

Breathe Deep: A Breath-Sensitive Interactive Meditation Coach

Ameneh Shamekhi
Northeastern University
Boston, Massachusetts
ameneh@ccs.neu.edu

Timothy Bickmore
Northeastern University
Boston, Massachusetts
bickmore@ccs.neu.edu

ABSTRACT

Mindfulness meditation has been demonstrated to be an effective approach for alleviating symptoms related to a variety of chronic health conditions, including pain, anxiety, and depression. Meditation takes practice and requires training, especially for novices, to learn mindfulness and emotion regulation. However, while face-to-face instructions can provide the best long-term results, many people cannot afford or schedule attendance at meditation classes. We present an automated conversational agent that acts as a virtual meditation coach that is interactive and adaptive to a user's breathing behavior, based on inputs from a respiration sensor in a meditation session. We designed and validated three interaction techniques based on the user's breathing. Results from two experimental studies demonstrate that users are highly receptive of the virtual coach technology, and appreciated the interactivity afforded by the respiration sensor. Participants also felt more relaxed when the meditation coach adapted the instructions to their breathing.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**;

KEYWORDS

Conversational Agents, Sensors, Meditation, Respiration

ACM Reference Format:

Ameneh Shamekhi and Timothy Bickmore. 2018. Breathe Deep: A Breath-Sensitive Interactive Meditation Coach. In *PervasiveHealth '18: 12th EAI International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth '18)*, May 21–24, 2018, New York, NY, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3240925.3240940>

1 INTRODUCTION

Mindfulness meditation is a form of attention training in which practitioners are taught to focus their mind on the present moment, to the exclusion of distractions and thoughts [23]. Recent fMRI studies have demonstrated a link between mindfulness meditation and mood [13]. Focusing less on persistent negative thoughts is key to alleviating many depressive symptoms, and may account for the effectiveness of mindfulness meditation in treating mood disorders such as anxiety and depression [8, 17]. A significant body

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PervasiveHealth '18, May 21–24, 2018, New York, NY, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6450-8/18/05...\$15.00

<https://doi.org/10.1145/3240925.3240940>



Figure 1: The virtual Meditation Coach

of literature has also demonstrated the beneficial effects of mindfulness meditation on a wide variety of physical [1] and mental disorders [17], such as chronic pain, post-traumatic stress disorder, and insomnia, in addition to treating symptoms of other chronic conditions, including cancer, hypertension, diabetes, HIV/AIDS [5].

Mindfulness meditation requires training and practice to be effective. Novices in particular, may need significant coaching and practice to become proficient and be able to regulate their thoughts and emotions. Unfortunately, cost and logistical barriers prevent many people from attending mindfulness meditation training courses. While many self-help books, videos, and applications [6, 20, 24] have been developed for meditation training, these methods lack the individualized instruction of a human trainer, and suffer the same high attrition rates of any longitudinal behavior change program, for those wanting to make regular meditation a habit.

In recent years, many HCI researchers have explored how technology could create and support a mindfulness experience [25, 33, 40]. Several studies have demonstrated promising results regarding stress reduction, and improving mindfulness for participants who have undergone meditation training using mobile applications [34], Virtual Reality [25], and sensory-based interfaces [29]. Zhu et al. reviewed 102 applications for digital mindfulness and proposed a framework for classifying this large number of commercial mindfulness applications into four categories including: digitalized, personalized, and quantified mindfulness applications [46]. The increasing number of these applications shows that both consumers and developers are very interested in leveraging technology to improve their mindfulness experiences. However, existing applications still largely lack individualized tailoring and instruction. Although most of the applications provide different types of mindfulness instructions and guidance through text, audio and video, only a limited number of them offer real-time sensing and adaptive feedback based on the user's performance [46].

To address these shortcomings, we have developed a virtual meditation coach (Figure 1) that guides users through meditation sessions. Virtual coaches that simulate counseling sessions with an expert human have now been successfully developed to address a range of health behaviors [3]. They have also been demonstrated to be effective in promoting patient adherence in longitudinal treatments [4]. In addition to simulating the verbal and nonverbal behavior of a human counselor, they can provide privacy and accessibility, decreasing the expense and eliminating distance barriers associated with attending human-led classes. Thus, these agents can be a useful medium particularly for mindfulness applications, because of their conversational interface, which reduces the cognitive burden of understanding instructions for practitioners.

The virtual meditation coach is an embodied conversational agent that appears on a touch screen desktop computer. Given the importance of breathing behavior in meditation, and to make the coach interactive and adaptive to users needs during meditation, we integrated a respiration sensor into the coaching system and incorporated user respiration into the system as an input modality and data to provide real-time, tailored feedback.

In the rest of this paper, we first review related work, then describe the design of the virtual meditation coach system, including novel respiration-based interactions and their validation. Then we describe two studies to evaluate the coach system, first comparing the respiration-based interaction to an equivalent interaction without the respiration sensor input, and second comparing the interactive coach to a non-interactive self-help video of a human meditation trainer. Finally we discuss the findings of both studies and future research.

The contributions of this work include: development and evaluation of the acceptance and feasibility of 1) an *embodied conversational agent* in guiding meditation instruction; 2) a chest-expansion-based *respiration as an interaction modality*; and 3) new interaction techniques using respiration (chest expansion), beyond simple biofeedback. This work takes a step towards the development of interactive and customizable calming applications [42], which can help individuals to improve their overall wellbeing and mindfulness, in addition to addressing a range of other health problems.

2 RELATED WORK

To better situate our work, we briefly review the existing research on interactive systems that have been developed to support meditation training.

2.1 Automated Meditation Training

In addition to a plethora of self-help books, videos, and audio-based media to help novices learn and practice meditation, there have also been some commercial applications developed to guide users through mindfulness meditation. Headspace [20] is one of the most popular mindfulness application that leads users through a succession of short meditations, guided by an audio recording of a human meditation instructor. In recent years, several research projects have explored ways to support mindfulness activities with automated and interactive interfaces such as mixed-reality environments [33, 40], haptic sensory systems [29] and wearable devices [35]. In one study, Niksirat and his colleagues presented an interactive mobile application for self-regulation and mindfulness exercise, called PAUSE. Their application detects the user's attention

without using any sensors and provides real-time feedback. PAUSE outperformed Headspace in attention improvement, and showed better mindfulness improvement in busy environments [34].

Other researchers have investigated how providing biofeedback based on physiological responses in Virtual Reality (VR) can improve a mindfulness experience. For instance, visual and auditory real-time feedback based on the level of the skin conductivity in a VR environment has been shown to be useful for pain management [18], and Roo showed the effectiveness of a mixed-reality environment that reflects the user's heart rate and breathing for mindfulness. RelaWorld is also a neuroadaptive VR meditation system [25] in which a head-mounted VR display is used to show a relaxing virtual environment. The system monitored users' brain activity, and provided adapted feedback for novices during a mindfulness meditation session. An evaluation study found that the combination of neurofeedback and the head mounted display resulted in significantly deeper levels of relaxation and increased feelings of presence compared to those of a control group without any feedback or VR display. Neither of these projects used a conversational agent interface, neither used adaptive verbal instructions and feedback, and neither used respiration beyond biofeedback.

