion regulation game, "Virtual Iraq" PTSD treatment environment). However, it is conceivable that users could gain access to these technologies without the supervision of a clinician, or that the supervising clinician might not be aware that a negative emotion of undesirable intensity was being induced.

One can go even further, and consider situations where an increasingly autonomous and sophisticated virtual affective agent, with a detailed affective model of the user, can misapply a therapeutic technique and induce an emotion in the user that will be traumatizing, or will undo previous gains made in treatment.

There are no easy answers to these ethical dilemmas and carefully monitored application of these technologies, along with extensive education of both the end-users and the professionals administering the technologies regarding the possible risks, will be essential.

Virtual Relationships

As stated earlier in this chapter, humans are "wired to relate", and one of the most powerful features of the affective agent technologies is precisely their ability to induce the "relational instinct" and attachment behavior in the human users (Reeves & Nass, 1996). Through the ensuing connection, and associated trust, agents can then provide support and coaching, even social companionship. By inducing attachment and trust, virtual affective agents have the potential to mimic aspects of human relationships and humans can thus, at least theoretically, enter into relationships with agents. These virtual relationships present yet another ethical challenge.

How does a human user know when to trust a virtual agent? If the agent appears "confident" (e.g., an eldercare companion agent telling the human to take a particular medication), is this an indication that it should be trusted? What about the risk of virtual relationships replacing actual human relationships? What if an elderly person finds their synthetic companion more compelling, and empathic, than his/her family members, and reduces contact with family, eventually investing emotions into a relationship with an entity not capable of experiencing real emotions, engaging in a emotionally responsible relationship, or, for that matter, bringing the elderly food or medicine? What if the agent uses its persuasive capabilities to induce beliefs or behavior which are not in the human user's best interests? One may think that these are outlandish possibilities, given the state-of-the-art of agent technologies. However, one has only to consider the fact that DSM-5 has identified both "Internet Addiction" and "Internet Gaming Disorder" as possible new Axis I diagnostic categories to realize that as the virtual agent technologies advance we may soon be facing a variety of new potential disorders to consider, including, perhaps, a "Virtual relationship addiction disorder."

Some of these issues have begun to be raised by several agent and affective computing researchers (Castellano et al., 2010; Bickmore, 2005) but the research community has a long way to go in adequately understanding and addressing the ethics of virtual relationships. However, attention is increasingly being focused on the issue of ethics in human−agent interactions, as evidenced, for example, by the 2015 workshop on "The Emerging Policy and Ethics of Human−Robot Interaction" held at the HRI conference in Portland Oregon, USA (http://www.openroboethics.org/hri15/).

FUTURE PROSPECTS

It is difficult, one might say even irresponsible, to make predictions in technological developments, and prognosticians in artificial intelligence, a core technology mediating both virtual affective agents and serious games, have a particularly dismal track record. It is also difficult not to be influenced by one's wishes (or concerns) when making predictions. With these caveats in mind, a discussion of near-term anticipated developments in virtual affective agent and serious gaming technologies is offered below.

Proliferation

Given the recent growth in both of these technologies, coupled with the recognition of their potential in behavioral health and the increasing need and acceptance of technologies in behavioral health, it is safe to assume that both affective agents and serious therapeutic games will continue to expand beyond research contexts and beyond the currently typical applications in coaching of health behavior. A good example of this trend is the increasing popularity of the Secret Agent Society game for children on the autism spectrum, which is integrated into a formal training program offered to clinicians worldwide. Use of therapeutic games will become more commonplace, particularly for children and adolescents, and gaming elements will be integrated into training, learning, and rehabilitation environments.

Formal Evaluations

Increasing emphasis will be placed on formal evaluations of therapeutic games and affective agents within these games. As these technologies mature, it will no longer be sufficient to demonstrate feasibility and user acceptance, and increasingly emphasis will be placed on formal, randomized controlled trials demonstrating effectiveness. These studies will in turn help determine the populations, disorders, and contexts within which these technologies are helpful. For example, while we may wish it to be the case that loneliness among the elderly can be reduced through the use of social robots, it is likely that these technologies will not be readily accepted by the current generation of elders.

