Published in final edited form as:

Patient Educ Couns. 2013 August ; 92(2): 160–166. doi:10.1016/j.pec.2013.05.007.

Virtual Training and Coaching of Health Behavior: Example from Mindfulness Meditation Training

Eva Hudlicka

Psychometrix Associates, Inc., Blacksburg, VA, US

Abstract

Objective—Computer-based virtual coaches are increasingly being explored for patient education, counseling, and health behavior training and coaching. The objective of this research was to develop and evaluate a Virtual Mindfulness Coach for training and coaching in mindfulness meditation.

Method—The coach was implemented as an embodied conversational character, providing mindfulness training and coaching via mixed initiative, text-based, natural language dialogue with the student, and emphasizing affect-adaptive interaction. (The term ‘mixed initiative dialog’ refers to a human-machine dialogue where either can initiate a conversation or a change in the conversation topic.)

Results—Findings from a pilot evaluation study indicate that the coach-based training is more effective in helping students establish a regular practice than self-administered training using written and audio materials. The coached group also appeared to be in more advanced stages of change in terms of the transtheoretical model, and have a higher sense of self-efficacy regarding establishment of a regular mindfulness practice.

Conclusion—These results suggest that virtual coach-based training of mindfulness is both feasible, and potentially more effective, than a self-administered program. Of particular interest is the identification of the specific coach features that contribute to its effectiveness.

Practice Implications—Virtual coaches could provide easily-accessible and cost-effective customized training for a range of health behaviors. The affect-adaptive aspect of these coaches is particularly relevant for helping patients establish long-term behavior changes.

1. Introduction

The past decade has witnessed an increased interest in the use of embodied conversational characters (ECAs) [1] as virtual coaches. These coaches provide assistance, training or coaching support across a variety of contexts; e.g., virtual coach Laura’s exercise coach that helps users establish and maintain an exercise program [2], advisors for patients with specific medical conditions (e.g., agent Greta who provides advice about eating disorders

© 2013 Elsevier Ireland Ltd. All rights reserved.

Corresponding author: Psychometrix Associates, Inc., 1805 Azalea Dr., Blacksburg, VA 24060 USA, Tel. 413 207 7066, FAX 413 549 0357, hudlicka@ieee.org.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

"I confirm all patient/personal identifiers have been removed or disguised so the patient/person(s) described are not identifiable and cannot be identified through the details of the story."

[3], and virtual counselors to provide problem-solving skills training and emotional support for caregivers [4].

In this paper we describe an ECA virtual coach that provides mindfulness meditation training, and the coaching support necessary to begin a regular practice, via a web-based application. The coach is represented as a female character, capable of showing a variety of affective and conversational expressions, and interacts with the user via a mixed-initiative, text-based, natural language dialogue. The coach adapts both its facial expressions and the dialogue content to the user's learning needs and motivational state. Findings from a pilot evaluation study indicate that the coach-based training is more effective in helping users establish a regular practice than self-administered training via written and audio materials.

Below we provide background information about mindfulness and embodied conversational characters. The remainder of the paper then describes the virtual mindfulness coach prototype, the empirical evaluation study, and the findings regarding the coach's effectiveness. The paper concludes with a discussion of technical challenges in creating effective virtual coaches, focusing on their multi-modal interaction capabilities and social and affective realism.

1.1 Mindfulness meditation

Mindfulness is best described as a state of awareness characterized by "full attention to, and awareness of, the internal and external experience of the present moment" [5], and a non-judgmental attitude towards whatever thoughts, images, feelings or sensations enter awareness. A typical mindfulness practice session lasts about 20 minutes. The meditator sits in a comfortable position, and begins the session with a concentration practice, by focusing his/her attention on some fixed, concrete entity, usually the breath, followed by mindfulness practice proper, where the meditator observes whatever emerges into consciousness, with full, non-judgmental awareness. Regular practice usually involves 20–30 minutes, done once or twice daily.