2.1.1 Respiration-based Feedback for Mindfulness. Several technology based mindfulness programs incorporated respiration data both for monitoring the practitioner's state of the mind [19], and providing biofeedback in interactive systems aimed at reducing stress. For example, Sonic Cradle is an interactive installation designed to facilitate meditation, reacting to users' breathing with audio feedback (music), in which users can control and shape an abstract sound experience by focusing on their own breathing [41]. Results of an evaluation study showed a significant increase in resting respiration length before and after a sonic cradle experience [40]. In a similar vein, Moraveji et al. presented a peripheral respiration pacing program and demonstrated that participants' respiration rate significantly decreased using the pacing program compared to a no-feedback condition [30]. Wongsuphasawat and Moraveji continued their work by developing a mobile application called "breathwear" to regulate breathing and reduce stress. The study demonstrated that auditory guidance and feedback are more effective than visual feedback on self-report levels of calm [43]. Yu et al. also developed "Breathe with Touch"; a haptic interface in which users receive feedback on their respiration by following the changes of the shape of the interface [45]. Their results showed significantly higher satisfaction with the new interface, however the stress reduction was not significant with the haptic interface. Sonne et al. developed ChillFish; a breath-controlled biofeedback game that can calm and distract children during a blood test [38].

While a fair amount of research has been conducted on breath as biofeedback in mindfulness systems, no prior systems have incorporated respiration as an input modality in directed meditation training systems. Several researchers have examined the use of breath as an input modality in entertainment applications [27], and assistive devices such as on-screen keyboards, a breath joysticks [12], and powered wheelchairs [44], but none of the meditation programs have used breath as an *intentional* input in an interactive system to promote mindfulness meditation, or any kind of relaxation or calming. We believe that respiration is a crucial input modality for meditation, given the centrality of breathing in meditation training,

and because of its unconscious and peripheral nature it is important for calming user interfaces which avoid overburdening users as they interact with computers. Our work extends the use of respiration in human computer interaction by focusing on mindfulness meditation as a common and effective way for stress reduction, and by introducing new functions using user respiration as both an intentional input to interact with the system, and a basis for providing biofeedback.

2.2 Embodied Conversational Agents as Virtual Coaches

There are numerous research papers describing the design, development and evaluation of embodied conversational agents that act as virtual coaches in training [2], education and behavior change. For example, Bickmore, et al., developed a series of virtual coaches for health counseling [3], showing the effectiveness of virtual coaches in health counseling, especially for patients with low health or computer literacy.

A few researchers have also explored the use of virtual coaches in meditation training, although these did not adapt their interaction to user breathing. Hudlicka developed a virtual coach displayed as a static image that provided audio-recorded guided meditation training, along with the coaching support necessary to begin regular practice [21]. Results of a 7-week evaluation study indicated the effectiveness of the coach system, with participants in the coach group practicing significantly more frequently, and for a longer time, compared to those in a control group. Shamekhi et al, showed the effectiveness of a conversational agent to alleviate chronic pain and stress using yoga and meditation instruction, in conjunction with human-led group visits [36].

Despite the huge potential of affect sensing in coaching systems, most of these virtual coach systems do not respond to sensed input from users. One exception is an embodied pedagogical agent developed by D’Mello and Graesser, that acted as a virtual tutor, adapting its instruction based on real-time assessments of student affective state. Evaluations indicate significant learning improvement for students who worked with Autotutor [11].

2.3 Summary of Innovation

Our work is unique in that we combine the use of an embodied conversational agent as a virtual meditation coach with adaptive responsiveness to user respiration. The agent provides a rich communication medium and is available to users of all levels of health and computer literacy, and the respiration sensor enables several novel breath-driven interaction techniques, such as the input modality, biofeedback and mirroring, described in the next section.

3 DESIGN OF THE VIRTUAL MEDITATION COACH

The design of our virtual meditation coach was inspired by our experiences with human meditation coaches, as well as by previous work on embodied conversational agents that played the role of coaches for wellness and healthcare applications. Our initial prototype was developed for use in a broader intervention to help individuals with chronic pain using yoga and nutrition counseling in addition to meditation training. This coach was deployed on touch screen tablet computers, and incorporated a scripted,

non-adaptive meditation session that required users to touch the screen at certain key times to indicate they were ready to proceed. Although we received positive feedback from patients about this coach, we felt that having to use the touch screen was disruptive to the meditative state and sought other input modalities we could use to make the experience more adaptive and natural.

The current virtual meditation coach is an embodied conversational agent, animated and rendered in a 3D game engine using custom animation software (Figure 1). The agent’s appearance is designed to be a racially ambiguous female in her mid-forties. The agent has a range of nonverbal behaviors such as iconic, emblematic, and deictic hand gestures, facial displays of emotion, posture shifts to mark topic boundaries, eyebrow raises for emphasis, and head nods. Nonverbal behavior is determined for each agent utterance using the BEAT text-to-embodied speech system [9] and is synchronized with the agent’s speech. The agent’s voice is generated by a commercial text-to-speech program¹. Since the quality of the instructor’s voice is crucial in relaxation practices, we manipulated the agent’s voice and intonation to be as calming as possible. We used the speech synthesis markup language (SSML), to slow the speaking rate, insert pauses, and lower the baseline pitch. A calming background music is also played at a low volume when the system is running to help the agent’s voice sound more calming. The overall system is deployed on an integrated touch-screen desktop computer. Previous research [43] showed that visual feedback (e.g., abstract imagery) during meditation, is less effective than audio feedback, thus we did not include visual feedback in our system.

4 INCORPORATION OF RESPIRATION INTO THE VIRTUAL MEDITATION COACH SYSTEM

A mindfulness meditation requires a high level of concentration and awareness. Requiring meditation practitioners to tap on a touch screen to interact with our coach is a source of distraction, since it interrupts the meditation session. Thus, we replaced the touch modality with a less distracting and more responsive interface. Given the close connection between meditation and breathing, the respiration process was an ideal choice around which to develop a new responsive interface.

Respiration is the most natural bodily function we perform continuously and is much less distractive than touching the screen or even speaking while a user is meditating. Therefore, in the development of this system we introduce "breathing" as an innovative way of interaction with a computer, through which people can simply signal the system by breathing with different patterns.

In order to sense user respiration, we first evaluated the commercial Spire sensor [39], but at the time, Spire did not provide real-time data output. We also considered the approach taken in Blui [31], in which users blow loudly into a computer microphone, but we felt blowing was not conducive to relaxation. We decided to use a chest expansion sensor to provide information about user respiration –including inhalation depth, respiration rate, and respiration pattern– and used this information in three different ways in the system. The sensor, from Thought Technology (TT), is a flexible rubber strap, which can be worn either on the chest or the abdomen

¹IVONA text to speech by IVONA Inc.