Improved Understanding of Suitable Applications and Contexts for Agents and Games

As agents and games proliferate, and are increasingly subjected to more formal evaluations with end-users (clients, patients, trainees, students), we will begin to develop a better understanding of the benefits and the limits of these technologies, including:

- types of roles agents can effectively play in different contexts and for different skills; e.g., trainer, mentor, coach, virtual therapist;
- types of therapies for which virtual agents and games are most suitable; e.g., highly structured, protocol-based therapies that emphasize the acquisition of specific skills and do not rely on free-form natural language interaction, and the associated nuanced understanding and interpretation of complex constructs and mental and emotional states;

- types of conditions and symptoms amenable to agent- and game-mediated treatment; e.g., in psychotherapy, Axis I disorders, such as anxiety, depression, PTSD, and substance use for which evidence-based treatment protocols exist; e.g., dialectical behavior therapy (DBT), exposure and systematic desensitization, possibly motivational interviewing. In contrast, therapeutic approaches focusing on Axis II disorders, requiring gradual transformations of cognitive-affective schemas, and the context of a strong therapeutic alliance created over the long term, are unlikely to be facilitated by these technologies (although they may be useful as adjuncts to face-to-face therapy to facilitate particular skill training, such as emotion regulation).

Agents: Empathy and Personality

Much recent work in affective agents has focused on the development of empathic agents, capable of establishing an affective loop between the agent and a human user (e.g., Paiva et al., 2004). While short-term pseudo-empathic behavior may be feasible, the establishment of true empathy between agents and humans will remain a challenge. It is unlikely that the near future will see virtual affective agents with either the degree of affective realism, or the degree of cognitive competence, required to convey empathy to a human user during long-term interactions.

However, creating agents with distinct personalities will be increasingly possible, and research in this area is yielding promising results (e.g., Hudlicka, 1998, 2003). These types of agents will make it possible to create customized social environments within therapeutic games, tailored to the specific needs of the user; for example, create feared social situations to support cognitive-behavioral treatments for social anxiety.

Improved User State Recognition, Affective User Modeling and Personalization

One of the most rapidly advancing areas in affective computing is emotion recognition. Advances in both nonintrusive sensing hardware, and classification algorithms capable of automatic feature extraction and unsupervised learning, contribute to increasing accuracy. This, in turn, increases the accuracy of affective user modeling and facilitates affect-adaptive interactions between the user and the gaming environment, including any affective agents within the game. Increasingly, emphasis will be placed not only on single-frame *emotion recognition* (e.g., "player is feeling frustrated") but on longer-term *emotion understanding* (e.g., "player is feeling frustrated because the last three game tasks were too anxiety-provoking and s/he needs some encouragement to improve her sense of social competence"). These developments will contribute to increasingly personalized therapeutic games, targeted for specific symptoms, emotions, or skills, which will, hopefully, enhance their effectiveness and patient outcomes.

Natural Language Understanding

Although progress will be made in statistically based natural language understanding, and agents' ability to converse in restricted gameplay contexts will be much improved; these systems will remain "brittle": they will not be able to carry on the types of wide-ranging conversations that humans engage in. This will continue to limit agent interactions to narrow contexts, within which they may be quite effective, but will not enable more complex natural-language interactions, such as those required for psychotherapy. Clinicians worried about agents and robots taking over the world, and their jobs, can relax.

New Types of Relationships and Improved Understanding of Relationships

Virtual affective agents, whether stand-alone or used in the context of games, will create opportunities for humans to create unique types of relationships: human—agent relationships. Researchers are attempting to develop relational agents (e.g., Bickmore, 2003), capable of long-term relationships, and are using research in psychology and sociology to determine the types of functionalities relational agents need to be engaging. Human—agent interactions will provide valuable test beds for exploring the nature of both human—human relationships, and also the unique nature of human—agent relationships. New types of relationships may be identified, and we may be surprised to find that particular attributes we value in humans may not be valued in agents. The assumption that effective human—agent relationships need to emulate human—human relationships may therefore need to be questioned, including, perhaps such deep-seated assumptions that effective and engaging virtual agents need to be empathic. While for certain applications this will likely be the case, and preliminary research suggests that empathic agents are often more effective, it is also possible that we will see entirely different directions of development. As humans begin to coexist with virtual agents, new sets of expectations for human—agent interactions are likely to develop, and different expectations will develop for different contexts. We should not assume that empathic and affective agents will always be more desirable and effective, and researchers are beginning to identify contexts where nonempathic agents may be more effective.

CONCLUSIONS

Therapeutic games have a unique ability to engage the user by providing an immersive environment, directly engaging the user in a pleasurable activity, and by providing immediate positive feedback via a well-defined reward structure. The often-quoted "no pain, no gain" assumption about skill acquisition does not apply to games, where the players can acquire and practice skills effortlessly, often without the exertion associated with conscious effort.