Regular practice of mindfulness meditation is associated with number of health benefits. Evidence from existing studies indicates that mindfulness practice enhances health-related quality of life, affecting both physical and psychological symptoms [6,7]. Benefits of mindfulness meditation have been demonstrated in stress reduction [8], pain reduction [9], enhanced immune responses [10], reduction of symptoms in anxiety disorders [11], prevention of relapse in major depression [12], improvement in a subjective sense of well-being [13], and improvements in cognitive functions [5].

In spite of these benefits, the relative ease of learning mindfulness techniques, and no costs associated with their practice, use of mindfulness meditation remains limited. This is due to several factors, including lack of adequately trained healthcare providers, lack of access to teachers and training programs, cost of training programs, misconceptions about the methods involved, and difficulties with establishing a regular practice.

The virtual mindfulness coach prototype described here was developed to address this problem, and to provide increased access to mindfulness training and coaching. The coach mindfulness training content is based on a widely-used mindfulness training program, the Mindfulness Based Stress Reduction (MBSR) [14]. The mindfulness meditation training consists of 4 lessons, with one lesson administered each week during the 4-week training period. Each lesson focuses on a specific topic, and the 4 lessons together are designed to provide a solid introduction to mindfulness techniques, address common problems beginners encounter (e.g., restlessness, boredom), and suggest how to integrate mindfulness into daily life.

1.2 Embodied conversational characters and virtual coaches

Virtual coaches use the emerging technologies of *embodied conversational agents* (ECA's) [1], and *relational pedagogical agents* [15]. ECA's are animated virtual characters, displayed on a computer or a mobile device screen. ECA's play the roles of teachers, mentors, advisors, social companions, and, increasingly, of virtual coaches [16–20]. The use of ECA's requires minimal or no training, due to their capability to engage in natural interaction with humans through dialogue and non-verbal expression. The latter includes facial expressions, gaze, and gestures, which together help control conversation flow and augment the ECA's visual and behavioral realism. *Relational pedagogical agents* represent a subset of ECA's, designed both to train a particular subject or skill, and to develop the type of a longer-term relationship with the user that helps facilitate coaching [21]. The use of ECA's and relational pedagogical agents has recently begun to be explored in healthcare settings, where these virtual entities act as coaches, trainers, and counselors.

Both the visual appearance and the multi-modal interaction capabilities of these characters vary greatly. The embodiments range from cartoonish, animated characters or faces, to fully-articulated full bodies. The interaction capabilities range from text-based, multiple-choice user input and text-based character output, to speech recognition and synthetic speech. The virtual characters also vary in the degree to which they can adapt to the user's knowledge, affective and motivational state, and display affective and social realism.

2. Methods

The Virtual Mindfulness Coach is implemented as an *embodied conversational agent*, “Chris”, who guides the student through the training material, and provides supportive coaching to help students establish a regular mindfulness practice (see figure 1).

Chris engages multi-modal verbal and non-verbal interaction with the user. Chris is represented by a drawn image of a face, which can show a variety of affective and conversational expressions (e.g., happy, sad, concerned, confused). The verbal interaction is conducted via a mixed-initiative, natural-language, text-based dialogue (see figure 2 (right)). The coach can also support limited non-verbal interaction, by displaying affective and conversational expressions that match the content of the dialogue, provide appropriate affective reaction to the user's utterances, and display conversational expressions that aim to communicate the coach's mental state (e.g., understanding, confusion) (see figure 2 (left)).

These capabilities enable the coach to provide customized coaching and affect-adaptive feedback to the user; e.g., smile when the user reports success, express concern and verbal encouragement when the user reports a problem.

The coach is able to support several types of pedagogical strategies necessary for effective training of mindfulness meditation, and the coaching required to initiate and maintain a regular practice: *didactic* (to convey the conceptual basis of mindfulness, information about the techniques, and their health benefits); *experiential learning* of mindfulness techniques (via the guided meditation recordings, narrated by an experienced mindfulness teacher); and *supportive and coaching interaction*, designed to maintain student motivation via empathic dialogue and customized, concrete advice.