Figure 2: Experiment setup. Participants worn respiration sensor on the chest over the clothing.

over clothing to measure chest expansion as a user breathes (Figure 2). Data from the sensor is processed in real time to provide inputs for the meditation coach to adapt its instruction.

Prior to each meditation session, the system is calibrated by asking users to breathe normally for 30 seconds (normal breathing calibration), and then take a few deep breaths for 20 seconds (deep breathing calibration). The average breath duration and the breathing rate are calculated for each participant based on the normal breathing calibration data. These calculated amounts as well as the maximum and minimum length of the chest expansion strap are stored after normal and deep breathing calibration to be later used for feedback and deep breath continuation during the meditation session. We will describe the three respiration-based interaction techniques in the following paragraphs.

4.1 Deep Breath as an Input Modality

In human-led meditation sessions, participants are often asked to take a deep breath at key points during their practice. We decided to use the “deep breath” as an acknowledgment and continuation signal to replace the need to tap on the touch screen. The agent dialog system allows users to signal that they are ready to continue by simply taking a deep breath. The system detects when a user takes a deep breath by processing the respiration data in real time. At several points during the meditation session (e.g., before changing to a new set of instructions) the agent asks the user to take a deep breath when he/she is ready to continue. This function allows users to adjust the pace of their practice, and the amount of time they spend on different parts of their practice, without being distracted by moving their hand to touch the screen or offering verbal inputs. This function uses the data collected during the calibration phase to detect a deep breath, defined as a threshold related to the stored maximum and minimum chest expansion. The threshold is set to 90% of the expansion range observed during the deep breath calibration for each user.

4.2 Respiration Rate Feedback

The respiration sensor is also used to compute a user’s breathing rate so that the coach can provide feedback on it at key points during the meditation. The system uses the normal breathing calibration data to calculate the baseline respiration rate. During the meditation

session, the number of user’s breaths per minute is calculated during certain time intervals, and the agent adapts the instructions she offers based on the respiration rate. For example, the virtual coach waits longer for the user to relax if they are still breathing fast (compared to the baseline), and asks her/him to breathe at a slower pace, or to just keep breathing. This second function improves the interactivity and adaptability of the system, as the agent guides users to become calmer and more relaxed during their practice. To be consistent with the nature of mindfulness meditation techniques, we tried to provide feedback in a nonjudgmental way and just to raise awareness.

4.3 Respiration Mirroring by the Virtual Coach

Mirroring is a pedagogical technique in which an instructor mimics a behavior of a student to raise awareness about that action [37] through meta-cognition. Mirroring is used as a way to promote awareness, behavior change and learning [32] through social interaction [16] in prior research. The virtual meditation coach is designed to mirror users’ breathing in order to help them be more aware of their breathing pace and encourage them to regulate their breathing. At certain times during the meditation session, the agent starts breathing with the user after stating that “I am breathing with you for the next couple of seconds”, and then appears to inhale and exhale in rhythm with the user.

To develop the mirroring module, we used an audio recording of a human female doing slow breathing, and split it into five parts: start inhale, inhale, transition from inhale to exhale, exhale, exhale end. During mirroring, the respiration sensor continuously outputs three breathing states based on the user’s breathing: inhale, exhale, and exhale to inhale transition. The audio clips are queued to play (repetitively, if necessary), to make the agent’s “voice” sound as if she is breathing in synchrony with the user.

In the final system, the user can signal the agent to move on to the next step by taking a deep breath (function 1) at seven points during the meditation. There are also three points that the mirroring function and feedback function are called consecutively in each session. For example the coach starts mirroring their breathing right after saying: “... Feel the breath moving in and out of your body... So lets empty our breath, empty our lungs with an exhale, then inhale. Let me breathe with you...” After the mirroring the coach monitors the users breathing and provides feedback based on their respiration rate (e.g. “continue breathing at a slower pace.”).

4.4 Sensor Validation Study

To validate the accuracy of the chest expansion sensor and respiration-based functions used in the virtual coach system, we conducted a validation study. We recruited 11 participants (63% female, mean age = 33 SD = 11.4), outfitted them with the sensor, and performed our calibration procedures.

To determine the accuracy of deep breath detection, we asked each participant to go through three 20 second cycles of breathing normally and breathing deeply, following prompts on a monitor. The system correctly detected the first deep breath 85% of the time, with the remaining 15% detected by the fourth deep breath. The system never classified a normal breath as a deep breath (85% sensitivity, 100% specificity).

We then asked participants to breathe “normally” for 20 seconds,

and “quickly” for 20 seconds, while being monitored both by the system and by a trained exercise physiologist using a stethoscope to obtain gold standard respiration rates [37]. The accuracy of the sensed breathing rate was 91.8%, and overall precision was ± 0.69 , relative to the standard measure.

4.4.1 Formative Pilot Study. Prior to running the main study, we conducted a preliminary pilot study with a limited version of the meditation coach system to assess users’ reactions to the responsive virtual meditation coach (previously published in [blind ref]). Overall, all nine participants were very satisfied with the virtual coach, rating it an average of 6.11 ($SD = 1.69$) on a 1 (not at all satisfied) to 7 (very satisfied) scale. Post-intervention interviews also indicated that participants enjoyed the experience. When participants were asked for areas of improvement, users often mentioned that they wanted even “*more feedback*” and interactivity from the virtual coach. There were also requests to change the way the agent communicated feedback. We made several changes to the system based on user feedback from the pilot study. Given that users were requesting even more feedback from the virtual coach, we doubled the frequency of occasions during a session in which the agent would give corrective feedback on user breathing rate.

5 STUDY 1: EVALUATION OF RESPIRATION FUNCTIONS

We conducted an experiment to evaluate the effectiveness of the respiration-based interaction in the virtual meditation coach system, by comparing the full system to an identical one in which the respiration sensing modality was removed, and participants were required to tap on the touch screen to advance the interactions. The study was a two-treatment, counterbalanced, within-subjects randomized experiment.

5.1 Method

Participants: Participants were required to be at least 18 years old, and be able to read and speak English to participate in the study. Overall, 21 participants were recruited from an online job posting site. Participants were: 62% female; mean age = 42, ($SD = 14$); 76% reported they use a computer regularly 57% had already tried mindfulness meditation a few times, and 23% meditated regularly.