Whether used as stand-alone tools to provide coaching, training, and support behavior change, or playing supportive roles in psychotherapy, these technologies can perform many functions and significantly enhance both the delivery and the training of behavioral healthcare interventions. Taking advantage of the innate human capabilities and desires to connect and to play, virtual affective agents and serious games provide immersive affective experiences, where skill acquisition happens naturally and effortlessly, through a series of engaging interactions.

A number of challenges exist, of course, both technological (e.g., development of affectively realistic agents) and ethical (ensuring privacy of affective data and understanding the benefits and potential undesirable side effects of virtual relationships). However, existing technological developments, the recognized need and desire by the research community to conduct formal evaluation studies, and, most importantly, the recognized need and promise of these technologies in behavioral health, suggest that both affective virtual agents and serious therapeutic games will represent a major growth area in the next few decades.

REFERENCES

André, E., Klesen, M., Gebhard, P., Allen, S., & Rist, T. (2000). *Integrating models of personality and emotions into lifelike characters* (pp. 150—165). *Affective interactions* Berlin: Springer.

André, E., Rist, T., & Müller, J. (1998). Integrating reactive and scripted behaviors in a life-like presentation agent. In K. P. Sycara, & M. Wooldridge (Eds.), *Proceedings of the second international conference on autonomous agents* (pp. 261—268). Minneapolis, MN: ACM Press.

Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, *37*(7), 122–125.

Bates, J., Loyall, A., & Reilly, W. (1992). Integrating reactivity, goals, and emotion in a broad agent. In *Proceedings of the 14th meeting of the cognitive science society.*

Beaumont, R., & Sofronoff, K. (2008). A multi-component social skills intervention for children with Asperger syndome: The Junior Detective Training Program. *Journal of Child Psychology and Psychiatry*, *49*, 743–753.

Becker, C., Nakasone, A., Prendinger, H., Ishizuka, M., & Wachsmuth, I. (2005). Physiologically interactive gaming with the 3D agent Max. In *Paper presented at the international workshop on conversational informatics at JSAI-05*, Kitakyushu, Japan.

Becker, J. C., Tausch, N., & Wagner, U. (2011). Emotional consequences of collective action participation differentiating self-directed and outgroup-directed emotions. *Personality and Social Psychology Bulletin*, *37*(12), 1587–1598.

Bickmore, T. (2003). *Relational agents: Effecting change through human-computer relationships* Cambridge, MA: MIT Press.

Bickmore, T. (2005). Ethical issues in using relational agents for older adults. In *Proceedings of the AAAI fall symposium*.

Bickmore, T., Gruber, A., & Picard, R. (2005). Establishing alliance in automated health behavior change interventions. *Patient Education and Counseling*, *59*, 21–30.

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

Bosse, T., Broekens, J., Dias, J., & van der Zwaan, J. (2014). *Emotion modeling: Towards pragmatic computational models of affective processes* Berlin: Springer.

Brezinka, V. (2013). Rick and the spider—A video game to support cognitive behavioral treatment of children with obsessive-compulsive disorder. *Clinical Neuropsychiatry*, *10*(3 Suppl. 1), 6–12 . Retrieved from: <http://www.clinicalneuropsychiatry.org/pdf/01Brezinka.pdf> .

Brezinka, V. (2008). Treasure hunt—a serious game to support psychotherapeutic treatment of children. *Studies in Health Technology and Informatics*, *136*, 71.

Brezinka, V., & Hovestadt, L. (2007). *Serious games can support psychotherapy of children and adolescents* (pp. 357–364). Berlin: Springer.

Broekens, J., DeGroot, D., & Kosters, W. A. (2008). Formal models of appraisal: Theory, specification, and computational model. *Cognitive Systems Research*, *9*(3), 173–197.

Burke, J. W., McNeill, M. D. J., Charles, D. K., Morrow, P. J., Crosbie, J. H., & McDonough, S. M. (2009). Optimising engagement for stroke rehabilitation using serious games. *The Visual Computer*, *25*(12), 1085–1099.

Castellano, G., Leite, I., Pereira, A., Martinho, C., Paiva, A., & McOwan, P. W. (2010). Affect recognition for interactive companions: Challenges and design in real world scenarios. *Journal of Multimodal User Interfaces*, *3*(1–2), 89–98.

Cassell, J., Sullivan, J., Prevost, S., & Churchill, E. (2000). *Embodied conversational agents* Cambridge, MA: MIT Press.