Coach Chris acts as a guide through the mindfulness lesson materials, which are displayed primarily via segments of text, and augmented by Chris' introductions and elaborations, that emphasize material Chris thinks is of particular interest to the user, based on previous dialogue. Chris delivers the mindfulness training in 4 lessons, which follow the same structure; introduction and overview, guided meditation, main lesson, Q&A, adaptive

segment, homework, good-bye. Chris also offers coaching, during which she prompts student to discuss specific problems, and offers suggestions and advice, including pointers to existing didactic material, which she displays for the student.

The coach thus guides the student through the didactic lesson material, including recorded guided meditations, and provides customized advice about meditation practice, based on the student's expressed concerns. For example, if the student repeatedly asks about meditation and its effects on stress, the coach follows-up with questions about the student's stress levels, and provides customized suggestions for practice.

2.1 Multiple types of interactions

Chris interacts with the user via text-based natural language (see figure 2), displayed in a dedicated pane on the coach user interface (refer to figure 1). Chris can engage in several types of dialogue with the user: *didactic*, *relational* and *adaptive/motivational*. Chris can also answer the user's free-form questions.

During the training segment, Chris engages primarily in *didactic interaction*, as she guides the user through the training materials, by introducing the different lesson segments, and displaying the didactic material in a dedicated pane of the coach user interface (see "Main Text Window" in figure 1). Throughout the lesson Chris checks-in with the user, to ensure understanding and motivation. For example, after displaying the material for the weekly lesson, Chris asks the user whether s/he understood everything, and offers to provide additional information for any items that require elaboration. Periodically, Chris engages in *relational interaction* to promote relationship building with the user, and user engagement. For example, at the beginning of each lesson, Chris asks the user how their meditation went the previous week, and expresses joy or concern, depending on the user's answer.

To facilitate the *adaptive/motivational interaction*, Chris collects information about the user throughout the interaction, and periodically asks specific questions (e.g., "Did your meditation go ok last week?" "Were you able to meditate 5 times?"). The information provided by the student is used to gradually construct a user model, which is then used to support *adaptive/motivational dialogue*, by customizing responses and questions to the student's knowledge and motivational state. (The term 'user model' refers to the information collected about a user regarding important user attributes and characteristics relevant for the particular interaction; e.g., level of expertise, known gaps in knowledge, preferences for interaction style.) [22] For example, Chris tracks how many questions the student asks about a particular topic, how well the practice is going, what specific problems the student is encountering, etc.

Based on this information, Chris then asks a series of 'adaptive questions' at the end of each lesson, to provide customized interaction. For example, if the student asked a number of conceptual questions about mindfulness (e.g., "What's the theory of mindfulness?"; "How does mindfulness work?"), Chris offers to provide additional information. If the user asks repeated questions about stress (e.g., "Is mindfulness good for stress?" "Will meditating help me feel less stressed?"), Chris first attempts to determine whether the student is feeling stressed, and then offers additional material on mindfulness and stress. Chris frames the questions in a tentative manner, to allow for the possibility that her understanding of the student's motivational state is incorrect; e.g., "You asked several questions about mindfulness and stress. Is this something of particular concern to you?". If the student answers 'Yes', Chris responds: "Many people practice mindfulness for stress reduction. Would you like to hear additional information about mindfulness and stress reduction?"

2.2 Natural language understanding and dialogue management

The user has two options for natural language dialogue with coach Chris: *free-form natural language input*, and *multiple-choice based input*. In the free form input mode, the user can interrupt the coach at any point and make a comment, or ask a question related to mindfulness practice. The natural language understanding component is currently relatively simple, and is implemented via keyword and template matching. User input is scanned and matched against a list of stored keyword templates, until a match is found. Chris then responds with the phrase associated with that template. If the user input is not understood, Chris expresses disappointment (“I am sorry, but I didn’t understand what you meant” and displays an expression of confusion or concern. Chris also asks the user if s/he would like the question to be added to the database, to improve future natural language understanding in future interactions.

This approach has several limitations, including lack of robustness (any new output must be explicitly included in the template database), the need for extensive template ‘tuning’, limited domain of discourse, and a relatively high rate of user utterances that are not understood or misunderstood. Nevertheless, by appropriately structuring of the coach dialogue, and providing information to the users about the coach’s limited NL understanding capabilities, effective dialogue between the coach and the student can take place. Figure 3 illustrates a segment of the coach-student dialogue, highlighting the mixed-initiative capabilities.