Measures: We assessed the level of mindfulness in participants’ everyday experiences as a baseline covariate using the Mindful Attention Awareness Scale (MAAS) trait version [7]. Participants rated the meditation instructor using a composite measure self-report questionnaire consisted of six items ($\alpha = .88$). sample items include: *How satisfied are you with your instructor?*, *How much would you like to continue working with the instructor?* on a 7-point scale (1=not at all, 7=very much). We also assessed how participants evaluated the mediation experience by asking them to indicate how much they found the experience *relaxing* and *helpful*. Lastly participants were asked to rate the instructor’s awareness of their breathing to determine whether they felt the interactivity provided by the sensors. We also assessed their State Anxiety [28], Mindfulness (Toronto Mindfulness Scale [26]), and Flow State [22] after each treatment, using self-report measures. User’s respiration rate was recorded throughout each session as an objective measure to assess the physiological response to the meditation exercise guided by the virtual coach.

Procedure: Participants conducted two 12-minute meditations during an experiment session, one in each treatment (Figure 3). Prior to the meditation session, participants were asked to wear the breathing sensor over clothing for monitoring respiration. The research assistant made sure that the respiration sensor was tight enough to detect breathing, while also ensuring that the user felt comfortable wearing the sensor. In order to compensate for the possible carryover effects of the first meditation on participants’ starting mood for the second meditation, we had participants engage in a brief task designed to elicit frustration prior to the start of each meditation practice. In addition to controlling carryover effects, we were also interested in how the virtual meditation coach might be able to assist a user in managing his or her feelings of frustration or anxiety, and how this might make a computer more effective as a tool for teaching meditation.

Frustration was elicited by having participants engage in a two-minute interactive math test. Since emotional states are extremely difficult to elicit consciously, we needed to make participants feel that they did not perform well on the math test. This was a deception. The frustration task contained 10 basic math problems. The computer displayed each problem for 4 seconds, followed by multiple-choice responses for 3 seconds. For the first 5 problems, audio tones were played based on the user’s correct or incorrect answers. However, for the last 5 math problems, the computer played the tone associated with a wrong answer regardless of the participant’s response. Pre-testing demonstrated that this was effective at consistently eliciting frustration, in addition to comments from many participants that they felt this part of the study was frustrating. This same technique has been used in previous studies (e.g., [10]). We explicitly asked participants how the math test made them feel in post-test interviews, and most of our participants said that the tests elicited negative feelings. Participants were debriefed about the math test at the end of the study.

Following the frustration elicitation, and prior to the start of each meditation session, baseline readings of respiration was recorded for 30 seconds while participants were prompted to “sit quietly with your eyes open” (normal breathing calibration). Participants were then prompted to breathe deeply for 20 seconds to provide a baseline for the deep breath detection.

We also added a break in between the two meditation practices to eliminate the effects of the first practice as much as possible. Participants were asked to go outside the study room, fill out a questionnaire, and walk around for 5 minutes, before beginning the next session. The respiration sensor and calibration protocol were not actually needed for the non-interactive control condition intervention, but were included in the procedure for consistency, and to collect data for comparison with the intervention condition.

In order to gain a better insight on participants’ attitudes toward our system, we also conducted a semi-structured interview at the end of both meditations, before they were dismissed. The study was approved by the Institutional Review Board (IRB) at our University, the whole session took about 90 minutes, and participants were compensated \$15 for their time.

5.2 Study 1 Results

Self-report: We examined participant’s subjective perception of the instructor, their meditation experience and how they found




Figure 3: Study procedure

the instructor aware of their breathing. Table 1 shows mean values, standard deviation of each condition, as well as the p-value in testing the effect of treatments. Participants rated the meditation experience higher –in terms of relaxation and helpfulness– when the agent used sensors to monitor their respiration, and provided tailored feedback ($p < 0.03$). They also perceived of the coach significantly as more aware of their breathing in the “sensor” condition ($p < 0.001$). Participants rated the instructor significantly greater than neutral in both conditions as demonstrated by a single sample Wilcoxon signed rank test (Sensor p-value = .008, NoSensor p-value=.012), however there is no significant difference between two groups for instructor evaluation. We found no significant differences between study conditions on State Anxiety, Mindfulness, or Flow State. There were no significant effects of treatment order on any self-report measure.

Respiration: We pre-processed respiration data for all participants. Pre-processing included cleaning and resampling the data. Each participant experienced two meditations, resulting in 42 respiration data sets. We eliminated one respiration data file that the data was not recorded correctly due to a loose sensor connection, leaving 20 respiration records for analysis. The respiration data was then resampled to 10 Hz prior to analysis. The outliers were removed to assure a more robust analysis. To prevent aliasing, the signal was passed through a low-pass Butterworth filter to remove the high-frequency artifacts We split the signal into two segments; the

frustration task, and calibration/meditation task. Breath rate was calculated by counting the number of peaks (inhalation end points) per minute. The peaks were detected using the “findpeaks” function in MATLAB, using thresholds for peak distance and prominence based on the signal’s features. We slide a 30-second window across the signal with 15-second overlaps and calculated the respiration rate (RR). The average respiration rate in baseline calibration across two study conditions was 13.76 breath per minute (SD=2.9), while the average breathing rate during the deep calibration interval was 9.6 bpm (SD=2.1).

We computed the respiration rate in the second half of the meditation session to see whether the average breathing rate had changed during meditation. The average respiration rate in the second half of the meditation across two conditions was 8.1 bpm (SD=2.1). We were interested to see whether the mindfulness meditation influences the respiration rate, so we looked at the changes of RR from the first baseline (RR during normal calibration of the first condition) and the second half of each meditation. We finally compared the decrease of the respiration rate (Delta = BaselineRR - Second-HalfRR (green shaded area in Figure 4)) between the two conditions to assess the effectiveness of the sensor system on breath regulation. We found a significantly greater decrease of respiration rate between the baseline and end of meditation in the sensor condition ($M=6.04$, $SD=3.5$), as compared to the no-sensor condition ($M=4.8$, $SD=3.01$), $t(19)=2.69$, $p=0.014$. Figure 4 illustrates the respiration rate (mean breath rate in 30 seconds intervals) for one participant during one meditation session.

Based on these results, we conclude that the integration of respiration sensor with the virtual meditation coach system provided a more relaxing experience and resulted in better respiration regulation. In the remainder of this section we delve into the interview data to find out how people perceived the interactive system and why they found it helpful and relaxing.

5.3 Study 1 Qualitative Feedback

All interviews were transcribed by an independent research assistant and analyzed using thematic analysis techniques. In this inductive approach, we first open-coded the data and then clustered similar codes to identify the main themes, which were related to effects of the interactive system on relaxation, and the pros and cons of the virtual coach versus a human coach.

5.3.1 Effectiveness of the Virtual Coach for Relaxation. Overall, all participants were highly receptive of the virtual meditation coach, and found the instructor “very helpful” in “concentrat[ing] on [the] body and just slow down and relax.” They indicated that the virtual coach was able to influence their mood even after a short session. This effect was frequently echoed in the interviews.