Coyle, D., Matthews, M., Sharry, J., Nisbet, A., & Doherty, G. (2005). Personal investigator: A therapeutic 3D game for adolecscent psychotherapy. *Interactive Technology and Smart Education*, *2*(2), 73–88.

Dautenhahn, K. (2007). Socially intelligent robots: Dimensions of human–robot interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1480), 679–704.

Dautenhahn, K., & Werry, I. (2004). Towards interactive robots in autism therapy: Background, motivation and challenges. *Pragmatics & Cognition*, *12*(1), 1–35.

D'Mello, S. K., Craig, S. D., Witherspoon, A., Mcdaniel, B., & Graesser, A. (2008). Automatic detection of learner's affect from conversational cues. *User Modeling and User-Adapted Interaction*, *18*(1–2), 45–80.

Ekman, P., & Friesen, W. V. (1978). *Facial action coding system* Palo Alto, CA: Consulting Psychologists Press.

Forbes-Riley, K., Rotaru, M., & Litman, D. J. (2008). The relative impact of student affect on performance models in a spoken dialogue tutoring system. *User Modeling and User-Adapted Interaction*, *18*(1–2), 11–43.
Gilleade, K., Dix, A., & Allanson, J. (2005). Affective videogames and modes of affective gaming: Assist me, challenge me, emote me. Retrieved from: <http://comp.eprints.lancs.ac.uk/1057/1/Gilleade_Affective_Gaming_DIGRA_2005.pdf>.
Goncalves, R., Pedrozo, A. L., Coutinho, E. S. F., Figueira, I., & Ventura, P. (2012). Efficacy of virtual reality exposure therapy in the treatment of PTSD: A systematic review. *PLoS One*, 7(12), e48469.
Gunes, H., Piccardi, M., & Pantic, M. (2008). *From the lab to the real world: Affect recognition using multiple cues and modalities* (pp. 185–218). Aarhus, Denmark: InTech Education and Publishing.
Hayes-Roth, B. (2009). Putting intelligent characters to work. *AI Magazine*, *29*(2), 43.
Hayes-Roth, B., & Saker, R. (2010). Automating individualized coaching and authentic role-play practice for brief intervention training. *Methods of Information in Medicine*, *49*(4), 406–411.
Horne-Moyer, H. L., Moyer, B. H., Messer, D. C., & Messer, E. S. (2014). The use of electronic games in therapy: A review with clinical implications. *Current Psychiatry Reports*, *16*(12), 520. Available from http://dx.doi.org/10.1007/s11920-014-0520-6.
Hudlicka, E. (1998). Modeling emotion in symbolic cognitive architectures. In *Proceedings of the AAAI fall symposium: Emotional and intelligent I*, Orlando, FL, AAAI Press.
Hudlicka, E. (2002). Increasing socially-intelligent architecture realism by modeling and adapting to affect and personality. In K. Dautenhahn, A. H. Bond, L. Canamero, & B. Edmonds (Eds.), *Multiagent systems, artificial societies, and simulated organizations*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
Hudlicka, E. (2003). *Modeling effects of behavior moderators on performance: Evaluation of the MAMID methodology and architecture* Phoenix, AZ: BRIMS-12.
Hudlicka, E. (2005). Affect sensing, recognition and expression: State-of-the-art overview. In *Proceedings of the first international conference on augmented cognition*, Las Vegas, NV.
Hudlicka, E. (2008a). Affective computing for game design. In *Proceedings of the 4th international North American conference on intelligent games and simulation (GAMEON-NA)*, Montreal, Canada.
Hudlicka, E. (2008b). What are we modeling when we model emotion? In *AAAI spring symposium: Emotion, personality, and social behavior*. Stanford University, CA; Menlo Park, CA: AAAI Press. Technical Report SS-08-04, pp. 52–59.
Hudlicka, E. (2008c). Modeling the Mechanisms of Emotion Effects on Cognition. In *AAAI fall symposium: Biologically inspired cognitive architectures*. Arlington, VA: Menlo Park, CA: AAAI Press. TR FS-08-04, pp. 82–86.
Hudlicka, E. (2009). Affective game engines: Motivation and requirements. In *4th international conference on foundations of digital games*. Orlando, FL: ACM Press.
Hudlicka, E. (2011a). Affective gaming in education, training and therapy: Motivation, requirements, techniques. In P. Felicia (Ed.), *Handbook of research on improving learning and motivation through educational games: Multidisciplinary approaches*. Hershey, PA: IGI Global.
Hudlicka, E. (2011b). Guidelines for developing computational models of emotions. *International Journal of Synthetic Emotions*, *2*(1), 26–79.
Hudlicka, E. (2013). Virtual training and coaching of health behavior: Example from mindfulness meditation training. *Patient Education and Counseling*, *92*(2), 160–166.
Hudlicka, E. (2014a). From habits to standards: Towards systematic design of emotion models and affective architectures. In Tibor Bosse, Joost Broekens, Joao Dias, & Janneke van der Zwaan (Eds.), *Towards pragmatic computational models of affective processes*. Berlin: Springer LNAI 8750, pp. 1–21.
Hudlicka, E. (2014b). Affective BICA: Challenges and open questions. *Biologically Inspired Agent Architectures*, *14*, 98–125.