The second mode of user input, via multiple-choice based options, was implemented to address the relatively high frequency of user utterances that were not understood by the coach. This enhancement was based on the success of multiple-choice based input in health coach Laura [15]. This user input alternative is emphasized during the coaching sequence, where the coach asks the student to select a particular problem encountered, or a goal on which to focus, from a set of multiple choice options (refer to “Multiple Choice user input in figure 1). For example, the following are some of the options provided when the coach asks the user whether s/he has any concerns about the practice: Can’t find the time; Meditation is too boring; Keep falling asleep; Feel too restless; Pain makes it difficult to sit still; Not sure meditation is the right thing for me; etc. Based on the user’s response, further, context-sensitive options are then displayed. The user’s responses are echoed in the dialogue pane, so that a continuous dialogue is displayed for the user (refer to figure 1). With careful and thorough analysis of the domain and the student’s likely needs, adequate sets of multiple-choice options can be developed to minimize the likelihood that a particular option will not be available.

This approach to augmenting the free-form text input proved to be effective, and reduced the users’ frustration with Chris’ inability to understand all of the student utterances. In fact, after this enhancement was implemented, the users reported that Chris’ ability to understand free-form natural language input improved, even though this component was not changed.

2.3 Affective modeling and adaptation

An important component of training and coaching interaction is the coach’s ability to adapt to the student’s knowledge and motivational state. An effective coach should be able to customize its feedback to the student’s specific knowledge needs (provide additional information or correct misconceptions), affective state (express empathy and understanding when the student is sad), and motivational state (provide support and encouragement when the student is lacking in self-confidence, and challenges when the student is becoming bored). A core component of this capability is the coach’s ability to recognize and adapt to the student’s emotional state.

To implement this type of adaptation, coach Chris collects information about the student from questions asked, comments, and answers to Chris' questions. Data collected include frequency of questions on a particular topic (e.g., stress, health benefits, restlessness), frequency of positive vs. negative comments, student's assessment of the overall experience, number of questions asked, etc.

Together, this information comprises the user model and the affective user model. Chris uses this information to deepen the relationship with the user (via customized relational dialogue), to provide adaptive training by providing specific information to enhance the user's knowledge of the training content, and to enhance engagement and motivation by providing appropriate emotional support and encouragement.

2.4 Implementation

The first version of the coach was implemented using proprietary software developed by LifeLike Solutions, Inc. A version 2 coach was later developed in-house, to provide us with the capability to modify and augment the software to introduce new functionalities (e.g., the multiple-choice based user input). The coach is currently deployed via the web. The user interface is defined in HTML, with embedded Javascript supporting the parsing of user input necessary to support NL understanding and student-coach interaction, and coach affective expressions. External data files define the keyword template-phrase pairs, the facial expressions, and the didactic material. The affective user model is currently stored in local variables via cookies.

3. Results

The coach's effectiveness was evaluated in a 7-week study, to assess the coach's ability to provide mindfulness training (weeks 1–4), and to support students in establishing a regular meditation practice (weeks 5–7). An experimental group used the coach, and a control group used written and audio materials. The content was identical. Participants (N=32; 75% female; mean age=38) were asked to meditate 5x/week, for 20 minutes each time. Participants were pseudo-randomly assigned to the two conditions, to provide gender- and age-balanced samples. During the 4 week training period, one lesson was administered per week. During the 3 week coaching period, the participants were free to interact with the coach (for the experimental group), or read the manual (for the control group), as needed.

Data were collected via web-administered surveys, using SurveyMonkey.com. The format of most questions was a 5-point Likert scale, with opportunities for free-form feedback. Weekly surveys were administered to track the meditation frequency, and the participants' overall experience, including any benefits and difficulties. At 5 weeks, two tests were administered to assess knowledge of mindfulness concepts and techniques. In week 7, the participants' sense of self-efficacy regarding their ability to continue mindfulness practice was assessed, via a single (Likert scale) survey item: "I feel confident that I can continue my practice after this study is finished". The participants' stage of change within the transtheoretical model of change [23] was also assessed, by eliciting their assessment of 3 negative ("Meditation takes a lot of time", "Meditation requires a lot of effort", and "It is difficult to find the time") and 2 positive aspects ("There are tangible benefits associated with meditation", "The benefits I have derived from meditation are worth the time and effort") of their overall experience with meditation.