“Overall it was really good. What I found interesting, at first I was actually nervous for the math questions... and then I was able to go from a fairly nervous state to a very relaxed state

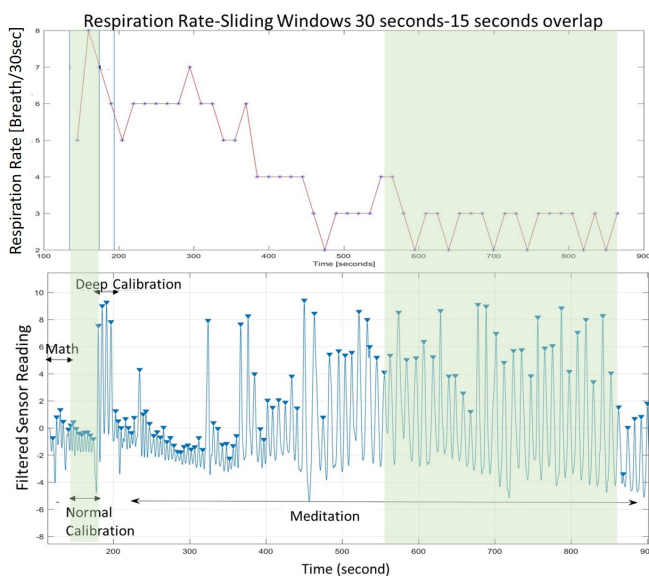


Figure 4: An example respiration signal from one participant with detected peaks during calibration and meditation (bottom), the chart in the top shows the mean breath rate as sliding a 30-second window over the signal.

Evaluation Scale	Study 1 (n=21)			Study 2 (n=24)			Test Type
	Sensor	No-Sensor	p-value	Agent	Video	p-value	
Instructor	5.05 (1.47)	4.88 (1.46)	0.32	4.72(1.39)	5.37(1.1)	0.01	Paired Sample ttest
Meditation Experience	6.02 (1.07)	5.19 (1.97)	0.03	5.35 (1.6)	5.66 (1.51)	0.41	Wilcoxon Signed Rank
Interactivity*	5.47(1.93)	3.57(2.29)	0.001	4.82(1.56)	3.89(1.99)	0.02	Wilcoxon Signed Rank

Table 1: Results of the instructor evaluation questionnaire, the experience evaluation and the system interactivity evaluation-study 1 and 2. *This scale had one item (instructor’s awareness) in study 1, and 2 items (awareness and tailored instructions) in study 2

because the meditation did a really good job, better than I was expecting.” [P8]

Overall, 61% of participants explicitly mentioned that they found the experience relaxing.

Comparison of sensor vs. no-sensor conditions: When asked to compare the two conditions, most participants indicated a stronger preference for the sensor condition, because they found it more “relaxing” and “interactive” when they were “participating with her [the coach]”. They found the feedback by the agent useful and appreciated the adaptive guidance in the sensor condition.

“The first one [sensor] was more guided and gave more direction and feedback at the same time, as well as monitored feedback through my breathing... it was very easy and very clear to do that and I think because I was being guided. I felt because I had the sensors she was able to easily sense what was going on, and in the second one it was very non-interactive. Not a lot of guidance.” [P14]

Also 86% of the participants found the agent more aware and “on point” with their breathing in the sensor condition. They appreciated how the coach noticed and acknowledged their deep breaths, and some said the right timing of the instructions and the instructor’s awareness improved their experience. Participants also commented on the specific respiration-based interaction methods.

Breath as a reciprocal signal: All participants considered the integration of respiration sensor as a useful technology for meditation training applications. They appreciated the use of respiration as an interaction modality, and showed strong preference for the deep breath continuation signaling over tapping on the touch screen.

“I liked the second one [senor] a lot because it gave feedback on how I was breathing. I was breathing correctly and instead of having to click in the first session this was much more relaxing because I wasn’t coming out of my zone to hit continue.” [P17]

In addition, many participants found tapping the touch screen “very distracting” versus the breathing that is “more fluid”. Participants described the deep breath signaling as a “natural” and “organic” signaling system, that “does not break their concentration”.

Breath signaling also had unanticipated benefits. Many appreciated the opportunity to occasionally take a deep breath, and also appreciated the coach’s recognition of their doing this successfully.

“I like how if you took a deep breath and if you did it in a satisfactory way that was good feedback because it was encouraging me to breathe.” [P8]

Thus, we argue that the deep breath works as a reciprocal messaging system between the user and the agent, through which users inform the coach to move on, and the coach provides a feedback by acknowledging their deep breath. However, one participant did not like the breath signaling specifically for this reason, because she felt she “was trying harder to please her[the coach]” in the sensor condition. We believe it is an important issue and such feelings

may contradict the nonjudgmental and mindful spirit of meditation exercises. For this particular participant it happened because the sensor was not tight enough at times during the session. It is important to first use a sensitive enough sensor to avoid missing input, and secondly to design the instructions not to be demanding as much as possible.

Adaptive feedback as an instructional guide for novices: Participants commented positively on the biofeedback they received from the virtual coach during the meditation session in the sensor condition. They found the feedback and guidance “useful” and “very helpful” to “stay on track”. Many participants also envisioned the feedback feature as a way to increase control in their practice, “I had more control in how well I was able to do” [P10]. Our participants commented on how such a program could be useful for novice practitioners as they may need more feedback to learn how to meditate. They indicated that they found the biofeedback more “helpful, especially if you are not used to meditation”. P10 added:

“I liked the feedback because never have done this before. I liked to know if I was doing it right, and liked to hear her exhaling loudly because I had no guidance otherwise. Second time [sensor] I felt I had more feedback, like she was understanding that I took a deep breath because she would say awesome.”

They also appreciated the agent’s mirroring feature as a guide.

“She said let me breathe with you so we were breathing together versus by myself which was helpful. It was really helpful to hear the breathing voice. She was with me in sync.” [P21]

However, some participants did not like the coach breathing sounds, especially when it was not synchronized perfectly with their breathing, finding it “creepy” and a “little distracting”.

5.3.2 Can a virtual coach ever outperform a human coach? Many participants commented on the voice and appearance of the virtual coach compared to a human instructor. Most participants found the synthetic voice acceptable. Some actually liked her synthetic voice and found it “relaxing” and “calming”, while a few said that her voice was “too robotic” and could potentially distract practitioners during meditation. Regarding the appearance of the coach, some participants found her appearance “calming”, and several of them appreciated the fact that they thought she was a woman of color.

Many of our participants compared the virtual coach to a human instructor, as well as to their previous meditation experiences. They acknowledged the personalized feedback, interactivity and the coach’s awareness as advantages of the virtual coach over a human instructor in a class.

“I was very impressed because it’s a very personal thing but when I go to yoga classes I can get quite irritated by the tone of instructors. Some people talk really fast so I really enjoy instructors who can talk calmly to you.” [P18]

Such comments illustrate the desire to have more customizable features for the coach, and we think that computerized meditation

programs offer the potential to provide personalized and customized experiences. Participants' comments regarding the comparison between the virtual coach and a human instructor suggest that although humans voice and look may be more preferable, a virtual meditation coach has the potential to actually outperform some aspects of a human coach in a public class. Thus, we decided to conduct a second experimental study.