Hudlicka, E., Lisetti, C., Hodge, D., Paiva, A., Rizzo, A., & Wagner, E. (2008). Artificial agents for psychotherapy. In *Proceedings of the AAAI spring symposium on emotion, personality and social behavior.* Menlo Park, CA: AAAI. TR SS-08-04, pp. 60–64.
Hudlicka, E., & McNeese, M. (2002). User's affective & belief state: Assessment and GUI adaptation. *International Journal of User Modeling and User Adapted Interaction*, *12*, 1–47.
Hudlicka, E., Payr, S., Ventura, R., Becker-Asano, C., Fischer, K., Leite, I., et al. (2009). Social interaction with robots and agents: Where do we stand, where do we go? In *Proceedings of the 3rd international conference on affective computing and intelligent interaction*, Amsterdam, Holland.
Kang, S., Gratch, J., Sidner, C., Artstein, R., Huang, L., & Morency, L.-P. (2012, June 4–8). Towards building a virtual counselor: Modeling nonverbal behavior during intimate self-disclosure. In Conitzer, Winikoff, Padgham, & van der Hoek (Eds.), *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Valencia, Spain.
Kapoor, A., Burleson, W., & Picard, R. W. (2007). Automatic prediction of frustration. *International Journal of Human-Computer Studies*, *65*(8), 724–736.
Lisetti, C., Amini, R., Yasavur, U., & Rishe, N. (2013). I Can Help You Change! An empathic virtual agent delivers behavior change health interventions. *ACM Transactions on Management Information Systems*, *4*(4) Article 19.
Loyall, A. B., & Bates, J. (1993). Real-time control of animated broad agents. In *Proceedings of the fifteenth annual conference of the cognitive science society* (pp. 664–669).
MacDorman, K. F. (2005). Androids as an experimental apparatus: Why is there an uncanny valley and can we exploit it. In *CogSci-2005 workshop: Toward social mechanisms of android science* (pp. 106–118).
Malbos, E., Boyer, L., & Lançon, C. (2013). Virtual reality in the treatment of mental disorders. *Presse Medicale (Paris, France: 1983)*, *42*(11), 1442–1452.
Mayer, J. D., & Salovey, P. (1995). Emotional intelligence and the construction and regulation of feelings. *Applied and Preventive Psychology*, *4*(3), 197–208.
Mcquiggan, S. W., Mott, B. W., & Lester, J. C. (2008). Modeling self-efficacy in intelligent tutoring systems: An inductive approach. *User Modeling and User-Adapted Interaction*, *18*(1-2), 81–123.
Mori, M. (1970). Bukimi no tani. *Energy*, *7*, 33–35.
Motraghi, T. E., Seim, R. W., Meyer, E. C., & Morissette, S. B. (2014). Virtual reality exposure therapy for the treatment of posttraumatic stress disorder: A methodological review using CONSORT guidelines. *Journal of Clinical Psychology*, *70*(3), 197–208.
Norman, T., & Badler, N. (1997). A virtual human presenter. In *IJCAI'97 Workshop on Animated Interface Agents*.
Ortony, A., Clore, G., & Collins, A. (1988). *The cognitive structure of emotions* New York: Cambridge University Press.
Ortony, A. (2001). On making believable emotional agents believable. In R. Trappl, & P. Petta (Eds.), *Emotions in humans and artifacts*. Cambridge, MA: MIT Press.
Ortony, A. (2002). On making believable emotional agents believable. In R. Trappl, P. Petta, & S. Payr (Eds.), *Emotions in humans and artifacts* (pp. 189–211).
Paiva, A., Dias, J., Sobral, D., Aylett, R., Sobreperez, P., Woods, S., et al. (2004). Caring for agents and agents that care: Building empathic relations with synthetic agents. In *Proceedings of the third international joint conference on autonomous agents and multiagent systems* (pp. 194–201). New York, NY.
Paiva, A., Dias, J., Sobral, D., Aylett, R., Woods, S., Hall, L., et al. (2005). Learning by feeling: Evoking empathy with synthetic characters. *Applied Artificial Intelligence*, *19*(3–4), 235–266.
Picard, R. (1997). *Affective computing* Cambridge, MA: MIT Press.
Popescu, A., Broekens, J., & van Someren, M. (2014). GAMYGDALA: An emotion engine for games. *IEEE Transactions on Affective Computing*, *5*(1), 32–44.