Participants in the experimental group filled out additional surveys assessing specific features and functionalities of the Virtual Coach Chris, and their experience with Chris (weeks 1, 3, 6, 7). The focus was on assessing Chris' affective and social realism, ability to effectively answer questions, and the adaptive and relational capabilities.

The results of the study indicate that the coach provides more effective training and coaching than a self-administered program using written and audio materials. Specifically, the coach group participants practiced more frequently (in days/week) ($M = 4.5$, $SD = .28$) compared to the control group ($M = 3.2$, $SD = .72$), $t(6) = 8.00$, $p = 0$ and for longer time (in minutes/meditation session) ($M = 19$; $SD = .85$) vs. ($M = 16$; $SD = 2.35$), $t(6) = 3.27$, $p = 0.01$ than the controls. This effect was particularly evident during the 3-week follow-up period, where the meditation frequency in days/week was higher for the coach group ($M = 4.3$; $SD = .37$) than for the control group ($M = 2.7$; $SD = .85$), $t(2) = 18.01$, $p = 0.00$ as was the length of meditating ($M = 18.5$; $SD = 1.05$) for the coach group, vs. ($M = 13.6$; $SD = 1.43$) for the control group, $t(2) = 22.41$, $p = 0.00$. The coach group participants found the experience more rewarding, enjoyable, beneficial, and engaging than the controls. The coach group also demonstrated a general increased sense of involvement, as demonstrated by the lower number of dropouts than the control group (3 vs. 6), and higher frequency of answering the surveys.

The coach group participants also appeared to be in more advanced stages of change, in terms of the transtheoretical model of behavior change, and felt more confident in their ability to maintain a regular meditation practice. The conceptual and practical knowledge about mindfulness meditation was equal for both groups.

4. Discussion and Conclusions

4.1 Discussion

These results provide support for the hypothesis that virtual coach-based training and coaching is not only effective for mindfulness meditation training, but is more effective in helping students establish a regular practice, than the use of written and audio materials. The fact that the results were statistically significant in spite of the small number of participants is particularly encouraging. However, a more extensive study is necessary to confirm these findings.

Interestingly, the assessment of several coach features and functionalities was not as positive as anticipated, or as the meditation frequency data would suggest. Participants felt neutral to mildly positive about the coach's ability to provide customized feedback and address questions or difficulties (.3 on a -2 to +2 Likert scale). They were slightly more positive about the coach's ability to provide support and encouragement (.9). The assessments of the coach's ability to adapting non-verbal expressions to the student's needs, and a feeling of 'personal connection' with the coach, were negative (-.6 and -.5, respectively).

The study data suggest that the most successful coach feature was its ability to provide customized feedback and support, in an interactive manner. This was the case in spite of the negative assessments of the coach's ability to conduct free-form, natural language dialogue, and the negative assessment of the coach's affective and social realism.

Several specific hypotheses should be further explored to explain these findings. *First*, it is likely that the negative assessment of the coach's social and affective realism was due to two factors: the frequent misunderstood user input, and the appearance of the coach's face. Additional experiments are necessary to identify the role of these two factors in the overall sense of the coach's affective and social realism. It is possible that the participants experienced a phenomenon referred to as the 'uncanny valley' [24, 25], where the sense of overall believability of synthetic agent or robot negatively correlates with the degree of visual realism of the agent / robot appearance. In other words, more 'cartoonish' agents are perceived as more believable and affectively realistic than more visually-realistic agents. 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 non-verbal interaction, and when these expectations are not met, we feel that the character is not believable. The role of the ‘uncanny valley’ phenomenon in the coach-participant interaction can be further explored, by varying the degree of visual realism of the coach’s face, to determine an optimal level of realism to establish an engaging relationship.