6 STUDY 2: COMPARATIVE EVALUATION STUDY

While Study 1 demonstrated the effectiveness of respiration-based responsiveness, we wanted to compare the responsive coach to widely available self-help media to evaluate the performance of our approach over one of the most common solutions people currently use to learn meditation. We conducted a summative evaluation of the full virtual meditation coach system by comparing it to passive observation of a publicly-available 11-minute video of a guided breathing meditation led by an experienced meditation instructor. Many self-help books, videos, and audio guides exist to provide meditation instructions. However, these instructional media are neither tailored to the user, nor interactive in real time. Thus, we felt that a study that compares the virtual meditation coach to one of the most common existing alternatives would provide a strong validation of the concept.

The study design is similar to the first study; it is a two-treatment, counterbalanced, within-subjects experimental design. In one treatment, the virtual coach guides the meditation session as she monitors user respiration, while in the other condition, participants watch an 11-minute video of a meditation instructor. As before, we used a frustration task and break (Fig. 3) to reduce carryover effects.

6.1 Method

Eligibility requirements and recruitment methods were the same as in the previous study. Overall, 24 participants were recruited with an average age of 40.9 years (SD = 14.5); 63% were male; 83% used a computer regularly and the rest rated themselves as computer experts; 54% had tried mindfulness meditation a few times before, 29% never tried meditation before, and 16% said they meditate regularly. None of the participants of our first study were recruited for this study.

Measures: We used the same self-report measures as in the first study, with the addition of an item to the last scale (*To what extent did you feel that the instructor tailored her instructions to your breathing?*) to assess the interactivity of the system.

Procedure: Participants conducted two 12-minute meditations in a single experiment session (one in each treatment). The order effect was controlled by counterbalancing the conditions. We followed the same protocol as in the prior study (Figure 3), except that we compared the full virtual meditation coach system to a meditation in which participants passively observed a video of a human instructor. We used a video that was published by Eckhart Tolle TV and has been viewed by more than 283K users on YouTube. We modified the dialog script used by the virtual meditation coach so that its language, timing, and synthetic speech prosody were as similar to the video as possible, to control for instructional content.

6.2 Study 2 Results

Self-report: Differences on the evaluation scales are shown in Table 1. Participants indicated that the virtual coach was significantly more interactive and adaptive to their breathing compared to the videotaped human coach. However, they rated the human coach significantly higher than the virtual coach on overall satisfaction. There were no significant differences between treatments on State Anxiety, Mindfulness, or Flow State.

Respiration: We preprocessed respiration data for all participants similar to the first study. The average respiration rate in baseline assessment was 14.3 breaths per minute (SD = 4.5). The analysis showed that the respiration rate significantly decreased in the second half of the meditation in both treatments (Agent condition: $M=7.7$, $SD=2.2$, $t(23)=8.7$, $p<.0001$, Video condition: $M=7.6$, $SD=2.4$, $t(23)=6.7$, $p<.0001$), however we did not find any significant difference between the two groups.

6.3 Study 2 Qualitative Feedback

Interviews were transcribed and analyzed as in Study 1. In general, many participants stated that the virtual coach was a comparable relaxation method to watching the video of the expert meditation instructor, and *"they were both pretty relaxing"*. Since the quantitative results showed a general preference for the human instructor, we focus here on the reasons participants gave for this preference.

6.3.1 Feedback; Encouraging or Judgmental? As in Study 1, participants commented positively on the biofeedback function of the virtual coach, and perceived the coach as more interactive.

"The first one [agent] gave me more advice on breathing. So it was more interactive than the second one... I felt like the first one was much more aware, but that's how she was giving me feedback. You know you're doing well. The other one it didn't do that, it kept instructing through the whole thing" [P2]

They found the feedback by the agent, *"encouraging"*, *"helpful"* and *"constructive"* and felt that the feedback could let them know how they were progressing.

Some participants indicated that the respiration monitoring made them focus more on their behavior. For instance, P16 found it very helpful because *"it makes (him) think that's what (he) should be doing."* On the other hand, a few participants raised the concern that the feedback could be distracting by making people think about whether they are breathing right or wrong.

Three participants mentioned that the video was less demanding, and that made it easier to relax. They thought a practice might not be as relaxing when someone is required to act in a certain way.

"I felt like the first one [video] she wasn't really talking directly at me or she really wasn't following what I was doing and that made it more organic for it to be comfortable with. The second one she was adjusting to things I was doing, but I had to do things to achieve the next things or to keep going." [P5]

In this specific case the participant preferred to follow instructions without being monitored. However, we argue that this is inconsistent with the nature of any meditation practice, since paying attention on purpose and trying to be in the present moment are the main aspects of mindfulness meditation [23].

6.3.2 Human versus Computer. Although many participants appreciated the feedback, signaling and interactivity of the coach, most of them preferred the human instructor in the video, when

they were asked about their overall preference simply because it was “*more human*”.

"I just always connect better with a human interaction and it's the same thing like I'd prefer an in-person conversation rather than a Skype conversation, it's just I feel like it's easier. I'm also not a computer person I don't really love animation." [P17]

There were also many participants who justified their preference for the human instructor by describing differences in the appearance and voice qualities between the human instructor and the agent.

"I think they both worked pretty well. But the human voice was a little more soothing so I think that allowed me to be slightly more relaxed." [P4]

Participants also indicated that the fact that they interacted with a “*real person*” versus an “*animated person*” affected their attitude towards the instructor, as P7 said: “*I think the human element in the first one [video] was the key component that I liked the most*”. Regarding the instructor’s awareness of their breathing, most participants found the virtual coach more aware of their breathing because they received feedback and acknowledgment from the coach. However a few participant associated “*being human*” with “*being more aware*” and “*being approachable*”. So, their preference for the human instructor who “*is breathing herself, and knows what it feels like to breathe in and to breathe out*” could be also be simply due to the feeling of being more connected and in touch with the human instructor.

7 DISCUSSION AND CONCLUSION

Our experience developing and evaluating a virtual meditation coach that is responsive to user respiration indicates that users clearly recognized the new input modality and that they appreciated the novel interactions that it afforded including the tailored feedback. In the first experiment, participants felt that the virtual meditation coach was significantly more aware of their breathing when it incorporated information from a respiration sensor in its tailored training. Participants also reported that their meditation was significantly more relaxing and helpful when the coach responded to their breathing. They also had a significantly stronger respiration regulation with the breath-sensitive coach as measured by their respiration rate during meditation. Our qualitative findings shed light on how the provided guidance by the responsive coach could improve the meditation experience, and highlighted the importance of the agent’s voice in mindfulness applications, while there may be less attention to the visual details. Thus we suggest the designers to attend closely to the coach’s voice, and leave the visual features with less details to provide more calming experiences in mindfulness applications.