Prendinger, H., & Ishizuka, M. (2004). *Life-like characters* (Heidelberg: *Springer* pp. 1−4).

Reeves, B., & Nass, C. (1996). The media equation: How people treat computers, television, and new media like real people and places. *Computers & Mathematics with Applications*, *33*(5) 128−128.

Reilly, W. S. & Bates, J. (1992). Building emotional agents. Pittsburgh, PA: School of Computer Science, Carnegie Mellon University.

Reilly, W. S. N. (2006). Modeling what happens between emotional antecedents and emotional consequents. In *Paper presented at ACE 2006*, Vienna, Austria.

Ricciardi, F., & De Paolis, L. T. D. (2014). A comprehensive review of serious games in health professions. *International Journal of Computer Games Technology, 2014*, 9.

Rickel, J., & Johnson, W. L. (1999). Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence, 13*, 343−382.

Rizzo, A., Difede, J., Rothbaum, B. O., Reger, G., Spitalnick, J., Cukor, J., et al. (2010). Development and early evaluation of the Virtual Iraq/Afghanistan exposure therapy system for combat-related PTSD. *Annals of the New York Academy of Sciences*, *1208*(1), 114−125.

Rizzo, A., Forbell, E., Lange, B., Buckwalter, J. G., Williams, J., Sagae, K., et al. (2012). *SimCoach: An online intelligent virtual agent system for breaking down barriers to care for service members and veterans. Healing War Trauma: A Handbook of Creative Approaches* New York: Routledge.

De Rosis, F., Pelachaud, C., Poggi, I., Carofiglio, V., & De Carolis, B. (2003). From Greta's mind to her face: Modelling the dynamics of affective states in a conversational embodied agent. *International Journal of Human-Computer Studies*, *59*(1), 81−118.

Sajjad, S., Hanan Abdullah, A., Sharif, M., & Mohsin, S. (2014). Psychotherapy through video game to target illness related problematic behaviors of children with brain tumor. *Current Medical Imaging Reviews*, *10*(1), 62−72.

Scherer, K. R. (1992). Emotions are biologically and socially constituted: A response to Greenwood. *New Ideas in Psychology*, *10*(1), 19−22.

Scherer, K., Schorr, A., & Johnstone, T. (2001). *Appraisal processes in emotion: Theory, methods, research* NY: Oxford University Press.

Schröder, M., Baggia, P., Burkhardt, F., Pelachaud, C., Peter, C., & Zovato, E. (2011). EmotionML − An upcoming standard for representing emotions and related states. In S. D'Mello, A. Graesser, B. Schuller, & J.-C. Martin (Eds.), *Affective computing and intelligent interaction* (pp. 316−325). Berlin: Springer.

Sourina, O., Wang, Q., & Nguyen, M. K. (2011). EEG-based "serious" games and monitoring tools for pain management. *Studies in Health Technology and Informatics*, *163*, 606−610.

Smith, C. A., & Kirby, L. (2000). Consequences require antecedents: Toward a process model of emotion elicitation. In J. P. Forgas (Ed.), *Feeling and thinking: The role of affect in social cognition*. NY: Cambridge University Press.

Sykes, J. (2004). Affective gaming. Retrieved May 2008, from: <http://www.jonsykes.com/Ivory.htm>.

Yannakakis, G. N., & Paiva, A. (2014). Emotion in games. In R. A. Calvo, S. D'Mello, J. Gratch, & A. Kappas (Eds.), *Handbook of affective computing* (pp. 459−471). NY: Oxford University Press.

Yannakakis, G. N., Hallam, J., & Lund, H. H. (2008). Entertainment capture through heart rate activity in physical interactive playgrounds. *User Modeling and User-Adapted Interaction*, *18*(1−2), 207−243.

Zeng, Z., Pantic, M., Roisman, G. I., & Huang, T. S. (2009). A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *Pattern Analysis and Machine Intelligence, IEEE Transactions On*, *31*(1), 39−58.