Since preference for faces are highly idiosyncratic, it is also possible that the reason for the negative assessments of the coach’s believability and non-verbal interaction was that the participants simply did not like the face of coach Chris, and would have felt a greater sense of believability, social realism and engagement with a different face. Future versions of the coach will explore this possibility, by providing the users with a set of alternative faces, including user-selected photographs.

Finally, there is also the possibility that the coach face did not play a major role in the effectiveness of the interaction, and that the dialogue itself, the linguistic coaching support, and the ability to rapidly select relevant material (e.g., ability to browse relevant segments of content), were responsible for the coach’s effectiveness. This hypothesis will be evaluated in future studies, by exploring the relative effectiveness of an embodied conversational agent coach (such as the one used for the prototype) vs. a coach capable of dialogue but not represented by a face.

Identification of the specific features of an embodied conversational agent that contribute to its believability, and to the user’s level of engagement, is an active area of research. Future studies will focus on a more extensive evaluation of individual coach features with respect to their effect on the overall coach believability, effectiveness, and a sense of engagement experienced by the user.

4.2 Conclusions

Results of an evaluation study of Chris’ effectiveness with 32 participants indicates that the virtual coach is superior to a self-administered mindfulness course in providing mindfulness meditation training and coaching, and results in more frequent meditation. A more extensive study with a larger sample size is necessary to confirm these preliminary findings.

These results are encouraging, and provide evidence that even a relatively simple coach structure (in terms of NL understanding, dialogue management, affective realism, affective user modeling and affect-adaptation) can have significant impact on the degree of student engagement during training, and on learning outcomes. (The term ‘affective user modeling’ [26] refers to models of users that include information about the user’s emotional state and affective profile; e.g., what types of interactions are likely to trigger which emotions or personality types.)

Future work will focus on enhancing the coach’s NL understanding and generation capabilities and dialogue management, and the affective and social realism of the coach’s embodiment and interaction, via more complex user emotion recognition and affective modeling [27–29]. In addition, several specific hypotheses will be explored in future studies to identify the key factors that contribute to student engagement and learning.

4.3 Practice Implications

An ability to deliver training and coaching via virtual characters has broad implications for healthcare, patient communication and counseling. The work reported here focuses on

mindfulness meditation, but the coach software could easily be adapted to a range of health behavior training and coaching – both the cessation of unhealthy behavior (e.g., smoking, overeating, substance abuse) and the acquisition of healthy behavior (e.g., exercise, disease-specific diet such as diabetes, etc.). Ability to deliver customized training could greatly enhance adherence to training protocol. An associated ability to deliver customized and affect-adaptive coaching could facilitate the actual acquisition of the necessary habits and help patients establish long-lasting regimens and behavioral changes. The ability of the coach to interact with patients via multi-modal interactions, including natural language, makes the coach accessible to a broad range of population, since no special training is required. Future extensions to both text-based and audio version further enhance accessibility. Finally, access via the web minimizes access costs and makes the training / coaching available in remote locations, and to patients who might be unwilling or unable to attend a hospital or clinic-based training program.

Acknowledgments

This work was performed under an NIH SBIR grant 1R43 AT005397-01, awarded to Psychometrix Associates by the National Center for Complementary and Alternative Medicine. The author would like to acknowledge the support of Dr. John Glowa of NIH, and thank the following individuals who contributed to this research: Kevin Griffin, who developed the mindfulness training materials; Jonathan Klein, who designed and developed the coach's face image and expressions, and developed the coach user interface; George Roberts, who designed and implemented the coach JavaScript parsers; and Drs. Noreen Klein and Marion Reynolds of Virginia Tech, who provided statistical consultations.