In the second study, although participants rated the virtual meditation coach as being more interactive and reactive to their breathing compared to the videotaped human, and well above neutral on satisfaction, they were significantly more satisfied watching a video of a human meditation instructor. The interviews indicate that a major source of displeasure with the virtual coach came from its lacking human-like features, including its synthesized voice. Although we tried to make the instructor’s voice as calming as possible, it is not nearly as effective as an expert human’s speech at relaxing people. This finding is not surprising, as previous research has shown that people rate virtual agents more engaging when

human recordings are used as their voice compared to synthetic text-to-speech [14]. Georgila et al. also explored the differences between recorded human speech and synthetic voices, regarding naturalness, conversational abilities, and likability, and found that human voices are perceived better in all aspects. However, the ratings of high quality synthesized voices were not far behind those of human voices in absolute terms [15]. In retrospect, it was not surprising to see that participants preferred a human voice over synthesized speech, and we believe that this is even more important for meditation than most other applications. Future research should also explore how allowing users to customize their virtual coach’s voice, speaking pace, and prosody could affect their satisfaction.

Overall, participants responded positively to the respiration-based interaction modalities, including the deep breath continuation and mirroring functions. Conducting two evaluation studies we found that it is crucial to keep the instructor’s language nonjudgmental and the input modality (breathing versus tapping versus no input) is more appreciated when requires minimum effort to proceed and has less distraction.

Taken together, these results indicate that the virtual meditation coach that is responsive to user respiration is an ideal medium for teaching guided meditation, although the appearance and voice should be more human-like for optimal effectiveness, at least with new users. One direction of future work will focus on how to improve the voice quality. Meditation training may be one class of applications in which using recorded human speech or video (especially from an expert coach) is especially important. Video of an expert human coach providing an introduction may be sufficient to establish an initial therapeutic relationship, after which calming imagery could be shown during the interactive part of the meditation: adaptive character animation (e.g., during mirroring) may not be necessary since participants often have their eyes closed during actual meditation sessions. Another promising approach may be to use segmented video clips of a human instructor to combine the interactivity and tailoring of the virtual coach with the appearance and voice of an expert human meditation instructor.

Of course, all of the effects we evaluated were from single, brief meditation sessions. These interventions must ultimately be evaluated in a longitudinal context. It may be that longitudinal adaptivity in which the virtual coach remembers a user’s past performance, and goals becomes crucially important in real use cases, and that the relatively superficial differences between an animated conversational agent and a videotaped human instructor become irrelevant. People may also acclimate to the synthetic voice over time.

7.1 Limitations

Our studies have several limitations beyond the small convenience samples used. We did not evaluate system use in a real (e.g. home) context over an extended period of time, which would have more ecological validity. We also used relatively short meditation sessions, which are appropriate for beginners but may not be adequate for more experienced practitioners. Although study conditions were counterbalanced to minimize the carryover effect, and participants had gone through a frustration task before each meditation, their comments indicate that there was some carryover effect with a few participants saying they felt more familiar with the system and meditation in the second meditation session compared to the first.