References

1. Cassell, J.; Sullivan, J.; Prevost, S.; Churchill, E. *Embodied Conversational Agents*. Cambridge, MA: MIT Press; 2000.
2. Bickmore T, Giorgino T. Health Dialog Systems for Patients and Consumers. *J Biomed Inform.* 2006; 39:556–571. [PubMed: 16464643]
3. de Rosi F, Pelachaud C, Poggi I, Carofiglio V, De Carolis B. From Greta's mind to her face: Modelling the dynamics of affective states in a conversational embodied agent. *Int J Hum Comput Stud.* 2003; 59:81–118.
4. Johnson, WL.; LaBore, C.; Chiu, Y-C. *Proc. of AAAI Fall Symposium: Health Dialog Systems*. Menlo Park, CA: AAAI Press; 2004. A pedagogical agent for psychosocial intervention on a handheld computer; p. 64-70.
5. Chambers R, Lo BCY, Allen NB. The Impact of Intensive Mindfulness Training on Attentional Control, Cognitive Style, and Affect. *Cognit Ther Res.* 2008; 32:303–322. (2008).
6. Lazar, S. Mindfulness Research. In: Germer, CK.; Siegel, RD.; Fulton, PR., editors. *Mindfulness and Psychotherapy*. NY: The Guildford Press; 2005.
7. Reibel DK, Greeson JM, Brainard GC, Rosenzweig S. Mindfulness based stress reduction and health-related quality of life in a heterogeneous patient population. *Gen Hosp Psychiatry.* 2001; 23:183–192. [PubMed: 11543844]
8. Shapiro SL, Schwartz GE, Bonner G. Effects of mindfulness-based stress reduction on medical and premedical students. *J Behav Med.* 1998; 21:581–599. [PubMed: 9891256]
9. Kabat-Zinn, J. *Full Catastrophe Living: Using the Wisdom of Your Body and Mind to Face Stress, Pain, and Illness*. NY: Delta Publishing; 1990.
10. Davidson RJ, Kabat-Zinn J, Schumacher J, Rosenkranz M, Muller D, Santorelli SF, Urbanowski F, Harrington A, Bonus K, Sheridan JF. Alterations in brain and immune function produced by mindfulness meditation. *Psychosom Med.* 2003; 65:564–570. [PubMed: 12883106]
11. Evans S, Ferrando S, Findler M, Stowell C, Smart C, Haglin D. Mindfulness-based cognitive therapy for generalized anxiety disorder. *J Anxiety Disord.* 2008; 22:716–721. [PubMed: 17765453]

12. Teasdale JD, Segal ZV, Williams JMG, Ridgeway VA, Soulsby JM, Lau MA. Prevention of relapse/recurrence in Major Depression by Mindfulness-Based Cognitive Therapy. *J Consult Clin Psychol.* 2000; 68:615–622. [PubMed: 10965637]
13. Brown KW, Ryan RM. The benefits of being present: Mindfulness and its role in psychological well-being. *J Pers Soc Psychol.* 2003; 84:822–848. (2003). [PubMed: 12703651]
14. Kabat-Zinn J. <http://www.mindfulnesscds.com/series1.html>.
15. Bickmore, T. PhD Thesis, Media Arts & Sciences. Cambridge, MA: MIT; 2003. Relational Agents: Effecting Change through Human-Computer Relationships.
16. Prendinger, H.; Ishizuka, M. Life-Like Characters. New-York: Springer; 2004.
17. Hayes-Roth B, Amano K, Saker R, Sephton T. Training Brief Intervention with a Virtual Coach and Virtual Patients. *Ann Rev Cyber Therapy Telemedicine.* 2004; 2:85–96.
18. Rickel J, Marsella S, Gratch J, Hill R, Traum D, Swartout W. Toward a new generation of virtual humans for interactive experiences. *IEEE Intell Syst.* 2002:32–38.
19. Clarebout G, Elen J, Johnson WL, Shaw E. Animated pedagogical agents: An opportunity to be grasped? *J Educ Multimedia Hypermedia.* 2002; 11:267–286.
20. Johnson WL, Rickel JW, Lester JC. Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *Int J Artificial Intell Education.* 2000; 11:47–78.
21. Bickmore T, Gruber A, Picard R. Establishing Alliance in Automated Health Behavior Change Interventions. *Patient Education and Counseling.* 2005; 59:21–30. [PubMed: 16198215]
22. Kobsa A. Generic User Modeling Systems. *User Model User-Adapt Interact.* 2001; 11:49–63.
23. Prochaska, JO.; DiClemente, CC. The transtheoretical approach. In: Norcross, JC.; Goldfried, MR., editors. *Handbook of psychotherapy integration.* NY: Oxford University Press; 2005.
24. MacDorman, KF. Androids as an experimental apparatus: Why is there an uncanny valley and can we exploit it?; *Proceedings of Cognitive Science Society Annual Meeting*; 2005. (Toward Social Mechanisms of Android Science Workshop).
25. Mori M. Bukimi no tani. *Energy.* 1970; 7:33–35.
26. Elliot, C.; Rickel, J.; Lester, J. Lifelike Pedagogical Agents and Affective Computing: An Exploratory Synthesis. In: Wooldridge, M.; Veloso, M., editors. *Artificial Intelligence Today, Lectures Notes in Artificial Intelligence.* Vol. 1600. Berlin: Springer-Verlag; 1999. p. 195-212.
27. Hudlicka, E. *Proc. of International Conference on Human-Computer Interaction.* Munich, Germany: 1999. Methodology for Developing Adaptive Affective Interfaces.
28. Hudlicka, E. Increasing Socially-Intelligent Architecture Realism by Modeling and Adapting to Affect and Personality. In: Dautenhahn, K.; Bond, AH.; Canamero, L.; Edmonds, B., editors. *Multiagent Systems, Artificial Societies, and Simulated Organizations.* Dordrecht, The Netherlands: Kluwer Academic Publishers; 2002.
29. Hudlicka E, McNeese M. User's Affective & Belief State: Assessment and GUI Adaptation. *User Model User-Adapt Interact.* 2002; 12:1–47.