REFERENCES

- [1] Albert J Arias, Karen Steinberg, Alok Banga, and Robert L Trestman. 2006. Systematic review of the efficacy of meditation techniques as treatments for medical illness. *Journal of Alternative & Complementary Medicine* 12, 8 (2006), 817–832.
- [2] Yulong Bian, Chenglei Yang, Dongdong Guan, Sa Xiao, Fengqiang Gao, Chia Shen, and Xiangxu Meng. 2016. Effects of Pedagogical Agent's Personality and Emotional Feedback Strategy on Chinese Students' Learning Experiences and Performance: A Study Based on Virtual Tai Chi Training Studio. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 433–444.
- [3] Timothy W Bickmore, Laura M Pfeifer, and Brian W Jack. 2009. Taking the time to care: empowering low health literacy hospital patients with virtual nurse agents. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 1265–1274.
- [4] Timothy W Bickmore, Kathryn Puskar, Elizabeth A Schlenk, Laura M Pfeifer, and Susan M Sereika. 2010. Maintaining reality: relational agents for antipsychotic medication adherence. *Interacting with Computers* 22, 4 (2010), 276–288.
- [5] Alison Brazier, Andrea Mulkins, and Marja Verhoef. 2006. Evaluating a yogic breathing and meditation intervention for individuals living with HIV/AIDS. *American journal of health promotion* 20, 3 (2006), 192–195.
- [6] Calm. 2016. Calm: Meditation techniques for stress reduction. <https://www.calm.com/>.
- [7] Linda E Carlson and Kirk Warren Brown. 2005. Validation of the Mindful Attention Awareness Scale in a cancer population. *Journal of psychosomatic research* 58, 1 (2005), 29–33.
- [8] Linda E Carlson, Zenovia Ursuliak, Eileen Goodey, Maureen Angen, and Michael Specia. 2001. The effects of a mindfulness meditation-based stress reduction program on mood and symptoms of stress in cancer outpatients: 6-month follow-up. *Supportive care in Cancer* 9, 2 (2001), 112–123.
- [9] Justine Cassell, Hannes Högni Vilhjálmsón, and Timothy Bickmore. 2001. Beat: the behavior expression animation toolkit. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 477–486.
- [10] Katarina Dedovic, Robert Renwick, Najmeh Khalili Mahani, Veronika Engert, Sonia J Lupien, and Jens C Pruessner. 2005. The Montreal Imaging Stress Task: using functional imaging to investigate the effects of perceiving and processing psychosocial stress in the human brain. *Journal of Psychiatry and Neuroscience* 30, 5 (2005), 319.
- [11] Sidney D'mello and Art Graesser. 2012. AutoTutor and affective AutoTutor: Learning by talking with cognitively and emotionally intelligent computers that talk back. *ACM Transactions on Interactive Intelligent Systems* 2, 4 (2012), 23.
- [12] Grigori Evreinov and Tatiana Evreinova. 2000. " Breath-Joystick"-Graphical Manipulator for Physically Disabled Users. *Proc. of the ICCHP*, 193 200 (2000).
- [13] Norman AS Farb, Zindel V Segal, Helen Mayberg, Jim Bean, Deborah McKeon, Zainab Fatima, and Adam K Anderson. 2007. Attending to the present: mindfulness meditation reveals distinct neural modes of self-reference. *Social cognitive and affective neuroscience* 2, 4 (2007), 313–322.
- [14] Katherine Forbes-Riley, Diane J Litman, Scott Silliman, and Joel R Tetreault. 2006. Comparing Synthesized versus Pre-Recorded Tutor Speech in an Intelligent Tutoring Spoken Dialogue System.. In *FLAIRS Conference*. 509–514.
- [15] Kalliroi Georgila, Alan W Black, Kenji Sagae, and David R Traum. 2012. Practical Evaluation of Human and Synthesized Speech for Virtual Human Dialogue Systems.. In *LREC*. 3519–3526.
- [16] György Gergely and John S Watson. 1996. The social biofeedback model of parental affect-mirroring. *The International Journal of Psycho-Analysis* 77, 6 (1996), 1181.
- [17] Madhav Goyal, Sonal Singh, Erica MS Sibling, Neda F Gould, Anastasia Rowland-Seymour, Ritu Sharma, Zackary Berger, Dana Sleicher, David D Maron, Hasan M Shihab, et al. 2014. Meditation programs for psychological stress and well-being: a systematic review and meta-analysis. *JAMA internal medicine* 174, 3 (2014), 357–368.
- [18] Diane Gromala, Xin Tong, Amber Choo, Mehdi Karamnejad, and Chris D Shaw. 2015. The virtual meditative walk: virtual reality therapy for chronic pain management. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 521–524.
- [19] Tian Hao, Chongguang Bi, Guoliang Xing, Roxane Chan, and Linlin Tu. 2017. MindfulWatch: A Smartwatch-Based System For Real-Time Respiration Monitoring During Meditation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 57.
- [20] Headspace. 2016. Headspace: Guided Meditation and Mindfulness. <https://www.headspace.com/>.
- [21] Eva Hudlicka. 2013. Virtual training and coaching of health behavior: Example from mindfulness meditation training. *Patient education and counseling* 92, 2 (2013), 160–166.
- [22] Susan A Jackson and Herbert W Marsh. 1996. Development and validation of a scale to measure optimal experience: The Flow State Scale. *Journal of sport and exercise psychology* 18, 1 (1996), 17–35.
- [23] Jon Kabat-Zinn. 2003. Mindfulness-based interventions in context: past, present, and future. *Clinical psychology: Science and practice* 10, 2 (2003), 144–156.
- [24] Barbara Ann Kipfer. 2006. *Self meditation*. New York: Workman Pub. 432 pages.
- [25] Ilkka Kosunen, Mikko Salminen, Simo Järvelä, Antti Ruonala, Niklas Ravaja, and Giulio Jacucci. 2016. Relaworld: neuroadaptive and immersive virtual reality meditation system. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 208–217.
- [26] Mark A Lau, Scott R Bishop, Zindel V Segal, Tom Buis, Nicole D Anderson, Linda Carlson, Shauna Shapiro, James Carmody, Susan Abbey, and Gerald Devins. 2006. The Toronto mindfulness scale: Development and validation. *Journal of clinical psychology* 62, 12 (2006), 1445–1467.
- [27] Joe Marshall, Duncan Rowland, Stefan Rennick Egglestone, Steve Benford, Brendan Walker, and Derek McAuley. 2011. Breath control of amusement rides. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 73–82.
- [28] Theresa M Marteau and Hilary Bekker. 1992. The development of a six-item short-form of the State scale of the Spielberger State-Trait Anxiety Inventory (STAI). *British Journal of Clinical Psychology* 31, 3 (1992), 301–306.
- [29] Henry Moller and Harjot Bal. 2013. Technology-enhanced multimodal meditation to treat stress-related symptoms: Clinical results from an observational case series. In *Virtual Rehabilitation (ICVR), 2013 International Conference on*. IEEE, 35–40.
- [30] Neema Moraveji, Ben Olson, Truc Nguyen, Mahmoud Saadat, Yaser Khalighi, Roy Pea, and Jeffrey Heer. 2011. Peripheral paced respiration: influencing user physiology during information work. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 423–428.
- [31] Shwetak N Patel and Gregory D Abowd. 2007. Blue: low-cost localized blowable user interfaces. In *Proceedings of the 20th annual ACM symposium on User interface software and technology*. ACM, 217–220.
- [32] Margot Phaneuf. 2005. The mirror effect: Mediator of knowledge and self-image. the Philosophy and Education Convention, Evora, Portugal.
- [33] Joan Sol Roo, Renaud Gervais, Jérémy Frey, and Martin Hachet. 2017. Inner Garden: Connecting Inner States to a Mixed Reality Sandbox for Mindfulness. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 1459–1470.
- [34] Kavous Salehzadeh Niksirat, Chaklam Silpasuwanchai, Mahmoud Mohamed Hussien Ahmed, Peng Cheng, and Xiangshi Ren. 2017. A Framework for Interactive Mindfulness Meditation Using Attention-Regulation Process. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, 2672–2684.
- [35] Corina Sas and Rohit Chopra. 2015. MeditAid: a wearable adaptive neurofeedback-based system for training mindfulness state. *Personal and Ubiquitous Computing* 19, 7 (2015), 1169–1182.
- [36] Ameneh Shamekhi, Timothy Bickmore, Anna Lestoquoy, and Paula Gardiner. 2017. Augmenting Group Medical Visits with Conversational Agents for Stress Management Behavior Change. In *International Conference on Persuasive Technology*. Springer, 55–67.
- [37] Amy Soller, Alejandra Martínez, Patrick Jermann, and Martin Muehlenbrock. 2005. From mirroring to guiding: A review of state of the art technology for supporting collaborative learning. *International Journal of Artificial Intelligence in Education* 15, 4 (2005), 261–290.
- [38] Tobias Sonne, Timothy Merritt, Paul Marshall, Johanne J Lomholt, Jörg Müller, and Kaj Grønbaek. 2017. Calming children when drawing blood using breath-based biofeedback. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. ACM, 725–737.
- [39] Spire. 2016. Spire: Breath and activity tracker. <https://spire.io/>.
- [40] Jay Vidyarthi. 2012. *Sonic Cradle: evoking mindfulness through 'immersive' interaction design*. Ph.D. Dissertation. Communication, Art & Technology: School of Interactive Arts and Technology.
- [41] Jay Vidyarthi and Bernhard E Riecke. 2013. Mediated meditation: cultivating mindfulness with sonic cradle. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2305–2314.
- [42] Mark Weiser and John Seely Brown. 1997. *The coming age of calm technology*. In *Beyond calculation*. Springer, 75–85.
- [43] Kanit Wongsuphasawat, Alex Gamburg, and Neema Moraveji. 2012. You can't force calm: designing and evaluating respiratory regulating interfaces for calming technology. In *Adjunct proceedings of the 25th annual ACM symposium on User interface software and technology*. ACM, 69–70.
- [44] Motoji Yamamoto, Takeshi Ikeda, and Yoshinobu Sasaki. 2008. Real-time analog input device using breath pressure for the operation of powered wheelchair. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*. IEEE, 3914–3919.
- [45] Bin Yu, Loe Feijs, Mathias Funk, and Jun Hu. 2015. Breathe with touch: a tactile interface for breathing assistance system. In *Human-Computer Interaction*. Springer, 45–52.
- [46] Bin Zhu, Anders Hedman, and Haibo Li. 2017. Designing Digital Mindfulness: Presence-In and Presence-With versus Presence-Through. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2685–2695.