Figure 1.
The Virtual Mindfulness Coach User Interface

Chris: Did the guided meditation go ok for you?

eva: Yes

Chris: That's great. Most people really enjoy the guided meditations.

eva: Next

Chris: OK. So let's get started with the main topics for today.

eva: Next

Chris: The left panel shows you a detailed outline for each lesson. [...]

eva: Next

Chris: Let's begin by defining mindfulness. Mindfulness is a non-judgmental awareness of what is happening to you, at the present moment.

eva: Is it hard to learn mindfulness?

Chris: It's hard to answer this question. Some people immediately enjoy the experience, and begin to see benefits quickly. [...]

eva: I have a hard time sitting still.

Chris: For some people it can be a challenge, in the beginning. But it gets easier with practice. Be sure to avoid caffeine [...]

eva: How long is this course?

Chris: This course is designed for 4 weeks. You should cover 1 lesson per week [...]

eva: Next

Chris: With mindfulness, we look at the whole range of experience, the pleasant and the unpleasant, with a non-reactive and curious attention.

Figure 2.
Examples of Coach "Chris" Expressions, and Sample Coach-Student Dialogue

Hello. My name is Chris and I will be your mindfulness meditation teacher for the next four weeks. I will be working with an experienced mindfulness teacher. His name is Kevin Griffin.

hi chris

Hello. How are you today?

i feel fine, thanks

I'm glad to hear that!

Today we will begin your meditation training with a bit of background about mindfulness meditation. Then I will tell you how to practice mindfulness.

Have you ever meditated before?

no

OK. So this will all be new. Let's get started then.

Let me tell you a bit about today's main topic. Today we will cover the basics. I will tell you about mindfulness, its benefits, and how to practice mindfulness.

is it like prayer?

Mindfulness is not like prayer, although there can be some shared qualities. Would you like to know more about the differences?

not right now, thanks

OK. Maybe we'll talk about this later.

is it ok to fall asleep?

Feeling sleepy is normal during meditation. We will talk about that later in the course.

Do you have any questions about today's topics?

where can you meditate?

You can meditate almost anywhere, but it should be a quiet place where you feel comfortable sitting with your eyes closed.

do you have to be buddhist to meditate?

Are you interested in Buddhism?

> yes

There is lot of information about Buddhism on the Web. Try Googling it sometime.

Coach follows lesson plan

User-initiated Interruption & interaction

Coach returns to lesson plan

user-initiated interruption & interaction

Coach attempts to re-direct the discussion by not elaborating the topic

Figure 3.

Sample of a Free-Form Text Dialogue Between Coach and Student

User-input is highlighted in yellow. Coach follows a lesson plan (red text), but allows for interruptions from the student to follow specific topics of interest.